

Could Women Be Analytically Oppressed?

Derek Ball
University of St Andrews

There are a number of conceptions of philosophical analysis, but a common thread in most of these conceptions is that the aim of analysis is descriptive: we take our words and concepts as they are, and try to draw some philosophical conclusion from them. Frank Jackson captures this picture of philosophy when he writes, “What we are seeking to address is whether free action *according to our ordinary conception*, or something suitably close to our ordinary conception, exists and is compatible with determinism, and whether intentional states *according to our ordinary conception*, or something suitably close to it, will survive what cognitive science reveals about the operations of our brains” (1998: 31); analysis is important, on Jackson’s view, precisely because it ensures that we do not stray too far from ordinary thought and linguistic usage. The upshot is much in the spirit of Wittgenstein: “Philosophy may in no way interfere with the actual use of language; it can in the end only describe it. [...] It leaves everything as it is” (1953: §124).

Other philosophers – across a range of subdisciplines – have advocated a critical and constructive philosophical project: we should not simply take our words and concepts as they come, but should aim to improve them if we can. Among the most provocative of such attempts is Sally Haslanger’s social constructivist analysis of race and gender terms. Haslanger proposes the following analysis of “woman”:

S is a woman iff

- i S is regularly and for the most part observed or imagined to have certain bodily features presumed to be evidence of a female’s biological role in reproduction;
- ii that S has these features marks S within the dominant ideology of S’s society as someone who ought to occupy certain kinds of social position that are in fact subordinate (and so motivates and justifies S’s occupying such a position); and
- iii the fact that S satisfies (i) and (ii) plays a role in S’s systematic subordination, that is, *along some dimension*, S’s social position is oppressive, and S’s satisfying (i) and (ii) plays a role in that dimension of subordination. (2012a: 234)

Call this claim (W) (and set aside the many interesting questions involved in working out the details). Haslanger concedes that (W), construed as an analysis, clashes with a number of aspects of ordinary and expert usage and belief. (For example, given Haslanger's view, feminists should attempt to eliminate women (Haslanger 2012a: 239).) But she nonetheless maintains that this is the analysis that we should accept – at least if we share her broadly feminist goals, such as identifying and explaining “persistent inequalities between females and males” and uncovering how “social forces [...] work to perpetuate such inequalities” (2012a: 226-7), with an eye toward ultimately undermining these forces for the sake of social justice.

Analyses like (W) are *revisionary* in that accepting such an analysis involves a departure from ordinary or expert usage and belief. But the idea that something could be both *revisionary* and an *analysis* seems to stand in considerable tension with the idea that analysis must be descriptive. This tension has led opponents of revisionary analyses to the accusation that such analyses are *changing the subject*; roughly, the idea is that if our analysis is not describing current usage of, say, “woman”, it must be describing an alternative usage of a homophonic word, a usage on which “woman” would take on a novel meaning. To take an example from popular politics, proponents of same-sex marriage are often accused of trying to redefine “marriage”, where the subtext of the accusation is that opponents of same-sex marriage are right about *marriage*, and proponents (presumably knowing this) want to stop talking about marriage and start using the word “marriage” to talk about something else. A similar worry motivates Jackson's emphasis on continuity with ordinary thought and usage; as Jackson sees it, the alternative to descriptive analysis is stipulated change of meaning, but such stipulation is of little philosophical interest: “If I say that what *I* mean – never mind what others mean – by a free action is one such that the agent would have done otherwise if he or she had chosen to, then the existence of free actions so conceived will be secured, and so will the compatibility of free action with determinism [...] I have turned interesting philosophical debates into easy exercises in deductions from stipulative definitions together with accepted facts” (1998: 31).

Philosophical discussion of revisionary analysis, from Carnap (1956: 7) on, generally takes for granted the idea that revisionary analyses change the subject in this sense. The controversy is only over whether (and when and how) it makes sense to advocate for particular concepts, for particular changes of subject. Thus proponents of particular revisionary analyses often see themselves as defending the replacement of one concept with another (e.g., Scharp (2013)), and proponents of “conceptual engineering” and “conceptual ethics” present themselves as defending the idea that we can study the questions like ‘Should we use concept C (over alternative A)?’

(Burgess & Plunkett 2013a,b), and that the answers can inform our understanding of a range of popular and philosophical debates (Plunkett & Sundell 2013; Thomasson 2016).

In my view, the idea that revisionary analysis requires the *replacement* of one concept with another, or with *changing* the meaning of our words, or with *developing new concepts*, is entirely misguided. There is a perfectly natural sense in which one can advance a revisionary analysis like (W) as an analysis of “woman” as we have always used it, the word and the concept that we have been employing all along – without changing the subject or engineering a new concept or anything of the sort. There is no conflict between the idea that such analyses are revisionary and the idea that they are descriptive: they are revisionary, in that they depart from present usage and beliefs, but are still descriptive in that they are making a claim about what we mean now and have meant all along. Further, there is a perfectly natural sense in which revisionary analyses might turn out to be correct – not that it might turn out to be correct to adopt a new term or concept, or that changing the subject might be the right thing to do, but that (W) might turn out to be the correct analysis of “woman” as we have used it all along. (So it might turn out that women are subordinated not just as a matter of fact, but analytically, and have always been so – no matter how widespread belief to the contrary may be or have been.)

Of course, none of this is to deny the obvious fact that we sometimes introduce new terminology (or technical uses of extant terminology), and correspondingly new concepts, in the course of theorising. The point is that this is neither the only, nor the most plausible, account of what proponents of same-sex marriage, or social constructivist analyses of gender, or most other related debates, are doing. A better picture – better, because (as I will show) it better explains our epistemic engagement with revisionary analyses – is that conceptual analysis can go beyond (and even overturn) extant belief and linguistic usage without changing of meaning, and without introducing new concepts. Our theoretical activity shapes what we mean, but it does so not by making us mean something new, but by shaping what we meant all along. So we are rarely if ever faced with two concepts and forced to choose which to use. We *are* often faced with new claims and new theories; we accept some and reject others, on good or bad grounds, and this plays a role in determining our meanings and our concepts. But this is just part of what it is to be a thinker. There are no interesting questions of conceptual engineering or conceptual ethics, over and above questions about when we accept (or should accept) certain claims or theories, of the familiar sort long studied by epistemologists and philosophers of science, as well as psychologists and philosophers of mind.

This paper aims to clarify and defend these claims by developing an ac-

count of revisionary analysis that makes sense of our engagement with revisionary analyses, and in particular, makes sense of the way we argue for and against them – our epistemic engagement with them. I begin by examining these arguments, and showing that the view that revisionary analysis involves meaning change or concept replacement cannot make sense of them. In section 2, I turn to Haslanger’s proposal, and suggest that her view is best formulated in terms of competing analyses (rather than in terms of competing concepts). Section 3 develops the relevant notion of analysis – a notion on which analytic claims fix the meanings of words. Sections 4 and 5 show how this notion can be developed to explain our engagement with revisionary analyses. A key idea is that analytic claims can emerge in the course of debate without change of meaning, so that our acceptance (perhaps late in the course of a discussion) of some analyticity can fix the meaning of a word as we used it all along.

1 Changing the Subject

Consider the view that accepting a sentence like “Two men can marry each other” amounts to changing the meaning of the word “marry”.¹ As Jackson might put it, the claim is that the proponent of same-sex “marriage” says, “this is what I will mean by ‘marriage’ – never mind what others mean”; or perhaps better, “this is what I will mean by ‘marriage’, and others should do the same – never mind what we have meant in the past”. If this is right, the proponent of same-sex “marriage” might go on to speak truly by saying things like “Two men can marry each other”. But if someone thinks that a man can only marry a woman, she should not regard the proponent of same-sex “marriage” as contradicting her. The proponent’s utterance of “Two men can marry each other” does not express the proposition that two men can marry each other; she has not changed her belief in the falsehood of this proposition, but only her beliefs about what sentences are best used to express it. Her choice to use the term “marriage” in an idiosyncratic way does not change the facts about marriage, any more than choosing to call

¹The case of same-sex marriage is complicated by the fact that it was until recently a *legal* impossibility everywhere for two men to marry each other. But opponents of same-sex marriage – especially the kind of opponent of same-sex marriage who claims that proponents are changing meaning – tend to think that same-sex marriage is impossible in a deeper sense (conceptually impossible), so that laws permitting same-sex “marriage” must be regulating some state other than marriage. The claim “Two men can marry each other” is to be read as advocating the conceptual possibility of same-sex marriage; so (in the dispute I have in mind) the proponent of same-sex marriage is claiming that it would make sense to have laws that allowed same-sex marriage (and not merely same-sex “marriage”). This claim could be true even if same-sex marriage were legally impossible (though of course the proponent would presumably also advocate changing such a law).

a tail a “leg” makes it the case that horses have five legs. (And of course this would be true even if the proponent of this way of speaking is able to convince everyone to adopt her usage.)

If we are talking about legs, and I try to stipulate that “leg” is to be applied to tails, I am trying to change the subject.² If my stipulation succeeds and we continue to speak of “legs”, then the subject has (at least partially) changed; we are no longer speaking (just) of legs, but of tails as well. (If words are individuated by their meanings, we are no longer using the same word; I have introduced, and gone on to use, a mere homonym of the ordinary English word “leg”.) When we say “Horses have five legs”, we should not take ourselves to be disagreeing with those who, unfamiliar with my stipulation, say, “Horses have four legs” – except, of course, about a matter of linguistic usage.

Let’s say that an analysis that involves a change of meaning in this way *changes the subject*. Self-avowed proponents of conceptual ethics and conceptual engineering, and some proponents of revisionary analyses, tend to regard revisionary analyses as changing the subject. Call this claim the *subject-change view* of revisionary analysis.

Proponents of the subject-change view have made concerted attempts to explain some aspects of disputes about revisionary analyses. For example, they point to other cases in which we have a sense of genuine disagreement even where it is clear that the parties are not asserting contradictory contents (e.g., “I like Steve”; “No, you’re wrong: I hate him”). And they claim that in many circumstances, a primary purpose of an assertion may be to communicate metalinguistic information about how words are to be used (Sundell 2011; Plunkett & Sundell 2013).

These observations are fair enough as far as they go. But there is a great deal more to explain. We do not simply express disagreement about revisionary analyses and leave it at that: we offer and respond to arguments for and against them. I claim that the subject-change view cannot explain our argumentative practice: it cannot make sense of the kinds of arguments we offer, and the way we respond to these arguments.

Consider, for example, disputes about same-sex marriage. One kind of argument, invoked especially by opponents of same-sex marriage, appeals explicitly to the proper use of a word; for example, “We have always used ‘marriage’ to mean *a union of one man and one woman*.” (This is, of course, enthymematic; the suppressed premise, apparently, is that we should con-

²It might be possible to form a *theory* that tails are legs; for example, perhaps animals with tails evolved from animals with five functional legs (so that tails are a sort of vestigial leg). In that sort of case, there need be no change of subject. But that is not the case I have in mind; the case I have in mind is one in which I simply stipulate (without argument) that I will use “leg” to apply to tails.

tinue to use the word as we have in the past). Call arguments of this kind – arguments that concern the use of a word, such as “marriage” – *metalinguistic arguments*.

Although opponents of same-sex marriage do advance metalinguistic arguments, this is not their only argumentative tactic; for example, they sometimes claim that the purpose of marriage is to produce children. This is (or at least seems to be) an argument about marriage (rather than an argument about “marriage”.) Call arguments of this kind – arguments that concern something other than the use of a word – *first-order arguments*.

1.1 The Argument Argument

As a matter of methodology, we should look for an interpretation of these arguments and our responses to them that *makes sense of what we are doing*. Of course, there may be cases in which our arguments just don’t make sense – in which the parties to the debate are just confused or irrational. But in general, when a debate seems unconfused, we should try to understand what the participants in the debate are doing in a way that makes them unconfused. In this kind of case, we want to be able to say that the parties to a debate are rational, that their contributions to the debate make sense and contributes in a recognisable way to their aims. This suggests two requirements on our interpretation of a debate about a revisionary analysis:

1. It must give parties to the debate reasonably good epistemic status with respect to what they say. In many typical cases, this will mean that the parties to the debate know, or at least justifiably believe, what they say. In other cases, the relevant epistemic status may be different. For example, if one party to the debate is attempting a *reductio* of her opponent’s position, she should at least know or justifiably believe that what she says is a consequence of her opponent’s position.
2. It must make assertions relevant to the debate. When a party to the debate makes an assertion, that assertion should serve some purpose: for example, by giving evidence for her position or evidence against her opponent’s position.

I claim that no such interpretation of first-order arguments for or against revisionary analyses is available on the hypothesis that such analyses change the subject. Let’s call the opponent of same-sex marriage *O* and the proponent of same-sex marriage *P*; and consider *O*’s assertion of (1):

- (1) The purpose of marriage is to produce children.

Now on the hypothesis that *O* is changing the subject, there are two ways of interpreting this utterance, since there are two meanings of “marriage” at

issue in the conversation. On one interpretation, O is using “marriage” in her own preferred way, to pick out a status that it is by definition impossible for same-sex couples to enter into; call this status *marriage*₁. On the other, O is using “marriage” in P’s way, to pick out a status that same-sex couples can enjoy; call this status *marriage*₂.

I claim that neither interpretation can make sense of the debate. Suppose first that O is using “marriage” in her own preferred way. Then (let’s suppose) her assertion of (1) might well express a truth – a truth that O is in a position to know. So we are able to satisfy our first requirement on an interpretation.

But we are not in a position to satisfy the second requirement. For the proposition that the purpose of *marriage*₁ is to produce children simply does not bear on P’s position. Indeed, P may accept this claim; it is, after all, entirely consistent with the idea that “marriage” should be used to mean *marriage*₂. So on the hypothesis that O is using “marriage” to mean *marriage*₁, her utterance of (1) just looks irrelevant to the debate.

Suppose instead that O is using “marriage” in P’s way. This may seem like a surprising way for O to behave, but it could make sense for O to take on P’s way of speaking for the sake of argument, in order to illustrate its consequences. For example, if you try to stipulate that “tail” be applied to all and only round squares, I might object “But then there are no tails!” In this kind of situation, my first-order objection makes sense: I argue the way I do because I take that you will be able to see that what I am saying is a consequence of your view, and I think it likely that you will regard this consequence as unacceptable.

But this cannot be what is going on in O’s assertion of (1). Though (1), on this interpretation, is something that P will regard as unacceptable, it is no consequence of her view. On the contrary: on the hypothesis that P is using “marriage” to mean *marriage*₂, she should regard (1) as an obvious falsehood; moreover, it is an obvious falsehood that is obviously not entailed by her position. So on this interpretation, O’s utterance of (1) fails to meet either of our criteria: it is an obvious falsehood that O could neither reasonably believe, nor reasonably believe bears any relevance to P’s position, and it is hard to imagine any reason that it might make sense for O to attempt to use it in the debate.

So the most natural interpretations of the debate on the subject change view fail to make sense of O’s assertion. As the debate continues, we will find the same pattern: the subject-change view simply cannot make sense of many of the argumentative moves. So even if the subject-change view can explain the fact that disputes about revisionary analyses involve genuine disagreement, it cannot explain how these disputes are conducted: it cannot explain the way we argue. And – given that these arguments really do make

sense – this constitutes a powerful argument against the subject-change view, which I call *the argument argument*.

1.2 Two Responses

I now want to consider two ways of responding to the argument argument. The first response has it that we are adopting too limited a view of the possible purposes of the debate. True, O’s utterance of (1) will not convince anyone of a first-order proposition. But there are other cases in which an utterance of an apparently first-order sentence primarily functions to communicate metalinguistic information. For example, a person newly arrived at an Antarctic research station may ask with a shiver, “Is this cold?”; we could answer the question by pointing at a salient thermometer and saying, “This isn’t cold”. The thermometer is clearly visible: there is no mystery about the temperature, and no point in trying to communicate information about what the temperature is. Our utterance – despite the fact that it is not overtly metalinguistic – functions primarily to communicate information about how the word “cold” should be used around here (Sundell (2011); Plunkett & Sundell (2013), and see also Barker (2002, 2013)).

I agree that our utterances often communicate metalinguistic information in the way described; and I also agree that there is *something* metalinguistic going on in the dispute: if I say, “Two women can marry each other”, one of the purposes of my utterance is to communicate information about how I think the word “marriage” ought to be used. But two points need to be made. First, even if our dispute is in part metalinguistic, this does not entail (or even suggest) that the subject-change view is correct. It is entirely possible (indeed, common) for someone who uses a word with a particular meaning to communicate metalinguistic information about how that word ought to be used to someone who uses it with the very same meaning. (Plausibly, this happens every time I use a sentence to make an assertion: I assert the proposition expressed by my use of the sentence, but I also communicate that that sentence is correctly used in these circumstances to express a truth.)

Second, it is not obvious how the claim that our dispute is metalinguistic helps us make sense of O’s utterance of (1). This utterance does not bear in any obvious way on the claim that “marriage” should be used to express marriage₂. Even if we concede that O’s main intent in uttering (1) is to convince P of a metalinguistic proposition, it is very hard to imagine any way that uttering (1) could reasonably be thought to aid in this goal. So the observation that the debate is or could be in part metalinguistic simply does not help meet our second desideratum. Even if the debate is metalinguistic, given the subject-change view, advancing arguments of this kind just doesn’t make sense.

So the subject-change view cannot make good sense of the role an assertion of (1) plays in the debate. The second response to the argument is to question whether making sense of the debate should really be one of our aims. Perhaps arguments of this kind are just confused; perhaps an utterance of (1) really would express either the obvious falsehood that the purpose of marriage₂ is to produce children, or (more plausibly) the possibly true but irrelevant claim that the purpose of marriage₁ is to produce children.

If this is what is going on, then we would expect P (to the extent that she is unconfused) to reply simply by pointing out the fact that O's assertion of (1) is irrelevant. This is, after all, how a comparable debate would proceed in a case where it is uncontroversial that the subject has changed. For example, if I stipulate that "leg" is to apply to tails, and you object that horses have only four legs, to the extent that I am unconfused, I will not see your objection as a threat to my view and I will not try to refute it; the right response for me to make is something like, "As you are using the term, I agree; but that fact – admitted on all sides – does not bear on my claim about the preferred use of 'leg'."

I admit that it would be possible in many cases for P to respond in something like this way. (P need not actually agree that (1) is correct on anyone's view of "marriage", but she could at least argue that its correctness is irrelevant to her view: "That may be so, as you are using the term; but true or no, it does not bear on my claim".) I also agree that the subject-change view is plausible *if this kind of response is made and accepted*. (In my view, responding in this way goes toward making it the case that we are using the word "marriage" with different meanings; see section 5 below.) But this is not the way debates about same-sex marriage typically proceed, and it is certainly not the only way they can proceed. When confronted with (1), it is very natural (and apparently rational) for P to feel a tension between this claim and her view. Given this tension, a typical, apparently reasonable response is to dispute the claim (for example, by pointing out that elderly opposite-sex couples can be married, even though they cannot produce children).

Now perhaps the subject-change view can make some kind of sense of this response; the most natural suggestion is probably that P is adopting O's way of speaking in order to show that her view has absurd consequences, as in the "But then there are no tails" case. But this does not fully explain the phenomenon, because it does not explain P's reaction to O's utterance, and therefore does not explain why P responds in the way she does. In a typical case, P will feel that O has launched an attack on her view that puts her view under pressure; if O's claim (1) is true, P thinks, her view has a real problem. So she feels obliged to argue against (1).

This reaction – the feeling that what has been asserted is incompatible

with one's view, so that one must offer first-order resistance – is something that the subject-change view cannot explain. Of course, it is possible that this reaction is grounded in confusion – that in fact there is no such incompatibility, no such pressure on P's view. No one thinks that debates over same-sex marriage (for example) are conducted under conditions of ideal rationality, and it is certainly possible that we make mistakes. But, as I have emphasised, this reaction is utterly normal. Most participants in this kind of debate – even philosophers raised on a strict regimen of distinguishing use and mention – feel that assertions like (1) constitute a threat to their view, and a corresponding need to engage in first-order argument. The claim that we are all so confused about what is going on is to be resisted as long as some alternative interpretation is available (and, as I will go on to show, one is available once we reject the subject-change view).

I conclude that the replies to the argument fail: the subject-change view simply make sense of the way we argue for and against revisionary analyses; and even in those cases where the subject-change view has a story to tell, it often fails to explain the whole phenomenon (for example, because it does not explain the participants' own take on the debate, and therefore fails to explain their reactions).

2 Concepts

Let me try to relate the way of thinking about the threat of subject change that we have been developing to some things Haslanger says. As Haslanger puts the point, the worry is that when we have changed the subject “The term [...] expresses a different concept than it did” (2012c: 394). The picture seems to be that concepts are or correspond to word meanings, so that distinct words (that is, distinct word types) express different concepts.

This way of putting the point suggests that terms always or at least typically express only one concept at a time (in the sense for a given person at a given time, there is a single concept expressed by all typical uses of the term). But this would sit awkwardly with other things Haslanger says about concepts. Haslanger distinguishes between two analytic projects (in addition to her own “ameliorative project”). One involves collecting our ‘intuitive’ reactions to hypothetical cases and principles and trying to reach a set of principles that best systematises these reactions; Haslanger regards this as producing an analysis of our *manifest concept* – “the concept I thought I was guided by and saw myself as attempting to apply” (2012c: 388). The other that involves collecting and systematising empirical information about our actual classificatory judgments; Haslanger regards this as producing an analysis of our *operative concept* – “the concept that best captures the type we are concerned with” (2012c: 387).

One's manifest concept can be distinct from one's operative concept, as the following example from Jennifer Saul illustrates:

If we are investigating our manifest concept of democracy, then, we will probably arrive at a standard that requires (at least) elections that are free of voter intimidation, the counting of all ballots, equal access to polling places for all voters, and so on. [...] If we are investigating our operative concept of democracy, [...] we may well find that we apply the term far more broadly than our manifest concept would suggest. It may turn out that our operative concept of democracy requires only regular elections and that all adult citizens be formally permitted to vote. This is compatible with substantial voter intimidation, great variation in access to polling places, and ballots going uncounted. We could find out that this is our operative concept by noticing that in fact we apply the term 'democracy' even when we know that voters were intimidated, ballots went uncounted, and so on. (2006: 124)

So on the one hand, Haslanger thinks that change of subject involves using a word to express a different concept; but on the other, she thinks that words are often associated with multiple concepts. And while this combination of views isn't incoherent, it seems awkward. Fortunately, the difficulty is not deep; resolving it just requires some terminological regimentation.

I take beliefs (desires, etc.) to be mental representations. There is a sense in which different thinkers never share mental representations: my representations are in my head, and yours are in your head. But there is another sense in which you and I might believe the same thing; we do not share a belief *token*, but we have a belief of the same *type*. To say that a thinker possesses a particular concept is to say that she has mental representations of a certain type.

There are many ways of grouping token representations into types – type-individuating them, in the jargon – and different ways may be appropriate for different purposes. For example, it may make sense to individuate mental representations in an extremely fine-grained way, so that different thinkers rarely if ever share concepts, and one's concepts very often change over time, as one's beliefs and inferential dispositions change (e.g. Block (1986)). But there are other purposes for which it makes sense to individuate concepts in a less fine-grained way, so that a thinker can retain the same concept over time despite changes in belief, and different thinkers often share the same concepts.

One such purpose is tracking the kind of facts about certain kinds of agreement and disagreement, as well as closely related facts about successful

communication and argument. We have already conceded that some genuine disagreements do not involve parties asserting contradictory contents, so that there can be genuine disagreement despite difference in meaning; but we have also seen that disagreements that support a full range of first-order arguments are not like this. The most natural way to make sense of these disputes will involve individuating concepts in a less fine-grained way.

Here is a simple way of developing the picture. Suppose that you tell me, “Fido is a dog”. You are expressing a certain belief – a mental representation of a particular type – and that I understand you only if I entertain a mental representation of the same type. If I believe you, I am justified only if the mental representation I form is a representation of the same type you expressed (at least in typical cases). My utterance of “Fido is a dog” expresses or constitutes agreement with your utterance of the same-sounding sentence just in case the two utterances are expressing beliefs of the same type. And so on.

Similar sorts of points can be made even if we are paying attention to a single thinker. If at one time I say, “Fido is a dog”, but at a later time I say, “Fido is not a dog”, have I changed my mind? The answer depends on whether I reject at the later time a thought of the type that I had earlier accepted, and this amounts to whether “Fido” and “dog” (and perhaps also “is” and “a”) express the same concepts in both utterances. If I am inclined to say both “Fido is a dog” and “Fido is not a dog” at the same time, am I incoherent? Again, the answer depends on what types of belief I would express by these sentences.

If we want concepts to do this kind of work, we have to individuate them so that thinkers can have the same concepts over time and across changes of view. I may disagree with an interlocutor about many weighty matters, matters central to our discussion; I may despite this understand her, and agree with much of what she says. This requires that we share concepts of the same type. I may change some of my beliefs on a given topic while still retaining others, and again this requires individuating concepts in such a way that I can retain concepts of the same type despite my change of view.

Let’s not dispute whether there are other ways of treating these phenomena; I am happy to concede that there could be. There are, after all, many ways of individuating mental representations. Our purpose now is simply to lay out some stipulations about one useful way of using the word “concept”; I am stipulating that concepts are individuated in at least a somewhat coarse-grained way, and indicating (admittedly roughly) how the individuation is to go.

If we are individuating concepts in this way, then we should not think of the manifest “concept” and the operative “concept” as being distinct *concepts* (at least in typical cases). If they were, then we should see no conflict be-

tween our judgment that democracy requires that all citizens have equal access to polling places, and our judgement that (say) the USA is a democracy, even if we discover that many regions of the USA have enacted voter identification laws specifically designed to reduce the ability of certain groups to vote. Since my judgments about the theoretical principle (which brings to bear my “manifest concept”) and my judgment about the USA (which brings to bear my “operative concept”) involve different concepts, there is no inconsistency between them. We could simply rest content with our view.

It is quite possible that some cases fit this model. But the most natural understanding of the “democracy” case does not seem like this. If we judge that the USA is a democracy, but fails to meet some of the general conditions we judge necessary for democracy, this would generate a real feeling of conflict. The judgements just seem inconsistent. Anyone who finds themselves in this position and fails to take steps to reconcile her view – coming to believe that the USA is not a democracy, modifying the principle that democracy requires equal access, etc. – is *pro tanto* irrational.

So if the result of an “analysis of our manifest concept of democracy” and the result of an “analysis of our operative concept of democracy” are not analyses of distinct concepts, what are they? On the way of individuating concepts that we are now considering, we should think of these not as analyses of different concepts, but *as competing proposed analyses of one and the same concept*. To get clear on why this makes sense, we will need to turn to the notion of analysis.

3 Analyticity

Haslanger proposes (W) as an analysis of “woman”, and I take it that this proposal entails that (W) is an analytic truth.³ But what exactly does it mean for a claim to be analytic, or to be proposed as an analysis?

Analysis and *analyticity* came under heavy scrutiny after Quine’s (1953; 1966a; 1966b) critique. Traditional accounts of analyticity combined metaphysical and epistemological elements, usefully distinguished by Boghossian as follows:

Epistemic analyticity A sentence S is epistemically analytic iff “mere grasp of S’s meaning by T suffice[s] for T’s being justified in holding S true.” (334)

Metaphysical analyticity A sentence S is metaphysically analytic if “in some appropriate sense, it owes its truth-value completely to its meaning, and not at all to ‘the facts’” (1997: 334)

³I will also extend the term “analytic” to obvious entailments of analyses, so that (for example) if (W) is an analysis, then it is an analytic truth that women are oppressed).

It's hard to see how metaphysical analyticity could be relevant to revisionary analyses (like Haslanger's), and that's a good thing: with Boghossian and most others, I take Quine's attack on metaphysical analyticity to have been decisive. Our meaning that Hesperus is Hesperus by "Hesperus is Hesperus" doesn't make it the case that Hesperus is Hesperus; what makes it the case that Hesperus is Hesperus is a thought- and language-independent fact.

Boghossian himself defends epistemic analyticity, and although this cannot be the right account of what it is to be analytic in the sense relevant to our debate – manifest analyses, for example, will typically be known on the basis of empirical investigation, not justifiably believable by anyone who grasps them – investigating Boghossian's case will help reveal another element in notion of analyticity that will provide a framework for understanding revisionary analyses like Haslanger's.

The foundation of Boghossian's defense of epistemic analyticity is a view about what it takes to possess certain concepts, such as the concepts of geometry: "grasp of the indefinables of geometry consists precisely in the adoption of one set of truths involving them" (1997: 348). More generally:

Adoption Grasp A term t is *grasped by adoption* (or *a-grasped*) just in case there is some set of truths S which is such that one understands t iff one accepts each member of S .

Boghossian's main interest is in logic; there, his idea is that (for example) one possesses the concept of conjunction just in case one accepts principles like, "If p and q , then p ". He claims that this account of grasp "generates" (1997: 348) the following account of how logical constants get their meanings:

It is by arbitrarily stipulating that certain sentences of logic are to be true, or that certain inferences are to be valid, that we attach a meaning to the logical constants. More specifically, a particular constant means that logical object, if any, which would make valid a specified set of sentences and/or inferences involving it. (1997: 348)

As before, we might want to consider views on which the meanings of other sorts of terms are fixed by implicit definition. We can generalise the proposal in the following way:

Implicit Definition A term t is *implicitly defined* iff we attach a meaning to t by arbitrarily stipulating that certain sentences involving t are to be true. More specifically, t has that meaning, if any, which would make true a specified set of sentences involving it.

If *Implicit Definition* is true, then there is a sense in which certain sentences are true by convention. But Boghossian argues convincingly that there is no threat that implicit definitions will result in conventional truths of the sort that Quine's arguments show to be problematic. Perhaps the simplest way to see the point is on the view that propositions – the meanings of sentences – are the primary truth bearers, so that a sentence is true only if it expresses a true proposition. We can imagine propositions having their truth values fixed antecedently to any human linguistic activity. Stipulation makes it the case that a particular sentence expresses a particular proposition; it does not make it the case that the proposition expressed by the sentence is true.

Consider an example. Suppose that the length of a certain rod r is l , so that the proposition that the length of r is l is true (and true independently of human thought and linguistic activity). Then we can introduce the term “meter” by stipulating: “The length of r is one meter”. Given the meanings of the other words in the sentence (“The”, “length”, “of”, etc.), the syntactic structure of the sentence, and the composition rules of the language, there are a variety of propositions that the sentence might express, depending on what “meter” means: the proposition that the length of r is h , the proposition that the length of r is l , the proposition that the length of r is m , etc. By stipulating that “The length of r is one meter” expresses a truth, we fix on a particular candidate – the only true candidate in the vicinity: the proposition that the length of r is l . This in turn fixes the meaning of “meter”: given the extant meanings of “The”, “length”, etc., and the facts about syntax and composition, “The length of r is one meter” expresses the proposition that the length of r is l only if “meter” picks out l .

It is clear in this example that our linguistic stipulations have not made true any non-linguistic facts. The facts about length were fixed prior to our stipulation. Something like this model of *Implicit Definition* provides the foundation of what we will develop in what follows. But in order to avail ourselves of this foundation, we will need to rethink what it takes to grasp the meaning of an implicitly defined term, as well as the notion of stipulation.

3.1 Implicit Definition and Adoption Grasp

Implicit Definition would fit neatly with *Adoption Grasp*. But the two are in principle separable; an *Implicit Definition*-style account of how words get their meaning is compatible with any number of accounts of how meaning can be grasped.

Here is an example. Consider Kripke's (1980) causal picture of meaning: a word is introduced at an initial baptism, and later uses get have their meaning in virtue of their connection to this initial baptism. A typical instance of such a baptism might involve pointing at a certain baby and saying

something like, “I hereby name this baby ‘Ansel’”. This is an arbitrary stipulation that attaches a meaning to the word “Ansel”. But it does not exactly fit the template described in *Implicit Definition*, since we are not stipulating that any sentence is true. So let’s reformulate our initial baptism: “I hereby stipulate that the following sentence is true: ‘The baby in room 110, Ninewells Hospital, at 10:00am GMT 25th April 2013, is Ansel.’” Plausibly, then, the meaning of “Ansel” has been fixed by arbitrarily stipulating that a certain sentence involving “Ansel” is true. So “Ansel” is implicitly defined.

But note that this does not entail that *Adoption Grasp* is true of “Ansel”. On the contrary: Kripke’s view has it that later speakers grasp “Ansel” in virtue of their causal connection to the initial baptism. He specifically rejects the idea that speakers’ competence with a term depends on entertaining the stipulation with which a word is introduced. So a word can be implicitly defined but not a-grasped.

It might be objected at this point that on Kripke’s story, the meaning of “Ansel” as it is used by later speakers is *not* determined solely by the stipulation. Instead, it is determined by the stipulation *and the later speakers’ causal connection to that stipulation*. If that is right, then “Ansel” is not implicitly defined.

This is a fair point. In general, we should allow that a stipulation can play a significant role in fixing the meaning of a word, while also allowing that other factors may also play a role. So we should allow for the possibility that stipulation *partially* fixes the meaning of a word, while other factors also play a role:

Partial Implicit Definition A term *t* is *p-implicitly defined* iff *t* has the meaning it does at least in part in virtue to an arbitrary stipulation that certain sentences involving *t* are to be true. More specifically, if *t* has a meaning, it has a meaning which would make true a specified set of sentences involving it.

What we have established so far, then, is that partial implicit definition and adoption grasp can come apart, at least in principle; a term can be *p-implicitly defined* but not *a-grasped*. And this is a good thing: Boghossian’s defence of epistemic analyticity has come under sustained attack, notably in the work of Timothy Williamson (2006; 2007). Williamson points out that experts with idiosyncratic theoretical views can dissent from even the most basic and plausible truths involving a word. For example, Williamson suggests that a thinker might reject an elementary logical truth like “Every vixen is a vixen” – an instance of a principle acceptance of which might well be held to be constitutive of understanding of “every” – because she maintains that “every” has existential import, and denies that there are vixens. Similarly, Williamson points to logicians who reject modus ponens. Such

thinkers seem to be counterexamples to the claim that logical terminology is a-grasped: these experts seem to understand “every” and “if”, while rejecting principles in which this understanding putatively consists.

There have been, of course, a variety of attempts to resist Williamson’s examples (e.g., Boghossian (2011)); but I, for one, continue to find the examples convincing. But there is need for care. Williamson’s examples refute the claim that Adoption Grasp is true of words like “every” and “if”. They do not refute the claim that these words are p-implicitly defined. For example, one could imagine a Kripke-style view on which “if” means what it does because of our causal connection to an ancient stipulation. (I do not claim that this is a plausible story, only that it is unscathed by Williamson’s examples.)

The crucial lesson of Kripke’s causal theory, for our purposes, is just that the stipulation that p-implicitly defines a word need not be grasped (known, believed, entertained) by competent users of the term. The meanings of my words can be fixed by stipulations that are spatially and temporally distant from me – stipulations that I am in no position to know of.

3.2 From Implicit Definition to Metasemantic Analyticity

A further thread in Quine’s critique is that the notion of analyticity is unclear, and that attempts to clarify or explain it are either circular or rely on equally unclear notions of meaning or synonymy. Carnap’s response was to offer an account that suited the formal languages that are the primary focus of his *Meaning and Necessity* (1956).⁴ In developing such a formal language, we may stipulate that certain sentences are to be regarded as *meaning postulates*. Carnap’s idea is that a sentence is analytic just in case it is entailed by the set of meaning postulates.

Carnap’s use of meaning postulates is meant to ground *metaphysical* analyticities: analyticity is understood as truth in virtue of meaning. But meaning postulates were also taken up by proponents of Montague grammar in the 1970s and 80s (see e.g. Montague (1974: 212-3), Dowty et al. (1981: 224-5)). In these works, the metaphysical ambitions have largely dropped away. The role of meaning postulates is to help to fix the meaning of certain terms. The basic idea is simple. In Montague grammar, expressions are assigned semantic values relative to a *model*. To say that a certain sentence is a meaning postulate is to say that we will only consider models that make that sentence true. Since logical truth is defined as truth at every model, meaning postulates will be logical truths.

Let’s consider a simplified example. Suppose that the semantic values

⁴Carnap characterises his project here as *explication*, which he sees as involving the replacement of one concept by another; but (of course) on my view we need not accept this characterisation.

of one-place predicates are sets of individuals (the individuals of which the predicate is true), so that a sentence like (2) is true if and only if the semantic value of “bachelor” is a subset of the semantic value of “unmarried”:

(2) For all x , if x is a bachelor then x is unmarried

By letting (2) be a meaning postulate, we restrict our attention to models relative to which the sentence is true, so models relative to which semantic value of “bachelor” is a subset of the semantic value of “unmarried”. The idea is that although this does not completely determine the meanings of “bachelor” and “unmarried”, it puts a constraint on, and so partially fixes, them.

The traditional metaphysical notion of analyticity had it that some sentences were true in virtue of their meanings. The use of meaning postulates that we are now considering flips this on its head: some expressions have the meanings they do in virtue of the truth of certain claims. Instead of *truth in virtue of meaning*, we have *meaning in virtue of truth*: by stipulating that we will only consider interpretations of our words that make a certain sentence true, we have put a constraint on, hence partially determined, the meanings of the words that compose that sentence.

It is no accident if meaning postulates remind one of partial implicit definitions. They play an extremely similar role; in effect, meaning postulates p-implicitly define the words that compose them. I suggest that sentences that p-implicitly define (some of) the words that compose them deserve to be thought of as *analytic* in the following sense:

Metasemantic Analyticity A sentence is metasemantically analytic with respect to a word (or a use of a word) iff the meaning of that (use of the) word is partially fixed by the stipulation that the sentence is to be true (i.e., iff the sentence p-implicitly defines the word).

Metasemantic analyticity has, perhaps, received less attention in the literature on analyticity than epistemic and metaphysical analyticity, in part because it plays a less important role in the metaphysical and epistemological debates to which analyticity has often been turned. (For example, there is no clear role for metasemantic analyticity in giving an account of metaphysical modality, or of apriori knowledge.) But there is nonetheless a notable tradition of using the word “analyticity” to pick out metasemantic analyticity. Meaning postulates are, after all, what is analytic on Carnap’s account, and one of their primary roles (the role picked up on by Montagovians) is to fix meaning. And metasemantic analyticity is the notion most relevant to the many discussions of analyticity in the work of Fodor and his collaborators. To take just one representative example, Fodor and Pylyshyn write, “analytic beliefs can’t be revised without changing the content of (some or

all of) their conceptual constituents; that is, they can't be changed without equivocating." (2015: 57). This passage only makes sense if what is at issue is something like metasemantic analyticity (applied to beliefs and the representations that compose them rather than to natural language sentences and words): the idea is that some of our beliefs *determine the content* of their representational constituents, so that changing these beliefs would correspondingly change the content of the constituents. What is at issue is how the contents of certain mental representations are determined: that is, the metasemantics of those representations.

So there is an established pattern of using "analyticity" to pick out metasemantic analyticities in the literature. On my view, proposed revisionary analyses, such as Haslanger's (W), should be seen as attempts at metasemantic analysis.

4 Timing

Let's pause to take stock. I began by showing that the subject-change view cannot explain how we argue for revisionary analyses, and I have now developed the notion of metasemantic analyticity, which I claim to be relevant to revisionary analysis. But at this point it may be unclear how my rejection of the subject-change view can be reconciled with my embrace of metasemantic analyticity. For isn't the proponent of the subject-change model precisely endorsing the claim that revisionary analyses are a kind of stipulation that introduces a new meaning for a word (and hence changes the subject)? After all, what makes it the case that post-revision uses of the word have a different meaning? Surely it is the revisionary analysis – the stipulation – but that is just to say that the revisionary analysis partially fixes the meaning of the word and so is metasemantically analytic.

I propose to grant that revisionary analyses – *if successful* – are metasemantically analytic. (We will return momentarily to the question of what makes a revisionary analysis successful when it is.) But we can grant this without embracing the subject-change view. On the *subject continuity model* that I advocate, there is no new meaning; a successful stipulation fixes the meaning of the word as it was used all along.

It must be admitted that the subject-change view would fit well with our standard conception of stipulation. Stipulated definitions are typically thought of as coming when a word is first introduced. But what motivates this conception? One possible motivation comes from Adoption Grasp: if users of a word must grasp the stipulated definition, then the stipulation must occur before the word is used. But as we have seen, there are popular and plausible views on which p-implicit definition (and hence metasemantic analyticity) can be prised apart from Adoption Grasp. On these views,

the stipulation that p-implicitly defines a word can be made by an unknown person in the distant past, and the stipulation need not be known by current users of the word as long as they are causally connected to it in the right way. So users of a word need not accept the sentence that p-implicitly defines that word; a word I use can get its meaning (at least in part) from a stipulation made by someone else, of which I am not even aware.

The subject continuity view is a development of this idea. Although the Kripkean causal picture (for example) rejects the view that users of a term must grasp the stipulation that introduced, it retains the idea that stipulations that fix the meanings of my words must be temporally prior to my use of those words: *first* someone stipulates, and that gives *later* uses of the words meaning. In my view, the key to understanding revisionary analyses is that they involve a stipulation that (partially) fixes the meaning of *prior* uses of the word: in this kind of case, *first* we go about using a word, then *later* we make the stipulation that gives meaning.⁵

On this view, in typical cases of revisionary analysis there is no change of meaning and no change of subject. Words are used univocally throughout a discourse; each word expresses the same concept throughout, even when one party proposes a revisionary analysis. What happens when a proposed revisionary analysis is successful is that stipulations that fix the meaning of a word (throughout the discourse) have emerged late in the discourse.

The subject continuity view may seem like a radical departure from familiar conceptions of stipulation and analyticity,⁶ but it is inevitable given the conclusions that we have already reached. The argument argument gives us reason to reject the subject-change view: we should maintain that words are univocal – that no meaning-change occurs and no new concepts are introduced – before and after revisionary analysis. But we have also granted that successful revisionary analyses are in fact (metasemantic) analyses; they partially fix the meanings of words. Since they do not introduce new meanings or new words, they must fix the meanings that words have had all along. And this is just the core claim of the subject continuity view.

5 Successful Analysis

But the subject continuity view may seem to leave us with a serious problem. On traditional views of stipulation, such as the subject-change view – views on which stipulation precedes use – in typical cases there will be no

⁵The general term for views on which the meaning of an expression or the content of a thought can depend on facts about later times is *temporal externalism*. For discussion and defence, see Jackman (1999, 2005), Ball (forthcoming).

⁶Perhaps it constitutes a somewhat revisionary analysis of stipulation and of analyticity. (Perhaps pointing this out is too cute.)

serious problem in determining whether a given stipulation was successful: if we make a stipulation, and then go on to use the word stipulatively introduced, the stipulation was successful.⁷ But on the subject continuity view, the picture is more complex. We may have multiple parties to a dispute, each proposing a competing revisionary analysis, all of which purport to fix the meaning of the same set of uses of a term. What makes it the case that one of these proposed analyses fixes the meaning, while the others do not?

In order to answer this question, I want to begin by considering a metase-mantic story that some philosophers have taken to support the subject-change view that we have just rejected: the view that meaning is use. As this slogan is typically interpreted in this debate, the view that meaning is use amounts to the view that the extension of a term as used by a speaker is determined by that speaker's dispositions to apply the term and to withhold application of the term. Call this *meaning is use now* (MIUN).⁸

MIUN suggests that revisionary analysts are changing the subject, since (at least to the extent their dispositions to apply the term are consistent with their proposed analysis) revisionary analysts will exhibit quite a different pattern of applying and withholding the term at issue; for example, the proponent of same-sex marriage will apply the term "marriage" to same-sex unions while her opponent will not, and Haslanger will withhold "woman" from females who are not subordinated (if such there be).

The problem with MIUN is similar to the problem with the subject-change view: it cannot make sense of our argumentative practice. We have already seen that the subject-change view has trouble making sense of the full range of first-order argument, and to the extent that MIUN is committed to the subject-change view, it will share this problem. I now want to emphasise an additional problem with the view that meaning is use now: it cannot make sense of how we respond to arguments. In particular, it cannot make sense of apparently rational and correct change of view. We may be convinced by a revisionary analysis. Suppose that I once believed that same-sex marriage was impossible (and said things like, "A man can only marry a woman and vice-versa"), but I was convinced by some argument that same-sex marriage is possible after all. To keep the case simple, suppose that this last belief is stable (and so not overturned by further argument), and that most or all of my fellow speakers come to share my view.

How would I describe my change of view? One thing that I would take to have changed is my metalinguistic beliefs; I use to take "Same-sex couples can marry" to express a falsehood, but I came to believe that it expresses a

⁷Perhaps we need to make some further allowances to prevent the stipulative introduction of "tonk" and the like, but these complexities are not relevant to our present concerns.

⁸See Sundell (2012) for a defence of this kind of view from a proponent of the idea that revisionary analysis requires meaning change.

truth. But this is neither the only, nor the most central, belief that I would take to have changed; I would also take myself to have believed that a man can only marry a woman, and to have come to reject this belief and to believe instead that same-sex couples can marry. (I might say things like, “I used to think that two women could not marry each other, but now I see that I was wrong.”) That is, I would take my first-order views about marriage to have changed.

In section 1.1, I suggested that as a matter of methodology, we should look for an interpretation of our debates about revisionary analysis that makes sense of what participants in the debate are doing. The same is true here: we should look for an account of the debates that makes sense of our informed judgments about them. Of course, we are not infallible about what we believed and meant. But other things equal – in cases when we are not obviously confused, irrational, suffering from memory failure, and so on – an interpretation that makes sense of our judgments is to be preferred to an interpretation that does not. And since all that is at issue here is remembering my past view, and judging whether my current view is inconsistent with them, the most obvious way an interpretation can make sense of my judgment is by making it correct. (Perhaps this is the only way to make sense of it: since nothing is wrong with my memory, it seems hard to explain how I could fall into error.)

This is something that MIUN cannot deliver. According to MIUN, my dispositions to apply “marriage” before I heard of the revisionary analysis make it the case that “marriage” as I used it picked out a relation that men can stand in only to women and vice-versa, while my dispositions after I accept the revisionary analysis make it the case that “marriage” as I use it picks out a relation that men can stand in to other men. So MIUN cannot make me right when I say, “I used to think that two women could not marry each other, but now I see that I was wrong”. As I use “marriage” at the time of this utterance, it isn’t true that I used to think that two women could not “marry” each other, and what I used to think wasn’t wrong. So in order to make me correct, what is wanted is a view that predicts not that the meaning of “marriage” has changed, but that what I meant all along is the relation that same-sex couples can stand in. In this case, the revisionary analysis has succeeded: claims like “Two men can marry each other” are metasemantically analytic of “marry” as I used it all along.

Of course, there are cases in which MIUN delivers more plausible results. But even in these cases, it gets the right answer for the wrong reason. Suppose I accept that being a woman is a matter of biological fact, and say things like “Necessarily, a person is a woman iff she has two X chromosomes” – until a philosopher puts forward (W) as an analysis of “woman”. I consider the analysis, and despite some intriguing arguments, I decide that the bal-

ance of considerations pull against accepting it. So I judge that (W) is false, and that “Necessarily, a person is a woman iff she has two X chromosomes” is true. To keep the case simple, suppose that my rejection of (W) and my continued acceptance of are stable (and so not overturned by further argument), and that most or all of my fellow speakers share my views (so also reject (W) and accept the biological account of sex). And suppose further that in the end even the proponent of (W) changes her mind and becomes convinced that (W) is false.

Again, on the assumption that I am being reasonable – not unduly stubborn, not neglecting relevant evidence, etc. – then my judgment is something we should aim to make sense of, and the obvious way to do this is by making it correct. So we want a view that predicts that in this case, “woman” as I use it picks out a biological concept that makes no reference to social status. (W) purports to be an analysis of “woman”, but in the end fails to do metasemantic work.

MIUN delivers this result; but we have already seen that it fails elsewhere, so that we need to look for a different view. One suggestion would be to appeal to a wider range of dispositions: not only dispositions to apply the term, but also dispositions to respond to new information and so on. For example, perhaps the difference between the described “marriage” case and the described “woman” case turns on my dispositions to respond to revisionary analyses.

This strategy may make sense in some cases. But in general, our dispositions are messy. I may be disposed to respond favourably to some arguments (in some moods, when presented by some interlocutors, in some tones of voice) but not others. I am not simply disposed to come to believe that same-sex marriage is possible no matter what: whether I come to believe this will depend on exactly which arguments I encounter, and in which circumstances. My dispositions just don’t fix my meaning – or at least, not until they are activated in particular circumstances. (That is why accepting a view can count as a stipulation, even though it is also a response to an argument.)

The obvious difference between the success of the revisionary analysis of “marriage”, and the failure of the revisionary analysis of “woman” (in the described cases) is that the successful analysis is *accepted* and the failed analysis is *rejected*. I suggest that this obvious difference is what does the metasemantic work. At least in kind of simplified case we are considering here (in which the debate is resolved), what makes a given putative metasemantic analysis successful is that it is accepted.

Further, I suggest that once we have given up the subject-change view (and so accepted that there a single meaning/concept throughout the discourse), this is an extremely natural conclusion. We seem to have two possible accounts of what is going on in simple cases like the ones we are dis-

cussing:

Beginning of Debate The facts that fix meaning – including the stipulations and beliefs that are metasemantic analyticities – must be in place at the beginning of a discourse. They fix the meanings of our terms and the content of our concepts throughout the rest of the inquiry. So revisionary analyses always fail; if, at the beginning of inquiry, we accept “Marriage is the union of one man and one woman”, then we cannot successfully advocate same-sex “marriage” without changing the subject.

End of Debate The facts that fix meaning can be determined by the discourse. So revisionary analyses can succeed; if everyone comes to accept on reasonable grounds that same-sex marriage is possible, this goes toward making it the case that “marriage” picks out a relation that can hold between two men or two women.

But surely the End of Debate picture is much to be preferred. It seems absurdly uncharitable – perverse, really – to maintain that if everyone goes from rejecting the possibility of same-sex marriage to accepting it (and regarding their prior first-order views as incorrect), they are making a mistake, even if this transition is grounded in apparently reasonable argument. It looks like we have a choice of preferring our views *prior to inquiry* or our views *after inquiry*: but surely (other things equal) our views after inquiry are to be preferred. (After all, it is only then that we have had a chance to consider the evidence.)

Of course, this is not to say that discussion can never go awry, or that we can never be mistaken or confused. It is only to say that the views of people in a better epistemic position are *ceteris paribus* to be preferred to the views of those in a worse epistemic position.

5.1 Agreeing that a Dispute is Verbal, and Other Complications

We have made simplifying assumptions about the cases we have discussed: that everyone comes to agree, so that the debate is settled once and for all. And of course in many cases this kind of assumption will be entirely unrealistic. What predictions do we make about more realistic cases? One kind of complication that we can easily make sense of. Some debates end not in agreement, but in acknowledgement that the disputants were talking past each other, so that the debate was merely verbal. We already mentioned that opponents of same-sex marriage sometimes claim that proponents are trying to redefine “marriage”. Typically (and reasonably, on the view we are developing) proponents resist this charge. But proponents might also (and again, reasonably) agree that they are using “marriage” differently, and proceed to dispute about how the word should be used. In this case, too, I think

that the considered judgments of the parties to the conversation should be taken seriously: if everyone comes to the conclusion that the subject has changed, that can make it the case that the subject has in fact changed.

The case where the (first-order) debate is resolved (or perhaps dissolved) when we agree that we are talking past each other reveals a final data point worth attending to. In this case, we may re-evaluate our earlier reactions to arguments and judge some of them to have been confused. To revisit our example from section 1.1, if P has decided that the dispute about same-sex “marriage” is merely verbal, then P should also decide that O’s assertion of (1) did not put any pressure on her view. Instead of arguing against it, she should have rejected it as irrelevant.

As before, an interpretation of the debate should seek to vindicate this considered judgement. And that is exactly the result we get when we take our end-of-debate claims to be metasemantically analytic. When we are evaluating (1), for example, there are three possibilities to consider:

1. If O convinces P – so that O and P end up agreeing that same-sex marriage is by definition impossible – then the assertion of (1) will seem to have been a good argument: an argument that bears on the issue, and expresses a reasonable belief. We vindicate this appearance: the assertion of (1) bears on the issue because it is inconsistent with the view that P held, and was a reasonable commitment for someone with O’s view. (If P and O end up agreeing on (1), it too may end up as a metasemantic analyticity, hence true.)
2. If P convinces O – so that O and P end up agreeing that same-sex marriage is possible – then the argument will seem flawed. If the assertion were true, it would have been relevant (since it is incompatible with P’s position), and perhaps O had good reason to think it true; but a fuller consideration of the evidence revealed it to be false. We vindicate this appearance: since there is no meaning change, the assertion will be relevant – it is inconsistent with P’s position – but since P’s position is correct, the assertion will turn out to be false.
3. If P and O agree that they are talking past each other, then the argument will seem irrelevant: either a truth that does not bear on P’s position, or a falsehood (and an obvious falsehood, given their current evidence). We vindicate this appearance: since in this scenario the word “marriage” is used with two distinct meanings in the conversation, there are two relevant interpretations of (1), one of which is an irrelevant truth and the other of which is an relevant falsehood.

But many typical disputes about revisionary analyses – at least in philosophy – do not resolve at all: we do not agree, and we do not agree that our

dispute is verbal either. And in this kind of case the considerations we have given so far make no definite prediction. The basic moral of our story is that facts about our future use can play a meaning fixing role – can be metasemantically analytic – with respect to our present use. Exactly how meaning gets determined in any actual case will no doubt be a complicated matter that requires weighing a variety of factors: facts about past, present, and future dispositions and uses, as well as (probably) other sorts of facts about the social and physical environment (which may include facts about metaphysical naturalness and so on). The view we have developed does not answer every metasemantic question. It only introduces a new factor to be considered – albeit a factor that is decisive in many significant cases.

6 Conclusion

This paper has defended the following claims:

1. Revisionary analysis does not typically involve change of meaning or the introduction of new concepts.
2. Revisionary analyses can be seen as a kind of (attempted) stipulation, and so as attempts at giving metasemantic analyticities; but if they are successful, they fix the meaning of words as we have always used them (even before the stipulation took place).

These claims together constitute a view that repudiates the idea that philosophy is purely descriptive – in making judgments, we are also making our meanings and concepts, not merely “leaving everything as it is” – but at the same time repudiates the idea that we are changing the subject, or choosing between different concepts. (Making judgments, forming beliefs, choosing theories? Yes. Conceptual engineering or conceptual ethics (understood as something other than this)? No.)

Let me conclude by returning to Jackson’s worry: does the stipulation that “a free action is one such that the agent would have done otherwise if he or she had chosen to” result in a change of subject that turns “interesting philosophical debates into easy exercises in deductions from stipulative definitions together with accepted facts”? Jackson’s worry is not completely unfounded, since it is *possible* that this stipulation would change the subject; it would change the subject if the parties to the debate treat it as changing the subject, for example by refusing to engage with first-order argument for and against it. But philosophers who make this kind of claim do not usually have this attitude. It is perfectly possible to introduce a particular account of free action, recognising it as a stipulation that will have a metasemantic effect if accepted, but still engaging in interesting philosophical debates. This is exactly what proponents of revisionary analyses usually do.

So could women be analytically oppressed? Yes: (W) could be metasemantically analytic. Are they analytically oppressed? It is up to proponents of (W) to convince the rest of us.

References

- Ball, D. (forthcoming). Relativism, metasemantics, and the future. *Inquiry*, forthcoming in a special issue edited by Henry Jackman.
- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25(1), 1–36.
- Barker, C. (2013). Negotiating taste. *Inquiry*, 56(2-3), 240–257.
- Block, N. (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10, 615–678.
- Boghossian, P. (1997). Analyticity. In B. Hale & C. Wright (Eds.), *A Companion to Philosophy of Language* (pp. 331–368). Oxford: Blackwell.
- Boghossian, P. (2011). Williamson on the a priori and the analytic. *Philosophy and Phenomenological Research*, 82, 488–497.
- Burgess, A. & Plunkett, D. (2013a). Conceptual ethics i. *Philosophy Compass*, 8(12), 1091–1101.
- Burgess, A. & Plunkett, D. (2013b). Conceptual ethics ii. *Philosophy Compass*, 8(12), 1102–1110.
- Carnap, R. (1956). *Meaning and Necessity, second edition*. Chicago: University of Chicago Press.
- Dowty, D. R., Wall, R. E., & Peters, S. (1981). *Introduction to Montague Semantics*. Dordrecht: Kluwer.
- Fodor, J. A. & Pylyshyn, Z. W. (2015). *Minds without Meanings: An Essay on the Content of Concepts*. Cambridge, MA: The MIT Press.
- Haslanger, S. (2012a). Gender and race: (what) are they? (what) do we want them to be? In Haslanger (2012b), (pp. 221–247).
- Haslanger, S. (2012b). *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.
- Haslanger, S. (2012c). What good are our intuitions? philosophical analysis and social kinds. In Haslanger (2012b), (pp. 381–405).

- Jackman, H. (1999). We live forwards but understand backwards: Linguistic practices and future behavior. *Pacific Philosophical Quarterly*, 80, 157–177.
- Jackman, H. (2005). Temporal externalism, deference, and our ordinary linguistic practice. *Pacific Philosophical Quarterly*, 86, 365–380.
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. New York: Oxford University Press.
- Kripke, S. A. (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Montague, R. (1974). English as a formal language. In *Formal Philosophy* (pp. 188–221). New Haven: Yale University Press.
- Plunkett, D. & Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13(23).
- Quine, W. V. O. (1953). Two dogmas of empiricism. In *From a Logical Point of View* (pp. 20–46). Cambridge, MA: Harvard University Press.
- Quine, W. V. O. (1966a). Carnap and logical truth. In Quine (1966c), (pp. 100–125).
- Quine, W. V. O. (1966b). Truth by convention. In Quine (1966c), (pp. 70–99).
- Quine, W. V. O. (1966c). *The Ways of Paradox and Other Essays*. New York: Random House.
- Saul, J. (2006). Philosophical analysis and social kinds: Gender and race. *Proceedings of the Aristotelian Society*, 106, 119–143.
- Scharp, K. (2013). *Replacing Truth*. Oxford: Oxford University Press.
- Sundell, T. (2011). Disagreements about taste. *Philosophical Studies*, 155(2), 267–288.
- Sundell, T. (2012). Disagreement, error, and an alternative to reference magnetism. *Australasian Journal of Philosophy*, 90(4), 743–759.
- Thomasson, A. L. (2016). Metaphysical disputes and metalinguistic negotiation. *Analytic Philosophy*, 57(3).
- Williamson, T. (2006). Conceptual truth. *Aristotelian Society Supplementary Volume*, 80, 1–41.
- Williamson, T. (2007). *The Philosophy of Philosophy*. Oxford: Blackwell.
- Wittgenstein, L. (1953). *Philosophical Investigations*. Oxford: Blackwell.