Aporia

Undergraduate Journal of the St Andrews Philosophy Society
Volume XVIII — Issue No. 2

Aporia *is funded by the University of St Andrews Philosophy Society, which receives funds from the University of St Andrews Department of Philosophy, the University of St Andrews Students' Association, and independent benefactors.*

Contents

Acknowledgements

We would like to thank the individuals and institutions that continue to support the journal. Special thanks go to Kirsty Hardwick, Oscar Westerblad, Ethan Landes, Quentin Pharr, and Dr Patrick Greenough for their helpful contribution to the review process, and Jackie Proven and Kyle Brady from the Online Journal Hosting Service at the University of St Andrews for their technical support. Most of all, we are thankful to everyone who submitted their papers for our consideration.

Kimon Sourlas-Kotzamanis
Editor-in-chief

Spencer Albert
Deputy editor

The Normativity Problem as a Serious Obstacle to Modelling Gender

Weston Richey*
westonleorichey@gmail.com

Abstract

In this paper, I explore Sally Haslanger's (2000) proposed approach to modelling gender which she intends to overcome several problems for such a project. I specifically focus on what Haslanger calls the normativity problem, in which definitions meant to overcome oppression only reinforce oppressive norms. I argue that the normativity problem is a serious one for defining gender and that Haslanger does not successfully overcome it with her definitions of *man* and *woman*. In §§1 and 2, I offer background for and explain her account of the problem before offering my own formal reconstruction of it as what I call the normativity argument that (a) we ought not marginalise individuals in our defining of social categories, (b) definitions encouraging normative behaviour do this, and (c) any model of gender encourages such behaviour. In §3 I then give an account of her proposed definitions of *man* and *woman* along with her theoretical objections to the normativity argument—suggesting that only certain kinds of marginalisation are undesirable within the constraints of a particular feminist project and that her definitions do not encourage normative behaviour. I then offer my responses to her objections in §4, suggesting that her definitions are normative and do marginalise in a way incongruous even with her particular feminist project. Before concluding, I briefly discuss in §5 where my criticisms of Haslanger's approach to defining gender fit into some existing criticisms, in order to give my position an even clearer shape. This paper concludes in §6 by sketching some possible ways forward in the philosophy of gender responding to this problem.

1 Introduction

An important project in the philosophy of gender is perhaps the most basic: that of developing a working model for gender. What does being a woman, or a man, or being of another gender or not being of one at all, consist in? And even prior to that question: how ought we go about developing such a model or definition? Sally Haslanger, at one juncture in her paper

*Weston Richey just received their BA in Philosophy and English and American Literature from New York University (Spring 2018). They are now an MFA candidate in Creative Writing at Rutgers University.

'Gender and Race' (2000), considers a problem that serves as an obstacle to developing a model of gender: what she calls the *normativity problem*. In this paper, I want to take that problem seriously, and suggest that it is indeed a serious obstacle for defining gender. I will do this by example, illustrating how the normativity problem poses a challenge for Haslanger's own definition of *man* and *woman*. First, I will explain the normativity problem for defining gender by reconstructing it as an argument—which I will call the *normativity argument*, to distinguish it from its broader, less-detailed formulation as a problem. I will also explain Haslanger's response to the normativity argument by explaining her proposed definition of gender—particularly, *man* and *woman*—and how she may use it to avoid the problem. I end by presenting replies that defend the normativity argument and concluding that, even within the scope of her project, our intuitions about normativity, marginalisation, and oppression, the normativity problem remains a challenge she has not successfully overcome.

2 The Normativity Problem and the Normativity Argument

Although the normativity problem is described by Haslanger in a broader way, I will reconstruct it more formally here as the normativity *argument*. Before I do this, I will give the broad, short account of the normativity problem. Haslanger herself presents the normativity problem very briefly: 'The normativity problem raises the concern that any definition of "what woman is" is value-laden, and will marginalise certain females, privilege others, and reinforce current gender norms' (37). I believe her account can be reformulated this way: *in defining what it is to be a* woman *or a* man*, one inevitably creates norms that, rather than disrupt oppression, only perpetuate it*. When we look to either Haslanger's formulation or my own, the problem seems to be expressed much too quickly to illustrate the mechanics of how it works. Understanding these mechanics will allow us to better understand the approaches Haslanger and her definition of gender take in responding to it—what's more, a greater level of detail can be extracted from this quick version when we see this problem as more of an argument whose conclusion is *we ought not define what it is to be a* woman *or* man.

I think three premises can be extracted for the normativity argument. The first of these is that if defining a social category of people causes the marginalisation of individuals, then we ought not define that category. Although this premise is not explicitly found in Haslanger's broad figuration of the problem, it is important for the implicit argument of problem, because the premise serves as a standard by which attempting to define gender should be judged—that, simply, marginalising people is bad and any definition should not engage in it. Put differently, this premise is important because it is what makes the problem a problem: there being an ethical standard that must be met that defining gender does not meet. On its surface, this seems relatively straightforward and easy to accept—especially with a subject like gender, which carries great stakes in feminist politics, it makes sense that one ought not create ontological categories (by way of definition) that marginalise, and therefore harm, others.

The normativity argument's second premise is that if defining a social category encourages

normative behaviour, then individuals will be marginalised under that definition. This premise is also not explicit. However, including it is important not only to the logical structure of the argument by connecting the alleged effect of defining gender—encouraging normative behaviour—with the marginalising of individuals, but it is also important to why this problem is the normativity problem. The spirit of this problem, while it is applied specifically to gender, has to do with the ways in which normativity brings about undesirable consequences, and it is with this premise that normativity is introduced into the problem (and our argument). We might believe that this premise is true because, when a standard to which a thing is expected to conform is established, it seems to naturally follow that those things which fail to meet that standard are considered lesser than those that do. An example that illustrates this is the phenomenon of gatekeeping in various cultures and subcultures—behaviour that seeks to regulate who rightly belongs or does not belong in such a culture by imposing rigid (often unreasonable) standards for belonging. Fan subcultures are one vivid instance of this sort of marginalisation, albeit one with admittedly less severe consequences than others. A person might claim that, say, one is not a 'true' fan of a particular band unless one owns and has listened to all of that band's multiple albums. This definition encourages normative behaviour—that is, there is a norm (owning and listening to all the albums) that one must meet in order to be a 'true' fan. Consequently, in a situation in which a subculture accepts a conservative, gatekeeping standard for being part of that culture, those individuals which do not behave normatively are of a lower status than those that do—leaving those fans who, say, only listened to one album marginalised and those that have listened to all of them privileged. This example is vivid in its commonality, but some might also question its seriousness. After all, someone might say, being seen as not as much of a music fan as most does not seem like the end of the world. But this example is meant strictly to illustrate is that where normative behaviours are engendered, so too is marginalisation. Even if one's material conditions are not terribly affected by being marginalised in this specific context, it seems clear that marginalisation in some form is happening—one is still seen. Moreover, we might not take marginalisation among fan subcultures terribly seriously, but I think this is merely because what we are lacking is a kind of ubiquity to the norm established. A social category like gender is far more pervasive, so much so that it is typically unheard of in the *status quo* to assert that someone has absolutely *no* gender category. Clearly, the material consequences of this kind of marginalisation will be greater, as I will discuss soon.

Indeed, it is at this point when the argument's third, final, and most obviously necessary premise ought to be introduced: simply, that defining what it is to be a woman or a man encourages normative behaviour. If this is the case, then, given our previous two premises, it must be the case that we ought not to define what it is to be a woman or a man. Decades of feminist criticism, theory, and philosophy provide innumerable examples of our current definitions of *man* and *woman* enforcing normative standards, but one instance I find relevant to bring in as an example is the present oppression of transgender people. This will also serve as a second example of how normative behaviour marginalises individuals—here in a more pressing, violent way. The predominant, 'simple view' of gender—where to be a man is defined as having a phenotypically male body, and to be a woman is to have a phenotypically female body—encourages a frame-

work in which males are expected to act and be treated as men, with all its attendant behaviours, and females are expected to act and be treated as women. Trans people, in having an internal gender experience (one's psychological identification with being, say, a man or a woman) that is at odds with their phenotypical sex (and so assumed gender), violate this expectation. And currently, trans individuals face statistically obvious material oppression: the unemployment rate among the transgender community is reportedly double that of the general US population, with 26% of unemployed trans people report being fired by virtue of their gender identity; 78% of trans students in primary and secondary education experience some form of harassment, with 35% experiencing violence of some form; and so on (Grant et al. 2011). They also face other forms of other, less obvious discrimination, such as the denial of their gender identity among friends, family, and colleagues. This oppression would seem to confirm not only that there are normative behaviours that arise from defining gender (our third premise), but that such normative behaviour directly leads to the marginalisation of individuals (our second premise) as a kind of punishment for violating normative expectations. This punishing relationship is evidenced by a growing body of sociological research that suggests associations between individuals with strong cis-heteronormative beliefs (that being cisgender and heterosexual is the norm, proper, or right) and individuals with hostile attitudes toward trans people (such as in Worthen 2016, 37–38, 45).

Thus we have our full normativity argument:

(P1) If defining P social category marginalises individuals, we ought not define P.

(P2) If defining P encourages normative behaviour, then it marginalises individuals.

(P3) Defining what it is to be a woman or a man encourages normative behaviour.

(C) We ought not define what it is to be a woman or a man.

3 Haslanger's Account of Gender and Response to the Normativity Argument

Haslanger proceeds with a definition of *woman* and *man*. As such, she rejects the normativity problem, and therefore objects to our extracted normativity argument. Because this argument is valid, Haslanger—or any objector to the normativity problem—must reject one or more of our premises. There are two likely approaches to objecting to this argument, both of which Haslanger takes: perhaps more expectedly, to reject (P3) and, perhaps more surprisingly, to reject (P1).

However, before considering these objections to the normativity argument, I will first give Haslanger's definition of *woman* and *man* (albeit in simpler terms), because the content of her definition is important both for the rejection of (P3) and (P1). Haslanger defines *woman* by saying that person Q is a woman if and only if they meet the following three conditions:

(1) Q passes as a person with a phenotypically female body (that is, Q is consistently recognized as female);

(2) Q is marked as a person who ought to occupy the position of being socially oppressed because Q passes as female; and

(3) Q is in fact oppressed because Q meets conditions (1) and (2).

Haslanger defines man in a nearly identical way, substituting *female* with *male*, and *oppressed* with *privileged* (Haslanger 2000, 42). With this definition in mind, I can now present the two objections mentioned earlier to the normativity argument—namely, rejecting (P3) and rejecting (P1).

One may, as Haslanger does, reject (P3) by contending that this definition of *woman* (or *man*) does not encourage normative behaviour. This contention is grounded in the presence of the words *oppressed* and *privileged* in the definitions, such that they would encourage behaviour that works against gender normativity rather than encourage normative behaviour toward their fulfilment. That is, building oppression into the definitions of *man* and *woman*, there is enough negative motivation to disrupt and resist normative standards of these genders, because what is quintessentially *man* or *woman* is to either oppress or be oppressed, and, through conceiving of one's gender as necessarily entailing such a hierarchy, people will recognise and seek to end such oppression (46). This second claim, that people seek to end oppression, is implicit (and, perhaps to most of us, obvious)—however, it is important to acknowledge it as critical to this response. If people seek to end oppression, and they see oppression as a necessary property of being a man or a woman, then it follows that such people will not aim to live up to the normative ideals of those genders.

However, one may also reject (P1). It is important to explain the way in which Haslanger frames her project of defining gender, because this response necessitates an understanding of it. The framework she adopts is that of an *analytic project* whose goal is the furthering of feminist politics and the combating of oppression (Haslanger 2000, 33, 36). This means two things: that Haslanger aims to define gender in terms of its potential utility for an end, and that that end is opposing injustice and oppression against women. With this project in mind, one may respond to marg by suggesting that it is overly general in its application of avoiding marginalisation. That is, Haslanger holds that a requisite aspect of any sort of metaphysical project, particularly an analytic one such as hers, necessarily involves giving some values and goals priority over others, and it is only to those values or goals that one should be beholden. Because her definitions of *woman* and *man* are constructed as part of a larger feminist project to critique and eventually end the oppression of women as a social class, the pertinent question is not whether her definition may marginalise individuals generally, but rather, whether her definition marginalises individuals in a way that conflicts with the feminist values she adopts (Haslanger 2000, 46). Put differently, even if one accepts that her definitions may lead to normative standards, and that such standards may marginalise certain individuals, the only thing that matters is whether the definitions are consistent with the aims of her specific project. And, in the case of the definitions given, she

believes that they are consistent, because just as her goal is to end and oppose oppression and injustice, particularly against women, this is exactly what her definitions of woman and man bring to the fore: the presence of these as being closely linked with gender, particularly in such a way that subjugates women. Haslanger grants that, even if there are cases in which people who we would intuitively call women are not women on her view, and so ostensibly would be marginalised, such people and such consequences are not relevant to her project, and so do not provide enough negative motivation against defining *woman* or *man* (Haslanger 2000, 46).

4 My Replies to Haslanger's Objections

I do not find either of these responses satisfying, and I will now offer replies to each of them in defence of the normativity argument. In the case of Haslanger's first response, rejecting (P3), I believe she both ignores the risk posed by the first condition to encourage normative behaviour, and mistakenly overemphasises the power of building-in oppression to mitigate such normative behaviour. If we recall the first condition of her definition of *woman*, it seems apparent that a clear normative standard is established without mention of oppression: one must *pass as a female* before anything else to genuinely be a woman (likewise for men, with passing as a male). This presence of a standard against which a person can be judged that is independent of one's oppressive status, even if it is only one component of being a woman, suggests that there certainly is a possibility of normative behaviour in the service of validly identifying as a woman being encouraged. Haslanger depends heavily on the fact that to be a woman on her view is to be oppressed in order to counter normative drives to perform womanhood properly. If you are a woman, then you are *ipso facto* oppressed on her account. This oppressive, marginalised status that comes with womanhood is intended to motivate the challenging of normative gendered norms. If one is to be oppressed by virtue of being a woman, then one will be inclined to kick against the normative standards that engender the oppression of women. However, I am unconvinced that this is enough. Rather, I believe that a normative standard of passing as phenotypically female can quite easily marginalise.

Earlier in this paper, I discussed the oppression of the trans community as at least in instance of punishment for the violation of norms on the basis of one's assigned gender and sex. On Haslanger's account of gender, not enough is done to relieve the oppressive pressure that gendered norms place on trans people. Although Haslanger ostensibly allows for a trans person to properly be a man or a woman—as her definition says, one's gender is only related to sex in how others recognise or imagine one's sex—this puts a disproportionate level of importance on passing as the sex associated with one's preferred gender in order to properly *be* that gender. Again, as with the simple view, one's gender identity is not relevant to one's 'actual' gender, but is subordinated to a normative standard—here, being able to *seem* to be one of the sex associated with that gender. Inevitably then, there still will be people who may, by their own lights and internal experience, be women, but who are not *genuinely* women on Haslanger's view by virtue of their inability (or unwillingness) to pass as female—and, as is the case for transgender people generally today, such people will be marginalised by Haslanger's definition. And so, by

our normativity argument, such a definition should not be accepted.

But if marginalisation is not something we ought to avoid wholesale, this ceases to be a problem. Thus my reply to Haslanger's second response: simply, even accepting her framework about what is important, her very project demands that individuals not be marginalised, particularly in the case of transgender women already alluded to above. As I've mentioned earlier, Haslanger's project is one that seeks to end oppression and injustice. And, in defining gender, that project makes use of feminist values to achieve that end—more simply, she pursues 'sexual justice' (37). However, it seems to me that, if I am right in my analysis of *passing as female* being a condition of being a 'true' woman, then her goal of opposing injustice and oppression isn't being met—rather, it becomes a new standard by which injustice (here, that injustice meted out to transgender women who fail to pass as female) is perpetuated. The fact that these non-passing transgender women, although not genuinely women by either Haslanger's definition or the simple definition, still share some psychological property, experience, or strong affinity with other women, should intuitively suggest this as well. In some sense, Haslanger's definition of *woman* still causes some kinds of women—or at least some people who have a strong affinity with womanhood—to be oppressed. And because this oppression would be rooted in their not being recognised *as* women, and so would not be considered oppressed *as women*, we cannot rely on our drive to end oppression to end this form of it, because it will not be recognised as such to begin with.

One might accuse me of question-begging here by critiquing Haslanger's definition of *woman* by presupposing another definition that includes trans women as women. But my point is that in an analytic project of gender that is meant to be progressive and to fight the oppression of women, empirical facts like *There is a group of people who claim identity as women* do not vanish, and neither should our intuitions about those facts. If Haslanger's project is to combat the oppression of women, I think the spirit of her project naturally ought to attend to individuals who also have close affinity to womanhood. Even if one does not take trans women to *be* women—a view which, as I've argued earlier, subjugates and oppresses these women—one intuitively cannot deny that trans women bear at minimum a uniquely closer relationship to womanhood from men, one that should qualify some sort of protection in a project aiming to define gender in the service of ending the oppression of women. But I think Haslanger's proposed definition does little to protect trans women who do not pass as cisgender are unjustly marginalised by the norm created in her definition. Indeed, her definition only lays the foundation for oppression. As such, this response, too, doesn't seem to erase the problem of normativity, because even according to Haslanger's project, unjust and unintended marginalisation occurs.

5 Other Criticisms of Haslanger's Account

Both Haslanger's conceptualisation of the *analytic project* and her proposed account of gender have been often cited and built upon. Of particular interest here, I believe, are two criticisms—one particular criticism by one particular philosopher, and one species of response that has

gained increasing support among social philosophers. Briefly discussing these will make clearer the scope and nuanced contours of the normativity argument against Haslanger's attempt at defining gender (and against defining gender generally) I offer in this paper.

Katharine Jenkins (2016) offers a criticism of Haslanger's account of gender with a very similar spirit to my own: that Haslanger wrongly marginalises trans people through her proposed definition. However, Jenkins's critique of Haslanger both complements and differs from my own in two ways. First, and perhaps most importantly, Jenkins does not focus on *normativity* as the agent of marginalisation, but *exclusion*. Haslanger discusses the normativity problem in her paper, but also what she calls the *commonality problem*: the claim that there isn't a property that sufficiently unites all the people we might want to call *women*, and so some people are (wrongly) left out (Haslanger 2000, 37). It is the commonality problem that Jenkins focuses in on. In discussing how Haslanger's account might bear on the status of trans women, Jenkins concludes as I do, 'trans women will be categorised as women by Haslanger's account only if they find themselves in scenario 3 [passing as phenotypically female] most of the time. Some trans women will never find themselves in scenario 3, and many trans women will find themselves in that scenario only some of the time. Therefore, many trans women will not be categorised as women according to Haslanger's definition' (Jenkins 2016, 401).

For Jenkins, the exclusion of trans women is the critical marginalising factor. By contrast, I hold normativity as the primary obstacle to a development of a non-marginalising model of gender. It is not strictly because Haslanger's definition excludes trans people that I take it to marginalise them, but rather because it creates a normative standard one must meet to validate one's gender identity. If I am right, then normativity proves problematic in defining gender in a more robust and diverse way. While the exclusion of trans people by Haslanger's definition is one of the more vivid ways in which her—or any—definition of man and woman might marginalise people through its normative aspects, normative marginalisation may manifest in many ways, as I stated in §3. Second, Jenkins does not, as I do, take the problem she discusses to be one that endangers Haslanger's project, but sees it as a prompt for a more rehabilitated instance of 'trans-inclusive amelioration' (407). I am not convinced that Haslanger's approach—or indeed any approach to modelling gender as a social category—can be successfully rehabilitated in light of the normativity problem. But I will touch on this again in §6.

Another tactic, popularised by Theodore Bach (2012, 2016), is to reject Haslanger's fundamental project of modelling *man* and *woman* as social constructs unified by some empirical property (i.e., all women share the property *being oppressed*), what Bach terms an *objective type* (2016, 179). The main thrust of this criticism is twofold: first, a use of objective types (as with Haslanger) risks undermining the very political goals that motivated their usage in the first place. Specifically, Bach believes that to define gender (or any social category) by a particular property—here, one's oppressed or privileged status—is to engage in 'an empirical and political gamble' (194). That is, one stakes a claim that it is one's oppressed status that is the only (or at least most) important fact to consider in a political project. However, this is not guaranteed. As I've argued in this paper, Haslanger wrongly stakes too big a claim on passing as a sex associated with a gender. And second, a use of objective types shuts out epistemological possib-

ilities: because Haslanger defines gender using a unitary, rigid property, it becomes impossible to 're-identify' a gender in more accurate terms as features of it may change over time (Bach 2016, 198). If being a woman is *ipso facto* being oppressed, then we cannot imagine a conception of womanhood that doesn't involve this oppression (195). Instead, Bach supports a model of gender that figures man and woman as natural rather than socially constructed. Moreover, Bach claims the unifying principle that makes gender (and all social categories) cohere is *historical essentialism*. The suggestion is that the essential feature of *man* or *woman* is an individual's connection to a historical lineage of other individuals of the same kind (Bach 2012, 242–43). By this Bach means that the essential feature that makes someone a man is to take on properties that 'make that individual a replication of ancestral men', and likewise for women (2016, 193). What makes this alternative model of gender so appealing on Bach's view is that the properties that mark this lineage can be varied and fluid: just as a species, another natural kind, can have some properties change over time while preserving the species on this view, so too can a gender or other social category. As such, Bach claims to circumvent the commonality or normativity problem in that one can fail to instantiate 'characteristic gender properties' while still properly being of the gender associated with those properties, as the essential feature of that gender is found in 'wider historical processes' (193).

A proper reply to Bach's approach necessitates a fuller discussion beyond this paper, but I will mark a broad, but hopefully illuminating, difference. As with Jenkins, Bach hopes to preserve a working model of gender in spite of the theoretical problems, and in fact cites conceptually preserving gender as a benefit of his approach over type-objectivism (195–96). I find this urge to preserve gender unmotivated and without obvious benefit, especially in the face of the innumerable problems that come from it (both conceptually as discussed here, but politically and ethically as discussed in the whole body of other feminist philosophy). Moreover, out of my normativity argument, if one *did* seek to properly model and preserve gender, attending to and avoiding the creation of norms would include norms favouring binary gendered systems. Marion Godman has shown that while Bach's theory can be applied to account for non-binary gendered systems, such an application risks losing some of the theoretical virtues Bach claims—such as unifying members of a gender across cultures under the same definition (Godman 2018). To put these differences in simplest terms: I take the normativity problem to be serious enough that I am not motivated to preserve a concept of gender in the face of it (or its other theoretical vices).

6 Conclusion and Implications

In this paper, I've taken the shortly—and broadly—described normativity problem for defining gender and expanded it into a larger argument with clearer mechanics, and have considered how Haslanger and her definition of *woman* and *man* have attempted to either reject or circumvent the problem. I conclude that these attempts fail when confronted with intuitions about normative behaviour and marginalisation that we can observe today within and regarding the transgender community. Haslanger suggests on one hand that her proposed definitions of *man*

and *woman* do not suggest the establishment of norms. However, I take her condition of passing as a particular gender to be enough of a norm to cause worry. And Haslanger asserts on the other hand that her analytic project and consequent definition are not and need not be concerned with marginalisation. However, given the goal around which her project is oriented, the ending of oppression and injustice, I maintain that her proposed definitions do still create normativity and marginalisation that create the injustice she seeks to end. And so, the normativity problem remains a problem.

How we move forward with modelling gender in response to this problem is a subject for further inquiry. If the normativity argument as I have reconstructed and defended it is right, then perhaps the natural conclusion is to do away with the project of defining gender. As I have argued for this strong version of the argument, I am inclined to accept such a conclusion. Conceptually, it seems, gender might run too high a risk of marginalising to properly define. This is a view that naturally lends itself to the yet-more ambitious stance of abolishing gender, doing away with it as an object in our social ontology entirely. But if the challenge is to create a model of gender that manages to circumvent the problems I've discussed (as Jenkins, Bach, and Godman attempt to do) then such alternative avenues of definition might be explored. But whether we are to abolish gender or find an innovative new model of it, we must reckon with the challenge of normativity and its capacity to marginalise.

References

Bach, Theodore. 2012. 'Gender Is a Natural Kind with a Historical Essence'. *Ethics* 122 (2): 231–72.

———. 2016. 'Social Categories Are Natural Kinds, Not Objective Types (and Why It Matters Politically)'. *Journal of Social Ontology* 2 (2): 177–201.

Godman, Marion. 2018. 'Gender as a Historical Kind: A Tale of Two Genders?' *Biology and Philosophy* 33 (3–4): 21.

Grant, Jaime M., Lisa A. Mottet, Justin Tanis, Jack Harrison, Jody L. Herman and Maria Keisling. 2011. 'Injustice at Every Turn: A Report of the National Transgender Discrimination Survey, Executive Summary'. *National Center for Transgender Equality and National Gay and Lesbian Task Force.* `endtransdiscrimination.org/PDFs/NTDS_Exec_Summary.pdf`.

Haslanger, Sally. 2000. 'Gender and Race: (What) Are They? (What) Do We Want Them to Be?' *Noûs* 34 (1): 31–55.

Jenkins, Katharine. 2016. 'Amelioration and Inclusion: Gender Identity and the Concept of Woman'. *Ethics* 126 (2): 394–421.

Worthen, Meredith G. F. 2016. 'Hetero-Cis-Normativity and the Gendering of Transphobia'. *International Journal of Transgenderism* 17 (1): 31–57. doi:10 . 1080 / 15532739 . 2016 . 1149538.

Conceptual Engineering in the Biological Image

Brian Kobylarz*

bwkobylarz@gmail.com

Abstract

Conceptual Engineering, and investigations into concepts generally, is becoming a focus of recent academic philosophy. The importance of this movement is upheld here, and I approach the conversation with a metaphysics of concept relations. I claim such relations between concepts feature properties expressible in the language of theoretical biology, particularly as presented by William Wimsatt. By synthesising the properties of robustness and entrenchment with themes in the conceptual engineering conversation, I introduce new concepts of my own and hope to show that a biology, or complexity orientation, can aid the conceptual engineer.

1 Introduction

The creation of concepts is an activity at least as old as philosophy itself, and meta-level considerations of this activity can be found throughout the history of philosophy. Recently, convergence has emerged on the importance of this meta-level, and attempts are being made at teasing out the various dimensions to concepts and their production. There exist, for instance, conceptual engineers and meta-philosophers who advocate the value of concept building and selection generally, with implications ranging from the epistemic to the normative, and with import into our day-to-day (Burgess and Plunkett 2013; Eklund 2017). This effectively amounts to a methodological shift from using whatever concepts cultural evolution has left lying around to the intelligent design of concepts. Another group, the cognitive penetrability community, suggests an entanglement between our conceptual apparatuses and our ability to perceive (Zeimbekis and Raftopoulos 2015). Here we find a successor to the theory-ladenness of observation tradition in the philosophy of science (Hanson 1958; Kuhn 1962; Thagard 1992). Psychologists and neuroscientists, through various brain imaging technologies, are exploring concepts on empirical grounds (Binder et al. 2009; Hofmann et al. 2012). Recently, a study at Carnegie Mellon

*Brian Kobylarz completed his BA in Philosophy at William Paterson University, and plans on submitting his transcripts to PhD departments in 2019. He is interested in various philosophical traditions but is currently most interested in concepts, how they grant agency, and how to model it. He is currently exploring intersections between Conceptual Engineering and Dynamical Epistemic Logic.

claims to have visually confirmed distributed neural activation signatures which corresponded to subjects thinking about particular concepts in physics (Mason and Just 2016). Along the meta-philosophical axis, Luciano Floridi presents a call-to-action of his own:

> Too much ink has been spilt on philosophy as conceptual analysis. The alternative view, that philosophy is at least as much, if not actually more, engaged with creating, refining and fitting together our conceptual artefacts [...] has received too little attention (Floridi 2011, 293).

A trans-disciplinary convergence exists here in terms of the significance granted to concepts. I too share many of these beliefs, and I am convinced further convergence is possible, desirable, and important. Concepts are given many roles, they are said to frame discourse, to grant agency, orient observation, compose into explanatory frameworks, and allow for reasoning generally. It is difficult to even speak of 'concepts' without employing the usage of the word 'concept.' As with any subject matter imbued with such gravity there exist less optimistic, sceptical voices. One such representative here, expressing important issues for those who wish to engineer concepts, is Herman Cappelen. For Cappelen our concepts—even though he refuses to label them as such—are largely at the whim of social forces; they are not things we can domesticate and sterilise nor intelligently design on par with what the conceptual engineer hopes. In his own words:

> In most cases the detailed mechanisms that underpin particular instances of conceptual engineering are too complex, messy, nonsystematic, amorphous, and unstable for us to fully grasp or understand (Cappelen 2018, 72).

I find Cappelen's portrayal to be the right portrayal, but I differ in that I believe a solution rather than a dead end is implied. I thus find myself on more optimistic grounds. I take the usage of 'complex' to be where I start. By 'complex' I mean a system with many moving parts, something that if modelled, would demand many variables to capture the relevant mechanisms such that correlations and causal relationships can be mapped, explored, and exploited. Often such systems exhibit emergent behaviours, and non-linear developments which are beyond prediction. If conceptual engineering has a complexity problem, then looking towards those sciences which have successfully managed to navigate complexity appears to be the reasonable step to take. This paper thus exists as an initial, and very rudimentary attempt at exploring some points of connection between conceptual engineering and the complexity/network sciences. William Wimsatt happens to be an exemplar in providing novel ways to speak about and understand complexity. As such, I will be appropriating some of Wimsatt's language to the conceptual engineering conversation. I believe this labour to be one which is in some sense continuous with Cappelen as he claims: 'We should object to any theory that makes it [conceptual engineering] easy' (70). By portraying concepts, and their relations in light of complexity, we can admit the messiness while welcoming various techniques and methods for navigating the messiness.

What I intend on importing here exactly, from Wimsatt to conceptual engineering, are two ideas: robustness, and entrenchment. Both ideas denote object relations. The modus operandi of this paper can then be said to be descriptive——a metaphysical investigation into concept relations where I hope to show that concepts relate as such, and that the consideration of both relations are important for the engineering, augmentation, and implementation of concepts. It may seem strange to apply this biological image to concepts, as it appears to demand more metaphysics than I have room for here, but as Wimsatt's work makes explicit: the properties under consideration are found beyond the biological sphere. Generally, a network that adapts and has dependency relations between its members will show the properties I will explore (Wimsatt 2007). Thus, it is equally fair to refer to these *biological properties* as *adaptive network properties* more generally. Once we grant adaptive network properties, it should be less controversial to associate robustness and entrenchment with concepts and their functionality.

2 Divisions of compositional labour

To frame, take for instance the concept 'ice-cold', this may be used to refer to a drink's temperature. Given the right conditions, this concept can also be employed to refer to a particularly cruel individual. For the concept 'ice-cold' to survive, and to be suited for a range of phenomena like the temperament of an individual or the temperature of a drink, it must change along various axis, including those denoting verbal fashion, the environmental context which influences its utterance, and the analogical skills of the speakers in the conversation where 'ice-cold' appears. I find it, then, non-controversial that concepts are adaptive, as their applicability, usage, and utility (in terms of possible inferences, and actions they enable) change relative to various influences.

Another influence on concepts can be concepts themselves. Here we see network or ecological properties present in terms of how our concepts relate to one another. Consider the concept 'fire', which when put into conjunction with the concept 'danger' can be altered such that it can convey a very particular warning, and with it a sense of urgency that can trigger the production of stress hormones. The same concept 'fire', when conjoined with 'barbecue', can have a welcoming effect such as the management of hunger, leading to inferences that dinner is on its way. Because of such relations between concepts and how they can influence one another, I then have no issue committing myself to what can be called *concept ecologies* or *concept networks*. Dependency relations between the concepts can be elucidated through this language. We use concepts, and they can loosely be said to mutate by forces both internal and external. The internal amounts to concepts operating on concepts, such as how they compose to enact different responses as when 'fire' is used in conjunction with 'danger'. The external can denote those environmental forces, such as an 'ice-cold' person, which influence what utterances are made (which concepts are enacted). I accept both divisions of labour. While we may have our differences as to how we weigh the internal labour and the external labour which shape our concepts, I find concept ecologies or concept networks to be permissible for both internalists and externalists, as well as those who uphold the necessity of the conjunction of both, such as

myself. This essay will, however, be largely focused on the internal relations, in other words: the concept-on-concept relations within a concept ecology. It is my belief, then, that for the kind of mastery over concepts the conceptual engineer desires, the complexity of the related systems must be considered. I will begin by outlining robustness, move to entrenchment, and conclude with some remarks as to what other problems arise within this framework.

3 Robustness and amelioration

Wimsatt defines robustness as 'the existence of multiple means of determination or access to an entity, property, process, result, or theorem that are at least partially independent' (359).

In biology, this lends itself well to expressing situations where one biological activity can be instantiated by various biological mechanisms. The imagery of networks can also be grasped through Wimsatt's characterisation. Here, nodes can represent an entity, and arrows between the nodes can represent the causal relationships between nodes in a network. The more arrows which converge upon a single node (entity), the higher the measure of robustness, the greater chances an entity can be instantiated by other entities within its ecology (network). The evolutionary benefit here is obvious, as mechanisms can pick up the slack of other mechanisms—a network feature which increases the survivability of an organism.

Let's take concepts then as our entities. Robustness can then be held as the measure which allows us to approximate the degree of effect concepts on other concepts, and what effects a network can produce in light of this relation. The concept designer, nearly by definition, is someone who is interested in the result(s) of a given concept. Take for instance a general situation where X is a desired effect believed to be arrived at through concepts 2, 3, and 4. In our concept ecology, concepts 2, 3, and 4 are present, but none of them are producing X. If X is not currently actualised by this network, then it could amount to an unnecessary and costly endeavour for the conceptual designer who chooses to build an entirely new concept to bring about X. The payoff may be arrived at through investing less into other—already existing—concepts by way of a kind of augmentation to them. Choosing to build a new concept here would be akin to nature building an appendage where one isn't needed, as the present appendages would be able, with slight modification, to complete the desired task. To understand robustness is then to both understand a property of concept ecologies, and to know when concept creation is appropriate—and when augmentation of already existing concepts is appropriate.

To make this less abstract, let's return to the concept 'fire' where we already found an example of robustness. Someone possessing the concept 'fire' has the initial materials which enables inferences to be arrived at for a robust range of situations where 'danger' and 'comfort' are appropriate. If their 'fire' concept is defective, lacking a degree of robustness in terms of possible inferences regarding the dangers of fire, then we may opt to educate them in terms of the dangers of fire, and fire safety more generally, thus expanding the possible inferences and actions the concept 'fire' can grant a subject. In the language of conceptual engineering, *amelioration* is the name given to the activity of augmenting concepts as such. As Herman Cappelen understands,

amelioration can come in a 'variety' of strategies, one such example Cappelen articulates:

> You have a word with a certain meaning and extension/intension. You keep the word and improve the meaning/extension/intension. In the standard case, the lexical item will be preserved, but the semantic value will be improved (Cappelen 2018, 35).

In the case of teaching fire safety, we find the lexical item preserved ('fire'), and the semantic value changed such that 'fire' denotes not only that which cooks, heats, and warms but also that which can harm, kill, and destroy. Teaching fire safety, in other words, makes for a more robust conception of 'fire', broadening the range of possible inferences regarding fire and thus the possible actions someone sufficiently informed can take in relation to fire. One can imagine such an expansion to allow other effects to be arrived at through combinatorial concept processes such as when 'fire' combines with 'danger' in a statement, allowing both concepts to interface with one another to grant the desired effect of successfully conveying a warning.

In closing this section, there are those who find amelioration to be too difficult. I find this position odd as it seems clear to me that while amelioration is difficult, there exists examples of it. My initial intuition is to consider education generally and how at all it would be possible if amelioration were beyond us. Propaganda also comes to mind, where it can be argued that significant alterations have been made throughout history to change the extensions of concepts such as 'enemy', 'solidarity', 'terrorism' and 'terrorist', all of which appear malleable in the hands of propagandists. While this amounts to a generally monstrous conceptual engineering, it also suggests a more optimistic reading: that concepts are within our control and that it can presumably be employed for good. A third example comes from psychology, where we find that cognitive behavioural therapy has for decades demonstrated success for the treatment of various mental illnesses (Hofmann et al. 2012). The patient, employing cognitive behavioural therapy, is directed to keep journals, to stay mindful of negative thinking habits, and to replace those negative thoughts with more positive conceptions. This amounts partially to a reordering of already existing cognitive architecture, and one which is conducted with success by an individual without the funding and mobilisation of propaganda and education. I believe one can frame conceptual engineering as a kind of cognitive behavioural therapy for epistemic and normative conepts. Amelioration, as far as I can tell, is a ubiquitous human activity.

4 Generative Entrenchment

As William Wimsatt presents, generative entrenchment is 'a measure of how many things depend upon an element and thus likely to change if it changes' (Wimsatt 2007, 355).

We can think of our anatomical arrangements as deeply entrenched—they depend upon a number of mechanisms such as our genetics. This amounts to a dependency between a higher order thing and a more foundational thing where a change to our genetics can have significant

ramifications to our anatomical arrangement. Entrenchment can also play a role in determining what entities can be integrated into an ecology offering a degree of resilience, and predisposition, to a network where whatever is commensurate with higher order entities must also be commensurate with lower level entities due to the dependency relations. In other words, entrenched entities operate as gatekeepers, in part determining what is integrated into our concept ecologies, and what is not.

As I mentioned earlier, concepts appear to depend on other concepts too, as does our agency itself. This important point is expressed by Burgess and Plunkett in their work 'Conceptual Ethics I':

> Arguably, our conceptual repertoire determines not only what beliefs we can have but also what hypotheses we can entertain, what desires we can form, what plans we can make on the basis of such mental states, and accordingly constrains what we can hope to accomplish in the world. Representation enables action, from the most sophisticated scientific research, to the most mundane household task. It influences our options within social/political institutions and even helps determine what kinds of people we can be, what sorts of lives we can lead (Burgess and Plunkett 2013, 1096–97).

We are presented with concepts as influencing one another such that our epistemic capacities are implicated by the concepts we possess, as are our capacities to act more generally. As the conjunction of entrenchment and concepts suggest here: to what degree are our concepts then entrenched by one another, and ultimately what does this mean for the constraint and freedom we have to know and act? The ability to see concept ecologies as such—to conceptualise concept ecologies as such—leads to investigations into some of the necessary conditions for concept implementation. The concept designer speculates in the form of: if concept 1 is to be introduced into concept ecology ϕ, then what entrenched concepts are present in concept ecology ϕ such that concept 1 may be successfully or unsuccessfully integrated? An example: One could imagine, for instance, a fairly reasonable individual with entrenched concepts which make said individual value virtues such as deduction and consistency. This individual's concept ecology may then be immunised against the addition of new concepts which do not follow, either explicitly or implicitly from the concepts present. In other words, our entrenched concepts can immunise us against some concepts, and leave us predisposed to others. If Burgess and Plunkett are correct, this greatly determines what we are capable of. Conceiving of concept relations with entrenchment can aid the concept designer in terms of what concepts will survive and what concepts will die off when integrated into a concept ecology.

It is common in biology to find mechanisms and processes which have multiple functions. While entrenched concepts can immunise against or predispose us to new concepts, they can also have large scale implications when tampered with, given the nature of their fundamentality. I am forced to consider, then, the possibility of *concept ecology collapse*. Would it be possible, once the entrenched concepts are discovered, to manipulate them or introduce concepts which cause the networks to fail? Are there concepts which entrench to such a degree that entire branches in

our *concept networks* would fail if the entrenched concepts were tampered with in an irresponsible way? The implications for this are as exciting as they are horrific. It may be useful to borrow another concept from biology here, that of the 'invasive species', where a species, traditionally foreign to an ecology, can be introduced with disastrous consequences, due to an ecology lacking an arrangement which can support such a species. The conceptual designer may want to consider *invasive concepts*, or the degree to which a designed concept is invasive and how such invasive concepts function relative to a concept ecology and the desired goal(s).

Perhaps here amelioration can find itself in some cases to be a supplement to concept creation and implementation where an amelioration is conducted on deeply entrenched concepts while new concepts are introduced, creating an environment where a new concept can be integrated with greater ease while diminishing potential externalities including those such as concept ecology collapse, or the rejection of a new concept given the entrenched concepts inherent to a concept ecology. Entrenchment allows further evaluations of the structure of a concept ecology, giving the concept engineer an idea of what the fundamental concepts are, and how to proceed.

5 Conclusion and future work

Experiments conducted in the 1980s lead to the understanding of fruit fly construction during the embryonic stage where the development of a fruit fly emerges from head to wings. This process follows an elaborate and delicate chemical activation of certain genes in a particular order (Shubin 2009). The activation is sequential, and thus the history of this activation is necessary to fully grasp and model the development of a fruit fly. . There exist numerous biological mechanisms throughout nature which exhibit sequential operations. Delicate operations, such as this, share some symmetry with conceptual engineering. For mastery over our concepts, precision seems necessary, and precision appears only possible when the multivariate and complex dynamics which shape our concepts are understood. This amounts to an advocacy of descriptive investigations into our concept-and-concept relations as well as the concept-and-world relations which determine our ecologies/networks. This of course demands much more work than what is possible in this paper, which as I mentioned in the introduction is an attempt to simply bridge complexity and conceptual engineering in an initially novel way.

When one looks towards biology and other sciences, which take seriously the hierarchical and complex relations between their objects, there exists measurements. Robustness and entrenchment can be defined quantitatively, if my characterisations have been successful then presumably measurement can be on the horizon for the conceptual engineer. Significant work is needed in order to arrive at such a methodology. This leads to the elephant in the room, that being what exactly a concept is, as before we can begin to model, map, measure and determine with empirical accuracy our concept ecologies/networks, we must have a satisfying definition of what is to be measured. Cappelen's cut is deeper here as he addresses the lack of convergence in terms of a universally accepted theory of concepts (Cappelen 2018). This attack on the presuppositions of conceptual engineering may condemn the project before it gets off the ground. It

is true that there exists a smorgasbord of different theories of concepts (Carey 2009; Machery 2009; Churchland 2012; Gärdenfors 2000), and that as such a pessimistic reading is within reason. I do not have the room here to develop a theory of concepts, but there are some presuppositions which this paper makes concerning the nature of concepts. One such presupposition is that concepts are representational devices as they approximate environmental conditions. On the brain level, these representations are seemingly distributed (Binder et al. 2009). Another presupposition present is that concepts have a combinatorial nature, they arrange themselves relative to one another and environmental context. I find that the compositionality of our concepts constrains and determines much of what our concepts can do. As such, this biology/complexity semantics I have explored seems necessary to consider beyond the neural-level. This amounts to a descriptive venture into a level beyond neurons (concepts and their relations) and the prescriptive claim that a science of such a level is necessary. Thus, I am at odds with the Eliminativist tradition. Both the significance of compositionality (see Werning, Hinzen and Machery 2012; Murphy 2004) and concepts as representational devices (Churchland 2012; Mason and Just 2016) are popular theories of concepts, and their performance. How both levels interface with one another is another problem. I suspect that as neuro-imaging technologies advance, concepts themselves will begin to be taken more and more seriously as observation of conceptual operations in the brain will come into clearer focus. The work I've presented here, I believe to be useful to the understanding of such an interface between neural and language based investigations/interventions into our concepts. While I am interested in the emancipatory capabilities of conceptual engineering, I then am also interested in a kind of concept-informatics which accompanies the engineering of our concepts.

I believe both would feedback into one another and that there are answers, with number values, to questions such as:

(1) How deeply entrenched is concept 1?

(2) What degree of robustness, in terms of possible inferences, does concept 2 possess?

Finally, I am of the position that conceptual engineering needs to be an interdisciplinary venture. I believe that organisational attempts to synthesise cross-discipline discoveries should be made, and that each discipline contributes to a greater image of concepts and their relations, which hang together as complex systems. The philosopher's labour enters in terms of the creation of concepts to grant greater degrees of agency, and the augmentation of concepts to make actual the unfamiliar potentials in familiar concepts. This extends to the necessary meta-labour needed to sort out conceptual engineering itself, which I have hoped to elaborate a bit further myself.

References

Binder, Jeffrey R., Rutvik H. Desai, William W. Graves and Lisa L. Conant. 2009. 'Where Is the Semantic System? A Critical Review and Meta-Analysis of 120 Functional Neuroimaging Studies'. *Cerebral Cortex* 19 (12): 2767–96. doi:10.1093/cercor/bhp055.

Burgess, Alexis, and David Plunkett. 2013. 'Conceptual Ethics I'. *Philosophy Compass* 8 (12): 1091–1101.

Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering.* Oxford University Press.

Carey, Susan. 2009. *The Origin of Concepts.* Oxford University Press.

Churchland, Paul M. 2012. *Plato's Camera: How the Physical Brain Captures a Landscape of Abstract Universals.* MIT Press.

Eklund, Matti. 2017. *Choosing Normative Concepts.* Oxford University Press.

Floridi, Luciano. 2011. 'A Defence of Constructionism: Philosophy as Conceptual Engineering'. *Metaphilosophy* 42 (3): 282–304.

Gärdenfors, Peter. 2000. *Conceptual Spaces: The Geometry of Thought.* MIT Press.

Hanson, Norwood Russell. 1958. *Patterns of Discovery: An Inquiry Into the Conceptual Foundations of Science.* Cambridge University Press.

Hofmann, Stefan G., Anu Asnaani, Imke J. J. Vonk, Alice T. Sawyer and Angela Fang. 2012. 'The Efficacy of Cognitive Behavioral Therapy: A Review of Meta-analyses'. *Cognitive Therapy and Research* 36 (5): 427–40. doi:10.1007/s10608-012-9476-1.

Kuhn, Thomas S. 1962. *The Structure of Scientific Revolutions.* University of Chicago Press.

Machery, Edouard. 2009. *Doing Without Concepts.* Oxford University Press.

Mason, Robert A., and Marcel Adam Just. 2016. 'Neural Representations of Physics Concepts'. *Psychological Science* 27 (6): 904–913. doi:10.1177/0956797616641941.

Murphy, Gregory L. 2004. *The Big Book of Concepts.* MIT Press.

Shubin, Neil. 2009. *Your Inner Fish: The amazing discovery of our 375-million-year-old ancestor.* Penguin.

Thagard, Paul. 1992. *Conceptual Revolutions.* Princeton University Press.

Werning, Markus, Wolfram Hinzen and Edouard Machery, eds. 2012. *The Oxford Handbook of Compositionality.* Oxford University Press.

Wimsatt, William C. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality.* Harvard University Press.

Zeimbekis, John, and Athanassios Raftopoulos, eds. 2015. *The Cognitive Penetrability of Perception: New Philosophical Perspectives.* Oxford University Press.

How to Be a Gricean Russellian

Martin Schlombs*

schlombs.martin@gmail.com

Abstract

According to many philosophers, many utterances of sentences contain unarticulated indexical expressions. In response to the problem of incompleteness, for instance, many Russellians hold that definite descriptions contain unarticulated indexical expressions which restrict the denotation of the description's nominal. I argue that such unarticulated indexical expressions pose a problem for Griceans who wish to explain meaning in terms of speakers communicative intentions. Roughly, the problem is that if there are unarticulated indexical expressions, then speakers can't rationally intend to have their audience entertain a single complete proposition by their utterance of a sentence.

Following Buchanan (2010), I suggest that the only way for the Gricean to meet this problem is to argue that often *restricted proposition-types* (and not propositions) are the objects of speakers communicative intentions. In particular, I argue that restricted proposition-types should be understood as proposition-types restricted by speakers' *act-coordinating intentions*. In uttering a sentence speakers intend their audience to entertain any token proposition of their communicatively intended proposition-type; furthermore, speakers intend that the proposition the audience entertains by means of their utterance disposes them to act in coordination with the speakers intended action.

1 Introduction

According to Russellians, definite articles (e.g. the determiner 'the' in English) semantically express *existence* and *uniqueness* of their nominals. So, for instance, for 'the beer is cold' to be true, according to Russellians, it must be the case that (among other things) 'beer' denotes a set with at least one member and at most one member (i.e. a singleton set). Since there are many cases in which speakers can intuitively say something true by uttering a sentence containing a description, although the description's nominal doesn't seem to denote a singleton set (Strawson 1950, 332), many Russellians suggest that often a description's nominal contains a hidden indexical element such that in context the nominal denotes a singleton set. So, for

*Martin Schlombs finished his Bachelor's in Philosophy at King's College London this summer. He intends to continue his philosophical studies at the University of Oxford.

instance, although $\{x : x \text{ is beer}\}$ isn't a singleton set, Russellains argue that 'the beer is cold' can be true in context since the set denoted by the description's nominal is the intersection of $\{x : x \text{ is beer}\}$ and a set denoted by an unarticulateed indexical expression. The intersection of these two sets is a singleton set.

I will argue that the above Russellian-indexical account of descriptions is problematic for Griceans. In many contexts there will be too many potential values the nominal's hidden indexical element might be thought of as referring to and hence it's hard to see how a speaker's communicative intention should fix the indexical element's value in context. I will further argue that the best way for the Gricean Russellian to meet this difficulty—known as the multiple-candidate problem—is to argue that sometimes the object of a speaker's communicative intention is a *restricted proposition-type* (and not a proposition). A restricted proposition-type, is a proposition-type restricted by the speaker's *act-coordinating-intention*.

The plan for this essay is as follows. In §2 and §3, I outline Gricean semantics and Russell's theory of descriptions. In §4 I present the multiple-candidate problem and in §5 I argue that this problem shows that speakers can't mean single complete propositions by their utterances of sentences containing incomplete descriptions. In §6, I endorse Buchanan's suggestion that the objects of speakers communicative intentions are *restricted proposition-types*. In §7, I try to make sense of what a restricted proposition-type is; I argue that it's a proposition-type restricted by a speaker's *act-coordinating intention*.

2 Gricean Semantics

According to Gricean semantics the fundamental notion of meaning is speaker-meaning. Speaker-meaning is what a speaker means by his utterance on an occasion of use. Speaker-meaning (for indicative speech) is metaphysically determined by the speaker's communicative intention that his audience entertains the proposition the speaker has in mind *via* the speaker's utterance. Roughly, a necessary and sufficient condition for (indicative) speaker-meaning is that the speaker has the following reflexive communicative intention (Grice 1989, 110):

(Speaker-meaning) Speaker S means proposition p iff in uttering u S intends:

(1) his audience, A, to entertain that p;

(2) A to recognise S's intention (1);

(3) A to recognise (1) by means of recognition of (2).

So, for my utterance 'Beer is nice' to mean *Beer is nice*, it's necessary that I intend my audience to entertain this proposition and that I intend my audience to recognise my intention by recognising that this is what I communicatively intend by my utterance. What's required for *communicative success* is that my audience entertains exactly this proposition and recognises that I intend them to entertain this proposition by my utterance.

Now, according to Griceans, regular speaker-meaning metaphysically determines *timeless meaning* (Grice 1989, 91). Timeless meaning is a word's or sentence's conventional meaning. So, for instance, the fact that speakers in a linguistic community regularly utter 'Beer is nice' to mean *Beer is nice* gives rise to something like a convention that 'Beer is nice' means *Beer is nice*.

Timeless-meaning isn't something like the content of an assertion or inscription[1]. Rather, according to Griceans, it's similar to an instruction or constraint to speaker-meaning: it tells speakers what *types* of propositions they can rationally intend to induce in their audience by a literal utterance of a sentence and it guides the audience in interpreting which proposition the speaker *might* literally mean by his utterance.

How does timeless-meaning constrain speaker-meaning? According to Griceans, it's a fact about human agency that speakers can only intend what they believe to be possible (Grice 1989, 98). Call this principle *Grice's principle*. Now, if speakers adhere to this principle, then they can't ordinarily use (say) the utterance 'Beer is nice' to mean *Wine is nice* because they don't believe it's possible for their audience to work out their intended meaning by exploiting the timeless-meaning of 'Beer is nice'[2]. Hence, a sentence's timeless-meaning constrains speaker-meaning by constraining what propositions speakers can rationally intend to communicate by their utterances.

How should 'possibility' be understood in the above formulation of Grice's principle? I suggest, following Buchanan, that for S to mean p by uttering u, S *must expect it reasonable* that A is capable of working out S's intended meaning from S's utterance (Buchanan 2010, 350). This formulation is admittedly vague, yet I think some precisification of Grice's principle along the lines of *what it's reasonable for speakers to expect* should be right. 'Possibility' can't be understood as physical or logical possibility because then the constraint on communicative intentions will be too weak such that Gricean semantics collapses into humpty-dumptyism—the view that you can use words to mean whatever you want them to. 'Possibility', however, also can't be understood as 'being very likely' because then the constraint on intentions will be too strong such that speakers can't use sentences non-literally to conversationally implicate.

3 Russell's Theory of Descriptions

According to Russell's theory of descriptions, definite articles (and possessive determiners such as 'my' or 'her') semantically express existence and uniqueness of their nominals (Russell 1905, 481). So for a sentence of the form 'The F is G' to be true, the description's nominal, 'F', must denote a set with at least one member and at most one member (i.e. a singleton set). Furthermore, this singleton set must be a subset of the set denoted by the VP 'is G'. A problem for Russell's theory, known as the problem of incompleteness (Strawson 1950, 332), is that

1. Griceans follow Strawson (Strawson 1950, 326) in maintaining that sentences don't express propositions; speakers *use* sentences to express propositions. See Neale (2004, 151).

2. There might of course be special contexts in which S can conversationally implicate *Wine is nice* by uttering 'Beer is nice'.

there are many cases in which speakers can intuitively say something true in uttering a sentence containing a description, although the description's nominal doesn't seem to denote a singleton set. Consider, for instance, Stephen telling Anita 'The beer is cold' at a philosophy conference. Intuitively, Stephen can say something true by this sentence if there's a single cold beer in the vicinity on the occasion of utterance[3]. According to Russell's theory, however, Stephen will almost trivially say something false because 'beer' doesn't denote a singleton set—there's more than one beer in the universe.

One popular response to this problem is to argue that the set denoted by 'beer' is restricted to a singleton set in context[4]. Stanley, for instance, suggests that for sentences such as the above, the description's nominal co-habits a node with a hidden indexical expression H[5] which denotes a contextually-salient set such that the sentence's logical form is as follows (Stanley 2002, 9):

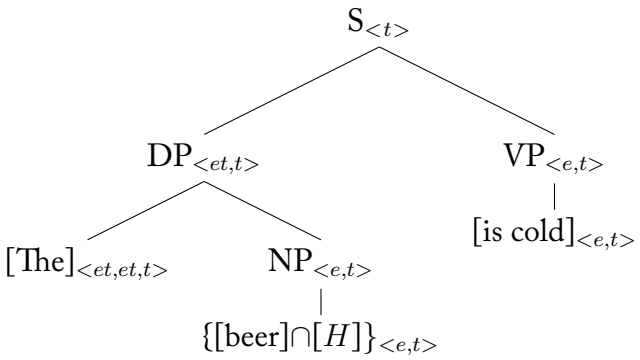

If the set of beers is intersected with the set denoted by H such that $\{\text{beer} \cap H\}$ denotes a singleton set, then Russell's theory predicts that 'The beer is cold' can be true in context although there are more than one beer in the universe.

4 Gricean Russellianism and the Multiple-Candidate Problem

A problem with Stanley's proposal is that there must be some explanation of how H's value is determined in context. For Gricean Russellians this explanation must be couched in intention-based terms. In particular, a Gricean will want to say that in uttering 'The beer is cold' the speaker communicatively intends his audience to entertain a complete proposition

3. 'The beer is cold' also has a generic-reading in which 'the beer' refers to a particular *kind* of beer (e.g. a particular brand or the set of beers at the conference). On this reading there needn't be a single cold beer in the vicinity on the occasion of utterance for 'the beer is cold' to be true. I will ignore generic definite descriptions in the confines of this essay. Ultimately, although I can't argue for it here, I think the problem I present in §3 also arises for generic-readings. See Carlson (1977) for a discussion of generic descriptions.

4. A different response is to say that 'The beer is cold' *semantically expresses* a false proposition, however, *communicates* (i.e. conversationally implicates) a true proposition (Blackburn 1984, 308–310). I won't have anything to say about this pragmatic strategy, although I think ultimately it gives rise to the same problem presented in §3.

5. Stanley's proposal is in fact more complex. According to him H is itself composed of f, a contextually-sensitive function from objects to sets, and i, a contextually-sensitive object f takes as argument to yield the restricting set. For purpose of this essay, I will ignore the nuances of Stanley's view.

$\exists!x(\ x$ is $\{\text{beer} \cap H\} \wedge x$ is cold$)$ [6], where H is a particular restricting set (e.g. {in this fridge}) which is a constituent of the speaker's communicatively intended proposition.

Now, Stanley's proposal raises a problem for Gricean Russellians—the multiple-candidate problem—because it seems implausible to think that in uttering 'The beer is cold' at the conference Stephen can have a complete proposition (with a definite value for H) in mind which he intends to communicate by his utterance (Schiffer 1995, 114–15).

That's because, as noted above, speakers can only intend what they believe to be possible. In the above example, however, there are many contextually-salient candidate propositions (corresponding to different values for H) which Stephen knows that Anita might reasonably interpret him as expressing by his utterance. For instance:

(1) $\exists!x(\ x$ is $\{\text{beer} \cap \text{you like}\} \wedge x$ is cold$)$
(2) $\exists!x(\ x$ is $\{\text{beer} \cap \text{from this corner shop}\} \wedge x$ is cold$)$
(3) $\exists!x(\ x$ is $\{\text{beer} \cap \text{for you}\} \wedge x$ is cold$)$
(4) $\exists!x(\ x$ is $\{\text{beer} \cap \text{at this conference}\} \wedge x$ is cold$)$
(5) $\exists!x(\ x$ is $\{\text{beer} \cap \text{in this fridge}\} \wedge x$ is cold$)$
(6) $\exists!x(\ x$ is $\{\text{beer} \cap \text{on the top shelf of this fridge}\} \wedge x$ is cold$)$
(...) [7] [8]

Given the sheer amount of candidate propositions in context, Stephen knows that he can't reasonably expect Anita to work out any proposition of (1)-(6) in particular by his utterance. Hence, he can't mean a definite proposition—say (1)—by 'the beer is cold' without violating Grice's principle. Communicative success wouldn't be a matter of chance or luck, but something close to a miracle. *Ipso facto*, Gricean semantics, according to which the objects of speaker's intentions are definite propositions, is flawed[9]; speakers can't always intend to communicate definite propositions by their utterances.

6. This representation of the logical form of 'the beer is cold' is a bit inaccurate. It would be more precise to represent it as $\exists!x(\ x$ is $\{\text{beer} \cap H\}) \wedge \exists y(\ y$ is $\{\text{beer} \cap H\} \wedge y$ is cold$)$. To avoid an excessively complicated presentation, I will ignore this matter in what follows.

7. The list of propositions the speaker could be understood as meaning in the above context could I think potentially go on indefinitely.

8. The anonymous reviewer and Kimon Sourlas-Kotzamanis note, that my description of the party scenario is somewhat underdescribed and that if we were to flesh out the situation and context, the list of potential can- didate propositions Stephen might be understood as meaning will be significantly smaller. I agree with them that fleshing out the context will reduce the list of potential candidate propositions. Yet it seems to me that even if we were to fully elaborate the context, the list of potential candidate propositions would still be far too large for Stephen to mean one definite proposition in particular.

9. One might think that this problem refutes Russell's theory of descriptions, rather than Gricean semantics.

I disagree. The multiple-candidate problem arises for many constructions besides definite descriptions. For instance, for non-sentential assertions such as 'Tiger!' (uttered when a tiger is coming) the explicit linguistic material also underdetermines the proposition expressed, such that *prima facie* speakers can't mean a single complete proposition by their utterance. This suggests that the multiple-candidate problem undermines Gricean semantics, rather than the theory of descriptions.

5 Gricean Replies to the Multiple-Candidate Problem

If the Gricean picture of communication is to be maintained, we must find a proposition corresponding to 'the beer is cold' the speaker can mean and the audience must have entertained for communication to be successful.

(A) One suggestion is to maintain that in contexts such as the above—contexts in which there are prima facie too many description completions—there's something like a convention that speakers usually mean object-dependent propositions by their utterances. Following Neale's Russell-friendly account of referential uses of descriptions (Neale 2004, 171), we could then maintain that in uttering 'the beer is cold' Stephen communicatively intends Anita to entertain $\exists!x(\, x \text{ is } \{\text{beer} \cap x = o\} \wedge x \text{ is cold})$, where o is a particular beer (the object of the proposition) which is the object of Stephen's communicatively intended proposition. If there's such a convention of usage, then Stephen can mean a definite object-dependent proposition by his utterance without violating Grice's principle—there's only one candidate proposition salient in context.

The main problem with this suggestion is that it lacks generality (Bowker 2017, 7). *Prima facie*, there are many cases for which the multiple-candidate problem arises where we can't plausibly say that the speaker meant an object-dependent proposition by his utterance.

Consider, for instance, Stephen and Anita on their way to the conference. Now, imagine Stephen uttering 'the beer, *whichever it is*, is cold'. In that case, there will be multiple propositions 'the beer, *whichever it is*, is cold' could express in context. Intuitively, however, none of them will be object-dependent propositions because *prima facie* Stephen is making a *general* statement about whatever it is which has the property 'beer'. This example suggests that the above response to the multiple-candidate problem is at best incomplete.

(B) Another similar response by Michaelson is to suggest that for sentences containing incomplete descriptions there's something like a convention that speakers usually mean propositions in which H refers to the speaker's *de dicto* mental state (Michaelson 2016, 13). So, in uttering 'The beer is cold', it's standard usage that Stephen will mean the proposition $\exists!x(\, x \text{ is } \{\text{beer} \cap$ in the speaker's mind$\} \wedge x$ is cold$)$. If that's the case, then Stephen can mean a definite proposition by his utterance of 'the beer is cold' without violating Grice's principle—$\exists!x(\, x \text{ is } \{\text{beer} \cap$ in the speaker's mind$\} \wedge x$ is cold$)$ is the only proposition salient in context. A nice thing about Michaelson's suggestion is that, unlike Neale's, it doesn't only meet the multiple-candidate problem for object-dependent propositions.

The problem with Michaelson's response[10] is that it leads to humpty-dumptyism and hence should be unappealing to Griceans. Consider, for instance, Stephen at the conference. For what ever reason he wants to communicate to Anita the proposition *The beer which is in the safe in Kripke's office is cold* and produces 'The beer is cold'. Can Stephen mean this proposition by

10. I don't mean to imply that Michaelson thinks that (B) is a satisfactory response to the multiple-candidate problem. In fact in personal correspondence he agreed with me that (B) isn't a viable response because it collapses into humpty-dumptyism.

his utterance without violating Grice's principle? According to Michaelson's suggestion we'd have to say yes, since Stephen has the proposition involving Kripke's safe in mind. Hence, Anita should be able to identify this proposition in a semi-opaque manner *via* Stephen's *de dicto* mental state.

This seems implausible. If Stephen can mean *The beer which is in the safe in Kripke's office is cold* by his utterance, then there's a real danger that Stephen can mean any proposition of the form $\exists!x(\ x \text{ is } \{\text{beer} \cap H\} \wedge x \text{ is cold})$ so long as he intends Anita to identify the value for H via his *de dicto* mental state. I take it, therefore, that Michaelson's suggestion isn't a viable response for Griceans.

6 Buchanan's Solution to the Multiple-Candidate Problem

Although I can't claim to have considered all responses to the multiple-candidate problem, I think the Gricean picture of communication can't be maintained as it is.

One helpful suggestion made by Buchanan, in response to this problem, is to say that sometimes the object of a speaker's communicative intention is a *restricted proposition-type* (and not a proposition) and that communicative success merely requires the audience to entertain at least one proposition which is a token of the speaker's communicatively intended restricted proposition- type (Buchanan 2010, 358–59).

A proposition-type is the property common to all propositions which a sentence with a discrete logical form can be used to express in different contexts. So, for instance, the proposition-type for 'The beer is cold' is $\exists!x(\ x \text{ is } \{\text{beer} \cap \_\} \wedge x \text{ is cold})$, where '_' is a gap to be completed by different values corresponding to different token propositions this sentence can express in different contexts. Propositions (1)-(6) are all tokens of the type $\exists!x(\ x \text{ is } \{\text{beer} \cap \_\} \wedge x \text{ is cold})$.

A *restricted* proposition-type is a proposition-type for which the value of '_' has been restricted in context by the speaker's preferences. So, for instance, for 'The beer is cold' at the conference, the value of '_' might be restricted such that (1) is a token of the speaker's restricted proposition-type, but not (say) $\exists!x(\ x \text{ is } \{\text{beer} \cap \text{belongs to George Bush}\} \wedge x \text{ is cold})$.

Now if Buchanan's account of speaker-meaning is correct, then Stephen doesn't violate Grice's principle when he utters 'The beer is cold' at the conference: it's perfectly reasonable for him to have the expectation that Anita works out at least one proposition which is a token of his restricted proposition-type. If Anita interprets Stephen as meaning $\exists!x(\ x \text{ is } \{\text{beer} \cap \text{at the conference}\} \wedge x \text{ is cold})$, but not the other tokens—say (2)-(6)—of Stephen's restricted proposition-type, then this is sufficient for her to understand Stephen and for communication to be successful. Hence, the multiple-candidate problem is overcome.

A problem with Buchanan's proposal is that he leaves a few things unexplained. According to Griceans, as noted in §2, successful communication consists in A entertaining proposition p and *recognising* that S intends them to entertain p *via* recognition of this intention. Now, if on Buchanan's picture A merely needs to entertain a *token proposition* of S's restricted *proposition-*

type, then presumably successful communication needn't consist in *A fully recognising S*'s communicative intention to have A entertain a token of his restricted proposition-type (Bowker 2017, 8). Hence, Buchanan will still have to work out the details a watered down version of **speaker-meaning** which presumably only involves S having the intention that *A partially recognises* the speaker's communicative intention by his utterance.

Another general concern, that I intend to address in the rest of this essay, is that Buchanan says very little about what restricts the speaker's proposition-type. One might think, following Michaelson, that a restricted proposition-type just is a cluster of propositions the speaker *has in mind*; each proposition of this cluster is a proposition the speaker has a representation of and is happy to be understood as meaning by his utterance[11].

This simple suggestion, however, won't do. To see this, consider again Stephen uttering 'the beer is cold' in order to inform Anita that they might drink beer after the conference. Furthermore, assume that the conference is held in a bungalow. Lastly, assume that for whatever reason Stephen doesn't possess the concept 'bungalow'. Since Stephen doesn't possess this concept, it seems that if a restricted proposition-type is a certain cluster of propositions the speaker *has in mind*, then $\exists!x(\ x \text{ is } \{\text{beer} \cap \text{in the bungalow}\} \wedge x \text{ is cold})$ won't be a token of Stephen's restricted proposition-type. Now, assume that for whatever reason Anita interprets Stephen as meaning $\exists!x(\ x \text{ is } \{\text{beer} \cap \text{in the bungalow}\} \wedge x \text{ is cold})$ by his utterance.

Is communication successful? According to Michaelson's interpretation of restricted proposition-types communication can't be successful, since the proposition Anita entertains by Stephen's utterance is a proposition Stephen doesn't have the concepts to represent; Anita's proposition, therefore, can't be a token of Stephen's restricted proposition-type.

This seems implausible. It's intuitive to think that communication may still be successful in the above scenario, so long as Anita grasps a proposition of type $\exists!x(\ x \text{ is } \{\text{beer} \cap \_\} \wedge x \text{ is cold})$ and knows from Stephen's utterance that she might drink beer with Stephen after the conference. Life will go on perfectly well: Stephen won't notice that Anita entertained a proposition he doesn't have the concepts to represent and after the conference they might both have beer together which seems to be the main purpose of Stephen's communicative act. This suggests that Michaelson's interpretation of restricted proposition-types can't be right; communicative success doesn't require there to be an overlap in the propositions speaker and audience *have in mind*[12].

7 Proposition-types restricted by Act-Coordinating Intentions

Following Paul (1999, 158), I maintain that (at least in part) communication is a means of coordinating actions between speaker and audience; speakers utter words or sentences in conversational contexts with higher-order intentions to coordinate their actions with their audience's.

11. Michaelson (2016, 13) seems to assume this interpretation when he says that according to Buchanan '[speakers have a] cluster of propositions [...] in mind [they intend] to convey'.

12. This argument is indebted to Michaelson (24)

Call these higher-order intentions *act-coordinating-intentions*. So, for instance, when Stephen utters 'the beer is cold' we might think of him as having the *act-coordinating intention that*, as the result of his utterance, Anita is *disposed to act* in coordination with his intended action[13] of (say) potentially drinking beer with her after the conference.

Now, to provide an improved account of restricted proposition-types which meets the above difficulty with Michaelson's interpretation, I suggest that for utterances of sentences with incomplete descriptions speakers communicatively intend their audience to entertain at least one proposition of their intended proposition-type. What restricts this proposition-type is the speaker's act-coordinating-intention; speakers intend their audience to entertain any proposition of their proposition-type the entertainment of which disposes the audience to act in coordination with the speaker's intended action[14]. So, for instance, when Stephen utters 'the beer is cold', then he intends Anita to entertain at least one proposition of the type $\exists!x(\ x$ is $\{$beer $\cap$ $\_\}\ \wedge\ x$ is cold$)$. He also intends—this is what restricts the proposition-type—that the proposition Anita entertains is a proposition the entertainment of which disposes her to act in coordination with Stephen's intended action of (say) potentially drinking beer with her after the conference[15].

How can the entertainment of a proposition dispose the audience to act in coordination with the speaker's intended action? I suggest that the entertainment of a proposition can supply the audience with sufficient information such that they're able to *infer* the speaker's intended action. Knowing the speaker's intended action (*via* this inference) is to have a disposition to act in coordination with it. So, for instance, if Stephen utters 'the beer is cold' and Anita thereby entertains the proposition $\exists!x(\ x$ is $\{$beer $\cap$ in this fridge$\}\ \wedge\ x$ is cold$)$, then she can use this proposition to infer Stephen's intended action. Since she believes that Stephen meant $\exists!x(\ x$ is $\{$beer $\cap$ in this fridge$\}\ \wedge\ x$ is cold$)$, she can reason that he probably intends to drink a beer with her after the conference—why else would Stephen mean the above proposition at the conference? Hence, she'll have a disposition to act in coordination with Stephen's intended action.

If act-coordinating intentions restrict proposition-types, then communication will presumably succeed in the conference scenario if Anita entertains $\exists!x(\ x$ is $\{$beer $\cap$ in this fridge$\}\ \wedge$ x is cold$)$, but not if she entertains $\exists!x(\ x$ is $\{$beer $\cap$ belongs to George Bush$\}\ \wedge\ x$ is cold$)$ by

13. The object of a speaker's act-coordinating intention may also be an intended *inaction*. This leaves it open whether omitting to do something is an action.

14. I suggest that act-coordinating intentions are reflexive just like ordinary Gricean communicative intentions. For S to have the intention that A be *disposed to act* in coordination with S's intended action, it's *prima facie* necessary that S also intends that A *recognises* S's act-coordinating intention. If there's no such recognition involved, it's hard to see how A's actions should properly be *coordinated* with S's actions.

15. Sourlas-Kotzamanis notes that act-coordinating intentions might also help to explain Grice's *maxim of relation* to be relevant to the communicative purpose. The idea might be that speakers (tacitly) agree on some act-coordinating goal in a communicative exchange and that the maxim of relation states that speakers should make conversational contributions which are relevant to this shared act-coordinating goal. I am sympathetic to this suggestion. Grice himself is rather reluctant in 'Logic and Conversation' to satisfactorily explicate the maxim of relation (Grice 1989, 27).

Stephen's utterance. The former proposition, unlike the latter, provides Anita with sufficient information to infer Stephen's intended action and therefore disposes her to act in coordination with it.

Communication will presumably also succeed if Anita entertains $\exists!x(\ x$ is $\{$beer $\cap$ in the bungalow$\} \wedge x$ is cold$)$ —a proposition Stephen doesn't have the concepts to represent—so long as the entertainment of this proposition disposes her to act in coordination with Stephen's intended action. Since act-coordinating intentions restrict proposition-types, Stephen needn't have the token propositions in mind which Anita *could* entertain in order to understand him. Hence, on my proposed view, communication may succeed although there's no overlap in the propositions speaker and audience have in mind.

Presumably, if successful communication consists in part in the audience being disposed to act in coordination with the speaker's intended action, then communicative success will have *degrees*[16]. That's because, presumably, acting in coordination with the speaker's actions has degrees. For instance, consider Stephen uttering 'the beer is cold' thereby intending Anita to entertain a proposition of type $\exists!x(\ x$ is $\{$beer $\cap$ _$\} \wedge x$ is cold$)$ and intending that the entertainment of this proposition disposes her to act in coordination with his intended act of drinking *Guinness* with her after the conference.

Now, if Anita were to entertain $\exists!x(\ x$ is $\{$beer $\cap$ in the bungalow$\} \wedge x$ is cold$)$ by his utterance, then presumably communication wouldn't be *fully* successful. That's because, on the face of it, the entertainment of $\exists!x(\ x$ is $\{$beer $\cap$ in the bungalow$\} \wedge x$ is cold$)$ in context doesn't provide Anita with sufficient information to infer Stephen's *exact* intended action. Presumably, she'll only expect, by entertaining the above proposition, that she might drink beer (but not necessarily *Guinness*) with Stephen after the conference.

I think folk-intuitions support the idea that communicative success has degrees. We often say things like 'he didn't understand me completely' or 'he roughly understood what I meant'. My proposal explains such talk as follows: the proposition A entertained by S's utterance doesn't fully dispose A to act in coordination with S's intended action.

8 Conclusion

I have argued that, at least for definite descriptions, unarticulated indexical expressions don't pose a problem for Griceans. We can all be Gricean Russellians so long as the objects of communicative intentions are restricted proposition-types. The multiple-candidate problem merely shows that sometimes speakers can't communicatively intend discrete propositions by their ut-

16. Because speakers intended actions often seem vague—when S utters 'the beer is cold' it seems plausible to think that it will be indeterminate whether S's intended action is (i) to potentially drink beer with A after the conference or whether it is (ii) to potentially drink beer with A *on the sofa* after the conference,—a speaker's act-coordinating intention will presumably often be vague too.

I don't think this vagueness is problematic, however. Some act-coordinating intentions may just be irredeemably vague, such that it's indeterminate whether communication was *fully* or *partially* successful.

terances.

I have argued that restricted proposition-types are proposition-types restricted by speaker's *act-coordinating-intentions*. This proposal explains how communication may succeed, although the proposition the audience entertains by means of the speaker's utterance is a proposition the speaker doesn't have the concepts to represent[17].

References

Blackburn, Simon. 1984. *Spreading the Word: Groundings in the Philosophy of Language.* Clarendon Press.

Bowker, Mark. 2017. 'Saying a Bundle: Meaning, Intention, and Underdetermination'. *Synthese.* doi:10.1007/s11229-017-1652-0.

Buchanan, Ray. 2010. 'A Puzzle About Meaning and Communication'. *Noûs* 44 (2): 340–71.

Carlson, Gregory N. 1977. 'Reference to Kinds in English'. PhD diss., University of Massachusetts, Amherst.

Grice, Paul. 1989. *Studies in the Way of Words.* Cambridge, MA: Harvard University Press.

Michaelson, Eliot. 2016. 'Vague Thoughts, Vague Quantifiers [draft]'. https://www.dropbox.com/s/mxn0e8pu6pl3srz/PuzzleMeaning3.pdf?dl=0.

Neale, Stephen. 2004. 'This, That and the Other'. In *Descriptions and Beyond,* edited by Anne Bezuidenhout and Marga Reimer, 68–182. Oxford: Oxford University Press.

Paul, Matthias. 1999. *Success in Referential Communication.* Springer.

Russell, Bertrand. 1905. 'On Denoting'. *Mind* 14 (56): 479–93.

Schiffer, Stephen. 1995. 'Descriptions, Indexicals, and Belief Reports: Some Dilemmas (but Not the Ones You Expect)'. *Mind* 104 (413): 107–131.

Stanley, Jason. 2002. 'Nominal Restriction'. https://semanticsarchive.net/Archive/jhhYjdmZ/NominalRestriction.pdf.

Strawson, P. F. 1950. 'On Referring'. *Mind* 59 (235): 320–44.

17. I would like to thank Ray Buchanan, Eliot Michaelson, Kimon Sourlas-Kotzamanis, Mark Textor and an anonymous reviewer for helpful comments and discussion.