


Issues in Objectivity and Mind-Dependence

by

Ekaterina Botchkina

A.B., Harvard College (2010)



Submitted to the Department of Linguistics and Philosophy
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy in Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2016

© Ekaterina Botchkina. 2016. All rights reserved.

Signature redacted

Signature of Author: _____
Department of Linguistics and Philosophy
September 2016

Signature redacted

Certified by: _____
Stephen Yablo, David W. Skinner Professor of Philosophy
Thesis Supervisor

Signature redacted

Sally Haslanger, Ford Professor of Philosophy
Thesis Supervisor

Signature redacted

Accepted by: _____
Roger White, Professor of Philosophy

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.



77 Massachusetts Avenue
Cambridge, MA 02139
<http://libraries.mit.edu/ask>

DISCLAIMER NOTICE

Due to the condition of the original material, there are unavoidable flaws in this reproduction. We have made every effort possible to provide you with the best copy available.

Thank you.

The images contained in this document are of the best quality available.

Issues in Objectivity and Mind-Dependence

by

Ekaterina Botchkina

Submitted to the Department of Linguistics and Philosophy
on September 2, 2016 in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy in Philosophy

Abstract

Reality and objectivity are often characterized in terms of independence from the mind: the first-pass idea is that what it takes for any particular subject matter to be real and objective is for facts about it to obtain independently of beliefs, linguistic practices, conceptual schemes, and so on. But if we take seriously the possibility that significant realms of reality, including social kinds, judgment-dependent properties, and mental phenomena themselves, stand in various dependence relations to the mental, then this first-pass characterization needs to be significantly revised. In this set of papers, I consider the special questions that metaphysically mind-dependent entities raise for issues of objectivity and realism.

In Part I, I substantiate the notion of metaphysical mind-dependence with a taxonomy of the various ways in which entities can stand in metaphysical relations of dependence to mental phenomena.

In Part II, I address the question of *realism* and mind-dependence: I argue that while certain entities stand in relations of significant, direct, and essential dependence on mental activity, they are nevertheless fully real. In making the argument, I elaborate a distinction between *enactive* and *essential* dependence on mental phenomena, arguing that both kinds of dependence may obtain without impinging on an entity's reality.

In Part III, I address the question of *objectivity* and mind-dependence: I argue that certain kinds of mind-dependence, in particular, dependence on judgments, have the effect of undermining the objectivity of the relevant domain. One consequence of the view I develop is that the objectivity of a subject matter can come apart from the reality of its associated entities; another is that objectivity is a feature that is relative, rather than absolute, and depends crucially on which perspectives are brought to bear for the purposes of evaluation.

Thesis Supervisors: Stephen Yablo, David W. Skinner Professor of Philosophy
Sally Haslanger, Ford Professor of Philosophy

Contents

I.	Varieties of Metaphysical Mind-Dependence	4
II.	Realism and Mind-Dependence	18
III.	Objectivity and Mind-Dependence	48

Part I: Varieties of Metaphysical Mind-Dependence

The notion of *mind-dependence* is by now familiar: the idea is that certain entities are best described as reliant, in some sense, on sensory responses, judgments, agreements, beliefs, or other mental phenomena. But there hasn't been, to my knowledge, much by way of systematic categorization of what specific forms this dependence relation might take; as things stand, different writers conceive of the landscape in very different ways.

Here is a small glimpse at how existing positions cross-cut and overlap, both with respect to differentiating types of mind-dependence and to drawing conclusions about what these types imply for the objectivity and reality of the entities they describe. Gideon Rosen identifies *causal* and *existential* mind-dependence as categories distinct from *conceptual* mind-dependence; he takes the former two notions to be unproblematically compatible with reality and objectivity.¹ Ralph Wedgwood, on the other hand, argues that the kind of mind-dependence that would threaten objectivity is *essential* mind-dependence,² a distinction that, as I'll discuss below, cuts across both of Rosen's former categories. Finally, Amie Thomasson's discussion of the issue doesn't explicitly differentiate along any of the lines above, but the objects whose objectivity and reality she questions belong to both the *causal* and *existential* categories of Rosen's.


In this brief study, I offer a broad outline of ways in which an entity might be metaphysically mind-dependent, in a way that would allow us to situate and

¹ (Rosen 1994)

² (Wedgwood 1997)

compare the various distinctions mentioned above while bringing into view other kinds of dependence that are often overlooked.

Metaphysical vs. non-metaphysical mind-dependence



I'll begin by saying a bit about what differentiates *metaphysical* mind-dependence, the kind I focus on in this set of papers, from other varieties. To that end, here is a distinction made by Ralph Wedgwood between two broad ways of thinking about mind-dependence: "one is an account of a certain way of *thinking* of things' being F; the other is an account of what it is for things to *be* F."³ The first variety may be elaborated as "a semantical account of the meaning of the term 'F'; [or] an epistemological or psychological account, of the way in which we think about, or achieve epistemic access to, an object's being F." The second variety, on the other hand, is "an *ontological* or *metaphysical* account, of what it is for things to be F. It is stated entirely at the level of reference, not at the level of sense."

Wedgwood argues that the second kind of account, the metaphysical one, is more apt for capturing our intuitive ideas about mind-dependence, and, in particular, better suited for articulating theses about the relation of objectivity and reality to mind-dependence: "the core of any adequate conception of response-dependence must be an account of *response-dependent properties*, not of response-


³ (Wedgwood 1997)

dependent concepts. This ... is the basic problem with the conceptions of response-dependence that have been developed by Mark Johnston, Philip Pettit, and Crispin Wright: they are accounts of response-dependent concepts, not of response-dependent properties.”

In what follows, I too will focus on the metaphysical conception of mind-dependence – not because I regard, like Wedgwood, conceptual accounts to be somehow inapt, but because the metaphysical conception alone leaves us more than enough material to work with, standing in need of considerable disambiguation and clarification.

One final note on Wedgwood before moving on to the taxonomy: having drawn a distinction between metaphysical and conceptual accounts and dismissed the latter as inadequate, Wedgwood settles on the “essential” conception of mind-dependence, on which “a property is response-dependent just in case it is an essential part of something’s being an instance of the property that it stands in some relation to some sort of mental response to that property.” This might suggest that the metaphysical conception of mind-dependence just *is* the essential conception: that to be metaphysically mind-dependent is to essentially stand in some relation to some mental response. But as I discuss below, the essential conception is far from exhaustive of the different kinds of strictly metaphysical varieties of mind-dependence that we can identify; in what follows, I’ll situate Wedgwood’s essential conception among several other varieties of metaphysical mind-dependence, and further disambiguate between different kinds of *essential* mind-dependence.

Contingent vs. Necessary Mind-Dependence



Contingent mind-dependence

In discussing paradigmatically mind-independent phenomena, writers often speak of entities that not only could have, but in fact do (or did), exist independently of minds: entities like mountains, dinosaurs, and distant stars. What tends to be overlooked is the category of the *contingently mind-dependent*: entities that do, as a matter of fact, owe their existence to mental phenomena, but could have existed independently of minds.


Consider the dog breed *German Shepherd*, understood as a set of dogs sharing a certain lineage and exhibiting a particular phenotype: domed forehead, square-cut muzzle, black nose, and so on. Perhaps it would have been possible, if very unlikely, for that very lineage and phenotype to have arisen without human intervention: just by certain dogs interbreeding in certain ways. If our best understanding of the kind *German Shepherd* indeed allows for that possibility, then it picks out entities that are contingently mind-dependent: mental activity figures in their actual causal etiology, but not their causal etiology in all possible worlds.⁴ Or take chloroflourocarbons, a family of gases synthesized for diverse household and

⁴ Contrast this with an account of dog breeds that makes essential reference to the kinds of behaviors that particular breeds are developed to perform, described in intentional (and thus mind-dependent) terms: retrievers are dogs of a particular lineage who were trained to retrieve; German shepherds are dogs that were trained to herd, and so on. In this picture, the relevant dog breeds are not contingently but *necessarily* mind-dependent: a phenotypically-identical dog that didn't issue from a lineage intentionally trained to perform particular actions wouldn't count as true German shepherd.

industrial applications, with no existing natural source. The actual causal etiology of CFCs involves mental activity, but it would be highly implausible to build that etiology into what it is to *be* a CFC – for the latter, it is sufficient to be something like “an organic compound that contains only carbon, chlorine, and fluorine, produced as a volatile derivative of methane, ethane, and propane.”⁵


Within the general category of the contingently mind-dependent, we may distinguish further subtypes by identifying the *kind* of relation to the mind that contingently obtains. To begin with, let’s distinguish between causal and non-causal dependence on the mind; the latter category will be further disambiguated below.

[Contingent] causal mind-dependence:



A kind is contingently causally mind-dependent if the actual causes, but not all possible causes, of that kind’s existence include mental phenomena. Our examples above of German Shepherds and CFCs are contingently causally mind-dependent.

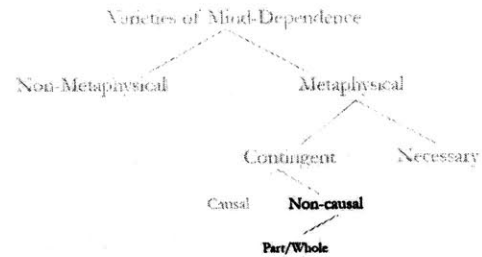
[Contingent] non-causal mind-dependence:



⁵ <https://en.wikipedia.org/wiki/Chlorofluorocarbon>

Not all relations of dependence on the mind are causal ones; below are varieties of contingent mind-dependence in which mental phenomena stand in some other relation to the entities depending on them.

[Contingent] part/whole mind-dependence:

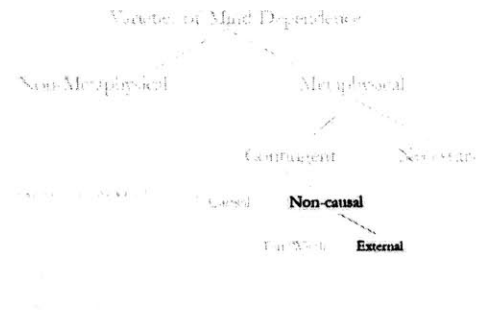


An entity is *part/whole* mind-dependent if it contains mental phenomena as one or some of its parts; it is *contingently* part/whole mind-dependent if its mental parts aren't necessary for its existence.⁶ Consider the security system in a particular apartment building, consisting of a network of cameras and monitors, a doorbell, a door-unlocking mechanism, and a doorman overseeing their operation. The doorman, as I'm conceiving of things, doesn't just *cause* the operation of the security mechanism: he's *part* of it. But we can suppose that fairly soon, the doorman's function will be performed perfectly well by a computer: he's part of a system that will persist, unfortunately, without him. If that's a plausible way to regard the identity conditions of that particular security system, then it is a *contingently part/whole mind-dependent* entity⁷.

⁶ So the possibility of such entities requires the fairly weak metaphysical assumption that an object can maintain its identity in the absence of some of its parts.

⁷ The doorman in the example is a proper part of the entity in question; are there contingently part/whole mind-dependent entities that consist *entirely* of mental phenomena? That depends on one's underlying ontology. If, for instance, kinds can be defined by external relations, like *Mary's favorite things*, then it seems like a kind could indeed contingently consist entirely of mental phenomena.

[Contingent] external-relation mind-dependence:



Not every non-causal dependence on the mind is of the part/whole variety: some involve entities that stand in a certain external relation, like recognition or disposition to provoke a response, to mental phenomena. Consider a kind whose identity conditions include, together with certain mind-independent criteria, being regarded a certain way by an observer. For instance, consider a possible kind, call it *ovalish*, with the following stipulated disjunctive identity conditions: to be ovalish is either to be a describable by a differential, simple, convex plane curve with two angles of symmetry, or to be such as to remind one of an egg. If a certain figure is ovalish in virtue of meeting *both* criteria, then it is contingently external-relation mind-dependent: its actual status of being ovalish is due, in part, to being regarded a certain way, but it would still be ovalish even if no one was disposed to be reminded of an egg when they saw it.


That concludes our overview of (at least some) types of contingent mind-dependence; but before moving on to the category of necessary mind-dependence, let's pause to return to the question of realism. Is it possible that the mind-dependence of contingently mind-dependent entities compromises, in any way, their reality?

Recall that what marks out the category of the contingently mind-dependent is that the entities in question *are*, but *need not have been*, dependent in some way on mental phenomena. That means that there are possible worlds where those very kinds exist but are fully independent of minds. Take, for example, the possible world in which German Shepherds exist as a result of a chance mating patterns, rather than by deliberate breeding activity – call the creatures in this world Natural Shepherds. Is there any room to argue that Natural Shepherds enjoy a degree of reality or objective existence that is somehow more robust than that of German Shepherds? It's not very easy to see how that could be the case. After all, we've granted that Natural Shepherds are the very same *kind* as German Shepherds. And of course we grant that German Shepherds *exist* around here. But if we grant that they *exist* and are the very *same kind* of creatures as in worlds in which they're not dependent on minds, it becomes difficult to leave room for the claim that their dependence on minds somehow affects their reality. To go that route, we'd have to argue that reality is not a function of what *kind of thing* an entity is, or whether it exists, but a function of what brings it about.

But if it's right that the contingently mind-dependent is, as a rule, just as real or objective as the mind-*independent*, then that would help to unify and explain many of the examples often cited as being obviously 'benign' or 'mundane' forms of mind-dependence, posing no threat to realism.⁸


⁸ For instance, an object being in a particular location (Jenkins 2005), global warming and the depletion of the ozone layer (Rosen 1994), synthetic gases (Mason 2016), and many more.

Necessary mind-dependence:



What's more commonly discussed in the literature on mind-dependence are entities that are *necessarily* mind-dependent: entities that depend on minds in any possible world in which those entities exist.

Taking to heart a lesson from Kit Fine, we might take care to distinguish necessary mind-dependence described in merely modal terms from necessary mind-dependence that is necessary because *essential*: in other words, from mind-dependence that stems from the essence or nature of the mind-dependent property:⁹




Entities whose mind-dependence is necessary in merely the modal sense might be mind-dependent for reasons that do not issue from their own natures and essences: for instance, the necessary mind-dependence of such Fs might be explained by citing an essential feature of *minds*, rather than by citing an essential

⁹ (Fine 1994)

features of Fs.¹⁰ By contrast, essentially mind-dependent entities are those whose very natures – in Wedgwood’s terms, whose real definitions – require dependence on the mind.

Setting merely modally necessary mind-dependence to one side, let’s focus on the category of essential mind-dependence. Here we can make some distinctions similar to the ones we encountered above, when considering the contingently mind-dependent.

[Essential] causal mind-dependence:



This category includes kinds that, to be the very kinds that they are, must have causes that include mental phenomena. Consider a theory of art on which an object counts as, say, a painting only if it was produced as a result of some intentional action. If some paint canisters, knocked about by the wind, dripped out a perfect physical replica of Pollock’s *Untitled (c. 1950)*, the result would nevertheless, on this view, fail to be a painting, since it would lack the right causal etiology.¹¹

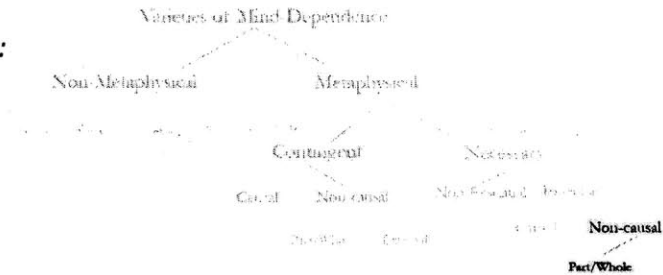
[Essential] non-causal mind-dependence: Once again, not every essential dependence relation will be a causal one. Rosen, for instance, contrasts causal

¹⁰ See (Jenkins 2005) for examples and further discussion of this distinction; Jenkins argues, like Wedgwood, that it is essential mind-dependence, rather than the kind described in merely modal terms, that is relevant to questions of objectivity.

¹¹ See also functional and intentional theories of artifacts, as those in (Levinson 1979) and (Bloom 1996).

dependence with a category he calls *existential mind-dependence*, characterizing entities that “could not exist or obtain if there were no empirical minds or mental activity.” But we can distinguish several non-causal senses in which objects “could not exist or obtain” without dependence on the mental.

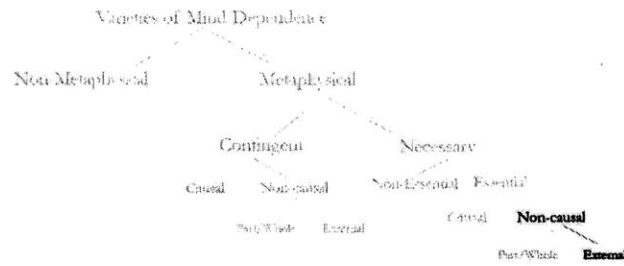
[Essential] part/whole mind-dependence:



This category includes entities that, to be the very entities that they are, must include, as parts, mental phenomena. This is a very different kind of relation from the causal dependence discussed just above: for instance, *Untitled (c. 1950)* depends causally on mental activity, but does not contain any mental activity as a proper part: the painting itself includes just the paint and the canvas. Examples of essential part/whole mind-dependence include both entities that have mental phenomena as proper parts, like democracies and revolutions, and those entities consisting entirely of mental phenomena, like sensations of redness and indeed, minds themselves.¹²

¹² Here we have a counter-example to a crude view on which *any* essential mind-dependence is inimical to realism: minds themselves are essentially part/whole mind-dependent, but the sorts of things that most antirealist accounts are careful to exclude.

[Essential] external-relation mind-dependence:



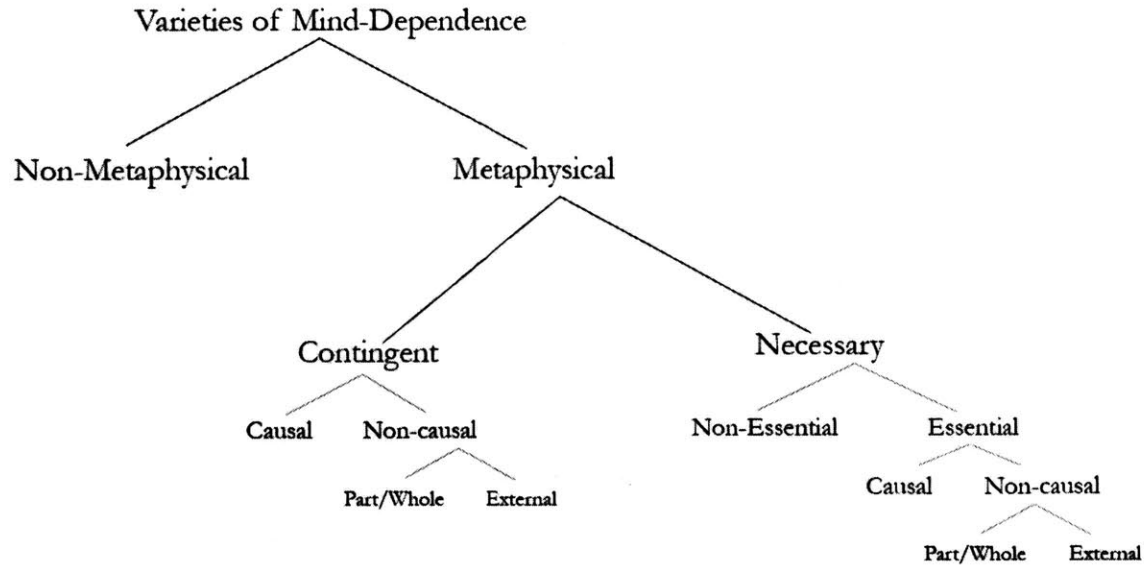
An entity is essentially, external-relationally mind-dependent if, to be the very entity that it is, it must stand in some external relation to some mental phenomena; for instance, it must be regarded (or acknowledged, or believed, and so on) to be a certain way, or it must have the disposition to induce a certain mental response.¹³ Many response-dependent accounts of colors¹⁴ and values¹⁵ fall under the category of essential, external-relational mind-dependence. Another example, to be discussed in much further detail below, is Haslanger’s account of gender of race, on which to be a woman, for instance, is to be subject to certain kinds of social treatment on the basis of perceived physical characteristics.¹⁶

¹³ The rudimentary theory of art that we considered above could be expanded to include external-relational mind-dependence, not just causal mind-dependence: on such a view, to be art would be, in part, to be regarded a certain way. One advantage of this expansion would be that it would allow us to include as art found objects, and not just ones produced for a certain purpose.

¹⁴ E.g. (Johnston 1989, 1993)

¹⁵ (McDowell 1985); (Brower 1993).

¹⁶ Sally Haslanger, 2000. “Gender and Race: (What) are they? (What) do we want them to be?”



Thus concludes our outline, with the important proviso that it makes no claim to be exhaustive of the different ways in which an entity might be mind-dependent: there are surely finer distinctions to be drawn within the broad categories – part/whole vs. external, causal vs. non-causal – that I have identified. But I have found it to be a useful first approximation of the various general approaches that a metaphysical mind-dependence account might take, and a helpful aid in bringing the relations between various existing accounts into view.¹⁷

¹⁷ For instance, I interpret Ian Hacking’s “looping effect” account of human kinds (Hacking 1996) as an interesting example of both essential external-relation mind-dependence *and* contingent causal mind-dependence. According to Hacking, kinds like “child abuse” and “schizophrenic” come into existence in part by being the objects of human classificatory practices. So part of what it *is*, on this view, to be, say, a child abuser is to fall under certain criteria marked as salient by certain social groups. But it turns out that these classificatory practices also play a *causal* role in helping to bring about further instances of the kind in question.

Part II: Realism and Mind-Dependence

I'll now turn to a concrete example from the literature of mind-dependence accounts – John Searle's theory of institutional kinds – and an interpretation of that theory by Amie Thomasson as yielding a picture of those kinds on which they are less than fully real. Part of my aim will be to clarify and adjudicate this particular exchange, which will involve an expansion and clarification of Searle's position and an identification of what I take to be a significant misreading by Thomasson. But I think that this example has broader significance to understanding the relationship between realism and mind-dependence in general, which I hope to draw out over the course of the essay.

1. Searle's account of institutional kinds

Searle's key idea is that institutional reality involves a special capacity of human societies: the capacity to imbue objects, persons, or other entities with the ability to perform certain functions which they would not have been able to perform merely in virtue of their intrinsic (typically physical) properties. Searle calls these special functions *status functions*, because they obtain in virtue of the fact that the community collectively accepts or recognizes¹⁸ the entity in question as having a certain deontic status; in other words, the entity is collectively recognized as playing a role in a system of "obligations, rights, responsibilities, duties, entitlements,

¹⁸ Searle uses 'acceptance' and 'recognition' interchangeably, to emphasize that the phenomenon in question needn't involve approval: "hatred, apathy, and even despair are consistent with the recognition of that which one hates, is apathetic toward, and despairs of changing." (Searle 2011).

authorizations, permissions, requirements and so on.”¹⁹ Institutional facts are, for Searle, facts about the existence of status functions.

A crucial aspect of Searle’s account is that the existence of these sorts of deontic statuses follows the logical form of what Searle calls a Declaration: a kind of speech act in which accepting or recognizing a fact as obtaining itself brings about, or constitutes, the obtaining of that fact. Thus an entity has a deontic status giving it certain powers – say, the positive power of the President to veto bills or the negative power of a citizen to not pay taxes – in virtue of the fact that there is collective acceptance or recognition that it has that deontic status. The basic point is that a system of institutional deontology – of rights, responsibilities, obligations, and so on – can only exist insofar as it is collectively accepted as existing.²⁰

Before moving on to Thomasson’s concerns, it is helpful to bring Searle’s account closer to the ground with a familiar example. Consider the border between New York and New Jersey: there is collective acceptance, codified in law and on official maps of the United States, that a certain imagined line on the ground marks the boundary where one state ends and the other begins. That region of space thereby acquires a certain deontic status – that of a state border – and becomes able to perform a countless variety of functions, broadly construed, that it couldn’t

¹⁹ (Searle 2004)

²⁰ There is a strong reading of Searle’s representation requirement that would see it as applying to all types of deontology, including moral facts: on this interpretation, it can only be morally wrong to F unless there is some representation of the wrongness of F. But it’s possible to reject this stronger claim – to see moral facts as wholly independent of human representations – while maintaining that the deontology of social obligations relevant to the formation of institutions does depend on representations.

perform before the collective acceptance: it determines where I can be arrested for an unpaid speeding ticket, where I can buy liquor on Sundays, and so on. The overall structure is this: the collective acceptance of the deontology adhering to state borders is (at least partially) constitutive of the existence of that deontology, which gives particular entities the ability to perform certain functions they couldn't have otherwise performed – status functions.²¹ Finally, that system of status functions, and others like it, constitute the institutional kind *state border*, one that couldn't exist without collective recognition of its existence.²²

II. Thomasson's antirealist interpretation

Thomasson²³ argues that, assuming Searle is right about institutional reality depending constitutively on the collective acceptance of constitutive rules, it turns out that we'll "require a substantively different ontology, epistemology, and semantics if we are to make sense of the objects studied by the social and human sciences as well as those of the natural sciences."²⁴ The idea is that the ontology, epistemology, and semantics associated with the natural sciences are realist – they're invested in a "picture [that] allows hope of genuine discovery of facts about the

²¹ Note that certain functions, like the function of preventing people from freely crossing from one country into another, might be over-determined – they might be performed both in virtue of physical features and in virtue of the collective acceptance of deontic statuses, as is the case when a political boundary follows a natural one like a mountain range or coast. In such a case we could still say that a status function obtains, in virtue of its counterfactual features – the function would still be performed even if the physical features of the entity were changed, for instance.

²² See below (p31) for a discussion of where Searle's account of institutions falls in our taxonomy of metaphysical mind-dependence; for now, we can say that it is a version of essential (necessary) metaphysical mind-dependence.

²³ (Thomasson 2003)

²⁴ *ibid.* 580

structure of the universe and the nature of kinds in it” – and there are certain criteria for realism that institutional kinds fail to meet.

Thomasson gives the following criteria for realism about any particular kind:

REALISM: “[Real kinds are those that] exist and have their boundaries quite independently of how our concepts and representations might happen to divide things up, in particular, independently of what we believe about the conditions relevant to drawing those boundaries”²⁵

While REALISM itself is phrased as a metaphysical requirement – it spells out what the structure of particular kinds must be in order to count as real – it leads to two further epistemological requirements:

IGNORANCE: “The conditions that determine whether or not something is of kind K [must be] independent of whether or not those conditions are accepted by anyone. *As a result, for all conditions determining the nature of the kind K, it is possible that these remain unknown to everyone.*”

ERROR: “*Any beliefs (or principles accepted) regarding the nature of Ks could turn out to be [...] wrong.*”²⁶

Where does that leave Searle and his intention-dependent institutions? Recall the basic structure of Searle’s account, on which the existence of some kind *K* constitutively requires a certain *K*-regarding belief: for example, the existence of the kind *state border* is constitutively dependent on the acceptance of a set of principles of the form, “*X counts as a state border in context C.*” In the case where this acceptance is successfully made, certain things will, in virtue of that fact, take on the status functions of political borders, and thereby they indeed will *be* state borders.

²⁵ *ibid.* 582

²⁶ *ibid.* 583

Thomasson tries to capture this general structure more formally in what she calls Searle's "Dependence Principle":

DEPENDENCE PRINCIPLE: Necessarily, there is some x that is [institutional kind] K if and only if there is some set of conditions C such that it is collectively accepted that: (if all conditions in C are fulfilled, there is something that is K) and all conditions in C are fulfilled.

Let's consider two possible readings of the Dependence Principle, one stronger and one weaker.

Strong, unconstrained dependence

On the first, stronger reading, the Dependence Principle states that the natures of institutional kinds K are a function of collective acceptance, with no further constraints on what forms that collective acceptance might take. On this reading, the nature of any K is subject to extreme variation: the existence conditions for any particular K can be virtually *any* set of conditions C, so long as they are the conditions that the relevant group in fact accepts as the conditions for K-hood.

Some of what Thomasson writes seems to suggest this stronger, unconstrained reading of the Dependence Principle. For instance: "Suppose we collectively accept that, for any y, fulfilling certain conditions C is sufficient for y to be K. In such a situation, it could not turn out that we are wrong, that being C is not really sufficient for being K. *For the dependence principle ensures that if we accept those conditions, and they are fulfilled, then the entity in question is thereby a K.*"²⁷

Here, Thomasson seems to be making it clear that it is the acceptance itself, whatever form it might take, that links the nature of K to the accepted conditions C. If

²⁷ *ibid.* 588, my emphasis.

there were further constraints on the collective acceptance, distinguishing in some way legitimate or successful collectively accepted conditions from illegitimate or unsuccessful ones, then it would not be the case, as Thomasson writes, that “it could not turn out that we are wrong” about what constitutes K. In cases in which our collective acceptance was of the unsuccessful variety, then we *would* be wrong about the real nature of K, since that unsuccessful collective acceptance would fail to make it the case that K is C.

Setting aside for a moment the question of whether this is indeed the reading that Thomasson intends, let’s consider what happens when it is put into effect. Here are two statements that Searle would be committed to if the stronger reading of the Dependence Principle were true:

Tax Brackets could be Fish: the institution of the American progressive tax schedule is such that, necessarily, if it is collectively accepted that (if X is a green fish in the Nile then X is a tax bracket), then every green fish in the Nile is thereby a tax bracket

Presidents could be Sundays: the institution of the United States Presidency is such that, necessarily, if it is collectively accepted that (if X is the third Sunday in March then X is the President of the United States), then every third Sunday in March is the President of the United States.

Could any remotely plausible account in the business of describing social reality casually commit itself to the (*necessary*, no less) modal claims above?

Some writers would take the above conclusions as immediate reductios of the strong Dependence Principle, on the grounds that they’re straightforward conceptual contradictions.²⁸ But one needn’t have a view of philosophy as conceptual analysis or believe in fixed analytic truths in order to find the above conclusions to be decisive

²⁸ Thomasson’s own commitment to strong analytic truths would seem to put her in just this camp. See (Thomasson 2008), (Thomasson 2014a), and (Thomasson 2014b).

considerations against (this interpretation of) Searle's view. One could deny that existing linguistic usage, or even existing common sense, imposes rigid constraints on what shape a theory could take: a good theory may ultimately arrive at conclusions that are surprising, revolutionary, or counter-intuitive. But it seems like radical breaches of intuition like those above can come only at a substantial price, perhaps by being independently motivated as unavoidable; it would be an uncharitable reading that suggested that Searle breezily accepts such statements as consequences of his account.

Weaker, constrained dependence

Given the unattractive consequences of the stronger reading of the Dependence Principle, on which collective acceptance functions as an unconstrained determiner of the nature of institutional kinds, let's consider whether a modified version would do any better.

On a weaker reading of the Dependence Principle, the nature of institutional kinds is still a function of collective acceptance, but the latter is now subject to independent constraints: only certain kinds of collective acceptance stand in a constitutive relationship to the natures of particular institutional kinds. So if a collective acceptance is made with respect to a kind K, that acceptance has the power to determine the nature of K only if it falls within a range of legitimate collective acceptances by meeting certain further constraints. I discuss below two general approaches to spelling out what these constraints might be: the *Explicitly Constraining* approach, and the *Rigidifying* approach.

On the first approach, one tries to give more or less explicit characterizations of the particular kinds of collective acceptance that are legitimate, either by specifying procedures for yielding legitimate collective acceptances (e.g. participants must be deliberating soundly)²⁹, or, less plausibly, by generating explicit lists of specific legitimate and illegitimate collective acceptances. Call this group of approaches *Explicitly Constraining*, since they attempt to give substantive characterizations of what legitimate collective acceptances are like.

Another approach, the *Rigidifying* approach, to constraining the range of collective acceptances is not to give substantive characterizations of successful instances, but to rigidify the Dependence Principle in a way that limits valid collective acceptances to only those that *actual* relevant participants would make. On this route, we let reference to actual participants do the work of constraining possible collective acceptance.³⁰

Both of these versions of Constrained Dependence have the benefit of helping us avoid at least some of the unattractive consequences of Strong Dependence. For instance, whether we go the explicit or rigidifying route, we now have the means to exclude such extreme violations of common sense as those described in *Tax Brackets could be Fish* and *Presidents could be Sundays*, by arguing that the kinds of collective acceptances figuring in those statements could not issue from the appropriate forms

²⁹ This is analogous to the approach Crispin Wright takes in spelling out the sorts of responses that are relevant to pinning down response-dependent properties (Wright 1998): he describes a substantive list of the sorts of procedures for generating responses that would undergird the relevant response-dependent properties.

³⁰ Compare this to Pettit's "rigid mode" for defining redness-fixing responses in (Pettit 1991).

of deliberation, or would not be the sorts of collective acceptances that relevant participants in the actual world would ever in fact make.

But I am still worried that even the weaker, constrained forms of the Dependence Principle still don't give a satisfactory account of how social reality depends on collective acceptance. In particular, I worry that both forms of the weaker Dependence Principle allow too much theoretical weight to rest on the selection of proper constraints on collective acceptance, leaving our accounts of particular social institutions uncomfortably dependent on the latter.

To see what I mean, suppose we go the Rigidifying route, and accept a version of the Dependence Principle rigidified to read something like this: "Necessarily, there is some x that is money if and only if there is some set of conditions C such that the relevant actual world participants would collectively accept that (if all conditions in C are fulfilled, there is money) and all conditions in C are fulfilled." But do we have an especially good grasp on what the relevant actual world participants would or would not collectively accept? What if it turns out that what actual participants would accept is very different – even shockingly different – from what we thought they would accept: would we be willing to stake our understanding of what the institutional kind *money* is on the outcome of that question?

That seems to get things the wrong way around. A better description of the situation would be that we already have a pretty good theoretical grasp of what money is, one informed by sophisticated social science; it seems that we would use this prior theoretical grasp to determine whether various possible collective acceptances do, or do not, accurately describe money, rather than the other way

around. In certain cases, it seems like we could confidently use this antecedent understanding of money to rule out even those collective acceptances that would be made by relevant *actual* world participants. But if this is right, then the Rigidified reading of the Dependence Principle fails to give a good account of our understanding of the relevant social institutions.³¹

The problem, it seems, with both the strong and weak versions of the Dependence Principle is that they draw too firm a connection between the very natures of institutional kinds and the outcomes of collective acceptances. I've been trying to suggest, by contrast, that the natures of institutional kinds are not quite so variable – indeed, we often have a much firmer theoretical grasp on those natures than we do on facts about what we would or would not collectively accept.

But if the natures of institutional kinds aren't directly constituted by the outcomes of collective acceptances, is there still a significant constitutive role left for those collective acceptances to play? And how should we interpret Searle's emphasis on these acceptances, if not in the direct nature-constituting way suggested by either reading of Thomasson's Dependence Principle?

In what follows, I develop a different structure for the dependence of social institutions on collective acceptances: one that assigns a significant metaphysical role

³¹ A similar argument holds for the Explicitly Constrained version of the principle. It seems like our grip on the nature of *money*, or other particular social institutions, is often stronger than our understanding of what sorts of procedures are legitimate ones for yielding collective acceptance; if the proposed procedures for yielding an effective collective acceptance turned out to commit us to a very unintuitive understanding of what, say, money is, it seems like that would be grounds for a reevaluation of the procedures, rather than of the nature of money. If that's the case, then the Explicitly Constrained reading of the Dependence Principle wouldn't be a good account of our understanding of the relevant institutions.

for collective acceptance in the creation of institutional kinds, but does not leave the very *natures* of those kinds dependent on those acceptances. I think the view I develop is a good reading of Searle's original intentions; either way, I hope it is an independently plausible account of social reality.

III. What are Constitutive Rules?

We've seen above that Thomasson's worry – that institutions as described by Searle violate basic requirements for realism – is based largely in the *constitutive* role that Searle accords to mental phenomena. Here's how Thomasson summarizes her position: "If we understand institutional entities as dependent on the acceptance of certain *constitutive rules* laying out (at least) sufficient conditions for their existence, ... we cannot conceive of investigations into the nature of our own institutional kinds as completely a matter of substantive and fallible discovery."³² To see whether this realism-compromising reading holds up, we should get clearer on what Searle means, exactly, by mental phenomena playing a constitutive role.

Here is one of the few passages in which Searle addresses the issue directly:

We can best understand constitutive rules if we contrast them with regulative rules. The *regulative* rule, "Drive on the right hand side of the road," for example, is a standing Directive. Its function is to bring about a certain form of behavior, and it is satisfied if the behavior matches the content of the rule. The rule has upward, or world-to-word direction of fit. By contrast, the *constitutive* rule, "The oldest surviving son counts as the new king," is a standing Declaration. Its function is to make it the case that a certain person becomes the new king on the death of the old king. [...] It has both directions of fit, word-to-world and world-to-word, simultaneously. It makes something the case by representing it as being the case.³³

³² (Thomasson 2003, 589), my emphasis

The passage above leaves us with the following, rather imagistic, understanding of what's special about constitutive rules: they *do* something. Instead of passively articulating a standard to be met, constitutive rules actively generate a certain state of affairs, or "make something the case."

Now Searle's own description above suggests that the difference between constitutive rules and merely regulative ones is, at least in part, one of grammatical form: constitutive rules are declarative, while regulative rules are imperative. But can the power of certain rules to generate institutional structures really be determined, or even tracked by, the grammatical form in which they are articulated?

Consider Searle's example of a merely regulative rule: "Drive on the right hand side of the road." Is the collective acceptance of that rule really all that different from the collective acceptance of a rule that reads, "The right hand side of the road counts as the correct side to drive on"? If the explanation by way of grammatical form is right, then the second articulation of the acceptance, but not the first, could yield a status-function generating constitutive rule. Or again: rephrase the constitutive rule cited in the passage in imperative form (so that it now reads: "Recognize the oldest surviving son to be the new king"), and it suddenly shrinks down to merely regulative size.

The ease with which the grammatical-form explanation would shift a metaphysically active articulation of an accepted rule into a metaphysically passive one should already give us pause – the grammatical form the accepted rule assumes seems too variable a matter to be given such weight. But what's worse is that on Searle's own account, only very few collectively accepted rules are ever explicitly

formulated; these rare cases include, for example, the written charter establishing a limited liability corporation. By contrast, for most of the institutions that Searle describes – like “a cocktail party, private property, and the adjournment of a meeting” – there doesn’t seem to even be a fact of the matter as to whether the institution-generating collective intention is stated in declarative or imperative form. Instead, the rule is arrived at by tacit mutual understanding in a way that would seem to leave grammatical form underdetermined.

Is there a better account of what differentiates a constitutive rule from a merely regulative one? For some guidance, let’s consider a related distinction emerging from Rawls’s “practice conception” of rules.³⁴ On the practice conception, one thinks of certain rules as *definitive of*, or *constitutive of*, some set of actions that go on to constitute a practice: these rules are “logically prior” to the actions that they govern, in the sense that “there cannot be a particular case of an action falling under a rule of a practice unless there is the practice,” and the rules are what establish the practice as what it is in the first place:

In a practice there are rules setting up offices, specifying certain forms of action appropriate to various offices, establishing penalties for the breach of rules, and so on. We may think of the rules of a practice as defining offices, moves, and offenses. Now what is meant by saying that the practice is logically prior to particular cases is this: given any rule which specifies a form of action (a move), a particular action which would be taken as falling under this rule given there is the practice would not be *described as* that sort of action unless there was the practice. In the case of actions specified by practices it is *logically impossible* [my emphasis] to perform them outside the stage-setting provided by those practices, for unless there is the practice ... whatever movements one makes will fail to count as a form of action which the practice specifies.³⁵

³⁴ (Searle 2010, 92)

³⁴ (John Rawls 1955); Indeed, some commentators rephrase his distinction as that between “constitutive” and “nonconstitutive” rules – see Tamar Schapiro’s “Non-ideal Conditions.”

Rawls's own concern in laying out the practice conception has to do with justification – he wants to illustrate the difference between justification for particular actions within a practice, and justification for the practice itself. But the stage he sets also provides materials for Searle to distinguish between regulative and constitutive rules: while regulative rules apply to elements within already existing institutions, taking the existence of those institutions for granted, constitutive rules serve to establish the overall structure within which regulative rules apply. So the difference between constitutive and regulative rules is not one of grammatical form, but one of logical priority: constitutive rules, in generating the institutions that fall under the scope of regulative rules, come first. So here is a first-pass account of constitutive rules: they are just those rules that serve to create institutional structures.

But what, exactly, is the sense of creation at work? I think that here Searle allows an important, and importantly misleading, ambiguity to persist unaddressed. To get a sense of what it is, consider the following two quotes; each is representative, respectively, of two different ways in which Searle tends to describe the “creation” of institutional reality:

i) “A football game, a stock market transaction, a cocktail party, private property, and the adjournment of a meeting are all [...] *brought into existence* by constitutive rules.”³⁶

ii) “There is no limit to the institutional realities we can create *just by agreeing*, in language, that we are creating them. We create universities, cocktail parties, and summer vacations, for example. [...] Institutional facts exist *only insofar as they are represented* as existing.”³⁷

³⁶ Ibid. 25

³⁶ (Searle 2010, 10); my emphasis

In both cases, there is a dependence relation posited between institutions and constitutive rules, but the relations in question are very different. The first quote makes no mention of any *acceptance* or *recognition* of constitutive rules: it just says that different institutional facts are defined, or captured, by different rules. The second quote talks not just about rules, but about our use of them: here, the claim is clearly that it is a specific mental phenomenon – the collective acceptance of certain constitutive rules, rather than the rules in the abstract – that “creates” the relevant institution. Both senses of constitutivity are crucial to Searle’s conception, and lack of explicit disambiguation makes the two easy to conflate. The confusion is further exacerbated by the fact that much “constitution”-related vocabulary – “creation,” “generation,” “establishment,” and so on – is equally applicable to both conceptions, making it difficult to discern what conception is at play at any particular occasion. In what follows, I’ll try to spell out both conceptions in a little more detail, and emphasize their differences.

The first conception seems to be the one that Rawls has in mind in the passage above, when he characterizes “the rules of a practice as defining offices, moves, and offenses.” Let’s call this the *essential* conception of constitutive rules: the idea is that certain rules establish the nature or essence of particular institutional kinds.³⁸ For

³⁸ *ibid.* 88, my emphasis

³⁸This conception might also have been called the *definitional* conception, taking a cue from the quoted Rawls. But that would run the risk of implying that constitutive rules describe their institutional kinds in a rigid or analytic way that precludes evolution or change. That’s not the conception I have in mind. The idea is just that, to the extent that *any* kind has a nature or essence that is captured by stating certain existence conditions, institutional kinds have a nature or essence that is, at least in part, captured by specifying the rules pertaining to them.

example, the rules defining certain rights, powers, and obligations define, in part, what it is to be the president of the United States. On the essential conception, constitutive rules are prior to regulative rules in that they fix or determine the nature of the structures that regulative rules invoke.

This is just what Searle seems to have in mind when he describes institutional kinds as consisting of particular status functions, which are in turn captured by rules specifying a particular deontology. What it is to *be* a particular institution is to be, in part, an entity with a set of status functions determined by certain deontology-establishing rules.³⁹ The relation that obtains between the rules, the deontology, and the resulting institution is therefore something like a logical one, obtaining even in the absence of any actual, concrete instantiations of the relevant institutional kind in the world. So the corresponding sense in which, say, a cocktail party is “brought into existence” or “created” by a set of rules is just that the rules capture, at least in part, what it is to *be* a cocktail party: the essential relation between those rules and the institutional kind *cocktail party* obtains regardless of whether anyone actually accepts the relevant rules, and so regardless of whether there actually *are* any cocktail parties.

Let’s turn now to the conception of constitutivity at play in the second set of quotes, the one that captures Searle’s recurring insistence on the idea that it is the *act* of representing an institutional kind that “brings it about.” On this conception of constitutivity, the existence of institutional reality requires a set of real, concrete

³⁹ I say “in part” because the nature or essence of an institution will likely not be wholly captured by a set of rules: it might also include information about what entities the institution applies to, its causal etiology, and so on.

events: the mental phenomena corresponding to collective acceptance of constitutive rules. Call this second conception – on which it is a concrete event, namely our acceptance of the rules, that brings about the existence of the corresponding institutional kind – the *enactive* conception. On the enactive conception, the sense in which constitutive rules “bring about” or “create” institutions is not the same as the sense in which they “bring about” or “create” institutions on the essential conception: the latter has to do with establishing an essence or nature, while the former has to do with concrete, spatiotemporally-defined enactment, or realization, of an entity with that nature in the world.

To illustrate the difference a little more concretely, consider a game of chess. The rules assigning allowable moves to each type of piece are, in part, *essentially constitutive* of the game of chess: a game described by substantially different rules – ones according to which minor pieces move the same way as the queen, for instance – would be not chess, but something else. But the rules assigning allowable moves to each type of piece are also *enactively constitutive* of a particular game of chess, in that two players must intend to be following those rules, at least roughly, in order to count as playing a game of chess at all – in order, that is, to be manifesting a particular chess game in a particular time and place.

Now that we have the two conceptions of constitutivity on the table – essential and enactive – we must complicate things slightly by considering ways in which the two conceptions interact. In particular, on Searle’s account it is *essential* to certain institutions that they be *enacted* by certain mental phenomena: part of what it *is* to be an institutional kind is not just to be governed by certain rules, but to be

such that collective acceptance of those rules metaphysically manifests a token of that institutional kind. Thinking back to our taxonomy of metaphysical mind-dependence, this places Searle's account in the category of the *essentially mind-dependent*.

Within that category, should it be further classified as *causally* or *non-causally* mind-dependent? Here, matters are a little bit less clear, and it seems like the best interpretation is: a little bit of both. Searle often describes the imposition of status functions using "in virtue of" language: "They are functions that a person or other entity has, not in virtue of physical structure, or at any rate not solely in virtue of physical structure, but in virtue of collective imposition and recognition of a status." But should "in virtue of" be understood causally, or non-causally? Sometimes the suggestion is strong that it is collective acceptance that *itself* manifests the relevant deontology, in an immediate, non-causal sense: "deontic powers are powers that exist only because they are recognized and accepted as existing." But many of the status functions that Searle describes seem to depend in a very straightforwardly *causal* sense on the collective intentions that underwrite them: for instance, it seems right to say that our intentions to treat this piece of paper in a certain way are, in part, what *causes* it to fulfill the status function of money, given that they cause it to elicit certain responses from people, cause certain behavior to be more likely, and so on. When Searle writes that collective acceptance is what is required "for the status functions to *actually work*,"⁴⁰ it is hard to make sense of an entirely non-causal reading.

With the distinction between essential and enactive constitutivity on the table, we can finally summarize Searle's account in the following way: it is one on which it is *essential* for social institutions that they be captured by particular deontology-establishing rules; it is also *essential* to them to be *enacted*, whether causally or non-causally, by certain mental phenomenon, namely the collective acceptance of particular rules: in certain cases, the sets of enactively constitutive and essentially constitutive rules coincide. This occurs in our chess example, in which the actual enactment of a particular chess game requires acceptance of the very rules that are essentially constitutive of playing chess.

But it is crucial to keep careful track of what falls within the scope of collective acceptance: *what is essential* to an institution is not *itself* within the scope of collective acceptance. Collective acceptance determines whether an institution is *enacted* or not, but has no power to determine what is *essential* to it: enactment falls within the scope of essentially, but not the other way around.

In what follows, I'll argue that this latter conflation is precisely what's behind Thomasson's antirealist criticism of Searle.

IV. Reinterpretation of Thomasson's critique

With this expanded understanding of Searle's position on the table, let's take another look at Thomasson's Dependence Principle:⁴¹

⁴¹ *ibid.* 8

⁴¹ For the purposes of the following discussion, the difference between the stronger and weaker reading of the Dependence Principle won't be relevant; for simplicity, I'll talk about the Dependence Principle tout court, assuming that either reading applies.

DEPENDENCE PRINCIPLE: Necessarily, there is some x that is [institutional kind] K if and only if there is some set of conditions C such that it is collectively accepted that: (if all conditions in C are fulfilled, there is something that is K) and all conditions in C are fulfilled.

Based on the distinction between the essential and enactive constitution of institutional kinds – i.e. the difference between rules that, in the Rawlsian sense, fix or determine the nature of institutional kinds and rules whose acceptance brings about, either causally or constitutively, the existence of particular, concrete members of that kind – it's clear that the Dependence Principle is invoking the enactive conception: the x that figures in the Dependence Principle is a concrete manifestation of an institutional kind, and it is the collective acceptance of the rule in question, and not the rule in the abstract, that is supposed to be bringing that institutional kind about.

But would the enactive conception, any more than the essential one, allow for the existence conditions of K to be a function of collective acceptance in the way the Dependence Principle describes? It wouldn't. As we've seen above, what the enactive conception is committed to is that whether or not an institutional kind is actually realized in the world is, in part, dependent on mental phenomena: insofar as the acceptance of constitutive rules is, in some sense, up to the group in question, the enactment of the institutional kind will also be up to the group in question. But what is *not* up to the group is *what kind* of institution, with *what kind* of existence conditions, will be enacted by the acceptance of particular rules: that's a matter that's fixed by the nature of the institution, by looking at what is essentially constitutive of it.

When Thomasson claims that “in the case of institutional kinds, those principles we accept regarding sufficient conditions for the existence of these entities must be true because ... the principles accepted play a stipulative role in constituting the nature of the kind,” she is, I’ve been suggesting, making a mistake about what falls within the scope of collective acceptance in Searle’s account. What Searle is committed to is that what rules we adopt, and therefore what practices we enact, depends on what we *do*: institutional kinds are established and maintained by our own, in part mental, activity. By adopting different rules – by setting different collective intentions, by making different agreements, by behaving in different ways – we would have given rise to potentially very different institutions. What he is *not* committed to is the question of “which institutions are enacted by which rules?” itself being a matter dependent on collective acceptance.

Notice that in what’s been said above, nothing special hinges on the fact that the activity in question is the mental one of accepting constitutive rules. Consider the following analogy:

Livestock Principle 1: The variety of livestock depends on human activity: if agricultural practices had developed differently, livestock might have been very different from what it actually is. There might not, for instance, have been domesticated cows, but instead domesticated reptiles.

Livestock Principle 2: The variety of livestock depends on human activity: if agricultural practices had developed differently, livestock might have been very different from what it actually is. Domesticated cows, for instance, might instead have been domesticated reptiles.

I am urging that the mistake at play in Livestock Principle 2 is the same one, at base, as the one behind the Dependence Principle, and the mistake expressed in statements like the following: “Suppose we collectively accept that, for any *y*,

fulfilling certain conditions C is sufficient for y to be K. In such a situation, it could not turn out that we are wrong, that being C is not really sufficient for being K. For the Dependence Principle ensures that if we accept those conditions, and they are fulfilled, then the entity in question is thereby a K.”⁴² I hope it’s clearer now what’s going wrong⁴³: the dependence relation between institutional kinds and collective intention could indeed ensure that, if we accept the conditions *relevant to K-hood*, then the entity in question would thereby be a K. But that dependence relation is very different from one that sees the very *nature* of institutional kinds as functions of collective acceptance.

Before moving on, let me try to say a bit to help situate the position that is emerging with respect to a wider landscape of possible views.

One way to characterize my argument is as accepting the spirit behind Thomasson’s Realism Principle – accepting, that is, that real kinds are those that have established natures or essences that are independent of what those natures or essences are taken to be – while insisting that institutional kinds are nevertheless real, by that Principle’s own lights. So, there has been implicit in my argument the same metaphysical view of natures or essences that Thomasson associates with the “ontological realist paradigm:” namely, a view of essences on which they are themselves objectively existing entities independent of human conceptual practices. One might therefore worry about the compatibility of my view of institutional kinds –

⁴³ (Thomasson 2003, 588)

⁴³ At the risk of belaboring the point, here is the corresponding livestock version: “... In such a situation, it could not turn out that having reptile-like qualities is not really sufficient for being a cow. For the Livestock Principle ensures that if an animal issues from our agricultural practices, then the animal in question is thereby a cow.”

one on which they depend essentially on human conceptual practices without having essences that are *functions* of human practices – with views that are less realist about essences altogether. Consider, for instance, a view on which the criteria for objecthood *in general* make essential reference to conceptual practices (as in Einheuser 2006), or one on which the essences or natures of all kinds are ‘conferred’ on those kinds by human activity (as in Sveindottir 2008).

It is true that my argument, as written, invokes a more realist conception of essences than these views would permit. But its central point would nevertheless survive transposition into the vocabulary of these more ‘essence-antirealist’ views. That’s because the main target of my argument is any position that stakes realism on independence from human conceptual practices, drawing corresponding distinctions between kinds that are thereby ‘more’ and ‘less’ real, or *real* and *unreal*. Against that position, I’ve been arguing that the essential dependence of institutional kinds on conceptual practices does *not* make them less real than kinds that lack such a dependence. But if, on the other hand, one adopts a view on which *all* objects, from numbers to chemical elements to distant stars, are generally dependent on human practices for their essences, then one is not making a specifically antirealist argument against institutional kinds of the sort that it has been my aim to dislodge; a view that assimilates, as far as their reality is concerned, social institutions to chemical kinds is one that I welcome as allied with my own.

V. Epistemic Concerns

Up until now, my argument has focused almost exclusively on the metaphysical relation between mental phenomena and the generation of institutional kinds. The central concern has been the general structure of the relation of mental phenomena to the nature and existence of institutional kinds, which doesn't seem to directly depend on the content of the mental phenomena in question, and in particular isn't affected by the specifically *K-regarding* nature of the relevant mental phenomena.

But those cases in which the collective intention *is* K-regarding do pose a special set of epistemological issues, if not metaphysical ones; so while our discussion of the (metaphysical) Dependence Principle proceeded in abstraction from this question, the discussion of the (epistemological) Ignorance and Error principles must take it into consideration. Let's look at these principles one more time:

IGNORANCE: "The conditions that determine whether or not something is of kind K [must be] independent of whether or not those conditions are accepted by anyone. *As a result, for all conditions determining the nature of the kind K, it is possible that these remain unknown to everyone.*"

ERROR: "*Any beliefs (or principles accepted) regarding the nature of Ks could turn out to be ... wrong.*"

Thomasson takes Searle's account to be violating both Ignorance and Error. Let's briefly see whether that's the case.

First, as I've been insisting above, care must be taken in distinguishing two senses in which the content of mental phenomena may be said to "determine" or "constitute" an institution: the essential sense and the enactive one. As they stand,

the Ignorance and Error principles are somewhat ambiguous as to whether the conditions “determining the nature of the kind K” should be understood enactively or essentially. Is the possibility of ignorance and error supposed to characterize our epistemic standing towards facts about particular, concrete, enacted tokens of institutional kinds? Or is it supposed to characterize our epistemic standing towards the *natures* or *essences* of those kinds, considered in the abstract from any particular tokens of those kinds? Here’s a version of the enactive sense of the question: can we all be wrong, of this particular piece of paper, that it’s a bit of money? Here’s a version of the essential sense: can we all be wrong, in general, about what the nature of money is?

If the principles are meant to be read in the *essential* sense, it is not obvious that they are violated at all by Searle’s account of institutions. Consider the first component of the Ignorance requirement: “The conditions that determine whether or not something is of kind K [must be] independent of whether or not those conditions are accepted by anyone.” As I’ve been emphasizing throughout, the essential properties, or natures, of Ks *are* independent of whether or not those conditions are accepted by anyone – what’s not independent of collective acceptance is whether tokens of those kinds are enacted. And while it may be *essential* to an institutional kind that it be enacted by collective acceptance, that fact itself does not fall under the scope of what is determined by collective acceptance. The question of what is essential to an institutional kind K – and, for that matter, what is essential to any other kind – is, I’ve been arguing, a real fact about the world, one that’s not

subject to decision or stipulation.⁴⁴ Furthermore, the content of a collectively accepted rule like, “bills printed by the Fed will count as money in the United States” needn’t say anything explicit about the *nature* of the institutional kind money: knowledge of the content of that rule might, together with some contextual data, directly lead to knowledge of whether this piece of paper counts as money, but it needn’t directly lead to knowledge of money’s nature. So familiarity with, and ability to engage in collective acceptances regarding, aspects of money’s role do not lead in any direct and infallible way to knowledge of money’s nature or essence – that is a quite separate question, one more likely addressed by social scientists and philosophers than by users of money.

But things are a little different if the principles are to be read in the enactive sense, explicitly rephrased as such below:

IGNORANCE (enactive): For all facts about particular tokens of kind K, it is possible that these remain unknown to everyone.

ERROR (enactive): Any beliefs (or principles accepted) regarding facts about particular tokens of kind K could turn out to be ... wrong.

But now, it indeed seems to be the case that in certain felicitous contexts of collective acceptance of rules, there can be no possibility of ignorance or error regarding the particular institutional facts that those rules describe, since the collective acceptance has itself made it the case that those facts obtain. For example, if it is collectively accepted that *X counts as Y in context C* for an institutional kind Y for which Y-regarding collective intentions are enactively constitutive, then the

⁴⁴ Though, of course, decision and stipulation may themselves figure among what’s (essentially) enactive of an institution. It is this difference that I’ve been trying to bring out throughout this essay.

collective acceptance will make it the case that X counts as Y, seeming to guarantee the truth of its own content.

But notice that a version of the description above will apply in general to almost *any* felicitous act of stipulation. For instance, in cases where certain background conditions are met, the statements “This child is named Mary” and “You are now husband and wife” will themselves bring about the truth of the state of affairs they describe, seeming to ensure their utterers against ignorance or error regarding the child’s name or the couple’s marital status. Take the naming example: it seems to violate Ignorance, since there *is* a particular fact – the name of the child – that is *not* possibly unknown to everyone: in particular, it is not possibly unknown to the person stipulating what that name will be, in the proper context of stipulation. For the same reason, it also violates Error: the namer, at the moment of the successful naming, has no possibility of being wrong about what the name will be, at least at the moment the naming occurs. But stipulations are ubiquitous and not especially mysterious; if they *all* constitute violations of Ignorance and Error, we should seriously question Thomasson’s claim that these principles are requirements of “standard realist epistemology.” So let’s take a closer look at what might be going wrong.

First, notice that Ignorance and Error, in their enactive versions, are violated (if at all) only by a very specific entity in a very specific context: namely, the particular person or group issuing the stipulation regarding fact F, and only in those contexts in which the F-enacting stipulative powers actually apply. Any F-regarding

statement made outside the context of felicitous stipulation, even by the original stipulator, might well be wrong.

But what exactly is going on *within* that very narrow context – the window in which the utterer seems ensured against ignorance or error by being in a felicitous context of stipulation? It seems to me that confusion looms once we start regarding the particular utterance made within the context of stipulation as evaluable *at all* with respect to ignorance or error regarding the facts that that utterance enacts.

Here is a useful thing Wittgenstein says on the matter: “‘Observing’ does not produce what is observed. (That is a conceptual statement.) Again: I do not ‘observe’ what only comes into being through observation. The object of observation is something *else*.”⁴⁵ I think it’s tempting to say something similar of ignorance and error as what Wittgenstein claims is true of observation. Just as observation (conceptually) requires that the observed object be held fixed with respect to the process of observation, so too does the (conceptual) possibility of ignorance or error regarding some set of facts require that set of facts to be held fixed, at least in principle, with respect to the particular belief or utterance being evaluated for ignorance or error. In those very special cases when the belief or the utterance itself has the function of *determining* the relevant set of facts, evaluation cannot proceed in the ordinary way, and the possibility of ignorance and error simply do not arise.^{46, 47}

⁴⁵Philosophical Investigations, §67.

⁴⁶ This is a considerably stronger position than the one described, for instance, in (Holton 1991). There, Holton is also concerned to deny the idea that response-dependent properties automatically and quite generally entail immunity from error. But while Holton concedes that certain specific circumstances nevertheless *do* involve a kind of immunity from error, the position I am suggesting here is that in *no* circumstances should the kind of privileged

As I see it, there are two fairly strong theoretical pulls, each in a different direction. On one hand, if the possibility of ignorance or error conceptually requires that the relevant state of affairs be held fixed with respect to the statement being evaluated, then successful stipulations like “this child is named Mary,” and “you are now husband and wife” aren’t even candidates for evaluation – that’s the possibility that I’ve been suggesting just above. On the other hand, those statements certainly *seem* to be evaluable for error: after all, they have the form of ordinary, descriptive, truth-evaluable sentences. Why not treat them as such?

I think the solution will come down to a matter of theoretical preference. If we’re more moved by the first set of considerations, then we can keep the Ignorance and Error principles exactly as they are, and continue to maintain that they capture an intuitive notion of what it takes for a state of affairs to be real. In this case, stipulations won’t violate either principle, since they won’t count as the sorts of things that can be candidates for ignorance or error in the first place: though they have the appearance of descriptive, truth-evaluative statements, that appearance is

position involved in mind-dependent phenomena be construed as conferring immunity from error, as a conceptual matter.

⁴⁷ This is the epistemological version of a point that can also be made in a more metaphysical mode. Consider Michael Dummett’s (1978) characterization of realism, on which realism means that “statements of the disputed class possess an objective truth-value, independently of our means of knowing it: they are true or false in virtue of a reality existing independently of us.” But there *is* a sense in which even essentially mind-dependent phenomena, including facts about states of affairs that essentially, enactively depend on mental activity for their existence, are nevertheless “true or false in virtue of a reality existing independently of us.” That’s because before the relevant mental activity occurs, there is simply no fact *there* yet to allow the corresponding statement to be evaluable for truth or falsity – the mental activity is in part what constitutes the very obtaining of the relevant state of affairs, and without the mental activity, there is no relevant state of affairs to be described either truly or falsely. But once the mental activity has occurred, playing its role in constituting the relevant state of affairs, then the truth value of any statement describing that state of affairs is no longer up to us in any way. In this sense, the *existence* of certain states of affairs is mind-dependent, but the truth or falsity of any descriptions of that state of affairs is not.

deceptive. What they really are is a kind of doing, not a kind of describing. On the other hand, if we're inclined to respect the outward descriptive grammar of those statements, then we'll need to make small adjustments to the Ignorance and Error principles to differentiate between ordinary descriptive statements and stipulative ones. This adjustment will not reflect some special irreality of institutional kinds in particular: instead, it will make explicit the ordinary understanding that stipulations shouldn't be construed as violations of realism about the subject in question. In either case, there is nothing specific to institutional or human kinds, as against any other instance of stipulation, that would lead us to have these worries.

* * *

In this paper, I hope to have elaborated a view of social and institutional kinds that achieves two aims. On one hand, it does justice to the special relationship that institutional and social reality, as contrasted with other types of entity, bear to mental phenomena: the former stands in a relation of essential enactive dependence on the latter. On the other hand, institutional kinds are nevertheless fully real, and their natures or essences are not functions of what we take those entities to be: there is an important sense in which the truth or falsity of statements describing institutional reality *is* wholly independent of us, making these kinds acceptable objects, both ontologically and epistemologically, of inquiry.

* * *

Part III: Objectivity and Mind-Dependence

In "Objectivity and Modern Idealism: What is the Question," Gideon Rosen argues for the reality and objectivity of mind-dependent entities. While various writers have argued that certain mind-dependent domains fall short of full reality or objectivity, Rosen concludes that they are, and indeed can be, nothing less than "a domain of fact that is *out there* ... [obtaining] *anyway*, regardless of what we may think."⁴⁸

In this paper, I will use Rosen's argument as a starting point from which to evaluate claims of the reality and objectivity of mind-dependent facts; I will argue that these features come apart, and that certain domains that we acknowledge as real may nevertheless be seen as less than fully objective. I will conclude by identifying some elements essential to the objectivity of a subject matter, as a feature over and above the reality of its associated entities.

I. *Rosen's argument for the reality and objectivity of judgment-dependent domains*

The bulk of Rosen's argument is devoted specifically to *judgment-dependent* entities. Some writers, including Amie Thomasson, have taken certain kinds of judgment-dependence to pose a unique threat to realism, one that distinguishes those entities from other, more benign forms of the mind-dependent.⁴⁹ But one of Rosen's central strategies is to diminish the apparent disanalogy between the special case of judgment-dependence and other kinds of mind-dependence. He concludes

⁴⁸ (Rosen 1994)

⁴⁹ (Thomasson 2003)

that judgment-dependent domains, like any other kind of mind-dependent domain, can be nothing short of objective.

Let's begin with a brief reconstruction of Rosen's argument for the reality of judgment-dependent entities and the objectivity of their associated domains.

Rosen starts with a rough construal of the mind-dependent in general: basically, any domain of subject matter enjoying a "dependence on our linguistic and social 'practices,'" or standing in some dependence relation to "some response of subjects which essentially and intrinsically involves some mental process."⁵⁰

Crucially, Rosen emphasizes that the relevant mental phenomena involved in the dependence relation are themselves nothing less than fully real: "a flexible and relatively undemanding naturalism functions for us as an unofficial axiom of philosophical common sense ... If we believe in minds at all, they are the embodied minds of human beings and other animals." And it seems to follow from this 'undemanding naturalism' about the mental that further facts about the relations in which the mental stands to other natural entities – for example, facts about dispositions to produce mental responses – are *also* fully real:

The point is the obvious one: dispositions to bring about mental responses would seem to be on a par, metaphysically speaking, with dispositions to produce merely physical responses in inanimate things. ... Absent a reason to construe mentality itself as less than fully real, the facts about the annoying, the embarrassing and the rest are no different from facts about the poisonous or the corrosive.⁵¹

With this slightly more expansive naturalism about the mental on the table, we turn specifically to the question of judgments: what is judgment, Rosen asks, if

⁵⁰ *ibid.* 290

⁵¹ *ibid.* 293

not just another kind of mental phenomenon, on a metaphysical par with “facts about the annoying, the embarrassing, and the rest”? Though issuing judgments may seem, from the perspective of the judger, to be a significantly different kind of mental activity from feeling itches or sensing smells, Rosen claims there is “no reason to think that the facts about what a certain group of people would think after a certain sort of investigation are anything but robustly objective, ... possess[ing] the same status as the facts about what any other collection of animals would do if prompted with certain stimuli, or set a certain problem.”⁵²

We now arrive at the crux of Rosen’s argument. Let’s call it his *Intuitive Metaphysical Principle*:

Intuitively, if the facts in the contested class can simply be read off in a mechanical way from the facts in an uncontroversially objective class, then there can be no grounds for denying the same status [w.r.t reality and objectivity] to facts in the contested area.⁵³

In other words, if we can establish that some set of facts *I* bears a certain relation *R* – Rosen describes it as “mechanical reading-off from,” but we’ll reconsider what that might mean a little further down – to some set of facts *M* that is uncontroversially real and objective, then *I* is thereby seen to enjoy the same real and objective standing that *M* has.

But if it’s true that judgment-dependent phenomena should be treated, for the purposes of ontological and epistemic evaluation, as indistinguishable from other kinds of mind-dependent phenomena, then given our Undemanding Naturalism and the Intuitive Principle, we get the result that any entity that is judgment-dependent –

⁵² *ibid.* 300

⁵³ *ibid.* 301, Rosen’s emphasis

that is, any entity bearing the special relation R to some set of judgments – is objective and real; Q.E.D.

Here's the argument once again, condensed:

P1. *Undemanding naturalism about the mental*: facts about mental phenomena, including facts about dispositions to produce them, are fully objective and real.

P2. *Assimilation of judgments*: Insofar as objectivity and reality go, higher cognitive phenomena like judgments are just like any other kind of mental phenomena.

P3. *Intuitive Metaphysical Principle*: facts that bear a certain relation R (roughly, "mechanical reading-off from," but to be further discussed below) to a set of uncontestedly real and objective facts are themselves real and objective

Conclusion (from 1, 2, & 3): Facts that bear relation R to judgments – that is, facts about judgment-dependent phenomena – are perfectly real and objective.

Before moving on, the time has come to finally address the mysterious R relation, for which we've so far left just a placeholder. Now, in the original form of his argument, Rosen takes the relation that facts about judgment-dependent entities bear to facts about judgments to be *analytic or conceptual entailment*: for some property F to be judgment-dependent is captured in the truth of the following a priori biconditional, which follows from facts about the concept F :

It is a priori that: x is F iff we would judge that x is F in certain appropriate circumstances.

Recalling Wedgwood's distinction between conceptual and metaphysical mind-dependence,⁵⁴ we see that Rosen goes the former route, treating judgment-dependence as, in the first instance, a matter of conceptual relations between facts about judgments, on one hand, and facts about the relevant judgment-dependent

⁵⁴ (Wedgwood 1997); discussed p3 above.

domain, on the other. In other words, he takes Relation *R* in the argument above to be conceptual or analytic entailment.

But while Rosen's original argument proceeds in the conceptual mode, I see no good reason against transposing its basic structure into a metaphysical register: no good reason, that is, against evaluating his argument in contexts where judgment-dependent entities stand in a relation of *metaphysical* dependence to judgments.

For one thing, Rosen's discussion took for its central examples, naturally enough, the several prominent response-dependence accounts being developed at the time, particularly those of Jackson, Pettit, and Wright; these were all stated in terms of conceptual relations. But Rosen has himself suggested⁵⁵ that if he were revisiting the matter of mind-dependence today, he would likely help himself instead to the (metaphysical) language of essence and grounding in order to capture what it means for certain domains to depend on mental phenomena.

But even setting these contextual considerations aside, I think that the text itself provides good reason for thinking that Rosen's arguments stand to be evaluated with respect to metaphysical, and not just conceptual, mind-dependence. First, his crucial Intuitive Metaphysical Principle is already stated in terms general enough to apply to a metaphysical reading: the "mechanical reading-off" (and, in other places, "supervenience") that he mentions as ensuring a transference of metaphysical status seems just as applicable to relations of essence or grounding as to conceptual entailment. There seems, in other words, to be nothing special about

⁵⁵ In email correspondence

conceptual or analytic entailment that justifies or underwrites the use of the Intuitive Metaphysical Principle on which Rosen's argument depends.⁵⁶

Second, there is good dialectical reason for considering Rosen's arguments with respect to the metaphysical case. That's because Rosen himself mentions several versions of metaphysical mind-dependence – what he refers to as “existential” and “causal” mind-dependence – and sets them aside as obviously posing no risk at all to either the objectivity or the reality of the entities they describe: “Objectivity, if the notion makes sense at all, is thus opposed to some sort of *non-causal, non-existential* (for short: non-empirical) mind-dependence.”⁵⁷ In other words, Rosen doesn't seem to take the metaphysical case as standing in need of defense in the first place: metaphysically mind-dependent entities are suggested to be obviously, unproblematically objective and real. But the objectivity and reality of such entities *has* been seriously questioned: if Rosen's defense is supposed to work in the (purportedly trickier) case of conceptual mind-dependence, we should, at the very least, see if it works in the (purportedly unproblematic) case of metaphysical mind-dependence.

To that end, I will, in what follows, apply the basic structure of Rosen's argument to the case of metaphysical mind-dependence, interpreting the *R* relation to be essential metaphysical dependence.⁵⁸

⁵⁶ Indeed, it might seem *more* plausible that a transference of metaphysical status from one set of facts to another should occur through a metaphysical, and not merely conceptual, relation between those sets of facts.

⁵⁷ *ibid.*, my emphasis.

⁵⁸ Including its causal and non-causal variants; see Part I: Varieties of Metaphysical Mind-Dependence, above.

II. *Different kinds of judgments*

We have seen that a large part of Rosen's strategy for arguing for the reality and objectivity of judgment-dependent entities is based on emphasizing the similarities between judgments and other mental phenomena, such as responding to irritants or having certain visual sensations. The analogy is facilitated by considering how things look from a certain anthropological remove: from a distance, it's hard to see what the difference might be between people making judgments and "what any other collection of animals would do if prompted with certain stimuli."

In this section, though, I want to explore some special features of judgments that make them importantly dissimilar from other kinds of mental phenomena, and that turn out, I'll suggest, to bear on questions of objectivity, in particular. This will lead to my broader aim of inserting a wedge between objectivity and reality.

Investigative vs. Legislative Judgments

To begin with, it's important to distinguish between several quite distinct mental attitudes that might be referred to equally naturally as "judgments." Consider the following examples:

- a) John established that the bookcase wouldn't fit through the door, so he had it shipped in pieces.
- b) John established that the bookcase would stand besides the fireplace, so he moved aside the woodpile.

In the first case, John's "establishing" is best understood as a report of an empirical discovery: he did some measuring and *found out* how things would go. But

in the second case, John's "establishing" is best understood as the expression of a decision or stipulation: he did some thinking and *chose* how things would go.

A lot of the language that we use to describe mental attitudes exhibits a similar ambiguity between, roughly, investigation and determination: think of words like *conclude*, *figure*, *suppose*, and, importantly, *judge*. In most cases, contextual cues are sufficient for resolving the intended sense, but in discussions of objectivity, realism, and judgment-dependence, it's important to be explicit about what attitudes we're trying to ascribe.

That's because a subject matter that's plausibly described as judgment-dependent in the investigative sense would often be extremely *implausibly* described as judgment-dependent in the determinative sense, and vice versa. For example, even those metaethical positions that treat normative properties as grounded in, reducible to, or expressive of our moral judgments typically acknowledge the latter to at least *seem*, to their subjects, to be responsive to an external moral reality:⁵⁹ a metaethical theory taking normative properties to be determined by moral sentiments must not be confused with one taking them to be determined by choice or stipulation. Something similar is true for accounts of aesthetic properties: while it is a familiar family of views that holds aesthetic value to derive, in whole or in part, from a set of responses on the part of certain observers, it would be a far more radical position that included among those responses things like free decisions or stipulations to regard something as beautiful, say, or worthwhile.⁶⁰

⁵⁹ See Geoffrey Sayre-McCord, "The Many Moral Realisms," in (Sayre-McCord 1988)

⁶⁰ See Cavell on Kant on beauty: "[Judgments about the beautiful] must – logically must, some of us would say –... "demand" or "impute" or "claim" general validity, universal agreement

Now, as I've been using these terms so far, the question of whether a judgment is investigative or determinative is answered with respect to the self-conception or perspective of the judging subject. Roughly, to be an investigative judgment, in this sense, is to be a judgment issued while taking oneself to be responsive to or constrained by *how things are* with respect to the matter one is judging about; to be a determinative judgment is to be a judgment issued while taking oneself to be choosing or deciding how *things shall be* with respect to the matter one is judging about.⁶¹

But as we've seen in our brief discussion above of aesthetics and metaethics, whether a judgment is investigative or determinative in *that* sense can come apart from the actual effects that that judgment has on the subject matter in question: an investigative judgment to the effect that 'x is F' can play the role of determining or making it the case that 'x is F.' One way of putting this is to say that certain investigative judgments can turn out to have determinative effects. So, for clarity, we need to start distinguishing between the ***first-personal mode*** of a judgment – what the judging subject takes his relation to the subject matter to be – and the ***metaphysical mode*** – what the relation of the judgment to the subject matter actually is. Now we can say things like: first-personally investigative judgments that *x is beautiful* are metaphysically determinative judgments that *x is beautiful*.

with them; and when we make such judgments we go on claiming this agreement even when we know from experience that they will not receive it. ... Kant says about this claim to universal validity, this voice, that it "so essentially belongs to a judgment by which we describe anything as *beautiful* that, if this were not thought in it, it would never come into our thoughts to use the expression at all." (Cavell 2002)

⁶¹ This ability of judgments to have specific first-personally investigative vs. determinative registers is one way in which they're different from "brute sensations" or "raw feels."

In summary, we see that to characterize a domain as judgment-dependent is often to give only a small part of the story: a full account would need to specify not just the metaphysical relation of the judgments to the domain in question, but also whether those judgments are first-personally investigative or first-personally determinative.

IV. Tension between first- and third-personal perspectives

Rosen himself notes that there is a potential tension in cases where one's judgment is relevant to determining the subject matter that one is judging about. He describes the tension in terms of two possible perspectives onto the relevant subject matter, *F*: on one hand, the 'engaged' perspective of the subject rendering the judgment, and on the other, the 'anthropologist' perspective from which the real nature of *F*-hood is clear. Sometimes, Rosen claims, these two perspectives can come into conflict when deliberating on the question of *F*.

The passage as written is somewhat puzzling, so it will be helpful to quote it in full, and then apply to it the distinctions we made above:

[The engaged judger] must still think of himself as constrained to consider the balance of reasons, the force of precedent, and so on, on both sides of the issue. If he steps back from himself momentarily, he may see this process as simply a matter of providing input to a system whose output constitutes the relevant fact. He can theorize about what this output will be, and from this perspective, the facts about [*F*] will strike him as fully objective. But what is he to think when he is "engaged"? In a certain sense he must concede that the facts do not constrain his decision at all, but rather flow from it. He runs no risk at all of failing to weigh the evidence and arguments 'correctly,' and so failing to make the right decision. And yet, in another sense, when he is actively engaged in his deliberations, it seems to me that he cannot possibly think this. He must think of himself as trying to figure out where the arguments point: which decision they indicate as correct. He must think of himself as being led rather than leading. He must think of himself as aiming

to conform his judgment to an independent fact in virtue of which it will be either correct or incorrect.⁶²

We now have the resources to state a little more clearly the situation that Rosen is trying to describe. Let us understand Rosen's references to the "engaged" perspective with respect to F to refer to the perspective of the subject whose judgments metaphysically determine how things stand with F . Let us understand the "anthropological" perspective with respect to F to be the perspective of a subject who knows full well the real criteria for F -hood; in particular, he knows that the 'engaged' subject's judgments are metaphysically determinative with respect to F .

Engaged perspective (w.r.t judgment-dependent domain F): perspective of the subject on whose judgments F metaphysically depends.

Anthropological perspective (w.r.t judgment dependent domain F): perspective of any subject who is familiar with the actual metaphysical conditions for F -hood.

We now also have the resources to explain why the engaged judge "must think of himself as aiming to conform his judgment to an independent fact:" he must think this if F is a property that depends metaphysically on first-personally investigative judgments. In such a case, unless the subject's judgments were first-personally investigative, we wouldn't have the proper metaphysical basis for the domain in question.⁶³

⁶² *ibid.* 304

⁶³ Compare: on a metaethical view on which the property *goodness* depends on our considered moral sentiments about what is good, merely looking at facts about what we *stipulate* to be good wouldn't be relevant to determining the distribution of facts about goodness.

Finally, we can now diagnose why the tension comes about: why, that is, in a “certain sense he must concede that the facts do not constrain him,” while “in another sense ... he cannot possibly think this.” The tension occurs because Rosen is describing a case in which the anthropological and engaged perspectives coincide in a single judging subject: a subject who, in order to issue a valid judgment, must take himself to be investigating an independent subject matter, but who also knows that his judgment is what metaphysically determines how things stand with F.

Now, Rosen concludes that the tension is a real one. He writes, “It seems to me a very interesting question what to say about the ‘phenomenology’ of the engaged perspective. ... Does rational deliberation involve an inevitable illusion to the effect that one’s verdict might fail to get things right? Obviously, this requires further investigation.” At the same time, however, he doesn’t take the tension between perspectives to be a real worry as far as investigation into *realism* and *objectivity* is concerned:

Nonetheless, I am convinced that this does not bear directly on the *metaphysical* question. It seems to me clear that a) the facts described by a discourse whose central concepts are first-personal judgment-dependent concepts are in principle describable from the anthropologist’s stance in an entirely third-personal idiom; and that b) the facts the anthropologist describes – the properties he attributes to objects – are, for all we have said, entirely *objective* in the sense that interests us.”⁶⁴

Two important conclusions emerge from this part of Rosen’s discussion. First, he seems to suggest that the tension that occurs when engaged and anthropological perspectives coincide is, at worst, a matter of tricky “phenomenology;” it doesn’t bear on any metaphysical questions about the subject at hand. Second, he transitions

⁶⁴ *ibid.* 304, my emphasis

smoothly – indeed, as he does throughout the paper – between speaking of the *reality* of the entities in question – that is, presumably, “the metaphysical question” – and their *objectivity*. In particular, his four-part argument is meant to establish, all at once, both the reality *and* the objectivity of any particular judgment-dependent domain.

In the following section, I want to challenge both of those positions. First, I will argue that the conflicts that come up as a result of misaligned perspectives are not limited to phenomenology: they can have a *metaphysical* bearing on the state of affairs in question.⁶⁵ But secondly, I want to suggest that there is good reason for starting to insert a wedge between the metaphysical side of things – whether the entities in question are *real* – and their objectivity: this latter question I will take up more directly in Section V. To anticipate a bit, I will argue that judgment-dependent entities are indeed fully *real*: one can concede that this is a plausible result of Rosen’s four-part argument.⁶⁶ But I will argue that that leaves untouched a further question about the *objectivity* of the relevant domain, and the judgments that purport to describe it. I will end by sketching out a conception of what the objectivity of a certain domain, over and above its reality, might require.⁶⁷

⁶⁵ Rosen claims that “from a metaphysical point of view, biography and autobiography are on a par.” One way to put my point is to say that, from a metaphysical point of view, autobiography can have a very real effect on biography.

⁶⁶ It is also what I have arguing in Part I above against worries raised by Thomasson.

⁶⁷In drawing a distinction between reality and objectivity, it’s important to consider what the object of those designations is: the reality and objectivity, in other words, of what? Different writers take different approaches to answering this question. Rosen’s discussion, for example, is in terms of *facts*, or, occasionally, “domains of facts.” Thomasson’s discussion, on the other hand, is largely in terms of *entities*, including both kinds and particular objects. Both of these approaches have certain drawbacks. Thomasson’s approach seems to work well for questions of realism, but less well for questions of objectivity: it makes sense to ask whether a certain object is real, but what does it mean for it – some particular thing – to be

Assumption #1 – a matter of mere phenomenology?

Let's take up Rosen's call for "further investigation," and work, step-by-step and rather slowly, through a toy example analogous to the case he describes above: a situation in which a) the anthropological and engaged perspectives coincide regarding a subject matter that depends metaphysically on b) first-personally investigative judgments. I'll argue that the tensions that arise in such a case are not limited to phenomenology, but can have very real effects on the metaphysical status of the domain.

Here is the set-up. Take some first-personally investigative judgment-dependent property *F*, such that the engaged subject issuing *F*-related judgments must be taking himself to be responsive to existing facts about *F*-hood. The metaphysical relation between the judgments and *F* is such that the engaged

objective? Rosen's approach in terms of domains of facts seems more amenable to discussions of objectivity: a "domain of facts" seems to incorporate talk of not just objects, but also their relations, and perhaps even perspectives onto those objects and relations. But there is a certain unnaturalness, it seems to me, in distinguishing between "real" and "unreal" facts. ("...The world is determined by the facts, and by these being *all* the facts," is how Wittgenstein begins the *Tractatus*.) What could an *unreal* fact be? In my discussion, I take a hybrid approach: I follow Thomasson in applying designations of reality or irreality to concrete objects or entities, and I roughly follow Rosen in applying designations of objectivity or non-objectivity to domains of subject matter. By subject matter, I intend to mean something very inclusive: concrete entities and objects are part of what a particular subject matter comprehends, but so are relations between these entities, facts about perspectives onto them, rules or laws governing their interactions, and so on. I am thinking roughly of what T. M. Scanlon calls "domains," in distinguishing broadly between physical, normative, and mathematical reality. I call these domains of *subject matter*, rather than of *fact*, to explicitly signal the possibility of our own conceptual practices contributing to a determination of what is, or is not, part of a particular domain.

Apart from these basic commitments – that reality is attributable to particular objects, while objectivity is attributable to a more inclusive realm that includes relations between objects, the laws governing them, facts about perspectives onto them, features taken to be essentially or conceptually definitive of them, and so on – I hope that the main elements of my discussion are compatible with multiple ways of spelling out these distinctions.

subject's judgment that x is F is, for any x , a necessary condition for x to be F . The following biconditional captures the metaphysical relation between being F and being judged by S to be F :

$$x \text{ is } F \text{ iff } x \text{ is judged by } S \text{ to be } F$$

Suppose, as in Rosen's example, that the anthropological and engaged perspectives coincide in S : that is, he is aware of the criteria for F -hood, and his judgment is determinative of whether F is the case. Let's see what happens when S tries to determine whether x is F .

Well, since there's only one condition for F -hood, S must begin by figuring out whether or not that condition is fulfilled – that is, he needs to figure out whether or not he judges that x is F . But how is he supposed to do that?

Of course, the natural route to determining what you yourself judge about some matter M is just to reflect directly on M , but here this route is closed: direct reflection on F only leads S to conclude that one element determining F -hood is still missing, namely his own judgment as to whether or not F obtains; he finds himself back at square one.

But maybe S could try to go about things in a more roundabout way, determining what he judges with respect to F not by reflecting on F directly but by gathering evidence about his judgment. To that end, perhaps he tries to recall notable features of himself on previous occasions that he judged that F , and checks if those features are currently present; or maybe he performs a kind of statistical analysis to see how frequently he judged that F in similar situations in the past.

The problem here, though, is that as long as S is occupying not just the engaged perspective but also the anthropological one – that is, as long as he is aware that it is his judgment, and nothing else, that determines whether or not x is F – then S will be aware that any process of trying to gather “evidence about his judgment” is in reality a process of *eliciting* some judgment. And once that is clear, then the purported evidential import of facts about his past behavior loses its purchase. That’s because statistical or circumstantial evidence about how you acted in the past doesn’t on its own provide a rational basis for determining how you will act in the future: its import for determining your future behavior comes instead from the potentially useful information it provides about other factors bearing on how you will act.

Suppose, for example, that one day you find yourself wondering whether you should walk to work on the north or south side of the street – usually you just start walking without giving the matter a second thought, but today the question presents itself to you explicitly. In this case, evidence about your past behavior can play an important role in providing rational grounds for making the decision: for example, if you’re told that nine prior times out of ten you picked the south side, that may give you evidence that it makes for a more pleasant walk. Or maybe you value consistency: in that case, the mere fact that you typically pick the south side itself gives you reason to stick to that side. But in every case in which evidence about your past behavior gives you reason to act a certain way, it does so in virtue of its relation to other factors: by providing evidence of further reason-giving facts, or by standing in some relation to prior preferences you might have. Facts about past behavior

considered wholly in isolation do not themselves provide a rational basis for acting one way or another.

But that is precisely the case in the situation that we're considering: here, when S seemed to have nothing F-related on which to base a F-regarding judgment, we invited him to consider evidence about his own past behavior, hoping to thereby gather evidence about how he is generally inclined to judge, and thereby secure some rational basis for the present judgment about F. But if this evidence is the only thing S has to go on, then it isn't linked in any way at all to any external factors arguing in favor of or against F (since, by hypothesis, there are no such factors – all that determines F-hood is S's judgment to that effect). So this kind of evidence becomes wholly irrelevant to the question that S finds himself facing: i.e. it becomes irrelevant to the question of whether S should currently judge x to be F or not.

The result seems to be that strict adherence to our initial assumptions – that is, a) coincidence of anthropological and engaged perspectives with respect to a subject that is b) metaphysically dependent on first-personally investigative judgments – yields the result that no first-personally investigative judgment on the matter is possible at all.⁶⁸

But this isn't, contra Rosen, merely a matter of the “phenomenology of the engaged perspective.” Because of the metaphysically determinative effects of S's

⁶⁸ Here we must be careful to avoid the temptation to tacitly shift our initial assumptions. Of course, in such a case S can certainly issue *something* that might be called a “judgment”: he can throw up his hands and just choose whether x is F, or he can momentarily put out of mind what the real conditions for F-hood are. But both those options would be to abandon the demands of our initial set-up, which required first-personally investigative judgments (in which the subject takes himself to be judging based on independent F-related facts), and coincidence of anthropologist and engaged perspectives.

judgment, the impossibility of issuing such a judgment has direct bearing on whether F is the case: it has direct bearing, that is, on what Rosen calls the metaphysical question. What this shows is that “merely” perspectival questions – how things appear, or must appear, from the engaged perspective – can’t be dismissed out of hand as a priori irrelevant to the metaphysical matters at hand: we have seen that at least one combination of first-personal perspective and metaphysical set up leads to a case where a judgment about F isn’t issuable at all, which itself has bearing on the distribution of F-related facts.

But here one might respond in the following way. It may be true, the response goes, that in certain carefully constructed situations like the one above, “merely” perspectival facts may have metaphysical effects, at least insofar as they can render the issuing of a judgment impossible. But while we can, in principle, set up a situation with these parameters, most *real* instances of judgment, even those in which first-personally investigative judgments have determinative effects, won’t look like the situation above. After all, what real domain – ethics, aesthetics, or anything else – would demand both first-personally investigative judgments having determinative effects, *and* perfect coincidence of anthropological and engaged perspectives at all times? So as long as we remove the implausible condition of perfect coincidence at all times of anthropological and engaged perspectives, or loosen the requirement of a strictly outward-looking, first-personally investigative stance, we can arrive at a functional account on which first-personally investigative judgments regarding some subject matter F metaphysically determine the facts about F.

When applied to ethics, for example, the picture would look something like this. Ethical facts, like the fact that torture is wrong, are metaphysically dependent on judgments made by members of the relevant community; these judgments are first-personally investigative, in keeping with our ordinary conception of what we're doing when we engage in ethical reasoning, but have a metaphysically legislative effect over the domain of ethical fact. Our metaethical theory describes this structure of metaphysical dependence: it is the anthropological perspective. But our first-order normative ethical theory lays out the reasons and principles that our ordinary first-personally investigative judgments actually rely on: it issues from the engaged perspective. When we're engaged in ethical reasoning, we allow those first-order normative principles and reasons to govern our thinking, and are therefore able to issue the judgments that (as evident to the anthropological point of view) turn out to be the metaphysical basis for the truth of the domain they describe. Call this whole view Judgment-Dependent Ethics.⁶⁹

The view, admittedly, would offer a number of advantages. First, it would allow us to be straightforward naturalists while maintaining full-blown realism about ethical facts: the metaethical theory ensures that normative facts really exist, because they're metaphysically grounded in the perfectly natural facts about what judgments are made by a particular group. Second, it would respect the common sense conviction that ethical judgments are first-personally investigative: they present themselves to the judging subjects as responses to an independent realm of ethical fact. Better yet, even as Judgment-Dependent Ethics maintained both of these

⁶⁹ There are some parallels between this position and Sharon Street's constructivism about reasons.

seemingly incongruous positions, it wouldn't even really count as an error theory! To be sure, there would be *some* kind of error attributable to the engaged perspective. But the error would not occur at the level of first-order normative judgments, like "stealing is wrong." Those judgments would actually be perfectly true, describing, as they do, a domain of real facts – real, because metaphysically dependent on the relevant judgments. The error would instead occur in a kind of mismatch between our first-personally investigative stance, and the truth that emerges in the metaethical theory: that our judgments are actually metaphysically determinative. But how damaging could that kind of error really be? After all, Judgment-Dependent Ethics holds that ethical facts exist and are real – and isn't that what would matter most to the engaged perspective anyway?

We are here approaching my objection to Rosen's second assumption: the easy transition from speaking of realism about the entities belonging to a given subject matter – the metaphysical question – and speaking of that subject matter's objectivity. Judgment-Dependent Ethics is, indeed, a position on which ethical properties are real: the way the facts are hooked up to our judgments ensures that that's the case. But would it be right to say that the subject matter of ethics, described that way, could be properly called *objective*?

Assumption #2: Reality = Objectivity?

One peculiar feature of Judgment-Dependent Ethics is that it makes the reality of ethical facts apparent only from a perspective that is inaccessible to the subject qua engaged judger: it is only the subject qua anthropologist who is in a position to

appreciate the reality of ethical properties and facts. Of course, the domain *appears* to be real to the engaged perspective as well, but that's owing to an illusion that that perspective must labor under if it is to issue any first-personally investigative judgments at all. So the engaged perspective finds itself in a kind of Gettier case with respect to the reality of the facts in question: they are real, as it turns out, but not for the reasons that he thought they were.

There's something very unsatisfying, I want to suggest, about classifying as 'objective' a domain in which we must posit that our perspective as engaged judges is to some degree misinformed and, if we are to continue forming the judgments on which the domain depends, necessarily so. To the extent that the engaged perspective is operating under the appearance of investigation when in fact the case is one of determination, the engaged perspective doesn't represent our full knowledge of the situation; to some degree, then, the engaged perspective doesn't really represent *us*.⁷⁰

V. Objectivity vs. reality: some case studies

In this section, my goal will be to investigate and sharpen this intuition, and to try to get a little bit clearer on what objectivity, as opposed to reality, might involve; I'll begin by looking at two more examples of cases in which reality and objectivity

⁷⁰ Compare McDowell's suggestion that what we want in vindicating a domain as real is, in part, "a style of explanation that makes sense of what is explained (in so far as sense can be made of it.)" For instance, if we are to characterize the property of fearfulness in a way that "is an 'attempt to understand ourselves'" as subjects who perceive certain things as fearful, "then merely causal explanations of responses like fear will not be satisfying." Similarly, I'm suggesting, we will not really "understand ourselves" as moral reasoners so long as we associate with moral reasoning a set of deep and necessary confusions. (McDowell 1985)

seem to come apart. From these examples, I'll try to tease out a few elements that seem to be components of an intuitive understanding of what it takes for a subject matter to be *objective*, as opposed to merely be describing entities that are *real*.

I should be clear that my goal is not a full descriptive analysis of 'the' concept of objectivity, if there is such a thing. Instead, it is to point out some tensions that arise when no distinction at all is drawn between reality and objectivity, and to use those tensions to identify some potentially interesting distinctions we might be attuned to; I think there is good reason for thinking that these distinctions do, in fact, already belong to our ordinary understanding of what it takes to be objective.

Case 1: The Property U

Consider the case, popularized by Nancy Mitford and rendered as a philosophical example by Philip Pettit, of the concept *U*. Suppose that to be *U* or not-*U* is to be judged to be so by a Sloane: "whether something is *U* or not is a matter of the say-so of those in the appropriate set; the members of that set have an authoritative, dictating role in regard to the concept."⁷¹ Translating this from the conceptual register into the metaphysical one in which we've been speaking, we can say that to be the *property U* is to be, at least in part, such as to be judged to be *U* by a Sloane.

Assume, further, that there is something a little shifty about the way judgments regarding *U*-ness are made. Pettit writes:

I assume that there is something distinctively collusive in the way Sloanes use the *U*-concept: that as they individually decide whether something is *U* or non-*U* they look over their shoulders to make sure they stay in step – the community is the authority – rather than looking to the thing itself to see what

⁷¹ (Pettit 1991, 611)

profile it displays. ... As the regular bourgeoisie try to get in on the game, Sloanes are notorious ... for shifting the extension of the U-concept.⁷²

What would Rosen's argument lead us to conclude about the property *U*?

Well, we have stipulated that to be *U* is, as a metaphysical matter, to be such as to be judged *U* by a Sloane. Now, recall Rosen's premise (2): that there is "no reason to think that the facts about what a certain group of people would think after a certain sort of investigation are anything but robustly objective." Since the facts about *U*-ness are metaphysically dependent on facts about the Sloanes' judgments – they can be "read off in a mechanical way" from the latter – we conclude from Rosen's basic argument that *U*-ness is itself real and objective.

Let us concede that *U*-ness constitutes a perfectly *real* property; I have been arguing, against worries raised by writers like Thomasson and in the spirit of Rosen's basic argument, that it cannot impinge upon the reality of *U*-ness that it metaphysically depends on mental phenomena in general, and so on judgments regarding *U*-ness in particular.⁷³

But is it very satisfying to conclude, as Rosen would seem to, that facts about *U*-ness are also thereby *objective*? Recall what Pettit specified about there being something 'distinctively collusive' about the way facts about *U*-ness are judged – the strong sense that something about *U*-ness is absurd, or underhanded, or not what it presents itself to be. To my ear, a very natural way to put this worry would be to say, "look, the thing about calling things *U* or not-*U* is that it's not *objective* – those guys are just making it up as they go." But by Rosen's lights, this natural-seeming criticism

⁷² *ibid.*

⁷³ See Part II: Realism and Mind-Dependence, above.

would be wholly off the mark: so long as U-ness depends, as it does, on judgments of U-ness made by the Sloanes, U-ness is as objective as anything else. So if we're to vindicate that natural-seeming criticism, we'll need to find a different criterion for objectivity than the one Rosen provides.

For our next case study, let's turn to a piece of actual critical social philosophy:

Case 2: Haslanger's ameliorative accounts of gender and race.

Haslanger has offered accounts of race and gender on which these categories are fundamentally grounded in social, rather than biological, facts.⁷⁴ On this account, roughly:

S is a woman iff S is systematically subordinated along some dimension, and S is "marked" as a target for this treatment by observed or imagined bodily features presumed to be evidence of a female's biological role in reproduction.⁷⁵

So we have in this view a species of the judgment-dependence accounts that we've been talking about throughout: the kind *woman* is unified by standing in a certain relation to the responses and classificatory practices of a particular group.⁷⁶

Now, Haslanger's account isn't meant to offer an analysis of our existing understanding – in her terms, the *manifest concept* – of what it is to be a woman;

⁷⁴ (Haslanger 2012)

⁷⁵ Haslanger offers analogous sets of conditions for what it is to be a member of a particular race.

⁷⁶ This isn't an example of pure judgment-dependence, however, because part of the basis for establishing an individual as a woman is how she is actually treated (i.e., whether she is subordinated), in addition to how she is judged (i.e. "marked as a target" for certain treatment.)

indeed, it contravenes many typical pretheoretical conceptions of gender and race. It is, rather, an *ameliorative* account: it is meant to draw attention to a particular kind that, whether it aligns perfectly with the manifest concept or not, is one that we have reason – political, theoretical, ethical – to pay attention to. In this case, Haslanger’s ameliorative concept of *woman* is informed by considering what kind we are *actually* tracking when we speak of “women:” in Haslanger’s terminology, it is informed by the *operative concept*.

So what should we make of the reality and objectivity of the kind *woman* as picked out by Haslanger’s ameliorative account? By now, I hope to have made my stance on the reality of this kind clear: as I’ve been arguing throughout, the pattern of judgments and responses on which, in Haslanger’s view, the kind *woman* depends, is the metaphysical basis for a kind that is just as real as any we might expect to encounter. Rosen’s four-part argument for establishing an entity’s reality is a way of making this especially clear: since the kind *woman* stands in a relation of metaphysical dependence to a pattern of (perfectly real) facts about human judgments and practices, the kind *woman* is itself real.

But even if we assume that *woman* is a fully real metaphysical kind, is it right to say that the subject matter of gender as a whole, including the practice of classifying people on that basis – of picking them out as men or women – is objective? As in the case of U-ness, I feel a strong inclination to resist that way of describing things. The account that Haslanger presents is naturally described as a *debunking* account, one that reveals the basis of our judgments about who is and isn’t a woman to be rooted in a kind of confusion and misapprehension of the actual facts,

a misapprehension that is only perpetuated by the continuing practice of classifying people into women and men. So while Haslanger's notion of *woman* picks out a real kind, there's an obvious danger in leaving the story just at that: that is, in understanding Haslanger as merely improving our understanding of a robust and unproblematic category into which we can continue classifying people, much as before.⁷⁷ I will return to these intuitions further down and try to make them more precise; for now, grant only the idea that we can fully appreciate the reality of the kind *woman* while resisting a characterization of gender-regarding subject matter as fully objective.

VI. Why less than full objectivity?

What the examples of *U* and gender show is that there are certain properties with respect to which we experience a kind of tension. On one hand, we want to acknowledge that they're real – they have, as Catherine MacKinnon has put it, “all the indeterminacy of a bridge abutment hit at sixty miles per hour.” At the same time, however, it doesn't sit quite right to call the subject matters to which judgments about these properties belong objective: we want to leave room to insist that something about them is capricious, or arbitrary, or confused, in a way that seems to impinge precisely upon their objectivity. Rosen's criteria, however, are insensitive to these concerns: his joint requirements for reality/objectivity rule that these domains

⁷⁷ Compare with Judith Butler: “Is the construction of the category of women as a coherent and stable subject an unwitting regulation and reification of gender relations? And is not such a reification precisely contrary to feminist aims?” That is, even if our new understanding, informed by feminist and moral theory, gives us a truer and clearer view of existing gender categories, we might still have good reason to resist the continuing application of these (admittedly improved) conceptions. (Butler 2003)

are perfectly real *and* perfectly objective. So while they help make intelligible the first half of the tension – the sense in which these domains describe a real set of facts – they neglect the second half: the sense in which something about these domains seems to fall distinctly short of objectivity.

But if Rosen’s criteria are insensitive to the ways in which facts about *U* or *woman* seem to fall short of full objectivity, we’ll need to look elsewhere for grounds on which to draw the distinction.

Criteria for objectivity – first pass

One candidate emerges from Pettit’s own discussion of how *U* differs from other, more realist, response-dependent properties like redness.⁷⁸ He suggests that the distinguishing feature is what he calls the wholly “dictatorial” nature of *U*-ness:

When we look to what determines the [*U*-ness- and redness-fixing] responses then we do indeed find a significant difference. *U*-responses are determined ... by the efforts of Sloanes to keep in step with one another in their classification of things. But clearly red sensations do not generally spring from such collusive machinations. ... When subjects see something as red ... they do so ... because it is a certain way: ... red.

Ask the Euthyphro question with *U*-ness and the unambiguous answer is that something is *U* because it evokes the *U*-response in suitable subjects. Ask the question with redness, and the answer is less straightforward: in one sense something is red because it looks red to normal observers, in another sense it looks red to normal observers because it is red.⁷⁹

⁷⁸ Pettit’s discussion is in terms of reality rather than objectivity – he wants to claim that taking seriously the existence of *U* things “undermines realism” in a way that taking seriously the existence of red things does not. But because I think we should agree with Rosen that both redness *and* *U*-ness have a full claim to reality, I will instead apply Pettit’s arguments to the question of objectivity: that is, I will consider whether what he takes to be realism-compromising about *U* is in fact what we should take to be objectivity-compromising.

⁷⁹ (Pettit 1997, 613)

Pettit argues that the key difference between redness and U-ness is that the case of redness involves a kind of two-way dependence: normal redness-responses are, like U-responses, essential in fixing the relevant extension, but, unlike U-responses, they fix it in such a way that alights on a property that is itself causally robust: the property that those responses pick out is one that *causes* those very redness responses. In the U case, on the other hand, there is no underlying property that can be said to cause the U-responses; instead, the causal explanation for U-responses are altogether unrelated to U itself, and instead include things like a desire to fit in with other Sloanes.

But while Pettit's distinction is perfectly intelligible – on one hand, properties that cause the responses fixing their extension; on the other hand, properties that *don't* cause those responses – it doesn't seem to apply to the cases we're considering.

In particular, it's false that the property U is not causally relevant to the production of U-fixing responses. Recall that by Pettit's own description, to be U is nothing more than to be such as to be judged U by a Sloane: it is, in Rosen's terminology, a "tacitly autobiographical" property: "what we are really talking about" when we talk about things being U or not-U is "what we would think of these things under certain specifiable conditions." But, *also* by Pettit's own lights, Sloanes asking themselves "what we would think of these things" is precisely what drives the generation of U-judgments. Remember that U-judgments are "distinctively collusive": as Sloanes decide whether something is U or non-U, "they look over their shoulders to make sure they stay in step." In other words, they are forming their U-judgments in part on the basis of whether other Sloanes have judged that item to be U: that is,

they are forming their U-judgments in part on the basis of whether that item *is* U. And if we're willing, as Pettit is, to adopt a broad notion of causal powers,⁸⁰ then we should have no trouble seeing how an item's being U – that is, an item's being such as to be judged U by a Sloane – can be perfectly causally efficacious, including where the production of U-judgments is concerned.⁸¹

But if Pettit's reality/objectivity-determining criterion doesn't rule out the reality/objectivity of U-ness, it certainly doesn't rule out Haslanger's account of gender. One causal effect of the existence of the practice of identifying *women* – of marking certain individuals for certain kinds of treatment on the basis of certain presumed characteristics – is that it keeps the practice going. An individual can be marked for a certain kind of treatment on the basis of presumed characteristics precisely because she *has* been marked for that treatment on the basis of those characteristics: the very practice of "marking" becomes evident in the way the marked individuals look, the way they act, the opportunities available to them, and so on. One element that emerges clearly from Haslanger's work is just how real and causally efficacious the kind *woman* really is; this causal efficacy extends to the issuing of further judgments about whether or not someone is a woman.

⁸⁰ Compare Pettit on the causal efficacy of colors: "For any sensation that a colour produces, it is true that that sensation will be attributable to more basic, microphysical properties of the object and of the light that falls on the object. But we can think of the colour as having a higher-level causal relevance to the sensation, provided that the object's having that colour more or less ensures that no matter how things are disposed at the micro-physical level, they will be disposed so as to produce the sensation. The colour may not "produce" the sensation in the most basic sense available for that term but it will be causally relevant provided that it "programs" for a process of basic production." (ibid. 614)

⁸¹ It can also cause, among other things, the price of that item to rise, or a reaction of desire for that item (or perhaps contrarian distaste for it) among non-Sloanes, or for Nancy Mitford to mention it in her article about U-ness.

So, in Pettit's terms, *woman* exhibits just the kind of two-way dependence that is supposed to ensure a response-dependent kind's full reality and objectivity: there is one sense in which something is a woman because it looks like a woman, and another sense in which it looks like a woman because it *is* a woman.

So much for both Pettit's and Rosen's criteria for objectivity: neither worked in spelling out what strikes one as not quite right in describing our example relevant subject matters as objective. I'll now try to develop a different understanding of what it is that accounts for this apparent lack of objectivity, even when the reality of the relevant entities is conceded.

Criteria for objectivity – second pass

Let's start by going back to the Haslanger example, and trying to put a finger on what strikes one as troubling about the objectivity of those kinds. To begin with, the kind *woman* has the peculiar feature that, built into its very identity conditions, is room for a kind of error: a significant mismatch between the kind's real nature, and the basis on which the relevant group issues its judgments regarding that kind. That's because what determines, in part, whether someone is a woman is whether she is *perceived* or *judged* to have certain physical features, a fact that comes apart from whether she actually *has* those features. The (possibly mistaken)⁸² perception

⁸² Usefully, Haslanger's account actually allows for two dimensions of error in the judgment that an individual possesses certain biological features associated with the female role in reproduction. First, one might mistakenly judge that an individual possesses certain specific features (e.g. certain secondary sex characteristics) that she in fact does not possess. But second, one might be mistaken about one's assumptions regarding what biological features are in fact associated with the capacity for a certain reproductive role. The first kind of

of the features is what motivates the judgment, but it is the judgment itself that metaphysically supports the existence of the kind.

What we have here is a version of the situation that we encountered in our toy example above of judging that *F*: a case in which certain first-personally investigative judgments actually play a metaphysically determinative role. The kind of judgment that metaphysically underwrites someone's status as a woman is typically a first-personally investigative one: it is based on perceptions, or at least purported perceptions, of bodily features. But such judgments themselves play a metaphysically determinative role: the judgment, whether veridical or not, is what forms (part of) the metaphysical basis for the kind's existence.

We should be careful to note that the problem here isn't merely that certain first-personally investigative judgments *also* happen to have legislative effects – that alone seems to pose no special problems with respect to objectivity. Consider, for instance, the case of an umpire demanding if a play was fair: he issues first-personally investigative judgments about certain neutral physical features of the play, and then those judgments underwrite whether the play was fair or not. There is no conflict between the investigative and legislative roles of his judgments in this case: what he's *investigating* are the physical features, and what those investigative judgments thereby *determine* is whether the play was fair – two different matters entirely. In Haslanger's example, on the other hand, the first-personally investigative judgments about physical features are themselves taken, by the engaged subjects, to be judgments about whether the individual is a *woman*: in other words, both the

mistake is connected with the conflation of gender and biological sex; the second kind of mistake is connected with an overly essentialist conception of biological sex itself.

first-personally investigative judgments and the metaphysically determinative ones have the very same subject matter.

And here, just as in our toy example, there would arise a considerable instability should anthropological and engaged perspectives on the kind coincide. Consider a case in which everyone in the community adopted the anthropological perspective – in other words, the community’s manifest concept of *woman* was replaced with Haslanger’s ameliorative one. Could our classificatory practices of sorting individuals into women and men proceed just as before?

In such a case, the possibility of issuing the relevant classificatory judgments wouldn’t be strictly impossible, as it was in our toy example, but it seems like it would be significantly hampered. Once it’s brought to light that it isn’t actual biological features that metaphysically underwrite the kind, but rather perceptions – possibly mistaken ones – of these features and subsequent subordination on the grounds of those perceptions, it becomes difficult to sustain the impression that that classificatory practice is particularly useful or desirable.

For one thing, casting the matter of subordination in such stark light may well have the effect of making such subordination less practicable, and therefore tend to undermine that part of the kind’s metaphysical basis. But another impediment to deploying the concept as before would be the revealed mismatch between the basis for our judgments and the actual nature of the kind. If it’s the biological *features* themselves that we’re interested in, at bottom, then why not switch our classificatory practices to a kind that would get at those features directly, rather than by detour of purported perceptions? If, on the other hand, it’s something like the *appearance* of

those features that we're trying to get at, then why not switch to concepts that directly target that appearance, without the misleading and irrelevant connection to actual biology reality? It seems like the general coincidence of engaged and anthropological perspectives with respect to *woman* – a case in which most speakers in the community adopted the ameliorative concept as their own – would be highly unstable over time, resulting in an abandonment of, or at least a significant shift away from, those classificatory practices.⁸³

VII. A *relative* notion of objectivity emerging

This is beginning to give us a clue as to the possible shape of an answer to our initial question: what is it about certain admittedly real kinds that nevertheless strikes us as less than fully objective? We have seen that one thing that seemed to undermine objectivity in the *woman* case was a mismatch between anthropological and engaged perspective as to how investigative, vs. determinative, the judgment in question actually is: the anthropological perspective revealed the judgment to be metaphysically determinative, while the engaged perspective took itself to be issuing judgments that are investigative. So is objectivity always a question of mismatch

⁸³ Of course, it is part of Haslanger's very point that the ameliorative concept is useful precisely in the context of a targeted, historically specific intervention: the ameliorative concept is a tool for theorists and activists to help identify and undermine the very classificatory practices that it describes, thereby working to eliminate the kind that those practices metaphysically ground. So it is no part of Haslanger's intention to help bring about a stable situation in which anthropological and engaged perspectives on *woman* coincide among the general population, and people go on classifying individuals into women and men just as they did before. Instead, one hopes, the anthropological perspective would reveal the engaged perspective to be ultimately impracticable.

between what the judges think they're doing and what they're actually doing? I think that's part of it, but not the whole story.

Consider the U case, once again. Here, unlike in the *woman* case, we can allow that the engaged perspective on U-ness (the perspective of the Sloanes) and the anthropological perspective (the perspective from which the real criteria for U-ness are apparent) unproblematically coincide: that is, the Sloanes might well know that their practice of classifying things as U or not-U is, in Pettit's terminology, wholly dictatorial; they might be perfectly aware of what they're doing as they "look over their shoulders to make sure they stay in step." Here, despite the consistency of anthropological and engaged perspectives, there still obtains, I think, the appearance of less than full objectivity.⁸⁴ Why might that be?

I think in this case it's crucial to consider the wider context in which judgments of U-ness take place: one in which "as the regular bourgeoisie try to get in on the game, Sloanes are notorious ... for shifting the extension of the U-concept." I think that what gives the appearance of less-than-full objectivity is that U-ness is presented, perhaps not to the Sloanes themselves, but to the highly contextually salient perspective of 'the regular bourgeoisie,' as something that it's not: it gives the appearance of an investigative matter, when it is a fully determinative one.

This suggests that a crucial component of objectivity is alignment between how investigative vs. determinative a certain classificatory practice *appears*, and how investigative vs. determinative it in fact *is*. But how a practice appears is a matter that

⁸⁴ Captured in Pettit's impression, among other things, that something about taking U-ness seriously would "undermine realism."

depends on whom it is appearing *to*: it isn't possible to specify the practice's appearance until we specify the audience in question.

In the examples we've considered, we've seen several different combinations of misalignment between how investigative vs. determinative a practice appears, depending on who's looking. In the toy F example, there was a strict incompatibility regarding investigation/determination between engaged and anthropological perspective, making a judgment impossible to issue. In the Haslanger examples, there wasn't a strict incompatibility between engaged and anthropological perspectives, but there was an important sense in which the determinative nature of the judgments was largely invisible to ordinary engaged judges. This made it the case, I argued, that general adoption of the anthropological perspective would likely undermine general issuing of engaged judgments. Finally, the U case was an example in which engaged and anthropological perspectives unproblematically coincide (in the Sloanes), but are markedly inconsistent with *another* highly relevant perspective (that of the "regular bourgeoisie," who treat U-ness as a more first-personally investigative matter than it in fact is.)

On the view that is emerging, then, objectivity is a matter that is *relative*, rather than absolute. Further, it is relative along two separate dimensions.

The first dimension of relativity that emerges on my account is that the degree to which F-directed judgments are *in fact* investigative/determinative doesn't on its own determine F's objectivity: what matters is the *difference* between how investigative/determinative the judgments really are and how investigative/determinative they appear to the relevant perspectives.

This helps to explain why there are plenty of domains in which our judgments are purely determinative, but for which questions of objectivity never so much as arise. Consider the properties of being a Criterion Collection movie or Michelin-starred restaurant:⁸⁵ these are paradigmatic instances of properties that depend almost exclusively on purely determinative judgments by certain groups. If mere dependence on determinative judgments were enough to undermine objectivity, then these properties should be natural objects of objectivity-based criticism and debate. But of course they're not – they're precisely the kinds of judgment-dependent properties that Rosen rightly dismisses as obviously and boringly benign. My account explains this fact by pointing out that the determinative status of the judgments underlying these properties is perfectly clear to all relevant perspectives. No one is confused about whether being chosen to be part of the Criterion Collection is a matter that supervenes directly on certain judgments issued by certain judges: the anthropological, engaged, and contextual perspectives all agree.

Where things become tricky, of course, is where the question tends to shift, sometimes quite imperceptibly, to matters on which there *isn't* uniform agreement as to the investigative/determinative nature of the judgments in question. One such tacit shift, for example, is to the question of whether this movie *merits* being singled out as an exemplary cultural object, or whether this restaurant *in fact* serves delicious and inventive food. Those questions *are* perfectly apt, and indeed, perfectly familiar, cases to consider in discussions of objectivity. But these are precisely the

⁸⁵ Taking care to distinguish these properties from the (perhaps related but definitely distinct) ones of *being a good movie*, or *being a good restaurant*!

cases in which agreement as to determination and investigation is difficult to come by.

There is also a second way in which objectivity is, on my account, a relative notion. As I've been suggesting, whether a domain is objective or not is a matter of alignment between perspectives onto it; therefore, the objectivity of a domain is a matter that's relative to which perspectives onto it one considers. This means that a domain that is perfectly objective when one considers a certain set of perspectives may cease to be objective when one considers a different set, and vice versa.

This helps explain the somewhat uneasy sense one may get in calling the domains to which the property U or the property *woman* belong straightforwardly objective or not-objective, full-stop: there is a sense in which one might want to say the relevant practice *is* objective, and a sense in which one wants to insist that it isn't quite. My account allows us to explain this by reference to what perspectives are relevant to the evaluation. In the context of a group consisting only of self-aware Sloanes (or only of detached, Nancy Mittford-like observers), it would be natural to describe U-ness as fully objective. But once we start taking seriously the perspective of the regular bourgeoisie, whose fascination and interest in U is largely what gives the property its special social significance, this characterization becomes much less appropriate. If this fascination is successfully dispelled – if the bourgeoisie joins the Sloanes in the less charmed, anthropological perspective onto U – then the objectivity

of the property would, it seems, be regained⁸⁶ (while the role it played in the community couldn't help, it seems, but change.)

Or consider a small group of Haslangerian feminist theorists observing the community's practice from a kind of theoretical remove. In this context, there's a perfectly clear sense in which to be a woman is a straightforwardly objective property. Want to see if someone is a woman? Just look at how she is treated, and on what basis. But what if those theorists must now present a talk on the subject to a group of average watchers of the evening news: people roughly familiar with concepts of sexism and oppression, but largely ignorant of basic feminist social theory? In this new context it would be natural, and perhaps even compulsory, to stress the sense in which the state of affairs regarding who is and isn't a woman is actually *not* objective: you all thought you were tracking straightforward biological reality, but what you're actually doing is sorting people into hierarchical systems of domination. My account allows for both understandings of the kind *woman* – objective on one hand, not objective on the other – to be seen as relevant and illuminating, depending on the context.

Conclusion

We can now formulate with a little bit more precision the set of intuitions that we've been pursuing throughout the discussion: objectivity and reality are two different features that we might attribute in evaluating a particular subject matter.

⁸⁶ U-ness would, in this case, come to seem more like the properties described above – those like 'Michelin-rated restaurant' and 'Criterion-collection movie' – transparently based on determinative judgments.

Reality is absolute: an entity is either real, or it isn't. Objectivity, on the other hand, is relative along two dimensions.

To determine whether a certain subject matter is objective, we need to a) contrast the degree to which judgments about that subject matter are *in fact* legislative vs. determinative with the degree to which they are *regarded* to be legislative vs. determinative. And that requires b) specifying the perspectives being evaluated, including which perspective(s) are being taken seriously for the purposes of the evaluation. By judgments "taken seriously," I mean just those that are taken, for the purposes of the evaluation, as contenders for an accurate representation for how things stand with respect to F.

More specifically, a subject matter F will fail to count as fully objective if one or more of those perspectives taken seriously regard F-directed judgments to be more investigative than the judgments in fact are; (that is, more investigative than the *anthropological* perspective regards them to be, where the anthropological perspective, recall, is what we call the perspective familiar with the actual facts regarding subject matter F.)

To see what's happening more clearly, let's return to the U example. If the only perspective taken seriously for the purposes of evaluation is that of the Sloanes, then U-related subject matter emerges to be fully objective: the Sloane perspective aligns with the anthropological perspective in regarding Sloane-issued U-directed judgments as legislative, and all other U-directed judgments as investigative. On the other hand, if among the perspectives taken seriously is that of the regular bourgeoisie, then U-related subject matter emerges *not* to be objective: the regular

bourgeoisie take all U-directed judgments to be investigative, while the anthropological perspective regards some of them to be legislative.

The question of which perspectives to take seriously for the purposes of evaluation is a potentially important one; it will be answered by considering what our objectives are in performing the evaluation in the first place. If the goal is abstract empirical inquiry, then the only perspective to take seriously is the anthropological one, describing how things really stand; our aim, in these cases, will be to align our own perspective as much as possible with the latter. On the other hand, if our aim is resolving disagreement, eliminating ignorance, or debunking misconceptions, then the perspectives to be taken seriously should include those that we aim to reconcile, enlighten, or correct. On this view, questions of objectivity aren't static, fixed, and ahistorical, but are instead dynamic and live, responsive to the changing shape of the discourse surrounding the subject matter at hand, and not just to its representational content.

This notion of objectivity doesn't fit very cleanly into either of two ways in which objectivity is frequently understood. On one hand, people sometimes speak of an *epistemic* notion of objectivity, which has to do with how inquiry into a set of facts proceeds.⁸⁷ This isn't quite the notion I'm after, because questions of perspectival orientation aren't just questions of inquiry into some subject matter held fixed: sometimes, facts about perspective are built in to the very metaphysical structure of

⁸⁷ See (Rosen 1994) and (Haslanger 2012). The epistemic notion is also at play in certain conceptions of "scientific objectivity," understood as a norm requiring that scientific inquiry proceed independently of personal bias or external normative commitments. For a historical discussion of this conception, see (Daston and Galison 2007). For classic criticism of it, see (Kuhn 1962); for criticism from a different direction, see (Longino 1990) and (Anderson 1995).

a domain (as in cases where domains are metaphysically dependent on first-personally investigative judgments.) On the other hand, that notion of objectivity is sometimes contrasted with a more metaphysical one, which has to do with whether a certain set of entities exists, and therefore tends to collapse into questions of realism.⁸⁸ As I've been trying to make clear, this isn't my notion either, since varying perspectives might issue conflicting verdicts of objectivity regarding a subject matter whose existence or reality is uncontroversial. Instead, the notion I've been pursuing lies somewhere in between, taking into account both the metaphysical structure of the domain in question and the structure of the perspectives that are brought to bear on it, while acknowledging the way in which those two elements sometimes interrelate.

* * *

Rosen left us with the suggestion that there just might not be room in the contemporary philosophical landscape for mind-dependent subject matter being anything less than fully objective: once all traces of Kantian trans-empiricism "give way to a naturalism according to which the only minds there are are *parts* of the natural world itself," we're left with no middle ground between the objective and the non-existent.

I'd like to suggest, in contrast, that such a middle ground is perfectly intelligible: debates regarding the objectivity of various domains are often necessary

⁸⁸ This is closer to the sense of "objective" as used in metaethical discourse; for instance, in distinguishing 'objective' from 'subjective' accounts of morality, where the former, but not the latter, are understood as positing the "real" or "actual" existence of moral properties or facts.

and fruitful. But to see this middle ground, we need to clear some space between a purely metaphysical notion of objectivity and a wholly first-personally epistemic one: we need to take into account the interrelation of contextually relevant perspectives onto the facts in question.

References

- Anderson, Elizabeth. 1995. "Knowledge, Human Interests, and Objectivity in Feminist Epistemology", *Philosophical Topics*, 23: 27–58.
- Appiah, Kwame A. 1995. "The Uncompleted Argument: DuBois and the Illusion of Race," in *Overcoming Racism and Sexism*, L. Bell and D. Blumenfeld (eds.), Rowman and Littlefield.
- Barnes, Elizabeth. 2016. *The Minority Body: A Theory of Disability*. Oxford University Press: Oxford.
- Blackburn, Simon. 1993. "Circles, Finks, Smells and Biconditionals." *Philosophical Perspectives* 7, 259-79.
- Bloom, Paul. 1996. 'Intention, History and Artifact Concepts.' *Cognition* 60, 1-29.
- Brower, Bruce W. 1993. "Dispositional Ethical Realism." *Ethics* 103, no. 2: 221-49.
- Burgess, Alexis and Plunkett, David. 2013. "Conceptual Ethics I." *Philosophy Compass*, 8: 1091–1101
- Butler, Judith. 2003. "Gender Trouble." in Cahill, Ann and Hansen, Jennifer (eds.) *Continental Feminism Reader*. Rowman and Littlefield.
- Cavell, Stanley. 2002. *Must We Mean What We Say*. Cambridge University Press: Cambridge.
- Daston, Lorraine & Galison, Peter. 2007. *Objectivity*. Zone Books.
- Dummett, Michael A. E. 1978. *Truth and Other Enigmas*. Harvard University Press.
- Einheuser, Iris. 2006. "Counterconventional conditionals." *Philosophical Studies* 127 (3): 459 - 482.
- Enoch, David. 2011. *Taking Morality Seriously*. Oxford University Press: Oxford.
- Epstein, Brian. 2014. "What is Individualism in Social Ontology? Ontological Individualism vs. Anchor Individualism," in Finn Collin and Julie Zahle (eds.) *Rethinking the Individualism/Holism Debate: Essays in the Philosophy of Social Science*, 17-38. Dordrecht: Springer.
- Fine, Kit. 1994. "Essence and Modality." *Philosophical Perspectives*, Vol. 8, Logic and Language, 1-16

- Hacking, Ian. 2000. *The Social Construction of What?* Harvard University Press: Cambridge.
- Harman, Gilbert. 2000. *Moral Agent and Impartial Spectator*. Oxford University Press: Oxford.
- Haslanger, Sally. 2012. *Resisting Reality*. Oxford University Press: Oxford.
- Holton, Richard. 1991: "Intentions, Response-Dependence, and Immunity from Error"; available at <http://www.ed.ac.uk/~rholtton/home.html>.
- 1992. "Response-Dependence and Infallibility." *Analysis* 52 (3):180 - 184.
- Jackson, Frank & Pettit, Philip. 2002. "Response-Dependence without Tears." *Noûs* 36 (s1): 97 - 117.
- Jenkins, Carrie. 2005. "Realism and Independence." *American Philosophical Quarterly*, 199-209.
- Johnston, Mark. 1998. "Are manifest qualities response-dependent?" *The Monist* 81 (1): 3-43.
- 1993. "Objectivity refigured: Pragmatism without verificationism." In John Haldane & Crispin Wright (eds.), *Reality, Representation, and Projection*. Oxford University Press 85--130.
- Kripke, Saul A. 1980. *Naming and Necessity*. Harvard University Press: Cambridge.
- Lewis, David. 1969. *Convention*, Harvard University Press: Cambridge.
- Longino, Helen. 1990, *Science as Social Knowledge*, Princeton University Press: Princeton.
- McDowell, John. 1998. *Mind, Value and Reality*. Harvard University Press: Cambridge
- 1985. "Values and Secondary Properties," in Ted Honderich (ed.), *Morality and Objectivity*. Routledge 110-129 (1985)
- Moran, Richard. 2001. *Authority and Estrangement*. Princeton University Press: Princeton.
- Nagel, Thomas. 1989. *The View from Nowhere*. Oxford University Press: Oxford.
- Pettit, Philip. 1991. "Realism and Response-Dependence." *Mind*, New Series, 100, no. 4: 587-626.
- Plunkett, David & Sundell, Timothy. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosophers' Imprint* 13 (23).

Putnam, Hilary. 1973. "Meaning and Reference." *The Journal of Philosophy*, Vol. 70, No. 19, pp. 699-711.

-----1975. "The Meaning of 'Meaning'." *Minnesota Studies in the Philosophy of Science*, 7:131-193.

Quine, W. V. O. 1960. *Word & Object*. The MIT Press: Cambridge.

Rawls, John. 1955. "Two Concepts of Rules." *Philosophical Review* 64 (1):3-32.

Richard, Mark. Draft of "Meaning as Species," available at:
<https://markrichardphilosophy.files.wordpress.com/2011/07/meanings-as-species-chap-1-3-21-june-16.pdf>

Rorty, Richard. 1989. "Solidarity or Objectivity?" in Michael Krausz (ed.), *Relativism: Interpretation and Confrontation*, University of Notre Dame Press.

Rosen, Gideon. 1994. "Objectivity and modern idealism: What is the question?" In John O'Leary-Hawthorne & Michaelis Michael (eds.), *Philosophy in Mind*. Kluwer 277--319.

-----1998. "Blackburn's essays in quasi-realism." *Noûs* 32 (3):386-405.

-----2001. "Nominalism, Naturalism, Epistemic Relativism." *Noûs* 35 (s15):69 - 91.

Sayre-McCord, Geoffrey (ed.) 1988. *Essays on Moral Realism*. Cornell University Press.

Scanlon, T.M. 2014. *Being Realistic about Reasons*. Oxford University Press.

Schroeter, Laura & Schroeter, Francois. 2009. *A Third Way in Metaethics*. *Noûs*, 43: 1-30

Searle, John. 1997. *The Construction of Social Reality*. Free Press.

-----2006. "Social Ontology: Some Basic Principles." *Anthropological Theory*. Vol. 6, no. 1, 12-29

-----2010. *Making the Social World: The Structure of Human Civilization*. Oxford University Press.

Schapiro, Tamar. 2003. "Compliance, Complicity, and the Nature of Non-Ideal Conditions." *Journal of Philosophy*. 100 (7):329-355.

Street, Sharon. 2008. "Constructivism about Reasons." *Oxford Studies in Metaethics* 3: 207-45.

Svavarsdóttir, Sigrún. 2001. "Objective Values: Does Metaethics Rest on a Mistake?" In Brian Leiter (ed.), *Objectivity in Law and Morals*. Cambridge University Press 144--193.

Sveinsdóttir, Ásta. 2008. "Essentiality conferred." *Philosophical Studies* 140 (1):135 - 148.

Tahko, Tuomas E. & Lowe, E. Jonathan. "Ontological Dependence," The Stanford Encyclopedia of Philosophy (Spring 2015 Edition), Edward N. Zalta (ed.), URL=<http://plato.stanford.edu/archives/spr2015/entries/dependence-ontological>

Thomasson, Amie L. 2003. "Realism and Human Kinds." *Philosophy and Phenomenological Research* 67, no. 3

-----2015. *Ontology made Easy*. Oxford University Press: Oxford.

Wedgwood, Ralph. 1997. "The Essence of Response-Dependence." *European Review of Philosophy* 3:31-54.

Wiggins, David. 1987. *A Sensible Subjectivism?* Basil Blackwell.

Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. Harvard University Press: Cambridge.

Williamson, Timothy. 2008. *The Philosophy of Philosophy*. Wiley-Blackwell

Wittgenstein, Ludwig. 1922. *Tractatus Logico-Philosophicus*. London: Routledge & Kegan Paul.

Wright, Crispin. 1988. "The Inaugural Address: Moral Values, Projection and Secondary Qualities." *Proceedings of the Aristotelian Society, Supplementary Volumes* 62: 1-26

-----1992. *Truth and Objectivity*. Harvard University Press.

Yablo, Stephen. 1995. "Singling out Properties." *Philosophical Perspectives* 9:477-502.

-----2002. Red, Bitter, Best. *Philosophical Books* 41 (1):13-23.

-----2002. "Coulda, Woulda, Shoulda." In Tamar S. Gendler & John Hawthorne (eds.), *Conceivability and Possibility*. Oxford University Press, 441-492.