Article

Michel-Antoine Xhignesse

Social Kinds, Reference, and Meta-Ontological Revisionism

Runner-up of the 2018 Essay Competition of the International Social Ontology Society

https://doi.org/10.1515/jso-2018-0013

Abstract: Julian Dodd has characterized the default position in metaphysics as meta-ontologically realist: the answers to first-order ontological questions are thought to be entirely independent of the things we say and think about the entities at issue. Consequently, folk ontologies are liable to substantial error. But while this epistemic humility is commendable where the ontology of natural kinds is concerned, it seems misplaced with respect to social kinds since their ontology is dependent upon the human social world. Using art and art-kinds as paradigmatic examples of social kinds, I argue that meta-ontological realism sets conditions that are too strict to apply to social kinds. Nevertheless, I argue that we should not be too quick to embrace the conclusion that our folk theories of social kinds cannot err substantially. By modelling the reference of social kind-terms on that of natural kind-terms, it becomes clear that in both cases, our sole epistemic privilege lies in our ability to pinpoint the subject of our inquiries.

Keywords: Art-kinds; Natural kinds; Reference; Folk-theoretic modesty; Revisionism; Descriptivism; Social kinds; Meta-ontological realism.

I

Julian Dodd has argued that attempts to determine the ontological status of musical works ought to defer to the results of metaphysics rather than to folk intuitions about music (2007, 2008, 2013). When we do so, he thinks, we discover that musical works are eternally existing and uncreated *types* rather than created works. Dodd bases his argument on a metaphysical principle he thinks we should all endorse:

Michel-Antoine Xhignesse, Department of Philosophy, University of British Columbia, 1866 Main Mall, Buchanan E370, Vancouver, Canada, e-mail: michel.xhignesse@ubc.ca

> *Meta-ontological realism* (MR)
> The correct answers to first-order ontological questions are in no way determined by what we say or think about these questions (Dodd 2013, p. 1048).

According to MR, the answers to questions concerning an entity's individuation conditions and ontological status have nothing whatsoever to do with our beliefs or intuitions about them. Our beliefs and intuitions might well have a great deal to do with the structure of our practices and with how their objects are *treated*, but they have nothing at all to do with whether those objects exist, or what shape that existence takes. To beat a dead horse, the correct answer to what water is has nothing to do with what we say or think about it, but rather with its chemical and physical properties: water is H_2O.

The rest of Dodd's work is done by a different principle which follows from MR:

> *Folk-theoretic modesty* (FM)
> Our common-sense ontological views might be substantially mistaken (Dodd 2013, p. 1048).[1]

So long as we concede that an entity *E*'s ontological nature is independent of what we think about *E*s, it follows that *E*s might turn out to be very different from the way we think they are (Dodd 2013, p. 1049). Once MR is granted, FM has been secured; FM is a straightforward cautionary principle enjoining us to bear in mind that common ways of talking and thinking about ontology may be misleading or wrong.

Such epistemic humility is certainly commendable and is clearly required by the ontology of natural kinds, since these resist interaction with the world of thought. But not all kinds are natural – some are best characterized as social kinds, since they depend on the human social world for their existence and properties. Art and art-kinds, for example, are paradigmatic social kinds: an entity's membership in the class of artworks, music, paintings, etc. depends entirely on human actions and interests, rather than on the entity's causal powers, homeostatic property clusters, or microstructural properties. By definition, social kinds resist MR. So what, then, should social ontologists make of FM? Does it follow that we cannot be substantially wrong about the ontological properties we ascribe to social kinds, since those kinds and properties depend on us for their existence?

I think not. While it is true that MR is misapplied to social kinds, I will argue that social ontology nevertheless mandates folk-theoretic modesty, using "art" and art-kind terms as characteristic examples. This is not because social kinds

1 Dodd characterizes FM in terms of *art*-ontological views, but I take this more general principle to be consonant with his views.

are wholly mind-independent, but rather because they have their origins in a relatively robust social network of conventions which exist and operate independently of our thinking. I begin, in §II, by considering the reference of natural kind terms, which helps to motivate MR for natural kinds. In §III I argue that we can profitably extend this account of the reference of natural kind terms to social kinds like "art" and art-kinds by anchoring it in the robust network of conventions that underpins the existence of social kinds. I return to meta-ontological realism in §IV, where I argue that while it sets plausible constraints on the ontology of natural kinds, it proves too strict for social kinds. Instead, I argue that a term's ability to play a particular explanatory role in a theory of a social kind should be cashed out in terms of the capacity of its referent to play certain kinds of roles in our practices. The upshot will be that although our social practices and beliefs are susceptible to error, we cannot err with respect to specifying the *subject* of our inquiries. Finally, having thus secured FM for the ontology of social kinds, I argue in §V that a social kind's ontological nature is inextricably tied to the network of cultural practices in which it is embedded. This is not because a social kind's cultural role determines its ontology, but rather because it sets the domain of discourse, and so determines the kinds of theories in which the kind-term has explanatory value.

II

The case for MR is borne out by closer consideration of the meaning of kind-terms – at least as far as natural kinds are concerned (I will deal with social kinds in §III). The question we have to ask ourselves is whether our beliefs about natural kinds play any role at all in determining what it is that these entities are – and, if so, just what that role is. One possibility is that it does: perhaps these concepts play an ineliminable role in fixing the reference of the kind-terms at issue (see, e.g. Thomasson 2005, p. 223, 2006, 2007, p. 189–190). This would mean that competent users of a kind-term cannot be mistaken about the basic or categorial properties of the kind in question; their categorial concepts could not fail to apply to members of the relevant kinds.[2]

A more plausible alternative comes from Hilary Putnam, who suggested that a natural kind term's reference is fixed by the world itself, not by our theories of the world (1975 and 1990). For Putnam, natural kind reference is determined by the

2 For defences of this infallibilist doctrine, see Ruben (1989) and Thomasson (2003, 2005, 2006, 2007). I offer some reasons to be skeptical of infallibilism in art-ontological contexts in Xhignesse (2018).

causal powers of the entity we pick out when we introduce our natural kind term. The fact that the referents of a term have their causal powers in common – as evidenced by the sameness of their microstructural properties – will then determine which empirical generalizations (theories) those terms can enter into (rather than the other way 'round, as descriptivism would have it). Natural kinds are thus those kinds which can play an explanatory role with respect to various phenomena in virtue of their entering into these kinds of empirical generalizations. Consider "gold", which has many different macro- and microstructural properties: e.g. it is the precious reddish-yellow element that is the most malleable of metals and has atomic weight 196.967, atomic number 79, and specific gravity 19.3 at 20°C. We can certainly *describe* gold in terms of these properties, but its reference is not fixed by that description. If it were, "gold" would change its reference every time we refined that description or proposed an entirely new one to take its place. In fact, it would turn out that for most of human history, dragons, thieves, and warlords have amassed hoards of some nameless stuff. Allow me to explain.

Pre-scientific peoples knew nothing at all about gold's atomic structure, let alone its specific gravity. Yet their ignorance of gold's microstructural properties (and of its actual causal powers) did not prevent them from caring deeply about whether their brooches were made of gold or fool's gold; only one of the two is valuable, after all. Yet the fact that they were ignorant of gold's microstructural properties and of its actual causal powers does not mean that their uses of the term failed to refer, or that they referred to something else entirely. In fact, pre-scientific peoples are not all that different from most of us today, at least as far as a thing's microstructural properties are concerned. Comparatively few people today, for instance, know much about gold's microstructural properties, beyond its being an element: its atomic weight, atomic number, and specific gravity are beyond most of our abilities to recite, let alone properly comprehend. Surely this does not mean that most contemporary users of the term "gold" are incompetent, and fail to refer when they deploy the term.

Certainly not. Language is a cooperative endeavour: individuals do not need to reinvent the wheel for every term in their lexicon. Hastings need not be a pedologist to acquire or make competent use of the term "muskeg"; he can glean at least some of its conditions of application from pictures and novels set in the Arctic.[3] Putnam's insight was that in using terms like "gold" or "muskeg", we rely on a *division of linguistic labour*. That is to say, we do not ourselves need to be in possession of the conditions needed to reliably distinguish between an instance of the kind and other things. All we need is to be in possession of a conventional

3 What he refers to, however, will be *muskeg*, not whatever entities satisfy his *conception* of muskeg, if there are any.

idea of what members of the kind typically look like, act like, or are like (a "stereotype" – Putnam 1975, p. 230;[4] e.g. gold is a yellow metal, and muskeg is a mossy northern bog), and we need to stand in the right kind of causal-historical relation to the referent of the term.[5] The important thing is that the linguistic community as a whole is in possession of more refined means of individuating the kind in question. In other words, when it comes time to distinguish gold from fool's gold and our folk concept or stereotype lets us down, we can defer judgement to some group of experts who can reliably distinguish between the two (Putnam 1975, p. 227–229). In this way, no one individual needs to have a full grasp of some term's application conditions, so long as these are present elsewhere in the linguistic community.

So far as gold is concerned, our competent use of the term today relies on the fact that experts elsewhere in our linguistic communities have a reliable method of recognizing which things are gold and which are not, based on gold's microstructural properties. But what about pre-scientific peoples, none of whom knew anything at all about gold's microstructure, and whose ideas about its causal powers and origin were substantially mistaken (e.g. chrysopoeia)? It turns out that even they had access to a linguistic division of labour which produced more or less reliable results thanks to, among other things, what is known as a "streak test". Streak tests work on the largely correct assumption that metals dragged on a touchstone leave behind a coloured streak. Gold, for example, leaves a golden streak whose hue matches the colour of the object being tested. Streak tests are not foolproof – modern chemistry can defeat them, as can plating the object, provided the assessor does not attempt to test a cross-section – but for the most part they are reliable. The result is that pre-scientific linguistic communities had access to a reliable method for identifying gold and other metals. Some of their beliefs about "gold's" extension certainly turned out to be false, but not nearly so many as if they had been relying on a simple folk theory to guide their use of the term.

So pre-scientific peoples were not incompetent users of natural kind terms, and neither are we. But did the *meaning* of "gold" change once we finally gained access to knowledge of gold's microstructural properties? The answer hinges on just what we mean by "meaning". The *sense* (intension) of "gold" certainly changed, but its *reference* (extension) did not (Putnam 1975, 1990, p. 60). People

4 Stereotypes are like the descriptions associated with proper names on the causal theory of naming: they give us one way of picking out an object, but they do not give us necessary and sufficient conditions for falling into the name's (term's) extension. It is the world around us that supplies its extension.

5 So that, e.g. if Hastings keeps trying to apply "muskeg" to rabbits, he will be all out of luck.

still cared about the same *stuff*; what they stopped caring about was the description of that stuff in terms of "things that streak gold," since they discovered that the description was insufficiently individuative. They shed one *description*, not one stuff, for another. Put another way, the thing pre-scientific dragons cared to hoard was some particular stuff which they individuated by means of the colour of its streak (or perhaps its susceptibility to bite marks); modern-day dragons care to hoard the very same stuff, namely gold-streaky stuff *with the right microstructure*, which they learned about from ancient dragons who stood in the right kind of causal-historical relationship to the term's referent.

The only difference is that modern-day dragons have more sophisticated methods for individuating gold: they no longer make the mistake of thinking that, say, *schmold* (which streaks gold) is gold. When Sigurðr came for Fáfnir's gold, he had to slay the dragon before the hoard could be his. If a modern-day Fáfnir were to discover that half his hoard was *schmold*, however, he would gladly give it up to Sigurðr's modern-day counterpart who, in turn, would presumably refuse it. Saga tales aside, the moral of this story is just that a natural kind term's reference is dependent on its actual nature, not an individual's psychological states or a linguistic community's beliefs. Our collective interests have a role to play in determining the kinds of explanatory *frameworks* in which natural kind terms occur (e.g. scientific vs. folk discourse – see, e.g. the discussion of "jade" in §V), but not in determining that term's extension. *That* much is supplied by the world around us. A term's extension is determined by its reference, not its description, and its reference is secured by the way the world is.

III

Trouble starts to brew when we consider social, rather than natural, kinds. A natural kind's hidden structure determines its kind-membership, but social kinds have no hidden structure in the first place. Consider Hastings, who is an Englishman – let us say he belongs to the kind *English citizen*. As a human being,[6] Hastings himself has microstructural properties: he is made of carbon and other elements, acts through the transmission of action potentials across synapses, etc. But none of these microstructural facts make him an English citizen. So while an individual instance of a social kind may have microstructural properties, the kind

6 Assuming, for the sake of argument, that Hastings is a real person and not just a fictional character. It remains an open question whether or in what sense the fictional character is a human being (see, e.g. Thomasson 1999).

itself does not, and so the instance's microstructure cannot figure in an explanation of its kind-membership.[7] So, too, for "art" – or at least, for visual art like paintings and sculptures, since these are obviously composed of some matter. The important point is just that the particular arrangement of matter is not what makes these entities *artworks*, or *paintings*, or *sculptures*.

Yet even though there is no microstructure for experts to discover and appeal to in their explanations of citizenship, or for them to use to distinguish fraudulent from genuine citizens, the fact of the matter is that there *does* exist an objective measure of citizenship. This measure is given by the complex network of conventions and institutions that figure in explanations of the existence and behaviour of nation-states. That network of conventions, in turn, determines which individuals living within a certain geographic area owe allegiance to the area's government and are entitled to its protection (along with determining the character of the allegiance and the type of protection). In this case, we can simplify by saying that Hastings possesses (or is entitled to possess) an English passport.

Notice that the meaning of "citizenship" is not determined by our folk theory of citizenship, even if citizenship is a social kind; it is codified in a series of legal conditions, documents, and institutions. Determining which individuals are citizens and which are not is thus not merely a matter of what we think, but of determining which individuals satisfy the conditions laid out by English law, such as being born in the right places or under the right conditions, having filed the appropriate paperwork, etc. *Citizenship* may have no microstructural properties, but this fact does not prevent us from dividing our linguistic labours and relying on the relevant "experts" – in this case, bureaucrats. The conditions for citizenship are somewhat arbitrary, to be sure, but this is only to be expected from a kind so thoroughly rooted in convention. The same will hold for other social kinds such as *art*, *money*, *music*, or *weed*.

Contrast this with a term like "chair", which Putnam thinks is not subject to the division of linguistic labour: when a speaker acquires the term "chair", she also thereby acquires something that contributes to fixing its extension (e.g. a functional role which chairs occupy; Putnam 1975, p. 227–228). This is just because chairs are neither natural kinds, nor sufficiently regimented in our social practices to require any kind of division of labour (much like "water" or "gold" for pre-scientific communities). Yet even so, it is not the chair-speaker's individual psychological states which do the work of fixing "chair's" extension; her use is embedded in, and contributes to, the general sociolinguistic state of her linguistic community, which in turn picks out whatever entities satisfy the particular

7 Note also that although Hastings instantiates (belongs to) a social kind, the man himself is not a social kind.

functional role that chairs perform (Putnam 1975, p. 228). In other words, the term's extension depends on the entities that actually serve the relevant functional role.

I should sound a brief note of caution at this point: while it may be tempting to classify kinds like "citizen" as institutional kinds (and thus as being dependent on the existence of the relevant institutions, which set their identity-conditions),[8] we should not proceed quite so quickly. There is surely something right in the idea that, once institutions are established, their members can confer certain kinds status on the entities falling under their purview (as with bureaucrats and citizens). Although institutional theories represent one way of articulating conventionalism, it remains an open question whether they properly characterize art-kinds in particular, let alone social kinds more generally[9] – it is not obvious, for example, that the kind "weed" is institutional, though it is surely social. Nor is it clear that the order of explanation offered by institutional theories is entirely appropriate to the subject matter of first-order ontology; as I have argued elsewhere, institutions are agglomerations of *conventions*, and it is these conventions which perform the ontological work of determining kind-membership, not the beliefs, concepts, or decisions of institutional agents (Xhignesse, forthcoming).[10]

So where does all this leave artworks? The first thing to notice is that "art" is not a natural kind term. It is more or less universally accepted that "art" is an artifactual kind term like "chair" or "pencil" and, thus, a social kind. So what determines membership in the kind *art*? The answer hinges on whether "art" exhibits a linguistic division of labour, and this is where things get tricky, because the evidence is mixed. On the one hand, no art-experts are capable of telling us definitively whether a particular entity is a work of art or not.[11] We certainly have a great many art-experts – artists, art historians, critics, curators, philosophers of art, etc. – and we do often ask them to supply reasons for a work's putative art-status. Consider, for example, the various expert witnesses called to testify before the U.S. Congress on behalf of the National Endowment for the Arts after the controversy over Robert Mapplethorpe's photographs. But these experts do not have access to a privileged method of recognizing art, since art has no properties that

8 See, e.g. Guala (2016, esp. Chs. 11 & 12).

9 See, e.g. Brand (2000), D. Davies (2004), and Walton (2007).

10 Note also that conventionalism has a readymade set of internal constraints on kind-membership (in the form of standard and contra-standard properties), while institutional theories place no internal constraints on kind-membership at all (see Friend 2012, p. 187–195).

11 Some theories of art do confer such a power upon art-experts. Dickie's version of the institutional theory, for instance, sees their verdicts as constitutive of art, and riven through with illocutionary force. Or consider Donald Judd's infamous dictum, "if [an artist] says it's art, it's art."

could ground these kinds of empirical generalizations. There is no art-ontological expertise akin to the expertise an entomologist wields with respect to insects, or that an astronomer has with respect to the formation of gas giants. Nor is it clear that we regularly defer to the judgements of art-experts, as should be evident from the perennial controversies over works of public art such as Richard Serra's *Tilted Arc* (1981). Indeed, in the hue and cry over *Tilted Arc*, expert judgements were defeated by the public uproar, and the work was destroyed.

On the other hand, we do sometimes defer to expert judgements when it comes to membership in relatively well-defined art-kinds, such as sonnets, fugues, film noir, odes, neo-classical ballets, sonnets, etc. Indeed, if Walton (1970) and Lopes (2014) are correct in thinking that categorization according to an art-kind or genre is a necessary first step in the identification of artworks, then it seems that this kind of division of labour plays a very important role in our linguistic community.[12]

My own view is that, all things considered, "art" is rather more like "citizen" than "chair". While it is true that we do not *defer* to the judgements of art-experts for determinations of art-status in general, we do consult them regularly when a work's art-status is in doubt, and we demand that putative experts supply us with reasons for their judgements. And they do so; not in terms of an objective measure of art-status, but by highlighting the kinds of conventions that govern – or have governed – our artistic practices, and showing how the work under consideration reproduces or responds to them. Determining which entities are artworks is not merely a matter of canvassing public opinion; it is a matter of comparing that opinion to the kinds of conventions that ground our institutions, that govern our practices, and that figure in our best theories of those institutions and practices.

The artworld supplies us with an alternative to folk theory in the form of a network of conventions that reproduce by weight of precedent. There *is* a "method of recognizing" present in our community, considered collectively; it just depends on an assembly-line rather than a craftsman model of the division of linguistic labour. Through their participation in the artworld, and through their applications of "art" and art-kind terms, speakers help to reinforce existing conventions by increasing the salience of various precedents.[13] Unlike craftsmen, speakers are not typically in possession of a complete picture of how we go about distinguishing art from non-art; they are alienated from the products of their labour. That is to say, they may not be fully aware of the ways in which they participate in the artworld and reinforce its conventions.

12 I am indebted to Eric Murphy for suggesting this line of thought.

13 I offer a more detailed sketch of artworld conventionalism in Xhignesse (forthcoming); for an illustration of the ways in which conventions govern literature, see Friend (2012).

IV

We are now equipped to return to the topic of folk-theoretic modesty (FM) about social kind terms. We saw in §I that FM follows straightforwardly from meta-ontological realism (MR), the view that the correct answers to first-order ontological questions are mind-independent. The considerations adduced in §II and §III all help to motivate FM for "art" and other social kinds; the question, however, is whether the investigation of social kinds is also constrained by MR.

While it is intuitively clear that MR applies to natural kinds, it is not obvious that MR is also true of social kinds like "art". Ian Hacking and Sally Haslanger, for example, have argued that the very act of classifying something in a social context can change the object of classification, or even establish and reinforce an entirely new classificatory scheme (Hacking 1995, 1999; Haslanger 1995, 2012). Hacking calls this phenomenon the *looping effect*. By way of example consider the category "queer", whose persistent pejorative use led Anglophone LGBTQIA communities to reclaim the term in the late 1980s. Today, it has largely lost its pejorative connotation and is instead predominantly used to refer to that community in a neutral manner by both its members and the general public alike. The possibility of this kind of interaction with a classificatory scheme means that the concomitant kind concepts can also change their extensions. Haslanger identifies socially constructed kinds as those for which the conditions of kind membership include social properties and relations; in other words, they are kinds whose nature and extension depends on just such a feedback loop (Haslanger 1995, p. 97–99). Natural kinds, by contrast, are "indifferent": they resist interaction (Hacking 1999, p. 104–109).

If Hacking and Haslanger are right, then the looping effect lends credence to the hypothesis that at least some kinds are mind-dependent, at least as far as our collective thoughts are considered. This, in turn, supports the conclusion that MR does not apply to social kinds: what we collectively do (where our agency reflects our thinking) can and does influence the answers to *some* first-order ontological questions – namely, those concerning social kinds. What is more, this suggests an avenue of response to Dodd which marries descriptivism to FM by way of rejecting MR: Dodd's remarks on the objects of our deference are appropriate to natural kinds but they do not apply to social kinds.[14]

14 This possibility was first suggested by S. Davies in his (2003); Dodd comments on it in his (2013, p. 1064–1066). Note that this is not quite the solution for which I argue here, according to which our theories may well be revisionary of our standard linguistic use of a term. On my account, the problem with MR is that it is *too strict*, not that metaphysics imposes no constraints on social ontology.

But I am getting ahead of myself: what *are* these mysterious ways in which thoughts can determine ontology? Ian Hacking and Muhammad Khalidi have argued that the mind-dependence of social kinds follows from the fact that their *existence* depends on human minds: without human intervention, they would not exist in the first place. From this fact, they infer that some of these kinds' *properties* are also mind-dependent. Haslanger, on the other hand, distinguishes between social kinds that are strongly and weakly "pragmatically constructed," where a kind is weakly pragmatically socially constructed if social factors only *partly* determine our use of it, and strongly pragmatically socially constructed if they wholly do so (Haslanger 1995, p. 100).[15] Only the existence of the latter is substantively mind-dependent.

While it is transparently true that the kinds "food," "queer," and "weed" all depend on human minds for their existence, Dodd rightly objects that we are not yet operating at the first order of ontology (2013, p. 1059). Although these *kinds* all depend on human minds for their existence and grouping, the underlying *entities* do not. The observation that some social kinds owe their existence to human minds is nothing more than the observation that we determine which properties count towards kind-membership relative to our own interests, or that some things only come into existence as a result of human thought and action. Our proper focus, instead, is on the *objects* of those kind-terms, namely grains (etc.), people, and the pesky plants themselves. All of *these* objects clearly exist and have their essential properties independently of any minds or thoughts about them.

There are two sets of countervailing considerations here. The first is that the boundary between natural and social kinds is not especially well marked. We already saw in §III that it cannot be drawn simply in terms of what does and does not call for a linguistic division of labour, since some social kinds *do* call for such a division.[16] The second is that this frontier is also populated by liminal kinds which owe their existence to human conceptions and actions but which, once instantiated, resist further interaction with the (bare) world of thought. Haslanger, for example, has suggested that "food" is one such kind, since "opinions about what is appropriate for humans to eat and so about what counts as 'food' have had a huge causal impact on the size, distribution, and behavior of animal

15 Note, however, that Haslanger later qualifies her position: the natural/social contrast "is better viewed as a spectrum from the non-social to the social within the natural" (2012, p. 213). Social kinds are ineluctably embedded in the natural world; the only properly "mind-dependent" things are thoughts and their ilk (see Haslanger 2012, esp. Ch. 6).

16 To this fact, we can add Putnam's later skepticism (1990) that all natural kinds are amenable to definition in terms of their microstructural properties, where those properties reflect the causal powers that entitle them to figure in powerful empirical generalizations.

populations" (Haslanger 1995, p. 104); to this we can add plant populations, too. Stephen Davies has offered the concept "weed" which, although it fails to capture a natural division, does depend on the mind-independent category of "plant" (Davies 2003, p. 6). Finally, both S. Davies and Khalidi have observed that even in chemistry, there may exist some reactionary, short-lived elements or compounds that can only exist in the laboratory, as a result of human conception and intervention (Davies 2003, p. 5; Khalidi 2010, p. 353–354).

These observations go some way towards reinforcing the view (suggested in §III) that natural and social kinds alike exist along a continuum, so that many putatively natural kinds have significant nominal properties not captured by definitions in terms of their microstructural properties, just as some putatively social kinds may also have some kind of real essence which can be captured by a definition. The important thing to notice is just that many social kinds, "art" and art-kinds among them, seem to make essential reference to the social factors that constitute them. If these observations are correct, then we should reject MR – not because deference to metaphysics is inappropriate for social kinds, but because MR stipulates conditions that are too strict: the way we think about *some* questions really does have *some* effect upon the answers to *some* first-order ontological questions.

The upshot for social kinds, I think, is that although our particular *theories* can be wrong, we cannot err about the *subject* of our inquiries. A theory of "gold" which does not allow us to discriminate between *Au* and FeS_2 is not a very good theory of gold, but it *is* a theory *of gold*: the object of our interest is the kind of entity that plays a particular explanatory role, and we treat this role as rigid across possible worlds.[17] My contention is simply that parallel considerations apply to the cases of "art" and art-kinds. Consider music: if a theory of music is so radically revisionary that it has the result that musical works are not actually the kinds of things that can play the kinds of cultural roles conventionally attributed to them, then it is not clear that what we have is a theory of *music* in the first place, rather than a theory of *schmusic*.

Compare Hilary Putnam's remarks on a similar hypothetical scenario in which we discover that the "pencils" on a twinned Earth are actually organisms:

> When we discovered this, we would not say: "some pencils are organisms". We would be far more likely to say: "the things on Twin Earth that pass for pencils aren't really pencils. They're really a species of organism" (Putnam 1975, p. 243).

17 It is worth noting that I think the prospects for giving a definition of "art" in terms of a Ramsified description of its socio-cultural role(s) are bright, although it lies well beyond the scope of the present paper to make such a case here. D. Davies paves the way for this kind of analysis of "art" in his (2009) and (2016), however.

The reason for this is just that our use of "pencil" is rigid: we have in mind not some particular description of pencils, but any and all things that share their nature (whatever it may be) with the things that, in our world, play the pencil-role. It would be more than passing strange to think that we could collectively be wrong about the cultural role that pencils occupy – speaking, at least, of our considered judgements rather than our bare intuitions. The cultural role that pencils occupy is just a matter of the uses to which we put them, after all. A theory of *pencils* – that is, a theory for which "pencil" designates rigidly – whose result was that pencils do not or cannot occupy this role would be absurd, since the result undermines the theory's ability to refer in the first place. Our untutored intuitions provide us with a defeasible starting point for our ontological investigations, nothing more. This is not because our thoughts determine the ontology of social kinds; rather, it is because our practices do the hard work of specifying which kinds of entities it is whose ontology interests us in the first place. Our practices set the limits of our domain of discourse, and offer us what David Davies has called "topic-specific constraints" which allow us to identify the entities that interest us in terms of the roles they play in our practices (Davies 2017). The result is an account that fixes the reference of social kind terms in the same way as natural kind terms.

V

And yet Dodd has argued, on the contrary, that an artifact's ontological nature supervenes on neither its function nor the network of practices in which it is produced and embedded. He tells us, for instance, that "Pencils […] could still be used for writing whether or not they turned out to be enduring entities, spacetime worms, or instantaneous temporal stages" (Dodd 2013, p. 1059). On the endurantist model that supplies the default "folk-theoretic" view, pencils are wholly present at every moment of their existence and have no temporal parts. On the perdurantist model, however, they are not: instead, in addition to their spatial parts they also have a temporal part at every instant in which they exist, and all of these parts can be strung together into a single spacetime worm.[18] Talk of pencils would thus be loose talk for either individual temporal parts of pencils, or the mereological sum that is a wormy pencil. Dodd's point here is just that the functional or cultural roles of pencils can be satisfied by many different ontological pictures; there is no close tie between the two.

18 Lewis (1986) was the first to distinguish between endurance and perdurance, but the distinction has its roots in McTaggart's (1908) characterization of the A- and B-series of time.

That much is true enough, but it misrepresents the character of Dodd's example. The fact that pencils can perform pencil-functions or play pencil-roles on either an endurantist or a perdurantist model might just as easily count as evidence in favour of the view that these ontologies are not substantially different. Both enduring-pencils and perduring-pencils are the same kinds of things – writing implements, artifacts, etc. – and each model leaves most of our ordinary ways of thinking and talking about pencils untouched. It is not as though pencils turned out to be *non-artifactual*, as in Putnam's example (where they are organisms). A different example should serve to make this clear.

In science, a term's ability to play a particular explanatory role is cashed out in terms of the referent's possession of certain causal powers in virtue of which the term can enter into powerful empirical generalizations about natural phenomena (see Putnam 1975, 1990; Davies 2017). So, in order to determine "gold's" reference it is not enough to know that *Au* stands in the right kind of causal-historical relationship to our introduction and use of the term "gold"; we also need to know that *Au* is capable of playing the right kind of explanatory role in our theoretical framework. In this case, the framework at issue is a scientific one. Compare this to the case of "jade", where both jadeite ($NaAlSi_2O_6$) and nephrite ($Ca_2(Mg, Fe)_5Si_8O_{22}(OH)_2$) stand in the right kinds of causal-historical relationship to our introduction and use of the term "jade". Yet these are distinct mineral species, each with a different microstructure and, thus, different causal powers. As a *scientific* term "jade" therefore fails to refer, since we cannot secure the uniqueness of its reference; it cannot play the right kind of explanatory role to figure in our empirical generalizations.[19] This is not to say that "jade" is no kind-term at all, or that there is no jade. In its disjunctive form, "jade" can serve perfectly well for ordinary, though not for scientific, purposes: it may well be a *phenomenal* kind, but it is not a *natural* kind. So long, of course, as natural kinds are the kinds that are supposed to serve an explanatory role *in science*.

With respect to social kinds like "art", the point is that a term's ability to play a particular explanatory role should be cashed out in terms of its capacity to play certain kinds of roles in our practices. So, for example, if a theory of "pencil" has the result that pencils are incapable of serving as writing implements (perhaps because they are organisms, as Putnam suggested, or because they are classes) then it is a bad theory of pencils. In fact, it is not clear that it is a theory of *pencils* at all, as opposed to some other phenomenon which resembles pencils. Similarly, a theory of "gold" that picked out all and only FeS_2 would be a very bad theory of gold; we would be much better off treating it as a theory of *fool's gold* instead.

19 Note, however, that individually $NaAlSi_2O_6$ and $Ca_2(Mg, Fe)_5Si_8O_{22}(OH)_2$ *can* do so.

The point is not that pencils' cultural role determines their ontological nature; as D. Davies puts it, the point, rather, is that our inquiries into specific entities are governed by certain topic-specific constraints. For our scientific inquiries into the nature of natural kinds, those topic-specific constraints are determined by the referent's causal powers; but because a social kind's explanatory value is not exhausted by its microstructural properties, the constraints must instead come from the role the referent plays in our practices.

The requisite proof, here, can be found in Dodd's own pudding. Dodd's ontology of music sets out to answer two questions about instrumental works of pure music: the *categorial* question (to which ontological category do these works belong?) and the *individuation* question (what are the identity conditions of musical works?) (Dodd 2007, p. 1). To answer these kinds of questions, he thinks, we must look to metaphysics. The problem, however, is that "works of instrumental pure music" are not the kinds of entities metaphysics usually concerns itself with: the category is too broad. So we must first determine what *kind* of entity a work of instrumental pure music is, and this means distilling such works to their essential properties, namely, audibility and repeatability (Dodd 2007, p. 3–5, 2013, p. 1053).[20] Dodd thus begins his inquiry by looking for the metaphysical kinds that support audibility and repeatability:

> Plausibly, musical works are in themselves both repeatable and audible: more precisely, such works can have multiple occurrences (e.g. performances), and we can listen to a work by listening to an occurrence of it. […] What kind of entity must a work of music be, given that it can have multiple occurrences? And how, given that an occurrence of a work is distinct from the work itself, is it possible to listen to the work in listening to an occurrence of it? An ontological proposal for works of music should try to answer these questions (Dodd 2013, p. 1053).

Dodd identifies types as the appropriate metaphysical kind; all that remains is to choose the most appropriate kind of type (viz. norm-types – Dodd 2007, p. 3) from the metaphysical menu.

My concern here is not to dispute Dodd's musical ontology, with which I actually agree. Instead, I wish only to observe that Dodd's methodology appears to stand at odds with his commitment to folk-theoretic modesty. Recall that FM stipulates that our common-sense art-ontological views might be substantially mistaken (Dodd 2013, p. 1048). Because they are essential properties of musical works, audibility and repeatability hold the key to answering first-order

20 "Audibility" is meant to capture the fact that we experience musical works by listening to them; "repeatability", the fact that musical works can be instantiated in many different ways, such as through performances.

ontological questions about musical works. But why should we believe that these properties are more essential to music than, say, concreteness (Mag Uidhir 2013), creatability (Levinson 1980, 2011), modal or temporal flexibility (Rohrbaugh 2003), or unrepeatability[21] (Goehr 1992, Ch. 7)? Dodd's answer is that both audibility and repeatability are common-sense views about music that any adequate theory must explain (Dodd 2007, p. 3). But, as such, they are susceptible to radical error. The problem is that the selection of audibility and repeatability as the fundamental properties of musical works proceeds by an appeal to folk intuition, while other plausible candidate properties are dismissed as mere folk intuition. Given the fallibility of folk intuitions, compliance with those intuitions is not an appropriate criterion for pinpointing a kind's essential properties – at least not if we are being folk-theoretically modest.

To be clear, I am not arguing that we are wrong to think that musical works are audible and repeatable; in fact, I think that audibility and repeatability form part of the topic-specific constraints which an ontological investigation of musical works must observe. Dodd's mistake is to single them out from all other candidate topic-specific constraints *pre-theoretically*; in this way, he arbitrarily privileges one set of folk intuitions about musical works over others. As D. Davies has put it,

> There is, *pace* Dodd, no *non sequitur* in the idea that our philosophical interests play a constraining role in ontological inquiry, simply a recognition of the need to particularize any ontological inquiry to the things about which we are inquiring, and a further recognition that what particularizes our ontological inquiry into the nature of artwork-kinds is the explanatory roles that such kinds are intended to serve. The interests, in other words, determine *what it is whose ontological status is at issue*, not *what that ontological status is* (Davies 2017, p. 122).

Appropriate topic-specific constraints must not presuppose the correctness of either our beliefs about artworks or our artistic practices. This means that the first step for an ontological investigation into a social kind like "art" is to consider the *totality* of our beliefs about the kind and our kind-related practices, and to subject these to philosophical scrutiny including, among other things, considerations of which metaphysical constraints are applicable. Some of these beliefs and practices will survive unrevised, while others will not. It is only once we have developed a reflective catalogue of these beliefs and practices that we will be in a position to begin asking and answering first-order ontological questions. We have no privileged epistemic access to the ontology of social kinds; our only

21 All I mean here is that individual compositions might be tailored to specific occasions, and not meant to be repeated for different occasions and audiences.

privilege lies in our ability to determine the proper subject of our inquiries. The result is that, in the end, even a revisionist like Dodd must rely on our collective (and reflective) cultural practices to fix the topic-specific constraints governing ontological investigations of musical works.

VI

Other theories of the reference of artifactual terms are available, of course. A descriptivist might, for instance, maintain that the reference of artifact and social kind terms is fixed by a description of their functional roles implicit in speakers' categorial intentions – a view which at least sounds close to the one I presented above. But I do not think that the prospects for such a view are very promising, for three main reasons. First, because it is often unclear just what a social kind's function *is*; indeed, where art and art-kinds are concerned, different theories of art have posited very different functions, including a characteristic *lack* of function. Moreover, applied to art-kind terms this strategy would seem to throw us back into the implausible clutches of the doctrine of medium specificity,[22] since different art-kinds would have to have different functions (otherwise their referents would be the same). Second, I do not think that this strategy latches on to quite the right phenomenon; a social kind's *functional* role, if it has any, is just one part of what is better characterized as its *cultural* role, which explains the ways in which it enters into cultural explanations. Finally, this kind of descriptivism would still stumble over cases where a social kind's cultural role has changed over time, or is variable across cultures, as is the case with "art" and art-kinds. A descriptive theory would have to class these as distinct terms, or find some means of denying that they showcase any significant changes in, e.g. the function of art or art-kinds.

These same considerations could be seen as evidence of reference shift thanks to speakers' historically-compounded errors, and thus as motivation for a hybrid theory like that suggested by Gareth Evans (1973 and 1982). Just consider, for example, the theory of art for art's sake. In conversation with an English friend who was studying Kant's aesthetics under Schelling's tutelage, Benjamin Constant's imagination was captured by (his misunderstanding of) Kant's notion of disinterested attention, of which he wrote in his journal: "l'art pour l'art et sans

22 The doctrine of medium specificity maintained that each artistic medium is uniquely suited to the communication of different kinds of aesthetic or conceptual content. See, e.g. Lessing (1905 [1766]) and Greenberg (1960).

but; tout but dénature l'art. Mais l'art atteint au but qu'il n'a pas."[23] For Kant, "disinterest" characterizes a distinctive kind of pleasure which accompanies judgements of *beauty*, not a theory of art. Nevertheless, Constant's error spread through Europe by way of the influential works of Victor Cousin and Théophile Gautier, as well as John Ruskin's criticism in England, and eventually came to dominate the way we talk about art's function and value, even today.[24]

This misunderstanding might well be taken to ground a shift in the reference of "art" from, say, objects of craftsmanship intended to perform particular functions to works of Fine Art created for no particular purpose. Because the shift in question would be just one of many, according to the hybrid theory the reference of "art" or of an art-kind term would be fixed by the dominant source of the body of descriptive information which speakers collectively associate with the term. I take it that this sort of strategy is generally friendly to the one I endorsed above. I would simply add that, once again, it is speakers' *reflective* judgements about the descriptive content which they associate with a term (that has a particular explanatory role) that do the work of fixing its reference.

I have argued that although meta-ontological realism sets plausible constraints on natural kinds, its formulation is too strict to apply to social kinds. The social world is not the result of rational deductions from a logical system; it is a slapdash amalgam of objects, practices, and works which have somehow or other captured our individual interests in particular ways. As a result, some of the data points it supplies are bound to be contradictory or otherwise incoherent. This means that we have precious few guarantees that, in deferring to the artworld data, for example, we are deferring to a body of knowledge that is substantially correct. What I have tried to show is that the appropriate response is neither to privilege the artworld data nor to surrender it to metaphysical convenience. We must beware our intuitions and experts who come bearing canons, and look instead to our collective reflective understanding of our artistic practices to supply the data to be explained by our theories. Our ontological investigations cannot pre-judge the issue in favour of one or another of the social world's data points. We must first clear the ground, and determine which of these properly constrain our subject matter. The solution to the social world's arbitrary and contingent nature is thus neither widespread mind-dependence nor wholesale revisionism; it is a pinch of epistemic humility.[25]

---

23 "Art for art's sake and without aim; all aims pervert art. But art achieves the aims it lacks" (my translation; Constant 1952, p. 58; 11 February, 1804).

24 John Wilcox traces the historical transmission of this mistake in his (1953).

25 This research was supported by the Social Sciences and Humanities Research Council of Canada. I am indebted to Emily Carson, David Davies, and Sherri Irvin for their many helpful comments on previous drafts.

Funding: Funder Name: Social Sciences and Humanities Research Council of Canada, Funder Id: http://dx.doi.org/10.13039/501100000155, Grant Number: 756-2017-0843.

Bibliography

Brand, Peg Zeglin (2000): "Glaring Omissions in Traditional Theories of Art." In: Noël Carroll (Ed.): *Theories of Art Today*. Madison, Wisconsin: The University of Wisconsin Press, p. 175–183.

Constant, Benjamin (1952): *Journaux Intimes*. Ed. A. Roulin and C.H. Roth. Paris: Gallimard.

Davies, Stephen (2003): "Essential Distinctions for Art Theorists". In: Stephen Davies and Ananta Ch. Sukla (Eds.): *Art and Essence*. Westport: Praeger, p. 3–16.

Davies, David (2004): *Art as Performance.* Oxford: Blackwell.

Davies, David (2009): "Dodd on the 'Audibility' of Musical Works". In: *British Journal of Aesthetics* 49. No. 2, p. 99–108.

Davies, David (2016): "The Function of Generalization in Art History: Understanding Art across Traditions". In: *Arts and Literary Studies* p. 8–19.

Davies, David (2017): "Descriptivism and its Discontents". In: *Journal of Aesthetics and Art Criticism* 75. No. 2, p. 117–129.

Dodd, Julian (2007): *Works of Music: An Essay in Ontology*. Oxford: Oxford University Press.

Dodd, Julian (2008): "Musical Works: Ontology and Meta-Ontology". In: *Philosophy Compass* 3. No. 6, p. 1113–1134.

Dodd, Julian (2013): "Adventures in the Metaontology of Art: Local Descriptivism, Artefacts and Dreamcatchers". In: *Philosophical Studies* 165. No. 3, p. 1047–1068.

Friend, Stacie (2012): "Fiction as a Genre". In: *Proceedings of the Aristotelian Society* 112. No. 2pt2, p. 179–209.

Goehr, Lydia (1992): *The Imaginary Museum of Musical Works: An Essay in the Philosophy of Music*. Oxford: Oxford University Press.

Greenberg, Clement (1960): "Modernist Painting". *University of Chicago Council on Advanced Studies*. http://cas.uchicago.edu/workshops/wittgenstein/files/2007/10/Greenbergmodpaint.pdf (accessed July 2013).

Guala, Francesco (2016): *Understanding Institutions: The Science and Philosophy of Living Together.* Princeton: Princeton University Press.

Hacking, Ian (1995): "The Looping Effects of Human Kinds". In: Dan Sperber, David Premack and Ann James Premack (Eds.): *Causal Cognition: A Multi-Disciplinary Debate*. Oxford: Oxford University Press, p. 351–383.

Hacking, Ian (1999): *The Social Construction of What?* Cambridge: Harvard University Press.

Haslanger, Sally (1995): "Ontology and Social Construction". In: *Philosophical Topics* 23. No. 2, p. 95–125.

Haslanger, Sally (2012): *Resisting Reality: Social Construction and Social Critique*. New York: Oxford University Press.

Khalidi, Muhammad Ali (2010): "Interactive Kinds". In: *British Journal for the Philosophy of Science* 61, p. 335–360.

Lessing, Gotthold Ephraim (1905 [1766]): *Laocoön: An Essay upon the Limits of Painting and Poetry.* London: George Routledge & Sons, Ltd.

Levinson, Jerrold (1980): "What a Musical Work Is". In: *Journal of Philosophy* 77. No. 1, p. 5–28.
Levinson, Jerrold (2011): "What a Musical Work Is, Again". In: Jerrold Levinson (Ed.): *Music, Art and Metaphysics*. New York: Oxford University Press, p. 215–263.
Lewis, David K. (1986): *On the Plurality of Worlds*. Oxford: Basil Blackwell.
Lopes, Dominic McIver (2014): *Beyond Art*. New York: Oxford University Press.
Mag Uidhir, Christy (2013): *Art & Art-Attempts*. New York: Oxford University Press.
McTaggart, J. M. E (1908): "The Unreality of Time". In: *Mind* 17. No. 68, p. 457–473.
Putnam, Hilary (1975): "The Meaning of 'Meaning'". In: *Mind, Language, and Reality: Philosophical Papers*. Vol. 2. Cambridge: Cambridge University Press, p. 215–271.
Putnam, Hilary (1990): "Is Water Necessarily H_2O?" In: James Conant (Ed.): *Realism With a Human Face*. Cambridge: Harvard University Press, p. 54–79.
Rohrbaugh, Guy (2003): "Artworks as Historical Individuals: Lessons from Photography". In: *European Journal of Philosophy* 11. No. 2, p. 177–205.
Ruben, David-Hillel (1989): "Dismantling Truth". In: Hilary Lawson and Lisa Appignanesi (Eds.): *Realism in the Social Sciences*. London: Weidenfeld, p. 58–75.
Thomasson, Amie L. (1999): *Fiction and Metaphysics*. New York: Cambridge University Press.
Thomasson, Amie L. (2003): "Realism and Human Kinds". In: *Philosophy and Phenomenological Research* 67. No. 3, p. 580–609.
Thomasson, Amie L. (2005): "The Ontology of Art and Knowledge in Aesthetics". In: *Journal of Aesthetics and Art Criticism* 63. No. 3, p. 221–229.
Thomasson, Amie L. (2006): "Debates about the Ontology of Art – What Are We Doing Here?" In: *Philosophy Compass* 1. No. 3, p. 245–255.
Thomasson, Amie L. (2007): *Ordinary Objects*. New York: Oxford University Press.
Walton, Kendall (1970): "Categories of Art". In: *The Philosophical Review* 79, p. 334–367.
Walton, Kendall (2007): "Aesthetics – What?, Why?, and Wherefore?" In: *Journal of Aesthetics and Art Criticism* 65. No. 2, p. 147–162.
Wilcox, John (1953): "The Beginnings of l'art pour l'art". In: *Journal of Aesthetics and Art Criticism* 11. No. 4, p. 360–377.
Xhignesse, Michel-Antoine (2018): "Fake Views – Or Why Concepts Are Bad Guides to Art's Ontology". *British Journal of Aesthetics* 58. No. 2, p. 193–207.
Xhignesse, Michel-Antoine (forthcoming): "What Makes a Kind an *Art*-Kind?" *British Journal of Aesthetics*.