

Natural Kinds, Mind Independence, and Defeasibility

Marc Ereshefsky^{*†}

A standard requirement on natural kinds is that they be mind-independent. However, many kinds in the human and social sciences, even the natural sciences, depend on human thought. This paper suggests that the mind independence requirement on natural kinds be replaced with the requirement that natural kind classifications be defeasible. The defeasibility requirement does not require that natural kinds be mind-independent, so it does not exclude mind-dependent scientific kinds from being natural kinds. Furthermore, the defeasibility requirement captures the idea that natural kind classifications are tools for investigating the empirical world.

^{*} To contact the author, please write to: Department of Philosophy, University of Calgary, Calgary, AB T2N 1N4, Canada; email: ereshefs@ucalgary.ca.

[†] My thanks to Mohammad Ali Khalidi, David Ludwig, Christophe Malaterre, Jay Odenbaugh, Matthew Slater, Ken Waters, and an anonymous reviewer for their helpful comments on earlier drafts of this paper. Audiences at the Philosophy Science Association, University of Toronto, University of Quebec at Montreal, and Concordia University provided helpful feedback. Financial support was provided by the Social Sciences and Humanities Research Council of Canada.

Natural Kinds, Mind Independence, and Defeasibility

1. Introduction

An underlying assumption of this paper is that a proper philosophical account of natural kinds should help us understand the success of natural kind classifications in science. Like many endeavors, a study of natural kinds should study best practices, and arguably those best practices are classificatory practices in science. This idea is far from new. It is line with numerous theories of natural kinds, from Locke (1894) to Boyd (1999). Naturalist approaches to natural kinds aim to learn from science's classificatory practices and offer a fuller understanding of those practices. However, philosophical theories of natural kinds tend not to be naturalistic enough. Consider a standard requirement that philosophers place on natural kinds, namely that natural kinds be independent of human thought (Bird and Tobin 2017, Lowe 2014, Devitt 2005, Psillos 2002, Searle 1995). Here is how Bird and Tobin (2017) articulate this requirement in the introduction to their *Stanford Encyclopedia of Philosophy* entry on natural kinds: "To say a kind is *natural* is to say that it corresponds to a grouping that reflects the structure of the natural world rather than the interests and actions of human beings." As some philosophers have pointed out, this requirement is problematic because numerous successful classifications in science depend on our thoughts and actions (Magnus 2012, Franklin-Hall 2015, Khalidi 2015).

Take, for example, classifications offered by such social and human sciences as psychology, medicine, sociology, economics, cognitive science, and political science. These disciplines posit classifications that depend on the psychological states and behaviors of humans, as well as the behaviors of groups of humans. These classifications help us understand human

and social phenomena, and in some cases, give us vital information for formulating health and social policy. If natural kind classifications are those classifications that help us understand and manipulate the world, then many classifications in the social and human sciences should be considered natural kind classifications even though the phenomena they classify frequently depend on our thoughts and behaviors.

Turning to biology and chemistry, a number of classifications in those disciplines classify entities that depend on us for their existence. Such kinds of entities include genetically modified organisms, artificially selected organisms, and synthetic chemicals. Classifications of such entities help us understand the world and allow us to affect the world in predictable ways. Though we cause such entities to exist, they have features that we study and further manipulate. If natural kind classifications help us understand and manipulate the world, then classifications of chemical and biological kinds that we create should be considered natural kind classifications. Once again, we have reason to resist the requirement that all natural kinds be mind-independent.

What should we do about the requirement that natural kinds be independent of human thought when many scientific classifications classify entities that depend on our thoughts? I propose that we drop the requirement that natural kinds be mind-independent. We should replace that condition with one that requires that natural kind classifications be defeasible. Roughly the idea is that if natural kind classifications are to help us understand and manipulate the world, then they should be vulnerable to disconfirming evidence. The notion of vulnerability here is akin to Popper's (1963) falsifiability requirement on scientific theories. Natural kind classifications should not be true *a priori*, but should be vulnerable to possible disconfirming empirical evidence. In what follows I will spell out the details of the defeasibility requirement on natural kind classifications. One thing to notice at the onset is that the defeasibility criterion

for natural kind classifications is not a metaphysical requirement on natural kinds. It is a requirement concerning which classifications to accept as natural kind classifications.

Consequently, the account of natural kinds offered in this paper is an epistemic one: it focuses on which natural kind classifications help science's epistemic endeavors, rather than the metaphysics of natural kinds.

Here is a breakdown of the paper's contents. The next section discusses various ways that classifications can be mind-dependent. Section 3 introduces the defeasibility requirement on natural kind classifications. Section 4 suggests some ways to refine that requirement. The last section reviews two recent attempts to avoid the mind independence requirement.

2. Varieties of Mind Dependence

There are several ways a classification can be mind-dependent. One way that classifications are mind-dependent that I won't discuss is simply that all classifications posited by scientists depend on human thought and action. Classifications are representations we construct, so they obviously depend on us. In addition, there is the fact that we choose which phenomena to classify. Among possible classifications of the world, we choose which classifications to construct according to our interests and preferences. These sorts of mind dependence are not the sorts of dependence that promoters of the mind independence requirement worry about. They worry about classifications that are mind-dependent when such dependence interferes with the objective study of natural kinds. Often that objectivity is aligned with realism. In a more metaphysically modest vein, I suggest that we should worry about mind-dependent classifications when such dependence does not promote the epistemic aims of scientific classification. That is, we should worry about mind dependence when such dependence causes classifications to be useless tools

for investigating the empirical world. Accordingly, let's sort the mind dependence of classifications into two categories: mind-dependent classifications that help us investigate the empirical world, and mind-independent classifications that do not help us investigate the empirical world.

Sociological classifications of socially caused racial kinds are examples of classifications whose categories are mind-dependent, yet those classifications help us investigate the world.

Michael Root offers the following introduction to such kinds.

Race is like marital status: no one would be married or single had we not invented matrimony; however, given that we did, we now divide ourselves along discernible boundaries, into categories like... "single" and "divorced" and treat each other differently depending on which of these categories we belong to. So too with race: we assign each other a race and treat each other differently depending on that race. As a result, epidemiologists can discover that rates of mortality or morbidity are different for one race than another... . In other words, race may be a biologically salient category even though there are no biological races, and race can mark the risk of a biological condition like diabetes or heart disease even though race is not itself a biological condition but a social status. (Root 2003, 1175)

Consider the social category Black American studied by sociologists. Numerous properties are associated with people in this category. One set of properties consists of racist beliefs held by many Americans about Black Americans, such as the belief that Black Americans have lower intelligence than White Americans, or work less hard. Such beliefs affect Black Americans, for instance, where they go to school, their income levels, and their health and mortality rates. The

relations between such beliefs and the lives of Black Americans are testable. A sociologist can test whether certain beliefs held by Americans are manifested in social practices, and whether those practices affect Black Americans' economic, disease, and mortality rates (Root 2000, Mallon and Kelly 2012). Furthermore, those interested in social policy use sociological classifications of race to ameliorate the negative effects of racism. The category Black Americans depends on human thought and is mind-dependent, but whether those thoughts affect Black Americans can be tested. That testability, I suggest, makes the classification of the social category Black American a useful tool for investigating and affecting the empirical world.

Let us turn to an example of a mind-dependent classification whose mind dependence prevents it from being a tool for investigating the empirical world. Here I focus on infallibilist social kinds (Guala 2015). Membership conditions in such kinds entirely depend on how we define those kinds rather than any investigation of the world. Searle (1995), Thomasson (2003), and Taylor (1971) hold infallibilist views of social kinds. As Thomasson writes, “[I]n the case of institutional kinds those principles we accept regarding sufficient conditions for the existence of these entities must be true” (2003, 590). Khalidi (2015) offers the social kind permanent resident as an example of an infallibilist kind. Khalidi calls such kinds ‘conventional kinds.’ The conditions for membership in the kind permanent resident are legislated by the appropriate governmental bodies. Khalidi suggests that a government could legislate that all permanent residents must be able to swim one hundred meters. Once that rule is set down, being classified a permanent resident is linked to the ability to swim one hundred meters. I want to suggest that classifications of such social kinds are poor candidates for natural kind classifications because of their infallibilist and conventional nature. For such kinds, we legislate which properties are associated with a kind rather than conduct an empirical investigation. As Khalidi (2015, 107)

writes, “no sociologist would be awarded a research grant to investigate the link between being a permanent resident and being a proficient underwater swimmer.” The reason this classification is not a natural kind classification is not the odd nature of the requirement placed on being a permanent resident, but the legislative and conventional nature of how that requirement is set.

In the first example, the case of sociological classifications of race, such classifications are mind-dependent, but not in a way that adversely affects their testability. In the second case, the legislated kind permanent resident, we see how mind dependence may render a classification untestable. The first classification helps us investigate the empirical world, and the second classification does not. The challenge, then, is to distinguish classifications that are mind-dependent yet help us investigate the world from classifications that are mind-dependent and do not help us investigate the world. The next section suggests a way to capture that distinction.

3. Defeasibility and Natural Kind Classifications

In what follows the notion of defeasibility is borrowed from epistemology to offer an alternative to the mind independence requirement on natural kinds. In epistemology, a belief is defeasible just in case it can be undermined by evidence. The evidence that undermines a belief is called a ‘defeater.’ There are different types of defeaters: undercutting versus rebutting defeaters, and external versus internal defeaters (Kelly 2014). For a defeasibility account of natural kind classifications, we don’t need to go through all four variations of defeaters. We can simply focus on rebutting external defeaters. Here’s an example of a rebutting external defeater. Suppose I believe that my name is Bert. My mom, however, tells me that I am mistaken and that my name is really Sue. She recounts how her and my dad were inspired to name me Sue by Johnny Cash’s song “A Boy Named Sue.” My mom’s testimony is a rebutting external defeater for my belief

that my name is Bert. Her testimony is rebutting in that it is evidence against my belief that my name is Bert. It is external because what she asserts is not among my pre-existing beliefs.

Before applying the notion of defeater to natural kind classifications let me mention a background assumption I will be making. The assumption is that natural kind classifications assert certain linkages among the properties of the members of a kind. For example, the category gold in the Periodic Table links having the atomic number 79 with conducting electricity. Similarly, the sociological category Black American links racist beliefs with the morbidity and mortality rates of Black Americans. With that background assumption in hand, let's apply the notion of defeasibility to natural kind classifications. A natural kind classification is defeasible if there could be external rebutting defeaters for the relations posited by that classification. For instance, the classification for gold is defeasible if it is possible that there is evidence indicating that entities with the atomic number 79 do not conduct electricity. Similarly, sociological classifications of race are defeasible if it is possible that there is evidence indicating that the socio-economic conditions of people of a certain race are not linked to social discrimination. In brief, the defeasibility requirement on natural kind classifications asserts that a classification is a natural kind classification only if the relations it posits among the properties of the members of a kind are defeasible.

It is important to note that the defeasibility requirement does not require that actual defeaters exist. The Periodic Table's classification of gold is defeasible not because there are actual defeaters of the relations it posits. It is defeasible because there could be such defeaters. The 'could' here means that the existence of a defeater is an empirical possibility. It means that the relations asserted by the classification are not true *a priori*; that is, they are not true by definition, math, or logic. Recall the motivation for placing a defeasibility requirement on

natural kind classifications. Namely, the assumption that natural kind classifications are those classifications that help us investigate the world. That in turn requires that such classifications should be empirically testable. Accordingly, the relations asserted by a classification should not be true by definition and could be defeated by empirical evidence.

With the defeasibility requirement for natural kind classifications in hand, let's return to the examples of mind-dependent categories discussed in Section 2. As we saw, some mind-dependent classifications help us investigate the empirical world, and some do not. The defeasibility requirement nicely distinguishes when a mind-dependent classification is useful for investigating the empirical world and when it is not.

Recall the social category Black American posited by sociologists. Social scientists posit that category to study how racist beliefs affect Black Americans' disease and mortality rates. Whether such racist beliefs affect the lives of Black Americans can be empirically tested, and we could be wrong that such beliefs affect the lives of Black Americans. Indeed, sociologists adjust or reject a classification of this category depending on the evidence obtained. So the mind dependence of the category Black American does not impede its ability to be a useful category for investigating the empirical world. The defeasibility requirement on natural kind classifications captures the testable nature of this category. The sociologist's category Black American passes the defeasibility requirement because the links posited among racist beliefs and the lives of Black Americans are defeasible.

Let's turn to a mind-dependent classification where that mind dependence impedes the classification from being a useful tool for investigating the empirical world. Recall Khalidi's example of the category permanent resident. Membership in that category is legislated by a government. Following Khalidi's example, suppose a government legislates that all permanent

residents of a country must have the ability to swim. That requirement is brought about by the legislators' decision, not by any discovery about the empirical world. As a result, that legislative decision is not open to empirical disconfirmation. The defeasibility requirement on natural kind classifications does a good job ruling out this construal of permanent resident from being a natural kind classification. The posited link between being a permanent resident and having the ability to swim is infallible and thus non-defeasible.

Stepping back from these examples, recall the problem we started with at the beginning of this paper. Traditional accounts of kinds require that natural kinds be mind-independent. However, that requirement excludes many scientific classifications in the human and natural sciences that help us understand and manipulate the world. The defeasibility account does not rule out such kinds. Furthermore, it distinguishes when a mind-dependent classification is useful for investigating the empirical world and when it is not.

4. Refinements

One might have concerns with the defeasibility requirement. In fact, when I have presented the requirement to various audiences several worries have surfaced. This section of the paper recounts those worries and suggests how the requirement should be refined to address them. The result is a fuller articulation of the defeasibility requirement.

One worry is that some putative natural kind classifications posit both defeasible and non-defeasible relations among the properties of the members of a kind. Consider the Canadian hundred-dollar bill. The Bank of Canada decides (legislates) that a certain configuration of plastic (Canadian bills are made of plastic not paper) should be Canadian currency. There's no empirical investigation involved in this action. At the same time, there should be certain social

processes and infrastructure in place so tokens of money have monetary value (Guala 2015). For example, only the Bank of Canada should make Canadian hundred-dollar bills, otherwise such bills lose their value as Canadian money. In addition, the Canadian government should accept Canadian hundred-dollar bills as payment for Canadian taxes, thereby creating a demand for that currency. The category Canadian hundred-dollar bill, in other words, is a mixed non-defeasible and defeasible category. On the one hand, a Canadian hundred-dollar bill's being a certain configuration of plastic is a non-defeasible property since it is legislated. On the other hand, that the Canadian government takes Canadian hundred-dollar bills as form of tax payment is defeasible since we could be wrong about that –perhaps that government no longer takes cash payment.

Does a classification's attributing both defeasible and non-defeasible property relations to the members of a category make it a bad candidate for being a natural kind classification? I will suggest that such classifications are still good candidates for being natural kind classifications. Such mixed defeasible and non-defeasible classifications are important for understanding the world. For instance, economists study the empirical aspects of monetary categories, such as the popularity of Canadian hundred-dollar bills and what percentage of them is in circulation. In general, mixed defeasible and non-defeasible classifications are frequently posited by the social and human sciences, and they are used to investigate the empirical world.

The example of money shows that the defeasibility requirement on natural kind classifications needs to be adjusted to allow mixed defeasible and non-defeasible classifications. As it stands now, the requirement says that the property relations among the members of a kind posited by a classification must be defeasible. Let's adjust that requirement so that it says that at least some of the property relations among the members posited by a classification must be

defeasible. This adjustment makes sense because sometimes categories originate by our positing non-defeasible property relations among the members of a category. Then social scientists conduct empirical investigations of other properties relations among members of that category. We saw this in the case of Canadian currency. The category permanent resident could share a similar pattern. Permanent residency for a country is first legislated, but suppose later social scientists empirically investigate other properties such residents might have, such as their household income.¹ In brief, we often posit categories whose members have non-defeasible and defeasible property relations, and it is the defeasible relations that make such categories candidates for natural kinds.

Another worry is that some classifications in the history of science are defeasible and defeated. For example, the classification of phlogiston is defeasible and defeated. Yet no contemporary chemist would say that the category phlogiston is a natural kind. As it stands, the defeasibility requirement allows a classification to be a natural kind classification when there are reasons to believe that the cited category is not a natural kind. At this juncture, it should be noted that the defeasibility requirement is a necessary condition not a sufficient condition for whether a classification is a natural kind classification. So as a necessary condition, the phlogiston example is not a counterexample to the defeasibility requirement. Nevertheless, one might wonder what conditions should be added to the defeasibility requirement to attain a more sufficient condition for when a classification is a natural kind classification. One condition is

¹ One might worry that here I say permanent resident could be a defeasible category, whereas earlier I said it was not a defeasible category. Whether the category permanent resident is a defeasible category is time indexed. If a classification of the category posits no defeasible property relations among its members, it is not a defeasible category and not a natural kind. But later, if defeasible property relations are posited, it can become a natural kind. The example of permanent resident shows that many categories that are not natural kinds have the potential to become natural kinds.

that a natural kind classification should not only be defeasible but also well tested and not defeated. (I leave ‘well tested’ as a placeholder to be filled in by the relevant scientists and philosophers.) In other words, the degree to which we should believe that a classification is a natural kind classification should turn on how well it has been tested and not defeated. So minimally, a classification needs to be defeasible to be a natural kind classification. After that, the degree to which we should believe that a classification is a natural depends on how well it has been tested and not defeated. This seems like a reasonable further requirement on natural kind classifications. It rules out outdated classifications as natural kind classifications, such as the classification of phlogiston. Consequently, the revised defeasibility requirement on natural kind classifications should read as follows. A classification is a natural kind classification so long as at least some of the relations it posits among the properties of the members of a kind are defeasible, well tested, and not defeated.

A third worry is that the defeasibility requirement is satisfied by classifications of categories that are obviously not natural kinds. Consider the category tall person. The classification with the category tall person could posit defeasible relations. For instance, it might assert that tall people on average are worse at limbo than short people, or that tall people on average are better at basketball. These assertions can be tested and defeated. Or consider the category pen. To be a pen, an object must have some weight. So there is the testable relation between being a pen and having weight. From these examples we see that the defeasibility requirement might be too lenient in that it allows that too many classifications are natural kind classifications. It allows that the categories tall people and pen are natural kinds, but that sounds like something an account of natural kind classifications should deny.

I want to suggest that the waters here are muddier than that. If the category tall people was investigated by science, then it would be a good candidate for being a natural kind. The same goes for the category pen. For example, suppose NASA was investigating which sorts of pens should be sent into space and that requires conducting tests on the weights of various pens. In both hypothetical cases, seemingly non-natural kind classifications are employed to investigate the empirical world. Recall a stated aim of this paper: to align our philosophical conception of natural kind classifications with classifications that help us investigate the world. If categories of entities don't intuitively seem to be natural kind categories but are nevertheless used in investigating the empirical world, then we should consider them natural kind categories. Assuming that natural kind classifications are tools for investigating the world, the question should be whether a classification is used for investigating the empirical world not whether a category in a classification fits our intuitions of what is a natural kind.

In light of this last worry, one could refine the defeasibility requirement on natural kind classification as follows. A classification is a natural kind classification so long as at least some of the relations it posits among the properties of the members of a kind are defeasible, empirically investigated, well tested, and not defeated. On further reflection, one will see that the 'empirically investigated' clause is unnecessary since if a category is well tested then it is empirically investigated.

5. Recent Attempts to Replace the Mind Independence Requirement

In this section I review two recent accounts of natural kinds that avoid the mind independence requirement on natural kinds. I also suggest that the defeasibility approach to natural kind classifications is preferable. Khalidi (2013, 2015) has written extensively on the mind

independence requirement on natural kinds. He suggests that we replace it by requiring that the links among the properties of the members of a natural kind be causal rather than conventionally established. The infallibilist social kinds discussed early –kinds where the properties of a category are legislated rather than discovered – are conventionally established and not natural kinds according to Khalidi. Given the similarity between our two approaches, one might wonder why do I suggest a defeasibility requirement on natural kinds rather than adopt Khalidi’s causal requirement?

Khalidi’s proposal has much to recommend it, but it places too much emphasis on causality. Many epistemically fruitful kinds in science posit correlations among the properties of a kind rather than causal relations. Such kinds can be found in microbiology (Ereshefsky and Reydon 2015), astrophysics (Ruphy 2010), and quantum physics (van Fraassen 1980). So Khalidi’s account leaves out significant areas of scientific classification. An account of natural kinds should be consistent with the various epistemic aims scientists have for positing categories, including the pursuit of non-causal correlative kinds. Defeasibility captures Khalidi’s motivation to reject infallibilist kinds while at the same time not ruling out scientific classifications that focus on correlation rather than causation.

Franklin-Hall (2015) also offers an account of natural kinds that avoids the mind dependence requirement. She locates natural kinds at the intersection of investigations conducted by different epistemic agents. In particular, Franklin-Hall writes “natural kinds are groupings that match those categories that well serve actual inquirers along with (what I call) ‘neighboring agents’ – those different somewhat from actual inquirers in their particular epistemic aims and cognitive capacities” (2015, 940). A virtue of Franklin-Hall’s account is that highlights the role inter-subjectivity plays in identifying natural kinds. However, the sort of

inter-subjectivity that Franklin-Hall requires is too abstract and other worldly for constructing an account of natural kind classifications based on actual classificatory practice in science. The sort of inter-subjectivity her account turns on employs “neighboring agents,” where those neighboring agents are, as seen in the quote above, not “actual inquirers.” Hers is not an account of what natural kind classifications are in actual scientific practice, but an in principle account of what natural kinds should be. An aim of this paper is to provide an account of natural kinds based on actual classificatory practice. Her account fails to do that. Moreover, it is a non-operational account: how could we check that non-actual inquirers pick out the same kinds as actual inquirers? With the defeasibility account we can check if a classification is defeasible, well tested, and not defeated.

The discussion of Franklin-Hall’s account of natural kinds brings to the fore a meta-issue concerning how philosophers should study natural kinds. A common view in philosophy is that an account of natural kinds should inform us about the fundamental nature of natural kinds. Franklin-Hall’s (2015) account falls into that camp, as do a number of philosophical theories of natural kinds (for example, Lowe 2006, Hawley and Bird 2011). The account offered in this paper, however, does not attempt to give an account of the fundamental nature of natural kinds. Instead, it targets the question of when we should think a classification is a natural kind classification. The account offered here is agnostic about the fundamental nature of natural kinds. Furthermore, it focuses on why scientists are successful in giving natural kind classifications, where ‘natural kind classifications’ means those classifications that help us understand, investigate, and manipulate the empirical world.²

² To avoid confusion, let me point out that the defeasibility requirement does not simply assert that if a category is useful for science it is a natural kind. Putative natural kinds are categories that are the *subject* of scientific investigation, not merely tools in scientific investigation. Pens

Why adopt such a practice-focused account of natural kinds rather than a metaphysical one that attempts to understand the fundamental nature of natural kinds? As mentioned at the start of this paper a study of natural kinds should study best practices, and when it comes to natural kinds those best practices are classificatory practices in science. There are at least two virtues of a practice-focused philosophical study of natural kinds. One is that it allows us to carefully scrutinize philosophical theories of natural kinds. For example, the focus on practice has revealed that natural kind essentialism is problematic in various areas of science (Dupré 1993). Another virtue of a practice-based study of natural kinds is that it focuses on studying why classificatory practices in science have been epistemically successful. This promotes a major aim of philosophy of science, namely to understand the success of science. Needless to say, the account of natural kind classifications offered in this paper subscribes to the practice turn in the philosophy of science (Soler et al. 2014, Kendig 2016). The practice turn endeavors to understand the epistemic success of scientific practices rather than offer a metaphysics of science. Similarly, the defeasibility approach to natural kind classifications aims to understand the epistemic success of classificatory practice in science rather than offer a metaphysics of science.

are useful for scientists but they are not putative natural kinds unless they are the subject of scientific investigation.

References

- Bird, Alexander and Emma Tobin. 2017. "Natural Kinds." In *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), ed. Edward N. Zalta, URL=<<http://plato.stanford.edu/archives/spr2017/entries/natural-kinds/>>.
- Boyd, Richard 1999. "Homeostasis, Species, and Higher Taxa." In *Species: New Interdisciplinary Essays*, ed. Robert A. Wilson, 141-85. Cambridge, MA: MIT Press.
- Devitt, Michael. 2005. "Scientific Realism." In *The Oxford Handbook of Contemporary Philosophy*, ed. Frank Jackson and Michael Smith, 766-90. Oxford, UK: Oxford University Press.
- Dupré, John. 1993. *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.
- Ereshefsky, Marc and Thomas Reydon. 2015. "Scientific Kinds." *Philosophical Studies* 172: 969-86.
- Franklin-Hall, Laura. 2015. "Natural kinds as Categorical Bottlenecks." *Philosophical Studies* 172: 925-48.
- Guala, Francesco. 2015. "Philosophy of the Social Sciences: Naturalism and Anti-Naturalism." In *The Oxford Handbook of Philosophy of Science*. Oxford, UK: Oxford University Press.
- Hawley, Katherine and Alexander Bird. 2011. "What are Natural Kinds?" *Philosophical Perspectives* 25 (1): 205-21.
- Kelly, Thomas. 2014. "Evidence." *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), ed. Edward N. Zalta. URL = <<https://plato.stanford.edu/archives/win2016/entries/evidence/>>.
- Kendig, Catherine, ed. 2016. *Natural Kinds and Classification in Scientific Practice*. New York, New York: Routledge.
- Khalidi, Muhammad Ali. 2013. *Natural Categories and Human Kinds*. Cambridge, MA: Cambridge University Press.
- 2015 "Three Kinds of Social Kinds." *Philosophy and Phenomenological Research* 90:96-112
- Locke, John. 1894. *An Essay Concerning Human Understanding*. London: Routledge and Sons.
- Lowe, Edward Jonathan. 2006. *The Four-Category Ontology: A Metaphysical Foundation for Natural Science*. Oxford, UK: Oxford University Press.

- 2014. “How Real Are Artefacts and Artefact Kinds?” In *Artefact Kinds: Ontology and the Human-Made World*, ed. Maarten Franssen, Peter Kroes, Thomas Reydon, and Pieter Vermaas, 17-26. Cham, Switzerland: Springer International Publishing.
- Magnus, P.D. 2012. *From Planets to Mallards: Scientific Enquiry and Natural Kinds*. Basingstoke: Palgrave Macmillan.
- Mallon, Ron and Daniel Kelly. 2012. “Making Race Out of Nothing: Psychologically Constrained Social Roles.” *The Oxford Handbook of Philosophy of Social Science*, ed. Harold Kincaid (ed.). Oxford UK: Oxford University Press.
- Popper, Karl. 1963. *Conjectures and Refutations*. London: Routledge and Keagan Paul.
- Psillos, Stathis. 2002. *Causation and Explanation*. Montreal: McGill-Queen’s University Press.
- Root, Michael. 2000. “How we Divide the World.” *Philosophy of Science*. 67:628-639.
- 2003. “The Use of Race in Medicine as a Proxy for Genetic Differences.” *Philosophy of Science*. 70:1173-1183.
- Ruphy, Stephanie. 2010. “Are Stellar Kinds Natural Kinds? A Challenging Newcomer in the Monism/Pluralism and Realism/Antirealism Debates.” *Philosophy of Science* 77 (5): 1109-20.
- Searle, John. 1995. *The Construction of Social Reality*. New York, NY: Free Press Publishing Co.
- Soler, Lena, Sjoerd Zwart, Michael Lunch, and Vincent Israel-Jost, eds. 2014. *Science After the Practice Turn in the Philosophy, History, and Social Studies of Science*. New York, NY: Routledge.
- Taylor, Charles. 1971. “Interpretations and the Sciences of Man.” *Review of Metaphysics* 25:3-51.
- Thomasson, A. 2003. “Realism and Human Kinds.” *Philosophy and Phenomenological Research* 68:580-609.
- Van Fraassen, Bas 1980. *The Scientific Image*. Oxford: Oxford University Press.