

Synthese, 2016.

NOMINALISM, REALISM AND OBJECTIVITY

Catherine Z. Elgin

Naïve realism is probably the metaphysical default. Many people seem to assume without much thought that the world is the way it is regardless of what anyone thinks, and that sentences are true just in case they describe the way the mind-independent world is. Naïve realism does not discriminate. It applies across the board. A little reflection reveals that it cannot be right. The truth values of some sentences pretty plainly depend on what some people think. Mental predicates are one source of difficulty. The truth value of the sentence ‘someone believes that aardvarks snore’ cannot be independent of what anyone thinks. If no one believes that aardvarks snore then the sentence ‘someone believes that aardvarks snore’ is false. Humility is also a problem. If I think that I am humble, then I am not. My being humble thus requires my not believing that I am. Then there is the problem of fashion. Manifestly I can be wrong about whether an outfit is fashionable. So can anyone else. But it is highly implausible that miniskirts or baseball caps could be fashionable if no one thought they were. Predicates like ‘fashionable’ seem depend on what people think, even if not on what any particular person thinks. Some philosophers believe that mind-dependence is characteristic of evaluative predicates in general. No more than fashionableness would properties like goodness and beauty be independent of what people think. Nor is it clear what roles such properties could play in the causal structure of the world. So, such philosophers contend, sentences like ‘The *Mona Lisa* is beautiful’ or ‘Giving to Oxfam is good’ should not be construed as describing the way the mind-independent world is.

Evidently, a retreat from naïve realism is necessary. The question is where we should stop retreating. Many draw the line at scientific realism. Conceding the points about fashion, ethics, aesthetics, and intentionality, scientific realists insist that natural science attempts to discover the way the mind-independent portion of the world is, and when it is successful, its success consists in its accurately characterizing that portion of the world. I believe that scientific realism is more problematic, and that an alternative position, which I call *constructive nominalism*, is more plausible than we might initially suppose. But the plausibility of constructive nominalism may be undermined by concerns about objectivity. A major goal of this paper is to show that objectivity is compatible with constructive nominalism.

Scientific theories go beyond the evidence they are based on. They may adduce substructures, superstructures, dispositions, causes or grounds to explain what occurs. Their explanations thus yield entailments and implicatures about other actual and possible cases, and about unexamined aspects of the phenomena they pertain to. Theories embed evidence in cognitive contexts that enable us to leverage our understanding of a domain. This is why they are valuable. But what makes them valuable also makes them vulnerable. Insofar as they go beyond the evidence, theories take risks. An explanation that accommodates all available evidence could be wrong. New evidence might turn up that indicates that, for example, estrogen prevents heart attacks after all. Even if no new evidence happens to turn up, a theory could still get things wrong. Although there is no evidence of them, perhaps there are additional, very rare blood types. Still, we may think, our epistemic vulnerability is due to the dearth of evidence. Were it not for practical impediments, we could get enough evidence to render our theories invulnerable.

Let us grant that *ceteris paribus*, more evidence is better than less. So imagine a best case scenario. Suppose that our theory accommodates all the evidence, past, present and future, and that there are no practical impediments to gathering any evidence we want. Suppose further that our theory displays the full panoply of theoretical virtues. It is simple, predictive, fruitful, robust, and so on. Even then we have a problem. For another theory might be equally good in all these respects. Because theories are underdetermined by evidence, more than one theory can accommodate all the evidence, however plentiful the evidence is. And although the requirement that they satisfy theoretical constraints enables us to dismiss *ad hoc* or *rococo* alternatives, there is no reason to think that a unique solution results. Multiple empirically equivalent accounts are apt to display an equal balance of theoretical virtues.

Such underdetermination should probably distress the realist. Realism is the view that things are the way they are, regardless of what anyone believes about the way they are, and that the true theory is the one that describes the way things are. But it follows from underdetermination that we cannot tell how things are. At best, we might conclude that a disjunction is true: 'Theory A or Theory B or Theory C or . . .' is true, but in principle, we will never be able to tell which one. This is not nothing, but it is disappointing. However good our particle theory, we will never be in a position to know, for example, that the world consists entirely of particles, since an equally good alternative says that it consists entirely of waves. This does not refute realism. It simply underscores an unfortunate epistemological consequence. Realism leads to skepticism, not because of the inevitable paucity of evidence, but because of the availability of alternative, equally good (although perhaps unformulated) theories.

The situation is this: Theory A (a particle theory) is an ideal theory. It answers to all the evidence – indeed to all possible evidence – past, present and future. It displays a maximal constellation of theoretical virtues as well. If Theory A were the only theory, we would accept it in a minute. It has everything going for it and nothing going against it – except for pesky Theory B (a wave theory), which is equally good in all relevant respects. If we conclude that we can accept the disjunction, but cannot accept either disjunct, we concede that scientific inquiry, however successful, cannot yield knowledge of the way things are. At best it yields knowledge of ways, for all we can tell, things might be. We cannot simply conclude that any particular disjunct is true, since we have no basis for rejecting the other. What should we do?

Bas van Fraassen thinks we should opt for constructive empiricism (1980). Which of the theories is true is determinate. But it is something that we can never know. Nor does it matter. For all theoretical, empirical, and practical purposes, any empirical equivalent is as good as the true theory. Pick one and work with it. It makes no difference if you pick a false theory so long as that theory is empirically equivalent to the truth. Science will proceed unscathed. But there is no reason to believe that science, even at the limit of inquiry, yields truth. This too is disappointing. We were hoping for more from science than a sophisticated device for saving the phenomena.

Another alternative is this: Rather than thinking that the adequacy of the two theories gives us grounds for accepting the disjunction, or flipping a coin to decide between them, why not conclude that it gives us grounds for accepting the conjunction? Following Israel Scheffler, let's call this view plurealism (2008). If we have sufficiently good reason to believe that A is true, and sufficiently good reason to believe that B is true, why not say that we have sufficiently good reason to believe that A&B is true?

Occam's razor evidently cuts off this option. We ought not multiply entities beyond necessity. But to endorse the truth of A&B seems to do just that. If the world is comprised at bottom of particles, there is no need to admit waves into our ontology; and if it is comprised of waves, there is no need to admit particles. Moreover, in a world comprised entirely of waves, there is no room for particles; and in a world comprised of entirely particles there is no room for waves.

Realism leads to skepticism; constructive empiricism settles for empirical adequacy; pluralism engenders severe overpopulation. All are disappointing if we seek through science (or any sort of systematic inquiry) to discover the way things are. But from the fact that a conclusion is disappointing, we should not infer that it is false. So perhaps the alternatives really are this bleak. Maybe even ideal science cannot disclose the way the world is.

Nevertheless, here is something weird about this predicament. Our problem is not a dearth of resources; it is an embarrassment of riches. If we had only one of the theories, we would have no reservations about considering it true and concluding that the world consists of whatever it says the world consists of. After all, that theory satisfies all the relevant theoretical and evidential requirements we can devise. What better reason could we have for thinking it is true? What more could we want? The problem is that a second, equally good theory is too much of a good thing.

This suggests that we are misconceiving the problem in such a way that success looks like failure. Rather than thinking of truth as the *ur*-value of a good theory – that which makes a good theory good – perhaps we should think of truth as constituted *within* a good theory, where the goodness of the theory is determined by other factors. How does such a revision reconceptualize theories?

Realists, and many ordinary people, think that a scientific theory is, or should be, a mirror of nature. It should reflect without distortion the way things actually are. Obviously, a theory must do more than that. To afford an understanding of a range of phenomena, it needs to find or make order, discriminate between relevant and irrelevant factors, point to some underlying or overarching regularities that the phenomena display, and so on. No mere mirror is so helpful. But however we spell out these details, ‘without distortion’ remains crucial. Theories should, we may think, yield neutral descriptions of the way the phenomena are and would be, whether they were so described or not. It’s a nice idea, and it explains why partiality and bias in theories are defects. But it may be too simple.

Perhaps rather than reflecting, theories refract. In that case, although no theory describes the phenomena as they are in themselves, each describes the phenomena as seen through a particular refractive lens. A characteristic of refraction is that the path of a light wave depends on the medium of refraction. This does not make the path less than objective. But it does mean that if we want to understand what happens to a light wave when it goes from one medium to another, we need to know the refractive indices of the media. Similarly I suggest, rather than thinking that we can just look through a theory and see the mind-independent world, our encounters with things are always mediated and modulated by cognitive commitments. The result is an understanding of the-world-as-modulated-through-a-particular-theory. Truth then is defined within a theoretical framework rather than outside of it. But this does not make the truth of a sentence less than objective. It simply means that a claim to truth contains a (perhaps tacit) reference to the theoretical framework in terms of which truth is claimed. Given that we are in a theoretical environment that recognizes waves, the sentence ‘under condition c , wave

form w occurs' is true, or it is false. The theoretical framework sets the conditions on truth; the phenomena comply or they do not. So the truth of p is independent of our belief that p is true. The world as described by the theory remains independent of our attitudes. But the kinds in terms of which we classify things, hence the sentences that are candidates for truth, depend on our cognitive constructs. This, I think, is part of what Hilary Putnam was getting at in 'Realism and Reason' when he characterized his position as *internal realism* (1978).

Realists might argue that what I am saying is trivial. Of course, if we had never invented the word 'rock', the sentence 'New England's soil is full of rocks' would not be true. But since all we are doing in introducing a word like 'rock' is making up a label for what is out there anyway, our contribution is negligible. What I have called 'the-world-as-modulated-through-theory-T', they would maintain, is simply the world as described in the language in which T is expressed. Their picture is this: a scientist looks into her microscope, sees a cellular mass and says to herself, 'Lo! A hitherto unobserved object! I need a word for it. I will call it *a mitochondrion*.' In that case, the label is just a tag for something that was there to be seen. The significant thing she has done is detect a constituent of the cell, which was there all along, even before she or anyone else noticed it. But this makes things too simple. Individuation and classification involve more than pasting on labels. To individuate a material object is to demarcate its spatial and temporal boundaries, and to specify the changes it can undergo while remaining the same thing. To classify objects is to set criteria for belonging to the same kind. How can items of this kind differ from one another and still count as the same kind of thing? Which things count as mitochondria and why?

A truly hardcore realist might maintain that every difference makes a difference to the identity of an object. The world does not discriminate between important and unimportant properties. So if the world determines identity, every difference between objects makes them different sorts of things, and every change an object undergoes makes it a different thing. Such a metaphysics might accurately reflect the way the world is, but it would be utterly unwieldy. We could not fathom it. We could not understand things in terms of it.

To determine what there is, we need to settle which differences matter, thereby marking out individuals and kinds. The scientist who discovers mitochondria does more than notice and name a blob on her slide. She explicitly or implicitly identifies that blob with some intracellular bodies and differentiates it from others. Of course this is not typically done in a flash, and may not be done by a single scientist. Maybe someone does look into a microscope, notice a previously overlooked blob and give it a name. But further work needs to be done. It has to be differentiated from other intracellular bodies. It has to be found in other cells. It has to be characterized in such a way that scientists can tell whether something is a mitochondrion; can tell whether, having undergone various changes, it is the same one; can tell when it ceases to exist as a mitochondrion; and so forth. This is not a matter of merely coining a term or labeling an antecedently recognized entity; it is a matter of (at least partially) reconceiving the interior of the cell.

The issue that divides realists and nominalists is whether – or perhaps to what extent – such scientific work is creation or discovery. Do the scientists draw the lines, or discover where nature draws them? Realists maintain that nature draws the lines. The scientists and other investigators simply find what is there to be found. The problem, though, is that there are too many things there to be found.

Mereology and set theory are exceedingly tolerant disciplines. Any individuals, however disparate, constitute a mereological sum – an individual whole. So there is an individual made up of the world's largest frog, Fred's 2014 income tax return, Rockefeller Center, and the Brazilian navy. And except for pathologically self-referential cases, every collection of objects constitutes an extension. There is, then, an extension consisting of Napoleon's nose, the number 12, and an electromagnet in Shanghai. The extension of 'grue' is as determinate as the extension of 'green'. The entities belonging to a given mereological sum are parts of the same whole; the members of an extension are alike in being members of that extension. However gerrymandered they seem, these items are genuine individuals and extensions. So neither mere part-whole relations, nor mere membership in a common extension solves our problem. Again we have an embarrassment of riches. We need a basis for, as Mary Kate McGowan says, privileging the individuals and kinds that matter (2003).

Realists maintain that the world does the privileging (Lewis, 1999). If so, there are individuals and kinds whose identity and integrity is independent of anything we know or think about them. *Green* is privileged over *grue* because *green* is a more natural kind than *grue*. That is a fact about the mind-independent world. And we know that *green* is more natural than *grue* because natural science tells us so. Let's set aside the small complication that secondary qualities like *green* do not have a particularly elevated scientific pedigree, and investigate the contention that the most plausible thing to say about scientific kinds is that nature marks them out.

Ordinarily, philosophers focus on kinds like water or gold, samples of which have (or at least are alleged to have) a common, stable underlying structure. That structure is

supposed to explain their behavior or characteristics. But in science and in daily life, there are other kinds that seem unlikely candidates for nature's privileging.

One is *toxicity*. Many sorts of things are toxic. Some are biological, some chemical, some physical. All are deleterious, but they affect different organisms and harm them in different ways. Some are toxic to cells, some to organs, some to organisms. Radiation, chlorine, arsenic, snakebites, and cigarette smoke do not have the same causes or the same specific effects. Yet they all fall under the label 'toxic'. The rationale for grouping them all together and for developing the science of toxicology is our abiding interest in not poisoning or being poisoned.

The concept of a *syndrome* also raises difficulties. A syndrome is a constellation of signs and symptoms that collectively indicate a disorder. Often there is a single, perhaps unknown, underlying cause. Marfan's syndrome, for example, is caused by a genetic anomaly. A realist could maintain that the common cause is the reason for coalescing the signs and symptoms of Marfan's syndrome together and saying that people exhibiting those signs and symptoms have the same malady. But syndromes need have no common cause. Chronic fatigue syndrome is currently thought to result from a variety of distinct causes. All that sufferers of chronic fatigue syndrome have in common is their signs and symptoms. And the only available treatment treats the undesirable symptoms. So the different causal trajectories that got the patients to their current state turn out to be medically irrelevant. Granted each has a causal trajectory, and each particular sign or symptom has a causal background. But the rationale for grouping them together into a single syndrome and for saying that people who exhibit them have something significant in common is that this constellation is a robust undesirable condition.

Both *toxic* and *syndrome* are tightly linked to an interest in health and well being. Were it not for that interest, the causal story about the behavior of each individual toxic agent and each individual victim of chronic fatigue could be told. But there would be no reason to group the members of each kind together. The criterion for membership would look as arbitrary as the criterion for calling things grue. Given our interests, however, there is nothing arbitrary about the classifications.

Still, a realist might argue, these classifications are largely medical. It is perhaps not surprising that our interest in health and well being figures in the privileging of kinds suited to medicine. Maybe medical kinds have more in common with predicates like ‘fashionable’, and less in common with predicates like ‘nucleotide’ than we thought. But, he would maintain, interests play no role in demarcating real scientific kinds, the kinds that figure in the hard sciences.

Again, this is not so obvious as the realist makes it seem. Biological taxonomy classifies organisms into kinds. It is plausible that members of a taxon have something biologically significant in common with one another that they do not share with members of other taxa. This is largely true. However, there exists the biological family *Protista*, which consists of all eukaryotic organisms that are not animals or plants or fungi. It comprises, among other things, amoebas, slime molds, and bacteria, whose only common biological property (besides having cell nuclei) is that they are all *not* animals, or plants, or fungi. It is hard to see how nature could privilege the property of being *not an animal or a plant or a fungus*. The rationale for recognizing *Protista* as a family is evidently a desire for an exhaustive, simple taxonomy. To achieve that goal, taxonomists simply introduce a ‘none of the above’ category – *Protista*.

These cases are not exceptions, they are characteristic of categorization in general. Our interests, policies and past practices inform the construction of the categories in terms of which we represent and understand things. This is so even in the case of the realists' favorite so-called natural kinds, water and gold. When we identify water with H₂O, we ignore the fact that in any sample of water there are always some dissociated molecules (Teller, ms). We are right to ignore this fact, but in doing so, we idealize away from the way the world really is. Even the best cases cannot be construed as merely reflecting the way things are. With the growth of science and other systematic enterprises, we learn what kind of kinds suits our cognitive and practical purposes. If we want an understanding of a certain sort, we need to carve up the domain in a particular way.

Science favors repeatable results. So it has good reason to seek to classify its objects in terms of categories whose membership is determinate and publicly recognizable, and where an object's standing as a member of a particular kind is not easily destabilized. This requires that there be a lower bound on precision, that there be clarity about what differences matter, and that relatively small variations in surrounding circumstances do not affect whether a given object is an instance of a particular kind. Choice of units keys to interests as well. What sort of understanding we seek may affect whether to focus on cells, organisms, populations, niches, ecosystems, or something else. It might seem that these are just questions of bookkeeping. Truths about one just map onto truths about the other. But in the absence of strong reducibility, there is no reason to think that this is so. The understanding that the ecologist seeks may gloss over distinctions that are crucial in cell biology, and may appeal to considerations that are, from the perspective of cell biology, irrelevant.

Such considerations make constructive nominalism plausible. The interests, practices, actions, and decisions of agents figure in delineating the individuals and kinds that constitute the world. This is not to say, as an idealist might, that there is nothing outside of us. But epistemic agency is involved in configuring whatever there is into individuals and kinds. And there are in principle multiple, equally acceptable, but mutually irreducible ways of doing so.

We partition our world into toxic and non-toxic substances, into primary colors, into species and genera, and into a host of other groupings, because in so doing we equip ourselves with the conceptual resources to pursue our various interests. Human interests overlap, compete, and range from the very crude to the extraordinarily fine-grained. So do the categories in terms of which we mark things out.

Nor are human beings the only creatures with kind constituting interests. Many other animals carve up the world as well. Their categories may be relatively crude as compared to ours, but they serve their purposes. Creatures discriminate between animals that are potential mates and those that are not, between predators and non-predators, between sources of nourishment and things that are not sources of nourishment. Some discriminations are, no doubt, hard wired; others are the fruits of experience. Moreover, non-human animals sometimes draw their lines in places where we do not. Monarch butterflies are noxious to many birds. Viceroy butterflies (which look quite similar) are actually quite tasty and nutritious. But birds who have been sickened from eating a Monarch scrupulously avoid both, treating them as the same sort of unpalatable thing. That we think there is a difference between Monarchs and Viceroys is nothing to them. Their kinds are not our kinds.

Nevertheless, the contention that we make the categories that fix the facts seems at odds with the stubbornness of fact. The world is largely independent of our will and largely impervious to our beliefs about it. The assumption that the stubbornness of fact is incompatible with constructive nominalism does not stand up to scrutiny. With rare exceptions having to do with self-ascription, once categories are framed and criteria for their instantiation established, what if anything, instantiates those categories is independent of anything we may think or want. Consideration of obviously constructed categories shows that this is so. Human beings invented the game of baseball, and set the criteria for winning. We can revise them if we like. But given the current criteria, the Chicago Cubs simply do not belong to the extension of 'winning team'. However much we might like them to, wishing does not make it so. The stubbornness of fact then does not require naturally privileged kinds.

Still, if the entities and kinds that constitute our world are even partly a product of our efforts, then it seems, had we not made those efforts, there would be no such entities or kinds. If, for example, no one had ever dreamed up the category *dinosaur* there would have been no dinosaurs. That seems incredible. Suppose a meteor destroyed the Earth after the Mesozoic but before rationality evolved. Dinosaurs would still have existed, even though no one ever thought of them as such. This is surely right. But it does not impugn constructive nominalism.

Although constructive nominalism is committed to the view that if there were no concept of a dinosaur, there would be no such thing as a dinosaur, it has the resources to construe the concept rigidly. Given that there is such a concept, it applies to all entities in a particular extension, under the usual range of historical, hypothetical, and counterfactual circumstances. Thus we can project the category onto the possible

scenario in which rationality never evolved. In that case there would have been dinosaurs but no rational agents. We can also construe the concept nonrigidly and acknowledge that if rational agents had never contrived the concept of a dinosaur (or a concept co-extensive with it), the beasts wandering around in the Mesozoic would not have been members of a single kind. Had rational agents partitioned their world differently, different entities would have been members of a single kind. Maybe, for example, the carnivores would have belonged to a different kind from the herbivores, or the large ones to a different kind from the small ones, or the early ones to a different kind from later ones, or non-avian ones to a different kind from avian ones, etc. There are a vast number of reasonable ways of organizing the domain.

Imagine an extreme scenario: what if no rational agents ever existed, and no categories contrived by rational agents applied? Wouldn't there still have been dinosaurs? Even on this scenario, there still would have been organisms that recognized conspecifics, predators and prey. Those organisms would be able to identify potential mates with whom they could breed. They would be able to identify predators to avoid, and prey to hunt. But with nothing to privilege the category 'dinosaur', nothing would make it the case that at a higher level of generality all the conspecifics, predators and prey were the same kind of thing. At that level, there would be no privileged entities or kinds. All configurations of entities and all collections of them would be on a par. There would have been dinosaurs and rockosaurs and fernosaurs and cloudosaurs and every other combination. Everything, in a sense would be there, but there would be no privileged order or organization.

There are then several ways of understanding the claims 'If no one had ever contrived the concept of a dinosaur, there would have been no dinosaurs' and 'If no one

had ever contrived the concept of a dinosaur, there still would have been dinosaurs.' Each is intelligible and in certain contexts plausible. Trouble arises only if we conflate or confuse them.

The question remains: If constructive nominalism is correct, how is objectivity possible? To answer this requires saying something about objectivity. The critical point to be preserved is that the world is largely independent of our attitudes. Wishful thinking does not make things the way one wishes them to be; nor does believing that they are so make them so. Realism accommodates the stubbornness of fact directly. The facts are what they are, realists believe, regardless of anything we think or do. But if we consider the various ways the world is, this position looks less compelling than it first appears. My challenge is to undermine the realist intuition in a way that shows that on the nominalist alternative the world is suitably resistant to our attitudes. Let's look at some cases.

A game is obviously a human contrivance. People establish rules and criteria. These rules and criteria make it possible to do such things as hit a home run. Without the institutional framework that the rules of baseball delimit, no matter what happened when a club propelled a sphere, there would be no home run. Without the rules, there would be no bats, balls, teams, or home runs. But within the framework that constitutes the game, there are such things. And given the definition of a home run, it is a determinate, objective, mind-independent fact that Ortiz hit a home run. That he did so is not a matter of opinion; nor does wishing make it so.

It is open to us to tinker with the rules to make home runs easier or harder to get. So we can, through our rule making activities, affect how frequently 'x hit a home run' is true. But once the rules are in effect, whether a given player hit a home run is a matter of

fact. This is so, even if the rules of a game are arbitrary. In that case perhaps there is a sense in which the truth is arbitrary too. But still it is an objective truth in that it is independent of what we will or believe.

A game and the facts it generates may seem pretty insular. Outside of baseball, the fact that Ortiz hit a home run does not seem to matter much. The game itself seems to mark the boundaries for the significance of the fact. But not all institutional facts are so insular. Consider promising. What justifies the contention, and makes it the case, that Fred promised Bill that he would wash the car, is that Fred's actions complied with the rules specified by the institution of promising. Only because there is such an institution, does Fred's saying, 'I'll wash the car' make it a fact that he promised. Had there been no such institution, his utterance would be, at best, a prediction. In that case, he would have been under no obligation to wash the car. But if in fact Fred promised Bill that he would wash the car, and Fred is sufficiently trustworthy, then this fact – which the institution of promising makes possible – gives Bill a lien on the future. He is entitled to count on Fred to wash the car, and he is entitled to a measure of confidence that Fred will do so.

Not all practices are fact-constituting. Astrology is not. There are complicated astrological rules for calculating the influence of celestial configurations on terrestrial events. Those rules might yield the verdict that today is an inauspicious day for me to take a trip. But this does not make it a fact. For despite the complicated calculations, there is no reason to believe that the correlations that astrology alleges obtain. My horoscope does not give me a lien on the future. I would be epistemically irresponsible to believe it.

Rawls distinguishes between justification *in* a practice and justification *of* a practice (1955). Justification *in* a practice is determined by the rules and procedures that

define and govern that practice. That Ortiz hit a home run, that Fred promised to wash the car, that today is an inauspicious day for me to take a trip are all justified *in* the practices that generate them. Where the three differ is in the justification *of* the practices that generate them. The justification *of* the practice of baseball presumably lies in the enjoyment the practice affords to participants and fans. Insofar as people enjoy baseball, and neither the practice nor its enjoyment has significantly deleterious effects, baseball is justified. The justification *of* the practice of promising is practical. It is useful to have a mechanism whereby people can voluntarily restrict their freedom in advance, so that actions can be coordinated and cooperative, temporally extended activities can be undertaken. There is, evidently, no justification *of* the practice of astrology – at least insofar as astrology makes claims about the causal structure of the world.¹ For there is no reason to believe, and good reason to doubt, that its deliverances and methods are reliable. I suggest that only practices that are themselves justified are fact-constituting. For only where practices are justified is there reason to countenance their deliverances.

Consider now the case of science. The acceptability of a particular scientific finding, such as the finding that CD38 regulates oxytocin secretion (Jin et al., 2007), depends on its being justified *in* the practice of science. That is, on its satisfying the rules, standards and criteria that science sets. There must be intersubjective agreement among qualified scientists that the evidence supports the finding. The methods used to generate that support must be recognized as reliable. The results must be reproducible and reach the level of statistical significance. The theory that embeds the finding must display the theoretical virtues – simplicity, precision, robustness, explanatory power, and so forth. A finding that has these features will be justified *in* the science that endorses

them. But to understand why these features are the features an acceptable theory ought to have, we have to consider justification *of* the practice of science.

Pretty clearly science is different from baseball and promising. As we saw, there is and need be little justification of the practice of baseball. People like playing and watching the game. That's enough. Promising is justified by its utility. The existence of the institution of promising makes certain sorts of cooperation feasible. It might seem that justification of the practice of science is that the factors it favors are truth conducive. In that case, the findings that satisfy the standards of science are more likely to be true than those that do not. This is an appealing answer. Unfortunately, it is too simple.

Some theoretical virtues, such as robustness, appear not to be truth conducive. A well established precarious finding is no less likely to be true than an equally well established robust one. Other values, such as precision, are even more problematic. A precise finding is less likely to be true than a rough one. That some transmembrane glycoprotein regulates oxytocin secretion is more likely to be true than that CD38, in particular, does. So if our goal is truth, we should probably eschew precision. But to scrap such scientific values would be ridiculous. Rather, I believe, we should scrap the idea that the justification of the practice of favoring theories with such virtues is a matter of truth conduciveness.

I suggest that the justification of the practice of science consists in its being the best available method for engendering the sort of understanding we seek.² Obviously, to make this non-vacuous, I need to say something about what sort of understanding this is and why it is worth having. What we want from science is a systematic understanding of a broad range of natural phenomena. That understanding should be grounded in empirical evidence. It should explain and predict the phenomena and perhaps enable us

to exert control over them. This characterization might seem to say just that we want to know the way the world is, and that empirical methods are, as far as we can tell, the best way of finding out. In that case, the justification of the practice of science is that it is truth-conducive. This is not the whole story, though. In extracting the promise, Bill wants more than that Fred wash the car. He wants to be *entitled to confidence* that Fred will wash the car. Analogously we want more than a cluster of empirical claims that explain, predict, and perhaps enable us to control the phenomena. We want to be *confident* and to be *entitled to confidence* that science yields (or at least has the best prospect of yielding) an empirically grounded, systematic understanding of a broad range of natural phenomena. I suggest that many of the methods and theoretical virtues of science have the role of underwriting entitlement to confidence.

Comprehensive worldviews are notoriously fallible. Not only are the mythic worldviews of our ancestors riddled with errors, so are the scientific theories of our immediate predecessors. So, probably, are our own. If we want an account we can have confidence in, we should build in resources for avoiding, detecting, hedging against, and correcting errors. Intersubjective agreement among acknowledged experts serves as a stay against carelessness, bias, and simple mistake. Reproducibility of results controls for accident; if results are reproduced, and suitably varied experiments yield correlative results, the likelihood that the outcome is due to random factors or undetected interference diminishes. Statistical significance affords evidence that the result is not a fluke.

Moreover, we seek an understanding that is not easily dislodged. This requires that it be secure, stable, and robust. An understanding is insecure if its grounds are readily undermined. This may be because the particular bits of evidence are dubious – as,

for example, when the sole source of testimony is a shady character. It may also be because the methodology is problematic. If, for example, the reliability of focus groups is doubtful, then an account based entirely on focus groups is insecure. A *secure* finding is one whose grounds are solid. One reason why science favors quantitative over qualitative methods is that quantitative methods yield greater security. A finding is *stable* to the extent that it can survive shifts in grounds and *precarious* to the extent that it is vulnerable under such shifts. This is clearly a matter of degree. If, for example, several relatively independent bodies of evidence support a given conclusion, then it is likely to remain tenable even if some of the evidence is undermined. If it is supported by just a single datum, it is precarious. Should that datum be discredited, we lose all reason to endorse the conclusion. A *robust* conclusion is one that can survive revisions in its grounds. Again this is a matter of degree. If a conclusion is acceptable only if every bit of its support, just as it stands, is acceptable, it is fragile. If, on the other hand, the account is acceptable so long as its evidence is in the right neighborhood, then it is robust. A robust conclusion withstands fluctuations in the data.

An idiosyncratic, irreproducible result could be true. If *truth* is our sole objective, then science's penchant for intersubjective agreement and reproducible results seems unduly demanding. An insecure, precarious, fragile finding could be true and justified, so long as the evidence is exactly as it appears to be. Even if science seeks *justified truths*, then many of science's requirements seem too exacting. But if we seek theories we are entitled to confidence in, the epistemic situation is different. Then we have very good reason to endorse the values and methods of natural science. In that case, the entities, properties, and kinds recognized by science should be such that science can serve its end: deliver an understanding of nature that we can have confidence in. What privileges

particular entities, properties and kinds is that they suit practices that are themselves justified.

What about truth? Here again I draw on Putnam (1983). If a theory is consistent, it is possible to define a truth predicate for it. That yields truth-in-the-theory. If that theory satisfies the highest epistemic standards we can devise, we have every reason to believe that truth-in-the-theory is truth. So, assuming that both our wave theory and our particle theory are ideal, we have every reason to believe that the entities they recognize are real, and the conclusions they sanction are true. But the conjunction of two truths is true. So if

(a) 'Light consists of waves' is true

and (b) 'Light consists of particles' is true

then (c) 'Light consists of waves & light consists of particles' is true.

Doesn't this land me back in the 'too much of a good thing' position that I disparaged earlier? It need not. Suppose

(d) 'Twelve first-year students got an A' is true

and (e) 'Seven women got an A' is true.

Then (f) 'Twelve first-year students got an A & seven women got an A' is true.

But we cannot infer that nineteen people got As, for some of the first-year students who got As might have been women. Because the truth of (f) might involve redundancy, not all the seemingly obvious inferences can be drawn. Similarly, I suggest, for (c). Even though, we are assuming, it is true that light consists of waves and it is true that light consists of particles, it is not true that light consists of waves and particles. The laws of logic then remain universal, but in drawing inferences we need to be sensitive to

redundancies. This is so, regardless of our metaphysics. It is no more a problem for constructive nominalism than for any other metaphysical theory.

I have argued that constructive nominalism is a more plausible theory than realism. It affords a uniform treatment of truths. It countenances truths about games, fashion, ethics, and science and, at a suitable level of abstraction, gives the same account of all such truths. This, of course, counts as a virtue only if one thinks that there are genuine truths about games, fashion, and ethics. Otherwise, one would want to sharply distinguish such enterprises from the realm where genuine truths obtain. That, by consensus, is (or at least includes) the realm of natural science. But constructive nominalism figures in a better account of actual science than realism does, for it can vindicate the theoretical values of science, as realism cannot. And, on a constructive nominalist account, just as on a realist account, facts are stubbornly independent of our attitudes toward them. So constructive nominalism allows for objectivity. Wishful thinking is not an effective methodology. Reasoning, observing, testing, revising, and all the other hard work that goes into figuring out how things are retain their importance. What differentiates the nominalist from the realist is not her endorsement of these activities, but her understanding of how much we contribute to what we find by engaging in them.

References

- Jin, Duo et al. (2007). 'CD38, a transmembrane glycoprotein, regulates oxytocin secretion,' *Nature* 446:41-45.
- Lewis, David K (1999). 'New Work for a Theory of Universals' in his *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press, pp. 8-55.
- McGowan, Mary Kate (2003). 'Realism, Reference and Grue' *American Philosophical Quarterly* 40: 47-57.
- Putnam, Hilary (1978). 'Realism and Reason' in his *Meaning and the Moral Sciences*. London: Routledge and Kegan Paul.
- _____ (1983). 'Models and Reality' in his *Realism and Reason*. Cambridge: Cambridge University Press, pp. 1-25.
- Rawls, John (1955). 'Two Concepts of Rules,' *Philosophical Review* 64:3-32.
- Scheffler, Israel (2008). *Worlds of Truth*. Malden MA: Wiley-Blackwell.
- Teller, Paul (ms). 'Truth and Fiction in Science'.
- Van Fraassen, Bas (1980). *The Scientific Image*. Oxford: Clarendon Press.

¹□

Astrology might be justified, as baseball is, as an enjoyable pass time. But that would require denying that it affords any sort of access to the world beyond its borders.

²□'Best available method' is significant. It allows that the scientific method evolves over time. When, for example, statistics develops more refined methods, science will demand better statistics before countenancing a result.