PAPER IN GENERAL PHILOSOPHY OF SCIENCE

Perspectival realism and norms of scientific representation

Quentin Ruyant [1]

Abstract

Perspectival realism combines two apparently contradictory aspects: the epistemic relativity of perspectives and the mind-independence of realism. This paper examines the prospects for a coherent perspectival realism, taking the literature on scientific representation as a starting point. It is argued that representation involves two types of norms, referred to as norms of relevance and norms of accuracy. Norms of relevance fix the domain of application of a theory and the way it categorises the world, and norms of accuracy give the conditions for the theory to be true. Perspectival realism could be made coherent by taking a realist stance towards one type of norm, and a perspectival stance towards the other, by assuming they are relative to a community. This provides two versions of perspectival realism, called relevance perspectivism and accuracy perspectivism. How each option fares with respect to the challenge of incompatible models is examined. Finally, the prospects of full perspectivism are evaluated.

Keywords Perspectivism · Scientific representation

1 Introduction

Various authors (Rueger 2005; Giere 2010a; Massimi 2012; Teller 2018) have defended a version of realism that they call "perspectival realism", the idea being that some aspects of scientific knowledge are dependent on a cultural or historical context, yet in a sense that does not impair the central realist tenet that scientific theories are at least approximately true in virtue of the world.

This article belongs to the Topical Collection: Perspectivism in science: metaphysical and epistemological reflections
Guest Editor: Michela Massimi

✉ Quentin Ruyant
  quentin.ruyant@gmail.com

[1] Universidad Nacional Autónoma de México (UNAM), Mexico City, Mexico

Positions known as perspectivism in philosophy of language claim that the truth-value of assertions is to be evaluated with respect to a context of evaluation (MacFarlane 2007), that is, that it is perspectival facts that make assertions with a given content true.[1] Perspectivism can be invoked about various types of contextual aspects, for example the time of an utterance for tense statements in A-theories of time (Hare 2010), and the intuitive notion of a visual perspective on an object associated with a vantage point is often offered as a helpful analogy.

However, perspectivists in philosophy of science do not generally refer to concrete, spatio-temporal perspectives. What they have in mind is rather relativity to a conceptual scheme, to a theoretical lexicon (Giere 2013), to scientific categories for classifying phenomena and interpreting our observations, to epistemic norms of justification (Massimi 2012), or to metaphysical tenets interpreted in the spirit of Kant's synthetic a priori. The associated context is that of a given epistemic community. Furthermore, perspectival realists combine this aspect with a commitment to realism, and refuse to associate their position with a purely relativist doctrine. But they share the general idea that there is no "view from nowhere", and that scientific knowledge cannot transcend a human perspective.

There are various motivations for perspectival realism. One stems from the idea that the empirical data available to a community, as well as the way this data is interpreted, depends on the methodology, beliefs and technical apparatus in use in this community. Giere (2010a) takes as a starting point the idea that colours are attributed to objects in a way that is not entirely subjective (we talk of colours as if they were objective characteristics of objects, we do not say that objects lose their colour in the dark for instance), although they depend on our sensory apparatus. He then extends this rationale to scientific measurements. In this context, perspectival realism purports to be a middle ground between an anti-realist reinterpretation of scientific statements in terms of direct observations and the standard realist idea that theoretical properties would directly refer to natural kinds (Teller 2019) (although Massimi (2017) endorses natural kinds).

Another motivation is to secure the inference from success to truth that characterizes realist positions in a context where incompatible models are just as successful. This can be understood either diachronically, with respect to successive theories about the same phenomena (in relation to the pessimistic induction argument against realism) or synchronically, with respect to various models of the same target (in relation to debates on idealisation, or on reduction and emergence). For example, a fluid can be modelled as continuous or as being composed of discrete particles, and both "perspectives", although incompatible, are successful in predicting complementary aspects of fluid dynamics. In the case of diachronic perspectivism, the relevant context is historical, while in the synchronic case, it is cultural, or perhaps even more local. In all cases, the idea is that incompatible models could both be true with respect to different perspectives.

---

[1] This is distinct from contextualism, which is the view that the content, or truth-conditions, of assertions depends on the context of utterance (the latter could be analysed in terms of indexicality or "hidden indexicality"). However, the literature on perspectival realism in philosophy of science does not really bear on this distinction (see for example van Fraassen (2008)).

Some difficulties have been highlighted by different authors. With respect to the first motivation of the perspectivist, Chakravartty (2010) notes that the corroboration of various accesses to the same phenomenon by different means (from different perspectives) is typically taken to support knowledge of non-perspectival facts. He also claims that even if what our instruments measure is relational rather than intrinsic, this does not preclude knowledge of intrinsic dispositions of the measured object that might be partial, but not perspectival, taking elements of the context as stimuli for the dispositions. With regard to the second motivation, incompatible models, he claims that similar arguments can be given. Morrison (2011) examines two cases of incompatible models (in fluid dynamics and nuclear physics) and argues that perspectivism is of no help, because either the models are really about different aspects of a target, in which case there is not any threat to standard realism (models are partial rather than perspectival), or they are about the same aspects, in which case perspectivist attempts to secure success-to-truth inferences would boil down to a re-branded version of instrumentalism (see Massimi (2018b) for a response).

In any case, a tension remains in this combination of "perspective" and "realism". These two aspects, the perspectival aspect and the realist aspect, could seem difficult to combine at first sight, or even plainly contradictory.

Standard scientific realism can be analysed as the conjunction of three theses: (i) an external reality (construed as a set of mind independent facts) exists, (ii) scientific theories are either true or false in virtue of this external reality and (iii) our best scientific theories are at least approximately or partially true (Psillos 1999). The second thesis is called semantic realism, and it can be defined as the combination of a truth-conditional semantics for theories and a non-epistemically constrained theory of truth. As such, it contradicts perspectivism quite directly: if it is exclusively mind-independent facts that make our theories true or false (for example, a structure of natural kinds), or, to say it differently, if the truth-conditions of our theories are independent from any epistemic vantage point, then perspectivism is false. It would be of no help to claim that the perspectival aspect lies in the "approximate" or "partial" qualification involved in the third clause, for this is a triviality that standard versions of scientific realism already accommodate. So perspectivism is rightly understood as a semantic thesis, as a matter of *interpretation*, and as such, it contradicts the tenets of scientific realism mentioned above.[2]

Now this characterisation of realism might be too strong, given the plethora of positions labelled "realism" that do not fulfil it, for example internal realism. We could relax it a bit and claim that it is *in part* mind-independent facts that make our theories true or false, thus making room for perspectival aspects. But the exact way to articulate the epistemic relativity, implicit in the notion of perspective, with the mind-independence that characterises scientific realism, or the respective contributions of mind-independent and perspectival facts, remains to be developed. This prompts the following question: can perspectival realism really achieve this middle ground?

---

[2] Massimi (2018a) develops a version of perspectivism about "truth conditions" (understood as standards of justification rather than propositional content), and claims that it concerns the epistemic clause of scientific realism. Although this is an interesting position, I am not convinced that it constitutes an alternative to standard scientific realism in the context of the traditional debate (notably because justification does not imply truth). In any case, I will not address this version of perspectivism here.

In this paper, I propose a way of making sense of perspectivist positions by taking as a starting point the literature on scientific representation. I distinguish various kinds of norms that are at play in scientific representation, and that fix, or convey, the interpretation of theories by their users. My suggestion is that perspectival realism would make sense if some of these norms were relative to an epistemic community, while others were more universal, hence vindicating the idea that a perspectival and a realist component can be combined. I examine various options in order to determine whether perspectivism, so conceived of, can respond to Chakravartty and Morrison's challenges.

2 Norms of representation

It is now standard to assume that scientific representation is at least a three-place relationship between a user, a vehicle and a target (Suárez 2004; van Fraassen 2008; Giere 2010b). It is also generally accepted that misrepresentation is possible: a model can represent a target inaccurately, but still represent it. This means that one can distinguish what I shall call *relevance*, by which a model counts as a representation of its target, and *accuracy*, by which this representation is faithful to the target.

Arguably, relevance has a denotational component (or, in Suárez's terminology, "representational force[3]"): it involves the user taking a model to be *about* the target. Callender and Cohen (2006) have argued for a conception of scientific representation that reduces it to mental representation, and in their account, what makes a model a representation of its target is merely stipulation by the user of the model. All other aspects, having to do with the role models play in science, would be a matter of pragmatics. However, this account has been criticised for not distinguishing between symbolic and epistemic representation. With symbolic representation, the connection between the vehicle and the target is purely conventional, and the role of the representation is simply to pick out the referent, while epistemic representations "allow their users to have access, in however simplified or specialized manners, to aspects of their targets, which fulfil, in a broad sense, an epistemic role" (Liu 2015). The connection between scientific models and their targets is not purely conventional, but depends on "licensing" by the scientific community, which is responsive to empirical aims (Boesch 2017). In order to know what a model represents, one has to look at the history of its construction, reception and use. According to these criticisms, these aspects are not merely pragmatic: they are constitutive components of the representation relation.

I will assume, in this article, that the notion of relevance, that is, what makes a model a representation of its target, is not merely a matter of stipulation, and that it can indeed be captured by communal norms, what I will call *norms of relevance*.

For a model to be an *epistemic* representation, it must allow its user to make inferences on the target, in the same way a map allows us to infer that it is possible

---

[3] The notion is meant to capture the fact that the target involved in modelling can be abstract or fictitious, which denotation does not allow. For example, a model can represent a thought experiment, engineers can represent a non-existent bridge that will never be realised, or physicists can represent the hydrogen atom as an abstract type of entity, without their model representing one atom in particular. However I will focus in this paper on the representation of concrete entities, assuming that this is the locus of the debate on scientific realism.

to travel from one point to another in such or such a way. One can view communal norms of representation as constraints not only on the kinds of targets that can be represented by a model, but also on the inferences that are viewed as legitimate. This concerns the appropriate uses of models, or the interpretation of the model in terms of the target (which component of the model denotes which property of the target), and the right interpretation of the theoretical lexicon used in the model. Note that the legitimacy of inferences does not mean that these inferences yield true conclusions, because the model could be inaccurate. However, it gives us *accuracy conditions*. We can understand accuracy in this way: the model is accurate if the inferences it affords yield true conclusions. The conditions that would make the model accurate are therefore given by the inferences that are licensed by the community. I will refer to these communal constraints as *norms of accuracy*.

Can we say more about the representation relation? Some authors have proposed substantial accounts, but according to Suárez (2004), whether a model represents a target depends on an epistemic context. Some vehicles or inferences are warranted within a community while others are not, and there can be various means of representation and standards of accuracy. This entails that no substantial account of representation, in the form of necessary and sufficient conditions, can be given. Suárez's claim can be interpreted as the idea that norms of relevance are relative to a community, which could be one component of a version of relativism. Perhaps a substantial account could be compatible with this kind of relativity, but the minimal characterisation of scientific representation presented so far, in terms of norms of relevance and norms of accuracy, will be enough for my purpose.

In sum, we can characterise scientific representation with the following norms:

- Norms of relevance specify the conditions for any model in a given context to count as a representation of a target (they tell us "what a given model can be about");
- Norms of accuracy specify the conditions for any model in a given context to be an accurate representation of its target (they tell us "what a given model says" about its potential targets).

I do not take these norms to be necessarily explicit in a given scientific community. Rather I consider them as an ideal reconstruction of the community's activity and intents: when its members consider that a model is apt or not to represent a target, what types of phenomena they purport to represent, and what their ideal standards of accuracy are (even if these standards are beyond reach: I wish to consider ideal, complete accuracy here, taken to be a property capable of grounding theoretical truth, rather than mere justification[4]). This bears, among other things, on the way they categorise phenomena, operationalise theoretical models and interpret their observations and the content of their models.

I would also stress that the direction of fit for these norms need not necessarily be from our representations to the world. Ensuring that a model does represent its

---

[4] Even if the realist accepts that our best theories are only approximately true, I presume that this is with respect to an ideal notion of truth, which is what I am interested in here. The "approximate" part is a matter of epistemology, while the "truth" part is a matter of semantics, and as already explained, I take perspectival realism to be a semantic position.

target, or that it is an accurate representation, could involve intervening in the world, for example through a preparation procedure, or through calibration of instruments, so as to make the target and the experimental context fit the model rather than the converse (I assume that in general, complying to representational norms involves a mutual adjustment between a model, a target and an experimental context). Scientific models are in general *performative*. This is where contextual factors could be implied.

As an illustration of the role of these norms, take the theory that water is the kind of substance composed of H_2O molecules. We would expect that a model of the theory (a certain configuration of H_2O molecules) represents its target only if this target is water (in a sense to be defined by the norms). This is a *relevance norm*. The model is accurate if the target is indeed composed of H_2O molecules. This is an *accuracy norm*.

The criteria by which a model is considered relevant or accurate can be more or less local. The reasons why a model is legitimate or not to represent a given target in a certain way and to make certain inferences can be multiple: this can depend on background knowledge entrenched in the community (this type of phenomenon must be modelled in this way, this apparatus must be used to measure this quantity), perhaps on accepted conventions, or on pragmatic aspects and practical aims that are more specific to an activity (the model must be computable, the expected level of precision is high for this quantity, this object does not matter). For example, if scientists merely want to predict the future value of a quantity, this value (and nothing else) will determine the accuracy of their model, so it cannot be a prerequisite for the model to represent the target that the value be correct. But if they want to *explain* this value, then it is a condition for the relevance of the model that it includes the value as an effect of whatever mechanism is posited. This is a mere question of salience, which depends on a local context.

However, it is quite plausible that other distinctions will be less dependent on the local context, but will appear at the level of an epistemic community. In the example of a theory of water given above, it would seem that the predicate "water" denotes, as a matter of generality, a type of object that must be present for the theory to apply at all, that it delimits the domain of application of the theory, while the term "H_2O" belongs to theoretical posits. And we could expect a certain consistency in the way the theory is applied by a community across various contexts, and in the way the theoretical lexicon is interpreted, so these norms must be shared by all local contexts. Even more local aspects could be captured by giving to cross-contextual norms the form of a function from local context to local constraints, assuming that some elements of the context (salient objects and properties, expected degrees of precision) can be formalised. By norms of accuracy and norms of relevance, I understand cross-contextual norms at the level of the epistemic community rather than local contextual criteria.

So far I have been talking about models, but perspectivist positions are generally presented as positions about theories. I shall assume that theories can be characterised by a collection of models. Let us provide a few connections between the two.

First, norms of relevance can help us specify a domain of application for the theory. The rationale is the following: if, according to a community of users of a theory, there is a model of this theory that is apt to represent a given target of

representation for at least one possible context of use, then this target belongs to the domain of application of the theory.[5] This same idea can help us *categorise* the domain of the theory, if the legitimate targets for various models or classes of models form a partition of this domain (models of a different class cannot apply to the same target).

Secondly, norms of accuracy give conditions for the theory to be true in its domain by fixing the circumstances under which its models would be accurate. It has became common-place, following the semantic conception of scientific theories, to characterise a theory as a collection of models. But a collection of models is not a truth-bearer, and for realist positions to make sense, one should also characterise the theory by a statement asserting something about how its models relate to their targets (this point has been made by early proponents of the semantic conception of theories; Giere (1991, p, 85) mentions "various hypotheses linking those models with systems in the real world", and Suppe (1989) talks about "a theoretical hypothesis claiming that real-world phenomena [...] stand in some mapping relationship to the theory structure"). This characterisation will presumably rest on the assumption that a theory is true if its relevant models are generally accurate (or perhaps, if it has at least one accurate model for any target in its domain), so that accuracy plays the role of the "mapping relationship" mentioned by Suppe. So we could characterise theoretical truth, as a first approximation, by the following formula, taking M to be any model of the theory Th and C a context of use, including a target T:

$$\text{Truth}(Th) := (\forall C, T)(\forall M \in Th)\ \text{Relevant}(C, T, M) \rightarrow \text{Accurate}(C, T, M)$$

The general idea expressed by this formula is that the theory is true if all its models are accurate insofar as they are relevant according to general norms (we cannot expect a model that does not represent a target at all to represent it accurately). The arrow in the formula can be read as a material conditional. This should hold whatever the context of use and whatever the target involved, hence the universal quantification on targets and contexts.

This formula could perhaps be refined, but the important point for our purpose is that theoretical truth depends on both norms of relevance and norms of accuracy. By analogy with the idea from philosophy of language that meaning can be analysed in terms of truth conditions, there is a sense in which communal norms specify "what the theory is about" and "what the theory says", by giving the conditions under which the theory would be true.

The distinction between norms of relevance and norms of accuracy and the involvement of both in any definition of theoretical truth is crucial for the following reasons:

- We do not want to trivialise theoretical truth (at least I assume that perspectival realism does not). Presumably, we would like to say that phlogiston theory is false

---

[5] The notion of possible context of use is required to make sense of the notion of domain of application, because the latter is not restricted to actual applications of the theory. This prompts the question of whether and in what sense various possible contexts of use are part of the same larger epistemic context or perspective. I think we can make sense of this hierarchy of contexts in terms of norm sharing, where norms are more or less local. I will come back to this idea later in this article.

for example, or at least that some theories are capable of being false within a perspective, and not merely irrelevant. This condition is directly associated with the possibility of misrepresentation mentioned at the beginning of this section. If norms of relevance and norms of accuracy were intensionally equivalent, that is, if it was not *conceivable* that one obtain and not the other, theoretical truth (per the formula above) would obtain as a matter of conceptual necessity, and not in virtue of the world. So the two kinds of norms must be intensionally distinct.

- We do not want to make theoretical truth too hard to achieve either (that would impair the epistemic component of scientific realism). For example, we would not want to say that a theory of optics is false because it does not accurately represent combustion phenomena. We do not expect a theory to represent the whole universe.

This means that norms of relevance must not be too strict, so as to become equivalent to norms of accuracy, otherwise any model would be accurate insofar as it represents something and all theories would be trivially true. Nor must they be so loose that they allow any model to be a legitimate representation of anything, otherwise all theories would be false. In order to achieve this balance, we must distinguish, at the theory level, the conditions that specify the right domain of application of the theory (so as to avoid that the theory is trivially false) and the conditions that are merely involved in the accuracy of a theory (so as to avoid that the theory is trivially true whenever it applies).

Intuitively, we would expect norms of relevance to depend on the general categories used to classify target phenomena, and norms of accuracy to correspond to states or dynamical aspects of these phenomena predicted by the theory. We could also expect norms of relevance to be associated with surface features that are more directly accessible, and by which scientists identify various targets of representation, or characterize the kind of phenomena they want to model, and norms of accuracy to correspond to theoretical posits. We cannot expect the same kind of scepticism with regard to our ability to know that a model represents something at all, as opposed to the scepticism we can legitimately entertain with regard to its accuracy. Presumably, accuracy is an ideal that is more difficult to achieve (and that is not necessarily achievable!) and it involves aspects that are less directly accessible. However things could be more complex, notably when theories concern unobservable entities, or when the accuracy of a theory is implicitly assumed in the application of another theory (this aspect will be examined at the end of this paper).

I do not claim that this distinction is always made explicit when presenting a theory, but it is something we might discover by asking a scientist: "What kind of phenomena is your theory about? What is its scope?" and "How does it describe or explain those phenomena? What are its posits?". If the distinction between these two kinds of questions makes sense, and if these questions can be answered, then there are such things as norms of relevance and norms of accuracy, and if this is not the case, then it is unclear how we could maintain that theories are either true or false in a non-trivial sense.

To end this section, let me note some caveats:

- Norms of representation are not necessarily known a priori: the best way to apply a theory to the world, to delimit its scope and to classify phenomena that must be modelled differently, can be learned by experience;

- They could evolve with time, for example when the inaccuracy of a theory leads to limiting its intended domain of application, or when the identification of a phenomenon acquires enough stability;
- Perhaps it is a matter of degree, implying that a model is more or less applicable to a target, or that the contour of the domain of application of the theory and its categorisation of targets are vague;

But again, without this distinction, it is not clear how one could be a realist: one either ends up assuming that no theory is false, some just do not apply to anything, or that theories apply in principle to anything, they have no restricted domain of application, so they are always false in many ways. And the problem is not necessarily restricted to realism: it is simply that we want our theories to be about something specific, and that we want them to say something substantial (potentially false) about this thing.

3 How can perspectivism and realism be combined?

In the previous sections, I have given a general framework that is a priori compatible with different views in the debate on scientific realism. Accepting this framework, at least part of this debate can be turned into a debate about norms: different positions about scientific realism will assume different norms of representation. To give an example, a verificationist could express norms of accuracy in terms of various operations and results on the target, while a realist could interpret these norms in terms of real objects and natural kinds (recall that norms of accuracy need not be achievable).

Obviously, expressing conditions of accuracy does not imply that theories are accurate: a realist would also like to say that these conditions are satisfied for the models of our best theories, or perhaps for most of them. I am merely addressing the semantic component of scientific realism here. But this is where the debate on perspectivism is located: once the question of what it means for a theory to be true or false is settled, there is only room for scepticism or optimism, and no place left for perspectivism (at least as I understand it).

If we accept this framework and its distinction between norms of relevance and norms of accuracy, there are three ways of making sense of scientific perspectivism, which are the following:

- Relevance perspectivism: norms of relevance are perspectival (relative to an epistemic community), but norms of accuracy are objective.
- Accuracy perspectivism: norms of accuracy are perspectival, but norms of relevance are objective.
- Full perspectivism: norms of relevance and norms of accuracy are both perspectival.

Let me say more about what I mean by perspectival and objective. We can assume that a norm of relevance could be expressed by a proposition such as "A model M can be used to represent a target T in a context C if and only if $\varphi(C, T)$", and similarly for norms of accuracy. Now claiming that such norm is relative to a community in the sense that other communities do not accept this norm is

necessary, but not sufficient for perspectivism. For example, the norm could rest on communal conventions regarding the interpretation of symbols appearing in the model. But variations in conventions among communities do not imply perspectivism, or at least not in a very interesting sense: the fact that the word "snow" is used in English, and "nieve" in Spanish to refer to snow does not entail that snow is a perspectival notion with respect to linguistic communities. A norm could also not be in used in a community because its members are not interested in a kind of target, and again, this does not entail any interesting sense of perspectivism.

In order to narrow the notion of perspective, let us say that a norm expressed by a proposition of the form mentioned above is perspectival if φ makes ineliminable reference to a specific epistemic community. This would be the case if φ contained indexicals or hidden indexicals, the meaning of which can only be elucidated by referring to an epistemic context, or if it made ineliminable reference to the abilities or conceptual scheme of the users of the representation. This would be the case in particular if it referred to particular actions, such as measurements, that could be executed by the members of the epistemic community only. Such norms are not accepted by other communities not because they use different conventions, but because they do not have the same epistemic resources. By contrast, a norm is objective if φ can be expressed without referring to the epistemic community (for example, by referring to natural kinds only).

It could seem that if the function φ that characterises a norm refers to aspects of an epistemic community, these aspects could in principle be naturalised, and φ could eventually be expressed in terms of natural properties, perhaps relational properties between the target and the users of the representation. In the end, the norm would not really be perspectival. But this move shifts the target of representation, from the initial target to the combination of it and the user or community. For this reason, it does nothing to show that the initial norms of representation, which concerned representation of the initial target alone, are not perspectival. All this move does is remark that it might be possible to move from a perspectival representation of something to an objective representation of something larger including the user of the representation. But these objective representations are not necessarily accessible to users, and perhaps the new representation would also be perspectival. Claiming that this is not the case merely begs the question against perspectivism.

Here is a good candidate for a perspectival norm of relevance: the way fishermen, chefs or biologists classify fish into families is distinct, although the extensions of these classifications (the set of classified objects) coincide. Presumably, a fisherman's categories depend on particular aspects of fishing activities and associated practical knowledge. In this sense, it would be relative to a perspective: one could not say whether or not a fish belongs to a category without referring to the users of this categorisation and their activities. But the norms of accuracy associated with the fisherman's representation could nonetheless be perspective independent if what fishermen believe about the members of each category happens to be objectively true, independently of this way of categorising fishes.

Now here is a candidate for a perspectival norm of accuracy: music magazines publish every year listings of the best albums of the year. All magazines agree on

what counts or not as an album, yet the rankings usually differ. We can assume that magazines use implicit norms for evaluating albums, and that according to these norms, their ranking is accurate. But the norms are relative (of course these examples could be challenged, but they are meant as an illustration).

Recall that theoretical truth depends on both types of norms. The idea behind the two first positions, relevant and accuracy perspectivism, is to combine the apparently conflicting perspectival and realist aspects of perspectival realism in a coherent way, by applying them to different types of norms. Now the picture could be more complex. Norms of representation could be more or less local, sometimes specific to a particular activity, and the accuracy—relevance distinction could be irrelevant to perspectivism. It could be interesting to examine whether standards of relevance and accuracy follow idiosyncratic norms or not in actual scientific practice to settle the issue. Nevertheless, I wish to explore, from the armchair, what the positions mentioned here would imply when applied to scientific theories, and whether they are plausible. I will examine in particular their ability to account for the success of incompatible models, using Morrison (2011)'s case studies.

3.1 Relevance perspectivism

Let us start with relevance perspectivism. The idea can be illustrated by Morrison's case study of fluid mechanics. Prandtl's model of water pictures it as a fluid with viscosity, while Euler's model pictures it as a fluid without viscosity. But the two models purport to apply in different contexts: Euler's model is relevant for large amounts of fluids away from boundaries and the Prandtl's model is relevant for fluids near solid boundaries, where friction becomes important and cannot be ignored. This difference does not obviously concern the users of the models and their activities rather than aspects of the target, but let us assume it does for the sake of the argument. In this case, it looks as if we have different rules for identifying *what the model represents*, that is, the kinds of targets that are legitimate, even if, in both cases, the target of representation is *water*. Context of use affects relevance. This is compatible with relevance perspectivism: different users classify objects and phenomena in a different way, but what they say about these objects is objectively true.

The problem is that the two models are incompatible: does water have viscosity or not? In our context, if norms of accuracy are objective, this would mean that the norms of accuracy of each model must give different results when applied to a concrete target: if one is satisfied, the other is not. This looks like a contradiction.

One answer could be that viscosity, although a natural property, is not actually part of the conditions of accuracy for Euler's model: the model can be accurate without water actually lacking viscosity. This strategy of taking idealisations to abstract away characteristics is often applied by the realist, and as Morrison observes, this gives no particular advantage to the perspectivist. A problem for the realist is when such idealisations play an indispensable explanatory role (Batterman 2005). This affects this perspectivist strategy as well.

However, another strategy is more distinctive. It is to observe that norms of accuracy will apply to models only insofar as they first satisfy norms of relevance (the two are involved in theoretical truth), so that what really determines the truth conditions of a theory is how conditions of relevance and conditions of accuracy are related. Assuming

that a natural property such as viscosity grounds accuracy, the model is accurate not in virtue of the natural property viscosity being present in water in general, but rather in virtue of it being present in the contexts where norms of relevance are satisfied. Here the conclusion would be different: Euler's model represents water as lacking viscosity, so water, from Euler's perspective, lacks viscosity, while from Prandtl's perspective, it has viscosity. Viscosity, which is a real property (out there, in the world), is related to one perspective on water but not to another, and these relational facts make the two models accurate, in an objective sense, even though they seem incompatible.

As we can see, relevance perspectivism is able to capture the idea that scientific theories are perspectival, but nonetheless true in a way that depends (in part) on a mind-independent reality. Now is this a good option?

What could give credit to this kind of perspectivism is the idea that norms of relevance are presumably more easily applicable than norms of accuracy (one cannot be sceptical in the same way with regard to the fact that a model represents something as with regard to its accuracy). This is explained by this kind of perspectivism: the reason would be that we access the world from a perspective, to which we have direct access, and from which we classify phenomena. But assessing the accuracy of our representation would be more difficult because accuracy is not perspectival, and less directly accessible.

One problem is that in the case of the models of water examined by Morrison, the differences of domain of application do not seem to depend on the epistemic community, but rather on the target itself. And even if it was the case, it is always possible to reinterpret situations of this kind as cases of polysemy: if the norms by which the identification of "water" differ for the two perspectives, why claim that the two models are about the same substance? Perhaps this could be solved by invoking shared norms attached to the identification of water in both contexts (associated with a broader epistemic community, the human community perhaps). But then the question is: does the extensions of concrete targets to which the incompatible models apply coincide? Is it possible for the norms of relevance of the two models to be satisfied at the same time? If the extensions are distinct, maybe these models are models of particular water conditions rather than models of water *simpliciter*. And then again, even assuming that the way these water conditions are identified is relative to a community, there is no advantage of perspectivism over standard realism. What we would have is the idea that models give partial representations of reality, because they represent different aspects or different targets, so their incompatibility is only apparent.

On the other hand, when the extension of targets coincide (norms of relevance can be satisfied for the two incompatible models in at least one situation) the problem of incompatible models cannot be avoided by this version of perspectivism. If conditions of accuracy are not perspectival, they must concern the target alone, not relations between the target and the context (this would mean switching the target), and if the target can be the same viewed from different perspectives, either conditions of accuracy are satisfied or they are not: two incompatible models cannot both be accurate.

For this reason, it is not clear that this form of perspectivism can solve the problem of incompatible models. This remark also applies to diachronic cases. The Newtonian theory of gravitation and general relativity were designed to account for the same class of phenomena, classified in more or less the same way (in terms of the presence of massive bodies), but described differently. It might be possible to claim that the way

they "carve nature" is actually distinct, and perspectival, even though their scopes coincide. But if one theory claims that there are forces of gravitation, and the other that there are only deformations of space-time, and if this matters for an objective notion of accuracy, then the theories are incompatible, and something has to go.

The fact that this version of perspectivism cannot account for incompatible models, or at least not in ways that are distinctive from standard realism, does not mean that it is false. However, it undermines a strong motivation for perspectivism.

3.2 Accuracy perspectivism

Accuracy perspectivism is perhaps better captured by Morrison's example of models of the nucleus mentioned above. In this case, various models apply to the same, real entity: the nucleus. Relevance norms, which specify what a model can represent, are not perspectival. Different communities could agree to apply their different, incompatible models in the very same circumstances. But how the nucleus is represented would be perspectival, and whether the model is accurate would depend on the perspective of evaluation, so that each community could rightly judge, by its own standards, that its model is accurate. The same real atom nucleus would resemble a drop from one perspective, and it would be like a shell from another, and both descriptions would be correct, but from different theoretical perspectives associated with different measurements on the nucleus.

This kind of perspectivism captures the notion of perspectival truth quite directly. It is closer to perspectivism in philosophy of language. One worry is that it looks like a slight departure from standard realism, in so far as the only thing the accuracy perspectivist is realist about is the set of accessible phenomena to be represented and their categorisation. The content of representation (what would make models accurate) is relative to the perspective. So is it, as Morrison suspects, a mere re-branding of instrumentalism? The main motivation for scientific realism is to explain the empirical success of scientific models and theories (including incompatible ones) by appealing to their correspondence to the world. If norms of ideal success are relative to a community, and if the correspondence to the world only concerns the categorisation of phenomena, it does not seem that the realist explanation will work.

A way to avoid this unwanted result is to claim that what makes a theory true, in this picture, is really a relation between the natural states of affairs that ground norms of relevance and an epistemic perspective. Then perhaps one is a realist about more than the classification offered by a theory, but also about the structure between members of this classification, which would be reflected in the way various classes are related to perspectival aspects (an argument very similar to the "upward path" towards structural realism (Psillos 2001)). This would bring perspectival realism quite close to structural realism (which is not necessarily a problem, see Wolff (2019)). One could consider that the notion of perspective merely provides an epistemological complement to the structural realist narrative, and explain how the "structure of reality" can be extracted from relations between perspectival representations of real entities.

This strategy is similar to the one suggested in the previous sub-section: what matters for theoretical truth is how conditions of relevance and conditions of accuracy are related, and the fact that norms of relevance are objective would prevent the position from collapsing into instrumentalism. This time, the perspective does not

correspond to the way we categorise the phenomena to be represented, but rather, so to speak, the way we look at them to assess the accuracy of our models (this is closer in spirit to Massimi (2018a)'s association of perspectives with standards of justification).

This version of perspectivism is subject to Chakravartty (2010)'s challenge: various perspectival characterisations of the nucleus could be reinterpreted as mere dispositions, taking the relevant context as a stimulus. This objection could be resisted by recalling that the said stimulus is expressed in a way that is relative to the user of the representation (perhaps using hidden indexicals or referring to idiosyncratic practices and concepts in use in a community). Then the dispositions that can be attributed are not "real", but perspectival, since they are never expressed in terms of "real" stimuli and manifestations. As argued before, attempting to naturalise these aspects means changing the target of representation, and claiming that it is always possible to eliminate perspectival aspects in this way is question-begging.

This version of perspectivism might be a good way of putting perspectivist ideas to the task of solving philosophical problems. The idea that norms of accuracy are more contextual than norms of relevance is also quite plausible: clearly, theoretical properties often have vague boundaries that are not easily extendible across contexts (such as "atom" in plasma, see Teller (2019)), while ways of categorising phenomena of interest are more easily shared among different communities. Taking norms of relevance to be objective seems to clash with some of the motivations for perspectival realism, notably the example of colours, and frequent claims to the effect that different perspectives "categorise the world differently". But it could be supported by a causal theory of reference for accessible kinds (after all, Kripke's story about baptism ceremonies for kinds, by which a name is assigned to a kind by direct ostentation of one of its exemplars, seems to apply better to accessible objects than to unobservable theoretical posits). Finally, this version seems more easily applicable to diachronic cases: Newtonian forces of gravitation might not be real entities, but still a good way of characterising gravitational phenomena from a theoretical perspective situated in flat space-times. Note, however, that the way the position solves problems is not necessarily distinctive from the strategy of structural realists.

3.3 Full perspectivism and hierarchies of contexts

There is a dilemma for the middle-ground positions presented above. Take accuracy perspectivism: if a model describes a gas as being composed of molecules, and if accuracy is perspectival, then it is only from a perspective that the gas is composed of molecules. But we also have models of these molecules. The conditions of relevance of these models seem to imply that models of gas are accurate. If the fact that these models apply is not perspectival, then molecules must really exist. We have a contradiction. The problem also affects relevance perspectivism: if models describing gases as composed of molecules are accurate in an objective sense, then molecules must exist, but then, how can it be perspectival that a model of molecule applies? And how can we make sense of the idea, alluded to at the end of Norms of Representation section, that the distinction between accuracy and relevance changes with time when the identification of a phenomena becomes sufficiently stabilised?

It might be tempting, for these reasons, to move towards full perspectivism, so as to take the advantages of both positions above, and achieve a more coherent position.

However I would urge the person who would do so to abandon her realist commitment: claiming this label as a badge of honour, as it were, is futile. If all norms involved in representation are perspectival, then what our theories describe are perspectives, not a mind-independent reality. Of course, the perspectivist can entertain the idea that "reality resists", and that what is true within a perspective somehow depends on real constraints external to the perspective. But as Chang (2017) argues, to the extent that these external constraints are unknowable, what we have is a noumenal conception of reality that is insufficient to sustain genuine scientific realism.

Perhaps at this point it could be useful to recall the main motivation for scientific realism: the inference from success to truth. Scientific realists are generally careful to explain that the success of a theory with regard to the phenomena it was designed to account for does not really count. They put emphasis on *novel predictions*. The idea is that the theory is successfully extended to new domains of experience. They also put emphasis on the corroboration of theories through various accesses to the same entities. How shall we understand this line of argument?

Perspectivism actually offers a good way to express it: the theory has been successfully transposed from one perspective to another, and as it happens, its models are also accurate in this new perspective. In the case of novel predictions, this seems to imply different norms of relevance for the theory (the domain of application is extended), and in the case of corroboration, different norms of accuracy (the same posited entity is accessed differently).

If we accept the inference from success to truth, then we should accept that there is some correspondence to something real that is preserved across perspectives. What is it? Relevance perspectivism has an answer: the accuracy of scientific models is grounded in real states of affairs, independent of the epistemic community, which explains the successful extension to a new perspective.[6] Accuracy perspectivism could make its case by endorsing some kind of structural realism: there is a structure underlying real entities, and the theory gets part of this structure right. But how could the full perspectivist address this question?

An aspect often put forth by perspectival realists is that the realist component of their doctrine rests on the possibility of cross-perspective assessment (Massimi 2018a). For a perspectivist, this is precisely what happens in the case of novel predictions and corroboration: a theory is assessed, and approved, from different perspectives. But if we do not want to be relativist and adopt incommensurability theses, this implies that the norms of one perspective can somehow be transposed to a new one. For corroboration to work, we must be able to claim that the represented entity is still of the same kind, even if the norms of accuracy (the way of measuring posited entities) are different, and for novel predictions to work, we must be able to say that the posited theoretical entities are still interpreted and measured in the same way, even if the norms of relevance (what kind of system the theory can represent) are different.

This points to the idea that both contexts share a "super-context", a common denominator of norms, from which the two contexts can be compared. And what

---

[6] This is for the case of novel predictions. In the case of corroboration, the relevance perspectivist could simply argue that corroboration through different accesses confirms her realism with respect to accurate properties.

the full-perspectivist could answer to the realist is this: novel predictions and corroboration show that the theory is true *relative to this super-context*. This means that the theory is now known to be less local and more universal than before: science has progressed. The same idea could be applied, in the other direction, to more local contexts: an epistemic context encompasses a set of more local possible contexts, associated with various activities that can be performed within this epistemic context, in the same way a "super-context" can unify various epistemic contexts. Perhaps epistemic contexts come to life by unifying various local activities in the same conceptual scheme, with the same applicable norms. This provides a nice picture of science as evolving towards more unification, encompassing more and more local activities into general schemes, while allowing local perspectives to enjoy a certain autonomy.

The notion of a super-context could seem contradictory with perspectivism. However, the idea is merely that contexts can be more or less broad, and that they can have part-whole relationships. It is not required that there be a unique, absolute super-context that contains all others, so the pluralist component of perspectivism can be maintained. This notion and its consistency with perspectivism should be analysed further, which cannot be done here for lack of space. At least it provides a potential solution to realist challenges.

However, a broader perspective is still a perspective. Truth is still relative. So there is a sense in which the idea of cross-context assessment does warrant some of the realist ideas (a notion of scientific progress towards more unification and more universality), and can provide an explanation for the impressive success of science in its own terms. But a full-perspectivist is still no realist.

4 Conclusion

I have argued that scientific representation involves two kinds of norms: norms of relevance, which fix the domain of application of the theory and the way it categorises the world, and norms of accuracy, which gives the conditions for the theory to be true, in particular in terms of the interpretation of posited entities. This gives two possible ways of articulating a perspectival and a realist component, relevance perspectivism and accuracy perspectivism, depending on which of these two kinds of norms is deemed perspectival, in the sense of being relative to an epistemic community, and which is deemed objective. The prospects for relevance perspectivism are not very good: it cannot really address the problem of incompatible models, unless these models are really about different targets. But then the strategy is available to standard realism and the notion of perspective does not play a role. Accuracy perspectivism is more apt to address the challenge of incompatible models, but the solution it provides is ultimately very similar to structural realism. One problem for both options is accounting for cases when one theory's accuracy is implicitly assumed when applying another theory's norms of relevance. Full perspectivism, which takes all norms of representation to be perspectival, is a more coherent position in this respect. It could explain the empirical success of science and account for a notion of scientific progress in terms of a hierarchy of contexts. However, full perspectivism is not a realist position.

References

Batterman, R. W. (2005). Critical phenomena and breaking drops: infinite idealizations in physics. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, 36*(2), 225–244. https://doi.org/10.1016/j.shpsb.2004.05.004.

Boesch, B. (2017). There is a special problem of scientific representation. *Philosophy of Science, 84*(5), 970–981.

Callender, C., & Cohen, J. (2006). There is no special problem about scientific representation. *Theoria: Revista de Teoría, Historia y Fundamentos de La Ciencia, 21*(1), 67–85.

Chakravartty, A. (2010). Perspectivism, Inconsistent Models, and Contrastive Explanation. *Studies in History and Philosophy of Science Part A, 41*(4), 405. https://doi.org/10.1016/j.shpsa.2010.10.007.

Chang, H. (2017). Is pluralism compatible with scientific realism. In J. Saatsi (Ed.), *The Routledge Handbook of Scientific Realism* (1st ed., pp. 176–186). New York: Routledge,. | Series: Routledge handbooks in philosophy: Routledge. https://doi.org/10.4324/9780203712498-15.

Giere, R. (1991). Explaining science: A cognitive approach. *Philosophical Review, 100*(4), 653–656.

Giere, R. (2010a). *Scientific Perspectivism* (Paperback ed.). Chicago: Univ. of Chicago Press.

Giere, R. (2010b). An Agent-Based Conception of Models and Scientific Representation. *Synthese, 172*(2), 269–281. https://doi.org/10.1007/s11229-009-9506-z.

Giere, R. (2013). Kuhn as perspectival realist. *Topoi, 32*(1), 53–57. https://doi.org/10.1007/s11245-012-9142-y.

Hare, C. (2010). Realism about tense and perspective: Realism about tense and perspective. *Philosophy Compass, 5*(9), 760–769. https://doi.org/10.1111/j.1747-9991.2010.00325.x.

Liu, C. (2015). Re-inflating the conception of scientific representation. *International Studies in the Philosophy of Science, 29*(1), 41–59. https://doi.org/10.1080/02698595.2014.979671.

MacFarlane, J. (2007). Relativism and disagreement. *Philosophical Studies, 132*(1), 17–31. https://doi.org/10.1007/s11098-006-9049-9.

Massimi, M. (2012). Scientific Perspectivism and its foes. *Philosophica, 84*, 25–52.

Massimi, M. (2017). Perspectivism. In J. Saatsi (Ed.), *The Routledge Handbook of Scientific Realism* (1st ed., pp. 164–175). New York: Routledge. | Series: Routledge handbooks in philosophy: Routledge. https://doi.org/10.4324/9780203712498-14.

Massimi, M. (2018a). Four kinds of perspectival truth. *Philosophy and Phenomenological Research, 96*(2), 342–359. https://doi.org/10.1111/phpr.12300.

Massimi, M. (2018b). Perspectival modeling. *Philosophy of Science, 85*(3), 335–359. https://doi.org/10.1086/697745.

Morrison, M. (2011). One phenomenon, many models: Inconsistency and complementarity. *Studies in History and Philosophy of Science Part A, 42*(2), 342–351. https://doi.org/10.1016/j.shpsa.2010.11.042.

Psillos, S. (1999). Scientific realism: How science tracks truth. In *Philosophical Issues in Science*. London ; New York: Routledge.

Psillos, S. (2001). Is structural realism possible? *Proceedings of the Philosophy of Science Association, 2001*(3), 13–24.

Rueger, A. (2005). Perspectival models and theory unification. *The British Journal for the Philosophy of Science, 56*(3), 579–594. https://doi.org/10.1093/bjps/axi128.

Suárez, M. (2004). An inferential conception of scientific representation. *Philosophy of Science, 71*(5), 767–779.

Suppe, F. (1989). *The semantic conception of theories and scientific realism.* Chicago: University of Illinois Press.

Teller, P. (2018). Referential and perspectival realism. *Spontaneous Generations: A Journal for the History and Philosophy of Science, 9*(1), 151. https://doi.org/10.4245/sponge.v9i1.26990.

Teller, Paul. 2019. "What is Perspectivism, and does it count as realism?" in *Understanding Perspectivism*, 49–64. Routledge. https://doi.org/10.4324/9781315145198-4.

van Fraassen, B. (2008). *Scientific representation: Paradoxes of perspective* (Vol. 70). Oxford: Oxford University Press.

Wolff, Johanna. 2019. "Representationalism in measurement theory. Structuralism or Perspectivalism?" in *Understanding Perspectivism*, 109–126. Routledge. https://doi.org/10.4324/9781315145198-7.