

Representing and Explaining: The Eikonic Conception of Scientific Explanation

Alisa Bokulich

Philosophy Department
Boston University
abokulich@bu.edu

Abstract:

The ontic conception of explanation, according to which explanations are "full-bodied things in the world," is fundamentally misguided. I argue instead for what I call the *eikonic* conception, according to which explanations are the product of an epistemic activity involving representations of the phenomena to be explained. What is explained in the first instance is a particular conceptualization of the explanandum phenomenon, contextualized within a given research program or explanatory project. I conclude that this eikonic conception has a number of benefits, including making better sense of scientific practice and allowing for the full range of normative constraints on explanation.

1. Introduction:

The vast majority of philosophical work on scientific explanation to date has been focused on *accounts* of explanation. An account of explanation is a claim about how explanations work: that is, do they work by deducing the explanandum statement from the relevant laws and initial conditions (such as in the covering-law account), or by tracing the relevant causes (as in a causal account), or by identifying the mechanistic components and their activities (as in the mechanistic account)? This paper is not about accounts of explanation; rather, it is about what might be called *conceptions* of explanation. A conception of explanation is a view about what explanations are. Currently, the orthodox conception of scientific explanation is the ontic conception, according to which explanations are "full-bodied things" (Craver 2014, 40); that is, explanations just are the concrete causes or mechanisms in the world themselves. Versions of the ontic conception have been defended by Wesley Salmon (1989), Carl Craver (2007, 2014), and Michael Strevens (2008). It is only in recent years that the ontic conception has come under scrutiny with significant criticisms being raised (e.g., Wright 2012, 2015; Illari 2013; Sheredos 2016; Bokulich 2016). However, the issue of conceptions of explanation remains so undertheorized that rival alternatives to the ontic conception have yet to be fully articulated and defended.

My aim in this paper is to outline and defend an alternative conception of scientific explanation that I call the *eikonic* conception (from the Greek word 'eikon' meaning representation). The eikonic conception can be understood as falling under the general umbrella of "epistemic" conceptions¹, a version of which has been defended

¹ As we will see, however, it is a different sense of "epistemic conception" than Salmon (1984, 1989).

(though typically only within the context of mechanistic explanations) by what might be called the "San Diego School" of explanation (e.g., Bechtel and Abrahamsen (2005), Bechtel (2008), Wright (2012, 2015), Sheredos (2016)). Like William Bechtel et al.'s version of the epistemic conception, I argue that "[scientific] explanation is fundamentally an epistemic activity performed by scientists" (Bechtel 2008, 18), with explanations themselves being an outcome or product of such an epistemic activity. The eikonic conception emphasizes the importance of representations and representational choices in constructing an explanation. As I will argue, representations are not just involved at the level of the explanans, nor just at the level of the 'explanatory text' (e.g., the particular diagram or equation); rather representations play a prior and more fundamental role in our very conceptualization of the explanandum phenomenon itself.

I will motivate the need for this new conception of scientific explanation in three ways: first, by (briefly) reviewing the rival ontic conception and its shortcomings (in Section 2); second, by examining scientific practice, with a couple of case studies (in Section 3); and third, by highlighting the philosophical benefits that come from moving to the eikonic conception of scientific explanation (Section 5). The eikonic conception is outlined in Section 4, paying particular attention to the different ways that an explanandum phenomenon can be represented, depending on the context and aim of the explanatory project. Although I will be rejecting the ontic conception, it is important to note that ontic constraints still play a central role on the eikonic conception. What form those ontic constraints take, however, depends on the particular context of the inquiry (see, e.g., Bokulich 2016). In what follows, I grant that explanation and understanding are "success terms", in that they require getting something right about the way the world is, and more generally, I take the eikonic conception of explanation to be compatible with a broadly realist approach to science. As will become clear, my primary concern here is with scientific practice, that is, with developing a conception of scientific explanation that can best help us understand what scientists are actually doing when they offer scientific explanations.

2. The Ontic Conception

An early version of the ontic conception of explanation was proposed by José Alberto Coffa and was elaborated by Wesley Salmon most clearly² in his book *Four Decades of Scientific Explanation* where he writes,

Coffa is a staunch defender of the ontic conception of scientific explanation. . . . For Coffa, what explains an event is whatever produced it or brought it about. . . . Explanations, in his view, are fully objective and . . . exist whether or not anyone ever discovers or describes them. Explanations are not epistemically relativized, nor . . . do they have psychological components, nor do they have pragmatic dimensions" (Salmon 1989, 133).

² Salmon also discusses the ontic vs. epistemic conceptions in his earlier (1984) work, but there the discussion is even more confusingly entangled with discussions of determinism vs. indeterminism.

Salmon similarly endorses the ontic conception, asserting that "[e]xplanations exist in the world" (Salmon 1989, 86).

Salmon contrasts the ontic conception with what he calls the epistemic conception.³ However, Salmon is not using the notion of a conception of explanation in quite the same sense as is being used here. For Salmon, the ontic-epistemic distinction provides a taxonomy for different *accounts*: causal and mechanistic accounts fall under the ontic conception, while Hempel's covering-law (nomothetic-inferential) account and van Fraassen's erotetic ("why questions") account fall under the epistemic conception. As I am using the term, however, a conception is a claim about what explanations are, independent of the choice of any particular account. So one could adopt either an ontic or epistemic conception of the causal account, or similarly adopt an ontic or epistemic conception of a nomothetic account. This early literature, which fails to clearly distinguish between accounts of explanation and conceptions of explanation is thus not particularly helpful for our present project.

More recently the ontic conception has been vocally defended by Carl Craver (2007, 2014), who offers a nice, clear statement of how the ontic conception of explanation should be understood:

Conceived ontically . . . the term explanation refers to an objective portion of the causal structure of the world, to the set of factors that produce, underlie, or are otherwise responsible for a phenomenon. Ontic explanations are not texts; they are full-bodied things. They are not true or false. They are not more or less abstract. They are not more or less complete. They consist in all and only the relevant features of the mechanisms in question. There is no question of ontic explanations being "right" or "wrong," or "good" or "bad." They just are. (Craver 2014, 40)

The ontic conception has been endorsed by several other scholars working on scientific explanation, such as Michael Strevens (2008). Indeed one could say that the ontic conception has become the orthodox view in the philosophy of scientific explanation today.

To summarize, on the ontic conception, explanations just are (in the sense of 'are identical to') the particular tree branch (that broke the window) or the particular electron (that ionized a water molecule). The claim is not just that these things are causes or causally relevant (which they undoubtedly are), but that they are further scientific explanations. On the ontic conception, scientists and scientific theorizing are not actually needed for there to be scientific explanations.

Surprisingly, it is only in recent years that the ontic conception has come under scrutiny. A striking problem for ontic theorists is that they have trouble talking consistently about explanation in a way that is true to the ontic conception.⁴ For example, in Craver's paper "The Ontic Account of Scientific Explanation" (which, on the terminology being urged here, should have been titled "The Ontic Conception of

³ He also contrasts both with the modal conception, which will not be discussed here.

⁴ For this criticism of Salmon's ontic conception see Wright (2015) and for a criticism of Craver's version see Illari (2013) and Bokulich (2016).

Scientific Explanation") he notes that two goals of an adequate theory of explanation are explanatory demarcation and explanatory normativity. If, however, one substitutes in "the causes and mechanisms in the world themselves" for the occurrences of "explanation" in his discussions, as the ontic conception requires, his statements become incoherent:

The theory should distinguish ~~explanations~~ [causes] from other forms of scientific achievement. ~~Explanation~~ [Cause] is one among many kinds of scientific success; others include control, description, measurement, and taxonomy. . . . The second goal is explanatory normativity. The theory should illuminate the criteria that distinguish good ~~explanations~~ [causes] from bad". (Craver 2014, 28).

Such statements and projects only make sense if one rejects the ontic conception of explanation and adopts an epistemic or representational conception instead. As Wright (2015) notes it is by slipping back and forth between representational and non-representational notions in unacknowledged ways that the ontic conception has arguably enjoyed an unwarranted success. Phyllis Illari (2013) has similarly urged,

In so far as Craver is interested in normative constraints on explanation, presumably he is not interested in constraints on mechanisms—the mechanisms themselves simple *are*—he is presumably interested in how ontic features . . . constrain explanatory *descriptions* of mechanisms. (Illari 2013, 243)

If Illari is right, then it becomes all the more important to start fleshing out what such an epistemic or representational conception would look like, so that we can begin to address the further question of what form those ontic constraints on our descriptions or representations should take. Some progress on this latter question has recently been made by Sheredos (2016) and Bokulich (2016). As shown next, the motivation for moving to an eikonic conception derives not just from the shortcomings of the ontic conception, but also from the need to better account for scientific practice.

3. Motivating the Eikonic Conception: Scientific Practice

When scientists investigate and seek to explain a complex entity or phenomenon in the world, they typically begin with a simplified representation. Different subfields of science will often represent a phenomenon differently, reflecting not only the different interests they may have in that phenomenon, but also the different theoretical resources and methodological tools available within that subfield. This point can be seen particularly clearly in Helen Longino's (2013) discussion of investigations into human aggression, for example. Longino examines five different subfields of science: quantitative behavioral genetics, molecular behavioral genetics, developmental psychology, and social-environmental studies. Each of these subfields conceptualizes and operationalizes human aggression differently. She writes,

[I have] surveyed five major approaches or families of approach to studying human behavior. . . . But how behavior is understood as an object of

investigation differs: it can be treated as disposition or as episode, as a dimension of variation in a population among populations, or as an individual characteristic. (Longino 2013, 103)

She notes that they not only conceptualize the phenomenon of human aggression differently, but also carve up the space of possible causal explanatory factors differently. So for molecular behavioral genetics, what she calls genotype 1 (allele pairs) and genotype 2 (whole genome) are part of the measured causal space for possible explanatory factors, while physiology (e.g., hormones) and anatomy (e.g., brain structure), intrafamily environment, and socio-economic status are unmeasured space. On the other hand, for social-environmental approaches, intrafamily environment and socioeconomic status are part of the measured causal space, while genotype 1 and 2, anatomy, and physiology are unmeasured, and hence not part of the causal space of resources for explanations. The other subfields (e.g., quantitative behavioral genetics and developmental psychology) carve up the measured causal space differently yet again.⁵

Longino concludes that this plurality of approaches is not in fact in competition, each one offering insights into different aspects of the phenomenon. She writes,

[E]ach approach can provide partial understanding of human behavior broadly understood, and . . . each can provide answers to the questions specific to it, but . . . certain features of the investigative space preclude their full integration or the elimination of one or more in favor of a single encompassing and unified approach. (Longino 2013, 125)

The plurality of representations of the explanandum phenomenon then is not a weakness of the investigation, but rather a source of strength. Confining ourselves to just one of these approaches would result in understanding less about the nature of human aggression, than we learn through this plurality of representations.

One can also find a plurality of representations of an entity or phenomenon within a single field of science. Here the plurality is not due to the different conceptual and methodological resources of different subfields, but rather due to the fact that certain representations are more useful for answering certain kinds of questions about an entity or phenomenon (while other representations lend themselves to answering other sorts of questions). An example of an entity of this sort is water.

Given the ubiquity and familiarity of water it may come as a surprise that explaining its properties—many of which are quite anomalous—is an active, indeed fraught area of scientific investigation:

Water is the most extensively studied molecule of unique importance to life. Yet our understanding of how this deceptively simple compound of just three atoms gives rise to the many extraordinary properties of its liquid phase is far from complete. The complexity of the water properties combined with multiple possible levels of approximation (e.g., quantum vs. classical, flexible vs. rigid)

⁵See Longino 2013, pp. 128-129, Figures 2, 3, and 4.

has led to the proposal of literally hundreds of theoretical and computational models for water. (Izadi et al. 2014, 3863)

No one model of water is able to account for all the experimentally well-known liquid-phase properties, and hence different models, involving different representations of water, have to be used in different contexts (Guillot 2002). Depending on which properties of water scientists want to explain, they will choose a different representation.

At the broadest level, there are three main classes of representations of water. First, there are continuum representations, where water is conceived of as a continuous substance with macroscopic properties taken to be well defined down to infinitesimal volume elements. Second, there are what are known as classical atomistic representations of water, which are the familiar “ball and stick” (or ball and spring) representations of water. And third, there are quantum representations of water. A challenge for quantum representations of water, however, is that one cannot exactly solve the Schrödinger equation for even a single water molecule.

Within each of these broad classes there are several possible further ways that water can be represented. Take, for example, classical atomistic representations of water, which can be further divided into three classes of models. There are, first, fixed-charge, rigid models, with fixed atom positions (i.e., fixed |OH| bond length and <HOH angle), which are used to explain properties such as the high static dielectric constant of water. Second there are flexible models that allow bond stretching and angle bending, which are needed to explain the vibrational spectra of water. And third, there are polarizable models, which are likely needed to explain phase transitions in water.

One can further zoom in on any one of these classes of models to again find a plurality of representations. So, for example, within the classical atomistic, rigid (nonpolarizable) models, one can find various n -site models.

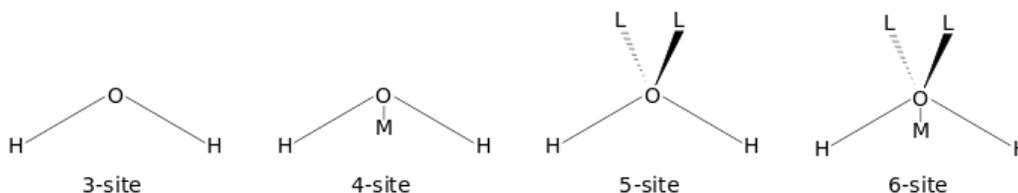


Fig. 1: N-site classical atomistic rigid models
(Source: https://commons.wikimedia.org/wiki/File:Water_models.svg)

There are the traditional 3-site models (e.g., the TIP3P model), which have three point charges corresponding to the atom positions. Then there are more advanced 4-site models (such as TIP4P), in which a “dummy” atom (labeled ‘M’ in Fig. 1) with a negative charge is included. There are also 5-site models (e.g., the TIP5P model) with two dummy atoms (labeled ‘L’ in Fig. 1) arranged in a tetrahedral geometry, which are used to better capture the charge distribution of the molecule. There are even 6-site representations of the water molecule, which are a combination of the 4-site and 5-site models.

It is instructive to examine briefly one of these representations of water, the OPC model. OPC is a 4-site, 3-charge, “optimal point charge” model introduced by Saeed

Izadi and Alexey Onufriev, in which the intuitive, on-nuclei placement of the point charges is abandoned in favor of a placement that optimally reproduces the required electrostatic charge distribution of a water molecule. They note the motivation for this new approach to modeling water is twofold: First, “none of these [hundreds of water] models faithfully reproduce[s] all of the key properties of bulk water simultaneously” (Izadi et al. 2014, 3863). Second, they note that increasing the realism of the models often does not yield better outcomes for these key properties. They were led to develop a model of water in the liquid phase that placed the three point charges in whatever locations best reproduced the electrostatic potential, hence the name optimal point charge (OPC) model. They note that “an accurate representation of electrostatic interactions is paramount for accurately accounting for hydrogen bonding and the properties of liquid water” (3864). However, by ‘accurate representation’ of the charge distribution of a water molecule, what they mean here is an accurate representation *within* the three-point-charge framework. Once the three-point-charge representational framework or vehicle is chosen, the question then becomes how to choose the values and positions of those three point charges such that they accurately produce the actual electrostatic charge distribution of the water molecule.

As noted before, understanding the electrostatic charge distribution of the water molecule (and hence understanding how the different hydrogen atoms will bond) is one of the most critical factors for explaining why liquid water has the distinctive properties that it does. After determining the optimized geometry of their 3-charge, 4-point OPC model, Onufriev and colleagues construct a ‘quality map’ of the various water models in the dipole (μ) and quadrupole (Q_T) space, in terms of how well they are able to explain the emergence of six key bulk properties of liquid water. The traditional water models, such as TIP3P and TIP5P mentioned earlier, did rather poorly when it came to accounting for the six test properties of water. Perhaps surprising, however, the models whose moments were determined by quantum-mechanical calculation also performed poorly. Onufriev and colleagues write,

The low quality of the test models in which the moments were close to the QM values . . . suggests that, within the 3-charge models explored here, an optimal fit of moments to QM predictions does not guarantee agreement with experimental liquid phase properties. (Izadi 2014, 3866)

This highlights an important point that has been largely overlooked in the philosophical literature on modeling. It is often implicitly assumed that the more realistic a model is, the more reliable the model output will be. In general, however, it will not be the case that increasing the realism of the model leads to a proportional increase in the accuracy of the model predictions. Once within the framework of a particular model representation, scientists often have to deliberately get things more wrong in the model in order for the conclusions to come out more right. This point was made by Bokulich and Oreskes (2017) in the context of models in the Earth sciences, where they note that this can occur also for physical “table top” models, such as in a scale model of a river, where instead of realistic river water, one might need to substitute a fluid of a different viscosity in order to compensate for the distorting scale effects.

In the case considered here, the failure of the water models with the “more realistic” quantum mechanical values for the moments can be understood as a result of the constraints of working within the rigid, non-polarizable, 3-charge representational framework. I would argue that the OPC model is still an accurate representation (to some relevant degree of accuracy) of the charge distribution of a water molecule (which as noted before is the key causal factor in explaining the peculiar properties of liquid water) but that it is an accurate representation with one set of representational quantities chosen as the vehicle, rather than another set that could have served as well.⁶

4. The Eikonic Conception

As we see in these case studies, scientists do not study phenomena in their full complexity, rather they study a simplified representation of the phenomenon contextualized within a particular field, research program, or explanatory project. This is in order to make the phenomenon tractable with the conceptual, theoretical, instrumental, methodological, etc. tools available within a particular subfield. This was seen in the case of human aggression, where different fields not only represented the explanandum phenomenon differently, but they also had a different space of causal resources available for the explanans. Even within a single subfield, one can often find a plurality of representations of an explanandum phenomenon, as was seen in the case of water. This is because some representations may be more appropriate for answering certain kinds of questions, while other representations are needed for answering other sorts of questions. When scientists explain a phenomenon (e.g., one of water’s anomalous properties), they do so within the context of one of these (several possible) representations. Often there is no “one best” representation that allows scientists to answer all relevant questions about that entity simultaneously. What we think of as the most realistic or fundamental representation does not always do the best job (as seen for example in the quantum models of water). As argued earlier, this plurality of representations is not a weakness of the investigation, rather it is a strength, allowing us to learn *more* about an entity or phenomenon than we would with any one representation alone.

Pace the ontic conception, explanations are not full-bodied things in the world. Rather, explanations are the product of an epistemic activity involving representations that are used to facilitate knowledge production and advance understanding. Representations are involved not only in the explanans (e.g., as in Bokulich’s (2011) model-based account of explanation), but also in the explanandum. What scientists explain in the first instance is not the phenomenon-in-the-world itself, but the phenomenon-as-represented; that is, the explanation is tied to one of the possible ways that the phenomenon can be conceptualized within the context of a particular research program. The phenomenon-as-represented is not some sort of subjective cognitive state, but a publically accessible representation shared by a community of researchers. Nor should the phenomenon-as-represented be equated with a particular diagram, equation, or linguistic description (what Craver (2014), for example, calls the “explanatory text”). It is not that the scientists explain the electron-as-represented-by-a-blue-dot on a particular

⁶ For a discussion of how (credentialed) ‘fictional’ representations can be a vehicle for truth see, for example, Bokulich (2016).

diagram. Rather, the eikonic conception calls attention to a prior, more fundamental conceptualization or representation of the phenomenon: Is the explanation conceiving of the electron as a classical particle, quantum particle, or excitation of a quantum field? These representational choices are both more basic and more important than the choice of any particular diagram or text. Which way the electron gets conceptualized depends on the particular explanatory context and purpose of the inquiry (and of course the world constrains which representations are, or are not, going to be adequate for a given explanatory context). And while the choice of conceptualization restricts the range of admissible diagrams, equations, or linguistic descriptions, it should not be identified with any one of these.

As noted above, the eikonic conception draws a distinction between the *explanandum phenomenon*, which is the phenomenon-in-the-world, and the *explanandum*, which is the particular conceptualization or representation of the explanandum phenomenon that is the immediate (though of course not ultimate) target of our explanations. Hence we can see that there are really two steps involved in evaluating a scientific explanation: First, is the explanans adequate in accounting for that particular explanandum (phenomenon-as-represented), and second, is that particular representation of the explanandum phenomenon an adequate representation (of the thing in the world) for the context of inquiry and question under investigation? While considerable philosophical attention has been paid to the first question, very little has been said about the second.

It is worth pausing briefly to distinguish the eikonic conception from other views on explanation, with which it bears some similarity. For instance, both the eikonic conception and van Fraassen's (1980) erotetic account emphasize the importance of context for determining what the appropriate explanation is. However, the eikonic conception, as a *conception*, is committed to none of the detailed machinery of van Fraassen's *account*, (e.g., the logic of questions, where a why question is understood formally as a triple of topic, contrast class, and relevance relation) nor to the antirealist perspective in which the account is situated. As noted in the introduction, the eikonic conception also bears some similarity to the "San Diego School" of mechanistic explanation, as expressed, for example, by Bechtel and Abrahamsen, who write,

[S]ince explanation is itself an epistemic activity, what figures in it are not the mechanisms in the world, but representations of them. These representations may be internal mental representations, but they may also take the form of . . . diagrams, linguistic descriptions, mathematical equations, physical models. . . (Bechtel and Abrahamsen 2005, 425)

Again, however, the eikonic conception is not tied to any particular account of explanation—in this case the mechanistic account. Epistemic conceptions have remained inadequately developed largely because they have not been disentangled from particular accounts of explanation. Another difference with Bechtel et al.'s approach is that the eikonic conception focuses on a different respect in which representations enter into explanations—namely community-level conceptualizations—rather than an individual's cognitive state or a particular diagram. Finally, the eikonic conception highlights the central role of representations not just in the explanans, but in the *explanandum*.

The eikonic conception directs our attention to the many different ways in which an explanandum phenomenon can be conceptualized and represented, prior to the choice of any particular explanatory text. Although the following list makes no claim to completeness, here are seven important ways in which an explanandum phenomenon gets simplified and conceptualized within an explanatory context. First, when a complex entity or phenomenon involves a number of different causes coming together to produce an effect—especially when the different causal factors belong to different subfields of science—scientists will typically focus on a *subset of the causal lines* and properties of the phenomenon as being the “important” ones to focus on (as seen for example in the different studies of human aggression).

Second, an entity or phenomenon is represented at a particular *spatial scale*, which can lead to different variables and theories being explanatorily relevant (e.g., at large scales, continuum representations of water and the Navier-Stokes equations are typically more relevant; see also Batterman (2013)). Third, an entity or phenomenon is represented at a particular *temporal scale*, which can affect which causal factors are important or negligible. For example, in trying to explain the features of a cliff on a short time scale, storm events are important and uplift is typically negligible; while on long time scales, storm events are negligible and uplift is important.

Fourth, there are decisions to be made about how the *boundaries of the system* or entity of interest are to be drawn and delimited. Is a factor or process taken to be internal to the system and to evolve with it (“autogenic”) or external to the system affecting it only from the boundaries (“allogenic”)? Fifth, a complex entity—or even a simple one as we’ve seen with the case of water—can be *ontologically decomposed* in different ways into different components (as in the various *n*-site models of water). Some explanations work better via one particular decomposition, rather than another.

Sixth, the explanandum phenomenon is represented at a particular *level of abstraction*. Consider, for example, when fluvial geomorphologists want to explain why rivers braid: some rivers flow straight, some meander, while others exhibit a braided pattern with a number of interwoven channels. There are many different levels of abstraction at which this phenomenon to be explained can be framed: In many cases scientists are not interested in why the Waimakariri river in New Zealand specifically (or the Platte river in Nebraska) braids; but instead, why do rivers braid in general? There are of course many different levels of abstraction at which the phenomenon can be framed between these extremes.⁷ Seventh, the explanandum phenomenon is represented as belonging to the *ontology of a particular theory* (e.g., is the electron conceived as a classical, non-relativistic-quantum, or quantum-field-theoretic entity?).

It is important both philosophically and scientifically to pay attention to all of these representational choices that are being made when one is constructing and assessing scientific explanations. As we have seen, representations of the phenomenon play a central role at the very outset of the investigation, even before a concrete explanatory text is given. Hence it is important to recognize, again *pace* the ontic conception, that it is not the full-bodied things in the world that enter into our explanations, but rather a

⁷ For a discussion of the various levels of abstraction at which an explanandum phenomenon can be framed (and how that relates to how-possibly versus how-actually explanations) see Bokulich (2014).

representation of those things, whose appropriateness to any explanatory context should be assessed, rather than assumed.

5. Conclusion

It is important to articulate and flesh out an epistemic conception of explanation, such as the eikonic conception, apart from the defense of any particular account. So far I have tried to motivate the eikonic conception by, first, briefly recalling the weaknesses of the rival ontic conception, and, second, by showing how the eikonic conception helps us better make sense of scientific practice. Although they cannot be fully defended here, I want to conclude by briefly highlighting some of the philosophical benefits that come from moving to the eikonic conception.

In addition to helping us understand scientific practice, the second philosophical benefit is that the eikonic conception—unlike the ontic conception—allows us to talk normatively about scientific explanations as good/bad, deep/shallow, etc. Third, the eikonic conception readily makes sense of explanatory pluralism. We can see that some (though not all) cases of scientific disagreement or competing explanations may in fact be only apparent, stemming from different conceptualizations of the explanandum (e.g., Longino's case of human aggression). Fourth, the eikonic conception has the resources to help us better understand the role of mathematics, diagrams, idealizations, and fictions in scientific explanation. Finally, the eikonic conception, unlike the ontic conception, allows us to recognize the full range of constraints on scientific explanation: ontic constraints, coming from the way the world is; communicative constraints, coming from the context of the inquiry; and epistemic constraints, coming from the fact that the aim of scientific explanation is to further the understanding of the scientific community, qua human knowers.

Acknowledgements:

I would like to express my deep gratitude to Alexey Onufriev for many stimulating discussions about water modeling research. I am also grateful to audience members at various venues (including PSA'16, Western University, Prague, University of Edinburgh, Cambridge University, and Durham University) for constructive feedback on earlier incarnations of this paper. I am most indebted, however, to Wendy Parker, whose uncompromising standards and incisive insights, always delivered in a generous and constructive spirit, make her an ideal editor.

Works Cited

- Batterman, Robert. 2013. "The Tyranny of Scales." In *The Oxford Handbook of Philosophy of Physics* ed. Robert Batterman. Oxford: Oxford University Press, pp. 255-286.
- Bechtel, William. 2008. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. NY: Routledge.
- Bechtel, William. and Adele Abrahamsen. 2005. "Explanation: A Mechanistic Alternative" *Studies in the History and Philosophy of Biology & Biomedical Sciences* 36: 421-441.
- Bokulich, Alisa. 2011. "How Scientific Models Can Explain." *Synthese* 180 (1): 33–45
- Bokulich, Alisa. 2014. "How the Tiger Bush Got Its Stripes: 'How Possibly' vs. 'How Actually' Model Explanations," *Monist* 97(3): 323-340.
- Bokulich, Alisa. 2016. "'Fiction As a Vehicle for Truth: Moving Beyond the Ontic Conception" *The Monist* 99 (3): 260-279.
- Bokulich, Alisa. and Naomi Oreskes. 2017. "Models in the Geosciences," *Handbook of Model-Based Science*, ed. by L. Magnani and T. Bertolotti. Springer, pp. 891-911.
- Craver, Carl. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: OUP.
- Craver, Carl. 2014. "The Ontic Account of Scientific Explanation" in M. Kaiser et al. (eds.), *Explanation in the Special Sciences: The Case of Biology and History*. Synthese Library 367. Dordrecht: Springer.
- Guillot, Bertrand. 2002. "A Reappraisal of What We Have Learned During Three Decades of Computer Simulations on Water" *Journal of Molecular Liquids* 101 (1-3): 219-260.
- Illari, Phyllis. K. 2013. "Mechanistic Explanation: Integrating the Ontic and Epistemic." *Erkenntnis* 78: 237-255.
- Izadi, Saeed, Ramu Anandakrishnan, and Alexey Onufriev. 2014. "Building Water Models: A Different Approach," *The Journal of Physical Chemistry Letters* 5: 3863-3871.
- Longino, Helen. 2013. *Studying Human Behavior: How Scientists Investigate Aggression and Sexuality*. Chicago : University of Chicago Press.
- Salmon, Wesley. 1984. "Scientific Explanation: Three Basic Conceptions" *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, Vol. 1984, Volume Two: Symposia and Invited Papers. (1984), pp. 293-305.
- Salmon, Wesley. 1989. "Four Decades of Scientific Explanation" in P. Kitcher and W. Salmon (Eds.) *Scientific Explanation*, Minnesota Studies in the Philosophy of Science, Vol. XIII. Minneapolis: U of MN Press, 3-219.
- Sheredos, Benjamin. 2016. "Re-Reconciling the Epistemic and Ontic Views of Explanation (Or, Why the Ontic View Cannot Support Norms of Generality)" *Erkenntnis* DOI: 10.1007/s10670-015-9775-5
- Strevens, Michael. 2008. *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
- van Fraassen, Bas. 1980. *The Scientific Image*. Oxford: Oxford University Press.

- Wright, Cory. 2012. "Mechanistic Explanation Without the Ontic Conception", *European Journal of Philosophy of Science* 2: 375-394.
- Wright, Cory. 2015. "The Ontic Conception of Scientific Explanation" *Studies in History and Philosophy of Science* 54: 20-30.