21
MODELS AS SYMBOLS

Catherine Z. Elgin

1. Introduction

Scientific models are a motley crew: some are concrete, others abstract; some are static, others dynamic; some represent states of affairs, others simulate processes; some have targets, others do not; some closely resemble their targets, others drastically distort. Nevertheless, scientific models of all sorts function epistemically. They embody and advance understanding. A critical question is how they do so.

The answer might seem obvious: models are similar to their targets; by investigating a model, we learn about its target. Sadly, this is too simple. First, it cannot accommodate models without targets. Phlogiston models, ether models, and caloric models turned out to lack targets. Nevertheless, their status as models was not rescinded. Nor are all targetless models the results of mistakes. Biologists invoke a model species with four sexes in order to investigate population dynamics.[1] Physicists devise models of perpetual motion machines to deepen our understanding of their impossibility (Weisberg 2013, 126–134). Second, when a model has a target, grounding modeling in similarity makes success too easy. Since any two items are similar in some respect, and each is maximally similar to itself, every item qualifies as a model of every item. Such ubiquity renders similarity epistemically inert. Moreover, if similarity suffices, the ubiquity of similarity makes it hard to see how a model can mislead. Accounts that ground modeling in isomorphism, homomorphism, and the like restrict the range of similarity to structural similarity (see Bartels 2006). Still, the same failings apply. They cannot accommodate models without targets, and too easily succeed if targets are available.

Giere (1988) attempts to evade the problem of easy success by maintaining that a successful model is similar to its target in relevant respects; irrelevant similarities are idle. However, problems remain. First, what we might call "accidental matching" is possible. A model designed to resemble its target in a specified, relevant respect fails to do so but happens to resemble it in unspecified, perhaps undiscerned, but nevertheless relevant respects. Since similarity is ubiquitous, this is a likely scenario. Second, a rococo model might include so much irrelevant information that it occludes relevant similarities. In that case, a

 DOI: 10.4324/9781003205647-26

relevant similarity obtains but is swamped by irrelevancies. This is a problem of too much information. Models streamline and simplify. They advance understanding by omitting what should be ignored. Moreover, models distort. Rice (2021) argues that some models are effective not despite but because of their pervasive, drastic distortions. If so, then however circumscribed, similarity seems the wrong metric.

Similarity approaches, whether selective or not, apparently assume that once we establish that a model stands in the proper relation to its target, the way the model affords understanding of its target will be evident. Models are construed as mirrors—as intentional replicas that reflect a portion of reality. But to a large extent, the value of models lies in their being vehicles for surrogative reasoning. Reasoning with and about a model enables scientists to better understand its target. An effective model fosters and facilitates epistemically fruitful surrogative reasoning. Rococo models include obtrusive, irrelevant information that impedes effective reasoning. Excessively simple models display relevant similarities but fail to facilitate reasoning. Both, however, mirror the phenomena they concern.

Advocates of relevant similarity might try to accommodate this by incorporating considerations pertaining to effective surrogative reasoning into the criteria on relevant similarity. Still, there is a problem. We resort to surrogative reasoning because reasoning directly about the target is too difficult, cumbersome, or time-consuming. The target is too obscure, too complex, entwined with confounding factors, mathematically intractable, or whatever. For a model to be an effective vehicle for surrogative reasoning, it must be suitably and often substantially dissimilar to the target.

A separate issue concerns the selection of reasoning strategies. Enabling the same reasoning we use when we directly confront the target is ill-advised. There is no basis for thinking that reasoning appropriate to the full complexity of the phenomena is equally appropriate when things are pared down. What sorts of reasoning are to be permitted? There need not be a one-size-fits-all answer to this question. But in any particular case, it should be clear what inferences are permissible. Is abductive inference allowed? Is analogical reasoning? Nothing about the similarity of a model to its target, or the structural relations between a model and its target settles, or even addresses, this issue.

We have uncovered a number of features that an adequate account of the epistemic function of models should accommodate. (1) Some models have no targets. Still, they seem to function epistemically. (2) Models can be ineffective because they provide too much information, even if that information is accurate. (3) Effective models distort in ways that are illuminating, not misleading. (4) Some models mislead. An adequate account of models should explain how such models impede understanding or foster misunderstanding. (5) Models are used for surrogative reasoning. Hence an adequate account should explain how the reasoning they promote figures in or advances understanding.

Models are not mirrors; nor are they transparent windows to the world. They are complex symbols whose epistemic contributions derive from multiple interacting symbolic functions. As symbols, they are subject to syntactic, semantic, and pragmatic constraints. They are artifacts—epistemic tools that equip us to understand the world in ways that otherwise we could not (see Knuuttila 2011). Drawing on Goodman (1968), the following sections begin by explicating a number of symbolic devices that figure in Hughes's DDI account (1997). This account will then be presented and extended. It will also be shown how the extended account satisfies the requirements listed above, and how models so construed embody and advance understanding.

2. Symbolic resources

Denotation is the relation of a name to its bearer, a predicate to the items in its extension, a portrait to its subject, a map to the terrain it maps, and a general picture (such as the picture of a sparrow in a field guide) to each of the things it depicts. If a model has a target, it denotes that target. If a symbol has no object, it does not denote. Thus, fictional names, such as 'Huck Finn', and fictional maps, such as the map of the route to Mordor, fail to denote. So do terms like 'phlogiston', 'the ether', and 'the Northwest Passage', which were once thought to denote but turned out to have no objects. Nevertheless, such symbols are not gibberish; nor are they parts of speech like prepositions or adverbs that play distinct, non-denotative grammatical roles. Despite having no referent, they are denoting symbols. 'Huck Finn' remains a name, even though it names no one. It is a denoting term. A model of the ether remains a model, even though it models nothing. It too is a denoting symbol. A critical question is how such symbols function.

Goodman distinguishes two dimensions along which a denoting symbol functions. One is the relation of the symbol to what it is a symbol *of* 'Feynman' denotes a particular physicist. The name is a representation of a particular person. The other dimension concerns what sort of representation it is. To mark the difference, Goodman introduces the concept of a *p-representation* (1968, 127–131). The term 'Feynman' is a symbol of the sort that is capable of denoting Feynman. It is a Feynman-representation, a physicist-representation, and so on. The formula 'O_2' is the sort of symbol capable of denoting oxygen. It is an oxygen-representation, a gas-representation, and so on. 'Representation of' is a two-place relation linking a representation with its object. Where its denotation is null, the symbol is not a symbol *of* anything. Still, such a symbol is of the same syntactic sort as symbols that successfully denote. Its grammar makes it capable of denoting. *'p-representation'* is a schema for a one-place predicate whose members all have the same putative object. It is a classification of denoting symbols themselves, without regard to what, if anything, they denote. In contextually relevant circumstances, 'Huck Finn' is a Huck-Finn-representation and a runaway-boy-description, just as 'Richard Feynman' is a Feynman-representation and an-expert-in-quantum-mechanics-representation. 'Phlogiston' is a phlogiston-representation and a-source-of-combustion-representation, just as 'oxygen' is an oxygen-representation and a-sustainer-of-combustion-representation. What qualifies various symbols to be members of the same class of *p*-representations is their relations to one another, not their relation, if any, to a denoted object.

Through the device of *p*-representation, we see how multiple representations qualify as being of the same putative item. A variety of terms in a novel coalesce to constitute a fictional character like Huck Finn, a variety of symbols in biology papers coalesce to characterize a fictional species with four sexes, a variety of nouns and pronouns, descriptions, and names in a factual report coalesce to characterize an actual avalanche. By being instances of the same *p*-representation, distinct terms and distinct uses of the same term count as being about the same real or ostensible item. The various instances of the same *p*-representation constitute a small genre—the genre of Hobbit-representations, phlogiston-representations, four-sex-species-representations, avalanche-representations (see Elgin 2010, 3). Over time, the genre evolves, as increasingly numerous and varied representations become recognized as members of a given class of *p*-representations. Thus, there were increasingly detailed phlogiston-models even though they turned out not to be models of anything. *p*-representations enable us to understand both why targetless models are representational, and how hypothetical representations function. At the outset, we may not know whether anything answers to a given symbol—that is, whether the symbol denotes anything. The putative item begins its career as a posit. To figure out whether anything answers to the posit requires elaboration—endowing it with more robust characteristics, incorporating it into models, and seeing what happens. The posit acquires a distinctive profile as it is elaborated, and increasingly detailed commitments are incorporated into it. The genre evolves over time, homing in on what it would take for something to constitute an answer to the posit—that is, what would be required for the symbol to denote. Elaborating a model that involves a posit then is a matter of extending, refining, and emending the *p*-representations that collectively come to constitute the identity conditions on the posited object.

Denotation and *p*-representation underwrite *representation-as*. For Winston Churchill to be represented as a bulldog is for a bulldog-representation to denote Churchill. For a nucleus to be represented as a liquid drop is for a liquid-drop-representation to denote the nucleus. Denotation can be affected by stipulation. A user can simply stipulate that *m* shall denote *n*, and *m* thereby comes to denote *n*. So any *p*-representation can, by stipulation, be used to denote any object. A bulldog-representation could represent the nucleus as a bulldog, and a liquid-drop-representation could represent Churchill as a liquid drop. If representation-as is to serve as a vehicle for modeling, further restrictions are required to exclude unwanted cases. This is where exemplification enters the picture.

Exemplification is a mode of reference by which an item refers to some of its own features, a feature being a property or relation at any level of abstraction. Exemplification involves both denotation and instantiation. For a symbol *s* to exemplify feature *t, s* must instantiate *t* and must refer to *t* via that instantiation (see Goodman 1968, 50–68, Vermeulen, Brun and Baumberger 2009). Commercial paint companies provide sample cards that exemplify the colors of the paints they sell. Problems worked out in textbooks exemplify the reasoning strategies students are expected to learn. Exemplars are not mere instances of features; they are telling instances. They highlight select features, making them manifest. Some exemplars, such as textbook cases and displays on paint cards, are highly regimented. Others are ad hoc. Anything can serve as a sample of any of its features, simply by being used as such. An ornithologist might point to a bird as an example of a goldfinch; if it is in fact a goldfinch, that bird comes to exemplify its species. It was, of course, a goldfinch anyway. What the ornithologist's gesture did was make it an example of its kind. Nor is it the case that exemplification is simply a vehicle for conveying what is already known. The chef samples the soup to see whether it needs more sage. Until he tastes it, no one knows. He is not especially interested in whether that particular spoonful of soup needs more sage. He treats the spoonful as a representative sample of soup in the pot it was drawn from. He draws inferences about the rest of the soup from what is exemplified by the spoonful he tastes.

Exemplars may require processing to bring the features they exemplify to the fore. Merely looking or tasting is not always enough to ascertain what an exemplar exemplifies. Like the chef, the mining inspector takes samples to exemplify something no one yet knows—in this case, the proportion of different gases at different levels of the mine. But unlike the chef who can trust his senses, the inspector needs to run his samples through a gas chromatograph to determine what the samples exemplify.

Exemplification is selective. To highlight some of an item's features requires bracketing, downplaying, or marginalizing others. In its standard use, a paint card exemplifies the colors on its face. It does not exemplify its position. In a non-standard use—for example, when used as a bookmark—the card exemplifies a place in a book, disregarding color completely. In the sciences, processing often requires more than a reorientation of emphasis. It often involves removing confounding factors. Then scientists work with pure samples rather than relying on what is found in nature. Processing may involve adding reagents to bring a particular feature to the fore or subjecting an item to extreme conditions, in order to highlight features that are not manifest in standard conditions. Experimentation is in large measure a matter of enabling items to exemplify features that are not ordinarily epistemically accessible.

Exemplification is not a matter of conspicuousness. To exemplify subtle factors, conspicuous features often need to be bracketed. A risk assessor may find that a manufacturing process exemplifies a subtle vulnerability to sabotage. To do so, he ignores the deafening din in the factory and the firm's annual production figures. Figuring out how to extract epistemically valuable information requires determining which aspects of the phenomena are significant, and which are irrelevant. Clearly, this is a contextual matter. Depending on the issue under investigation, and the conceptual, instrumental, and methodological resources available, the same phenomena can be interpreted as exemplifying any of a variety of features. What is a signal in one investigation may be noise in another.

In principle, an item can exemplify any of its features. But not all features are easily exemplified. Some are semantically unmarked; we have no readily available labels for them. When this is so, it may be far from obvious how far the exemplified feature extends. Even when a feature is semantically marked, the way it is represented may be unintelligible to those who seek to access it. Innovation is needed to bring it to the fore. In January 1986, the Challenger space shuttle exploded because its O-rings failed to seal due to the low temperature at the launch site. Hearings were held during which scientists presented myriads of relevant information. The Congressmen conducting the hearings did not understand the scientists' charts, graphs, equations, and explanations. Then Richard Feynman dropped an O-ring into a glass of ice water and showed that it became brittle in the cold (Feynman 2001, 146–153). His demonstration exemplified to scientific novices what the more learned explanations could not effectively convey. It displayed the connection between the temperature, the resulting brittleness of O-rings, and their inability to expand to form a seal. In this case, the epistemic limitation was only on the side of the lay audience. In other cases, the limitation is general. A situation may be so complicated that no one knows how to handle it in its full complexity. The task then is to exclude irrelevant details in order to focus on telling features. This is one reason we resort to models.

3. Models as symbols

Scientific models are schematic representations that systematically and rigorously omit irrelevancies. They make no pretense of being accurate. I have characterized epistemically effective models as felicitous falsehoods (see Elgin 2017). Some distort. A model representing planets as point masses ignores the breath of each planet and the fact that its mass is not evenly distributed. For certain purposes, such factors are irrelevant. Only the center of gravity and overall mass need to be exemplified. Other models augment. Maxwell's 'idle wheels' are fictional devices that forge an analogy between electromagnetic and mechanical systems, thereby exemplifying an abstract structure that electromagnetic and mechanical systems share (see Nersessian 2008, 19–61). Still others exaggerate. According to Kepler's first law, the Earth travels around the sun in an elliptical orbit. Diagrammatic models typically represent the major axis as considerably longer than the minor axis. In fact, the two axes are almost equal in length. But the models are effective because they exemplify only the property of *being elliptical*, not the precise shape of the ellipse. Statistical models may be true or true enough in the aggregate, but nowhere near true of any particular. Although there are no rational economic agents, irrational idiosyncrasies cancel out. Models that represent populations as infinite elide the effects of chance that finite populations are subject to, exemplifying the role non-random factors play in the behavior of the phenomena.

Patterns emerge when details are excluded. The Lotka–Volterra model is a pair of differential equations that characterize the interdependent dynamics of predator and prey population sizes. It is a simplified model that represents predators as insatiable and prey as immortal unless eaten. By bracketing the question of how the populations modulate their sizes, it reveals a pattern that holds of rabbits and foxes, mollusks and starfish, fish in the Adriatic, and even loan sharks and their victims. The bracketed mechanisms make no difference (see Strevens 2008). The model thus exemplifies a widespread pattern. To be sure, there are limits. The pattern plainly breaks down if the predators drive their prey to extinction. It is considerably more complicated if the predators are themselves prey, if multiple species target the same prey, and so forth. The model thus operates against background assumptions.

Qualms about its epistemic status may persist. The Lotka–Volterra model involves assumptions that are inaccurate. No members of any species are insatiable. No members of any species are immortal unless eaten. So how does a model that describes the population dynamics of such fictional species tell us anything about the dynamics of real populations? The contention that a distortion, simplification, or amplification is not a difference maker at best assures that we make no mistake in resorting to it; this does not yet explain how it advances understanding. To answer that, we need to look in more detail at how models function.

Effective models foster understanding by facilitating fruitful reasoning that illuminates the phenomena they concern. The liberties they take, the divergences from overall accuracy, are justified by their epistemic payoffs. A number of philosophers of science have emphasized that models are things we think with; they are neither windows nor mirrors, but vehicles for surrogative reasoning (see Suárez 2009). Hughes (1997) connects the referential and inferential roles. Drawing on Goodman (1968), he characterizes a model as a complex symbol that performs three interanimating functions: denotation, demonstration, and interpretation. His discussion is schematic. Here it has been elaborated to bring out features that he sketched.

Denotation, as we have seen, is the relation of the model to whatever it is a model *of*. The harmonic oscillator, being a model *of* a spring, denotes a spring; the Phillips–Newlyn machine, being a model *of* an economy, denotes an economy. *Demonstration* consists of reasoning with the model according, as Hughes says, to 'its own internal dynamic'. *Interpretation* consists in identifying the fruits of that reasoning and imputing them to the target. Denotation has already been discussed. Demonstration and interpretation require explication.

The *demonstration* phase of the modeling process is the locus of surrogative reasoning. A model's internal dynamic sets limits on permissible modes of inference, facilitating informative, fruitful, relevant, non-trivial inferences pertaining to its target while impeding misleading, irrelevant, and idle inferences. Just how the fruits of that reasoning pertain to the target depends on how they are interpreted. Before turning to that, more needs to be said about demonstration.

A model's internal dynamic specifies the resources that can permissibly be deployed and the uses to which they can permissibly be put. These resources both facilitate and rein in reasoning. They include background assumptions, auxiliary hypotheses, forms of inference, categories, standards of relevance and precision, and so forth. The recognition that the model is designed to afford epistemic access to a particular target and answer specific questions about that target guides the choice of constraints. Descriptions, inferences, and actions that take reasoning too far afield are sidelined.

'Inference' is construed broadly. In addition to familiar, rigorous modes of inference, a model's internal dynamic may (but need not) license analogical reasoning, associative reasoning, and/or abductive reasoning. It issues more focused licenses as well. It may license simplifications or distortions, such as treating a discrete function as continuous, ignoring or focusing on what happens in the limit, representing finite populations as infinite, or treating huge objects as point masses. It determines the choice of scale and grain. Reasoning according to an internal dynamic involves action as well as deliberation. Using a Phillips–Newlyn machine to figure out the effects of tweaking economic policy requires physically manipulating a flow of water, for it is by seeing how the water flows through the apparatus that one draws conclusions about the flow of money through an economy. Nor are practical inferences solely the province of material models. The internal dynamic of a purely abstract model or of a computer simulation licenses certain actions when particular results are reached. One important action is terminating demonstration—ceasing to draw further inferences. The internal dynamic determines when to stop. A model's internal dynamic thus specifies the range of permissions and prohibitions for reasoning with it.

Chains of inference are, in principle, endless. Further conclusions could always be drawn. Opportunities for inference radiate out in all directions. To properly use a model, we need to know what direction to take in drawing inferences and when to stop. Unrestricted inference licenses would generate a plethora of disparate conclusions, with no obvious way to tell which ones could be legitimately imputed to the target. It follows from $pV = nRT$ that $1 \neq 0$, that either $pV = nRT$ or Shanghai is in Spain, that if $pV = nRT$ then (q or $\sim q$), and so forth. Such inferences, although sound, are idle. The proper use of the model brackets them; it takes them offline. If a model's demonstration phase promoted drawing valid inferences indiscriminately, irrelevant inferences would swamp and likely deflect our thinking. To function as an effective device for surrogative reasoning, a model must block irrelevant and unproductive inferences.

Objects can be described in indefinitely many ways. Most are irrelevant to the purposes for which the model is to be used. So the internal dynamic also constrains representation. It dictates that model-representations are to take a particular form, grain, semantic character, and orientation.

The internal dynamic channels both inference and representation via exemplification. Models are exemplars. Like paint samples, they are designed to make some of their features salient. The features may be monadic or polyadic, static or dynamic, abstract or concrete. By representing a population as infinite, the Hardy–Weinberg model exemplifies the extent to which allele redistribution is insensitive to random fluctuations. By ignoring reproductive mechanisms, the Lotka–Voltera model exemplifies a widespread pattern in predator–prey dynamics. Exemplification, as we have seen, is selective. To highlight some features, an exemplar marginalizes or occludes others.

The inferences that a model's internal dynamic licenses are vehicles of exemplification. They show how changes in one parameter affect changes in others, how a system evolves over time, and how robust or fragile linkages are. They disclose patterns and discrepancies that might otherwise be hard to discern. A model does not exemplify the results of irrelevant inferences; its internal dynamic does not license them. So even when they are logically impeccable, they are idle. By functioning as an exemplar, the model constrains and directs reasoning toward features that can responsibly be imputed to the target. It facilitates relevant, informative inferences while blocking or bracketing irrelevant ones.

In the demonstration phase, features are exemplified only in the model. The molecule-representations in the model-gas-representation are represented as spherical, as perfectly elastic, and more generally as exemplary of the pattern displayed by '$pV = nRT$'. What remains is to link the results to the target.

Interpretation involves identifying the features exemplified in the model's demonstration phase and imputing them and only them to the target. Hughesian interpretation is not literal denotation. We know perfectly well that gas molecules are not spherical. So, in imputing sphericality to the molecules in the target gas—in representing actual gas molecules as spherical—we do not maintain that they really *are* spherical. Rather, we construe actual gas molecules in effect as spheres with distortions. In general, in imputing features of a model to a target, we represent the target as having the features exemplified by the model, distended, distorted, or overlaid by confounding factors. We then ignore the confounds as circumstantially irrelevant.

Frigg and Nguyen are sympathetic to this approach but consider it incomplete (2020, 159–204). Their reservations concern the lack of explicit rules of interpretation. Following Hughes, context and established practice may be allowed to determine how the fruits of demonstration are to be interpreted so as to illuminate the target. Because Frigg and Nguyen favor further regimentation, they have added a key. This yields the DEKI model (DEKI = Denotation, Exemplification, Key, and Interpretation). The key specifies the correlation between the features exemplified in the model and the features of the target. The question is whether the key needs to be separately articulable and specifiable independently of its use. It is doubtful that this is the case. An articulable key may be heuristically valuable for a novice learning to use a certain sort of model, but once a scientist has mastered a particular modeling strategy, it is obvious to her what, and with what precision, results of the demonstration are to be read onto the phenomena. Still, the addition of a key highlights the fact that interpretation is subject to public standards.

A model is designed to make particular features of its target salient. Its effectiveness depends on whether the features it exemplifies illuminate the target, enabling model users to understand the phenomena it bears on. By exemplifying a feature, a model affords epistemic access to it. The model equips users to discern factors that may have been overlooked and to appreciate their significance. $pV = nRT$ exemplifies the relation between temperature, pressure, and volume, omitting any mention of attractive force. If the results of the inferences drawn in the demonstration phase hold up when imputed to the target, we have reason to think that intermolecular forces play no significant role in the thermodynamics of the system we are investigating at the grain at which we are investigating it. We know, of course, that every material object attracts every other. So, we do not conclude from the effectiveness of the model that there is no attraction. Rather, we conclude that for the sort of understanding we seek, at the level of precision that concerns us, for the phenomena that concern us, intermolecular attraction is negligible. It is not a difference-maker. Similarly, representing gas molecules as spherical does no harm. Indeed, it helps. By representing the molecules as spheres, we omit the delicate, dynamic differences in the molecules' actual shapes, which would make calculations intractable and impede our understanding of the interdependence of pressure, temperature, and volume in a gas. The effectiveness of the model lies in its being fruitful to think of the target in terms of the features it exemplifies. A model invites us to think of actual gases as ideal gases with distortions, of springs as harmonic oscillators with friction as a confounding factor, of investors as rational economic agents with (perhaps irrational but anyway irrelevant) quirks, and so forth.

Because models omit, distend, distort, and amend, they are context- and purpose-relative. An inaccuracy that is illuminating in one context or for one purpose may be misleading in or for another. A psychologist interested in why people are drawn to conspiracy theories would not represent her subjects as rational agents. Such a model would obscure the very features that she sought to investigate. Devising an appropriate model requires recognizing what factors are and what factors are not likely to be difference-makers for the question one is investigating. Figuring this out may be an iterative process where models with a variety of internal dynamics are tested against one another. To use a model correctly requires understanding how it functions—what phenomena it denotes, what range of features it can exemplify, what modes of inference it licenses, what sorts of features it imputes, what assumptions it makes, and what scaffolding it relies on.

Models distort (see Rice 2021). When they are effective, the distortions illuminate. The fact that, for a given range of purposes, it makes no difference that gas molecules are not spherical reveals something significant about gases. Illumination may be indirect. An effective species-with-four-sexes model exemplifies allele distributions that differ in specific, significant ways from the allele distributions found in otherwise-similar two-sex species. Scientists can discover something important about an actual case by investigating a suitably constructed counterfactual.

The very same phenomenon can be modeled in mutually inconsistent ways, each of which is appropriate for a different range of problems. One model represents the nucleus as a rigid shell; another as a liquid drop. A shell model exemplifies features that depend on the stability of nuclides. A liquid drop model exemplifies features that bear on binding energy (see Massimi 2022, 94–110). The selectivity of exemplification explains why the features that the liquid drop model highlights are appropriately absent from the rigid shell model (see Elgin 2017, 249–272). Each facilitates some inferences and blocks others. The question for the user is which, if either, suits her current epistemic purposes. An effective model is a felicitous falsehood. It is false in that it misrepresents features that are non-difference-makers. Its doing so enables it to exemplify features that make a difference. This is what makes it felicitous.

Streamlining is epistemically valuable. The omission of irrelevancies figures in a model's capacity to advance understanding of its target. Strevens argues that it is permissible to omit these (irrelevant) factors since they are not difference-makers (2008). However, in omitting these factors, models exemplify something about the phenomena that we otherwise would not, or not easily, appreciate.

Models figure in the understanding of a range of phenomena when it is epistemically fruitful to represent the phenomena as if they had the features the model imputes to them, whereas something is epistemically fruitful only if it either fosters or challenges the integration of the behavior of the phenomena into our evolving understanding of the world. For example, because it is as if the traffic on the highway was a continuous fluid, we can use fluid flow models to understand the movements of traffic. The model explains why the traffic flows more smoothly in the center lanes than at the edges of the road.

It makes no difference that, rather than actually being a continuous fluid, the traffic consists of discrete cars.

Every object has indefinitely many properties and stands in indefinitely many relations to other things. The vast majority of these are of no interest. Some of the interesting and important ones are neatly labeled by our literal vocabulary. They can be directly and literally represented. Others are semantically unmarked. There is, for example, no term capable of accurately describing the exact shape of a carbon dioxide molecule. If properties and relations that lack literal labels are to be recognized, they need to be indicated indirectly. One way to do so is by characterizing the objects that display them *as-if-ishly* (see Vaihinger 2009). It is as if gas molecules were spheres, or as if predators were insatiable, or as if the moon were falling toward the earth. Such as-if-ish representations can capture something epistemically important. The reason is not just that it won't be wrong in a particular context to think of gas molecules as spherical or predators as insatiable or the moon as a falling body; the important point is that *the fact that it won't be wrong* discloses something significant about the phenomenon. The effectiveness of the model discloses that a particular aspect of things—for example, the molecule's shape being somewhat spherical—is significant. The model then provides emphasis and focus. It affords insight not only into what properties the object has but also into which of its properties are worth registering.

4. Conclusion

The account of models presented satisfies the requirements set out above. Models without targets are bereft of denotation. Ether-models are not models of the ether because 'ether' turns out to fail to denote. 'Four-sex-species-models' are not a model of a species with four sexes because 'four-sex-species' fails to denote. Scientists once thought 'ether' denoted; they were wrong. They never thought 'four-sex-species' had a non-null denotation; there was no mistake. In both cases, however, reasoning in the demonstration phase can be carried out. The models have their own internal dynamics, which constrain and channel reasoning, enabling scientists to explore the implications of the items they posit. They investigate what would happen if items of the sort posited behaved in the ways the dynamic mandates. Since 'what would happen if…?' is often a good question, models without targets are often epistemically valuable.

Because exemplification is selective, it enables us to evade the problem of too much information. An enormously complicated phenomenon can be idealized, bracketing the information that makes no difference to the question being examined. So an effective model excludes irrelevancies and focuses on what, in a given context, is significant.

Although models simplify, amplify, streamline, and distort, they illuminate their targets when the features they exemplify can be imputed to their targets in such a way that the problems at issue can be fruitfully addressed. When the effects of intermolecular attraction are negligible, a model that sets them aside enables scientists to appreciate the interrelation of pressure, temperature, and volume in an actual gas. When, however, they are non-negligible, $pV = nRT$ misleads. Misleadingness can take different forms. If intermolecular forces are significant, $pV = nRT$ can be imputed to the target, but its imputation does not supply enough relevant information to be useful. The result is an interpretation that is unacceptably sparse. It incorrectly suggests that no additional information is required. If a model is just irrelevant, imputation simply fails. A population of mice cannot plausibly be represented as an ideal gas. There is no non-arbitrary way to impute the pattern exemplified in $pV = nRT$ to the mice. A misleading use of a model exemplifies features that cannot be fruitfully imputed to its target. A misleading model of a given target exemplifies features that cannot plausibly be imputed to the target at all. Still, such a model, construed as targetless or imputed to a different target, would not necessarily mislead. Whether a model is misleading then depends on how it is used.

This chapter began by saying that modeling is a powerful epistemic tool. The power lies in its ability to simultaneously generate representations that afford focus and show why that focus (even when provided as-if-ishly) is valuable. In effect, models not only say, 'This is what you should be looking at', they also say, 'This is why you should be looking at it this way and ignoring factors that interfere with doing so.' They thereby extend our epistemic range.

Acknowledgment

This chapter is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2021 research and innovation program (Grant Agreement No. 101043071) with a project on Model Transfer and Its Challenges in Science: The Case of Economics.

Note

1 Weisberg discusses three-sex models. As it turns out, there are species that have three sexes. Since the point concerns the epistemic value of targetless models, I changed the number to four. Regardless of the number (n) of sexes actual species have, it is fruitful to be able to consider how alleles would redistribute if there were (n + 1) sexes. Such a targetless model can be informative.

References

Bartels, Andreas. 2006. "Defending the Structural Concept of Representation." *Theoria* 21(1): 7–19.

Elgin, Catherine. 2010. "Telling Instances." In *Beyond Mimesis and Convention: Representation in Art and Science,* edited by Roman Frigg and Matthew Hunter. Berlin: Springer, 1–18.

———. 2017. *True Enough*. Cambridge: MIT Press.

Feynman, Richard. 2001. *What Do You Care What Other People Think?* As told to Ralph Leighton, New York: W.W. Norton.

Frigg, Roman, and James Nguyen. 2020. *Modelling Nature: An Opinionated Introduction to Scientific Representation*. Cham: Springer.

Giere, Ronald. 1988. *Explaining Science: A Cognitivist Approach*. Chicago: University of Chicago Press.

Goodman, Nelson. 1968. *Languages of Art*. Indianapolis: Hackett Publishing Co.

Hughes, R.I.G. 1997. "Models and Representation." *Philosophy of Science* 64: S325–S336.

Knuuttila, Tarja. 2011. "Modelling and Representing: An Artefactual-Based Approach to Model Based Representation." *Studies in the History and Philosophy of Science* 42(2): 262–271.

Massimi, Michaela. 2022. *Perspectival Realism.* Oxford: Oxford University Press.

Nersessian, Nancy. 2008. *Creating Scientific Concepts*. Cambridge: MIT Press.

Rice, Collin. 2021. *Leveraging Distortion.* Cambridge: MIT Press.

Strevens, Michael. 2008. *Depth*. Cambridge: Harvard University Press.

Suárez, Mauricio. 2009. *Fictions in Science. Philosophical Essays on Modelling and Idealization*. London: Routledge.

Vermeulen, Inga, Georg Brun, and Christoph Baumberger. 2009. "Five Ways of (Not) Defining Exemplification." In *Nelson Goodman. From Logic to Art*, edited by Gerhard Ernst, Jakob Steinbrenner and Oliver R. Scholz. Frankfurt a.M.: Ontos. 219–250.

Vaihinger, Hans. 2009. *The Philosophy of As-If*. Mansfield: Martino Publishing Co.

Weisberg, Michael. 2013. *Simulation and Similarity.* Oxford: Oxford University Press.