Representation at the Nanoscale

Otávio Bueno[†‡]

In this paper, I provide an account of scientific representation that makes sense of the notion both at the nanoscale and at the quantum level: the partial mappings account. The account offers an extension of a proposal developed by R. I. G. Hughes in terms of denotation, demonstration, and interpretation (DDI). I first argue that the DDI account needs some amendments to accommodate representation of nano and quantum phenomena. I then introduce a generalized framework with the notions of unsharp denotation, unsharp interpretation, and partial mappings. I conclude by arguing that the resulting view not only accommodates nano and quantum representations but also makes sense of representation by imaging in (scanning tunneling) microscopy.

1. Introduction. Scientific models are often taken to be representational devices: they are used to represent the phenomena. But in what sense do they represent? We can say that scientists select certain features of the phenomena to be described in the model, specify appropriate relations between the model and the phenomena, and provide a reading of the model in which it's clear in virtue of which features of the model those relations hold. Bas van Fraassen and Jill Sigman have provided an insightful general format for representation:

> Representation of an object involves producing *another object* which is *intentionally* related to the first by a certain *coding convention* which determines what counts as *similar* in the *right way*. (van Fraassen and Sigman 1993, 74; italics added)

Following this proposal, we can say that an account of scientific representation involves three basic components: (i) an *intentional act*, (ii) a *coding convention*, and (iii) a *mechanism of representation*. The intentional

act relates two objects: the *target*, which is the object we aim to represent (e.g., a given phenomenon), and the *source*, the object that is used to represent (e.g., a model). The coding convention specifies how to read the relationship between the source and the target so that the right similarity between them is established. Finally, the mechanism of representation determines the way in which the target is being represented. This involves (a) representation *as*, where the target is represented *as* such and such (e.g., the phenomena are represented as being continuous), and (b) representation *in virtue of*, where the target is represented in virtue of certain features of the source (e.g., the phenomena are represented as continuous given the existence of a continuous function in the model).

In this paper, I examine the issue of scientific representation at the nanoscale, that is, of objects in the 1–100 nanometer (nm) range. In particular, I discuss whether there is something *special* about the notion of representation at that scale. I then develop an account of representation for scientific theories and indicate, according to that account, what is special about representation of nanoscale objects.

The account I offer—the partial mappings account—generalizes R. I. G. Hughes' DDI account of representation, which is articulated in terms of denotation, demonstration, and interpretation (Hughes 1997). By introducing the notions of unsharp denotation, unsharp interpretation, and partial mappings, I argue that some limitations faced by the DDI account can be overcome. But the partial mappings account also has an independent motivation: it yields a perfectly natural account of representation by imaging in microscopy. To support this claim, I indicate how that account accommodates the way in which a scanning tunneling microscope (STM) represents.

2. The Nanoscale. As is well known, the nanoscale deals with objects in the 10^{-7}–10^{-9} meter scale (1–100 nm scale). Nanoscience is then the study of basic properties of matter on that scale. Of particular interest are the cases in which these properties *differ* from the properties of bulk matter and atoms. These differences emerge in two crucial ways: as the result of (a) changes in *shape* or (b) changes in *size*. With regard to (a), the chemical reactivity of nanomaterials, such as gold nanorods, is *shape* dependent (Caswell et al. 2003). In fact, various studies have been developed of the way in which the shape of nanoparticles affects their chemical properties. For example, *cubic* gold nanoparticles (on an NaCl substrate) spontaneously evaporate, whereas particles with a *spherical* morphology don't. With regard to (b), nanoparticles are highly sensitive to *size* variation. The reactivity of a metal nanoparticle toward oxidation increases with the decrease in the particle's size. As a result, noble metals end up becoming very reactive in the nanoscale (Jana et al. 2002). Although we are

dealing with very small objects, the nanoscale ranges over a huge domain. Objects with substantially different properties are found in the 1–100 nm scale. For example, on the larger side of that scale, which includes objects on the 100 nm range, we find "things" as diverse as individual components on an Intel Pentium III processor chip, wavelengths of visible light, very large molecules, and bumps on a compact disk. If we consider smaller objects, such as those included on the 10 nm scale, we find viruses, such as the adenovirus responsible for the common cold, and medium-sized molecules. Finally, at the smallest side of the nanometer scale, when we consider objects on the 1 nm scale, we have things such as DNA and medium-small molecules. This is the smallest we can get, or so the story goes, before we obtain quantum effects. So, if we consider things on the 0.1 nm scale, not only are we then outside the nanoscale, but at this level we also have atoms, wavelengths of x-rays, and all the bizarre behavior of the quantum world. This is indeed often taken to be a significant feature of the nanoscale: the existence of small, fairly robust structures, such as molecules and DNA, which need not be subject to quantum effects. As I will indicate below, things are not so simple. But to argue for this point, we need first to consider how to represent nanophenomena.

3. Representation by Denotation. As noted, according to the DDI account, representation involves three major components: denotation, demonstration, and interpretation (Hughes 1997). First, a model is specified, and certain features of the phenomena are *denoted* in the model. Using the mathematical techniques available in the model, one draws new results (*demonstration*). These results are then *interpreted* back in the phenomena (Figure 1). As an illustration, let's consider an example that doesn't bear on the nanoscale (Hughes 1997, S326–S329).

In the third day of his *Discourses Concerning Two New Sciences*, Galileo examines the issue of naturally accelerated motions (Galilei [1638] 1974). And he establishes a result that relates the distance traveled by two objects during the same time: one is a uniformly accelerating object starting from rest; the other is an object moving with uniform speed. According to Galileo, the two distances will be the same if the final speed of the accelerating object is twice the uniform speed of the other. He establishes the conclusion using a geometrical diagram. His strategy is to start with a problem in physics and represent it geometrically. He then solves the geometrical problem and reads off the solution (of the physical problem) from the geometrical representation (Hughes 1997, S327).

The three components of the DDI account are clearly found in Galileo's reasoning: (i) Time intervals are *denoted* by distances along a vertical axis in the diagram, increases of speed by lengths of horizontal lines. (ii) *Demonstrations* (proofs) are formulated in the geometrical model: a geomet-

Phenomena   **Model**
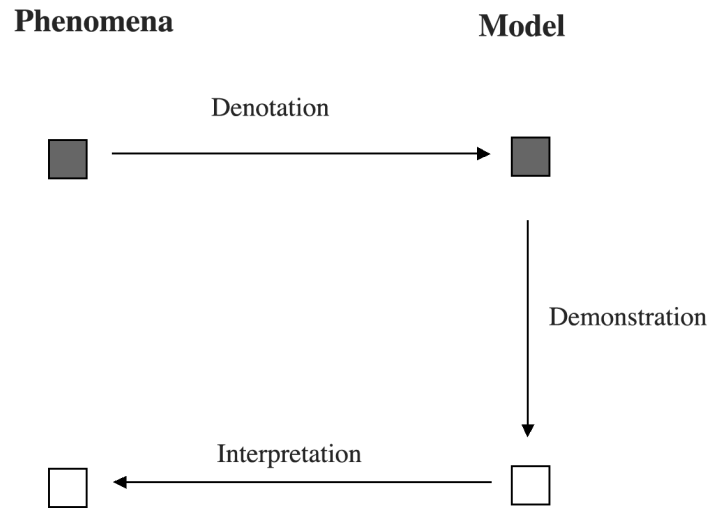
Denotation

Demonstration

Interpretation

Figure 1. The DDI account of representation.

rical reasoning establishes the equality of two areas in the diagram. (iii) The established result is then *interpreted* kinematically: the equality of the two areas is associated with the equality of the distances traveled.

This is an elegant account of scientific representation. Being very general, it should also be able to accommodate the vagaries of the latter notion. But does it?

4. Two Difficulties: When Denotation Fails and the Lack of a Mechanism of Representation. Despite its significant features, the DDI account faces two difficulties. First, it's unclear that it provides an explicit mechanism of representation. The account yields a general framework, a *scheme*, to address the issue, but not a specific proposal. After all, representation is achieved, according to the account, by denotation, demonstration, and interpretation, but no particular characterization of these notions is put forward. The only constraint is that denotation and interpretation should be different from each other (Hughes 1997, S333). As a result, the account doesn't specify (a) how the phenomena can be *represented as* such and such and (b) *in virtue of* which features of the model the representation is achieved. Thus, it's unclear how the mechanism of representation can be articulated.

In response, it might be argued that there is a well-established account of denotation (and interpretation), and that account provides a mechanism of representation for the DDI view. The account is the standard

model-theoretic conception of denotation. According to this conception, denotation is formulated in terms of a mapping from objects in a given domain to certain elements of a model. Typically, the elements in question are linguistic terms of the language in which the *model* is formulated, and the objects that are denoted are elements of a given set (e.g., the set of the *phenomena*). (A similar account can be provided of the notion of interpretation, which would essentially be a reverse mapping from the model to the phenomena.)

By means of this account of denotation and interpretation, a mechanism of representation for the DDI view is immediately provided: the relevant mappings do the representational work. In other words, representation is ultimately a matter of selecting certain features of the phenomena and *mapping* them into the model. When (mathematical) consequences are drawn from the model, new results are then generated. These results are, in turn, mapped back into the phenomena. The particular choice of the mappings specifies how the phenomena are being represented (representation *as*), and *in virtue of* which features of the model the representation is achieved (representation *in virtue of*).

But this response generates the second difficulty: How sensitive are the notions of denotation and interpretation, understood model-theoretically, to the *size* of the objects under study? Do these notions still apply when we consider increasingly small nanoparticles? The difficulty becomes particularly pressing in the context of objects whose properties are highly sensitive to variations of size and shape, such as nanoscale objects. Schrödinger has pressed a related point very vividly:

> As our mental eye penetrates into smaller and smaller distances and shorter and shorter times, we find nature behaving so entirely differently from what we observe in visible and palpable bodies of our surroundings that *no* model shaped after our large-scale experiences can ever be true. (Schrödinger 1952, 17)

As Schrödinger notes, a crucial difference between macroscopic and microscopic phenomena—that emerges from their distinct sizes—is that the concept of *identity lacks any sense with regard to elementary particles* (17–18). We may be able to count such particles, but we can't individuate them. Swapping electrons, for instance, typically doesn't change the state of the quantum system the electrons are in. Even though we may know that there are a definite number of electrons in a certain region, we are unable to distinguish them as individuals (Krause and French 1995).

The trouble for the DDI account is that without identity—without being able to individuate the objects we want to refer to—there's no *denotation*. We can't denote an object whose identity conditions are *not* defined for the simple reason that it is not clear which object we are referring to.

Thus, for objects without well-defined identity conditions, the model-theoretic notion of denotation cannot be applied. And without denotation, it's unclear how we could apply the DDI account. Does this mean that quantum physicists adopt a *different* notion of representation?

But wait! We were supposed to be talking about the *nanoscale*, and I suddenly raised a potential problem for the DDI account that relies on a particular interpretation of *quantum mechanics*—in fact, a particular "metaphysical package" for the theory (Krause and French 1995). But the quantum level, as noted above, deals with objects of a much "smaller" size than those at the nanoscale. And perhaps no difficulty about denotation arises for the latter. Unfortunately, things are not so simple. If we consider the small part of the nanoscale—the one that deals with nanoparticles in the 1–10 nm size range—it will become clear that quantum effects emerge. Not surprisingly, a particular group of nanoparticles for which these effects are detected are called *quantum dots*, and it's worth considering them here (Murphy and Coffer 2002).

As is well known, crystalline inorganic solids are divided into three groups: *metals*, *semiconductors*, and *insulators*. All of them have bands, which are nearly continuous energy levels. Metals have a partially filled band. Semiconductors have a filled band (called the *valence band*), which is separated from the mostly empty *conduction band* by a bandgap E_g. Finally, insulators have the same electronic structure as semiconductors, except that they have a larger bandgap E_g. This means that, in the case of insulators, more energy is required to move, say, an electron from the valence band to the conduction band. And that's why insulators are not very good energy conductors. *Quantum dots* are semiconductors with all three dimensions in the 1–10 nm size range.

Murphy and Coffer (2002, 17A) consider what happens to such a semiconductor when it's irradiated with light of energy $hv > E_g$. In this case, they note:

> An electron will be promoted from the valence band to the conduction band, leaving a positively charged "hole" behind. This hole can be *thought of* as the absence of an electron and acts as a particle with its own effective mass and charge in the solid. One can *calculate* the spatial separation of the electron and its hole (an "exciton") using a modified *Bohr model*:
>
> $$r = \epsilon h^2 / \pi m_t e^2$$
>
> where r is the radius of the sphere defined by the three-dimensional separation of the electron-hole pair; ϵ is the dielectric constant of the semiconductor; m_t is the reduced mass of the electron-hole pair; h is Planck's constant; and e is the charge on the electron. (Italics added)

After setting up the model, Murphy and Coffer are then in a position to draw some significant conclusions:

> For typical semiconductor dielectric constants, the *calculation* suggests that the electron-hole pair spatial separation is ∼1–10 nm for most semiconductors. In this size range, when the exciton is created, the physical dimensions of the particle confine the exciton in a manner *similar to the particle-in-a-box problem of physical chemistry*. Therefore, the *quantum effects such as quantization of energy levels can be observed in principle*. (17A; italics added)

In other words, despite being completely located in the nanoscale, quantum dots are subject to quantum effects—and this includes the unexpected behavior of the identity of some of their parts. Moreover, as the passage above makes it clear, the *description* of these nanoparticles crucially depends on the erratic behavior of electrons. The whole model is set up in terms of the latter.

There are several significant features in this piece of modeling. Note, to begin with, that Murphy and Coffer are clearly representing relevant features of quantum dots by *mapping* certain aspects of the *phenomena* (namely, the behavior of semiconductors that have been irradiated with light of a certain energy level) into a given *model* (a modified Bohr model). Our authors then *draw results* from the model (by calculating, e.g., the electron-hole pair spatial separation), and they *interpret* the result back into the phenomena (e.g., by expecting the detection of quantum effects). Clearly, we have here the broad features of the DDI account at work: mappings from the phenomena to the model, demonstrations in the model, and mappings from the results obtained back into the phenomena.

However, although such mappings are invoked, clearly they are *not* denotations and interpretations—in the model-theoretic sense. How could they be? If, as noted above, according to a given interpretation of quantum mechanics, electrons have no well-defined identity conditions, we couldn't establish *denotations* from electrons to the model, nor could we provide *interpretations* from the model back to electrons. But this seems to be exactly what Murphy and Coffer are doing! What is going on then?

One could claim that given that Murphy and Coffer seem to be denoting electrons, it follows that they are *not* adopting a Schrödinger type nonindividuality interpretation of quantum mechanics. But I don't think this is right. As I argue below, it's possible to make perfect sense of what Murphy and Coffer are doing *even if* we suppose that they have adopted a nonindividuality interpretation. The point here is that our notion of scientific representation shouldn't rule out a priori the possibility of certain interpretations of quantum mechanics, such as the nonindividuality view—especially if nothing in scientific practice has excluded them.

5. Unsharp Mappings. Noting the way in which Murphy and Coffer refer to electrons makes it clear how to accommodate the notion of *denotation* at the nanoscale and at the quantum level. What we need is a *broader notion of denotation* (and interpretation). Denotation is definitely a mapping. This much is right with the model-theoretic account. But the mapping can be *unsharp*, in that it assigns certain features of the phenomena to the model *without* the requirement that the objects in question have well-defined identity conditions. More formally, the mapping is *unsharp* given that instead of being defined for each *individual* element of its domain, it's defined for *equivalence classes* of elements of the domain. And having such an *unsharp mapping* is enough to connect the phenomena to the model. Note Murphy and Coffer's description: "*An* electron will be promoted from the valence band to the conduction band, leaving a positively charged 'hole' behind" (2002, 17A). At no point do the authors identify *the* electron they are referring to—*any* electron will do. And they continue: "This hole can be thought of as the absence of *an* electron and acts as a particle with its own effective mass and charge in the solid." Here, again, identity conditions for electrons are never asserted, nor is it even presupposed that electrons have identity conditions. Instead of a precise denotation, all that is needed is reference to an element of a particular class, independently of the identity of this element. And that's precisely what reference to an equivalence class does—in this case, the equivalence class of electrons determined by their indistinguishability (Krause and French 1995). In other words, an *unsharp mapping* from equivalence classes of electrons to the Bohr model (*unsharp denotation*) suffices. Similarly, a reverse unsharp mapping from elements of the Bohr model to equivalence classes of electrons (unsharp interpretation) will also do. So, suitably reformulated with the introduction of the notions of *unsharp denotation* and *unsharp interpretation*, the DDI account provides a broad enough account of representation. I'll call the generalized proposal, which includes the unsharp notions and emphasizes the need for suitable mappings between the model and the phenomena, the *partial mappings account*. (To complete the presentation of the account, I'll introduce in the next section the specific notion of a partial mapping.)

In particular, the partial mappings account accommodates representation at both the quantum and the nano levels, given that unsharp denotation and interpretation allow us to cover both domains. Moreover, the partial mappings account also indicates in what sense there is something special about representation at the nanoscale. Representation becomes *contextual*. When one gets closer to the quantum level from the nanoscale (say, with quantum dots), quantum effects are expected, and so the need for unsharp denotation and interpretation emerges. But there's no need for such unsharp notions when we deal with "larger" nano-objects

(e.g., those closer to the 100 nm size range, and those whose behavior is not explicitly described in terms of quantum particles). Thus, the partial mappings account helps to illuminate a significant aspect of the relation between the quantum and the nano levels: the context sensitivity of representation.

6. The Partial Mappings Account at Work: Representation and Imaging. How does the partial mappings account handle the issue of the mechanism of representation? The idea is that representation is ultimately established in terms of selecting *appropriate mappings* between the phenomena and the model. As we'll see, conceptualizing the process of representation in terms of unsharp denotation and unsharp interpretation is crucial at this point, and it has advantages independently of what happens in the context of quantum mechanics.

As noted, the "unsharpness" emerges from the fact that the domains of the denotation and the interpretation functions are characterized by *equivalence classes* of certain objects rather than by individual objects. But there is an additional way in which a mapping can also be *unsharp*: by mapping *partially* only some of the elements (in the equivalence class), and their corresponding relations, from a domain into another. In this case, there is only a partial isomorphism, or an even weaker partial homomorphism, between the source structure and the target structure. That is, there is a full isomorphism (or a full homomorphism) between *only part* of the source structure and the target one (Bueno 1997; Bueno, French, and Ladyman 2002; da Costa and French 2003).

Now, establishing an appropriate partial isomorphism (or a partial homomorphism) between the relevant structures provides a mechanism of representation, since it's in virtue of the (partial) sameness of structure that one can use the structure provided by the source to represent the structure at the target. Of course, given that there are several partial mappings between the source and the target, it's crucial to select which of them is the relevant one. At this point, introducing an intentional notion is crucial. To set up a scientific representation, it's vital to decide *which* of the various (partial) mappings should be selected. This is, clearly, a thoroughly context-dependent issue. Depending on which aspects of the target we are trying to represent, different mappings will be adequate. But, in each case, a particular mechanism of representation is provided: the selected mappings highlight how the phenomena are being represented, and in virtue of which aspects of the model the representation is achieved.

To spell this out and to show how the partial mappings account makes sense of a fundamental aspect of representation at the nanoscale, I'll consider the issue of representation by imaging in microscopy. I'll consider, in particular, the case of the scanning tunneling microscope (STM). In

fact, the partial mappings account accommodates this sort of representation very naturally. Imaging is a clear case in which the mechanism of representation is provided by establishing appropriate mappings (such as a partial isomorphism or a partial homomorphism) between the sample under investigation and the image that is produced. And once again, particularly when we are dealing with the smaller part of the nanoscale, we will face quantum effects. Indeed, the principle upon which the STM rests is fundamentally quantum mechanical: a tunneling effect. As a result, with the presence of quantum effects, the mappings in question will have to be unsharp.

The idea behind the STM is very ingenious (Binnig and Rohrer 1983; Chen 1993). It explores the tunneling effect: given two conducting, or semiconducting, solids that are very close to each other and are separated by vacuum, there is a passage, or tunneling, of electrons between them. Owing to the wavelike properties of electrons, it follows from quantum mechanics that electrons will appear as electron clouds that overflow a little from the surface of each solid. As a result, there is a definite probability that electrons will tunnel through the vacuum. Exploring this effect, the STM probe tip scans the sample's surface. And because the tunnel current between the tip and the surface is kept constant, a constant tip-to-surface distance is maintained. As a result, as the probe moves along the sample, it accompanies the contour of the sample's surface, yielding a two-dimensional image of the topography of the surface. A three-dimensional image can then be produced by assembling a whole group of such scans.

We find here a layer of models and mappings in the small region between the tip of the STM and the final image that is generated. At least four steps are involved: (1) At the first level, when the interaction between the STM tip and the sample occurs, a tunneling effect is established. If the microscope is properly calibrated, this detection should be systematic and reliable. Now, on its own, this doesn't establish the existence of an *object* that is responsible for the production of the effect. With quantum mechanics, though, such a supposition becomes much more plausible. The second step then emerges: (2) There is a conversion (translation) of the information regarding the tunneling effect into *topographic* information about the sample. Note that this also assumes quantum mechanics, which is required to make sense of the tunneling effect. But an additional step is still needed: (3) The *topographic* information about the sample needs to be converted into *visual* information. And this requires certain coding conventions about images; for example, the brightness in the image is a measure of the altitude in the sample. (4) Finally, micrographs are superimposed on the STM image as an aid to determine the geometrical structure of the atoms in the sample. But different geometrical configu-

rations are compatible with the same data (the same STM image). As a result, underdetermination emerges.

If we now return to the general format for representation provided by van Fraassen and Sigman (1993), we can see how the partial mappings account makes sense of representation by STM imaging. Recall that representation involves three basic components: (i) an intentional act, (ii) a coding convention, and (iii) a mechanism of representation. So, the representation provided by an STM is ultimately a matter of intentionally determining a source structure that, given a particular coding convention, is similar to the target structure in a specified way. In particular, in the STM case, the source structure—the image generated by the STM—is intentionally produced so that the topographic features of the target (e.g., the geometrical structure of the atoms in the sample under investigation) are represented by the geometry found in the STM image. Given a coding convention—such as the brightness in the image as a measure of the altitude in the sample—it becomes clear in which way the representation is achieved. The similarity between the STM image and the sample is specified *by construction* (given the way in which the STM image is generated) through the coding convention. This establishes a partial mapping from the sample to the image. This mapping, in turn, provides a mechanism of representation, for it specifies the features of the STM image (the particular configuration of shapes and shades in the image) in virtue of which the representation is obtained. Of course, there isn't a full mapping between every aspect of the sample and the corresponding STM image: only some aspects of the sample are selected for representation (such as the geometrical structure generated by the atoms). The reverse process of (partial) interpretation is articulated with the use of micrographs. They help scientists to study the geometrical configuration of the sample from the STM image. But, as noted, given that multiple micrographs fit the data, we face underdetermination. As a result, also here, there is only a partial mapping (a partial preservation of structure) from the STM image—interpreted via a given micrograph—to the sample.

It should now be clear that to provide an account of representation by imaging in microscopy, it's crucial to have unsharp denotations (to accommodate things such as the tunneling effect) and partial mappings (to provide a mechanism of representation). Hence, it's also possible to make sense of three important features: (a) There are aspects of the sample that are *left out* in the representation process, since the mapping between the sample and the image is only partial; these are the aspects of the sample that are *not* depicted in the STM image. (b) But there are aspects of the sample that *are* partially preserved, since there is some mapping between the sample and the STM image (e.g., regarding the relevant geometrical structure between the atoms). (c) Finally, misrepresentation is still pos-

sible: some aspects of the images are artifacts of the representation, and a plurality of micrographs adequately fit the data. Of course, although this plurality is unlikely ever to be completely removed, good training and appropriate STM imaging techniques allow researchers to identify various artifacts of the representation.

7. Conclusion. As we saw, the partial mappings account not only provides the resources to make sense of representation at the nanoscale and at the quantum level, but also accommodates crucial features of representation by imaging in microscopy. Although there's much more to be said, I hope I said enough to indicate how, properly conceptualized, the partial mappings proposal yields a reasonable account of scientific representation.

REFERENCES

Binnig, Gerd, and Heinrich Rohrer (1983), "Scanning Tunneling Microscopy," *Surface Science* 126: 236–244.

Bueno, Otávio (1997), "Empirical Adequacy: A Partial Structures Approach," *Studies in History and Philosophy of Science* 28: 585–610.

Bueno, Otávio, Steven French, and James Ladyman (2002), "On Representing the Relationship between the Mathematical and the Empirical," *Philosophy of Science* 69: 497–518.

Caswell, K., James Wilson, Uwe Bunz, and Catherine Murphy (2003), "Preferential End-to-End Assembly of Gold Nanorods by Biotin-Streptavidin Connectors," *Journal of the American Chemical Society* 125: 13914–13915.

Chen, Julian (1993), *Introduction to Scanning Tunneling Microscopy*. New York: Oxford University Press.

da Costa, Newton, and Steven French (2003), *Science and Partial Truth*. New York: Oxford University Press.

Galilei, Galileo ([1638] 1974), *Two New Sciences*. Translated by S. Drake. Madison: University of Wisconsin Press.

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

Jana, Nikhil, Latha Gearheart, Sherine Obare, and Catherine Murphy (2002), "Anisotropic Chemical Reactivity of Gold Spheroids and Nanorods," *Langmuir* 18: 922–927.

Krause, Décio, and Steven French (1995), "A Formal Framework for Quantum Non-individuality," *Synthese* 102: 195–214.

Murphy, Catherine, and J. Coffer (2002), "Quantum Dots: A Primer," *Applied Spectroscopy* 56: 16A–27A.

Schrödinger, Erwin (1952), *Science and Humanism*. Cambridge: Cambridge University Press.

van Fraassen, Bas, and Jill Sigman (1993), "Interpretation in Science and in the Arts," in G. Levine (ed.), *Realism and Representation*. Madison: University of Wisconsin Press, 73–99.