

Explaining Understanding: New Essays in Epistemology and the Philosophy of Science ed. Stephen Grimm, Christoph Baumberger, and Sabine Ammon. New York: Routledge, 2016, pp. 76-91.

Exemplification in Understanding¹

Catherine Z. Elgin

Examples are ubiquitous. Philosophers adduce examples to support their theories. Students learn from examples and display their learning by providing examples. If a student purports to understand a theory but can provide no examples of how the theory applies, her claim is at least suspect. Why are examples important? A single example is, after all, statistically insignificant. So, it might seem, the ability to provide a single example should count for virtually nothing. But often it counts for a lot. The reason, I suggest, is that the example *displays* an understanding of the subject. It is not just an instance, it is a telling instance.

When an item serves as a sample or example, it exemplifies: it functions as a symbol that makes reference to some of the properties, patterns, or relations it instantiates (Goodman 1968, Elgin 1996). Let us call anything that exemplifies an *exemplar*, and all of an item's properties, as well as all of the patterns and relations it figures in *its features*. Let us take a tolerant approach to properties, recognizing a property corresponding to each extension an item belongs to, regardless of whether that extension is semantically marked or metaphysically privileged.² A property then is just that which members of an extension share. Patterns and relations receive analogously tolerant treatment. Thus exemplified features may be dynamic or static, monadic or

¹This chapter was made possible through the support of a grant from The Varieties of Understanding Project at Fordham University and The John Templeton Foundation. The opinions expressed here are those of the author and do not necessarily reflect the views of The Varieties of Understanding Project or The John Templeton Foundation.

²See McGowan (2003) on metaphysically privileging.

relational, and may be at any level of generality or abstraction.

Exemplification requires instantiation, so only something that actually has a feature can exemplify that feature. Only a bit of Swiss chocolate can be a sample of Swiss chocolate; a bit of Belgian chocolate, however similar in flavor, is disqualified. But merely instantiating is not enough. The authentic Swiss chocolate I ate yesterday did not exemplify Swiss chocolate or any other feature. It just satisfied my sweet tooth. To exemplify, an item must refer to the feature in question, and must do so by virtue of instantiating that feature.

An exemplar can simultaneously exemplify multiple features. The candy might exemplify the properties of being Swiss chocolate, rich in flavor, and an expensive indulgence. But an exemplar cannot simultaneously exemplify all its features. It can point up, make manifest, display, or convey some of its features only by marginalizing, downplaying, bracketing, or muting others. Even if a sample logic problem is the only problem in the book that mentions a dog, it does not (at least under its standard interpretation) exemplify the property of being the only problem that mentions a dog. In a context designed to exemplify a logical form, the problem downplays its being the sole mention of a dog. Exemplification is selective.

Because an item can in principle exemplify any feature it instantiates, the range of features it has the capacity to exemplify is vast and heterogeneous. Remei Capdevila's discussion of Alighiero Boetti's *Rosso Gilera*, *Rosso Guzzi* brings this out:

Boetti's piece consists of two square, nearly identical panels, one next to the other, whose meager distinction from each other is a slight variation in their red paint, and the raised names and code numbers that identify the different paints, which are inscribed on the panel. Like a color palette, Boetti's piece juxtaposes two different kinds of reds, and in that way it is possible to distinguish between them. In a certain sense, then, the work functions as a paint sample; they exemplify two different synthetic reds, whose commercial codes (60 1232 and 60 1305) and names ("Rosso Guzzi" and "Rosso Gilera"). However the work does not only function as a simple color sample, but exemplifies other properties that a chip of paint in a color palette does not exemplify. "Rosso Guzzi" is the red used to paint Guzzi motorcycles, and "Rosso Gilera" is the one

used for Gilera motorcycles, the two rival Italian motorcycle manufacturers. Put side by side, the two panels not only exemplify a slight difference in color but stand for the divide between passionate advocates of each brand. That is to say, since each kind of red possesses the property of being used to paint a specific kind of motorcycle, they can further exemplify the two different brands, and via a chain of reference, the rivalry between the companies. In addition, since the difference in reds is barely noticeable, the piece can further symbolize the negligible distinction that sustains this rivalry. The illusion of difference is also exemplified by the fact that both paints come from the same manufacturer, as the serial numbers indicate. And the work's symbolization is not exhausted yet. Boetti's piece is made out of synthetic commercial paint intended to lacquer vehicles instead of common fine arts materials, and in that way, the artistic properties of this paint – glossiness, brightness, viscosity, or the dripping left when applying it on the panel – are exemplified. This can also symbolize the introduction of mass production and consumer goods into the art world. (Capdevila: 130-131)

As Capdevila's interpretation shows, a single item can, in the right context, exemplify any and many of its features, enabling the interpreter to forge a variety of epistemically valuable connections across a variety of domains.

Exemplars make the features they exemplify salient. They thus afford epistemic access to those features. This is not a matter of conspicuousness. An effective exemplar may marginalize conspicuous features in order to exemplify subtle, difficult to discern ones. The most conspicuous feature of a manufacturing process may be its din, while in the context of a safety inspection what the process exemplifies is a barely detectable vulnerability to sabotage. The vulnerability to sabotage then stands out, while the noise, although still deafening, fades into the cognitive background of things to be ignored.

Some exemplars belong to regimented systems, others are ad hoc. One of the great advantages of exemplification as a referential device is that we can improvise exemplars at will. Simply adducing something as an example typically makes it one.

Interpretation can be tricky, particularly where there is no regimented system. An ornithologist identifies a bird as an example of a junco. He may do nothing more, expecting his audience to figure out what extension it typifies – that is, which other birds are juncos. Or he

may underscore its size, coloration, markings, and the shape of its beak. Had he ignored the bird, it still would have had all these features. But it would not have symbolized them. By pointing it out, he exploits features the bird had anyway, bringing the bird to exemplify them. Moreover, the same item can exemplify different features in different contexts. A bird may exemplify being a junco in one context and slate gray in another. When exemplars are improvised, we have no regimented system to rely on. We draw on context, background assumptions and, where available, collateral information instead.

Exemplification figures crucially in the advancement of understanding. A mining inspector extracts air samples from a mine shaft to find out something no one yet knows about the distribution of gases in the mine. If the samples are properly taken, he has reason to believe that the gases his samples exemplify are typical of the gases at different levels in the mine. There are, of course, no guarantees. The inspector must project from a limited sample. It is a brute fact of inductive life that even well chosen, well taken samples are sometimes misleading. But if the samples are well taken and the background theory is well founded, there is epistemically good reason to project to a wider class of cases.

Goodman (1968) maintains that anything that instantiates a feature can exemplify it. In principle this is so, but in practice things are not so simple. For not all instances of a feature constitute *good* exemplars. A feature may be camouflaged: a tiger who blends into the surrounding jungle instantiates being striped, but would be unlikely in that circumstance to exemplify stripedness. (It might, however, exemplify how stripes – even orange stripes – can camouflage in a jungle of green.) A feature may be obscured: a bald man wearing a hat is ill positioned to exemplify his baldness. A feature may be overshadowed: the fearsome timbre and volume of a lion's roar may block its effectively exemplifying its pitch.

Considerable stage setting is often required to generate an effective exemplar. Sometimes we need to remove distractors. Even if peanut butter is typically accompanied by jelly, a good sample of peanut butter should probably omit the jelly. We may even need to tamper with the substance being sampled. An experiment designed to discover whether water conducts electricity would not use ordinary rainwater or tap water. Such liquids contain impurities. Rather, the experiment would begin with distilled water – water from which, as far as we can tell, all impurities have been removed. This involves filtering.

Filtering requires factoring. Before we can remove irrelevant factors, we need to conceptualize the item in question as composed of components – those we seek to exemplify, and those we do well to set aside. Our prior understanding of the domain frequently enables us to do so. Rainwater = H_2O + impurities. But things are not always so simple. For a different experiment, we might want to distinguish between the components of rainwater that are due to pollution and those that are not. In that case, pollen is a component of rainwater rather than an impurity. In yet other cases we might take rainwater to consist of whatever liquid falls from the sky. Then even sulfuric acid is a component of rainwater.

Reconceptualization can highlight features that obtain but are obscured under standard characterizations. Rather than calling an enclosed curve an ellipse, it might be fruitful to consider it a perturbed circle. The curve may be an ellipse, just as the raindrop may contain both sulfuric acid and pollen. The critical question is what features of the item we want to be in a position to focus on. Since exemplification is selective, to treat something as an exemplar is to selectively disregard some of its features. By factoring and filtering, we put ourselves in a position to selectively disregard features we have no current interest in.

[Berkeley's criticism of Locke's discussion of abstract general ideas illustrates the power

of such selectivity. According to Locke, 'the general idea of a triangle . . . must be neither oblique nor rectangle, neither equilateral, equicrural, nor scalene, but *all and none* of these at once.' (Locke, Book IV, ch. 7, sect. 9 quoted in Berkeley §13). Berkeley balks. He can, he insists, frame no such idea. Nor does he think that anyone else can. The definition of the abstract triangle is, as Locke admits, inconsistent. How are we supposed to frame an idea of such a thing? But if we cannot frame such an idea, how is it possible to prove theorems that hold of triangles in general? Berkeley's answer is this:

[T]hough the idea I have in view whilst I make the demonstration be, for instance, that of an isosceles rectangular triangle whose sides are of a determinate length, I may nevertheless be certain that [my proof] extends to all other rectilinear triangles of what sort or bigness soever. And that is because neither the right angle, nor the equality, nor the determinate length of the sides are at all concerned in the demonstration. It is true the diagram I have in view includes all these particulars, but then there is not the least mention made of them in the proof of the proposition. It is not said the three angles are equal to two right ones, because one of them is a right angle, or because the sides comprehending it are the same length. Which sufficiently shows that the right angle might have been oblique, and the sides unequal, and for all that the demonstration would have held good. (Berkeley §16).

By selectively disregarding the rightness of the angle, the equality of the sides, and the exact length of the sides, Berkeley brings the triangle he is working with to exemplify features it shares with all Euclidean triangles.

It might seem that exemplification is not necessary to achieve this goal: the fact that the precise measures of the sides and the angles are ignored suffices. This is not so. The triangle is *capable* of symbolizing generally because these particulars can be selectively disregarded. But to be capable of symbolizing generally is not enough. A picture of my cat does not become a picture of cats in general merely because in discussing it I omit mention of the distinctive color of her eyes, and the markings on her fur. Rather, Berkeley insists, 'An idea which, considered by itself, is particular, becomes general *by being made to represent or stand for* all other particular

ideas of the same sort.' (Berkeley, §12 italics mine). The mode of representation is exemplification. The isosceles right triangle in his proof not only *is* an enclosed three sided figure on a Euclidean plane, the proof is effective because *the triangle refers to the fact that it is* such a figure. By so doing, it stands for other such figures, and discloses properties they share.]

Exemplification figures prominently in empirical science. An experiment is no mere matter of bringing nature indoors. It is a controlled manipulation of events, designed and executed to make some particular phenomenon salient. Natural entities are multifaceted. Important properties and relations are often masked by the welter of complexities that embed them. In experimenting, a scientist isolates a phenomenon from many of the forces that typically impinge on it. To the extent possible, she eliminates confounding factors. She holds most ineliminable factors fixed, effectively consigning them to the cognitive background of things to be taken for granted. This enables the effect of the experimental intervention on the remaining variable to stand out. This strategy enables her to cast into bold relief factors that might typically be hidden from view.

Suppose a population of wild mice who were accidentally exposed to bisphenol-A subsequently exhibited a high rate of liver cancer. To conclude that exposure to bisphenol-A caused their disease would be premature. Those mice might have been peculiarly susceptible to liver cancer or been exposed to a carcinogen that scientists failed to notice. To glean direct, non-anecdotal evidence of a connection between exposure to bisphenol-A and liver cancer, scientists place genetically identical mice in otherwise identical environments, exposing half of them to massive doses of the chemical while leaving the rest unexposed. The common genetic endowment and otherwise identical environments neutralize a multitude of genetic and environmental factors believed to standardly influence the incidence of cancer. This blocks rival

explanations that might be proposed for the elevated rate of cancer in the wild population. If the exposed mice show a significantly higher incidence of cancer than the controls, the experiment exemplifies a difference that correlates with exposure to bisphenol-A.

The result of the experiment exemplifies the difference (if any) in the incidence of liver cancer between the two groups of mice. It not only instantiates the difference, it also highlights that difference. If the difference is statistically significant, then the result exemplifies a correlation between exposure to bisphenol-A and the incidence of liver cancer. Although correlation does not imply causation, a robust correlation is often evidence of causation. Here the background assumption that moves us from a mere correlation to a causal judgment is the well founded conviction that the experiment was so tightly designed and executed that nothing but the exposure to bisphenol-A could have caused the difference. That being so, the result may also exemplify a causal relation.

So far we are just talking about the particular mice in the experiment. But the goal of the investigation is not primarily to discover their medical fates. It is to use their medical fates to learn something more general. Since the mice in the experiment were chosen arbitrarily from the class of mice with particular genome, it is straightforward to extrapolate to other mice of the same strain. The experiment then also exemplifies the increased propensity of mice of that strain to develop liver cancer when exposed to bisphenol-A. Moreover, the mice are model organisms, so there is independent reason to think that what holds for them also holds for the organisms they serve as models for – mammals, including humans. So if the background assumptions legitimating treating the mice as model organisms are sufficiently accurate and adequate, it is reasonable to treat the experiment as exemplifying a causal connection between exposure to bisphenol-A and cancer in mammals.

The interpretive path I have just sketched is reminiscent of Capdevila's interpretation of *Rosso Galera, Rosso Guzzi*. Here the connecting links derive from medical science rather than Italian motorcycle culture. My interpretation is mediated by background assumptions that derive from our current best understanding of carcinogens, mammals, and methods for investigating carcinogenicity. That understanding may be wrong in ways that would vitiate the interpretations. So they are fallible. But if the understanding is sufficiently accurate and adequate (even if not true in every respect), the interpretations are too. We then are right to think that bisphenol-A is carcinogenic.

What if we are wrong? Exemplification requires instantiation. Suppose that, although we have no reason to think so, there is an abrupt threshold. Exposure below a certain level is causally inert. Above that level, cells go wild. Then the consequence of a small animal's abrupt exposure to massive doses of bisphenol-A over a short period of time is not indicative of what happens to a large animal exposed to small doses over a long period of time. Given that humans are never exposed to the spiked levels that the mice were, the experiment does not exemplify a danger to humans. This could be so. An important mediating assumption may be false. In that case, the result does not exemplify anything about human vulnerability to cancer.

Still, the result exemplifies a connection that justifies our thinking that exposure to bisphenol-A increases the likelihood of developing cancer. Even if the result is misleading, it affords insight into the structure of our current understanding of the subject. The attribution of the difference in the incidence of cancer to exposure to the chemical is reasonable to the extent that the scientists manufactured a situation where rival explanations of the difference between the exposed mice and the control group have been blocked. The experiment takes place against a cluster of fallible background assumptions. So it does not afford conclusive evidence. But

because of its rigorous controls, it affords stronger and more direct evidence than a mere correlation between exposure and cancer in a wild population would.

The features of experiments I have emphasized are well known. I mention them to highlight how distant many scientific experiments and their results are from the natural phenomena they illuminate. The items experimented upon are often artifacts constructed expressly for experimentation. The circumstances in which they are placed are artificial; they are carefully contrived situations, often ones that do not naturally occur but that are designed expressly to exemplify telling features of the phenomena. For an experiment to disclose something about a range of phenomena, it must exemplify features it shares with those phenomena. But it may, and in some cases must, diverge from the phenomena in other important respects.

In standard experiments, scientists simplify, streamline, manipulate and omit, so that the effects of potentially confounding factors are minimized, marginalized, or canceled out. An experiment deliberately departs from nature in order to advance an understanding of nature. Rather than invalidating the experiment, this departure is what enables it to disclose barely detectable, or normally overshadowed aspects of the phenomena.

Thought experiments involve further distancing. They are not actual, and often not even possible, experiments. They are imaginative exercises designed to disclose what would happen if certain, perhaps unrealizable, conditions were met.

Sometimes an actual experiment of the sort envisioned *cannot* be carried out. It is impossible or impracticable. By imagining a person's experience while riding in a uniformly accelerating elevator in the absence of a gravitational field and his experience while at rest in the presence of a gravitational field, Einstein shows the equivalence of gravitational and inertial

mass. To actually run the experiment would require placing an unconscious subject in a windowless enclosure, sending him to a region of outer space distant from any significant source of gravity, restoring him to consciousness, and querying him about his experiences. This is morally, practically, and physically unfeasible. Still, the recognition that we cannot do a real experiment does not by itself legitimate stopping short. The infeasibility of performing an experiment can translate into the infeasibility of finding out a particular fact. The reason Einstein's thought experiment is effective is that it takes the form of a challenge: Suppose the specified conditions were met. How could a subject tell whether he was in one situation or the other? If our best efforts to identify a way to tell the difference fail, and fail for scientifically principled reasons, we have evidence of the equivalence. Collectively, our failures exemplify that, if our theories are close to correct, there is no difference to detect.

Sometimes the imaginative rehearsal reveals that an actual experiment *need not* be carried out. The mental run-through itself discloses the relevant information. Even without physical implementation, Galileo's thought experiment discredits the Aristotelian contention that the rate at which bodies fall is proportional to their weight. Imagine a composite object consisting of a boulder tethered to a pebble. Being composed of two rocks and some rope, the composite object is heavier than either rock alone. If Aristotle is right, it should fall more quickly than the boulder. But since, according to Aristotle, the pebble falls more slowly than the boulder, once the two are tied together, the pebble should retard the boulder's fall. Hence the rate at which the composite object falls should be between that of the boulder and that of the pebble. The composite object cannot fall both more quickly and more slowly than the boulder, so the Aristotelian commitments are inconsistent. By exemplifying the inconsistency, Galileo's thought experiment demonstrates that the Aristotelian account cannot be correct.

One might argue that Galileo's thought experiment discredits my analysis.³ Exemplification, I said, requires instantiation. Real mice display an increased incidence of cancer. So it is reasonable to think that by exemplifying that increase, the experiment affords epistemic access to a correlation between exposure to bisphenol-A and cancer, enabling us to recognize it and appreciate its significance, not only in the experimental setting but also outside of it. In Galileo's thought experiment, however, nothing actually falls. A thought experiment, not being material, cannot exemplify material properties. This is so. The sequence of ideas or representations that constitutes Galileo's thought experiment does not instantiate material properties of falling bodies. But the *rate* at which bodies fall and the *independence of that rate* from the weight of those bodies are abstract properties. They can be instantiated by material and immaterial items alike. So there is no bar to saying that via exemplification thought experiments afford epistemic access to abstract properties that are instantiated in material objects. A thought experiment is a representation – a re-presentation – of abstract features, an imaginative re-embodiment of them. We are to imagine – that is mentally, verbally or pictorially present – a situation where the abstract features are realized. In effect, we are to investigate what would happen in a virtual reality where certain constraints are said to hold.⁴

Philosophers sometimes think that we resort to thought experiments only when, for one reason or another, a real experiment cannot be carried out. Perhaps Galileo could not have conducted a real experiment to conclusively demonstrate his point. Maybe he did not have sufficiently accurate timers or a high enough tower from which to run the test. Maybe he did not have the resources to eliminate the effects of air resistance, and so on. Now, however, we could

³I am grateful to Georg Brun and Christoph Baumberger for raising this criticism.

⁴This is consonant with Platonism but does not require it. Perhaps abstract features exist only if instantiated, but instantiations, whether material or virtual, can be created or emerge naturally.

conduct the experiment. Shouldn't we? Probably not. Rather than concluding that the thought experiment was a second-best strategy resorted to because of circumstances beyond the scientist's control, we should recognize that a real experiment would not have made Galileo's case any more forcefully than his thought experiment did. Indeed, it would simply have muddied the waters. Once we start dropping objects from towers, we face the problem that cancer-ridden wild mice pose for medical scientists. How do we know that unrecognized confounding factors do not explain our finding? By deploying an austere thought experiment where the distance and duration of the fall, the presence or absence of air resistance, and a host of other potential sources of interference are omitted, Galileo blocks such challenges. The thought experiment demonstrates an inconsistency in the Aristotelian position – an inconsistency that would obtain regardless of the conditions under which the experiment was conducted. The thought experiment is preferable to an actual experiment because it is invulnerable to a host of potentially misleading challenges that an actual experiment would face.

Scientific models also function as exemplars. They instantiate and refer to features they share with their targets, but diverge from their targets elsewhere. By representing a gas as composed of dimensionless, perfectly elastic spheres that exhibit no mutual attraction, the ideal gas model – $pV=nRT$ – highlights a relation between temperature, pressure and volume that obtains in real gases, but that is typically overshadowed because of the complex geometry of real gas atoms, the gravitational attraction between them, and the propensity of atoms to bond. The model effectively brackets those factors, thereby making the relation between pressure, temperature and volume manifest. Inasmuch as the relation highlighted really does obtain, we understand something about thermodynamics by means of it.

Once we recognize that models and other exemplars sideline features that are referentially

insignificant, we can exploit this capacity through factor analysis. We can construe the phenomenon of interest as factorable into components, distinguish between relevant and irrelevant ones, then sideline the irrelevant ones. We can, for example, represent gas molecules as perfectly elastic spheres with distortions, an oxygen molecule being represented as an elongated sphere cinched at the waist. We do not deny the complex geometry, we simply sideline the confounding factors that for current purposes do not matter.

So why should we resort to experiments, thought experiments and models which are, as I have insisted, to some extent inaccurate? Answer: 'Pay no attention to the man behind the curtain' might be good advice if only we could take it. But often we cannot. The man behind the curtain is too conspicuous to ignore. So we compensate by devising representations of the phenomena from which he is missing. We see what happens then. Maybe we need to introduce correction factors to accommodate the simplifying assumptions we made in our exemplars; maybe not. But if we recognize that the representation serves to illuminate the phenomena by exemplifying features it shares with them, and that it makes no commitment to the realism of unexemplified features, we can see how such exemplars embody, advance and convey an understanding of the world.

References

- Berkeley, George. (1710/1957). *A Treatise Concerning the Principles of Human Knowledge*. Indianapolis: Bobbs Merrill.
- Capdevila-Werning, Remei (2009). *Construing Architecture, Constructing Philosophy*. Dissertation presented to the Departament de Filosofia, Universitat Autònoma de Barcelona.
- Elgin, Catherine Z. (2011). 'Making Manifest: Exemplification in the Sciences and the Arts', *Principia* 15: 399-413.
- Elgin, Catherine Z. (1996)/ *Considered Judgment*. Princeton: Princeton.
- Goodman, Nelson (1968). *Languages of Art*. Indianapolis: Hackett.
- Locke, John (1690). *An Essay Concerning Human Understanding*.
- McGowan, Mary Kate (2003). 'Realism, Reference and Grue', *American Philosophical Quarterly* 40: 47-57.
- Strevens, Michael (2010) . 'Varieties of Understanding' draft.