THE ISAAC LEVI PRIZE 2020

HEMPEL'S RAVEN REVISITED*

Carl G. Hempel introduced an important puzzle about the confirmation of generalizations. A *prima facie* plausible way of thinking about it takes it to proceed by "instance confirmation." We see some particular *x*'s; we note that they have a certain property *P*; and we take ourselves to have found confirming evidence that, in general, *x*'s are *P*. However, as Hempel pointed out, attempts at formalizing this procedure generate some counterintuitive consequences.[1] Consider the generalization "ravens are black." The generalization is often formalized, as Hempel formalizes it:

$$\forall x(Rx \rightarrow Bx)$$

Observations of particular black ravens would count as evidence for the truth of this generalization. Applying a simple rule of inference, contraposition, we get:

$$\forall x(\sim Bx \rightarrow \sim Rx)$$

This statement has the same "semantic content" as the first formalization—it differs only in logical form. So, whatever counts as evidence

---

[1] For Jean Nicod's attempt at formalizing instance confirmation, see Jean Nicod, *Geometry and Induction* (Berkeley: University of California Press, 1970). For the critical commentary in which Hempel develops his problem, see Carl G. Hempel, "Studies in the Logic of Confirmation (I.)," *Mind*, LIV, 213 (January 1945): 1–26.

---

for this second statement should count also as evidence for the first. However, observations of white shoes would count as evidence for this second statement. We see a pair of shoes and notice that they are not black and not ravens—true antecedent, true consequent, therefore true conditional. It is indeed odd to think that observations of white shoes should count as evidence for the generalization that ravens are black. Yet that is what the logic seems to dictate.

In noting this clash between logic and intuition, you might think that we are at least on the right track with our inductive logic, and the fact that it clashes with our intuitions just means we need to shore up the logic. The goal, then, is to rework our system of inductive logic such that it no longer yields counterintuitive results the likes of which Hempel put his finger on. That is, as Goodman might say, to solve the problem is to bring our system of inductive logic into "reflective equilibrium" with our intuitions regarding which inductive inferences we should license and which ones we should not.[2] I am going to approach the paradox another way—by understanding the relevant intuitions.[3] This "intuitions first" approach has a decided advantage over the "logic first" reflective equilibrium approach that has been standard in philosophical commentary on Hempel's Paradox. On the latter approach, the intuitions are left as a kind of black box and it is taken on faith that, whatever the structure of the intuitions inside that box might be, it is a structure for which we can construct an isomorphic formal edifice—a system of inductive logic. By popping open the box, we can see whether that faith is misplaced.

Perhaps a word on this methodology is in order. In a sense, what I want to do here is nothing more than garden-variety empiricist philosophy. Just as Hume proposed to vet the idea of causation by interrogating the (supposed) character of our actual understanding of that relation, I propose to consider Hempel's Paradox in a similar way, namely, by trying to understand how we actually understand generalizations like the ones that generate Hempel's problem. In other words, my plan is to investigate the character of our intuitive, pre-theoretical understanding of statements like "ravens are black." Understanding what this amounts to can give us a clearer understanding of Hempel's Paradox. It enables us to say not just *that* the particular consequence that generates the paradox is counterintuitive but, also, *why* it is so.

---

[2] Nelson Goodman, *Fact, Fiction, and Forecast* (Cambridge, MA: Harvard University Press, 1983).

[3] For another "alternative approach" to the paradox, see Israel Scheffler, *The Anatomy of Inquiry: Philosophical Studies in the Theory of Science* (New York: Routledge, 2014).

Another methodological note: I take cases in which we come to terms with apparent *exceptions* to be particularly informative with respect to our pre-theoretical intuitive understanding of generalizations such as "ravens are black." I consider a number of such cases throughout the paper. I have found it helpful to do so in light of two hypotheses regarding the structure of our intuitive understanding—a *taxonomic* and a *nomological* one. I tentatively characterize these hypotheses at the outset of my inquiry, further refining them and developing the distinction between them as the analysis proceeds. Noting differences between how we would respond to apparent exceptions to generalizations if these hypotheses were true and how we *in fact* respond to such exceptions throws into relief our actual pre-theoretical understanding.

Rather than taxonomically or nomologically, my analysis suggests that we actually understand such generalizations *mechanically*. Though I am not developing a view on causation, and though my analysis does not depend upon hers, this mechanical understanding is "thick" in the same way that Nancy Cartwright argues our causal concepts are.[4] Just as Cartwright tells us that there are many distinct thick causal concepts, there are many distinct ways in which we pre-theoretically understand generalizations like the ones at hand mechanically. As my analysis suggests, we intuitively understand "ravens are black" to mean *ravens are black by some indeterminate yet characteristic means.*[5] Likewise, we intuitively understand "coal is black" to mean *coal is black by some indeterminate yet characteristic means.* Indeterminate though they may be, we understand the means by which ravens are black to be different from the means by which coal is black. This does not bode well for the prospect of explicating the inductive confirmation of generalizations in terms of a single monolithic formal system that depends upon abstracting away from just those differences. To borrow from Cartwright, when we do "we…lose a vast amount of information that we otherwise possess."[6] But we do not lose it entirely. In fact, we do not really

[4] Nancy Cartwright, "Causation: One Word, Many Things," *Philosophy of Science*, LXXI, 5 (December 2004): 805–19. Her notions of "thick" and "thin" are discussed in more detail, and further analyzed, in section III. While I am quite sympathetic to her view, and while engaging with her work has certainly been inspirational with respect to the thoughts expressed here, I do not take my analysis to depend in any substantive way upon the details of Cartwright's, nor do I commit myself to her conclusions regarding causality. For the purposes of this paper, I borrow her terms "thick" and "thin" to help clarify my own view and leave articulating the relationship between her view and my own for the future.

[5] I struggled to find the correct term to use in this formulation. I thank an anonymous reviewer for raising worries about "underspecific" and giving me the nudge that led me to settle on "indeterminate."

[6] *Ibid.*, pp. 814–15.

lose it at all. Rather, we decide not to countenance it in our formal theory. That we *retain* this information explains why we find certain consequences of formal theories of induction to be counterintuitive.

Sections I, II, and III of the paper develop what is meant by "means," "indeterminate," and "characteristic," respectively. This, however, is just a matter of emphasis. The significance of each term can only be appreciated in light of the analysis running through the entire paper. I close by spelling out the relationships between the taxonomic, nomological, and mechanical hypotheses and bring the analysis to bear on Hempel's problem.

I. RAVENS ARE BLACK BY SOME INDETERMINATE YET CHARACTERISTIC *MEANS*

Let us begin our analysis by asking ourselves how a generalization like "ravens are black," or as it is standardly formalized, $\forall x (Rx \rightarrow Bx)$, might be understood. Here are two possible ways:

(1) *The Taxonomic Understanding*: The conditional specifies a necessary condition on being a member of the category *raven*, namely, being black.

(2) *The Nomological Understanding*: The conditional states that a law-like relation holds between being a raven and being black.

At first blush, these may seem rather different. On the taxonomic understanding, "ravens are black" *defines* ravenhood partially in terms of blackness. On the nomological understanding, there is a notion of ravenhood defined independently of blackness, and the generalization states that this prior notion is related to blackness by law. Perhaps the most important difference between these two ways of understanding the generalization is that only on the nomological understanding do we get Hempel's Paradox. If the statement is a matter of definition, there is no question of evidence bearing on it in the first place and, hence, no paradox.

Some readers may doubt that these ways of understanding the generalization are really as different as I seem to be suggesting. After all, even if it is in some sense a matter of definition, as the taxonomic understanding would have it, that definition was not just stipulated from the armchair. Rather, taxonomists of the biological world have painstakingly built their categories working in the field. What is the difference, one may ask, between such taxonomic fieldwork, part of the goal of which would be to determine whether blackness is a necessary condition on being a member of the category "raven," and that empirical activity the end of which is to determine whether ravenhood and blackness are nomologically related? *Evidence* in the form of ob-

servations of actual ravens seems to come into the picture in both cases.

I sympathize. In fact, here are some further reasons why we might think that these two ways of understanding the generalization are not as different as they may seem on their faces.

- First, obviously, they both license the inference from "*x* is a raven" to "*x* is black."
- Second, it seems, both ways of understanding the generalization support the corresponding counterfactual. That is, assuming that some particular *x* is in fact not a raven, it is true, in light of the truth of the generalization, that "if *x* were a raven, then it would be black."
- Relatedly, third, both understandings take the universal generalization to be true without exception.
- Fourth, on neither understanding is it the case that one can know the truth of the generalization just by reflecting on the terms used in formulating it. To put this in terms of the other side of the same coin, no contradiction—like that involved in the idea of a married bachelor—is involved in the idea of a non-black raven (on neither understanding is the sentence "analytic"). This is a consequence of what I mentioned above, namely, that observations of actual ravens are relevant whether we are trying to determine if blackness is a necessary condition on ravenhood or if ravenhood and blackness are nomologically related. Besides, Hempel's Paradox is a problem for the justificatory status of *empirical* generalizations. Even if there were some way of understanding the sentence as analytic, it would not be relevant here.
- Fifth, confronting something that looks for all the world like a raven and yet is not black would require one of two responses similar across the taxonomic and nomological understanding. On the taxonomic understanding, we could: (1) stick to our category specification and argue that this white thing is not in fact a raven, or (2) we could modify our category specification to include the white thing within the raven category. The latter would require striking blackness from the list of necessary conditions for ravenhood. Either tack would likely necessitate some argument against, on the one hand, liberals who do not think the category should discriminate on the basis of color and, on the other, conservatives who insist that blackness is of the essence of ravenhood and that the category must remain pure. Likewise, on the nomological understanding, in encountering a white raven, we could (1) hold fast to the truth of the law and argue that this white thing is not in fact a raven, or (2) we could acquiesce in the law's falsity and modify it accordingly. One could wonder whether there is a substantive difference between holding fast to a category specification according to which blackness is a necessary condition on being a raven

and holding fast to the truth of a universally binding law that says ravenhood and blackness are nomologically related. Likewise, one could wonder whether there is a difference between modifying the specification of the raven category and modifying the law to accommodate observations of white ravens.

I return to the question of how to differentiate the taxonomic and nomological hypotheses in section III. For now, note that regarding the last bullet point above, one may object that the strategies for coping with being confronted by a white raven are only as limited as I have made them out to be because I have unfairly characterized the taxonomic and nomological ways of understanding the generalization as implying that it is *exceptionless*. We simply do not—nor should we—understand it like this. Categories have vague boundaries, and laws are always *ceteris paribus*. Think of it this way: imagine you enter a walk-in refrigerator with your little nephew. You hug your elbows, exaggerate a shiver, smile down at your nephew, and say, "Brrrr, it's freezing in here!" Apparently annoyed, your nephew looks up at you and says, "Freezing is 32 degrees Fahrenheit. It's *not* freezing in here." Understandably, you roll your eyes and reflect on how your nephew is turning into a little brat. Later that day, sitting at the table helping your niece with her homework she asks, "What color are ravens?" You respond, "Ravens are black." Just then your nephew walks in and says, "Oh yeah, what about *albino* ravens?" Again, brat. The exceptionless ways of understanding "ravens are black" that I have characterized, goes the objection, make it out that your brat-of-a-nephew's smarmy over-literality is paradigmatic of how such generalizations are understood. Yet, it is just his over-literality that is diagnostic of his being a brat. In fact, that we *do not* understand such generalizations as statements of exceptionless fact is a precondition on his ability to exercise this particular form of bratty-ness. That we do not intuitively understand them as exceptionless is precisely what affords him the opportunity to be a brat in this particular way.[7]

Again, I sympathize. Pre-theoretically, we simply do not understand these generalizations as stating exceptionless truths. In fact, the above example gives us a palpable sense of how far an exceptionless understanding of these generalizations departs from our intuitive understanding of them—precisely as far as your nephew's comment reflects a deviation from norms of linguistic comprehension. This is

---

[7] This first exchange regarding "freezing" is meant to establish the nephew's know-it-all character. He is *that guy*. (We all know *that guy*.) Thank you to my nephew, Wesley Bollhagen, for inspiring this example. We are all glad he has grown out of this phase.

significant for our understanding of Hempel's Paradox. In order to formulate the paradox, we must take our intuitive non-exceptionless understanding of "ravens are black," formalize it as the exceptionless $\forall x(Rx \rightarrow Bx)$ such that, by contraposition, $\forall x(\sim Bx \rightarrow \sim Rx)$ is its semantic equivalent. This process involves a departure from our pre-theoretical, intuitive manner of understanding them although, perhaps, a less obnoxious departure than your nephew's.

So, we do not intuitively understand such generalizations as exceptionless. Here is a case to help us see another aspect of our intuitive understanding of such generalizations. Take this example: "objects fall"—$\forall x(Ox \rightarrow Fx)$. Now imagine you are walking in the park with a group of friends. You come around a corner and see a park bench floating unsupported above the ground. One of your group points to it and says, "Well, I guess *that's* not an object" and continues along her merry way. Your friend's reaction would be odd indeed, yet such a response would be just what we would expect if the relevant generalization were intuitively understood taxonomically as providing a partial specification of the category *object*. Equally strange would be for one's response to this situation to consist in whatever simple act of mind is involved in acquiescing in the "law's" falsity—"Welp, I guess *that's* false."

Of course, any real person's cognitive response to this situation would be more complex than either of these alternatives suggest. One natural response would be wonderment or perhaps horror, "*Oh. My. God.*"[8] A less affective, more curious response—one more relevant to our inquiry here—would be to investigate to the end of determining by what *means* this remarkable phenomenon is occurring: "*How in the world…?*" That is, a natural response would be to try and understand the relevant *mechanism*. This does not involve denying that the bench is an object, nor does it involve denying that objects fall. In fact, recognizing that, in general, objects fall is crucial in attempting to construct a model of what is going on mechanically here. That seeking after the mechanism is the natural response to confronting such a phenomenon indicates that the generalization to which the phenomenon is an apparent exception is, correspondingly, understood with respect to mechanism rather than with respect to the sharply bounded categories into which the world breaks down or the exceptionless laws that govern it.[9]

---

[8] For more on "wonder" and "horror" in philosophy and science, see the magisterial Lorraine Daston and Katherine Park, *Wonders and the Order of Nature: 1150–1750* (New York: Zone Books, 1998).

[9] It should be noted that I do not intend my use of the term "mechanism" here to allude to any particular account of mechanism in the philosophical literature. At this

We can say more, however, regarding what it means to say that we understand generalizations like the ones in question "mechanically." For help, let us consider Michael Strevens's recently developed view on the semantics of causal generalizations. His view is akin to the one I am developing here. Like the view I am developing, his semantics has it that generalizations like the ones at hand are not understood as exceptionless. Further, on his view, there is a sense in which they are understood "mechanically." His account, however, flirts with a serious difficulty that I call the *anamnesis problem*. In brief, Strevens's exposition suggests a reading of his view on which we already know—albeit "opaquely"—determinate details regarding the mechanism relating, for instance, ravenhood to blackness. The empirical discovery of those details thus consists in making apparent what we already tacitly knew about that mechanism, turning mechanistic inquiry into an exercise in anamnesis. Nonetheless, it is a helpful contribution, and its shortcomings illuminate an important aspect of my analysis, namely, the significance of the term "indeterminate" in my formulation of how we intuitively understand generalizations like the ones at hand.

II. RAVENS ARE BLACK BY SOME *INDETERMINATE* YET CHARACTERISTIC MEANS

To illustrate his view, Strevens considers a situation comparable to our park bench anomaly.[10] A group of scientists studying ravens—and believing that ravens are black—one day discovers a population of gray ravens. Sometime later, these scientists discover that the gray

---

point in the analysis, I use it simply to distinguish this third hypothesis regarding how we pre-theoretically understand generalizations like the ones at hand from the taxonomic and nomological hypotheses. In other words, I do not mean to suggest that one's response to confronting the floating park bench would be to seek an account specifically in terms of "entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions" (Peter Machamer, Lindley Darden, and Carl F. Craver, "Thinking about Mechanisms," *Philosophy of Science*, LXVII, 1 (March 2000): 1–25) or in terms of "a structure performing a function in virtue of its component parts, component operations, and their organization" (William Bechtel and Adele Abrahamsen, "Explanation: A Mechanist Alternative," *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, XXXVI, 2 (June 2005): 421–41). These conceptions of mechanism are philosophically downstream from this point in our analysis. The response, as I stated above, would be to seek the *means by which* the exception to the generalization is occurring. It may be that, after inquiring into the matter, we settle on an account that would count as a mechanistic explanation by the lights of the philosophers alluded to above. However, if we found out that it was a matter of divine intervention, then *that* would be the relevant means—or, the *mechanism*—in the broader sense of "mechanism" that I intend.

   [10] Michael Strevens, "*Ceteris Paribus* Hedges: Causal Voodoo That Works," this JOURNAL, CIX, 11 (November 2012): 652–75.

ravens' habitat contained a previously unknown industrial pollutant that blocks a certain metabolic pathway preventing the development of the ravens' normal black plumage. In such a case, says Strevens, the scientists will not simply regard their original hypothesis as false. Why? According to Strevens,

> ...they will regard themselves as having discovered that the gray ravens were all along irrelevant to the blackness hypothesis, because the blackness hypothesis was intended to describe the consequences of the natural coloration mechanism, and the grayness of the ravens was no more a product of that mechanism than the whiteness of bleached ravens....In other words, they will regard their hypothesis as having had an implicit rider saying, among other things, *Provided that there is no significant amount of* ABC *in the environment*....This is a rider that they were incapable of spelling out at the time; it therefore gave their hypothesis opaque content.[11]

On Strevens's account, generalizations like "ravens are black" generally hold only *ceteris paribus.* This is illustrated in the example. The scientists do not take their discovery of gray ravens to falsify the generalization as they would if they took their generalization to be exceptionless. Rather, "*CP*-ravens are black" is what they meant all along. For Strevens, as alluded to in the example, a CP-clause renders a generalization non-exceptionless by introducing a reference to a particular mechanism. In the example above, the mechanism referred to is the "natural coloration mechanism" for raven plumage. Further, according to Strevens, the CP-clause refers to the conditions that enable the mechanism's operation. This would be the broader system supporting the pathway from DNA transcription to, ultimately, the expression of the black plumage phenotype. Finally, it states that the mechanism operates properly. That is, nothing interferes with the operation of the natural coloration mechanism itself. So, the truth conditions for "CP-ravens are black" are:

- *When condition O holds, then by way of the target mechanism M, the conditions Z and the property F bring about the property G.*[12]

Here, Z stands for, for instance, the broader system supporting the pathway from genotype to phenotype (the enabling conditions), O is the condition specifying that the mechanism, M, itself operates properly, F is being a raven, and G is having black plumage. These truth conditions purport to account for why the scientists do not reject

[11] *Ibid.*, pp. 662–63.
[12] *Ibid.*, p. 660.

the generalization as false. Their discovery that ABC interfered with the proper operation of the natural coloration mechanism was tantamount to discovering that condition *O* did not hold. Thus, in the case of these scientists, this condition endowed the generalization with "opaque content" in the form of a rider—*unless there is a significant amount of* ABC *in the ravens' environment.* It is the presence of such opaque content that renders generalizations like the ones at hand non-exceptionless.

Although Strevens's analysis is aimed at the semantics of CP-hedged generalizations, he concludes his paper stating, "Does it matter whether or not a causal generalization is explicitly hedged?...Let me conclude by endorsing this view: the sole semantics for causal generalizations is the semantics I have given for hedged hypotheses."[13] Thus, even if not explicitly CP-hedged, *all* statements of empirical generalization like "ravens are black" involve a tacit *ceteris paribus* clause—laws are always *ceteris paribus*.

As I mentioned, I am largely sympathetic. In fact, though I am not developing a semantic theory, there is considerable overlap between Strevens's view and my own analysis as it has been developed so far. First, for Strevens, a proper account of the semantics of generalizations like the ones at hand shows them to be not a matter of category or law but a matter of *mechanism*. Second, Strevens does not take such generalizations to be understood as statements of exceptionless truth. Finally, his view can be deployed to rationalize your reaction to your nephew's comment about albino ravens in the example above. The existence of albino ravens does not "falsify" the generalization "ravens are black" because, on Strevens's view, that generalization has a tacit CP-clause endowing it with an opaque rider to the effect of, "*provided that there are no genetic defects that prevent the normal production of melanin.*" On Strevens's understanding, that bit of content was not something "which [you] were capable of spelling out at the time" of your nephew's comment—it was opaque—though you may become capable of articulating it by Googling "albinism" and learning a little about how the genetic disorder works. Nonetheless it was there, like a hole awaiting an appropriately shaped peg, prompting you to roll your eyes at your nephew's over-literality and reach for your Googling device. Thus, although Strevens's example involves scientists in practice, the fact that his semantics applies so readily to the case of the bratty nephew—a case illustrating our pre-theoretical, intuitive understanding of our generalizations—shows that his semantics square with our pre-theoretical, intuitive understanding as well. In fact, Strevens's

[13] *Ibid.*, p. 675.

view could be taken as one possible answer to the question of what it means to say that we intuitively understand generalizations like "ravens are black" mechanically.[14]

That said, let us consider Strevens's view more closely. There are two ways of understanding it, one much more plausible than the other. The implausible view goes like this. When the raven researchers discovered that the gray ravens inhabit an environment containing ABC and recognized that this pollutant inhibits the development of normal black plumage, they genuinely discovered something *about their understanding* of the generalization "ravens are black." In other words, they found something that was, albeit tacitly, in their understanding all along, namely the bit of opaque content corresponding to the facts about the mechanism that they discovered. As I put it above, this content was always there like a hole into which only a *there-is-a-significant-amount-of-*ABC-*in-the-gray-ravens'-environment*-shaped peg could fit. The researchers' discovery consisted in finding this peg and, simultaneously, realizing that there was a hole in their understanding of the generalization into which, uncannily, this peg, and only this peg, fits. Discovering the peg and the hole and inserting the former into the latter is the act by which the generalization is saved from falsification upon the discovery of gray ravens. This way of spelling out the view suggests that there are many more bits of content—many more "holes"—scattered opaquely throughout the semantics of the generalization, waiting to be brought to light as the scientists discover further exceptions and come to understand why they are exceptions in the course of empirical inquiry. You do not have to be tryptophobic to find this view aversive. The fact that it turns mechanistic inquiry into an exercise in anamnesis is enough.[15]

---

[14] But there is a certain *direction* to Strevens's analysis, a *loop* that it follows, that from the point of view of my analysis, is superfluous. As I have argued, we already *pre-theoretically* understand generalizations like "ravens are black" as countenancing exceptions. We do violence to that understanding when we formalize it in a way that renders it exceptionless—that is, when we explicate it as $\forall x(Rx \rightarrow Bx)$. Tacking a CP-clause in front of this in reparation would be unnecessary if the violence had not been committed in the first place. In the end, Strevens's analysis seems to come full circle when he claims that his view applies even to unhedged generalizations suggesting that they are tacitly CP-hedged to begin with. But we had to pass through a cycle of violence and reparations to get there. Even then, characterizing our intuitive understanding of such generalizations in terms of their being *tacitly CP-hedged* is a way of characterizing it that still bears the mark of that cycle.

[15] We can make this same point without going all the way back to the *Meno*. Imagine that as the scientists in Strevens's example leave their lab to set out and study ravens in the field, a security guard in the building asks them, "Are ravens black?" The scientists respond, "Indeed, ravens are black." The security guard, who is a compulsive gambler and a persuasive man, convinces them to put fifty dollars on the matter. They all shake

That said, the passage quoted above suggests that a more plausible version may be in the offing. The passage reads that the scientists "…*will regard themselves* as having discovered that the gray ravens were all along irrelevant to the blackness hypothesis… [and] *they will regard their hypothesis* as having an implicit rider…" (my emphasis). The implausible view elaborated above has it that the scientists are *correct* in how they regard the status of their own discovery—as if in fact their updated understanding of "ravens are black" is *really* just the same as their antecedent understanding only now with some formerly implicit content made explicit in an act of anamnesis. Let us think through a modified example to help see our way to a more plausible version of the view.

Suppose that rather than discovering that the ravens' environment contained an industrial pollutant that blocked a metabolic pathway, the researchers piece together a plausible evolutionary story. They find good evidence suggesting that the population of gray ravens initially belonged to a homogeneously black population. At some point, a subpopulation became geographically isolated in a new selective environment that, for one reason or another, favors gray plumage. The story they settle on is that the ravens evolved gray plumage by natural selection. In light of finding this to be the case, how will these researchers view their original hypothesis that ravens are black?

The quote from Strevens above suggests an answer: "…the blackness hypothesis was intended to describe the consequences of the *natural coloration mechanism*" for raven plumage (my emphasis). Supposing these scientists believe that natural selection is the mechanism responsible for ravens being black, they will now see the grayness of the gray ravens as being a product of the natural mechanism for raven coloration. Strevens's view suggests, then, that the researchers in the modified example will view their original hypothesis that ravens are black as false because they have come to view gray plumage as me-

---

on it. Then, the situation plays out as Strevens describes. The scientists discover the gray ravens, find that their habitat contains the industrial pollutant, and learn that the pollutant blocks a metabolic pathway and prevents the normal development of black plumage. They return home to their lab and present their findings to their colleagues. After their talk, the security guard approaches them, holds out his palm, and asks for his fifty dollars. The scientists refuse to pay, saying that while they were not in a position to appreciate it at the time, when they said "Indeed, ravens are black," that statement contained an opaque rider, and so, what they actually meant was, "Ravens are black *provided there is no significant amount of* ABC *in the ravens' environment.*" They were unaware of this bit of content at the time—that rider was opaque to them—and they argue that bets made under conditions of "ignorance" are not binding. Understandably, the security guard rolls his eyes and reflects on how these scientists remind him of his bratty nephew. The implausible version of Strevens's view takes the scientists' rationalization to mirror the actual truth of the matter.

chanically consistent with the normal operation of the natural coloration mechanism—natural selection.

For my purposes, however, the issue of where these scientists will settle with respect to the truth of "ravens are black" is beside the point. The point of my modified example is to show that, rather than the scientists appealing to the *ontogenetic* mechanism to which they appealed in Strevens's own version of the example, the scientists in the modified version will appeal to a *phylogenetic* one—natural selection—as relevant for updating their understanding of the relationship between ravenhood and blackness.[16]

The more general point that emerges from the juxtaposition of the two examples is that what mechanism for raven coloration—what specific *means by which ravens are black*—they will appeal to depends upon the facts they uncover downstream in their empirical research. If the facts point in one direction, the mechanism they will come to attend to may be ontogenetic, as in Strevens's example. If the facts go another way, it may be phylogenetic, as in my modified version. However it goes down, on pain of running headlong into the anamnesis problem, it cannot be said that, prior to the facts coming in, the CP-clause that they tacitly understood to qualify "ravens are black" referred determinately to one or the other kind of mechanism. Further, for the same reason, the particular facts they find downstream cannot be thought to correspond to bits of doppelganger content lurking all along in the shadowy corners of the semantics of "ravens are black."

So, again, what could it mean to say that we intuitively understand generalizations like "ravens are black" mechanically? The lesson we have learned from our discussion of Strevens is that we can understand the sentence to pick out neither a particular mechanism nor even a particular kind of mechanism (for example, phylogenetic versus ontogenetic) without running into the anamnesis problem. This is why I include the term "indeterminate" in my formulation of how we pre-theoretically understand generalizations like "ravens are black"—

---

[16] Given the complexity of issues regarding taxonomy in biology and the concept of "species," the way such findings would lead researchers to update their understanding of the relationship between ravenhood and blackness is likely much more complicated than my brief discussion here suggests. Since my purposes are to indicate how my analysis departs from Strevens's and, in particular, to indicate the significance of the term "indeterminate," I confine my discussion to match the scope of his. For more on taxonomic practice in biology and the concept of "species," see Marc Ereshefsky, "Species, Taxonomy, and Systematics," in Mohan Matthen and Christopher Stephens, eds., *Philosophy of Biology* (Amsterdam: North-Holland, 2007), pp. 403–27; and David Hull, *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science* (Chicago: University of Chicago Press, 1988).

ravens are black by some *indeterminate* yet characteristic means. Keeping the lesson we have learned from our discussion of Strevens in mind, let us turn back to the taxonomic and nomological hypotheses I set out at the beginning of the inquiry. As I stated in my introduction, these constitute hypotheses regarding how we intuitively understand generalizations like the ones at hand. By further developing these hypotheses and noting the contrast between what they suggest and how we actually understand these generalizations, the latter is thrown into further relief.

III. RAVENS ARE BLACK BY SOME INDETERMINATE YET *CHARACTERISTIC* MEANS

Above, I blurred the distinction between the taxonomic and nomological ways of understanding generalizations like the ones in question, suggesting that they are not as distinct as they might at first appear. I gave a few reasons for thinking so. Here is another one, this one more historical than logical.

In tracing an important thread running through the history of analytic philosophy, Robert Brandom writes:

> The Early Modern philosophical tradition was built around a classificatory theory of consciousness and (hence) of concepts....The paradigmatic cognitive act is understood as classifying: taking something particular as being of some general kind. Concepts are identified with those general kinds.[17]

On Brandom's reading of the history, the development of Frege's logic—analytic philosophy's moment of conception—provided the tools for a mathematical explication of the core theory of mind that had animated the early moderns. On this view, thought was understood to be essentially *taxonomic*—concepts are identical to *classes* and reasoning consists in reasoning about these classes, the relations between them, and what belongs in them. Frege's logic lays bare the structure of taxonomic thought and reasoning and so lays bare "the structure of sapience itself."[18]

The logical positivists were enthusiastic about the prospects of using the new mathematical logic for explicating the structure of scientific theories. In their hands, a formal sentence—for instance, $\forall x(Rx \rightarrow Bx)$—that was originally conceptualized, as Brandom tells us, as specifying the structure of taxonomic judgments (being black

---

[17] Robert Brandom, "How Analytic Philosophy Has Failed Cognitive Science," *TAP-2009 Towards an Analytic Pragmatism* (2009): 121–33.

[18] *Ibid.*, p. 121.

is a necessary condition on being a raven) was redeployed to explicate the content of empirical generalizations or laws. This represents progress, arguably. There may be some *intuitive* difference between understanding a statement as specifying conditions for category membership and understanding it to specify an empirical law, but when we discipline our unruly intuitive understanding with the new logic, we can appreciate that whatever intuitive difference there may be is illusory. Logic, not our messy human intuitions, is the final arbiter of content. Since "ravens are black" is properly explicated as $\forall x(Rx \rightarrow Bx)$, and since this formalism recognizes no distinction between a taxonomic and nomological reading, then there just is no such distinction. On a progressive reading of the history, Frege's logic brings our inherently flawed human understanding to heel and allows us to *discover* that, our erring intuitions aside, there is in fact no difference between the taxonomic and nomological ways of understanding generalizations of the relevant sort.

One may object, however, claiming that I have missed the point here. "May I remind you," goes the objection, "that your task was to interrogate the structure of our *pre-theoretical, intuitive* understanding of such generalizations. In your telling of the history—a highly cursory and probably false one, I might add—you suggest that there may be a pre-theoretical, intuitive difference between understanding a generalization taxonomically and understanding it nomologically, and you claim that developments in logic allowed us to discover that that difference is merely an illusion. However, illusory or otherwise, *that difference* is the target of your inquiry, and simply calling it an illusion does not an adequate account make."

This objection is spot on. Simply calling some counterintuitive consequence of our logic an illusion does not constitute an adequate treatment. If this were the case, Hempel's Paradox could be solved with a word. So, thanking our objector for the reminder, let us return to the topic at hand.

In section 1, while blurring the distinction between the taxonomic and nomological ways of understanding I suggested that, though it may be that the taxonomic way takes the generalization to be a matter of category and the nomological way, a matter of law, they are both, so to speak, *a posteriori*. What this means is that whatever difference there may be between them, it is not going to be explicable in terms of the way in which the knowledge was acquired—both require fieldwork, after all. This suggests that the place to look is in how the knowledge is structured once gained. What is the difference between structuring this knowledge taxonomically and structuring it nomologically?

Let us start with the taxonomic way. One simple way to think of it is in terms of a Venn Diagram:
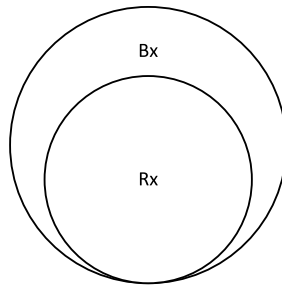
Figure 1. The Venn Diagram.

The Venn Diagram encourages us to think of the relation between being black and being a raven expressed in "ravens are black" as being one of *containment*. The category "raven" is *contained within* the broader category of "black things." This makes sense. We frequently talk about items being "in" a category.

Notice, importantly, that on the taxonomic understanding the raven category is not special vis-à-vis blackness. Just as a real container can contain any number of things indifferently, the category "black things" contains all sorts of subcategories other than the raven one. The raven category holds no pride of place in the "black things" category. That is, *there is nothing that the relation between blackness and ravenhood has*, on this taxonomic way of understanding the generalization, *that the relation between blackness and any other black thing does not have*.

This is perhaps a bit of a subtle point, so here is another way of making it. In mathematics we say that two numbers are *commensurable* when they can be shown to be composable out of the same base units. 5 and 10 are commensurable. We can show this by, for instance, choosing "5" as our base unit. 5 is composed of one "5," and 10 is composed of two "5"s. We can also show this by choosing "1" as our base unit. 5 is composed of five "1"s, and 10 is composed of ten "1"s. The primeness of prime numbers like 7 consists in the fact that they can only be shown to be commensurable with natural numbers less than themselves by appeal to a base unit of "1."[19]

"Properties," understood as things that some particular must bear in order to *count as* a member of some category, are the "base units" out of which members of a category can be composed. Thus, on the

---

[19] Of course, 7 is commensurable with 14, 21, 28, and so on. In this case, 7 is being construed as the base unit—"7"—with 14 being two of these base units, 21 being three of them, and so on. Any number will naturally be commensurable with multiples of itself. Thanks to Adam Fox for this point and to an anonymous reviewer for forcing me to clarify it.

taxonomic hypothesis, when we understand "ravens are black" we understand the property "blackness" to be a "base unit" without which a raven cannot be constructed but, equally, a property that can be used to construct indifferently any number of black things. Just like "1," construed as a base unit, bears no special relation to any of the numbers it can compose—there is nothing that base-unit-1's relation to 4 has that its relation to 9 and 368 does not have—"blackness" bears no special relation to ravenhood on this understanding. Again, there is nothing that "blackness"'s relation to ravenhood has that its relation to any other black thing does not have. In other words, "properties" are the interchangeable parts—base units—out of which the world is composed insofar as we "cut it at its joints" taxonomically. The blackness of a raven can be switched out for the blackness of coal, which can be switched out for the blackness of the sky at nighttime.[20] The traditional way of putting this is that blackness is the same "universal" instantiated in ravens, coal, and the night sky.

To see this yet another way, consider how we evaluate statements in predicate logic using "models." We build a "universe of discourse" that consists of a set of "objects." We further specify "predicates" and determine which predicates are true of which objects in the model. The model may have it that some predicate, F, is true of two objects, a and b, in the universe. From the point of view of evaluating a sentence with a model, there is just no difference between, on the one hand, the relation between a and F and, on the other, the relation between b and F. In other words, there is nothing that a's relationship to F has that b's relationship to F does not have and vice versa. It would make no sense to ask after something like a difference in the "manner of connection" between a and F and that between b and F as, from the point of view of evaluating formal statements using models, the relation between object and predicate is just *predication* indifferently.

However, it is clear that this taxonomic way of understanding generalizations is not how we intuitively understand them. Consider that we certainly do not respond to coming across exceptions to generalizations in the way we would if we understood them taxonomically. In coming across a white raven in a park, we do not intuitively diagnose it in terms of a failure with respect to the instantiation of a universal or a failure of predication—"I guess blackness failed to instantiate in that bird," or, "What we have here, friends, is a failure of predication." Someone who was thusly pathological in their diagnostic inclinations would be an utterly useless partner in inquiry. Further, one who was

---

[20] This is just another way of saying that the "raven" category, the "coal" category, and the "night sky" category all fall under the more general "black things" category indifferently.

so inclined would have no resources to diagnose differentially a white raven and a white piece of coal. From the point of view of the taxonomic understanding, both could only be understood as essentially the same failure, a failure of the universal "blackness" to instantiate in the particulars in which it normally and, as I mentioned above, *indifferently* instantiates. Intuitively, we *do* differentially diagnose these two situations. In being confronted with a piece of white coal, no one would ever think to diagnose it with albinism. Our pre-theoretical, intuitive understanding of the relationship between ravenhood and blackness (or coalhood and blackness, for that matter) is therefore more diagnostically fine-grained—more *intelligent*—than it would be if it were taxonomic in structure.

C. I. Lewis also recognized this matter of grain but from a slightly different point of view.[21] Transposing his argument into the key of my own analysis, imagine that you point to a box and say of the object inside it, "If that is a raven, then it is black." It turns out that it is not a raven. Nonetheless, adopting the subjunctive mood, you insist that you had uttered something true: "Yes, it is not a raven, but what I said is still true. After all, if it *were* a raven, then it *would* be black." Plausible enough. But why?

One might think that the plausibility we find in continuing to maintain the truth of the conditional derives precisely from the fact that the antecedent is false. After all, per the truth tables for "→," if the antecedent is false, then the conditional is true. However, if this were the case, then *any* conditional starting with "if this is a raven...," independent of its consequent, would be equally supported by the fact that the thing in the box is not a raven. For instance, you could say, "If this is a raven, then my bratty nephew will grow out of it," or "If this is a raven, then the sky will fall." If you understood the relation between ravenhood and blackness asserted in your initial conditional to be merely that summed up in the truth tables for "→," you would be as inclined to come to the defense of any of these conditionals, given the falsity of the antecedent, as you in fact were in coming to the defense of the conditional you actually uttered. In fact, there would be no reason for you to even privilege the conditional you actually uttered over any other one with the same antecedent or, further, to limit what you take to be justified in asserting to any one such conditional.

However, clearly, no actual person would respond to realizing that the antecedent of a conditional they just uttered is false by unspooling an infinite list of conditionals with the same antecedent. That would

---

[21] C. I. Lewis, *An Analysis of Knowledge and Valuation* (La Salle, IL: Open Court, 1946), pp. 524–32.

reflect a gross malfunction of human cognition. Thankfully, our actual understanding of the conditional is, as a matter of fact, more constrained—more *intelligent*—than that. The plausibility of the defense you gave of your initial statement—the plausibility we find in the idea that the conditional holds even when its antecedent is contrary to fact—relies on more than simply the fact that its antecedent turned out to be false. The relationship between ravenhood and blackness asserted in "if that is a raven, then it is black" is therefore narrower or more fine-grained than what "→" means per its truth table.

Lewis takes this to show that our understanding of generalizations like the ones at hand involves the idea of a "real connection."[22] For Lewis, the fact that we understand there to be a "real connection" between ravenhood and blackness is what explains why, after finding the antecedent to be false, you would be willing to defend your initial claim, "If that is a raven, then it is black," as still true but *not even think* to defend "If that is a raven, then my bratty nephew will grow out of it." We do not understand there to be the sort of connection between ravenhood and our nephew's unfortunate behavior that would have to be in place in order for such a conditional to even occur to us as relevant under these circumstances. Such a connection *is* in place between ravenhood and blackness, however—or so we understand it to be according to Lewis.

For Lewis, this "real connection" is what is not countenanced by the formal explication $\forall x(Rx \to Bx)$ of "ravens are black." For *me*, this difference between understanding the generalization as involving a "real connection" between ravenhood and blackness and understanding that relation as consisting in no more than what the formal "→" captures, *marks the difference between the nomological and taxonomic hypotheses.* On the taxonomic understanding, the relation between ravenhood and blackness expressed in "ravens are black" is merely "→." Likewise, on the taxonomic understanding, the relation between a raven and its blackness expressed in Ba, where a is a particular raven, is merely *predication*—precisely the same relation understood, taxonomically, to

---

[22] Here is a representative quote from Lewis expressing his core idea: "The relation of '*A*' to '*E*' in "If *A* then *E*" is *not* justly interpreted as the relation of material implication which many current developments of logic take as fundamental; it is *not* what is called in *Principia Mathematica* a formal implication; and it is *not* a strict implication or entailment such that '*E*' is, assertedly, deducible from '*A*'. It *is* generically the same kind of relation which Hume had in mind when he spoke of 'necessary connections of matters of fact'. And Hume, as we remember, affirmed that no such relation can be established as holding; though, as we also remember, he made no doubt that it is such a relation which is *asserted* as often as any objective fact or anything going beyond presentations of sense is taken to be empirically known. And whatever objections may be raised against Hume's analysis on other points, on this point of meaning he was exactly right." *Ibid.*, p. 212.

obtain between a given piece of coal and its blackness as expressed in
Bc, where c is a particular piece of coal. Both "predication" and "$\rightarrow$"
are *thin*, indifferent formal relations compared to the *thick* relation
of "real connection" that the nomological understanding takes to tie
objects and their properties together.

However, while I am decidedly sympathetic to Lewis's view, the mere
presence of "real connection" in our intuitive understanding of such
generalizations, thick though it may be, is insufficient to explain the
fact that we differentially diagnose white ravens and white pieces of
coal. Perhaps we understand there to be a "real connection" between
ravenhood and blackness, on the one hand, and coalhood and black-
ness on the other, but the fact that we would not think to diagnose
a white piece of coal with albinism tells us that our intuitive under-
standing recognizes *differences between* the manners in which they are
really connected. This is why, I submit, our intuitive pre-theoretical
understanding of the relationship between ravenhood and blackness
consists not merely in their being really connected but, rather, in a
sense of the *characteristic means* by which ravens are black. Mirroring
the park bench example from section I, in confronting a white raven,
our intuitive diagnosis is that there has been a failure with respect to
the *indeterminate yet characteristic means by which ravens are black*. Our
understanding of these means differs from our understanding of the
indeterminate yet characteristic means by which coal is black that is con-
stitutive of our pre-theoretical intuitive understanding of the relation
between coalhood and blackness. Albinism, we understand, is just not
a way in which the means by which coal is black can break down and so
it just would not occur to us to go there diagnostically in confronting
a white piece of coal.

Thus, our intuitive, *mechanical*, pre-theoretical understanding of the
relation between ravenhood and blackness expressed in the general-
ization "ravens are black" takes it to mean that *ravens are black by some
indeterminate yet characteristic means*. To put this, with some reluctance,
in terms that more closely mirror the structure of the relevant for-
mal sentence, $\forall x(Rx \rightarrow Bx)$ is intuitively understood to mean *ceteris
paribus—for all these xs, they are black by some indeterminate yet charac-
teristic means, namely, ravenhood.*[23] This same analysis applies to "partic-

---

[23] For an explanation of my reluctance, see footnote 12. Also, note that a primary
point of my discussion is to point out how the formal explication of the structure of
a sentence like "ravens are black" does not mirror the structure of our intuitive un-
derstanding of it. It should not be surprising, then, that attempting to articulate the
structure of our intuitive understanding of such sentences in terms that mirror the
structure of the relevant formal sentence will yield an awkward locution.

ular" statements as well. That is, rather than understanding it to be one of "predication," our intuitive pre-theoretical understanding of the relation between a particular raven and its blackness expressed by a sentence like "this black thing is a raven" is that this thing is black by some indeterminate yet characteristic means, namely, ravenhood.

In my introduction, I mentioned that this should be understood on analogy with how Nancy Cartwright has argued we ought to understand causal concepts. While I wish to remain agnostic with respect to her conclusions regarding causality, I borrow her terminology and maintain that, in the same sense as that in which Cartwright thinks causal concepts are thick, the *means by which* ravens are black that is involved in our pre-theoretical, intuitive understanding of "ravens are black" is thick. However, I am not merely importing the idea from Cartwright. Our inquiry has given us resources to provide further analysis of thickness itself.

To see this, let us look at Cartwright's view on the matter of "thickness." In writing on causation, she argues that "accounts that suppose that there is one thing—one characteristic feature—that makes a law a *causal* law" are mistaken. She offers an alternative:

> Just as there is an untold variety of quantities that can be involved in laws, so too there is an untold variety of causal relations....that we represent most immediately using content-rich causal verbs: the pistons *compress* the air in the carburetor chamber, the sun *attracts* the planets, the loss of skill among long-term unemployed workers *discourages* firms from opening new jobs....These are genuine facts, but more concrete than those reported in claims that use only the abstract vocabulary of 'cause' and 'prevent'.[24]

The theories Cartwright sets herself up against—ones that explicate causality in univocal terms of, for instance, Bayes-nets—overlook the distinctions captured by our rich vocabulary of causal verbs, and in the process, "[W]e...lose a vast amount of information that we otherwise possess..."[25] To illustrate, Cartwright considers Judea Pearl's discussion of

> an experiment in which soil fumigants, X, are used to increase oat crop yields, Y, by controlling the eelworm population, Z, but may also have direct effects, both beneficial and adverse, on yields beside the control of eelworms....[F]armers' choice of treatment depends on last year's

---

[24] Cartwright, "Causation: One Word, Many Things," *op. cit.*, p. 814.
[25] *Ibid.*, pp. 814–15.

eelworm population, Z_0....The quantities Z_1, Z_2, and Z_3 denote, respectively, the eelworm population, both size and type, before treatment, after treatment, and at the end of the season....B, the population of birds and other predators.[26]

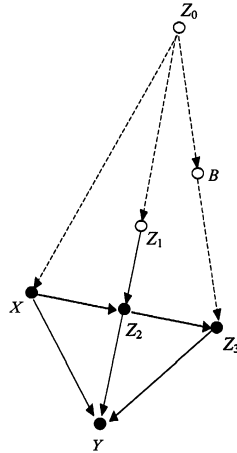The causal diagram that Pearl uses to represent this situation is presented in Figure 2:

Figure 2. "A causal diagram representing the effect of fumigants, X, on yields, Y. Variables: X: fumigants; Y: yields; B: the population of birds and other predators; Z_0: last year's eelworm population; Z_1: eelworm population before treatment; Z_2: eelworm population after treatment; Z_0: eelworm population at the end of the season."[27]

Here is the "thicker" description that Cartwright says we could give of the causal laws operating in this experiment:

Perhaps the soil fumigant *poisons* the infant eelworms, or perhaps it *smothers* the eelworm eggs....Perhaps the fumigant *enriches* the soil or *clogs* the roots. Instead Pearl gives an even thinner description. He replaces all the thick descriptions by one single piece of notation—the arrow.[28]

For Cartwright, moving from a description involving multiple, distinct, content-rich causal verbs to a univocal description given in terms of the formal language of the diagram is a move from a thick

---

[26] Judea Pearl, "Causal Diagrams for Empirical Research," *Biometrika*, LXXXII, 4 (December 1995): 669–88 at p. 669.

[27] Cartwright, "Causation," *op. cit.*, p. 816.

[28] *Ibid.*

to a thin description of the relevant causal reality. Likewise, I claim, when we move from our *thick* pre-theoretical and *mechanical* understanding of "ravens are black" to the formal $\forall x(Rx \rightarrow Bx)$, we lose thickness. However, in light of the preceding analysis, we can say even more than that.

When we move from our mechanical, pre-theoretical understanding of the generalization to $\forall x(Rx \rightarrow Bx)$, the *means by which* ravens are black, in terms of which we understand the relationship between ravenhood and blackness intuitively, drop out.[29] What is abstracted away from when we move from Cartwright's thick description to Pearl's thin one, however, is different. The relations specified by the arrows in Pearl's diagram are indeed thin, as Cartwright states, but in a specific sense, namely, that they abstract away from *differences between* the means by which the realities represented in the diagram affect one another—differences captured by verbs like "poisons," "smothers," "enriches," and "clogs." This leaves the relations between them understood in the manner of C. I. Lewis's homogenous notion of "real connection" discussed above. In other words, the arrows in the diagram capture precisely what Lewis points out is missing in the semantics of "$\rightarrow$" as defined by its truth table—hence the causal graph's ability to support counterfactual reasoning about the system in a way that our discussion of Lewis showed that material implication cannot. However, as Cartwright is keen to point out, the graph captures nothing more than this. Specifically, in my terms, it does not capture the *characteristic* means by which the elements of the system affect one another.

Thus, the move from thick to thin formulations can occur in two ways, each corresponding to a different dimension of thickness. We could drop the characteristic-ness of the relevant means, as we do when we move from our intuitive pre-theoretical understanding of "ravens are black" to the nomological understanding that takes ravenhood and blackness to be merely "really connected" by a counterfactual supporting law. This is analogous to the move from Cartwright's thick description of the system to Pearl's thin representation. On the

---

[29] Recalling the discussion of C. I. Lewis above, this is precisely why we find it counterintuitive to think that the falsity of a conditional's antecedent justifies believing any conditional with that antecedent whatsoever. Logic tells us that we are so justified:

(1) P $\rightarrow$ Q   (2) Not-P   (3) Therefore, P $\rightarrow$ Ψ,

where Ψ is any sentence letter, is a deductively valid inference. Abstracting away from *means* ("real connection" if you are Lewis) is precisely what enables (3) to follow from (1) and (2). Yet, it is precisely for this reason that the idea that (3) follows from (1) and (2) is counterintuitive.

other hand, we could drop the relevant means altogether as we do when we move from our intuitive understanding of "ravens are black" to $\forall x(Rx \rightarrow Bx)$. This is analogous to moving from Cartwright's thick description of the system to a version of Pearl's diagram in which the arrows in the diagram are understood as "$\rightarrow$" per its truth table.[30]

Both of these are different than the move from determinate to indeterminate means. Imagine we had an exhaustive account of the ontogenetic mechanism by which ravens are black—every detail of every gene and every turn of every biological pathway from genotype to phenotype laid out in perfectly *determinate* fashion before us.[31] The move from that description to the *indeterminate yet characteristic means by which ravens are black* that is involved in how we intuitively understand "ravens are black" is not a move from thick to thin. It is a move from *determinate* to *indeterminate*, with no "thickness" being stripped away in the process.

With all the pieces on the table, we are now in a position to understand the relationship between the various hypotheses regarding how we understand generalizations like "ravens are black" that I have distinguished throughout the paper—the taxonomic, the nomological, and the mechanical. The mechanical understanding of "ravens are black" takes the generalization to mean that *ravens are black by some indeterminate yet characteristic means.* As we move from this mechanical understanding to the nomological understanding, we retain the *means* but lose their characteristic-ness. As a result, we end up with an understanding of "ravens are black" that involves the idea of "real connection" between ravenhood and blackness but a "real connection" that is not distinct from that between, say, coalhood and blackness that is involved in a nomological understanding of "coal is black." As we move from the mechanical to the taxonomic we lose *means* altogether, determinate or indeterminate, characteristic or otherwise. This corresponds to the move from our intuitive pre-theoretical understanding of "ravens are black" to $\forall x(Rx \rightarrow Bx)$. Thus, the move from the

---

[30] Perhaps it is helpful to think about this on analogy with pain. We could think of pain as having two dimensions: (1) the qualitative "painfulness" and (2) "aversiveness." You can get rid of the "aversiveness" by taking an opiate, but you are still left with pain. If you just got rid of the pain, trivially, you would get rid of the aversiveness. Likewise, you can think of "thickness" as having two dimensions: (1) means and (2) characteristic-ness. You can get rid of characteristic-ness and be left with means like in Pearl's diagram. If you got rid of means altogether, you would, trivially, lose characteristic-ness.

[31] The ontogenetic mechanism in the example could be replaced with a phylogenetic one like natural selection. Imagine we had an exhaustive account of the phylogenetic mechanism by which ravens are black—every turn of the actual, historical evolutionary trajectory laid out before us in perfectly *determinate* fashion...

mechanical to the nomological, on the one hand, and from the mechanical to the taxonomic, on the other, are moves made in different directions—the former from characteristic to non-characteristic means, the latter from means-involving to not-means-involving. Both moves, however, take their departure from the pre-theoretical and intuitive point of view that, as I have argued, is a *mechanical*, means-involving one.[32]

So where does this leave us with respect to Hempel's Paradox? The problem is one that arises in an attempt to squeeze our intuitive, "thick," means-involving understanding of generalizations like the ones at hand into the "thin" formal apparatus of predicate logic, and anyone who has seen the comedy classic *Tommy Boy* knows what happens when you try to squeeze a fat man into a little coat—it rips.

ANDREW BOLLHAGEN

University of California, San Diego

---

[32] Hume thinks we understand the world nomologically after custom and habit have had their effect. Hume would thus endorse the nomological hypothesis. He does not think this is a *justified* worldview, but he thinks it is the one we have nonetheless. But Hume is fundamentally incorrect. We do not understand the world nomologically; we understand it mechanically. As Hume's problem of induction, and his skeptical solution to it, depend for their formulation upon the substantive but, as I argue, ultimately false hypothesis that we understand the world nomologically, my analysis has, I believe, broader implications for our understanding of induction and its classic problem. I leave this for future work.