6
Confirmation

When evidence supports a hypothesis, philosophers of science say that the evidence "confirms" that hypothesis. Bayesians place this confirmation relation at the center of their theory of induction. But confirmation is also closely tied to such epistemological notions as justification and reasons. Bayesian epistemology offers a systematic theory of confirmation (and its opposite, disconfirmation) that not only deepens our understanding of this relation but also provides specific answers about which hypotheses are supported (and to what degree) in particular evidential situations.

Since its early days, the analysis of confirmation has been driven by a perceived analogy to deductive entailment. In Chapter 4 we discussed epistemic standards that relate a body of evidence (represented as a proposition) to the doxastic *attitudes* it supports. But confirmation—though intimately linked with epistemic standards in ways we'll presently see—is a different kind of relation: instead of relating a proposition and an attitude, it relates two propositions (evidence and hypothesis). Confirmation shares this feature with deductive entailment. In fact, Rudolf Carnap thought of confirmation as a generalization of traditional logical relations, with deductive entailment and refutation as two extremes of a continuous confirmational scale.

In the late nineteenth and early twentieth centuries, logicians produced ever-more-powerful syntactical theories capable of answering specific questions about which propositions deductively entailed which. Impressed by this progress, theorists such as Carl Hempel and Carnap envisioned a syntactical theory that would do the same for confirmation. As Hempel put it:

> The theoretical problem remains the same: to characterize, in precise and general terms, the conditions under which a body of evidence can be said to confirm, or to disconfirm, a hypothesis of empirical character. (1945a, p. 7)

Hempel identified various formal properties that the confirmation relation might or might not possess. Carnap then argued that we get a confirmation relation with exactly the right formal properties by identifying confirmation with positive probabilistic relevance.

Fundamentals of Bayesian Epistemology 2: Arguments, Challenges, Alternatives. Michael G. Titelbaum, Oxford University Press. DOI: 10.1093/oso/9780192863140.003.0006

This chapter begins with Hempel's formal conditions on the confirmation relation. Identifying the right formal conditions for confirmation will not only help us assess various theories of confirmation; it will also help us understand exactly what relation philosophers of science have in mind when they talk about "confirmation".[1] We then move on to Carnap's Objective Bayesian theory of confirmation, which grounds confirmation in probability theory. While Carnap's theory has a number of attractive features, we will also identify two drawbacks: its failure to capture particular patterns of inductive inference that Carnap found appealing, and the language-dependence suggested by Goodman's "grue" problem. We'll respond to these problems with a confirmation theory based on Subjective Bayesianism (in the normative sense).

Confirmation is fairly undemanding, in one sense: we say that evidence confirms a hypothesis when it provides *any* amount of support for that hypothesis, no matter how small. But we might want to distinguish bodies of evidence that *strongly* support hypotheses from those that provide only weak support. Probabilistic theories of confirmation offer a number of different ways to measure the strength of confirmation in any particular case. In Section 6.4.1 we'll survey these different measures of confirmational strength, assessing the pros and cons of each. Finally, we'll apply probabilistic confirmation theory to provide a Bayesian solution to Hempel's Paradox of the Ravens.

6.1 Formal features of the confirmation relation

6.1.1 Confirmation is weird! The Paradox of the Ravens

One way to begin thinking about confirmation is to consider the simplest possible cases in which a piece of evidence confirms a general hypothesis. For example, the proposition that a particular frog is green seems to confirm the hypothesis that all frogs are green. On the other hand, the proposition that a particular frog is not green disconfirms the hypothesis that all frogs are green. (In fact, it *refutes* that hypothesis!) If we think this pattern always holds, we will maintain that confirmation satisfies the following constraint:

Nicod's Criterion: For any predicates F and G and constant a of $\mathcal{L}$, $(\forall x)(Fx \supset Gx)$ is confirmed by $Fa \,\&\, Ga$ and disconfirmed by $Fa \,\&\, {\sim}Ga$.

Hempel (1945a,b) named this condition after Jean Nicod (1930), who built his theory of induction around the criterion. We sometimes summarize the

Nicod Criterion by saying that a universal generalization is confirmed by its **positive instances** and disconfirmed by its **negative instances**. Notice that one can endorse the Nicod Criterion as a sufficient condition for confirmation without taking it to be necessary; we need not think *all* cases of confirmation follow this pattern.

Yet Hempel worries about the Nicod Criterion even as a sufficient condition for confirmation, because of how it interacts with another principle he endorses:

Equivalence Condition (for hypotheses): Suppose H and H' in $\mathcal{L}$ are logically equivalent ($H \dashv\vDash H'$). Then any E in $\mathcal{L}$ that confirms H also confirms H'.

Hempel endorses the Equivalence Condition because he doesn't want confirmation to depend on the particular way a hypothesis is formulated; logically equivalent hypotheses say the same thing, so they should enter equally into confirmation relations. Hempel is also concerned with how working scientists *use* confirmed hypotheses; for instance, practitioners will often deduce predictions and explanations from confirmed hypotheses. Equivalent hypotheses have identical deductive consequences, and scientists don't hesitate to substitute logical equivalents for each other.

But combining Nicod's Criterion with the Equivalence Condition yields counterintuitive consequences, which Hempel calls the "paradoxes of confirmation". The most famous of these is the **Paradox of the Ravens**. Consider the hypothesis that all ravens are black, representable as $(\forall x)(Rx \supset Bx)$. By Nicod's Criterion this hypothesis is confirmed by the evidence that a particular raven is black, $Ra \,\&\, Ba$. But now consider the evidence that a particular non-raven is non-black, ${\sim}Ba \& {\sim}Ra$. This is a positive instance of the hypothesis $(\forall x)({\sim}Bx \supset {\sim}Rx)$, so by Nicod's Criterion it confirms that hypothesis. By contraposition, that hypothesis is equivalent to the hypothesis that all ravens are black. So by the Equivalence Condition, ${\sim}Ba \,\&\, {\sim}Ra$ confirms $(\forall x)(Rx \supset Bx)$ as well. The hypothesis that all ravens are black is confirmed by the observation of a red herring, or a white shoe. This result seems counterintuitive, to say the least.

Nevertheless, Hempel writes that "the impression of a paradoxical situation... is a psychological illusion" (1945a, p. 18); on his view, we reject the confirmational result because we misunderstand what it says. Hempel notes that in everyday life people make confirmation judgments relative to an extensive corpus of background knowledge. For example, a candidate's performance in an interview may confirm that she'd be good for the job, but only relative to a great deal of background information about how the questions asked relate to the job requirements, how interview performance predicts job performance, etc. In assessing confirmation, then, we should always be explicit about the background we're assuming. This is especially important because background knowledge can dramatically alter confirmation relations. For example, in Section 4.3 we discussed a poker game in which you receive the cards that will make up your hand one at a time. At the beginning of the game, your background knowledge contains facts about how a deck is constructed and about which poker hands are winners. At that point the proposition that your last card will be the two of clubs does not confirm the proposition that you will win the hand. But as the game goes along and you're dealt some other twos, your total background knowledge changes such that the proposition that you'll receive the two of clubs now strongly confirms that you'll win.

Nicod's Criterion does not mention background corpora, so it's tempting to apply the criterion with full generality—that is, regardless of which corpus which use. Yet there are clearly some evidential situations in which a positive instance does *not* confirm a universal generalization. For example, suppose I know I'm in the Hall of Atypically Colored Birds. A bird is placed in the Hall only if the majority of his species-mates are one color but he happens to be another color. If my background corpus includes that I'm in the Hall of Atypically Colored Birds, then observing a black raven *disconfirms* the hypothesis that all ravens are black.[2] Because of such examples, Hempel thinks the Nicod Criterion should be applied only against a **tautological background**. A tautological background corpus contains no contingent propositions; it is logically equivalent to a tautology T.

When we intuitively reject the Nicod Criterion's consequence that a red herring confirms the ravens hypothesis, we are sneaking non-tautological information into the background of our judgment. Hempel thinks we're imagining a situation in which we already know in advance (as part of the background) that we will be observing a herring and checking its color. Relative to that background—which includes the information $\sim Ra$—we know that whatever we're about to observe will have no evidential import for the hypothesis that ravens are black. So when we then get the evidence that $\sim Ba$, that evidence is confirmationally inert with respect to the hypothesis $(\forall x)(Rx \supset Bx)$.

But the original question was whether $\sim Ba$ & $\sim Ra$ (taken all together, at once) confirmed $(\forall x)(Rx \supset Bx)$. On Hempel's view, this is a fair test of the Nicod Criterion only against an empty background corpus (since that's the background against which he thinks the Criterion applies). And against that corpus, Hempel thinks the confirmational result is correct. Here's a way of understanding why: Imagine you've decided to test the hypothesis that all ravens are black. You will do this by selecting objects from the universe one at a time and checking them for ravenhood and blackness. It's the beginning of the experiment, you haven't checked any objects yet, and you have no background information about the tendency of objects to be ravens and/or black. Moreover, you've found a way to select objects from the entire universe at random, so you have no background information about what kind of object you'll be getting. Nevertheless, you start thinking about what sorts of objects might be selected, and whether they would be good or bad news for the hypothesis. Particularly important would be any ravens that weren't black, since any such negative instance would immediately refute the hypothesis. (Here it helps to realize that the ravens hypothesis is logically equivalent to $\sim(\exists x)(Rx \ \& \ \sim Bx)$.) So when the first object arrives and you see it's a red herring—$\sim Ba \& \sim Ra$—this is good news for the hypothesis (at least, *some* good news). After all, the first object could've been a non-black raven, in which case the hypothesis would've been sunk.

This kind of reasoning defuses the seeming paradoxicality of a red herring's confirming that all ravens are black, and the objection to the Nicod Criterion that results. As long as we're careful not to smuggle in illicit background information, observing a red herring confirms the ravens hypothesis to at least a small degree. Hempel concludes that on his reading of the Nicod Criterion, it is consistent with the Equivalence Condition.

Nevertheless, I.J. Good worries about the Nicod Criterion, even against a tautological background:

> [T]he closest I can get to giving [confirmation relative to a tautological background] a practical significance is to imagine an infinitely intelligent newborn baby having built-in neural circuits enabling him to deal with formal logic, English syntax, and subjective probability. He might now argue, after defining a crow in detail, that it is initially extremely unlikely that there are any crows, and therefore that it is extremely likely that all crows are black. "On the other hand," he goes on to argue, "if there are crows, then there is a reasonable chance that they are of a variety of colors. Therefore, if I were to discover that even a black crow exists I would consider [the hypothesis that all crows are black] to be less probable than it was initially." (1968, p. 157)[3]

Here Good takes advantage of the fact that $(\forall x)(Rx \supset Bx)$ is true if there are no ravens (or crows, in his example).[4] Before taking any samples from the universe, the intelligent newborn might consider four possibilities: there are no ravens; there are ravens but they come in many colors; there are ravens and they're all black; there are ravens and they all share some other color. The first and third of these possibilities would make $(\forall x)(Rx \supset Bx)$ true. When the baby sees a black raven, the first possibility is eliminated; this might be a large enough blow to the ravens hypothesis that the simultaneous elimination of the fourth possibility would fail to compensate. In other words, the observation of a black raven might not confirm that all ravens are black, violating the Nicod Criterion even against a tautological background.

6.1.2 Further adequacy conditions

We have already seen two general conditions (Nicod's Criterion and the Equivalence Condition) that one might take the confirmation relation to satisfy. We will now consider a number of other such conditions, most of them discussed (and given the names we will use) by Hempel. Sorting out which of these are genuine properties of confirmation has a number of purposes. First, Hempel thought the correct list provided a set of adequacy conditions for any positive theory of confirmation. Second, sorting through these conditions will help us understand the abstract features of evidential support. These are features about which epistemologists, philosophers of science, and others (including working scientists and ordinary folk!) often make strong assumptions—many of them incorrect. Finally, we are going to use the word "confirmation" in subsequent sections as a somewhat technical term, distinct from some of the ways "confirm" is used in everyday speech. Working through the properties of the confirmation relation will help illustrate how we're using the term.

The controversy between Hempel and Good leaves it unclear whether the Nicod Criterion should be endorsed as a constraint on confirmation, even when it's restricted to the tautological background. On the other hand, the Equivalence Condition can be embraced in a highly general form:

Equivalence Condition (full version): For any H, H', E, E', K and K' in $\mathcal{L}$, suppose $H \mathrel{=}\joinrel\mathrel{|}\joinrel\mathrel{|}\joinrel\mathrel{=} H'$, $E \mathrel{=}\joinrel\mathrel{|}\joinrel\mathrel{|}\joinrel\mathrel{=} E'$, and $K \mathrel{=}\joinrel\mathrel{|}\joinrel\mathrel{|}\joinrel\mathrel{=} K'$. Then E confirms (/disconfirms) H relative to background K just in case E' confirms (/disconfirms) H' relative to background K'.

Here we can think of K as a conjunction of all the propositions in an agent's background corpus, just as E is often a conjunction of multiple pieces of evidence. This full version of the Equivalence Condition captures the idea that logically equivalent propositions should enter into all the same confirmation relations.

Our next candidate constraint is the

Entailment Condition: For any consistent E, H, and K in $\mathcal{L}$, if $E \,\&\, K \vDash H$ but $K \nvDash H$, then E confirms H relative to K.

This condition captures the idea that entailing a hypothesis is one way to support, or provide evidence for, that hypothesis. If E entails H in light of background corpus K (in other words, if E and K together entail H), then E confirms H relative to K. The only exception to this rule is when K already entails H, in which case the fact that E and K together entail H does not indicate any particular relation between E and H.[5] Notice that a tautological H will be entailed by every K, so the restriction on the Entailment Condition prevents the condition from saying anything about the confirmation of tautologies. Hempel thinks of his adequacy conditions as applying only to empirical hypotheses and bodies of evidence, so he generally restricts them to logically contingent Es and Hs.

Hempel considers a number of adequacy conditions motivated by the following intuition:

Confirmation Transitivity: For any A, B, C, and K in $\mathcal{L}$, if A confirms B relative to K and B confirms C relative to K, then A confirms C relative to K.

It's tempting to believe confirmation is transitive, as well as other nearby notions such as justification or evidential support. This temptation is buttressed by the fact that logical entailment is transitive. Confirmation, however, is not in general transitive. Here's an example of Confirmation Transitivity failure: Suppose our background is the fact that a card has just been selected at random from a standard fifty-two-card deck. Consider these three propositions:

A: The card is a jack.
B: The card is the jack of spades.
C: The card is a spade.

Relative to our background, A would confirm B, at least to some extent. And relative to our background, B would clearly confirm C (in accordance with the

Entailment Condition). But relative to the background that a card was picked from a fair deck, *A* does nothing to support conclusion *C*.

The failure of Confirmation Transitivity has a number of important consequences. First, it explains why in the study of confirmation we take evidence to be propositional rather than objectual. In everyday language we often use "evidence" to refer to objects rather than propositions; police don't store propositions in their Evidence Room. But as possible entrants into confirmation relations, objects have an ambiguity akin to the reference class problem (Section 5.1.1). Should I consider this bird evidence that all ravens are black (against tautological background)? If we describe the bird as a black raven, the answer might be yes. But if we describe it as a black raven found in the Hall of Atypically Colored Birds, the answer seems to be no. Yet a black raven in the Hall of Atypically Colored birds is still a black raven. If confirmation were transitive, knowing that a particular description of an object confirmed a hypothesis would guarantee that more precise descriptions confirmed the hypothesis as well. Logically stronger descriptions (it's a black raven in the Hall of Atypically Colored Birds) entail logically weaker descriptions (it's a black raven) of the same object; by the Entailment Condition, the logically stronger description confirms the logically weaker; so if confirmation were transitive anything confirmed by the weaker description would be confirmed by the stronger as well.

But confirmation isn't transitive, so putting more or less information in our description *of the very same object* can alter what's confirmed. (Black raven? Might confirm ravens hypothesis. Black raven in Hall of Atypically Colored Birds? Disconfirms. Black raven mistakenly placed in the Hall of Atypically Colored Birds when it shouldn't have been? Perhaps confirms again.) We solve this problem by letting propositions rather than objects enter into the confirmation relation. If we state our evidence as a proposition—such as the proposition *that a black raven is in the Hall of Atypically Colored Birds*—there's no question how the objects involved are being described.[6]

Confirmation's lack of transitivity also impacts epistemology more broadly. For instance, it may cause trouble for Richard Feldman's (2007) principle that "evidence of evidence is evidence". Suppose I read in a magazine that anthropologists have reported evidence that Neanderthals cohabitated with *homo sapiens*. I don't actually have the anthropologists' evidence for that hypothesis—the body of information that they think supports it. But the magazine article constitutes evidence that they have such evidence; one might think that the magazine article therefore also constitutes evidence that Neanderthals and *homo sapiens* cohabitated. (After all, reading the article seems to provide me with some justification for that hypothesis.) Yet we cannot adopt this "evidence of evidence is evidence" principle with full generality. Suppose I've randomly picked a card from a standard deck and examined it carefully. If I tell you my card is a jack, you have evidence for the proposition that I know my card is the jack of spades. If I know my card is the jack of spades, I have (very strong) evidence for the proposition that my card is a spade. So your evidence (jack) is evidence that I have evidence (jack of spades) for the conclusion spade. Yet your evidence is no evidence at all that my card is a spade.[7]

Finally, the failure of Confirmation Transitivity shows what's wrong with two confirmation constraints Hempel embraces:

Consequence Condition: If E in $\mathcal{L}$ confirms every member of a set of propositions relative to K and that set jointly entails H' relative to K, then E confirms H' relative to K.

Special Consequence Condition: For any E, H, H', and K in $\mathcal{L}$, if E confirms H relative to K and $H \,\&\, K \vDash H'$, then E confirms H' relative to K.

The Consequence Condition states if a set of propositions together entails a hypothesis (relative to some background), then any evidence that confirms every member of the set also confirms that hypothesis (relative to that background). The Special Consequence Condition says that if a *single* proposition entails some hypothesis, then anything that confirms the proposition also confirms the hypothesis (again, all relative to some background corpus).

The Special Consequence Condition is so-named because it can be derived from the Consequence Condition (by considering singleton sets). Yet each of these conditions is a bad idea, as can be demonstrated by our earlier jack of spades example. Proposition A that the card is a jack confirms proposition B that it's the jack of spades; B entails proposition C that the card is a spade; yet A does not confirm C. We can even create examples in which H entails H' relative to K, but evidence E that confirms H *dis*confirms H' relative to K. Relative to the background corpus K most of us have concerning the kinds of animals people keep as pets, the evidence E that Bob's pet is hairless confirms (at least slightly) the hypothesis H that Bob's pet is a Peruvian Hairless Dog. Yet relative to K that same evidence E disconfirms the hypothesis H' that Bob's pet is a dog.[8]

Why might the Special Consequence Condition seem plausible? It certainly looks tempting if one reads "confirmation" in a particular way. In everyday language it's a fairly strong claim that a hypothesis has been "confirmed"; this suggests our evidence is sufficient for us to accept the hypothesis. (Consider the sentences "That confirmed my suspicion" and "Your reservation has been confirmed.") If we combine that reading of "confirmed" with Glymour's position that "when we accept a hypothesis we commit ourselves to accepting all of its logical consequences" (1980, p. 31), we get that evidence confirming a hypothesis also confirms its logical consequences, as the Special Consequence Condition requires. But hopefully the discussion to this point has indicated that we are not using "confirms" in this fashion. On our use, evidence confirms a hypothesis if it provides *any* amount of support for that hypothesis; the support need not be decisive. For us, evidence can confirm a hypothesis without requiring or even permitting acceptance of the hypothesis. If your only evidence about a card is that it's a jack, this evidence *confirms* in our sense that the card is the jack of spades. But this evidence doesn't authorize you to *accept* or *believe* that the card is the jack of spades.

Another motivation for the Special Consequence Condition—perhaps this was Hempel's motivation—comes from the way we often treat hypotheses in science. Suppose we make a set of atmospheric observations confirming a particular global warming hypothesis. Suppose further that in combination with our background knowledge, the hypothesis entails that average global temperatures will increase by five degrees in the next fifty years. It's very tempting to report that the atmospheric observations support the conclusion that temperatures will rise five degrees in fifty years. Yet that's to unthinkingly apply the Special Consequence Condition.

I hope you're getting the impression that denying Confirmation Transitivity can have serious consequences for the ways we think about everyday and scientific reasoning. Yet it's important to realize that denying the Special Consequence Condition as a *general* principle does not mean that the transitivity it posits *never* holds. It simply means that we need to be careful about assuming confirmation will transmit across an entailment, and perhaps also that we need a precise, positive theory of confirmation to help us understand when it will and when it won't.

Rejecting the Special Consequence Condition does open up some intriguing possibilities in epistemology. Consider these three propositions:

E: I am having a perceptual experience as of a hand before me.
H: I have a hand.
H': There is a material world.

This kind of evidence figures prominently in G.E. Moore's (1939) proof of the existence of an external world. Yet for some time it was argued that E could not possibly be evidence for H. The reasoning was, first, that E could not discriminate between H' and various skeptical hypotheses (such as Descartes's evil demon), and therefore could not provide evidence for H'. Next, H entails H', so if E were evidence for H it would be evidence for H' as well. But this step assumes the Special Consequence Condition. Epistemologists have recently explored positions that allow E to support H without supporting H', by denying Special Consequence.[9]

Hempel's unfortunate endorsement of the Consequence Condition also pushes him toward the

Consistency Condition: For any E and K in $\mathcal{L}$, the set of all hypotheses confirmed by E relative to K is logically consistent with E & K.

In order for the set of all hypotheses confirmed by E to be consistent with E&K, it first has to be a logically consistent set in its own right. So among other things, the Consistency Condition bans a single piece of evidence from confirming two hypotheses that are mutually exclusive with each other.

It seems easy to generate confirmational examples that violate this requirement: evidence that a randomly drawn card is red confirms both the hypothesis that it's a heart and the hypothesis that it's a diamond, but these two confirmed hypotheses are mutually exclusive. Hempel also notes that scientists often find themselves entertaining theoretical alternatives that are mutually exclusive with each other; experimental data eliminates some of those theories while confirming all of the ones that remain. Yet Hempel is trapped into the Consistency Condition by his allegiance to the Consequence Condition. Taken together, the propositions in an inconsistent set entail a contradiction; so any piece of evidence that confirmed all the members of an inconsistent set would also (by the Consequence Condition) confirm a contradiction. Hempel refuses to grant that anything could confirm a contradiction! So he tries to make the Consistency Condition work.[10]

Hempel rightly rejects the

Converse Consequence Condition: For any E, H, H', and K in $\mathcal{L}$ (with H' consistent with K), if E confirms H relative to K and H' & $K \vDash H$, then E confirms H' relative to K.

The Converse Consequence Condition says that relative to a given background, evidence that confirms a hypothesis also confirms anything that entails that hypothesis.

Here's a counterexample. Suppose our background knowledge is that a fair six-sided die has been rolled, and our propositions are:

E: The roll outcome is prime.
H: The roll outcome is odd.
H': The roll outcome is 1.

In this case E confirms H relative to our background, H' entails H, yet E refutes H'. (Recall that 1 is not a prime number!)

Still, there's a good idea in the vicinity of Converse Consequence. Suppose our background consists of the fact that we are going to run a certain experiment. A particular scientific theory, in combination with that background, entails that the experiment will produce a particular result. If this result does in fact occur when the experiment is run, we take that to support the theory. This is an example of the

Converse Entailment Condition: For any consistent E, H, and K in $\mathcal{L}$, if $H \,\&\, K \vDash E$ but $K \nvDash E$, then E confirms H relative to K.

Converse Entailment says that if, relative to a given background, a hypothesis entails some evidence, then that evidence confirms that hypothesis relative to that background. (Again, this condition omits cases in which the background K entails the experimental result E all on its own, because such cases need not reveal any connection between H and E.)

Converse Entailment doesn't give rise to examples like the die roll case above (because in that case E is not *entailed* by either H or H' in combination with K). But because deductive entailment is transitive, Converse Entailment does generate the **problem of irrelevant conjunction**. Consider the following propositions:

E: My pet is a flightless bird.
H: My pet is an ostrich.
H': My pet is an ostrich and beryllium is a good conductor.

Here H entails E, so by the Converse Entailment Condition E confirms H, which seems reasonable.[11] Yet despite the fact that H' also entails E (because H' entails H), it seems worrisome that E would confirm H'. What does my choice in pets indicate about the conductivity of beryllium?

Nothing—and that's completely consistent with the Converse Entailment Condition. Just because E confirms a conjunction one of whose conjuncts concerns beryllium doesn't mean E confirms that beryllium-conjunct all on its own. To assume that it does would be to assume the Special Consequence Condition, which we've rejected. So facts about my pet don't confirm any conclusions that are about beryllium but not about birds. On the other hand, it's reasonable that E would confirm H' to at least some extent, by virtue of eliminating such rival hypotheses as "beryllium is a good conductor and my pet is an iguana."

Rejecting the Special Consequence Condition therefore allows us to accept Converse Entailment. But again, this should make us very careful about how we reason in our everyday lives. A scientific theory, for instance, will often have wide-ranging consequences, and might be thought of as a massive conjunction. When the theory makes a prediction that is borne out by experiment, that experimental result confirms the theory. But it need not confirm the rest of the theory's conjuncts, taken in isolation. In other words, experimental evidence that confirms a theory may not confirm that theory's further predictions.[12]

Finally, we should say something about disconfirmation. Hempel takes the following position:

Disconfirmation Duality: For any E, H, and K in $\mathcal{L}$, E confirms H relative to K just in case E disconfirms $\sim H$ relative to K.

Disconfirmation Duality pairs confirmation of a hypothesis with disconfirmation of its negation. It allows us to immediately convert many of our constraints on confirmation into constraints on disconfirmation. For example, the Entailment Condition now tells us that if E & K deductively refutes H (yet K doesn't refute H all by itself), then E disconfirms H relative to K. (See Exercise 6.2.) We should be careful, though, not to think of confirmation and disconfirmation as exhaustive categories: for many propositions E, H, and K, E will neither confirm nor disconfirm H relative to K.

Figure 6.1 summarizes the formal conditions on confirmation we have accepted and rejected. The task now is to find a positive theory of which bodies of evidence confirm which hypotheses relative to which backgrounds that satisfies the right conditions and avoids the wrong ones.

Name	Brief, somewhat imprecise description	Verdict
Equivalence Condition	equivalent hypotheses, evidence, backgrounds behave same confirmationally	accepted
Entailment Condition	evidence confirms what it entails	accepted
Converse Entailment Condition	a hypothesis is confirmed by what it entails	accepted
Disconfirmation Duality	a hypothesis is confirmed just when its negation is disconfirmed	accepted
Confirmation Transitivity	anything confirmed by a confirmed hypothesis is also confirmed	rejected
Consequence Condition	anything entailed by a set of confirmed hypotheses is also confirmed	rejected
Special Consequence Condition	anything entailed by a confirmed hypothesis is also confirmed	rejected
Consistency Condition	all confirmed hypotheses are consistent	rejected
Converse Consequence Condition	anything that entails a confirmed hypothesis is also confirmed	rejected
Nicod's Criterion	$Fa \,\&\, Ga$ confirms $(\forall x)(Fx \supset Gx)$	???

Figure 6.1 Accepted and rejected conditions on confirmation

6.2 Carnap's theory of confirmation

6.2.1 Confirmation as relevance

Carnap saw that we could get a confirmation theory with exactly the properties we want by basing it on probability. Begin by taking *any* probabilistic distribution Pr over $\mathcal{L}$. (I've named it "Pr" because we aren't committed at this stage to its being any *kind* of probability in particular—much less a credence distribution. All we know is that it's a distribution over the propositions in $\mathcal{L}$ satisfying the Kolmogorov axioms.) Define Pr's background corpus K as the conjunction of all propositions X in $\mathcal{L}$ such that $\Pr(X) = 1$.[13] Given an E and H in $\mathcal{L}$, we apply the Ratio Formula to calculate $\Pr(H \mid E)$. Two distinct theories of confirmation now suggest themselves: (1) E confirms H relative to K just in case $\Pr(H \mid E)$ is high; (2) E confirms H relative to K just in case $\Pr(H \mid E) > \Pr(H)$. In the preface to the second edition of his *Logical Foundations of Probability*, Carnap calls the first of these options a "**firmness**" concept of confirmation and the second an "**increase in firmness**" concept (1962a, p. xvff).[14]

The firmness concept of confirmation has a number of problems. First, there are questions about where exactly the threshold for a "high" value of $\Pr(H \mid E)$ falls, what determines that threshold, how we discover it, etc. Second, there