

Reorienting

6. Reorienting the Logic of Abduction

John Woods

Abduction, still a comparatively neglected kind of premiss-conclusion reasoning, gives rise to the questions I want to consider here. One is whether abduction's epistemic peculiarities can be accommodated happily in the mainline philosophical *theories of knowledge*. The other is whether abduction provides any reason to question the assumption that the goodness of drawing a conclusion from premisses depends on an underlying relation of logical *consequence*. My answer each time is no. I will spend most of my time on the first. Much of what I'll say about the second is a promissory note.

6.1	Abduction	138
6.1.1	Peirce's Abduction	138
6.1.2	Ignorance Problems	138
6.1.3	The Gabbay-Woods Schema	139
6.1.4	The Yes-But Phenomenon	140
6.2	Knowledge	141
6.2.1	Epistemology.....	141
6.2.2	Losing the <i>J</i> -Condition	142
6.2.3	The Causal Response Model of Knowledge	142
6.2.4	Naturalism	143
6.2.5	Showing and Knowing	143
6.2.6	Explaining the Yes-Buts	144
6.2.7	Guessing	144
6.2.8	Closed Worlds.....	146
6.3	Logic	148
6.3.1	Consequences and Conclusions	148
6.3.2	Semantics	148
	References	149

Three facts about today's logic stand out:

1. Never has it been done with such technical virtuosity
2. Never has there been so much of it
3. Never has there been so little consensus about its common subject matters.

It would seem that the more we have of it, the less our inclination to get to the bottom of its sprawlingly incompatible provisions. There is nothing remotely like this in real analysis, particle physics or population genetics. There is nothing like it in the premiss-conclusion reasonings of politics and everyday life. Left undealt with, one might see in logic's indifference to its own rivalries some sign of not quite knowing its own mind.

It could be said that one of logic's more stimulating events in our still-young century is the revival of the idea that it is a universal discipline, that when all is said and done there is a core structure to which all the multiplicities of our day are ultimately answerable. If the historical record is anything to go on, the cornerstone

of that core structure is the relation of *logical consequence*. It occasions some sensible operational advice: If in your work you seek to *enlarge* logic's present multiplicities, have the grace to say why you think it qualifies as logic, that is, embodies logic's structural core. This is not idle advice. I hope to give it heed in the pages to follow, as we turn our attention to the logic of abduction.

Although logic's dominant focus has been the consequence relation, in the beginning its centrality owed comparatively little to its *intrinsic* appeal. Consequence was instrumentally interesting; it was thought to be the relation in virtue of which premiss-conclusion reasoning is safe, or whose absence would expose it to risk. Reasoning in turn had an *epistemic* motivation. Man may be many kinds of animal, but heading the list is his cognitive identity. He is a knowledge-seeking and knowledge-attaining being to which his survival and prosperity are indissolubly linked, indispensable to which is his capacity to adjust what he believes to what follows from what. We might say then that as

long as logic has retained its interest in good and bad reasoning it has retained this same epistemic orientation. Accordingly, a logic of good and bad reasoning carries *epistemological* presuppositions that aren't typically explicitly developed.

It would be premature to say that abduction by now has won a central and well-established place in the research programs of modern logic, but there are some hopeful signs of progress (important sources include [6.1–13]). In the literature to date there are two main theoretical approaches, each emphasizing the different sides of a product-process distinction. The logical (or product) approach seeks for truth conditions on abductive consequence relations and of such other properties as may be interdefinable with it. The computational (or process) approach constructs computational models of how hypotheses are selected for use in ab-

ductive contexts. It is not a strict partition. Between the logical and computational paradigms, abductive logic programming and semantic tableaux abduction occupy a more intermediate position. Whatever its precise details, the logic-computer science dichotomy is not something I welcome. It distributes the theory of abductive reasoning into different camps that have yet to learn how to talk to one another in a systematic way. A further difficulty is that whereas abduction is now an identifiable research topic in logic – albeit a minority one – it has yet to attain that status in computer science. Such abductive insights as may occur there are largely in the form of *obiter dicta* attached to the main business at hand (I am indebted to Atocha Aliseda for insightful advice on this point). This leaves us awkwardly positioned. The *foundational* work for a comprehensive account of abductive reasoning still awaits completion.

6.1 Abduction

6.1.1 Peirce's Abduction

Although there are stirrings of it in Aristotle's notion of *apagogē* [6.14], we owe the modern idea of abduction to Peirce. It is encapsulated in the *Peircean abduction schema*, as follows [6.15, CP 5.189]:

“The surprising fact C is observed.
But if A were true, C would be a matter of course.
Hence there is reason to suspect that A is true.”

Peirce's schema raises some obvious questions. One is how central to abduction is the factor of surprise. Another is the issue of how we are to construe the element of suspicion. A third concerns what we are expected to do with propositions that creep thus into our suspicions. A fourth is what we are to make of the idea that an occurrence of something is a matter of course. Like so many of his better ideas and deeper insights, Peirce has nothing like a fully developed account of abduction. Even so, the record contains some important ideas, seven of which I'll mention here:

- P1 Abduction is triggered by surprise [6.15, CP 5.189].
- P2 Abduction is a form of guessing, underwritten innately by instinct ([6.16, p. 128], [6.15, CP 5.171], [6.17, CP 7.220]).
- P3 A successful abduction provides no grounds for believing the abduced proposition to be true [6.16, p. 178].
- P4 Rather than believing them, the proper thing to do with abduced hypotheses is to send them off

to experimental trial ([6.15, CP 5.599], [6.18, CP 6.469–6.473], [6.17, 7.202–219]).

- P5 The connection between the truth of the abduced hypothesis A and the observed fact C is subjunctive [6.15, CP 5.189].
- P6 The inference that the abduction licenses is not to the proposition A, but rather that A's truth is something that might plausibly be suspected [6.15, CP 5.189].
- P7 The *hence* of the Peircean conclusion is ventured defeasibly [6.15, CP 5.189].

Let us note that P3 conveys something of basic importance. It is that successful abductions are *evidentially inert*. They offer no grounds for believing the hypotheses abduced. What, then, is the good of them?

6.1.2 Ignorance Problems

Seen in Peirce's way, abductions are responses to ignorance problems. An agent has an ignorance problem in relation to an epistemic target when it can't be attained by the cognitive resources presently at his command, or within easy and timely reach of it. If, for some proposition A, you want to know whether A is the case, and you lack the information to answer this question, or to draw it out by implication or projection from what you currently do know, then you have an ignorance problem with respect to A.

Two of the most common responses to ignorance problems are (1) *subduance* and (2) *surrender*. In the first case, one's ignorance is removed by new knowl-

edge, and an altered position is arrived at, which may serve as a positive basis for new action. In the second case, one's ignorance is fully preserved, and is so in a way that cannot serve as a positive basis for new action (new action is action whose decision to perform is lodged in reasons that would have been afforded by that knowledge). For example, suppose that you've forgotten when Barb's birthday is. If her sister Joan is nearby you can ask her, and then you'll have got what you wanted to know. This is subduance. On the other hand if Joan is traveling incognito in Peru and no one else is about, you might find that knowing Barb's birthday no longer interests you. So you might rescind your epistemic target. This would be surrender.

There is a third response that is sometimes available. It is a response that splits the difference between the prior two. It is abduction. Like surrender, abduction is ignorance-preserving, and like subduance, it offers the agent a positive basis for new action. With subduance, the agent overcomes his ignorance. With surrender, his ignorance overcomes him. With abduction, his ignorance remains, but he is not overcome by it. It offers a reasoned basis for new action in the presence of that ignorance. No one should think that the goal of abduction is to *maintain* that ignorance. The goal is to make the best of the ignorance that one chances to be in.

6.1.3 The Gabbay–Woods Schema

The nub of abduction can be described informally. You want to know whether something A is the case. But you don't know and aren't in a position here and now to get to know. However, you observe that if some further proposition H were true, then it together with what you already know would enable you to answer your question with regard to A . Then, on the basis of this subjunctive connection, you infer that H is a conjecturable hypothesis and, on that basis, you release it provisionally for subsequent inferential work in the relevant contexts.

More formally, let T be an agent's epistemic target at a time, and K his knowledge base at that time. Let K^* be an immediate successor of K that lies within the agent's means to produce in a timely way. Let R be an attainment relation for T and let \rightsquigarrow denote the subjunctive conditional relation. $K(H)$ is the revision of K upon the addition of H . $C(H)$ denotes the conjecture of H and H^c its activation. Accordingly, the general structure of abduction can be captured by what has come to be known as the Gabbay–Woods schema [6.6, 19, 20]:

1. $T! E$ [The ! operator sets T as an epistemic target with respect to some state of affairs E]

2. $\neg R(K, T)$ [fact]
3. Subduance is not presently an option [fact]
4. Surrender is not presently an option [fact]
5. $H \notin K$ [fact]
6. $H \notin K^*$ [fact]
7. $\neg R(H, T)$ [fact]
8. $\neg R(K(H), T)$ [fact]
9. $H \rightsquigarrow R(K(H), T)$ [fact]
10. H meets further conditions S_1, \dots, S_n [fact]
11. Therefore, $C(H)$ [sub-conclusion, 1–7]
12. Therefore, H^c [conclusion, 1–8].

It is easy to see that the distinctive epistemic feature of abduction is captured by the schema. It is a given that H is not in the agent's knowledge set K . Nor is it in its immediate successor K^* . Since H is not in K , then the revision of K by H is not a knowledge-successor set to K . Even so, $H \rightsquigarrow R(K(H), T)$. But that subjunctive fact is evidentially inert with respect to H . So the abduction of H leaves the agent no closer than he was before to achieving the knowledge he sought. Though abductively successful, H doesn't enable the abducer to attain his epistemic target. So we have it that successful abduction is ignorance-preserving. Of course, the devil is in the details. Specifying the S_i is perhaps the hardest open problem for abductive logic. In much of the literature it is widely accepted that K -sets must be consistent and that its consistency must be preserved by $K(H)$. This strikes me as unrealistic. Belief sets are often, if not routinely, inconsistent. Also commonly imposed is a minimality condition. There are two inequivalent versions of it. The simplicity version advises that complicated hypotheses should be avoided as much as possible. It is sometimes assumed that truth tends to favor the uncomplicated. I see no reason to accept that. On the other hand, simplicity has a prudential appeal. Simple ideas are more easily understood than complicated ones. But it would be overdoing things to elevate this desideratum to the status of a logically necessary condition. The other version is a form of Quine's maxim of minimum mutilation. It bids the theorist to revise his present theory in the face of new information in ways that leave as much as possible of the now-old theory intact. It advises the revisionist to weigh the benefits of admitting the new information against the costs of undoing the theory's current provisions. This, too, is little more than prudence. No one wants to rule out Planck's introduction of the quantum to physics, never mind the mangling of old physics that ensued. Another of the standard conditions is that $K(H)$ must entail the proposition for which abductive support has been sought. In some variations inductive implication is substituted. Both I think are too strong. Note also that none of the three – consistency, minimality or implica-

tion – could be thought of as *process* protocols. The S_i are conditions on hypothesis selection. I have no very clear idea about how this is done, and I cannot but think that my ignorance is widely shared. Small wonder that logicians have wanted to offload the *logic of discovery* to psychology. I will come back to this briefly in due course. Meanwhile let's agree to regard line (10) as a promissory note [6.21, Chap. 11].

6.1.4 The Yes-But Phenomenon

Perhaps it won't come as much of a surprise to learn of the resistance with which the ignorance-preservation claim has been met when the Gabbay–Woods schema has been presented to (what is by now a sizable number of) philosophical audiences. There are those who think that precisely because it strips good abductions of evidential force, the G–W schema misrepresents Peirce. Others think that precisely because it is faithful to Peirce's conditions the G–W schema discredits the Peircean concept of abduction. Of particular interest is the hesitation shown by philosophers who are actually inclined to *accept* the schema, and *accept* the Peircean notion. It may be true, they seem to think, that abduction is ignorance-preserving, but it is not a truth to which they take kindly. Something about it they find unsatisfying. There is a conventional way of giving voice to this kind of reticence. One does it with the words, *Yes, but . . .* So we may speak of this class of resisters as the ignorance-preservation *yes-but*s.

Some philosophers are of the view that there are at least three grades of evidential strength. There is evidential strength of the truth-preserving sort; evidential strength of the probability-enhancing sort; and evidential strength of a weaker kind. This latter incorporates a notion of evidence that is strong in its way without being either deductively strong or inductively strong. It is, as we might say, induction's poor cousin. Proponents of this approach are faced with an interesting challenge. They must try to tell us what it is for premisses nondeductively to favor a conclusion for which there is no strong inductive support. If the weak cousin thesis is false, lots of philosophers are nevertheless drawn to it. So perhaps the better explanation of the yes-but's resistance to the ignorance-preservation claim is that they think that it *overstates* the poor cousin thesis, that it makes of abduction a poorer thing than it actually is. The poor cousin thesis says that abduction is the weakest evidential relation of the family. But the ignorance-preservation thesis says that it is an evidential relation of no kind, no matter how weak. Accordingly, what the yes-but's are proposing is tantamount to retention of the G–W schema for abduction *minus* Peirce's clause P3. This would allow success-

fully abduced hypotheses the promise of *poor-cousin* evidential backing; but it wouldn't be backing with no evidential force. It is an attractive idea, but it cuts too far.

There are too many cases in which successful reasoning, indeed brilliant reasoning, has the very characteristic the reformers would wish to suppress. A case in point is Planck's quantum hypothesis. In the physics of 1900s, black body radiation lacked unifying laws for high and low frequencies. Planck was disturbed by this. Notwithstanding his lengthy acquaintanceship with it, the disunification of the black body laws was a surprising event. It was, for physics, not a matter of course. Planck wanted to know what it would take to ease his cognitive irritation. Nothing he knew about physics answered this question. Nothing he would come to know about physics would answer it either, as long as physics was done in the standard way. Planck recognized that he would never attain his target until physics were done in a new way, in a way sufficiently at odds with the present paradigm to get some movement on this question; yet not so excessively ajar from it as to make it unrecognizable as physics. That day in 1900 when he announced to his son that he had overturned Newton, Planck was drawn to the conditional that if the quantum hypothesis Q were true then $K(Q)$ – that is, physics as revised by the incorporation of Q – would enable him to reach his target. So he put it to work accordingly. At no stage did Planck think that Q was true. He thought it lacked physical meaning. He thought that his reasoning provided no evidence that Q was true and no grounds for believing it to be true. Peirce wanted a logic that respected this kind of thinking. This is what I want too. The poor cousin thesis doesn't do this, and cannot.

Ignorance removal is prompted by the reasoner's desire to know something he doesn't now know, or to have more knowledge of it than he currently does. What are the conditions under which this happens? It seems right to say that without an appreciation of the general conditions under which a human reasoner is in a state of knowledge, this is a question without a principled answer. If, as I aver, there are abductive modes of reasoning prompted by the desire to improve one's epistemic condition which, even when wholly successful, do not fulfill that objective, there must be two particular considerations thanks to which this is so. One would have to do with abduction. The other has to do with knowledge. A fair part of this first factor is captured by the Gabbay–Woods schema (or so I say). The second is catered for by the right theory of knowledge, if there is one. We asked why, if a philosopher accepted the Gabbay–Woods schema for abduction, would he dislike its commitment to the ignorance-preservation claim? The possibility that we're now positioned to con-

sider is that his yes-but hesitancy flows from how he approaches the general question of knowledge. That is to say, it is his *epistemology* that makes him nervous,

not his *logic*. If so, the *yes* part of *yes, but ...* is directed to the logic, but the *but part* is directed to the epistemology.

6.2 Knowledge

6.2.1 Epistemology

I said in the abstract that epistemological considerations affecting the goodness or badness of premiss-conclusion reasoning are little in evidence in mainstream logic. In so saying, I intend no slight to the now large growing and prospering literature on epistemic logics [6.22–24]. For the most part these logics construct formal representations of the standard repertoire of properties – consequence, validity, derivability, consistency, and so on – defined for sentences to which symbols for *it is known that*, and *it is believed that* function as sentence operators. A central task for these logics is to construct a formal semantics for such sentences, typically on the assumption that these epistemic expressions are modal operators, hence subject to a possible worlds treatment. Notwithstanding their explicitly epistemic orientation, it remains true that there is in this literature virtually no express contact with any of the going epistemologies. So here, too, if they operate at all epistemological considerations operate tacitly as part of an unrevealed epistemological background information. I intend something different here. I want to bring epistemology to the fore, which is precisely where it belongs in logics of premiss-conclusion reasoning of all kinds.

I want also to move on to what I think may be the right explanation of the yes-but's dissatisfactions. Before getting started, a caveat of some importance should be flagged. The explanation I'm about to proffer attributes to the yes-but's an epistemological perspective that hardly anyone shares; I mean by this hardly any epistemologist shares, a notable exception is [6.25]. There is a good chance that whatever its intrinsic plausibility, this new explanation will lack for takers. Even so, for reasons that will appear, I want to persist with it for awhile. Here is what it proposes.

The Right-Wrong Thesis

While the Gabbay–Woods schema gets something right about abduction, it nevertheless gets ignorance-preservation *wrong*. What it gets right is that good abductions are evidentially inert. What it gets wrong is that this lack of evidential heft entails a corresponding failure to lift the abducer in any degree from his present ignorance.

Corollary 6.1

There are abductive contexts in which knowledge can be attained in the absence of evidence.

The idea of knowledge without supporting evidence isn't entirely new or in the least shocking. There is a deeply dug-in inclination to apply this characterization to quite large classes of cases. Roughly, these are the propositions knowledge of which is a priori or independent of experience; or, as with Aristotle's first principles, are known without the necessity or even possibility of demonstration; or, as some insist, are the immediate disclosures of sense and introspection. Disagreements have arisen, and still do, about whether these specifications are accurate or sustainable, but it would be a considerable exaggeration to call this sort of evidential indifference shocking, and wildly inaccurate as a matter of historical fact to think of it as new.

In truth, apriorism is beside the point of the right-wrong thesis and its corollary. The knowledge that falls within their intended ambit is our knowledge of contingent propositions, whether of the empirical sciences or of the common experience of life. The right-wrong claim is that there are contingent propositions about the world which, without being in any way *epistemically privileged*, can be ignorance-reducing by virtue of considerations that lend them no evidential weight. So what is wanted is a theory of knowledge that allows this to happen.

The historically dominant idea in philosophy is that knowledge is true belief plus some other condition, usually identified as justification or evidence. This, the *J*-condition, has been with us at least since Plato's *Theaetetus*, and much scholarly ink has been spilled over how it is best formulated and whether it might require the corrective touch of some further condition. But, as a general idea, the establishment bona fides of the *J*-condition are as rock solid as anything in philosophy.

The account of knowledge I am looking for arises at the juncture of two epistemological developments. One is the trend towards naturalism [6.26] and the other is the arrival of reliabilism [6.27]. It is a theory in which the *J*-condition fails as a general constraint on epistemically unprivileged contingent knowledge. Accordingly, my first task is to try to downgrade the condition, to deny it a defining role. Assuming some success with the

first, my second task will be to find at the intersection of these trends an epistemological orientation – perhaps I would better call it an epistemological *sensibility* – which might without too much strain be reconciled to the loss of the *J*-condition. For ease of reference let me baptize this orientation, this sensibility, the *causal response turn*.

Whereupon task number three, which is to identify those further features of the causal response model that link up the notions of evidence and knowledge in the heterodox ways demanded by the right-wrong thesis.

6.2.2 Losing the *J*-Condition

The *J*-condition has attracted huge literature and underwritten a good deal of strategic equivocation. On *engaged* readings of the condition, a person's belief is justified or evidenced only if he himself has produced his justification then and there, or he has presented the evidence for it on the spot. On *disengaged* readings, a person is justified in believing if a justification exists but hasn't been invoked, or evidence exists but hasn't been adduced or even perhaps found. The engaged and disengaged readings raise an interesting question. How deeply engaged does one have to be to meet the *J*-condition on knowledge? Most epistemologists formulate the engaged-disengaged distinction as one between internalist and externalist justification.

Engagement here is a matter of *case making*. The two readings of *J* define a spectrum, but for present purposes there is little that needs saying of what lies within. It suffices to note that in its most engaged sense a belief is justified or evidenced only if the believer can himself make the case for it here and now. At the other extreme, the belief is justified or evidenced if a case for it is available in principle to someone or other. In the first case, the individual in question has a high degree of case-making engagement. In the other, his engagement is a gestural, anonymous and proxied one: it is engagement in name only.

Suppose the following were true. Suppose that, for every piece of epistemically unprivileged contingent knowledge *p*, there were a structure of facts in virtue of which *p* is the case. Suppose further that for every such *p* a person knows, it would be possible in principle to discern this structure of the facts and the in-virtue-of-relation it bears to *p*'s truth. (I don't think there is any realistic chance of this being so, but let's assume it for the point at hand.) Suppose, finally, that we agreed to say that when in principle knowledge of that structure and that relation exists with respect to a *p* that a subject *S* knows, there exists a justification of *S*'s belief that *p*. For ease of reference, let's call these *factive* justifications. Factive justifications are justifications at their most

disengaged. They stand in radical contrast to highly engaged justifications, which we may call *forensic*.

By construction of the case presently in view, factive justification will be the constant companion of any piece of epistemically unprivileged contingent knowledge that *S* chances to have. But we have in this constancy not conditionhood but concomitance. Factive justification is a faithful accompaniment of such knowledge, but it is not a constituent of it. Forensic justification is another story. We might grant that if, when *S* knows that *p*, he has a forensic justification for his belief, then his justification will have made a contribution to this knowledge. But in relation to all that *S* knows it is comparatively rare that there is a forensic justification. Here is a test case, with a tip of the hat to Peirce: Do you know who your parents are? Of course you do! Very well, then, let's have your forensic justification.

This is troublesome. If we persist in making forensic justification a condition on knowledge, the result is skepticism on an undesirable scale. If, on the other hand, we decide to go with factive justification, then justifications exist whenever knowledge exists, but they aren't conditions on this knowledge. They are not a structural element of it. Whereupon we are met with the *J*-condition dilemma.

J-Condition Dilemma

Depending on how it is read, the *J*-condition is either an irrelevant concomitant of knowledge, or a skepticism-inducing discouragement of it.

The forensic-factive ambiguity runs through all the idioms of *J*-attribution. Concerning his belief that *p* there might be evidence for *p* that *S* adduces or there may be evidence for *p* that exists without attribution. There may be reasons for it that *S* gives, or reasons for it that exist without being given. Like confusions repose in careless uses of *have*. If we allow that *S* has a justification or has evidence or has reasons whenever these things exist factively, we mislicense the inference from the factive to the forensic, allowing, in so doing, *S* to have justifications that he's never heard of.

6.2.3 The Causal Response Model of Knowledge

The causal response (CR) model of knowledge is rightly associated with reliabilism. In all the going forms of it, the *J*-condition is preserved [6.28]. One of the few places in the reliabilist literature where we see stirrings of the pure version of the causal model is Alvin Goldman's first reliabilist paper, which appeared in 1967. It is a rare place in Goldman's foundational corpus where the *J*-condition, if there at all, is given shortest shrift. In some versions, the *J*-condition is

satisfied when one's belief has been reached by reliable *procedures*. In others, the condition is met when the belief was reliably produced, that is, produced by belief-forming *mechanisms* that were working reliably. In contrast to the standard versions, the *pure* version is one in which the *J*-condition is eliminated, rather than reinterpreted along reliabilist lines. As a first approximation, the pure theory characterizes knowledge as follows:

“*S* knows that if and only if *p* is true, *S* believes that, the belief was produced by belief-forming devices, in good working order, operating as they should on good information and in the absence of Gettier nuisances and other hostile externalities.”

Fundamental to what I've been calling the pure theory is the conviction that knowledge is not in any essential or general way tied to case making, that *knowing* is one thing and *showing* another. This is not to say that case making is never implicated in knowledge. There are lots of beliefs that would not have been had in the absence of the case makings that triggered their formation. Think here of a mother's sad realization that her son is guilty of the crime after all, or a nineteenth century mathematician's grudging acknowledgment of the transfinite. But as a general constraint, case making is rejected by pure causalists; by causalists of the sort that Goldman was trying to be in 1967.

6.2.4 Naturalism

Epistemology's naturalized turn supplies a welcoming habitat for the CR model. Naturalism comes in various and competing versions, but at the core of them all is the insistence that human knowledge is a natural phenomenon, achieved by natural beings in accordance with their design and wherewithal, interacting in the causal nexi in which the human organism lives out his life. Unlike the *J* theorist, the CR theorist is a respecter of the passive side of knowledge. He knows that there are large classes of cases in which achieving a knowledge of something is a little more than just being awake and on the scene. Even where some initiative is required by the knower, the resultant knowledge is always a partnership between doing and being done to. So even worked-for knowledge is partly down to *him* and partly down to his *devices*.

It would be wrong to leave the impression that, on the CR model, knowing things is just a matter of doing what comes naturally. There are ranges of cases in which knowledge is extremely difficult to get, if gettable at all. There are cases in which knowledge is unattainable except for the intelligence, skill, training and expertise of those who seek it. Everyone has an

aptitude for knowledge. But there are cases galore in which aptitude requires the supplementation of vocation and talent – and training. CR theorists are no less aware of this than their *J* rivals. The difference between them falls in where the emphasis falls. Among *J* theorists there is a tendency to generalize the hard cases. Among CR theorists there is a contrary tendency to keep the hard cases in their place.

Let me say again that *J*-theories give an exaggerated, if equivocal, place to the role of showing in knowing. Contrary to what might be supposed, the CR model is no disrespecter of the showing-knowing distinction, albeit with a more circumscribed appreciation of showing. I want to turn to this now.

6.2.5 Showing and Knowing

Consider the case of Fermat's Last Theorem. The theorem asserts that for integers x , y , and z , the equation $x^n + y^n = z^n$ lacks a solution when $n > 2$. Fermat famously left a marginal note claiming to have found a proof of his theorem. I want to simplify the example by stipulating that he did not have a proof and did not think or say that he did. The received wisdom is that Fermat went to his grave not knowing that his theorem is true. The received wisdom is that no one knew whether the theorem is true until Andrew Wiles' proof of it in 1995. If the forensically conceived *J* model were true, this would be pretty much the way we would expect the received wisdom to go.

If the *J* model is hard on knowledge, the CR model is a good deal more accommodating. It gives to knowledge a generous provenance. But I daresay that it will come as a surprise that, on some perfectly plausible assumptions, Fermat did indeed know the truth of his theorem, never mind (as we have stipulated) that he was all at sea about its proof. Fermat was no rookie. He was a gifted and experienced mathematician. He was immersed in a sea of mathematical sophistication. He was a mathematical virtuoso. Fermat knew his theorem if the following conditions were met: It is true (as indeed it is), he believed it (as indeed he did), his highly trained belief-forming devices were in good order (as indeed they were) and not in this instance misperforming (as indeed they were not), and their operations were not compromised by bad information or Gettier nuisances (as indeed was the case). So Fermat and generations of others like-placed knew the theorem well before its proof could be contrived.

We come now to a related point about showing and knowing. Showing and knowing mark two distinct goals for science, and a corresponding difference in their satisfaction conditions. Not unlike the law, science is in significant measure a case-making profession –

a forensic profession – made so by the premium it places on *knowing* when knowledge has been achieved, rather than just achieving it. This has something to do with its status as a profession, subject to its own exacting requirements for apprenticeship, advancement and successful practice. These are factors that impose on people in the showing professions expectations that regulate *public announcement*. Fermat may well have known the truth of his theorem and may have had occasion to say so to a trusted friend or his mother. But he was not to say it for publication. Publication is a vehicle for case making, and case making is harder than knowing. Journal editors don't give a toss for what you know. But they might sit up and notice if you can show what you know.

6.2.6 Explaining the Yes-Buts

The ignorance-preservation claim is rooted in the idea of the no evidence-no knowledge thesis.

The No Evidence–No Knowledge Thesis

Since successful abduction is evidentially inert, it is also epistemically inert. But this is justificationism: No advance in knowledge without some corresponding advance in evidence.

The CR model jettisons justificationism. It denies the very implication in which the ignorance-preservation thesis is grounded. It is not hard to see that the evidence, whose abductive absence Peirce seizes upon, is not evidence in the factive sense. Peirce insists that we have no business believing a successfully abducted hypothesis. Peirce certainly doesn't deny that behind any plausibly conjectured hypothesis there is a structure of facts in virtue to which it owes its truth value. Peirce thinks that our track record as abductive guessers is remarkably good. He is struck by the ratio of right guesses to guesses. He is struck by our aptitude for correcting wrong guesses. The evidence whose absence matters here is forensic, it is evidence by which an abducer could vindicate his belief in the hypothesis at hand. But Peirce thinks that in the abductive context nothing vindicates that belief.

We come now to a critical observation. There is nothing in Peirce's account that tells us that abducted hypotheses aren't believed as a *matter of fact*. Some clearly are not. At the time of their respective advancements, Planck didn't believe the quantum hypothesis and Gell-Mann didn't believe the quark hypothesis. But it takes no more than simple inspection to see that there are masses of cases to the contrary, that abductive success is belief-*inducing* on a large scale.

There is in this commonplace fact something for the CR theories to make something of. Let H be one

of those successfully abducted hypotheses that happen to be true and, contrary to Peirce's advice, believed by its abducer S . What would it take to get us seriously to propose that, when these conditions are met, S 's belief-forming device's are malfunctioning or are in poor operating order. Notice that a commonly held answer is not available here, on pain of question begging. It cannot be said that unevidenced belief is itself evidence of malfunction and disorder. That is, it cannot be said to the CR-theorist, since implicit in his rejection of justificationism is his rejection of this answer.

Is there, then, any reason to suppose that the arousal of unevidenced belief might be some indication of *properly* functioning belief formation? Ironically enough, there is an affirmative answer in Peirce himself. Peirce is much taken with our capacity for right guessing. Our facility with guessing is so impressive that Peirce is driven to the idea that good guessing is something the human animal is built for. But if we are built for good guessing, and good abduction is a form of guessing, how can the abduction of true hypotheses not be likewise something we're built for? Accordingly, there is a case for saying that.

Knowledge Enhancement

In the CR model of knowledge, there are numbers of instances in which successful abduction is not only not ignorance preserving, but actually *knowledge enhancing*.

Part of what makes for the irony of Peirce's enthusiasm for right guessing is his insistence that guesses not be indulged by belief. In this he is a justificationist. Abducers have no business in believing unevidenced propositions, never mind their abductive allure. This is enough of a basis to pin the ignorance-preservation thesis on Peirce, but *not* on a CR theorist who accepts the Gabbay–Woods schema. What this shows is that theirs is not a disagreement about abduction. It is a disagreement about knowledge.

There isn't much likelihood that yes-buts will flock to this accommodation. The reason is that hardly anyone (any philosopher anyway) thinks the CR model is true in its pure form. There is no space left to me to debate the ins and outs of this. Suffice it to say that it offers the abductive logician the very relief that the yes-buts pine for. Besides, the CR theory just might be true [6.21].

6.2.7 Guessing

In line (10) of the G–W schema the S_i occur as placeholders for conditions on hypothesis selection. Previously, I said that I didn't know what these conditions

are [6.7]. In point of fact there are two things that I don't know. One is the normative conditions in virtue of which the selection made is a worthy choice. The other is the causal conditions that enable the choice to be made. It is easy to see that there are a good many H s that could serve as antecedents in line (9)'s $H \rightsquigarrow R(K(H), T)$ without disturbing its truth value. It is also easy to see that a good many of *those* H s would never be abductively concluded, never mind their occurrence there. It is clear that a reasonable choice of H must preserve the truth of (9). It is also clear that this is not enough for abductive significance. A reasonable choice must have some further features. I am especially at a loss to describe how beings like us actually go about finding things like that. Perhaps it will be said that my difficulty is a reflection on me, not on the criteria for hypothesis selection. It is true that the number of propositions that could be entertained is at least as large as the number of H s that slot into the antecedent of (9) in a truth-preserving way. Let's think of these as constituting the hypothesis-selection space. Selection, in turn, is a matter of cutting down this large space to a much smaller proper subset, ideally a unit set. Selection, to this same effect, would be achieved by a search engine operating on the hypothesis-selection space. Its purpose would be to pluck from that multiplicity the one, or few ones, that would serve our purposes.

There is nothing remotely mystifying or opaque about search engines (why else would we bother with Google?). So isn't the problem I'm having with the S_i that I'm not a software engineer? Wouldn't it be prudent to outsource the hypothesis-selection task to someone equipped to perform it? To which I say: If that is a doable thing we should do it. There is no doubt that algorithms exist in exuberant abundance for search tasks of considerable variety and complexity. There are algorithms that cut down a computer system's search space to one answering to the algorithm's flags. Perhaps such an arrangement could be said to model hypothesis selection. But it is another thing entirely as to whether, when we ourselves are performing them, our hypothesis selections implement the system's algorithms. So I am minded to say that *my* questions about the S_i are not comprehensively answerable by a software engineer.

Here is where guessing re-enters the picture, which is what Peirce thinks that hypothesis selection is. Peirce is struck by how good we are at it. By this he needn't have meant that we have more correct guesses than incorrect. It is enough that, even if we make fewer correct guesses than incorrect, the ratio of correct to incorrect is still impressively high. We get it right, rather than wrong, with a notable frequency. Our opportunities for getting it wrong are enormous. Relative to the propositions that could have been guessed at, the number of

times that they are rightly guessed is amazing; so much so that Peirce is led to surmise that our proclivity for right guesses is innate. Of course, not all good guessing is accurate. A good guess can be one that puts the guessed-at proposition in *the ball park*, notwithstanding that it might actually not be true. Here, too, good guesses might include more incorrect ones than correct. But as before, the ratio of correct to merely good could be notably high. So the safer claim on Peirce's behalf is that beings like us are hardwired to make for good, although not necessarily correct, guesses with a very high frequency. It is lots easier to make a ball-park guess than a true one; so much so that the hesitant nativist might claim a hardwired proclivity for ball-park, yet not for truth, save as a welcome contingency, which in its own turn presents itself with an agreeable frequency. Thus the safe inference to draw from the fact that H was selected is that H is in the ball park. The inference to H 's truth is not dismissable, but it is weaker.

Needless to say, nativism has problems all its own. But what I want to concentrate on is a problem it poses for Peircian abduction. At the heart of all is what to make of ball-park guesses. The safest thing is to propose is that, even when false, a ball-park hypothesis in a given context is one that bears serious operational consideration there. There might be two overarching reasons for this. *One* is that ball-park hypotheses show promise of having a coherently manageable role in the conceptual spaces of the contexts of their engagement. Take again the example of Planck. The quantum hypothesis was a big wrench to classical physics. It didn't then have an established scientific meaning. It entered the fray without any trace of a track record. Even so, for all its foreignness, it was a ball-park hypothesis. What made it so was that $P(Q)$ was a theory revision recognizable as *physics*. Contrast Q with *the gold fairy will achieve the sought-for unification*. Of course, all of this turns on the assumption that Peirce got it right in thinking that hypothesis selection is guessing, and to note that good guessing is innate. Call this the *innateness hypothesis*. The *second* consideration is that the frequency of true hypotheses to ball-park hypotheses is notably high.

Whether he (expressly) knows how it's done, when an abductive agent is going through his paces, there is a point at which he selects a hypothesis H . If the innateness thesis holds, then the agent has introduced a proposition that has an excellent shot at being ball-park, and a decent shot of being true. On all approaches to the matter, an abduction won't have been performed in the absence of H ; and on the G-W approach, it won't have been performed correctly unless H is neither believed nor (however weakly) evidenced by its own abductive success. On the other, our present reflections suggest

that the very fact that H was *selected* is evidence that it is ball-park, and less good but not nonexistent evidence that it is true. Moreover, H is the antecedent of our subjunctive conditional (9) $H \rightsquigarrow R(K(H), T)$. If H is true so is $R(K(H), T)$ by modus ponens; and if $R(K(H), T)$ holds the original ignorance problem is solved by a form of subduance. In which case, the abduction simply *lapses*. It lapses because the nonevidential weight lent to a successfully *abducted* hypothesis is, on the G–W model, weaker than the evidential support given it by way of the innateness hypothesis as regards its very *selection*.

If, on the other hand, H is not true, but ball-park – hence favorably evidenced – and being evidenced is closed under consequence, then the reasoning at hand also goes through under the obvious adjustments.

The problem is that there are two matters on which Peirce can't have his cake and eat it too. If he retains the innateness thesis he can't have the ignorance-preservation thesis. Equally, if he keeps ignorance preservation he must give up innateness, which *nota bene* is not the thesis that guessing is innate but that *good* guessing is innate. Yet if we give up innateness we're back to where we started, with less than we would like to say about the actual conditions for which the G–W S_i are mere placeholders. I leave the innateness-ignorance preservation clash as an open problem in the abduction research program. Since, by our earlier reasoning, there is an epistemology (CR) that retains ignorance preservation only as a contingent property of some abductions, my present uncertain inclination is to retain G–W as modified by CR and to rethink innateness. But I'm open to offers. I'll get back to this briefly in the section to follow.

Having had my say about the epistemological considerations that play out in the logic of abduction, I want to turn to the question of how, or to what extent, a logic of abduction will meet universalist conditions on logic. I want to determine whether or to what extent abductive theories embody the structural core assumed by universalists to be common to any theory that qualifies for admittance to the province of logic.

Whatever the details, abduction is a form of premiss-conclusion reasoning. By the conclusions-consequence thesis, whenever the reasoning is good the conclusion that's drawn is a consequence of those premisses. As logics have proliferated, so too the consequences, albeit not exactly in a strict one-to-one correspondence. If today there are more logics than one can shake a stick at, there is a concomitant plenitude of consequences relations. Much of what preoccupies logicians presently is the classification, individuation, and interrelatedness of this multiplicity. Whatever their variations, there is one distinction to which they all an-

swer. Some consequence relations are truth-preserving; all the others aren't. Truth-preserving consequence is (said to be) monotonic. (It isn't. To take an ancient example, Aristotle's *sylogistic* consequence is truth-preserving but nonmonotonic.) Premises from which a conclusion follows can be supplemented at will and the conclusion will still follow. One way of capturing this point is that truth-preserving consequence is impervious to the openness of the world. As far as consequencehood is concerned, the world might as well be closed. Once a consequence of something, always a consequence of it. It is strikingly otherwise with non-truth-preserving consequence. It is precisely this indifference to the openness of the world that is lost.

6.2.8 Closed Worlds

When we were discussing the J condition on knowledge, we called upon a distinction between the factive justification of a belief and its forensic justification. In a rough and ready way, a factive justification is down to the *world*, whereas a forensic justification is down to *us*. We find ourselves at a point at which the idea of factivity might be put to further good use. To see how, it is necessary to acknowledge that the distinction between open and closed worlds is systematically ambiguous. In one sense it marks a contrast between *information states* at a time, with the closed world being the state of total information, and open ones states of incomplete information. In the other sense, a closed world can be called factive. A closed world at t is everything that is the case at t . It is the totality of facts at t . A closed world is also open at t , not with regard to the facts that close it at t , but in respect of the facts thence to come. We may suppose that the world will cease to open at the crack of doom, and that the complete inventory of all the facts that ever were would be logged in the right sort of Domsday Book. It is not, of course, a book that any of us will get to read. Like it or not, we must make do with openness. Both our information states and the world are open at any t before the crack. But the diachronics of facticity outpace the accuracy of information states. When there is a clash, the world at t always trumps our information about it at $t-n$.

At any given time the world will be more closed than its concurrent information states. At any given time the state of the world outreaches the state of our knowledge of it. When we reason from premisses to conclusions we are not negotiating with the world. We are negotiating with informational reflections of the world. We are negotiating with information states. Given the limitations on human information states, our representations of the world are in virtually all respects

open, and most premises-conclusion relations are susceptible to rupture. Truth-preserving consequences are an interesting exception. The world can be as open as openness gets, but a truth-preserving consequence of something is always a consequence of it, never mind the provisions at any t of our information states. Nonmonotonic consequence is different: Today a consequence tomorrow a nonconsequence.

We might think that the more prudent course is to cease drawing conclusions and postpone the decisions they induce us to make until our information state closes, until our information is permanently total. The ludicrousness of the assumption speaks for itself. Cognitive and behavioral paralysis is not an evolutionary option. Thus arises the closed world assumption. Given that belief and action cannot await the arrival of total information, it behooves us to draw our conclusions and take our decisions when the likelihood of informational defeat is least high, at which point we would invoke the assumption that for the matter at hand the world might just as well be closed.

The key question about the closed world assumption is the set of conditions under which it is reasonable to invoke it. The follow-up question is whether we're much good at it. I am not much inclined to think that we have done all that well in answering the first question. But my answer to the second is that, given the plenitude of times and circumstances at which to invoke it, our track record is really quite good; certainly good enough to keep humanity's knowledge-seeking project briskly up and running. Even so, the closed world assumption is vulnerable to two occasions of defeat. One is by way of later information about later facts. Another is by way of later information about the facts now in play. It is easy to see, and no surprise at all, that new facts will overturn present information about present facts with a frequency that matches the frequency of the world's own displacement of old facts by new. Less easy to see is how we manage as well as we do in invoking closure in the absence of information about the present destructive facts currently beyond our ken. Here, too, we have a cut-down problem. We call upon closure in the hopeful expectation that no present unannounced fact will undo the conclusions we now draw and the decisions they induce us to make. Comparatively speaking, virtually all the facts there are now are facts that no one will ever know. That's quite a lot of facts, indeed it is nondenumerably many (for isn't it a fact that, for any real number, it is a number, and is self-identical, and so on?).

There is a point of similarity between hypothesis selection and the imposition of world closure. Our good

track record with both invites a nativist account each time. Oversimplified, we are as good as we are at selecting hypotheses because that's the way we were built. We are as good as we are at closing the world because that too is the way we were built. I suggested earlier that in abductive contexts the very fact that H has been selected is some evidence that it is true (and even better evidence that it is ball-park). But this seems to contradict the Peircian thesis that abductive success confers on H nothing stronger than the suspicion that it might be true. Since Peirce's account of abduction incorporates both the innateness thesis and the no-evidential-support thesis, it would appear that Peirce's account is internally inconsistent. I said a section ago that I had a slight leaning for retaining the no-evidence thesis and lightening up on the innateness thesis. Either way is Hobson's choice. That, anyhow, is how it appears.

In fact, however, the appearance is deceptive. There is no contradiction. Peirce does not make it a condition on abductive hypothesis-selection that H enter the fray entirely untouched by reasons to believe it or evidence that supports it. He requires that the present support-status of H has no role to play in the abductive process. That H is somewhat well supported doesn't, if true, have any premissory role here. Moreover, it is not the goal of abduction to make any kind of case for H 's truth. The goal is to find an H which, independently of its own epistemic status, would if true enable a reasoner to hit his target T . But whatever the target is, it's not the target of wanting to know whether H is true. It is true that, if all goes well, Peirce asserts that it may be defeasibly concluded that there is reason to suspect that H might be true. But, again, abduction's purpose is not to make a case for H , no matter how weakly. The function of the suspectability observation is wholly retrospective. It serves as a hypothesis-selection vindicator. You've picked the (or a) right hypothesis only if the true subjunctive conditional in which it appears as antecedent occasions the abducer's satisfaction that that, in and of itself, would make it reasonable to suspect that H might be so. In a way, then, the G-W schema misrepresents this connection. It is not that the abduction implies H 's suspectability, but rather that the abduction won't succeed unless the truth of line (9) induces the suspectability belief [6.21] (for more on the causal role in inference, readers could again consult [6.21]). And that won't happen if the wrong H has been selected, never mind that it preserves (9)'s truth. For the point at hand, however, we've arrived at a good result. The innateness thesis and the no-support thesis are both implicated in the Peircean construal of abduction, but are in perfect consistency.

6.3 Logic

6.3.1 Consequences and Conclusions

I said at the beginning that for nearly two and a half millennia the central focus of logic has been the consequence relation. More basic still was a concomitant preoccupation with premiss-conclusion reasoning. For a very long time logicians took it as given that these two matters are joined at the hip.

Conclusions and Consequences

When someone correctly draws a conclusion from some premisses, his conclusion is a consequence of them.

Corollary 6.2

If a conclusion drawn from some premisses is not a consequence of them, then the conclusion is incorrectly drawn.

If this were so, it could be seen at once that there is a quite intuitive distinction between the consequences that a premiss set has and the consequences that a reasonable reasoner would *conclude* from it. Any treatment of logic in which this distinction is at least implicitly present, there is a principled role for agents, for the very beings who draw what conclusions they will from the consequences that flow from the premisses at hand. In any such logic there will be at least implicit provision for the nature of the agent's involvement. In every case the involvement is epistemically oriented. People want to know what follows from what. They want to know how to rebut an opponent. They want to know whether, when this follows from that that, they can now be said to know that. In a helpful simplification, it could be said that logic got out of the agency business in 1879. It is not that agency was overlooked entirely, but rather that it was scandalously short-sheeted. For consequence, the having-drawing distinction would fold into having; and having, it would be said, would be the very things drawn by an ideally rational reasoner. Of course, this downplaying of cognitive agency was never without its dissenters. Indeed today we are awash in game theoretic exuberance, to name just one development of note.

6.3.2 Semantics

Consequence derives its semantic character from its attachment to truth, itself a semantic property in an odd baptismal bestowal by Tarski. In the deductive case, it is easy to see how truth is implicated in consequence and how, in turn, consequence assumes its status as a semantic relation. Not only does truth ground the very definition of consequence, but it makes for a re-

lation that is also truth-preserving. The monotonicity of consequence provides the sole instance in which a consequence is impervious to the informational openness of the world. It is the one case in informational openness at t that is indifferent to the world's factive closure at t , to say nothing of its final closure at the crack of doom. It has long been known that logicians, then and now, harbor an inordinate affection for deductive consequence. It's not hard to see why. Deductive consequence has proved more responsive to theoretical treatment than any of the nondeductive variety. But more centrally, it is the only consequence relation that captures permanent chunks of facticity.

Whatever else we might say, we can't say that nonmonotonic relations are relations of *semantic* consequence. If B is a nonmonotonic consequence of A it holds independently of whatever makes for the *truth* of A and B . Sometimes perhaps it holds on account of probability conditions on A and B , but probability has nothing to do with truth. If there is such a thing as probabilistic consequence – think here of Carnap's partial entailment – it is not a semantic relation. We may have it now that the evidence strongly supports the charge against Spike in last night's burglary. We might come to know better tomorrow. We might learn that at the time of the offense Spike was spotted on the other side of town. So the world at t didn't support then the proposition that Spike did do it, never mind the state of information the day after t .

No one doubts that yesterday there existed between the evidence on hand and the charge against Spike a relation of epistemic and decisional importance, a kind of relation in whose absence a survivable human life would be impossible. But a fair question nevertheless presses for attention: Where is the gain in conceptualizing these vital premiss-conclusion relations as relations of logical *consequence*? Where is the good of trying to construe nonmonotonic relations on the model of attenuated and retrofitted monotonic consequences? My own inclination is to say that talk of nonmonotonic consequence misconceives the *import* of nonmonotonicity. We tend to think of it as a distinguishing feature of consequence relations, when what it really is is the defining feature of nontruth preservation.

When premiss-conclusion reasoning is good but not truth-preserving, it is made so by an underlying relation. Any theory of premiss-conclusion reasoning had better have something to say about this, about its nature and how it operates. We should give it the name it both deserves and better reflects how it actually does function. Let's call it *conclusionality*. Conclusionality is an epistemic or epistemic/prudential relation. It is a relation

that helps rearrange our belief states, hence possessing decisional significance. Any struggle to discern whether it is also a consequence relation seems to me to be sailing into the wind.

Abductive conclusions are on the receiving end of this relation; they are occupants of its converse domain. If our present reflections can be made to stand, there is no relation of abductive consequence; and it will cause us no end of distraction trying to figure out how to make it one. It hardly needs saying that depriving a logic of abduction of its own relation of abductive consequence must of necessity rearrange how abductive logic is conceptualized. There are plenty of logicians more than ready to say that a logic without consequence relations is a logic in name only – a logic *façon de parler*, hence a logic that fails universalistic prescriptions. I am otherwise minded. Logic started with *conclusionality* relations. It was adventitiousness, not essence, that brought it about that the ones first considered were also consequence relations. Logic has had a good innings right from the beginning. In a way, this has been unfortunate. The success we've had with consequence has obscured our view of conclusionality. It has led us to think that the more we can get conclusionality *in gen-*

eral to be a species of consequence, the faster we'll achieve some theoretical respectability. We would be better served to place conclusionality at the core of logic and to place consequence in an annex of less central importance. If we did that, we could reinstate the logic of abduction and equip it for admittance into universalist respectability. But we could also reinvest to good effect all that energy we've devoted to consequentializing the conclusionality relation, in a refreshed effort to figure how conclusionality actually works in the epistemically sensitive environments in which, perforce, the human organism must operate.

Acknowledgments. I would know a good deal less than I presently do about abduction without stimulating instruction from Dov Gabbay, Lorenzo Magnani, Atocha Aliseda, Ahti-Veikko Pietarinen, Peter Bruza, Woosuk Park, Douglas Niño and more recently – especially in relation to sections 11 and 12 – Madeleine Ransom. To all my warmest thanks. My student Frank Hong has also pitched in with astute suggestions; equal gratitude to him. For technical support and everything else that matters, Carol Woods is my go-to gal. Without whom not.

References

- 6.1 J.R. Josephson, S.G. Josephson (Eds.): *Abductive Inference: Computation, Philosophy, Technology* (Cambridge University Press, Cambridge 1994)
- 6.2 T. Kapitan: Peirce and the structure of abductive inference. In: *Studies in the Logic of Charles Sanders Peirce*, ed. by N. Houser, D.D. Roberts, J. Van Evra (Indiana University Press, Bloomington 1997) pp. 477–496
- 6.3 J. Hintikka: What is abduction? The fundamental problem of contemporary epistemology, *Trans. Charles S. Peirce Soc.* **34**, 503–533 (1998)
- 6.4 P.A. Flach, C.K. Antonis: *Abduction and Induction: Essays on Their Relation and Interpretation* (Kluwer, Dordrecht 2000)
- 6.5 L. Magnani: *Abduction, Reason and Science: Processes of Discovery and Explanation* (Kluwer, Dordrecht 2001)
- 6.6 D.M. Gabbay, J. Woods: *The Reach of Abduction: Insight and Trial, A Practical Logic of Cognitive Systems*, Vol. 2 (North-Holland, Amsterdam 2005)
- 6.7 S. Paavola: Peircean abduction: Instinct or inference?, *Semiotica* **153**, 131–154 (2005)
- 6.8 A.-V. Pietarinen: *Signs of Logic: Peircean Themes on the Philosophy of Language, Games and Communication* (Springer, Dordrecht 2006)
- 6.9 A. Aliseda: Abductive reasoning. In: *Logical Investigation into the Processes of Discovery and Evaluation*, (Springer, Dordrecht 2006)
- 6.10 P.D. Bruza, D.W. Song, R.M. McArthur: Abduction in semantic space: Towards a logic of discovery, *Logic J. IGPL* **12**, 97–110 (2004)
- 6.11 G. Schurz: Patterns of abduction, *Synthese* **164**, 201–234 (2008)
- 6.12 P. Bruza, A. Barros, M. Kaiser: Augmenting web service discovery by cognitive semantics and abduction, *Proc. IEEE/WIC/ACM Int. Joint Conf. on Web Intelligence and Intell. Agent Technol.* (IET, London 2009) pp. 403–410
- 6.13 L. Magnani: *Abductive Cognition. The Epistemological and Eco-Cognitive Dimensions of Hypothetical Reasoning* (Springer, Heidelberg 2009)
- 6.14 Aristotle: Categories. In: *The Complete Works of Aristotle*, ed. by J. Barnes (Princeton Univ. Press, Princeton 1985)
- 6.15 C.S. Peirce: *Pragmatism and Pragmaticism*, *Collected Papers of Charles Sanders Peirce*, Vol. 5 (Harvard Univ. Press, Cambridge 1934), ed. by C. Hartshorne, P. Weiss
- 6.16 C.S. Peirce: *Reasoning and the Logic of Things: The Cambridge Conference Lectures of 1898* (Harvard University Press, Cambridge 1992), ed. by K.L. Kettner
- 6.17 C.S. Peirce: *Science and Philosophy*, *Collected Papers of Charles Sanders Peirce*, Vol. 7 (Harvard Univ. Press, Cambridge 1958), ed. by A.W. Burks
- 6.18 C.S. Peirce: *Scientific Metaphysics*, *Collected Papers of Charles Sanders Peirce*, Vol. 6 (Harvard Univ. Press, Cambridge 1935), ed. by C. Hartshorne, P. Weiss

- 6.19 J. Woods: Peirce's abductive enthusiasms, *Protosociol.* **13**, 117–125 (1999)
- 6.20 J. Woods: Cognitive economics and the logic of abduction, *Review of Symbolic Logic* **5**, 148–161 (2012)
- 6.21 J. Woods: *Errors of Reasoning: Naturalizing the Logic of Inference*, Vol. 45 (College Publications, London 2013), Studies in Logic Ser.
- 6.22 P. Gochet, G. Gribomont: Epistemic logic. In: *Logic and the Modalities in the Twentieth Century*, Handbook of the History of Logic, Vol. 7, ed. by M. Dov Gabbay, J. Woods (North-Holland, Amsterdam 2006) pp. 99–195
- 6.23 R. Fagin, J.Y. Halpern, Y. Moses, Y.M. Vardi: *Reasoning About Knowledge* (MIT, Cambridge 1995)
- 6.24 J. Van Benthem: *Logical Dynamics of Information and Interaction* (Cambridge University Press, New York 2011)
- 6.25 N.E. Sahlin, W. Rabinowitz: The evidentiary value model. In: *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, Vol. 1, ed. by D.M. Gabbay, P. Smets (Kluwer, Dordrecht 1998) pp. 247–265
- 6.26 W.V. Quine: Epistemology naturalized. In: *Ontological Relativity and Other Essays*, ed. by W.V. Quine (Columbia University Press, New York 1969)
- 6.27 A. Goldman: What is justified belief? In: *Justification and Knowledge Philosophical Studies Series in Philosophy* **17**, ed. by G. Pappas (Reidel, Dordrecht 1979) pp. 1–23
- 6.28 A. Goldman: A causal theory of knowing, *J. Philosophy* **64**, 357–372 (1967)