

THE AVAILABILITY HEURISTIC AND INFERENCE TO THE BEST EXPLANATION

Michael J. SHAFFER

ABSTRACT: This paper shows how the availability heuristic can be used to justify inference to the best explanation in such a way that van Fraassen's infamous "best of a bad lot" objection can be adroitly avoided. With this end in mind, a dynamic and contextual version of the erotetic model of explanation sufficient to ground this response is presented and defended.

KEYWORDS: inference to the best explanation, explanation, scientific progress, heuristics, erotetic logic, contextualism

1. Introduction

The programs respectively associated with bounded and ecological rationality (BER) and the heuristics and biases program (HBP) have been regarded as having significant implications for many areas of philosophy and psychology. The HBP is an empirically motivated project that focuses on demonstrating why human cognitive performance with respect to tasks like probabilistic reasoning and decision-making often violates (or appears to violate) classical norms of rationality.¹ On a more positive note, those working in the context of this program have argued that human cognitive performance involves using variety of simple heuristics rather than conformity to the classical norms of rationality (i.e. the probability calculus, classical first-order logic, orthodox decision theory, etc.). The BER project is also an empirically minded project aimed at showing that human cognitive performance is actually rational despite the fact that such behavior often does not satisfy classical standards of rationality. BER specifically focuses on the

¹ See Daniel Kahneman, Paul Slovic and Amos Tversky, *Judgment under Uncertainty* (Cambridge: Cambridge University Press, 1982) and Gerd Gigerenzer, *The Adaptive Tool Box* (Oxford: Oxford University Press, 2000). Also, see Ken Manktelow, *Thinking and Reasoning* (New York: Psychology Press, 2012) for an excellent overview and Johnathan Howard, *Cognitive Errors and Diagnostic Mistakes* (New York: Springer, 2019) for discussion of heuristics and cognitive biases in medicine.

computational and environmental features of real cognitive performance as the key to understanding how humans are rational in terms of this alternative, heuristic-based and environmentally sensitive, account of rationality.

BER is a reaction to the pessimistic interpretation of the results of the HBP which were sometimes alleged to show that humans are badly irrational when judged against classical norms of rationality.² The defenders of the BER project effectively disputed this more pessimistic conclusion and argued that facts about human cognitive performance are better understood as evidence that the traditional norms of rationality are not the correct norms by which human cognitive performance should be judged. The opposition between these two camps is ongoing and it has led to some heated exchanges.³ But, these ideas can be usefully combined to support an alternative and empirically grounded conception of rationality as adherence to heuristic rules that are normatively appropriate in certain ecological contexts and given human cognitive limitations.⁴

In this paper this sort of empirically based and fallibilistic approach to rationality is used to justify inference to the best explanation (IBE) and this justification is specifically based on the availability heuristic. This strategy also involves the central contention that IBE involves the more general notion of problem or question substitution.⁵ In its relevant form, the availability heuristic is the claim that certain inferences and decisions are made on the basis of psychologically familiar factors, as opposed to all relevant factors.⁶ Problem or question substitution is just the tactic of substituting and solving an easier version of a problem when a given problem is itself too difficult to solve. So, the availability heuristic is just a special case of problem substitution.⁷ The contention here then is that it is rational to accept the best psychologically available

² See Richard Nisbett and Eugene Borgida, "Attribution and the Psychology of Prediction," *Journal of Personality and Social Psychology* 32 (1975): 932-43 and Massimo Piatelli-Palmarini, *Inevitable Illusions* (New York: John Wiley, 1994).

³ See Richard Samuels, Stephen Stich, and Michael Bishop, "Ending the Rationality Wars: How to Make Disputes about Human Rationality Disappear," in *Common Sense, Reasoning and Rationality*, ed. Renee Elio (Oxford: Oxford University Press, 2002), 236-268, Daniel Kahneman and Amos Tversky, "On the Reality of Cognitive Illusions: A Reply to Gigerenzer's Critique," *Psychological Review* 103 (1996): 582-591 and Gerd Gigerenzer, "On Narrow Norms and Vague Heuristics," *Psychological Review* 103 (1996): 592-596..

⁴ A version of this hybrid view antedates both HBP and BER and was defended in Herbert Simon, *Models of Man* (New York: Wiley, 1957).

⁵ See Daniel Kahneman, *Thinking Fast and Slow* (New York: Farrar, Straus and Giroux, 2011).

⁶ See Kahneman, *Thinking Fast and Slow*.

⁷ See Kahneman, *Thinking Fast and Slow*, ch. 9 for discussion of this connection.

explanation of psychologically available data when we frame this sort of inferential practice in terms of a more naturalistic and realistic conception of rationality. In other words, it is often perfectly rational to substitute simpler explanatory problems for more complex ones. This is due to our cognitive limitations and environmental constraints. Such substitution does carry with it the possibility of cognitive bias and error, but this is no surprise when we recognize that explanatory reasoning involves uncertainty and limited cognitive resources. However, as we shall see, such reasoning also involves the possibility for the correction of such errors and the refinement of our explanatory understanding.

The model proposed here for IBE is founded on a theory that combines insights from epistemic contextualism and the erotetic theory of explanation. One important implication of this work is that it provides an answer to van Fraassen's infamous criticism of IBE.⁸ This critical attack on IBE is based on the contention that the conclusions of such inferences should not be taken to be likely (and hence should not be accepted). This is supposed to be because such inferences are always based on a set of available hypotheses that constitutes only a small sub-set of all of the possible hypotheses that are potential explanations of a given phenomenon. So, as van Fraassen has argued, it appears to be the case that it will always be much more likely that the true explanation is among the set of unconsidered (and mostly unformulated) hypotheses. The alternative model of IBE presented in this paper neatly avoids this criticism and renders rational the acceptance of the conclusions of such inferences. In part this is because the model of IBE introduced here is both dynamic and contextual thus providing for the possibility of error correction and it is based on the insight that contextual factors fix the sets of hypotheses and evidence that are appealed to in such inferences.⁹

⁸ See Bas van Fraassen, *Laws and Symmetry* (Oxford: Clarendon, 1989).

⁹ The theory developed here has much in common with Hintikka's view of abduction as the search for correct explanations (i.e. as abductive search) as presented in Jaakko Hintikka, "What is Abduction? The Fundamental Problem of Contemporary Epistemology," *Transactions of the Charles S. Peirce Society* 34 (1998): 503-533. He concludes that abduction is not a form of inference at all. The view defended here is that IBE is the terminal step in abductive search and that IBE is indeed a form of inference involved in that process. But, abductive search also involves seeking evidence and constructing sets of theories that are used as inputs in IBEs. In other words, abductive search includes the construction of the sample space of theories and the marshalling of relevant evidence, which are then employed in IBE inferences. This aligns with much of Jonah Schupbach's criticism of van Fraassen's objection to IBE from "Is the Bad Lot Objection Just Misguided?" *Erkenntnis* 79 (2014): 55-64. Schupbach argues that van Fraassen's criticism of IBE is misguided in that it confuses the issue of the probity of IBE inferences with the matter of the completeness and appropriateness of the input into IBE inferences. See Kyle Stanford, *Exceeding our Grasp* (Oxford: Oxford University Press, 2010) and Finnur Dellsén,

2. Constructing a theory of IBE

IBE is perhaps the most basic form of reasoning that humans engage in. Perhaps more crucially, IBE plays a central role in scientific inquiry. For example, McMullin and Lipton contend that it is *the* central form of inference in science.¹⁰ But, there has been much critical discussion of this sort of explanatory reasoning and considerations of the probity of explanatory reasoning as a distinct form of inference are most notably traceable back to Peirce's work on abduction.¹¹ On this basis, it should be clear that any suitable account of IBE must satisfy (at least) three important desiderata. First, the account must incorporate a plausible theory of explanation. It is straightforwardly obvious that we must know what an explanation simpliciter is if we are to hope to come to know what the best explanation of anything is. Second, the account must provide an explication of what it is for one explanation to be better than another explanation. Finally, the probative nature of this form of inference must be accounted for. This last aspect of any adequate account of IBE is especially important, as IBE arguments must provide warrant for their conclusions in such a way that we are entitled to provisionally accept such theoretical claims.¹² If this final desideratum is not satisfied, then it is obvious that IBE would be of no use in solving the problem of the acceptance of theoretical claims in a substantial and normative sense.

2.1 The Questions of Explanation

The 20th century history of the philosophy of science is replete with examples of attempts to provide adequate theories of explanation, and this fact is well-represented and summarized in Salmon's classic 1989 survey.¹³ The most well-

"Reactionary Responses to the Bad Lot Objection," *Studies in History and Philosophy of Science Part A* 61 (2017): 32-40 on this issue and others related to the bad lot objection.

¹⁰ See Ernan McMullin, *The Inference that Makes Science* (Marquette: Marquette University Press, 1992) and Peter Lipton, *Inference to the Best Explanation*, 2nd ed. (London: Routledge, 2004).

¹¹ See C. S. Peirce, *Collected Papers of Charles Sanders Peirce*, eds. Charles Hartshorne, Paul Weiss, and Arthur Burks, 8 vols. (Cambridge: Harvard University Press, c.1901/1931-1958).

¹² This is the general gist of van Fraassen's *Laws and Symmetry* criticism of IBE. See Samir Okasha, "Van Fraassen's Critique of Inference to the Best Explanation," *Studies in the History and Philosophy of Science* 31 (2000): 691-710, Stathis Psillos, "On Van Fraassen's Critique of Abductive Reasoning," *The Philosophical Quarterly* 46 (1996): 31-47, Stathis Psillos, *Scientific Realism: How Science Tracks the Truth* (London: Routledge Press, 1999), Timothy Day and Harold Kincaid, "Putting Inference to the Best Explanation in Its Place," *Synthese* 98 (1994): 271-295 and Stanford 2010 for extensive discussion of van Fraassen's argument.

¹³ See Wesley Salmon, "Four Decades of Scientific Explanation," in Phillip Kitcher and Wesley

known theory of course is the deductive-nomological model of explanation. However, there are numerous well-known counter-examples to this account of explanation, and, for the most part, this theory has been rejected.¹⁴ But, this need not worry us as there is a readily available alternative account of explanation that can be used to ground IBE. This model takes an explanation to be the answer to an explanatory question. As such, the best explanation will turn out to be the best answer to such a question. This account of explanation is promising because it ties explanation directly to understanding without begging any specific questions about what types of explanations are legitimate. In point of fact, it is compatible with the view that different kinds of explanations are perfectly legitimate in different contexts within a particular discipline, or in different disciplines, or at different times, etc. As such, it is perfectly compatible with the idea that methodological standards can vary with context. As we shall see this is a significant virtue of the account of IBE presented here. The modern work on erotetic logic that gave rise to the general idea of an erotetic model of explanation can be traced back to the work of Åqvist via the more or less independent work of Belnap and Steel, Hintikka, and Bromberger.¹⁵ But, the best-known and more contemporary erotetic accounts of explanation are those presented by van Fraassen and Tuomela.¹⁶ However the

Salmon (eds.), *Scientific Explanation* (Minneapolis: University of Minnesota Press, 1989), 3-219 and Wesley Salmon, *Scientific Explanation and the Causal Structure of the World* (Princeton: Princeton University Press, 1984).

¹⁴ See Phillip Kitcher and Wesley Salmon, eds., *Scientific Explanation* (Minneapolis: University of Minnesota Press, 1989) and Bas van Fraassen, *The Scientific Image* (Oxford: Clarendon, 1980) for detailed consideration of the problems with the D-N model of explanation. This is not to say, of course, that other accounts of the nature of explanation are not also problematic. For example, as shown in Michael Shaffer, "Unification and the Myth of Purely Reductive Understanding," *Organon F* (forthcoming), the unificationist view of explanation is also afflicted with serious problems related to IBE. The unificationist view is most famously defended in Phillip Kitcher, "Explanatory Unification," *Philosophy of Science* 48 (1981): 507-531, Phillip Kitcher, *The Advancement of Science* (New York: Oxford University Press, 1993) and Michael Friedman, "Explanation and Understanding," *The Journal of Philosophy*, 71 (1974): 5-19.

¹⁵ See Lennart Åqvist, *A New Approach to the Logical Theory of Interrogatives, Part 1: Analysis* (Uppsala: Filosofiska föreningen i Uppsala, 1965), Noel Belnap and Thomas Steel, *The Logic of Questions and Answers*. New Haven: Yale University Press, 1976), Jaakko Hintikka, *The Semantics of Questions and the Questions of Semantics* (Amsterdam: North Holland, 1976), Sylvain Bromberger, *On What we Know we Don't Know* (Chicago: University of Chicago Press, 1992) and Sylvain Bromberger, "Why Questions," in Robert Colodny (ed.) *Mind and Cosmos: Essays in Contemporary Science and Philosophy*, vol. 3 (Pittsburgh: University of Pittsburgh Press, 1966.: 75-100.

¹⁶ See van Fraassen, *The Scientific Image* and Raimo Tuomela, "Truth and Best Explanation," *Erkenntnis* 22 (1985): 271-299.

Michael J. Shaffer

theory of IBE developed here will be more specifically based on Hintikka's account of the logic of questions and answers, though the account presented is ultimately also rather like that proposed by Tuomela.¹⁷ However, before turning to the relevant details of that account it will be instructive to first outline some of the characteristic and general features of the erotetic model of explanation.

2.2 The Multiplicity of Explanation and Context

It has been widely acknowledged for quite some time now that a given body of data can be explained by a potentially infinite number of theories. This is just the familiar point about the underdetermination of theory by evidence. However, there is another sort of ambiguity inherent in the activity of explanation that is accentuated in the erotetic model of explanation. This is the following sort of pedagogical phenomenon. Even mild acquaintance with science and how it is generally taught should make us aware of the kind of situation in which an explanation of some phenomenon is presented, where that explanation is later revealed to be incomplete or not quite correct. For example, classical mechanics is generally taught before quantum mechanics or relativistic mechanics, and, typically the latter types of explanation of the very same phenomena are regarded as more complete and more correct. However, in general, this does not impugn the simpler explanation either as worthless or as non-explanatory. Quite the opposite is true in practice. The explanation of many phenomena in terms of classical mechanics is often retained because it is appropriate *in certain contexts*. This issue raises an aspect of explanation that has not received as much attention as it deserves from philosophers of science. This is just the context dependence of explanation.¹⁸ It is however helpful for the purposes of this paper that sensitivity to context dependence has become commonplace in contemporary epistemology, and this provides us with some guidance on the matter.

The sense in which explanation appears to be context dependent is then relevantly similar to the sense in which the terms 'knowledge' and 'justification' have been said to be context dependent in relatively recent discussions in epistemology. Specifically, Keith DeRose and David Lewis have famously defended this sort of view.¹⁹ The basic idea behind the concept of context dependence of

¹⁷ Gilbert Harman, "Inference to the Best Explanation," *The Philosophical Review* 74 (1965): 88-95, and Lipton, *Inference to the Best Explanation*, 2nd ed.

¹⁸ Ironically, the theory presented in Bas van Fraassen, *The Scientific Image* (Oxford: Clarendon Press, 1980) incorporates the contextual aspects of explanation most straightforwardly.

¹⁹ See Keith DeRose, "Contextualism: An Explanation and Defense," in John Greco and Ernest Sosa (eds.) *The Blackwell Guide to Epistemology* (Malden: Blackwell, 1999), 187-205.

epistemological concepts like knowledge is that assumptions about the epistemic standards involved in a given situation vary from context to context and so our attributions of knowledge may also vary as a result. For example, in everyday discussion skeptical hypotheses are ignored as irrelevant while in the context of a philosophical discussion about the nature of knowledge skeptical hypotheses are taken to be relevant. As such, one may have the knowledge that there is a hand before one's face in the former context, but not in the latter context *without contradiction*. This is supposed to be the case because the standards that govern the philosophical context are much stronger than those that are in place in more ordinary, everyday, contexts. This then is the crux of the contextualist view of knowledge. Whether a particular person knows a particular proposition depends on certain contextual features of the person's epistemic situation.

What will be suggested here is that explanation has a similar sort of context dependence that has gone largely unnoticed by most philosophers of science. For example, what counts as an acceptable explanation of a phenomenon in a high school physics class is different from what counts as an acceptable explanation of that phenomenon in a graduate level physics seminar, and our theory of explanation needs to reflect this fact about scientific practice. The epistemic standards that are presupposed in the latter context are much more stringent than those at work in the former, and that makes an important difference with respect to which theory we ought to accept in a given context. The main feature of the view defended here is that context determines what kinds of explanatory standards are in place in a context, the body of explanatory hypotheses to be considered and the body of evidence to be explained. Different degrees or depths of explanatoriness are then appropriate to different contexts much like different standards of evidence apply in different contexts according to epistemic contextualism about knowledge.

In terms specific to the erotetic model of explanation, this will amount to regarding the best explanation as the best answer to some why-question or how-question given some specified explanatory context. Of course this means that we will have to say something about what contextual factors need to be taken into account in general when assessing what explanation is best in a fully specified explanatory context. However, as epistemic context appears to be highly plastic and variable, it may turn out that there is not very much of interest that we can say about general epistemic standards across contexts. So, one interesting aspect of the

Keith DeRose, *The Case for Contextualism: Knowledge, Skepticism and Context, Volume I*. (Oxford: Oxford University Press, 2009) and David Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74 (1996): 549-567.

theory of explanation presented here will concern the extent to which we can claim that there are any non-contextual methodological standards that all explanations must meet. The specific view defended here is that there are some such invariant standards, but they are rather weak. This acknowledgement of the relative plasticity of explanatory contexts then in turn helps to explain the variety of explanatory practices of practitioners in different disciplines, the variety of explanatory practices at different times in the same discipline, etc.

One might be immediately tempted to object to this general account of explanation due to the perceived relativity that it imposes on the concept of explanation, and there are at least *prima facie* reasons to be sympathetic to this initial reaction. However even though such worries appear cogent it will be argued here that they are ultimately not serious worries. For the most part, this sort of worry is the result of baggage left over from previous accounts of explanation. Going back to Hempel's classic work on explanation, 'explanation' has generally been taken to be a success term and one of the chief desiderata of an adequate explanation is that it be true. So, for example, as explanation is traditionally understood, the Ising model of magnetism in solids cannot explain anything because the Ising model of magnetism is, strictly speaking, false. Given this long-standing desideratum of theories explanation it might appear that the theory of explanation sketched above will be unacceptable as it would seemingly appear to allow both that false theoretical claims can be explanations provided the correct context is present. But this problem is really a non-issue.

This is because what does not vary is whether or not a particular theoretical claim is a *potential* explanation of a phenomenon. Whether a particular theoretical claim is, or is not, a potential answer to a given scientific question is purely a matter of erotetic logic. There may be an infinite number of such answers that can be formulated with respect to any scientific question, but this does not in any way entail relativism of any sort in and of itself. Again, on the view developed here what most importantly varies with context are the epistemic standards by which we judge the *superiority* of explanations relative to one another. This involves the acceptability of the epistemic standards in question. Should the same context arise on more than one occasion, then the same evaluative ranking in terms of 'bestness' of explanation should result provided we are considering the same set of theoretical claims with respect to the same body of evidence and background knowledge. As such, substantive worries about the relativity of explanation seem largely unfounded. Such relativity as there is in this account is simply a function of the fact that the epistemic standards for acceptance of theoretical claims can vary across epistemic situation types. But, what it really indicates is just that

explanatoriness comes in degrees and that evidential standards can vary and nothing more radical than that.²⁰

In line with this, it is well-known that IBE is a form of nonmonotonic inference.²¹ For nonmonotonic inferences of this sort then a given theoretical claim T_i might be the best explanation of a body of evidence e in context B_k , while T_j might be the best explanation of e & f in B_k or of e in B .²² It is in this sense that inference to the best explanation is then a kind of ampliative and defeasible inference, and it seems as if we might be able to represent this property of IBE while at the same time allowing for a sense in which it is probative. So, we need then to determine how to represent such inferences and when we can regard instances of IBE as “good” in a clear sense. But first there are some important other factors concerning IBE that need to be examined. First and foremost, in these sorts of inferences we typically restrict our attention only to some factors that make up a relatively well-defined inferential context. In these restricted contexts evidence is typically limited to some sub-set of the total known evidence e , where we limit the set of theoretical claims considered to a sub set of \mathcal{T} —the set of all competing theoretical claims with respect to some phenomenon, and/or where we fix other particular methodological features that govern inferences. If information is added to our premises or contextual factors change, then what inferences are considered to be warranted can also change. As a result, this version of IBE reflects the defeasibility of IBE and this account of IBE squares well with the fact that, in actual practice, scientists accept theories but never make such inferences from complete bodies of evidence or from exhaustive sets of theoretical claims. This is primarily because of cognitive and computational limitations.

2.3 IBE

Preliminaries aside, we can then introduce this account of IBE. An *explanatory scientific problem* S will be taken to be a quintuple consisting of one or more why- or how-questions Q_n , a set of all competing theoretical claims \mathcal{T} indexed to elements of Q_n that minimally fulfill a set of logical criteria EXP for what counts as an answer to a given question q , where $q \in Q_n$, the total body of relevant evidence E and a context B . So, the i -th ideal explanatory scientific problem will be written as $S = \langle Q_n, \mathcal{T}, E, B, \text{EXP} \rangle$. However, as most scientific problems are complex there

²⁰ See Peter Railton, “Probability, Explanation, and Information,” *Synthese* 48 (1981): 233-256.

²¹ See Gerhard Brewka, Jurgen Dix, and Kurt Konolige, *Nonmonotonic Reasoning: An Overview* (Stanford: CSLI, 1997) and Henry Kyburg and Choh Man Teng, *Uncertain Inference* (Cambridge: Cambridge University Press, 2001).

²² See Lipton, *Inference to the Best Explanation*, 2nd ed., 92.

will be several members of Q_n , but in the simplest case—what we will call a *simple problem*— Q_n will be a singleton and $q_i = Q_n$. Where \mathcal{S} is complex there will be an appropriate number of T indexed to the elements of Q_n , and B will be similarly indexed. The solution to a given simple explanatory scientific problem—a given \mathcal{S} where Q_n is a singleton—is then T_i , the element of T which satisfies EXP and fares best in terms of E and the various standards encoded in B . More realistic and contextually restricted explanatory scientific problems will involve restrictions of T and of E . In a given context B a research group trying to answer a given explanatory question q_i may limit consideration to T_n —a few select members of T such that $T_n \supset T$ —or they may limit consideration to some sub-set ϵ_k of the total relevant known evidence E_k . For example, one crucially important way that T is restricted by B is via the introduction of idealizing assumptions.²³ In such cases, when a given idealizing assumption I is imposed in a given context it effectively rules out of consideration all theoretical claims that fail to hold under I . In other words doing so restricts consideration to I -simplified theories. Other ways of limiting T are common and include restricting consideration to extant theories, or restricting consideration to highly plausible theories, or simple differential comparisons of just two competitors, etc. So, one example of a more realistic construal of the i -th simple explanatory scientific problem can be written as $\mathcal{S} = \langle q_i, T_n, \epsilon_k, B, \text{EXP} \rangle$. Typically this reflects the fact that real scientific research concerning a simple explanatory problem involves a finite set of theories and some sub-set of the known relevant evidence in a fixed context that determines which methodological standards will be used to evaluate the competing theories. It is here that the work on bounded and ecological rationality will ultimately play an important role in understanding the probative nature of this complex form of inference. However, let us turn our attention at this point to saying a bit more about questions and their role in scientific explanation.

Following Åqvist and Hintikka, the sorts of questions we are interested in can be analyzed in terms of epistemic imperatives to bring about certain epistemic states.²⁴ So, we can analyze questions as requests by an agent to some external source of information to bring it about that the agent knows the answer. All well-formed questions of these sorts implicitly incorporate the *presupposition* of that question. The question ‘Is ϕ the case?’ presupposes that ϕ is the case or that it is not the case that ϕ , and the question ‘Why is ϕ the case?’ presupposes that ϕ is the

²³ See Michael Shaffer, *Counterfactuals and Scientific Realism* (New York: Palgrave-MacMillan, 2012).

²⁴ See Åqvist, *A New Approach to the Logical Theory of Interrogatives, Part 1* and Hintikka, *The Semantics of Questions and the Questions of Semantics*.

case. A question admits of satisfactory answers only if the presupposition of that question is true, or at least approximately true. In general we will indicate the presupposition of a given question with an expression of the form $PR(q)$. Minimally acceptable answers to questions are then propositions that allow us to understand the presupposition of that question to some degree. So, a minimally acceptable answer—or a *potential answer*—to a given simple scientific problem is a theoretical claim that at least in part explains the presuppositions of a given scientific problem. Acceptable answers to specifically scientific problems are theoretical claims that allow us to understand a phenomena or the law that the question is about.

This view then naturally looks very much like an erotetic approach to Peircean abductive/explanatory inference. However, Hintikka criticized the common view that abduction is a distinct and bona fide form of inference at all.²⁵ Against this common view Hintikka suggested that abduction is really a *search strategy* in the epistemic attempt to discover truth, as opposed to a form of inference. As Hintikka ultimately saw it, abductive search is the search for true answers to why-questions and why-questions are simply requests for explanations. So, according to Hintikka, abductive search is erotetic—it is a form of explanatory inquiry—but there is no such thing as abductive *inference* per se. The view defended here is, to a significant degree, in agreement Hintikka's. As it will be understood here, abductive search is the dynamic process of searching for explanatory answers to why-questions. But, the contention made here is that IBE is the terminal and inferential stage of abductive search. So, the position defended here is that abduction is not precisely the same thing as IBE. However, against Hintikka in particular, the view defended here is that inference to the best explanation is a form of inference employed in the broader process of abductive search, even if *abductive search* itself is not a form of inference. In any case, the attempt to construe how the members of T are demarcated with respect to some problem S requires that we address explicitly what constitutes EXP, the set of logical requirements that a given theoretical claim must fulfill in order to be considered a member of T in the context of some scientific problem.

2.4 Potential Explanations

We can now turn our attention to satisfying one of the three desiderata for an account of IBE mentioned earlier. Specifically, we can address what it is for one claim to be explanatory with respect to another. As this conceptual issue does not

²⁵ See Hintikka, "What is Abduction."

incorporate any evaluative or comparative elements the minimal requirements for membership in the set of potential answers to a given scientific problem are neither especially strong nor especially interesting. In point of fact, it will be suggested here that in an ideal world where there were no computational or physical limitations on scientific practitioners, the evaluation of which explanation is best with respect to a scientific problem would be purely a matter of logic, probability and statistics in the more formal sense. However, as has been stressed in earlier sections of this paper we do not live in such a world, and so we are often forced to simplify things by limiting our concern to those relevant theoretical claims that have been formulated and which satisfy certain additional contextual constraints, and to the relevant evidence of which we are aware. In any case we can now turn to discussion of the minimal criterion that a theoretical claim must satisfy in order to be included in the set of potential answers to a given explanatory scientific problem. As we saw earlier, for a given answer to an explanatory scientific problem to be counted as an explanation it must satisfy the basic principle EXP. EXP is then understood here as follows:

(EXP) With respect to background knowledge B and where $T_i \in B$ and $PR(q) \in E$, theoretical claim T_i is a member of the set of potential answers to a simple problem S , or $T_i \in \mathcal{T}$, if and only if (1) $P(PR(q) \mid T_i) > P(PR(q))$ and (2) for all T_j $\neg[P(PR(q) \mid T_i \& T_j) \leq P(PR(q) \mid T_j)]$.²⁶

EXP is by no means especially novel and has been assumed to be a basic tenet of theories of explanation for some time. As was alluded to earlier, we should be aware here the epistemic imperative to bring it about that the agent *knows* that p used in the erotetic analysis of explanation will have to be weakened somewhat. In the context of why-questions and recognizing that explanation comes in degrees, it seems that we really need only know that a theoretical claim raises the probability of the phenomena or law in question and that there is no other theoretical claim that wholly accounts for this increase in probability in order for a theoretical claim to be counted as a potential explanation of some data or of some lower level theoretical claim

Notice however that EXP does not narrow the range of explanations very much at all. As we noted and stressed earlier, it is well known that a non-finite number of theoretical claims can be arbitrarily constructed that satisfy EXP with respect to any problem S simply by taking a theoretical claim T_i and disjoining it with arbitrary strings of expressions. This just tells us that the purely logical aspects

²⁶ The second conjunct on the right hand side of the bi-conditional in EXP is included in order to rule out pseudo-explanations. See Alan Goldman, *Empirical Knowledge* (Berkeley: University of California Press, 1991).

of explanation are not very interesting and that they presuppose a sort of informational omniscience with respect to evidence and theory, and that we are forced by computational, cognitive and physical constraints to consider only those theoretical claims that we deem to be relevant from among those that have been explicitly formulated. In the unrestricted case T has the form $\{T_1 \vee T_2 \vee T_3 \vee T_4 \vee \dots \vee T_n\}$, while in real cases we only consider T_n of finite, and often quite small, cardinality and which hold only under idealizing assumptions. These more realistic cases of confirmation of competing theoretical claims are then often themselves cases of epistemic/methodological idealization where we are simplifying the confirmational context by reducing the number of theories that are being considered as serious candidates for confirmation by some given body of evidence that is itself restricted. As should then be obvious, the real substance of the account of theory acceptance developed here is to be found in B , the contextual factors that determine the epistemic standards in terms of which a given scientific problem is considered. In particular we must pay careful attention to those standards in addition to EXP that impact the ranking of explanations in given context. So, context determines which theoretical claims are taken to be relevant, what idealizing assumptions are allowed with respect to a given scientific problem and what factors will be used to rank explanations in addition to EXP. Context thereby determines T_n , e_n , I and the evidential and explanatory standards that characterize that explanatory scientific problem.

2.5 The Contextual Aspects of Explanation

Now we can focus our attention squarely on what might be the most interesting aspect of this account of IBE, its contextual aspects. More specifically, we can consider how epistemic context relates to epistemological standards operative in explanation. Finally, we can move on to consider in detail how we evaluate which explanation is best in a given context, and with this established we can formulate a general rule of theory acceptance based on those evaluative standards.

So, what is an epistemic context? Answering this question is of central importance in explicating the sort of account of IBE offered here, and we can get some help from looking at epistemic contextualism. There are at least two forms of contextualism and we can follow DeRose's terminology in order to locate the sort of contextualism appropriate to the sorts of explanatory endeavors in the physical sciences that we have been considering. Most crucially, DeRose distinguishes between *subject* contextualism and *attributor* contextualism.²⁷ On the one hand,

²⁷ See DeRose, "Contextualism: An Explanation and Defense."

subject contextualists hold that features of the (physical) context of the subject of knowledge vary (e.g. location), and so whether the subject knows something or not depends on those contextual factors. Certainly environmental facts about computation and cognition can impact whether we know something or not. Also, facts about the environment in which we are located can impact whether we know certain things. When, for example, a subject inhabits an environment littered with fake barns or robot cats, we might say that he does not know that he sees a barn or a cat when she is the subject of particular sensory stimulations. When a type identical subject with type identical sensory experiences inhabits an environment that is relatively free from these sorts of deceptions, we might say that he does know that he sees a barn or a cat. On the other hand, attributor contextualism holds that contextual features of the conversational context of the attributor of knowledge *to some other subject* vary, and so whether we are warranted in saying of someone that they know varies with these contextual factors. What will vary in this sort of contextualism are the epistemic standards by which we judge of someone that they are warranted in making a knowledge attribution.²⁸

By and large, however, this distinction is superficial and it is not really necessary to opt exclusively for one or the other. This is simply because both kinds of contextual features are epistemically important. They are both essentially elements of what has typically been referred to as background knowledge. The former kinds of contextual factors are empirical facts about our cognitive limitations, computational capacities, physical environments, etc., and the latter kinds of contextual factors are pragmatic factors about how we are going to apply the term 'explanation' in light of our physical and epistemic situation. Furthermore, in a sense we are all both attributors and subjects of epistemic attributions, and being aware of one's environmental context as well as being aware of one's conversational context may make one's own attributions of knowledge, or of justification, to others—or even to one's self—different. In any case, the kind of contextualism that characterizes explanatory situations involves both aspects of attributor contextualism and aspects of subject contextualism. The view developed here will be framed in terms of attributor contextualism as that view will allow us to subsume the kinds of factors that are of interest in subject contextualism. So, what we are interested in determining is when, in context *B*, an attributor *a* is justified in claiming of some subject *b* that *b* has explained *e* or *T*_{*i*} to some other agent *c*. In terms of the erotetic model of explanation outlined above, we are then ultimately interested in examining when in context *B* an attributor *a* is

²⁸ See DeRose, *The Case for Contextualism* and David Lewis, "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8 (1979): 339-359.

justified in claiming of some subject b that b has provided an acceptable answer to a why-question about e or T_i to some other agent c . In other words, we want to know when b has met the imperative implicit in a scientific explanatory request, at least to some degree.

2.6 Best Explanation and Problem Substitution in the Sciences

So, now we can turn our attention to the issue of when are we justified in claiming of someone that they have provided the *best* answer to someone's request for explanatory information in a given specific context? This is essentially the question of when in context B of an attributor a , b has explained e or T_i to c . Given this understanding of the erotetic model of explanation and our understanding of the contextual aspect of scientific explanation, we can claim that in context B a is justified in claiming of b that b has explained e (or has explained T_i) to c if and only if c has made a request 'Why e ?' or 'Why T_i ?' to b and b has conveyed to c that ' T_j ' where $T_j \in \mathbf{T}$ and T_j satisfies EXP. More importantly, we can now see that IBE can be presented in a similar manner. In context B , an attributor a is justified in claiming of some subject b that b has best explained e (or T_i) to c if and only if c has made the request 'Why e ?' or 'Why T_i ?' to b and b has conveyed to c that ' T_j ' where $T_j \in \mathbf{T}$, T_j satisfies EXP, and T_j satisfies BEST. With respect to an ideal explanatory scientific problem involving \mathbf{T} and a given body of evidence e , BEST is then characterized as follows:

(BEST) If T_j satisfies EXP, then T_j is *the best (purely logical) explanation* of e in B if and only if $\neg(\exists T_i)[(T_i \in \mathbf{T}) \& (P(e \mid T_i \& B) > P(e \mid T_j \& B))]$.²⁹

What defenders of IBE assert uniformly is that if this sort of principle is satisfied, then we are *defeasibly* warranted believing that T_j . In terms of the contextualist view of explanation presented here, what we are really allowed to say of a theory that satisfies BEST is that we are warranted in believing that T_j *in*

²⁹ This is to be understood as a partial empirical analysis of the logical aspects of explanation in the sense articulated in Carl Hempel, *Fundamental of Concept Formation in Empirical Science* (Chicago: University of Chicago Press, 1952). Also, in Lipton's 2004 terminology, best or "loveliest" explanation is not being completely identified here with likeliest explanation. The conjecture about what explanation is best offered here is that it is the theory that is most highly ranked from among competitors based on the total set of criteria present in a given context. This is meant to stave off criticisms of (virtual) triviality that apply to stand-alone account of IBE based solely on criteria like BEST. See Christopher Hitchcock, "The Lovely and the Probable," *Philosophy and Phenomenological Research* 74 (2007): 433-440 for this criticism. See Peter Achenstein, *Evidence and Method* (Oxford: Oxford University Press, 2013) for some additional criticisms of IBE.

context B . For our purposes here, notice that if we adopt BEST as a core component of a rule of theory acceptance, it allows us to assess the confirmational status of theories that are more or less realistic and it can easily be applied to cases where we are dealing with restricted sets of theories or restricted bodies of evidence.

With respect to a more realistic explanatory scientific problem involving the restriction of theories considered to T_n and to a given body of evidence e , BEST can be modified to reflect this as follows:

(BEST') If T_j satisfies EXP, then T_j is *the best (purely logical) explanation* of e in B if and only if $\neg(\exists Ti)[(Ti \in T_n) \& (P(e \mid Ti \& B) > P(e \mid T_j \& B))]$.

This then means that we can still maintain a coherent and normative sense of inference to the best explanation with respect to both ideal and realistic contexts. In what follows we will primarily deal with BEST, and we will simply acknowledge at this point that BEST' can be substituted for BEST when dealing with more realistic cases of theory confirmation. Finally, one might then define the differential degree of confirmation of theoretical claim based on a measure of explanatory power as follows.³⁰ With respect to an ideal explanatory scientific problem involving T , a given body of evidence e , and where T_j satisfies BEST and T_i is the second most likely theory relative to e ,

(CN) $Cn(T_i) = \text{diff}[P(e \mid T_j \& B), P(e \mid T_i \& B)]$.³¹

So, on this particular view the differential degree of confirmation of a given best explanation is the degree to which it is more likely than the next most likely explanation of the same evidence.³² Of course this can be similarly defined for more realistic scientific problems by replacing BEST with BEST'. Real scientific problems then can be formally understood as follows: $S = \langle q, T_n, e, B, EXP, BEST' \rangle$. As we shall see, however, there is typically much more to rules of theory

³⁰ This is but one possibility and is in no way a necessary component of the theory defended here.

³¹ See Johnah Schupbach, "Comparing Probabilistic Measures of Explanatory Power," *Philosophy of Science* 78 (2011): 813-829 and Jonah Schupbach and Jan Sprenger, "The Logic of Explanatory Power," *Philosophy of Science* 78 (2011): 105-127 for discussion of other measures of explanatory power.

³² There may also be other measures of the degree of confirmation or evidential support, but this one seems reasonable and (importantly) it is suitably differential. See Edward Erwin and Harvey Siegel "Is Confirmation Differential?" *British Journal for the Philosophy of Science* 40 (1989): 105-119 for discussion of the differentiability of inference to the best explanation. One related alternative that looks similarly promising has been articulated by Kyburg and Teng (*Uncertain Inference*, 103). It is derived from the work in John Kemeny and Paul Oppenheim, "Degree of Factual Support," *Philosophy of Science* 19 (1952): 307-324. This differential measure can be stated as follows: $Cn^*(Ti \mid e) = P(e \mid Ti) - P(e \mid \neg Ti) / P(e \mid Ti) + P(e \mid \neg Ti)$.

acceptance at work in given contexts than EXP and BEST and this is part of the background knowledge present in such cases. But, more importantly, why should we regard this sort of inferential scheme as probative? If we cannot justify the probative nature of this account, then we are not entitled to hold that such inferences have normative force. So, why is inference to the best explanation a probative form of inference?

2.7 The Probative Nature of IBE

Many philosophers have raised objections with respect to IBE for a variety of reasons, but they have typically done so without explicitly acknowledging that IBE is nonmonotonic, that it is dynamic, and that such inferences often depend on simplifying assumptions with respect to the evidence entertained and the theories considered in those inferences. With respect to this latter feature, it is crucial to understand that typical cases of IBE are normative and depend (at least) on three simplifying assumptions. The first assumption is that scientists consider only a finite set of relevant theoretical claims when assessing what is the best explanation of some phenomenon or lower level theoretical claim.³³ Second scientists consider only a subset of the total known evidence relevant to a scientific explanatory problem. Thirdly, scientists typically deal with theoretical claims that hold only under one or more idealizing assumption. As we shall see, all of these assumptions are fixed by contextual factors.

That said, the standard and supposedly damning criticism of IBE in the literature is, of course, due to van Fraassen. The primary worry that he infamously raised about inference to the best explanation concerns the idea that we have no good reason to accept the best explanation of some phenomenon from among a finite set of actually formulated theoretical claims unless we have reason to believe that the true explanation is a member of the set we are considering. Of course, van Fraassen claims that we only ever deal with very small sets of such theoretical claims when those sets are compared to the set of logically possible, but unformulated, theoretical claims. So, van Fraassen concludes that IBE is not probative because it is more likely that we are accepting the best of a bad lot, and if we are just accepting the best of a bad lot then IBE does not track the truth. In other words, as he sees it, it is irrational to accept the conclusion of any actual IBE as likely to be true. Van Fraassen entertains three potential types of responses to this line of argument and he refers to these three general strategies as follows: the privilege strategy, the force majeure strategy and the retrenchment strategy.

³³See especially van Fraassen, *Laws and Symmetry*.

The privilege response essentially involves the idea that we have some special ability to track the truth and so are entitled to believe that the true theory is among those we consider in inferring the best explanation from sets of known theories. As van Fraassen puts it, the privilege strategy depends on the dubious assumption that "...we are predisposed to hit on the right range of hypotheses."³⁴ The privilege response takes both naturalistic and rationalistic forms, but neither is at all compelling. There is simply no good reason to believe that the set of known hypotheses we deal with must contain the truth. The force majeure response involves the basic idea that we simply have no alternative and so must infer the best explanation from among the relevant set of known alternatives. But, van Fraassen rejects this response because forced choices are not necessarily rational choices. So, from the fact that we must infer the best explanation from among known explanations it does not follow that the best alternative is true. The retrenchment response involves rejecting inference to the best explanation and replacing it with an alternative account of theory acceptance. So, ultimately, he claims we are not entitled to believe in the truth of our best explanations and that we should engage in radical retrenchment in epistemology. In doing so, he rejects the appeal mysterious powers, and he is right to do so. However, his argument against the probativity of IBE is flawed and his negative assessment of the probativity of IBE is over-stated. The contention made here is that this is the case because his argument against IBE is based on an uncharitable understanding of the actual practice of inferring best explanations as it is done in actual practice.³⁵ The defense against van Fraassen's argument mounted here is then best understood as a sophisticated version of the *force majeure* response, and we shall see that it is one that enjoys considerable support from the HBP as well as the BER program.

The sense in which IBE is probative needs to account for the idea that IBE is nonmonotonic and that in inference to the best explanation we deal with incomplete information (i.e. evidence) and incomplete sets of explanatory theories.³⁶ In accord with these ideas, the appropriate notion of "goodness" for IBE is nonmonotonic and is a form of ideal case reasoning. What we are entitled to

³⁴ van Fraassen, *Laws and Symmetry*, 143.

³⁵ Specifically, it involves all the elements of abductive search as understood in Hintikka, "What is Abduction."

³⁶ So, in his "Is the Bad Lot Objection Just Misguided?" Schupbach is correct to note that van Fraassen simply misses the point when he criticizes IBE as a probative form of inference in criticizing the quality of the inputs to which IBEs are applied. When coupled with Hintikka's understanding of the dynamic nature of abductive search from his "What is Abduction?" all of van Fraassen's worries go away. IBEs are simply inferences made in dynamic contexts where we are constantly updating the sets of hypotheses and bodies of evidence to which IBEs are applied.

assert when we use IBE is that in worlds that are more epistemically perfect than but still similar to the actual world, it is the case that (at least) one of the theoretical claims in T is more likely to be true than the others. The sense in which these worlds are ideal or perfect is that in such worlds we know of all the alternative theories, we know all the relevant evidence and we are able to assess those theories in terms of BEST (and whatever other norms are in place in a given context). Since that ideal case claim is true with respect to ideal worlds, we *should* employ IBE in actual practice and so it is an appropriate norm with respect to real world science. This is a sort of Kantian approach to normativity and it is based on the following sort of argument.³⁷ A fully rational scientist would select the best explanation from among all possible alternatives on the basis of all evidence. If a fully rational scientist would select the best explanation from among all possible alternatives on the basis of all evidence, then an imperfectly rational scientist ought to select the best explanation from among all possible alternatives on the basis of all evidence. Therefore, an imperfectly rational scientist ought to select the best explanation from among all possible alternatives on the basis of all evidence. Actual scientists are, of course, imperfectly rational. Therefore, actual scientists ought to select the best explanation from among all possible alternatives on the basis of all evidence. But, we can only be reasonably expected to obey norms to the degree that we can actually do so. So, we can further reason as follows. If actual scientists ought to select the best explanation from among all possible alternatives on the basis of all evidence but they are not capable of doing this at time t , then actual scientists ought only to do their best to select the best explanation from among all possible alternatives on the basis of all evidence at time t . Therefore, actual scientists ought only to do their best to select the best explanation from among all possible alternatives on the basis of all evidence at time t . So, the best actual scientists can hope to achieve in any given context at a given time is to select the best explanation of a phenomenon from among known hypothesis on the basis of known evidence. That is typically the best that we can do in our imperfect circumstances. We are limited beings in environments that constrain our abilities to reason and so we must often substitute more easily solvable problems for those that are beyond our abilities in a given context.

So, the purely probabilistic rule BEST (in conjunction with any additional norms in our background knowledge) tells us how to evaluate theories on the basis

³⁷ The argument presented here depends heavily on the interpretation of Kant from Robert Holmes, *Basic Moral Theory*, 4th ed. (New York: Cengage, 2006). See Michael Shaffer, "Bealer on the Autonomy of Philosophical and Scientific Knowledge," *Metaphilosophy* 38 (2007): 44-54 for discussion of ideal case counterfactuals.

of evidence in such situations, and in such cases we are warranted in accepting the theoretical claim that maximizes likelihood *even if we do not actually meet the preconditions of the ideal case claim*. We can be governed by the ideal norm and yet also be warranted in following its real world correlate because we cannot do any better. The normatively correct acceptance of theories in real world contexts then amounts to our being warranted in accepting the best of a *known* lot of hypothesis on the basis of *known* evidence in a given context. In other words, it is rational for us to employ the availability heuristic. In such cases we are entitled to accept the theory that maximizes likelihood from among known theories on the basis of known evidence, at least pending the introduction of more evidence, or the introduction of new theoretical claims, or other changes in context. In essence, we must settle and accept that if the restricted set of theoretical claims *were* the set of all possible theoretical claims and the evidence of which we are aware *were* all of the evidence, then we would be entitled to accept that theoretical claim which maximizes likelihood on that evidence as true in that context. What else could we do in such a situation? In fact, to claim that IBE of this sort is irrational would commit us to wholesale skepticism about explanation and about science and it would be totally at odds with actual practice. The history of scientific practice just is the history of explaining to the degree that we currently are able and so problem substitution is the bread and butter of explanatory science. We seek to solve simpler explanatory problems first and then attempt to deal with their more complex incarnations.

However, it is clear that in typical scientific contexts there are more norms at work than just BEST. Since we do science in the actual world and not in normatively perfect worlds, we also have to do our best to close the gap between the actual world and the normatively ideal world. Properly conducted science typically requires us to attempt to gather more evidence, to generate new and better evidence using new methods, and so on. It also typically requires us to formulate and consider new competing hypotheses. As such, science is typically conducted under the assumption of the following two additional norms, the norm of evidential generation and the norm of theoretical innovation:

(EVG) We should gather and generate evidence using the best means available.

(THI) We should formulate and consider hypotheses.³⁸

³⁸ These norms are part of the more broad process of abductive search as understood in Hintikka, "What is Abduction" and IBE can them be understood as the terminal and inferential stage of such abductive inquiry.

These are then norms of bias correction that allow us to alleviate worries about the kinds of biases that can arise from the kind of problem substitution that the availability heuristic involves. EVG and THI then allow us to offer an answer to van Fraassen's worries about IBE based on the nonmonotonic and dynamic practice of inferring explanations on the basis evidence. Dynamic and contextual IBE is a defeasible but probative form of inference that says that we should always accept the best available explanation of the available evidence in a given context, *but that is by no means the end of the story at all*. We should also strive to satisfy EVG and THI so that we come closer to satisfying the ideal case norm by correcting biases over time. So, while it is true that in some context at some time we may be accepting the best of a bad lot this need not be true in the long run. From the fact that actual conditions are not normatively perfect, it does not follow that it IBE is irrational and it does not follow that it does not track the truth in the long run. In effect, what we can see is that real scientific problems are dynamic in nature. So, real dynamic scientific problems are sequences of problems with the following form: $S = \langle q, T_n, \epsilon_k, B, EXP, BEST', EVG, THI \rangle$. They are instances of the application of problem substitution involving the availability heuristic to ideal problems of the form: $S = \langle Q_n, T, E, B, EXP, BEST \rangle$. Given EVG and THI such sequences of S s will involve sets T_n and ϵ_k that are being expanded sequentially as we become aware of new evidence and new theories in our search for the truth. Typical, environmentally situated, members of such sequences will be simplified version of a complete and far more complex problem. But, solving the simpler problems very often yields insight into the answers to those complete problems. The simpler explanation provide partial understanding of the very same phenomena that the more complex explanations more fully explain. There are however some other aspects of this theory of explanation that are in need of a bit more detailed discussion, especially as they pertain to the robust evaluation of what theory is the best explanation in a given context.

2.8 The Variety of Explanatory Practices

As stressed at the beginning of this paper what is then important to recognize is that given this very general account of explanation, we can account for the variety of explanatory practices in the various sciences and their respective sub-fields in terms of the different additional methodological norms that are elements of the contexts that characterize those disciplines. So, the standards required for the confirmation of the existence of a particle in high-energy physics may be very high, this need not be true for the confirmation of a claim that a patient has a particular psychological disorder in clinical psychology. Moreover, some scientific

contexts may require that acceptable explanations are causal/mechanical, while others may require only statistical models. Some contexts may allow black box explanations, while others may not. Similarly, in some scientific contexts that characterize problems in physics or chemistry general laws may be required to explain, whereas in others such as biology or archaeology only singular causal explanations may be required to explain. Finally, we may find that more general methodological norms like simplicity, predictive novelty, conservativeness and so on characterize scientific practice in different contexts. What is of great importance is that we recognize that this aspect of the contextual theory of IBE is an asset as opposed to a problem. This is because, while the theory developed here ties explanation to understanding in a minimal and partial way via EXP and BEST and thereby unifies explanatory practice in a normative way at a very generic level, it is compatible with the observed variety of explanatory practices in the sciences and the variety of additional methodological norms that characterize individual contexts. This means then that BEST is not a full account of IBE. It is merely a core part of the theory of what counts as the best explanation in a given context and this rule can be supplemented with all sorts of additional criteria that might be elements of our background knowledge. How these additional features count in ranking hypothesis beyond the ranking imposed on the set of potential answers to a given scientific problem will itself be a function of the background knowledge present in the context of that explanatory problem. This then further suggests that there are different epistemically virtuous senses of understanding as well that correspond to the satisfaction of different sets of scientific and methodological desiderata and also that there are different degrees of explanatory understanding. So, as suggested earlier, this view is particularly well suited to the naturalistic studies of the sciences and the study of the diversity of methodological practices that we find therein. With respect to the theory developed here, what this amounts to is just the idea that we cannot really assess the confirmational status of theoretical claims absent some serious understanding of the methodological features of actual scientific contexts. Nevertheless, once we have established the details of a given context the confirmational status of a given theory can be assessed in terms of EXP, BEST' and whatever additional norms happen to characterize that context.

3. Rational Heuristics, Ecological Rationality and Explanatory Contextualism

What is then worth emphasizing here is that, from the perspective of the voluminous literature on the psychology of human reasoning, the quasi-formal and philosophical view of explanation developed in this paper enjoys considerable

empirical support. This is secured via its natural relationship to the expansive body of work on fast and frugal reasoning heuristics for problem solving and some of its close relatives, including the BER. In particular the work of Gerd Gigerenzer and Daniel Kahneman, Paul Slovic and Amos Tversky are of special importance here.³⁹ As noted throughout this discussion, one core idea behind the concepts of the HBP and of BER is that real agents do not have unlimited computational capacities, time, complete information, etc. and that the heuristic rules of inference and decision-making that real agents use are normatively appropriate only relative to specific environments for which they have been evolutionarily developed. The idea then is that we need to explore the manner in which real inferences and decisions are made by actual cognizers in order to see how it is that such reasoning is done quickly and frugally based on our actual abilities. The second core idea relevant here is the concept of ecological rationality. The idea here is that real reasoning is not the result of a generic, domain-independent, capacity to deliberate and reason in accordance with some universal rules of rationality cashed out in terms of informational omniscience. As a result, the heuristics for reasoning and decision-making advocated by this approach are the results of and work only in the specific environments in which they are generated, presumably by evolutionary adaptation.

What is then important for the purposes of this paper is that the formal model of explanation developed here is readily compatible with this more general and realistic model of reasoning and decision-making. This is primarily because of two reasons. First, inferring best explanations from known sets of hypotheses and data can be understood to be a normative heuristic guided process that reflects our finite epistemic abilities. It crucially involves problem substitution and the availability heuristic. The availability heuristic is an epistemic norm that we ought to follow, but, more importantly, it is one which we can follow. It is normative in the short run in the sense that the best available explanation of the available evidence is the most likely explanation *from that set*. It is normative in the long run in the sense that we ought to continue to gather new and better evidence and to formulate new and better theories in order to combat the kinds of biases that the availability heuristic can introduce in its short run applications. So, the dynamic aspects of the account allow for the idea that such inferences are normative but revisable in light of newly acquired evidence and newly formulated theories. The process of explanatory reasoning is dynamically rational in the nonmonotonic sense. Second, the central role that contextuality plays in the account of IBE

³⁹ See Kahneman, Slovic and Tversky, *Judgment under Uncertainty* and Gigerenzer, *The Adaptive Tool Box*.

developed here is simply a way of formally representing the ecological aspects of real-world reasoning. We infer best explanations in real contexts governed by a variety of constraints that are the result of our epistemic finitude, our real environments and our background knowledge. So, explanatory contextualism is usefully be understood to be a formal analog of the ecological facts that constrain actual human reasoning that motivate problem substitution. Facts about our abilities and the environments we inhabit constrain us in the process of abductive search in general and specifically in the ultimate stage of such inquiry, IBE. It is virtually platitudinous to assert that we can only reason in terms of what is psychologically available to us given our computational abilities. But, we can ultimately be successful in explaining and understanding the world when we realize that IBE is also dynamic. Having the best explanation of some phenomenon in one simplified context is by no means the end of abductive inquiry. The employment of the availability heuristic opens the door to bias and incompleteness, but such biases and lacuna are correctable because reasoning is dynamic and problem contexts change over time. This allows us to search for deeper and more complex explanations as context changes and we are able to contend with greater complexity or become aware of new theories and evidence.

4. Conclusion: Dynamic Contextual IBE and Abductive Search for the Truth

So, by taking the HBP and BER conception of rationality seriously—specifically by appeal to the availability heuristic and the more general notion of problem substitution—we can see that IBE, the terminal inferential stage of abductive search, is rationally grounded. Moreover, this approach to IBE allows for a more sophisticated understanding of IBE as a dynamic and contextual sort of reasoning that functions in the context of the search for explanations. So understood IBE can be defended against van Fraassen’s “best of a bad lot” objection to IBE and, contrary to van Fraassen’s claims, it is rational to accept the conclusions of IBEs even if we are not in possession of the total set of logically possible explanatory theories of some body of evidence. But, IBE is not a static kind of inference and it yields provisionally true conclusions that hold relative to the context in which they are made, but context can change and so the specific standards used to judge bestness of explanations, the set of theories considered and the body of evidence explained can change. All of this reflects actual explanatory practice in the sciences much more accurately than does the static view of IBE.