LUDWIG FAHRBACH

UNDERSTANDING BRUTE FACTS

ABSTRACT. Brute facts are facts that have no explanation. If we come to know that a fact is brute, we obviously don't get an explanation of that fact. Nevertheless, we do make some sort of epistemic gain. In this essay, I give an account of that epistemic gain, and suggest that the idea of brute facts allows us to distinguish between the notion of explanation and the notion of understanding.

I also discuss Eric Barnes' (1994) attack on Friedman's (1974) version of the unification theory of explanation. The unification theory asserts that scientific understanding results from minimizing the number of brute facts that we have to accept in our view of the world. Barnes claims that the unification theory cannot do justice to the notion of being a brute fact, because it implies that brute facts are gaps in our understanding of the world. I defend Friedman's theory against Barnes' critique.

1. INTRODUCTION

I propose the following definition schema of what it is to be a brute fact: A fact is brute when an explanation for it does not exist. If we insert our pretheoretic understanding of the notion of explanation into the schema, we obtain a pretheoretic understanding of the notion of being a brute fact. Furthermore, there are several analyses of the notion of explanation which can be inserted into the schema, giving us several different notions of what it is to be a brute fact. In this way, many theories of explanation have, as an addition, an auxiliary notion of what it is to be a brute fact.[1] Moreover, I will assume that the notion of explanation inserted into the schema determines what is meant by the term "fact". Hence, the term "fact" is a placeholder that can denote events, facts, phenomena, etc, depending on the respective notion of explanation.[2]

As an illustration of what it is to be a brute fact consider "the theory of everything", TOE, which physicists hope to find one day. Let us assume that TOE exists and can be captured in a single formula. In order to apply my proposed definition schema, we must choose an account of explanation, so let us choose the DN-model of explanation. In that case, the truth of TOE is a brute fact relative to the DN-model. The reason is simple. TOE cannot be deduced from any more encompassing theory, because it already is the most fundamental and encompassing theory. Hence,

Synthese **00:** 1–18, 2003.

there is no correct DN-argument which has TOE as its conclusion. Even if we knew everything about TOE, we wouldn't have an explanation for it, because there isn't one.

To say that a fact is brute is not to say that we don't *know* its explanation. Rather to say that a fact is brute is to say that there *isn't* an explanation for it. I will sometimes (following Barnes 1994, 64) call brute facts "ontologically brute facts" to distinguish them from what Barnes calls "epistemically brute facts". The latter are facts for which there exists an explanation, but we presently don't know it. For example, when my computer crashes I seldom know the explanation, although there normally is one. Hence, the crash is only an epistemically brute fact for me, not an ontologically brute fact, notwithstanding my strong impression to the contrary.[3]

Whether a fact is epistemically brute or not is relative to the knowledge state of a person or group of persons (e.g., the scientists of a historical period). If Smith has knowledge of more explanations than Jones, then fewer facts are epistemically brute relative to Smith's knowledge state than relative to Jones'. By contrast, whether a fact is ontologically brute or not may or may not be relative to the knowledge state of a person. This is determined by the notion of explanation adopted. I will assume throughout this paper that the correctness of explanations doesn't depend on the knowledge state of any individual or group of individuals. Hence, the same is true for the bruteness of facts, i.e., ontologically brute facts don't result from a lack of knowledge on the part of any individual. In other words, whether we know nothing, something or everything, the same facts are brute.[4]

In the present essay, I will proceed as follows. I will first compare my approach to brute facts with Barnes' approach. Afterwards, I will provide a detailed analysis of the epistemic process occurring in our belief system when we learn that a fact is brute. I will suggest that brute facts reveal a difference between explaining and understanding: A brute fact is inexplicable, but can be fully understood. I conclude with a discussion of Eric Barnes' (1994) attack on Friedman's unification theory of explanation.

2. BARNES' APPROACH TO BRUTE FACTS

Barnes' approach to brute facts differs from mine. He defines ontologically brute facts as facts that have no "explanatory basis". It's not clear precisely how Barnes understands the notion of an "explanatory basis", but I will ignore this problem here, because I want to focus on a different claim of his. According to Barnes, "brute facts are perfectly explainable: their

explanation consists of the stipulation that such facts are ontologically brute" (Barnes 1994, 65). Barnes' idea is that the explanation of a brute fact consists in the statement that it doesn't have an explanatory basis.[5]

Barnes' claim that brute facts are explicable has two undesirable consequences. First, many philosophers think that explanations provide answers to corresponding why-questions. However, consider the decay of a particular neutron at a certain time. Let us suppose that this fact is brute. If we ask "Why did this neutron decay at that time?", then this why-question has no answer, since an answer should have the form "Because so-and-so". ("Because it has no explanatory basis" is certainly not an acceptable answer.) Hence, on Barnes' approach some explanations, namely those of brute facts, don't provide answers to the corresponding why-questions.

Second, on Barnes' view all facts are explicable. A fact either has an explanatory basis or it hasn't. In the first case, citing the explanatory basis (or some relevant part of it) explains the fact. In the second case, we have a brute fact, which is also explicable for Barnes. Hence, necessarily, every fact has an explanation, and this is a consequence that is certainly hard to accept.

To avoid these consequences of Barnes' approach, I chose a different one. My definition schema precludes the explicability of brute facts along the lines of Barnes, because that would simply lead to a contradiction: The statement that a fact is inexplicable explains it.[6] In addition, my definition schema doesn't have the two undesirable consequences that follow from Barnes' account. First, it doesn't imply that some explanations are not answers to why-questions, and second, it doesn't imply that all facts are explicable, but leaves the question of whether there exist inexplicable facts open.

My definition schema has further advantages. First, it is largely neutral with respect to the different models of explanation. It thereby allows us to discuss brute facts and make claims about them without having to settle on a specific model of explanation. This is an advantage, because the debate over which model of explanation is correct is going on with little sign of an early resolution.

Second, it may very well be that there are several types of explanation, each with its Appropriate model. If this is so, my definition schema has the advantage of generality. It offers a notion of being a brute fact that can be combined with all or most models of explanation. For example, it allows us to describe both the truth of the "Theory of Everything" as a brute fact (relative to a covering law model of explanation) and the properties and behavior of elementary particles as brute facts (relative to mereological models of explanation, for example).

In what follows, I'll describe the features of brute facts that are sufficiently general to be compatible with, and plausible for, most models of explanation. Nevertheless, we should not expect that we can always achieve full generality. We should count the definition schema and some associated general claims a success even if some models of explanation fall outside of their scope. For some theories of explanation, my definition schema of brute facts works poorly or does not work at all. In such cases we must correct the definition schema and the associated general claims, or else acknowledge the limits of their scope.

3. TWO SPECIAL KINDS OF BRUTE FACT

In this section I will briefly discuss two special sorts of brute fact. I mention them to help to avoid mixing them up with the main type of brute facts with which this paper is primarily concerned. Additionally, the second kind of brute facts leads to interesting answers to the question of which facts are brute in our world.

The first sort of brute fact derives from the interest relativity of explanation. Whether something is a good explanation depends in part on the interests of the person requesting an explanation. For example, a psychologist might only be interested in those causal factors of a car accident that are mental. If there are no mental factors causally relevant to the car accident, he may consider the accident to be a brute fact. Hence, some facts may be considered brute because of interest relativity.[7]

For the rest of this section, let us disregard brute facts that are due to interest relativity. In order to discuss the second special kind of brute fact, let us briefly consider the question which facts are brute in our world. I'll restrict myself to a few remarks on this question, because, as we will see, a serious discussion of it would lead to large issues such as the nature of mental causation, the relationship between "basic" and supervenient states, and the status of explanations in the special sciences. So, which facts are brute in our world? We already encountered TOE, the theory of everything. Next, according to most interpretations of quantum mechanics, events such as the decay of a neutron at a particular time are brute facts. Furthermore, the Big Bang and the values of constants and initial conditions at the beginning of the universe may well be brute.[8] In a mereological conception of explanation where the properties of wholes are explained by properties of their parts, brute facts may relate to features of fundamental particles or of space-time points. Of course, whether any of these are genuine examples of brute facts depends not only on the makeup of our world, but also on the theory of explanation adopted.[9]

All of the examples we have considered so far are drawn from fundamental parts of physics and cosmology. The question therefore arises whether there are brute facts in empirical sciences other than physics and cosmology. If not, then physicalism would be true in the sense that all explanation comes to an end in physics and cosmology and nowhere else. If physicalism is true in this sense, then our world has the remarkable feature that brute facts are very rare in it, so that we seldom encounter them in science or every-day life. But again, the answer to the question of whether there are brute facts in non-physical sciences depends on the notion of explanation we employ in our analysis of bruteness. An interesting example of an understanding of the notion of explanation that leads to a proliferation of brute facts is David Owens' account of coincidences.

Owens (1992) holds that coincidences are inexplicable. A coincidence is an event that can be decomposed into several constituent events which have independent causal explanations, i.e., their causal histories don't share any causal factors.[10] Consider the collision of two particles P and Q at location *l* at time *t*. Assume that the two trajectories leading to the Collision are independent in the sense that they don't share any common causal factor. Then, according to Owens's account there is a causal explanation of why particle P is at location *l* at time *t*, and there is a causal explanation of why particle Q is at location *l* at time *t*, but there is no causal explanation of why both particles are at the same place at the same time. Owens denies what he calls "agglomerativity" of explanations, i.e., he maintains that the explanations of the constituent events cannot be combined to yield an explanation of the simultaneous appearance of both particles at l at t. Consequently, Owens maintains that the collision cannot be explained. It is a brute fact. If, on the other hand, two particles R and S are shot at each other by physicists in a particle collider, then their collision is not a coincidence, because there are common causal factors, e.g., certain intentions of the physicists.[11]

Owens uses his account of coincidences to argue that the special sciences like biology, chemistry, and so on are autonomous in the sense that they provide explanations that are not underwritten by physical explanations. Consider the causal explanation of a certain occurrence of inflation by reference to an increase in money supply. This is a perfectly good economic explanation. Owens assumes that supervenience is true so that the two events, the occurrence of inflation and the increase in money supply, each supervene on a large set of physical conditions. Call the two sets 'MM' and 'INF' (for "more money" and "inflation"). Then Owens holds that MM cannot explain INF, because (among other reasons) the causal processes leading from elements of MM to elements of INF are

mostly independent of each other, and independent explanations cannot be agglomerated to yield an explanation of the composite event INF. Thus, INF is a physical coincidence; it has no physical explanation. Nevertheless, INF can be explained economically, by citing the increase in money supply. In this way, economics can explain events that are inexplicable for physics, and explanation does not only come to an end in physics, but in other sciences as well. This shows that the question which facts are brute in our world has no easy answer, but heavily depends on the account of explanation chosen.

4. TWO KINDS OF EPISTEMIC CHANGE

When we learn the explanation of a fact, an epistemic change occurs in our belief system. Let p be a fact whose explanation we don't know. When we request an explanation of p, we are in an "epistemically imperfect state vis-à-vis the fact that p, and there is something we need to know, or at any rate something epistemic we must acquire, if we are to ameliorate our epistemic situation".[12] When we receive an explanation of p, a change takes place in our system of beliefs: The explanation imparts an *epistemic gain* on our belief state leading to an improved knowledge state. Moreover, the epistemic gain imparted by an explanation is a gain in our *understanding* of the world. When we don't know the explanation of a fact, we suffer a gap in our understanding. When we are given an explanation of the fact, the gap in our understanding is closed and our understanding of the world is improved.

In order to gain clarity about the nature of explanations Kim (1994) poses what he calls the Epistemological Question: "What is it that we know – that is, what exactly is our epistemic gain – when we have an explanation of p?" Kim maintains that this is "the central question of any theory of explanation".[13] He discusses how different models of explanation answer this question. In the case of DN-explanations, for example, the epistemic gain consists in the grasping of the nomic expectability of the event (Kim 1994, 56). For the purposes of my discussion, however, the Epistemological Question is more interesting than its answers, because it suggests an analogous question about brute facts, which will provide a fruitful basis for examining the notion of being a brute fact.

Let us now analyze the epistemic change that occurs in a person's belief system when he learns that a fact is brute. Let X be a fact that is known to be brute (relative to some given model of explanation). For example, X may be the theory of everything, TOE, that physicists hope to find one day, or X may be the fact that there was a Big Bang which, let us assume, has

no cause. Let Smith be somebody who doesn't know these things about X, who is puzzled by X and who wants to know why X is the case. For example, Smith wants to have a DN-explanation of TOE or wants to know the cause of the Big Bang. Smith asks, "Why does X obtain?" expecting as answer an explanation of X. He asks that why-question because – in the parlance of Kim – he thinks he is in an epistemically imperfect state vis-à-vis X and hopes to improve his epistemic situation by receiving the epistemic gain conveyed by an explanation.

But now Smith receives an unexpected response to his question. To his question "Why X?" he doesn't get an answer of the form "Because so-and-so". Instead, he is told that no explanation for X exists. He is told, for example, that the Big Bang has no cause and hence no explanation, that it is a brute fact, which simply came to be. Let us assume that Smith accepts this information. Then, inspired by Kim's Epistemological Question about explanations, we can ask an analogous question about brute facts:

> When Smith comes to know that X is a brute fact, what kind of epistemic gain does he acquire, i.e., what kind of knowledge does he gain? Does he achieve any epistemic gain at all?

In the following section, I want to highlight two features of the change in Smith's epistemic system, a negative feature and a positive feature.

5. TWO FEATURES

To appreciate the negative feature consider the following pitfall. We may be tempted to claim that when Smith learns that X is brute, he does receive an explanation of X after all: The explanation consists in the statement that X is brute. For example, the statement that the decay of a neutron at a particular time happened by chance may sound like an explanation of the decay. But, of course, this is a contradiction. X has no explanation, and learning that this is so cannot be learning an explanation of X. The negative feature reminds us of this pitfall. It simply consists in the epistemic process not being one of learning of an explanation. It implies that Smith has to do without the pleasant feeling that usually accompanies the grasping of an explanation.

The negative feature can also be described in terms of Kim's general notion of epistemic gain imparted by an explanation: When Smith is told that X is brute, he doesn't acquire the kind of epistemic gain which one acquires when one is given an explanation. For example, in the case of DN-explanations of singular facts, the epistemic gain consists in our grasping

	Facts for which Smith thinks he has an explanation		Facts for which Smith thinks there exists an explanation, which he does not know	Facts about whose explicability Smith has no opinion	Facts that Smith thinks are ontologically brute
Facts that have an explanation	Smith has the right explanation	Smith has a wrong explanation			
Ontologically brute facts			alpha		beta

Figure 1. The thickly bordered area contains the epistemically brute facts for Smith, i.e., those facts for which there exists an explanation, but Smith doesn't know it.

the nomic expectability of the event to be explained. In the case of X, there is no set of laws and antecedent conditions based on which expectations about X could possibly be formed; hence, no grasping of the nomic expectability can take place.

The negative feature cannot be the whole story. If Smith had asked for an explanation of a fact that has an explanation, but science currently doesn't know it, he also wouldn't have received an explanation.[14] Hence, this epistemic process would also have the negative feature. In neither case would Smith get what he wants. But from now on, the two cases differ. The reason why Smith doesn't receive an explanation for X is, of course, that none exists. Science knows that this is so (we assumed) and this information is conveyed to Smith. This is a negative sort of information about X (something doesn't exist), nevertheless it constitutes an epistemic gain for Smith.[15] He learns that what explains other events or phenomena doesn't exist in the case of X. This epistemic gain is the positive feature of the epistemic process. It is valuable for Smith, because it implies that he need not further search for an explanation of X.

We can get an overview of epistemic processes akin to the one occurring in Smith's belief system with the help of Figure 1. In it, every fact falls in one of nine cells. The vertical dimension of Figure 1 represents the objective perspective, i.e., whether a fact has an explanation or is ontologically brute. The horizontal dimension represents Smith's subjective perspective, i.e., his opinion about a fact. Epistemic processes akin to Smith's consist in the movement of a fact from one cell to a different cell in the same row.[16] Thus, Smith's learning that X is brute moves X from the cell labeled "alpha" to the cell labeled "beta".

6. THE ORDER OF NATURE

As I have described it so far, what happens in Smith's epistemic system is not different from what generally happens when somebody asks a question that has no answer. For example, when somebody asks who committed the murder, in a situation where no one was in fact murdered, there is also a negative feature (the expected information isn't forthcoming) and a positive feature, which settles the matter (the questioner receives the information that the expected information doesn't exist). However, the information that X is brute is special in at least two respects.

In order to discuss the first respect in which brute facts are special, I have to introduce two further notions.

For many models of explanation, an explanation can be seen as conveying information about, or in some other way referring to, a *dependency relation* between the explanandum fact and other facts.[17] The *order of nature* is defined as the complete relational structure of all such dependencies.[18] Simplifying somewhat, we may think of the order of nature as a huge graph where the nodes are the facts or events, and arrows connect some nodes with another node if that node depends on the nodes where the arrows start. If our world is highly unified (as many believe it is) then the order of nature is one huge connected graph. If our world is not that unified (as many others argue) then the order of nature decomposes into several unconnected or loosely connected subgraphs.

Let us now consider these two ideas – of the dependency relations and of the order of nature – in the different models of explanation and consider their metaphysical status according to those theories. In the case of Hempel's DN-model, a representation of the order of nature is given by amalgamating all correct DN-arguments into one huge text. Here, the dependency relations derive from the logical relations of the DN-arguments between explanandum sentences and explanans sentences. Hence, Hempel's order of nature is not really a part of the world, but projected by us onto the world. The same is true for unification theories of explanation, to which I will return at the end of the paper.

Other theories of explanation conceive of the order of nature as an objective part of the world. According to most causal theories of explanation, explanations convey information about causal relations which are part of the external world. In this case, the dependency relations can be identified with the causal relations between events, and the order of nature can be identified with the whole causal net. If explanations are seen more generally as referring to objective dependency relations in the external world,

the dependency relations are these objective dependency relations and the order of nature is the totality of these objective dependency relations.

The dependency relations that hold between the explanandum facts and other facts are meant to be independent of the interests of any individuals, e.g., independent of the interests of the explainee (the person requesting an explanation). We can then account for one sort of interest-relativity of explanations in the following way: The explainee's interests carve out from the order of nature a piece to which the requested explanation has to refer. For example, the interests of a psychologist seeking an explanation of a car accident in terms of psychological factors determine a subset of the order of nature, namely the subset of psychological facts or events and the dependency relations between them. In what follows, I'll focus on brute facts that don't result from interest relativity. Nevertheless most of what I say holds mutatis mutandis for brute facts that do result from interest relativity. I'll also disregard brute facts that are coincidences in the sense of Owens (1992, ch. 1).[19]

Now, to see the first respect in which the information that a fact is brute is special, let us consider brute facts from a more global perspective using the idea of the order of nature. Consider once more X, the brute fact that Smith encountered. What is the place of X in the order of nature? Because X does not stand in the dependency relation to any other facts, X is a *starting point* in the order of nature;[20] it belongs to the *brink* of the order of nature, where the brink consists of all brute facts and all facts that don't explain any other facts. For example, the Big Bang is a starting point in the causal structure of the world, which belongs to the brink of the causal structure of the world, which, in turn, consists of all uncaused events and all epiphenomena (i.e., events that don't cause other events). In the case of mereological explanations, the brink of the order of nature may consist of elementary particles or space-time points and the whole universe.

Returning to Smith, we see that when he learns that X is brute, his epistemic gain consists in the information that X is a starting point in the order of nature. It makes him better informed about the order of nature and the place X occupies in it. This is how the information that X is brute is special. Interestingly, all of this holds independently of whether the order of nature is conceived of as part of the world or as projected by us onto the world.

Adopting such a global perspective, where the whole order of nature is envisaged, may also help to dispel any remaining feelings of puzzlement or mystery about brute facts. Smith might still harbor such feelings about X, because he might insist that, in spite of everything, he still doesn't know why X happened or occurred. Seen from a global perspective, it's

plausible that some facts have to be brute, and if this is so, it follows that all residual feelings of puzzlement or mystery about brute facts lack any rational basis.[21]

On the other hand, arguments for the mere existence of brute facts don't tell us which facts exactly are brute. This leads us to the second respect in which the information that a fact is brute is special. It is not easy to obtain such information. That is to say, it is not easy to find support for the claim that a given fact is brute. This problem is worsened, if our notion of explanation and our world are such that brute facts are very rare in it. In this case, it is prima facie rather improbable that a given fact is brute. One might even doubt that we can ever know that a fact is brute, because we can never exclude the possibility that there is an explanation of that fact which we haven't discovered yet.

Let us briefly consider some examples. If there is such a thing as a Theory of Everything that is true of our world, it seems possible that physicists could accumulate sufficiently diverse and extensive evidence to confirm that it indeed encompasses everything and consequently is brute. In the case of the decay of a neutron, we have to deal with the problems of Quantum Mechanics, its different interpretations, etc. In general, we have to assess the empirical support of the theories from which it follows that a given fact is brute, and find out whether those theories are not reducible to more fundamental theories. Although I don't see why it should be impossible in principle to accomplish these tasks, this issue is clearly in need of much further discussion.[22]

7. BRUTE FACTS AND UNDERSTANDING

Let us now relate the epistemic process that occurs in Smith's epistemic system when he learns that X is brute to the notion of understanding. This will allow us to compare the concept of explanation with the concept of being a brute fact.

It is plausible to assume that every case of learning something about the order of nature results in an increase in our understanding of the world. Hence, the epistemic process that occurs in Smith when he learns that X is brute results in an improvement in his understanding of the world. When Smith requests an explanation of X, it is because his understanding of the world suffers from a gap. It was incomplete with respect to X. When he gains the information that X has no explanation, the gap in his understanding is closed. For example, upon learning that there is no cause for the Big Bang, Smith's understanding of the Big Bang improves.

Now, let us assume that Smith not only learns that X is brute, but comes to know everything else about X and its place in the order of nature. If this is the case and if his epistemic state is optimal with respect to X, then it is plausible to assert that his understanding of X is complete.[23] For example, if a neutron decays at a particular time and Smith knows everything about neutrons in general and about that neutron in particular, then he has a complete understanding of its decay. It no longer represents a scientific mystery for him. So brute facts can be completely understood. Nevertheless, they have, of course, no explanation. Hence, the concept of being a brute fact provides a basis for distinguishing between the concept of explanation and the concept of understanding. Brute facts are inexplicable, but can be fully understood.

We are now able to compare the two kinds of epistemic changes, namely learning of an explanation versus learning that a fact is brute. On the one hand, Kim's Epistemological Question about explanations and my analogous question about brute facts have different answers, that is to say, the epistemic gain imparted by an explanation is different from the epistemic gain imparted by the statement that a fact is brute. If we don't observe this difference, we run into the contradiction mentioned above, namely, that the inexplicability of a fact explains that fact. On the other hand, the epistemic gain imparted by an explanation is a gain in understanding of the world, and now we see that the same is true for the epistemic gain imparted by the information that a fact is brute. Hence, we have two different kinds of epistemic processes and associated with these two kinds of processes two different kinds of epistemic gain, but in spite of this difference they are closely related, because both close gaps in our understanding of the world.

8. BARNES' CRITICISM OF FRIEDMAN

For Friedman, explanatory power is a holistic feature of a belief system. It is not a property of particular items like arguments or derivations, but a property of the belief system as a whole. It results from the unification of the phenomena we accept, (more precisely from the unification of the law-like sentences we accept, where law-like sentences represent phenomena).[24] Unification, in turn, is achieved by organizing the accepted phenomena in the right way into a system of derivations. The derivations of such a system are the explanations, and their conclusions are the facts-to-be-explained. Thus, Friedman's account distinguishes between explanations, which are derivations, and explanatory power, which is a

property of the whole system. Those phenomena for which no derivation inside the respective system exist are the brute facts.[25]

Now, how do we organize the accepted phenomena in the right way to achieve unification? According to Friedman, we progress towards unification by reducing the number of independent phenomena, i.e., the number of brute facts, that we have to accept. According to Friedman, "science increases our understanding of the world by reducing the total number of independent phenomena that we have to accept as ultimate or given" (Friedman 1974, 15). Hence, a theory is genuinely explanatory if it helps reduce the number of brute facts that we must posit. For example, the kinetic theory did this by implying other phenomena that could not be derived before the kinetic theory was introduced. As Friedman says: "Where we once had three independent brute facts – that gases approximately obey the Boyle–Charles law, that they obey Graham's law, and that they have the specific-heat capacities they do have – we now have only one – that molecules obey the laws of mechanics" (Friedman 1974, 14).

Barnes contends that Friedman's unification theory has a problem with the notion of being a brute fact. For the unification theory the goal of minimizing the number of brute facts is constitutive of explanation. Barnes seems to think that this goal implies that we should try to get rid of every brute fact that we encounter, and that we should always regard brute facts as something undesirable that should be eliminated. Barnes also contends that for the unification theory of explanation brute facts always represent a lack of scientific understanding. He claims that "Friedman ... clearly views ontologically brute facts as unexplainable scientific mysteries ... the absence of mystery being construed as equivalent to the presence of understanding" (Barnes 1994, 62) This would contradict the claim on which I agree with Barnes hat if we know everything about a brute fact, including that it is brute, then it doesn't represent a lack of scientific understanding and isn't a scientific mystery.

9. A DEFENCE OF FRIEDMAN

Barnes' criticism of Friedman is unsound. To see this we must look more carefully at Friedman's theory of unification. As Friedman's formal account is rather involved, I will give a simplified version of it, which nevertheless suffices to rebut Barnes' criticism.

Friedman conceives of the unification of the set of accepted law-like sentences (which represent phenomena) as some kind of axiomatization where all accepted law-like sentences are derivable from a set of axioms. Such axiomatizations compete for the title of best unification. According

to Friedman, the axiomatization that has the fewest axioms deserves the title of best unification. However not all axiomatizations are allowed to compete. In counting axioms, we are not entitled to count the conjunction of a set of axioms as a single axiom. For that reason, Friedman requires that the axioms of the axiomatizations be "independent", i.e., they should not be equivalent to conjunctions of law-like sentences. Friedman tries to spell out this idea by using the notion of "independent acceptability", the details of which need not concern us here. The only interesting thing to note is that the notion of "independent acceptability" is based on the notion of "good grounds for belief", which is an epistemic notion, although an objective one (Friedman 1974, 8).

Two clarifications are in order. First, because Friedman wants the notion of explanation to be an objective one, it should not depend on the knowledge state of a subject. In spite of this, Friedman's analysis refers to the set of phenomena accepted by a subject at a particular time. To get an objective notion of explanation we have to look at the phenomena that really exist, not the phenomena accepted by a subject or community at a given time. On the other hand, unification concerns the organization of beliefs or sentences and is based on logical and epistemic notions such as "good grounds to believe". In Kim's words, unification is something on the representational side, not on the worldly side of things. Therefore, in order to define explanatory power, we should look at a belief system of an ideal subject who accepts all and only the phenomena that exist in our world. We can think of this subject as a scientist at the ideal limit of inquiry who knows all true law-like sentences. For this belief system, the best unification is determined in the way just described, namely by first allowing for competition only axiomatizations that have "independent" axioms, and second, by choosing from these axiomatizations the one with the smallest number of "independent" axioms. This unification, then, is the best unification of the phenomena of our world. It is the order of nature defined above, i.e., the relational structure of all explanatory dependency relations.[26]

The second clarification concerns Friedman's use of the term "brute fact". Friedman uses that term more broadly than I do. His "brute facts" are the axioms of *any* axiomatization of *any* belief system, where the axiomatization need not be the best one and the beliefs need not be true. Hence, they are relative to a belief system and relative to an axiomatization of that belief system. By contrast, I do not consider axioms of axiomatizations whose number of axioms aren't minimal to be proper brute facts; they are only candidates for brute facts. The proper brute facts are those of the best axiomatization of all the phenomena of *our world*.

With the help of these clarifications, we can rebut Barnes' criticism. Consider the ideal belief system where all the phenomena of our world are known. Then the goal of minimizing the number of brute facts concerns brute facts in Friedman's sense, i.e., axioms. It is the goal of selecting the axiomatization with the smallest number of axioms. Barnes misconstrues this goal, when he takes it to imply that every single brute fact (i.e., axiom) is undesirable and should be eliminated. When we have found the best unification with the fewest number of brute facts, the remaining brute facts cannot be eliminated. (Actually, they are what I just called the proper brute facts.) The goal of unification is strictly holistic. It concerns the *number* of brute facts and cannot be taken to imply that every one of them should be eliminated.

An analogy may help to see the point more clearly. For the unification theory, explanatory power is like thrift. The goal of minimizing the number of brute facts is like the goal of minimizing the number of dollars that we have to spend when we want to buy something. In the case of thrift, if we find the best deal, we have achieved the goal of thrift and the dollars we then have to spend should not distress us. (We might still resent the money that we have to spend on the best deal, but such a feeling is not appropriate. Similarly any remaining feeling of puzzlement about the proper brute facts in the best unification is not appropriate.) The goal of thrift is truly holistic, it concerns the *number* of dollars spent. It cannot be taken to imply that for every single dollar we should avoid spending it, because we cannot avoid spending some dollars even if we find the best deal. The holistic nature of thrift can also be recognized in the fact that we are only interested in the number of dollars rather than in individual coins or banknotes. Only numismatists are interested in individual coins or banknotes. Similarly, when unifying our picture of the world, we are only interested in the number of brute facts, not in any one of them individually.

The second part of Barnes' criticism of Friedman concerns the notion of understanding. Barnes contends that Friedman is committed to the claim that brute facts always represent a lack of scientific understanding of the world and forever remain scientific mysteries. To assess this criticism we have to distinguish between Friedman's theory (as described above) and a specific remark he makes. Firstly, his theory is clearly compatible with the view that brute facts don't represent a lack of scientific understanding. His theory can be easily supplemented by the claim that if we have found the best unification of the ideal belief system and know we have found it, then our understanding of the world is complete and the remaining brute facts, i.e., the proper brute facts, neither represent scientific mysteries nor gaps in our understanding. This supplement also agrees well with the result

of the two previous paragraphs that the proper brute facts should not be considered undesirable. Hence, Barnes criticism doesn't harm Friedman's theory.

Although Friedman's theory is untouched by Barnes' criticism, Friedman himself makes a statement that suggests that brute facts do represent a lack of scientific understanding of the world. Friedman states that "[a] world with fewer independent phenomena is, other things being equal, more comprehensible than one with more" (Friedman 1974, 15). Barnes directly challenges Friedman's remark, and rightly so. He holds that "all possible worlds ... are equally understandable in principle" (Barnes 1994, 65), and that we completely understand a world, if we know the explanations of all facts of that world, and if we moreover know that no further explanations exist. Hence, Friedman's remark has to be corrected. But, as we just saw, it can easily be corrected without harming Friedman's theory.

10. CONCLUSION

I have analyzed the epistemic process that occurs in our belief system when we learn that a fact is brute. The epistemic process has two features, a negative feature and a positive feature. The negative feature is simply that we don't get an explanation for that fact. The positive feature is the epistemic gain imparted by the statement that an explanation for that fact doesn't exist. As brute facts are starting points in the order of nature, our learning that a fact is brute informs us about the order of nature, and thereby contributes to our understanding of the world. I compared this sort of epistemic gain (learning that a fact is brute) with the sort of epistemic gain imparted by an explanation of a fact. These two sorts of epistemic gain should be carefully distinguished. Nevertheless, they are closely related, because both exhibit the place of the fact in the order of nature and both close gaps in our understanding of the world.

NOTES

[1] I will mainly deal with the following theories of explanation: covering law theories, causal theories, objective dependency and objective determination theories and unification theories.

[2] This paper grew as a reaction to Barnes' stimulating paper "Explaining Brute Facts". Nevertheless, my approach differs significantly from Barnes'. I will later say where I differ and why.

[3] Figure 1 in the section 5 shows precisely which kinds of facts are epistemically brute.

[4] Accounts for which the correctness of explanations depends on the explainee's knowledge state include Hempel's inductive-statistical model (compare Coffa 1974) and Achinstein's account (1984, 287).

[5] Barnes' main concern is to criticize several models of explanation on the grounds that they cannot do justice to the notion of being a brute fact. In the last sections of the paper I'll return to Barnes' criticism of Friedman's unification theory of explanation and rebut it.

[6] The contradiction only arises if some brute facts exist.

[7] I will later propose a way of separating the interest-relative and interest-independent aspects of explanation.

[8] See Parfit (1998) for some interesting speculations about the limits of explanation in the matter of general features of the universe and its very existence.

[9] This is particularly true for the case of probabilistic explanations of indeterministic events, e.g., the decay of a neutron. For an overview of the main approaches to probabilistic explanation see Ruben (1990, 12–39).

[10] Owens (1992, 6). Owens' official definition of coincidences is more complicated. It is part of a larger project. He analyzes causal explanation in terms of coincidences, and coincidences in terms of necessary and sufficient conditions.

[11] Elder (2001) distinguishes coincidences from accidents. Both kinds of facts have independently produced constituent facts, but coincidences are conjunctions of facts, e.g., that P is at *l* at *t* and Q is at *l* at*t*, whereas accidents are relational facts, e.g., that P and Q are at the same place at the same time. He argues along the lines of Owens that both kinds of facts are inexplicable.

[12] Kim (1988, 54), also compare Ruben (1990, 7).

[13] Ruben (1990, 7) poses a similar question.

[14] Such a fact is epistemically brute relative to current scientific knowledge.

[15] Railton (1981, 175) describes the negative feature, but misses the positive feature. Compare also van Fraassen's notions of rejection and correction of illegitimate why-questions (1980, 111, 140, 145). van Fraassen is concerned with the pragmatics of brute facts, whereas I focus on epistemic issues.

[16] Facts can only be moved horizontally, because their status as either explicable or brute cannot change. This is so, because of the presupposition made in the introduction that the notion of explanation is not relative to the knowledge state of the subject.

[17] Here, the term "fact" is once more a placeholder.

[18] What I offer here is an obvious generalization of Kim (1974, 1988, 226; 1994, 67–69), Ruben (1990, 125), Kitcher (1989, 436), and others.

[19] Smith will have no problem distinguishing between interest relative brute facts, coincidences and brute facts proper. He knows his interests and can therefore determine whether X is an interest relative brute fact or not. Furthermore, he can determine whether X is a candidate for a coincidence, because coincidences are by definition decomposable into two or more events.

[20] Compare Ruben (1990, 126).

[21] For some considerations concerning whether all facts are explicable or some brute facts have to exist compare Ruben (1990, 125–129).

[22] See, for example, Lange (2002, chap. 4).

[23] See Barnes (1994, 65).

[24] The unification theory deals only with phenomena, i.e., general regularities, not with singular facts or events, hence, in this section, all facts-to-be-explained and all brute facts are phenomena.
[25] Friedman also calls them "independent phenomena".
[26] Because the notion of "good grounds to believe", on which Friedman bases his account of unification, is an epistemic notion, the order of nature is not part of the world, but projected by us onto it. Thus, Friedman's unification theory is an anti-realist theory of explanation (compare Kim, 1994, 63). Put very crudely, Friedman's order of nature is determined by both world and rationality.

REFERENCES

Achinstein, P.: 1984, 'The Pragmatic Character of Explanation', *Proceedings of the Philosophy of Science Association 1984*, Vol. 2, pp. 275–292.
Barnes, E.: 1994, 'Explaining Brute Facts', *Proceedings of the Philosophy of Science Association 199*, Vol. 1, pp. 61–68.
Coffa, A.: 1974, 'Hempel's Ambiguity', *Synthese* **28**, 141–163.
Elder, C.: 2001, 'Mental Causation versus Physical Causation: No Contest', *Philosophy and Phenomenological Research* **LXII**(1), 111–127.
Friedman, M.: 1974, 'Explanation and Scientific Understanding', *Journal of Philosophy* **71**, 5–19.
Kim J.: 1974, 'Noncausal Connections', *Nous* **8**, 41–52.
Kim, J.: 1988, 'Explanatory Realism, Causal Realism, and Explanatory Exclusion', in Howard K. Wettstein and Peter French (eds), *Realism and Antirealism. Midwest Studies in Philosophy*, *12*, University of Minnesota Press, Minneapolis, pp. 225–239.
Kim, J.: 1994, 'Explanatory Knowledge and Metaphysical Dependence', in E. Villanueva (ed.), *Philosophical Issues*, 5, *Truth and Rationality*, Ridgeview Press, pp. 51–69.
Kitcher, P.: 1989, 'Explanatory Unification and the Causal Structure of the World', in P. Kitcher and W. C. Salmon (eds), *Minnesota Studies in the Philosophy of Science*, Vol. 13, University of Minnesota Press, Minneapolis, pp. 410–505.
Lange, Marc: 2002, *An Introduction to the Philosophy of Physics: Locality*, *Fields*, *Energy*, *and Mass*, Blackwell Publishers, Oxford.
Owens, D.: 1992, *Causes and Coincidences*, Cambridge University Press, Cambridge.
Parfit, D.: 1998, 'Why anything? Why this?', *Times Literary Supplement*, July 3, 3–5.
Railton, P.: 1981, 'Probability, Explanation, and Information', *Synthese* **48**, 233–256.
Ruben, D.-H.: 1990, *Explaining Explanation*, Routledge, London.
van Fraassen, B.: 1980, *The Scientific Image*, Clarendon Press, Oxford.

Department of Philosophy
University of Konstanz
Hoheneggstrasse 102
78464 Konstanz
Germany
E-mail: ludwig.fahrbach@uni-konstanz.de