Metaphysics // Fall 2017

Handout 20

Ground: Plato, general remarks

Grounds of piety. In the selection before us the question is about the nature of piety, about what piety *is*. Euthyphro's proposal, his third in the dialogue, is to identify piety with being loved by (*all*) the Gods. ('All' is needed because Euthyphro got earlier into trouble with the fact of possible disagreements between different Gods.) But the question he poses at the outset is: 10a

> Is the pious loved by the Gods because it is pious, or is it loved by the Gods because it is pious? (20-1)

Now, after the analogy with 'being carried' and 'being seen' has been developed, Euthyphro is called to consider these propositions: 10d

(A) What is pious is loved by the Gods because it is pious. [T]
(B) What is God-loved is God-loved because it is loved by the Gods. [T]
(C) What is God-loved is loved by the Gods because it is God-loved. [F]
(D) What is pious is pious because it is loved by the Gods. [F]

He assents to, and dissents from, these claims as indicated. Now we have to say that there are two assumptions Socrates makes in this discussion. First, to put it anachronistically, to give a definition of a term is the same as to specify the nature, essence of what the term refers to. So Euthyphro's claim that to be pious is to be loved by the Gods should be read both as a definition and as a statement of what piety is. Second, it is supposed to be accepted by all that 'God-loved' and 'being loved by the Gods' are synonymous. They name the same fact, if you will, by using different grammatical constructions: passive participle and passive verb respectively.

With these assumptions, we get (A) ⇒ (C), and (B) ⇒ (D), which is a *reductio ad absurdum*. Hence 'piety' and 'God-loved' are not the same, and thus also 'piety' and 'being loved by the Gods' is not the same.

Explaining how. When we reflect on the grammatical features of the expression 'η explains how ξ', we see that the explanandum term occupies a sentence-position. Thus we have sensible constructions such as:

η explains how Trump won the election.
η explains how diesel engines work. (20-2)
η explains how cuckoos migrate south.

In all of these instances there is a sequence of events unfolding in time. We can distinguish its temporal stages succeeding each other from one moment in time to another. In The explains-how expression, however, is fairly selective: it disallows many declarative sentences. The constructions such as:

η explains how snow is white.
η explains how $E = mc^2$. (20-3)
η explains how one plus one is two.

are not sensible, bordering on nonsense, though grammatically well-formed. This is hardly surprising. We explain how a certain event happened, or how a mechanism works. None of the explananda in (20-3) is an event, or a description of a mechanism. There is a deeper difference. When we ask to explain how an event happens or how a mechanism works, we ask how a change was possible. We ask for a cause of the change, or for a mechanism governing the change.

Grounding explanations. A major motivation of grounding is to provide explanations: a fact X is grounded in a fact Y when Y explains X. As often emphasised, the kind of explanation in question is non-causal and metaphysical. That explanations can be non-casual is a commonplace in many, possibly all, accounts of explanation. The ground theorists, however, insist further that the grounding relation exhibits a distinct kind of metaphysical explanation. One idea I want to explore here is that the notion of metaphysical explanation is spurious. The ground theory has failed to articulate any new kind of explanation. Instances of ground, as conceived by the ground theorists, provide at most an illusion of explanation.

So far as the ground approach is piecemeal and so far as the intelligibility of ground is usually defended with examples, it is only too natural to follow in that the ground theorists and work with examples. A standard example involves Socrates and the singleton set {Socrates}. Intuitively, we are invited to agree, Socrates grounds {Socrates}, but not the other way around. We explain the existence of {Socrates} by citing the existence of Socrates. But we do not explain the existence of Socrates by citing the existence of {Socrates}.

Or again: the existence of the Everest is grounded in the existence of the parcels of land composing it. Why does the Everest exist? We are after a specification of a causal process culminating in the formation of the Everest. We are a specification of elementary facts that are *currently* responsible for the existence of the Everest. Thus we say that the Everest exists in virtue of the existence of

The fact P that Beijing is a city of 21 million people is said to be grounded in more fundamental facts $\langle p_i \rangle$ about the number of people residing in a particular geographic location, the facts of administrative and employment records. Why would anyone be attracted to this mode of thinking? As I see it, there are two ways towards endorsing the metaphysics of ground: the just-is way and the in-virtue-of way. The first way leads through ontology. What is it, we ask, for Beijing to be a city of 21 million? It is just to have a certain number of people in particular locations etc. The ground theory thus provides an answer to what seems to be a traditional ontological question. The question is about the constitution, structure, of the particular fact. The second way leads through explanation. We now ask, why does Beijing have 21

million people? This is so, the reply goes, in virtue of there being a certain number of people in particular locations etc. Here the ground theory deals with a no less traditional demand for explanation. The second way is indispensable. The project of fundamental ontology is attractive, primarily so, because of the links it claims to establish between ground and explanation. Other than an intuition about what there is, what makes a certain fact a suitable candidate for being a grounding fact is its explanatory role: P is explained by $\langle p_i \rangle$.

Hence Kit Fine writes:

> We take *ground* to be an explanatory relation: if the truth that P is grounded in other truths, then they *account* for its truth; P's being the case holds *in virtue of* the other truths' being the case. There are, of course, many other explanatory connections among truths. But the relation of ground is distinguished from them by being the tightest such connection. Thus when the truth of P causally explains the truth of Q, we may still maintain that the truth of Q consists in something more (or other) than the truth of P. Or again, the fact that someone broke a promise may "normatively" account for his having done something wrong, but that is still compatible with his wrongdoing's consisting in something more than his having broken the promise. There is, however, no explanatory connection that stands to ground as ground stands to these other forms of explanation. It is the ultimate form of explanation...

To specify the underlying ontology *is* to provide an explanation, because ground itself does the explanatory job. A ground explanation is a species of ontic explanation, where the explanans is a feature of the world, a cause or a fact, rather than a representation thereof, such as a theory or a statement. Explananda are in the same category of objective, subject-independent features of the world.

ONTIC EXPLANATIONS. The distinctive characteristic of the ontic conception of explanation is its claim about the explanans. A plausible way of examining this claim is by asking whether the sentence 'η explains why ξ' is well formed. If this construction is found to be not well formed, or not intelligible, then we are running into trouble. Testing this strategy with respect to explananda, we get the expected results: many entities are unsuitable to be explananda. Among them are properties, individuals, natural kinds:

'η explains why red' is not well formed.
'η explains why London' is not well formed. (20-4)
'η explains why water' is not well formed.

Looking through the list, we speculate that explananda terms cannot occur in either name- or predicate-positions. Should they occur in sentence-positions? This seems to be the case:

'η explains why snow is white' is well formed and sensible.
'η explains why $F = ma$' is well formed and sensible. (20-5)
'η explains why Trump won the election' is well formed and sensible.

Let us now turn to the original question: what kind of entity can the explanans be? Initially we may think that explanans and explananda belong in the same category. Thus we may think that the terms for the explanans occur in sentence-positions. But we will be disappointed:

'Snow is white explains why ξ' is not well formed.
'$F = ma$ explains why ξ' is not well formed. (20-6)
'Trump won the election explains why ξ' is not well formed.

The disappointment is predictable. Grammar dictates that the verb 'explains' be preceded by a noun-phrase.

Perhaps there is another way to preserve the symmetry between explananda and explanans. If explananda were in the sentence-position, explanans should be represented by names of sentences. This move fails too. It is, I think, no longer clear that the resulting constructions are ill-formed. But the least we can say is that they are not sensible:

'Snow is white' explains why ξ.
'$F = ma$' explains why ξ. (20-7)
'Trump won the election' explains why ξ.

Once again, the failure is unsurprising. Sentences *themselves* cannot do the explaining. It is strange to think that a fragment of a language is explanatorily apt by itself.

This suggests that the terms for explanans should properly occur in the name-position. But not every name is a suitable candidate. We quickly verify that properties, individuals, and natural kinds cannot be the explanans:

'Red explains why ξ' is not sensible.
'London explains why ξ' is not sensible. (20-8)
'Water explains why ξ' is not sensible.

All of these sentences can perhaps be used as abbreviations of a proper explanatory locution. In isolation, they border on nonsense. What at first seems more natural is to use names of theories. Thus consider:

General Theory of Relativity explains Mercury's precession. (20-9)

This statement appears sensible. Suppose, though, we think of theories as collections of sentences. Then again it seems strange how sentences on their own can engage in explaining.

These observations direct us towards an obvious remedy. Sentences, cities, or colour properties do not explain. But people do. Modifying (20-9), we get:

> Eddington explained why Mercury's precession occurs. (20-10)

The statement (20-10) is a truncated version of a longer sentence. Eddington perhaps did explain something, but he did not explain anything *on his own*. Nor did he explain in solitude. The full version might run as follows:

> Eddington explained *to* the scientific community why Mercury's precession occurs *by* using GTR. (20-11)

Explanation is a pragmatic act. It puts a speaker in a certain relation with the audience. The speaker explains by conveying informational content to the audience. Various representational products—descriptions, thoughts, theories—are used in that content. When we say that this content itself explains, this is elliptic for saying that the content is used in the act of explanation.

One might think that this line of argument is wrongheaded. People explain by discovering the explanatory features of the world. Eddington could explain precession because GTR captured the causes of Mercury's motion. Properly speaking, causes explain. The talk of people explaining piggybacks on the explanatory features of the world. Theoretical constructs explain too, but only when they identify real laws. In turn, laws explain, but only when they are conceived as, for example, relations between universals, and not as *our* conceptual or linguistic representations of them.

This is not a good objection. It is strange to think that explaining is done in absence of any conscious agent. The extinction of dinosaurs was caused by a meteorite crash. It is odd to say that the crash, all those millions of years ago, had the property of explaining the extinction. Objects, facts, reality and its chunks, do not explain. We cannot even say that the causal connections *themselves* are used in explaining the extinction. Instead we should say that representations of those connections are so used. Hence Kitcher on scientific explanation:

> The most obvious way in which to categorize explanation is to view it as an activity. In this activity we answer the actual or anticipated questions of an actual or anticipated audience. We do so by presenting reasons. We draw on the beliefs we hold, frequently using or adapting arguments furnished to us by the sciences.

Liberal metaphysical explanation. At this point the ground theorist might retreat—but only a little. He might concede that a grounding relation is not itself an explanatory relation. But he might insist that grounds are the material for the most satisfactory, ultimate explanation. Fine himself sometimes leans towards this conception:

> [T]he relationship of ground is a form of explanation; in providing the ground for a given proposition, one is explaining, in the most metaphysically satisfying manner, what it is that makes it true.

On a charitable reading, Fine is non-committal on the nature of the explanatory relation. He only claims that the best possible explanation of P will have to cite the grounds of P. This claim in turn rests on a general claim about the conditions for a successful explanation—that any explanation is good, real, or successful, only when the explanandum (the fact to be explained) is grounded in the explanans (the facts doing the explanation).

This more liberal stance on metaphysical explanation faces a number of problems. In the first place, as we learned from van Fraassen and others, explanations using a false theory are not pointless. They increase our understanding of the world. Newtonian mechanics explains—and we can admit that all the while denying that it is true. In fact we can argue that explanations better grounded in reality—that is, explanations using more detailed and sophisticated theories we believe are true—would have less explanatory power. Their detailed descriptions will interfere with our understanding of the world.

Comment 1. Another comment is that grounds have to be described in our language, with our concepts, with all the limitations that follow. If this is so, then I see no reason to think that a possibility of a final explanation is so much as intelligible, and the talk about it is helpful:

> [T]o ask what reality is *really* like, however, apart from human categories, is self-stultifying. It is like asking how long the Nile really is, apart from parochial matters of miles or meters. Positivists were right in branding such metaphysics as meaningless. (Quine)

Comment 2. As the ground theorist might think about it, the maximally successful explanation is the maximally specific one. As such, it will be able to specify what grounds are for the given fact. But in providing explanations, it is essential to leave some things out. Theoretical reasoning, in furnishing explanations, sorts available evidence into relevant and irrelevant. The very activity of creating theories involves leaving things out. For example: pendulum bobs swing with the period proportional to the square root of the length of the string it is hanging from.

So we could say: the fact that a pendulum bob has the period $\theta \propto l$ is grounded in the fact of its constitution plus a myriad of other facts about gravity. Among the facts of the bob's constitution is the fact of it having a certain colour—or, if you like, the fact of it being covered by a material emitting light of a certain wavelength. Initially this seems right. Every bob is coloured. So there is no fact [A bob is swinging]. There is only a fact [A coloured bob is swinging]. So the fact that this black bob is swinging with the period θ is grounded in the fact that it is black plus many other facts.

But now the problem is that the colour is irrelevant and contributes nothing to the explanation. Indeed, much of its material constitution contributes nothing the explanation why the bob is swinging with the period θ. I think this example may show that the ontological sense of ground, the sense of fundamentality, is parting ways with the its hermeneutic sense, in which it is charged with explanation. Every theory, in order to provide explanations, must of necessity omit certain facts. The maximally successful explanation will not contain the maximally detailed description of the world.

The theoretical basis of grounding. Putting this concern aside, therefore, let us ask this question: Why to believe, in each putative example of grounding, that an explanatory link exists between the ground and the fact grounded? If explanation improves our understanding, then we must believe that, having been presented with a grounding relation, we understand the grounded fact—understand it better than before. So the question is, what feature of the grounding relation is responsible for this improvement of understanding.

I believe the reason is in our possession of a theory T with certain characteristics. First, such a theory provides a set of instructions: anyone wishing to establish conclusively the fact P should begin by establishing the facts $\langle p_i \rangle$. I can of course derive the information about the population of Beijing from a secondary source (Wikipedia perhaps), and in general a chain of secondary sources. But the primary source, if trustworthy at all, should be able to establish the grounding facts $\langle p_i \rangle$. The reverse does not hold. Knowing the number of Beijing residents does not help in learning the facts of various administrative records of the city.

Secondly, the procedure, according to the theory T, is general. To establish the fact Q that Tokyo's population is 13 million, one has to go through a similar kind of facts $\langle q_i \rangle$ about the residence and other administrative records of Tokyo. Instead of having disconnected, mutually irrelevant facts P and Q, we see how they are related to each other through the facts $\langle p_i \rangle$ and $\langle q_i \rangle$. The similarity of these lower-level facts is not a mere article of faith. Ultimately it'll rest on the similarity of experiences: anyone having to establish the facts of administrative records will have to go through qualitatively similar experiences.

Thirdly, the theory so construed should deliver successful predictions. Here I am taking on board Hempel's dictum that explanations are 'potentially predictive'. But what are we able to predict once made aware of the grounding relation between two facts? It would be silly, though not wrong, to say that, upon learning the facts about administrative records of a city at a certain time, we are able to predict the number of its inhabitants at that time. Nor does the knowledge of these facts allow us predicting the number of Beijing inhabitants at a future time. More importantly, I think, in our case these predictions delivered by the theory concern not the facts themselves, but the use of concepts. The theory does not allow us to predict the population of Tokyo or Beijing at a future time. But it allows us to predict what, on a future occasion of use, will count as 'population of Tokyo' or 'population of Beijing'. Thus suppose an alternative proposal is raised, according to which the same fact P of the population numbers in Beijing is grounded in the facts of physical residence, however temporary. Administrative records are ignored. Then, according to this proposal, every visiting tourist can change the number of Beijing inhabitants. A defender of T plausibly protests that the population numbers *cannot* depend on the visiting tourists. But the protest merely exploits an alternative conceptual use.

Fourthly, and most significantly, the number of brute facts is reduced. This model of explanation is of a piece with scientific explanation, in particular with the unification model. Ground explains by opening a window into another area of discourse. The administrative records of the city, combined with certain facts about residence, explain the fact of the city's population, because they allow us to establish connections with further facts, say, of literacy. Suppose we read in a chronicle that in the year 1600 Beijing population was five million, but that one year later it dropped to four million. If, improbably, we were to take these two facts as brute, ungrounded, then we would have left it at that. If, on the other hand, we think they *are* grounded in the way described, then we would enquire whether perhaps the records were more adequately kept in 1600 than one year later. If they were, we would further ask whether this was due to greater literacy. A whole new line of enquiry is initiated. And if no difference in literacy is discovered, war or epidemic seem likely causes affecting facts of physical residence. Therefore, the facts of city population may be said to be explained by other facts if these facts integrate them into further descriptions of the world that previously appeared irrelevant.

So my claim is that, contrary to frequent announcements, the ground theory does not advance any new kind of explanation. The grounding relation itself does not explain. To be told merely that the ground is more fundamental than the grounded fact does not improve our understanding. At best identifying the grounding relation is a preliminary step towards explanation. The explanatory job itself is done by the greater simplicity and unity in our description of the world. This is a central plank of the unification approach to explanation—but also one that it shares with the Deductive-Nomological model and with many other approaches to explanation. Objects, facts, reality and its chunks, do not explain. Explanations are performed by our own representations.

YSB