

Article

Metaphysical Explanation: An Empirical Investigation

Andrew J. Latham ¹ and Kristie Miller ^{2,*}¹ Aarhus Institute of Advanced Studies, Aarhus University, 8000 Aarhus, Denmark; andrew.latham@cas.au.dk² Department of Philosophy, University of Sydney, Sydney, NSW 2050, Australia

* Correspondence: kristie.miller@sydney.edu.au

Abstract: The literature on metaphysical explanation contains three widely accepted assumptions. First, that the notion of metaphysical explanation with which philosophers are interested is a notion with which the folk are familiar: it is at least continuous with the folk notion. Second, that metaphysical explanations are propositions of a certain form that are true (or false), *simpliciter*. Third, that it is at least the case that mostly, if x metaphysically explains y, then y does not metaphysically explain x. On the basis of the empirical investigations that we pursued, we argue that at least two of these assumptions are false.

Keywords: asymmetry; experimental philosophy; fundamentality; grounding; metaphysical explanation

1. Introduction

In what follows, we suppose that metaphysical explanations are true propositions of the form $\ulcorner x \text{ because } y \urcorner^1$, where $\ulcorner x \urcorner$ and $\ulcorner y \urcorner$ are sentences, and ‘because’ expresses a particular kind of explanatory connection (i.e., one that is metaphysical rather than causal)².

One assumption we take to be common amongst philosophers is that the notion of metaphysical explanation with which philosophers are interested is deeply connected with an ordinary folk notion of explanation. Many philosophers suppose that metaphysical explanations have been a staple of human interaction for a long time, and are not a unique feature of philosophical discourse. Schaffer [3] claims that metaphysical explanation is “a natural and intuitive notion, for which there exist clear examples, and clear formal constraints” and Dasgupta [4] notes that metaphysical explanation is “intuitive and familiar. . . examples are ubiquitous. Why is a faculty meeting occurring? Because the faculty are gathered in a room discussing matters of importance to the department, etc. Why is this water hot? Because its mean kinetic energy is high. Why have I lost this game of chess? Because my king is in check-mate.” Indeed, it is “an everyday concept used by the masses. When I explain the concept to non-philosophers they recognize it immediately and talk intelligibly about it, offering examples of [metaphysical] explanations in their own fields of biology, economics, journalism, or cooking. To them it is not a new concept.” Further, Glazier [5] notes that “[f]or its enthusiasts, [metaphysical] explanation is both ubiquitous in ordinary life and central to many of philosophy’s biggest questions”. Let us call this *the assumption of folk–philosophical continuity*. This is the assumption that there is, at the very least, continuity between the folk and philosophical notions of metaphysical explanation. Hence, if there is a single notion at play then this assumption is true. If there are two notions that are closely connected, where, for instance, the philosophical notion is a slightly ‘tidied up’ version of the folk notion, then the assumption is true. If there are two quite different notions at play, then the assumption is false.

A second assumption we take to be even more common is that propositions of the relevant form are true, or false, *simpliciter*. It is not the case, for instance, that a proposition of that form is true for one party, and false for another if, say, the former gains some sort of epistemic status by knowing that proposition while the latter does not. More generally, it is not the case that propositions of that form are true relative to some standards and



Citation: Latham, A.J.; Miller, K. Metaphysical Explanation: An Empirical Investigation. *Philosophies* **2024**, *9*, 85. <https://doi.org/10.3390/philosophies9030085>

Academic Editor: Frederick Kroon

Received: 14 March 2024

Revised: 30 May 2024

Accepted: 4 June 2024

Published: 11 June 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

not others, or relative to some contexts of assessment and not others, or relative to some communities of knowers and not others. Rather, they are simply true, or false³. Let us call this *the assumption of truth simpliciter*.

The third common assumption is that mostly, if x metaphysically explains y, then y does not metaphysically explain x. Often, philosophers have endorsed the even stronger claim that metaphysical explanation is asymmetric⁴. Technically, properties like asymmetry are properties of relations, not propositions. Still, we will say that metaphysical explanation is asymmetric just in case for any proposition of the form 'x because y', p because q, if p because q is true, then q because p is not. We will say that a particular token metaphysical explanation, p because q, is asymmetric just in case p because q is true, and q because p is false, and is symmetric just in case both p because q and q because p are true. Then, metaphysical explanation is non-symmetric just in case some token propositions of the form 'x because y' are symmetric, and some are asymmetric.

The idea that metaphysical explanation is asymmetric is in part motivated by reflection on our judgements about a range of purported cases of metaphysical explanation. In all these cases, it is argued, philosophers tend to judge that there is a metaphysical explanation of x because y, but not the reverse. In turn, that is thought to give us reason to think that metaphysical explanation is asymmetric. Part of that same idea appeals to the idea that we tend to judge that circular explanations are not explanations. Recently, both these motivations have come under scrutiny.

Some have challenged the idea that there is something problematic about circles of explanation⁵. If a explains b, and b explains c and c explains d, and d explains a, there need not, it is argued, be anything vicious, or otherwise concerning, about this circle of explanation. Still others have provided examples of what are argued to be particular symmetrical metaphysical explanations⁶. For instance, Rodriguez-Pereya [21] argues that truthmaking (which he takes to be a species of metaphysical explanation) is not asymmetric. Consider proposition A: <A exists>. That proposition says of itself that it exists. If true, then, it is made true by the existence of that very proposition. So this proposition is made true by itself. Hence, there are reflexive metaphysical explanations and, it follows, there are symmetric metaphysical explanations since asymmetry entails irreflexivity. A related example comes from Thompson's [11] truth-telling pair. P says that 'Q' is true, while Q says that 'P' is true. This, too, appears to be a case in which P explains Q, and Q explains P.

These examples of symmetry are *recherché* to say the least. So even if they show that metaphysical explanation is non-symmetric, they suggest (or at least are consistent with) a claim that almost all philosophers accept: that most token metaphysical explanations are asymmetric: mostly, if p because q is true, then q because p is not. Call this *the assumption of widespread asymmetry*.

Of course, it is not the case that all philosophers working in this area explicitly sign up to all three of these assumptions, though we think that many, and perhaps most, accept all three. At the very least, it is reasonable to think, extrapolating from the literature, that this is a common trio of assumptions.

If all three of these widely held assumptions can be true, then the resulting picture is one on which *non-philosophers* will typically judge that if p metaphysically explains q, then it is not the case that q metaphysically explains p (assumption of widespread asymmetry). That is because if the assumption of folk-philosophical continuity is correct, then non-philosophers will, by and large, make the same judgements about metaphysical explanation as will philosophers. And since philosophers judge that metaphysical explanation is typically asymmetric (the assumption of widespread asymmetry) it should be that non-philosophers also make this judgement. We aim to empirically test if this is so.

At this point, it is worth making something clear. It might be thought that philosophers and non-philosophers have some sort of tacit concept of metaphysical explanation, and that what it means to say that the assumption of folk-philosophical continuity is correct, is to say that they have the same, or similar, such concept. We take it that this is the right way to conceive of this claim. We will take it that this commits those who defend the assumption

of folk–philosophical continuity to the claim that philosophers and non-philosophers will *by and large* make the same judgements about cases. In particular, we take it that it commits those who also defend the widespread asymmetry assumption, to the claim that philosophers and non-philosophers alike, will tend to judge that if *p* metaphysically explains *q*, then *q* does not metaphysically explain *p*. We make this additional assumption because, as noted previously, when philosophers argue that metaphysical explanation is asymmetric, they do so by reflecting on cases: they note that they are inclined to judge that if *p* metaphysically explains *q*, then the reverse is not the case. They *then* infer, on that basis, that metaphysical explanation is asymmetric. So, we take it to be plausible that people’s tacit concept of metaphysical explanation is *displayed* in their judgements about what metaphysically explains what, something that our interlocutors clearly take to be true, because they use their own judgements to inform their view of metaphysical explanation. That is, we think that it is fair to assume that our interlocutors, who hold both the assumption of folk–philosophical continuity and the widespread asymmetry assumption, are thereby committed to the idea that philosophers and non-philosophers will both tend to judge that if *p* metaphysically explains *q*, then *q* does not metaphysically explain *p*.

Likewise, we take it that those who defend both the assumption of folk–philosophical continuity and the assumption of truth *simpliciter* are committed to the idea that both philosophers and non-philosophers will judge that claims of the form *p* because *q*, are true, or false, *simpliciter*: that is, they will not judge that they are true at some contexts, and false at others.

It is these claims about non-philosophers’ judgements that we aim to test in this paper. To do so, we probe the judgements of non-philosophers about three cases that are widely discussed in the philosophical literature. In all three cases, the philosophical intuition is that the token explanation is asymmetric. So in each case, philosophers judge either *p* because *q* is true, *simpliciter*, and that *q* because *p* is false *simpliciter*, or the other way around. If the folk notion of metaphysical explanation is continuous with that of the philosophical notion, then we would expect to find that non-philosophers’ judgements, by and large, mirror those of philosophers in this regard.

We begin, in Section 2, by outlining relevant empirical research and developing our hypotheses. We then describe this study’s methodology in Section 3 and then the results in Section 4. Our results show that non-philosophers judge all three cases to be ones in which the metaphysical explanations in question are symmetric. We consider the upshot of these results in Section 5.

2. Empirical Evidence and Hypotheses

Suppose we accept the assumptions of widespread asymmetry, of folk–philosophical continuity, and of truth *simpliciter*. Call this *the assumption trio*. In what follows, we will outline a series of predictions that issue from accepting the assumption trio.

Given the truth of the assumption trio, we would expect that in non-recherche cases non-philosophers (a) will judge that either ‘*x* because *y*’ is true, or that ‘*y* because *x*’ is true, but not both and (b) will agree with philosophers’ judgements about which of these is true and (c) will have judgements about which is true that are insensitive to features of context. (c) follows from the assumption of truth *simpliciter* in conjunction with the assumption of folk–philosophical continuity; (a) and (b) follow from the assumption of widespread asymmetry in conjunction with the assumption of folk–philosophical continuity.

It is important to note, however, that the assumption of truth *simpliciter* is consistent with there being acts of metaphysical explanation—speech acts in which one party utters a proposition of that form—such that a speech act of that form only counts as an act of explanation if the proposition uttered is true, and, for instance, its utterance performs some epistemic role for the recipient. We might expect judgements about whether a speech act is an act of explanation to be sensitive to features of context.

That is why in the studies that we report, we ask participants to make two separate judgements. One is a judgement about the truth of a proposition of the form ‘ x because y ’. The second is a judgement about whether when one party uttered a proposition of that form, what they said *was an explanation*, for some specified party. We take the latter to probe judgements about whether the speech act is an explanation for the party in question, and the former to probe judgements about whether the proposition is true (i.e., is a metaphysical explanation) or not.

What should we predict about non-philosophers’ judgements about acts of explanation? There is little empirical work on non-philosophers’ judgements about metaphysical explanation⁷. There is, however, a good deal of empirical work on explanation more generally. That research suggests that non-philosophers’ judgements about whether there is an act of explanation present are sensitive to the epistemic goals of explanation recipients [23], whether the recipient comes to understand the phenomenon [24,25], the recipient’s capacity to successfully intervene in the world on the basis of the explanation [23] and the salience of various causal factors [26] including whether or not they are factors over which we have control [27] and whether they cohere with existing explanations [28,29].

In the studies that follow, we focus on just one such factor: the presence of a salient intervention.

The idea that there is some connection between interventions and metaphysical explanations is not new. Miller & Norton [30] make this case, while Schaffer [2] and Wilson [1,31] have independently developed accounts of grounding that make use of the same formalism as interventionist accounts of causation.

Research on causal explanation shows that there is an intimate relationship between causal explanations, the imagining of interventions, and the representation of interventional affordances (that is, representations that intervening on one part of the world will thereby intervene on some other part of the world). There is ample evidence that imagining interventions is vital in causal reasoning⁸ and in causal explanation⁹. People tend to judge that $[x]$ is a causal explanation for $[y]$ only when they judge that intervening on $[x]$ is a way to intervene on $[y]$, but that intervening on $[y]$ is not a way to intervene on $[x]$ [34].

This suggests that if we describe a scenario in which intervening on $[y]$ in order to intervene on $[x]$ is salient, and another in which intervening on $[x]$ in order to intervene on $[y]$ is salient, we might find that this has an effect on non-philosophers’ judgements about whether an utterance of ‘ x because y ’ is an act of explanation. Let us call which intervention is salient in a scenario the *salience condition*.

Suppose that in a scenario, intervening on $[y]$ in order to intervene on $[x]$ is salient. If that scenario is one in which a proposition of the form ‘ x because y ’ is uttered, (and then evaluated by participants) we will say that it is a *congruent condition*. A congruent condition is one in which the direction of the salient intervention aligns with the direction of explanation expressed by the proposition uttered in the scenario. By contrast an incongruent condition is one in which the direction of the salient intervention does not align with the direction of explanation expressed by the proposition uttered in the scenario.

What should we predict about non-philosophers’ judgements about acts of explanation given the truth of the assumption trio? We should predict that in non-recherché cases non-philosophers (a) will judge that at most, an utterance of ‘ x because y ’ is an explanation, or an utterance of ‘ y because x ’ is an explanation, but not both and (b) judgements about whether an utterance of ‘ x because y ’ or ‘ y because x ’ is an act of explanation will be sensitive to which intervention is salient, and to whether or not the condition is congruent. (a) follows from the assumption of widespread asymmetry in conjunction with the assumption of folk-philosophical continuity, alongside the assumption that an utterance is an act of explanation only if the proposition expressed is true. If, in non-recherché cases, only one of ‘ x because y ’ and ‘ y because x ’ can be true, (as demanded by the widespread asymmetry assumption) then it follows that at most, an utterance of one of them can be an act of explanation. (b) is not entailed by the assumption trio. But given what we know about non-philosophers’ judgements about causal explanation it is plausible to suppose that if ‘ x because y ’ is true,

then non-philosophers' will be more inclined to judge that an utterance of that proposition is an act of explanation in a context in which intervening on [y] in order to intervene on [x] is made salient (i.e., a congruent condition) than in an incongruent condition in which intervening on [x] in order to intervene on [y] is made salient.

In what follows, we focus on three cases of metaphysical explanation. Since familiarity with certain kinds of explanations (or with explanans and explanandum) might play a role in non-philosophers' judgments, we chose cases that vary along this dimension.

The first, mind/brain case, is one that is relatively familiar to non-philosophers. Philosophers systematically judge that the properties of brain states explain the properties of mental states, but not vice versa. In the vignette we present, the brain state is having high levels of cortisol, and the mental state is being stressed. Given what we have said so far, then, if the assumption trio is true we predict that non-philosophers:

- (a) Will judge that 'stressed *because* cortisol' ¹⁰ is true, and that 'cortisol *because* stressed' is not true and
- (b) Will have judgements about the truth of those propositions that are insensitive to which intervention is made salient, and to whether the condition is congruent.

Further, we predict that non-philosophers:

- (c) Will judge that an utterance of 'cortisol *because* stressed' is not an explanation and
- (d) Will be more likely to judge that an utterance of 'stressed *because* cortisol' is an explanation in a condition in which intervening on cortisol to intervene on stress is salient, than in the incongruent condition in which intervening on stress in order to intervene on cortisol is salient.

The second case is the set/singleton case. This was chosen to be a case that is unfamiliar to non-philosophers. Again, it is a case in which philosophers have asymmetric judgements: judging that the existence of the member of the singleton set explains the existence of the set.

Given what we have said so far, then, if the assumption trio is true, then we predict that non-philosophers:

- (a) Will judge that 'set *because* member' is true, and that 'member *because* set' is not true and
- (b) Will have judgements about the truth of those propositions that are insensitive to which intervention is made salient, and to whether the condition is congruent.

Further, we predict that non-philosophers:

- (c) Will judge that an utterance of 'member *because* set' is not an explanation and
- (d) Will be more likely to judge that an utterance of 'set *because* member' is an explanation in a condition in which intervening on the member to intervene on the set is salient, than in the incongruent condition in which intervening on the set in order to intervene on the member is salient.

Finally, our third case is the Euthyphro case, in which even philosophers do not always agree about which direction explanation proceeds, although, notably, each philosopher tends to think it proceeds in only one direction. Given this disagreement amongst philosophers we predicted that:

- (a) Some non-philosophers will judge that 'God approves *because* good' is true, and others will judge that that 'good *because* God' approves is true.
- (b) Non-philosophers will have judgements about the truth of those propositions that are insensitive to which intervention is made salient, and to whether the condition is congruent.

(a) seems like a reasonable prediction to have, given the data at hand. (b) follows from the conjunction of the folk-philosophical continuity assumption and the truth *simpliciter* assumption.

Further, we predict that:

- (c) Some non-philosophers will judge that an utterance of 'good *because* God approves' is an explanation and some will judge that 'God approves *because* good' is an explanation.

- (d) Non-philosophers will be more likely to judge that an utterance of ‘good *because* God approves’ is an explanation in a condition in which intervening on God in order to intervene on good is made salient, than they are in the incongruent condition in which intervening on good to intervene on God is made salient, and will be more likely to judge that an utterance of ‘God approves *because* good’ is an explanation in a condition in which intervening on good in order to intervene on God is made salient, than in the incongruent condition in which intervening on God in order to intervene on good is made salient.

3. Experimental Design

3.1. Method

3.1.1. Participants

A total of 1279 people participated in this study across three experiments. Participants were U.S. residents, recruited and tested online using Amazon Mechanical Turk,¹¹ and were compensated for their time. A total of 412 participants had to be excluded for failing to follow task instructions. This means that they failed to answer the questions (210) or failed an attentional check question (202). The remaining sample was composed of 867 participants (aged 18–99; 339 female, 6 transgender/non-binary). Mean age 36.23 ($SD = 12.07$). Ethics approval for this study was obtained from the University of Sydney Human Research Ethics Committee. Informed consent was obtained from all participants prior to testing. The survey was conducted online using Qualtrics. There were no significant differences between experiments for age $F(2, 864) = 1.012, p = 0.364$, or gender $\chi^2(4, N = 867) = 2.741, p = 0.602$.

3.1.2. Materials and Procedure

Experiment 1: Mind/Brain

Participants were divided into four groups, each of which saw a single vignette and responded to one set of statements.

Vignette 1(a) Brain/Mind: Congruent

Fred is experiencing a lot of stress. He learns that whenever someone is stressed, they have very high cortisol levels, and that whenever anyone has very high cortisol levels, they are stressed. He learns that stress and high cortisol levels go together because what it is to have one, is to have the other.

Fred goes to his pharmacist, Maria. Maria tells Fred that a new drug has just been released—*depressor*—and that this drug works by preventing cortisol from being released into the system. Maria tells Fred ‘Fred, you are stressed because you have high cortisol levels’ and she recommends that Fred takes depressor.

Vignette 1(b) Brain/Mind: Incongruent

Fred is experiencing a lot of stress. He learns that whenever someone is stressed, they have very high cortisol levels, and that whenever anyone has very high cortisol levels, they are stressed. He learns that stress and high cortisol levels go together because what it is to have one, is to have the other.

Fred goes to his pharmacist, Maria. Maria tells Fred that a new drug has just been released—*depressor*—and that this drug works by preventing cortisol from being released into the system. Maria tells Fred ‘Fred, you have high cortisol levels because you are stressed’ and she recommends that Fred takes depressor.

Vignette 1(c) Mind/Brain: Congruent

Fred is experiencing a lot of stress. He learns that whenever someone is stressed, they have very high cortisol levels, and that whenever anyone has very high cortisol levels, they are stressed. He learns that stress and high cortisol levels go together because what it is to have one, is to have the other.

Fred goes to his therapist, Maria. Maria tells Fred that a new therapy technique has been developed—*unstressor*—and that this therapy, consisting of meditation and mindfulness, works to alleviate the experience of stress. The therapist tells Fred ‘Fred, you have high cortisol levels because you are stressed’ and she recommends that Fred undergo the *unstressor* therapy.

Vignette 1(d) Mind/Brain: Incongruent

Fred is experiencing a lot of stress. He learns that whenever someone is stressed, they have very high cortisol levels, and that whenever anyone has very high cortisol levels, they are stressed. He learns that stress and high cortisol levels go together because what it is to have one, is to have the other.

Fred goes to his therapist, Maria. Maria tells Fred that a new therapy technique has been developed—*unstressor*—and that this therapy, consisting of meditation and mindfulness, works to alleviate the experience of stress. The therapist tells Fred ‘Fred, you are stressed because you have high cortisol levels’ and she recommends that Fred undergo the *unstressor* therapy.

Table 1 below provides a summary of the experimental design of experiment 1, showing each of the four possible combinations of metaphysical explanation and intervention target, and whether that intervention target is congruent or incongruent with the metaphysical explanation provided in the vignette.

Table 1. Experimental design of experiment 1.

Condition	Metaphysical Explanation	Intervention Target
1(a)	Stress <i>because</i> high cortisol	Congruent (high cortisol)
1(b)	High cortisol <i>because</i> stress	Incongruent (high cortisol)
1(c)	High cortisol <i>because</i> stress	Congruent (stress)
1(d)	Stress <i>because</i> high cortisol	Incongruent (stress)

After seeing one of these vignettes all participants were asked to respond to Maria’s assertion on two different Likert scales. One of the Likert scales ran from 1 ‘Completely sure that what Maria says is false’ at one end (either the far left or the far right, determined randomly: this is so for all the Likert scales henceforth described) to 7 ‘Completely sure that what Maria says is true’ at the opposite end of the scale via 4 ‘I am indifferent between these two options’. The other Likert scale ran from 1 ‘Completely sure that the statement is *not* an explanation for Maria’ at one end to 7 ‘Completely sure that the statement *is* an explanation for Maria’ via 4 ‘I am indifferent between these two options’.

After having done so, participants were taken to a new page that did not have either the vignette or Likert scales on it and were asked an attentional check question: “*In the vignette you were asked to read, what were Fred and Maria talking about?*” to which they could answer (1) abstract objects; (2) bicycles; (3) god or (4) minds and brains. Participants who choose incorrectly were eliminated.

Experiment 2: Sets/Members

This experiment had the same methodology as experiment 1, but with different vignettes.

Vignette 2(a/b) Member/Sets: Congruent/Incongruent

Sets are abstract objects, like numbers, or functions (like addition, and subtraction, and division). But unlike number or functions, sets have members. Even though sets are abstract objects, their members can be concrete things, like dogs, chairs, electrons, and people. In fact, for any bunch of things, there is a set containing just those things. *Singleton* sets are sets that contain only one member. So, the singleton set containing the number 2, is

the set that only contains 2. The singleton set containing Eminem, is the set that contains only Eminem as a member. Sets only exist when their members do, and their members only exist when the set does. So, if Eminem exists, then the singleton set containing Eminem exists. If the singleton set containing Eminem exists, then Eminem exists. If one exists, then so does the other.

Fred is a member of an elite secret time team. This team of people can travel anywhere in time, and can erase people and events from the timeline, so that those people and events never existed at all. Fred is talking with his supervisor, Maria, about the suggestion that they ought to erase Eminem from the timeline. Fred says to Maria ‘but what about the singleton set containing Eminem?’ Maria responds that if they erase Eminem, then they will erase the singleton set as well. At the end of their discussion, Maria tells Fred that [the singleton set containing Eminem exists because Eminem exists]/[Eminem exists because the singleton set containing Eminem exists].

Vignette 2(c/d) Sets/Members: Congruent/Incongruent

Sets are abstract objects, like numbers, or functions (like addition, and subtraction, and division). But unlike number or functions, sets have members. Even though sets are abstract objects, their members can be concrete things, like dogs, chairs, electrons, and people. In fact, for any bunch of things, there is a set containing just those things. *Singleton* sets are sets that contain only one member. So the singleton set containing the number 2, is the set that only contains 2. The singleton set containing Eminem, is the set that contains only Eminem as a member. Sets only exist when their members do, and their members only exist when the set does. So if Eminem exists, then the singleton set containing Eminem exists. If the singleton set containing Eminem exists, then Eminem exists. If one exists, then so does the other.

Fred is a member of an elite secret abstract-objects team. This team of people can travel anywhere in the abstract realm and can intervene on any abstract object. The team of people can, for instance, intervene on numbers, and functions and sets. Fred is talking with his supervisor, Maria, about the suggestion that they ought to erase the singleton set containing Eminem, from the abstract realm. Fred says to Maria ‘but what about Eminem?’ Maria responds that if they erase the singleton set containing Eminem, from the abstract realm, then they will erase Eminem as well. At the end of their discussion, Maria tells Fred that [Eminem exists because the singleton set containing Eminem exists]/[the singleton set containing Eminem exists because Eminem exists].

Table 2 below summarises the experimental design of experiment 2. It is organised in the same manner as Table 1.

Table 2. Experimental design of experiment 2.

Condition	Metaphysical Explanation	Intervention Target
2(a)	Set <i>because</i> member	Congruent (member)
2(b)	Member <i>because</i> set	Incongruent (member)
2(c)	Member <i>because</i> set	Congruent (set)
2(d)	Set <i>because</i> member	Incongruent (set)

Participants then responded to the same statements as those in experiment 1.

Experiment 3

This experiment had the same methodology as experiments 1 and 2, but with different vignettes.

Vignette 3(a/b) God/good: Congruent/Incongruent

Fred is not sure whether or not God approves of eating snails. Fred finds snails really tasty, and hopes that God does approve of eating snails, since, Fred knows, God approves

of things only when they are good, and everything that is good, is something God approves of. One day Fred meets up with his old friend, Maria, and the two of them visit Maria’s prayer group: the Change God’s Mind prayer group.

The group aims to pray to God, to change His mind about what he approves of. Maria suggests that Fred prays, with the group, to bring it about that God approves of eating snails. In fact, God exists, and He does approve of eating snails. After the prayer meeting, Maria tells Fred that [eating snails is good because God approves of it]/God approves of eating snails because it is good].

Vignette 3(c/d) good/God: Congruent/Incongruent

Fred is not sure whether or not God approves of eating snails. Fred finds snails really tasty, and hopes that God does approve of eating snails, since, Fred knows, God approves of things only when they are good, and everything that is good, is something God approves of. One day Fred meets up with his old friend, Maria, and the two of them visit Maria’s prayer group: the Change What is Good prayer group.

The group aims to pray to God, to get God to change which things are good. Maria suggests that Fred prays, with the group, to bring it about that eating snails is good. In fact, eating snails is good, and God exists, and He does approve of eating snails. After the prayer meeting, Maria tells Fred that [God approves of eating snails because eating snails is good]/[eating snails is good because God approves of it.]

Table 3 below summarises the experimental design of experiment 3. It is organised in the same manner as Tables 1 and 2.

Table 3. Experimental design of experiment 3.

Condition	Metaphysical Explanation	Intervention Target
3(a)	Good <i>because</i> God approves	Congruent (God)
3(b)	God approves <i>because</i> good	Incongruent (God)
3(c)	God approves <i>because</i> good	Congruent (good)
3(d)	Good <i>because</i> God approves	Incongruent (good)

3.1.3. Analyses

Call participants’ levels of agreement regarding whether Maria’s assertion is true, their *levels of truth agreement*. Then, participants whose level of truth agreement is *higher* are participants who more strongly agree that Maria’s assertion is true. Call participants’ level of agreement regarding whether Maria’s assertion is an explanation for Maria, their *levels of explanation agreement*. Then, participants whose level of explanation agreement is higher are participants who more strongly agree that Maria’s assertion is an explanation for her.

For each experiment, we compared levels of truth agreement and levels of explanation agreement between conditions using separate 2 (salience) × 2 (congruence) between-subjects ANOVAs. We also tested whether people’s levels of truth agreement and levels of explanation agreement differs significantly from indifference (a score of 4 on the Likert scale) by running separate one-sample *t*-tests for each condition.

4. Results

4.1. Experiment One

Before reporting the statistics let us summarise our main findings. We predicted that if the assumption trio is true, then non-philosophers would judge that ‘stressed *because* cortisol’ is true, and that ‘cortisol *because* stressed’ is not true. We did not find this. People, overall, judged that it is true that ‘stressed *because* cortisol’ and also true that ‘cortisol *because* stressed’. We also predicted that non-philosophers would have judgements about the truth of those propositions that are insensitive to which intervention is made salient, and to whether the condition is congruent. We did not find this. There was a significant interaction effect between salience and congruence on peoples’ levels of truth agreement.

We also predicted that non-philosophers would judge that an utterance of ‘cortisol because stressed’ is not an explanation. We did not find this. People, overall, judged that utterances of ‘cortisol because stressed’ and ‘stressed because cortisol’ was an explanation ¹². Finally, we predicted that non-philosophers would be more likely to judge that an utterance of ‘stressed because cortisol’ is an explanation in a condition in which intervening on high cortisol to intervene on stress is salient, than in the incongruent condition in which intervening on stress in order to intervene on high cortisol is salient. Instead, we found that salience and congruence had no significant effects on people’s levels of explanation agreement.

Table 4 below summarises the descriptive data from the experiment of participants’ levels of truth agreement. The ‘Yes’ column represents the proportion of participants who reported that what Maria said is true (5, 6, or 7). The ‘No’ column represents the proportion of participants who reported that what Maria said is false (1, 2, or 3). The ‘I’ column represents the proportion of people who reported being indifferent between these two options (4). We also included the results of the one-sample *t*-tests for each condition. The results of these tests show mean levels of truth agreement are significantly above 4 across all conditions.

Table 4. Descriptive data from all conditions of participants’ levels of truth agreement.

Condition	% Yes	% No	% I	Mean	SD	<i>t</i> -Test	<i>p</i> -Value
Stress because cortisol congruent (<i>n</i> = 74)	71.6	16.2	12.2	5.09	1.48	6.353	<0.001
Stress because cortisol incongruent (<i>n</i> = 75)	65.3	18.7	16.0	5.39	1.83	6.563	<0.001
Cortisol because stress congruent (<i>n</i> = 71)	77.5	7.0	15.5	5.46	1.34	9.212	<0.001
Cortisol because stress incongruent (<i>n</i> = 81)	54.3	28.4	17.3	4.81	1.94	3.786	<0.001

To compare participants’ level of truth agreement across conditions, we tested level of truth agreement with a 2 (salience) × 2 (congruence) between-subjects ANOVA. The results of this analysis found a significant two-way interaction between salience and congruence $F(1, 297) = 5.932, p = 0.015$.

Simple effects tests using a Bonferroni correction were carried out on the two-way interaction between salience and congruence. In stress because cortisol conditions, there was no significant difference between congruent conditions and incongruent conditions ($p = 0.288$). In cortisol because stress conditions, levels of truth agreement were significantly higher in congruent conditions than in incongruent condition ($p = 0.018$).

In congruent conditions, there was no significant difference between stress because cortisol conditions and cortisol because stress conditions ($p = 0.185$). In incongruent conditions, levels of truth agreement were significantly higher in stress because cortisol conditions than in cortisol because stress conditions ($p = 0.034$).

Table 5 below summarises the descriptive data from the experiment of participants’ levels of explanation agreement. The ‘Yes’ column represents the proportion of participants who reported that what Maria said is an explanation for her (5, 6, or 7). The ‘No’ column represents the proportion of participants who reported that what Maria said is not an explanation for her (1, 2, or 3). The ‘I’ column represents the proportion of people who reported being indifferent between these two options (4). We also included the results of the *t*-tests for each condition. The results of these tests show that mean levels of explanation agreement are significantly above 4 across all conditions.

Table 5. Descriptive data from all conditions of participants’ levels of explanation agreement.

Condition	% Yes	% No	% I	Mean	SD	t-Test	p-Value
Stress because cortisol congruent (n = 74)	71.6	16.2	12.2	5.00	1.54	5.604	<0.001
Stress because cortisol incongruent (n = 75)	57.3	28.0	14.7	5.01	1.91	4.586	<0.001
Cortisol because stress congruent (n = 71)	81.6	8.5	9.9	5.28	1.35	7.975	<0.001
Cortisol because stress incongruent (n = 81)	56.8	27.2	16.0	4.81	1.83	4.005	<0.001

To compare participants’ level of truth agreement across conditions, we tested level of truth agreement with a 2 (salience) × 2 (congruence) between-subjects ANOVA. The results of this analysis found no significant effects. That is, we found no evidence that salience and congruence (or their interaction) have any effect on people’s levels of explanation agreement.

4.2. Experiment Two

Before reporting the statistics let us summarise our main findings. We predicted that if the assumption trio is true then non-philosophers would judge that ‘set because member’ is true, and that ‘member because set’ is not true. We did not find this. People, overall, judged that it is true that ‘member because set’ and that ‘set because member.’ We also predicted that non-philosophers would have judgements about the truth of those propositions that are insensitive to which intervention is made salient, and to whether the condition is congruent. This prediction was confirmed. Salience and congruence had no effect on peoples’ levels of truth agreement.

We also predicted that non-philosophers would judge that an utterance of ‘member because set’ is not an explanation. We did not find this. People, overall, judged that the utterances of ‘member because set’ and ‘set because member’ were both explanations. Finally, we predicted that non-philosophers would be more likely to judge that an utterance of ‘set because member’ is an explanation in a condition in which intervening on the member to intervene on the set is salient, than in the incongruent condition in which intervening on the set in order to intervene on the member is salient. We did not find this. Salience and congruence had no effect on peoples’ levels of explanation agreement.

Table 6 below summarises the descriptive data from the experiment of participants’ levels of truth agreement. The ‘Yes’ column represents the proportion of participants who reported that what Maria said is true (5, 6, or 7). The ‘No’ column represents the proportion of participants who reported that what Maria said is false (1, 2, or 3). The ‘I’ column represents the proportion of people who reported being indifferent between these two options (4). We also included the results of the one-sample t-tests for each condition. The results of these tests show mean levels of truth agreement are significantly above 4 across all conditions.

Table 6. Descriptive data from all conditions of participants’ levels of truth agreement.

Condition	% Yes	% No	% I	Mean	SD	t-Test	p-Value
Set because member congruent (n = 56)	85.7	5.4	8.9	5.66	1.33	9.380	<0.001
Set because member incongruent (n = 59)	57.6	30.5	11.9	4.92	1.94	3.621	0.001
Member because set congruent (n = 64)	71.8	14.1	14.1	5.17	1.60	5.862	<0.001
Member because set incongruent (n = 68)	63.3	23.5	13.2	5.10	2.01	4.528	<0.001

To compare participants’ level of truth agreement across conditions, we tested level of truth agreement with a 2 (saliency) × 2 (congruence) between-subjects ANOVA. The results of this analysis found no significant effects. That is, we found no evidence that saliency and congruence (or their interaction) have any effect on people’s levels of truth agreement.

Table 7 below summarises the descriptive data from the experiment of participants’ levels of explanation agreement. The ‘Yes’ column represents the proportion of participants who reported that what Maria said is an explanation for her (5, 6, or 7). The ‘No’ column represents the proportion of participants who reported that what Maria said is not an explanation for her (1, 2, or 3). The ‘I’ column represents the proportion of people who reported being indifferent between these two options (4). We also included the results of the one-sample *t*-tests for each condition. The results of these tests show mean levels of explanation agreement are significantly above 4 across all conditions.

Table 7. Descriptive data from all conditions of participants’ levels of explanation agreement.

Condition	% Yes	% No	% I	Mean	SD	<i>t</i> -Test	<i>p</i> -Value
Set because member congruent (<i>n</i> = 56)	78.6	12.5	8.9	5.45	1.50	7.213	<0.001
Set because member incongruent (<i>n</i> = 59)	64.4	22.0	13.6	5.10	1.68	5.042	<0.001
Member because set congruent (<i>n</i> = 64)	75.0	17.2	7.8	4.95	1.58	4.832	<0.001
Member because set incongruent (<i>n</i> = 68)	60.3	27.9	11.8	4.85	1.95	3.609	0.001

To compare participants’ level of truth agreement across conditions, we tested level of truth agreement with a 2 (saliency) × 2 (congruence) between-subjects ANOVA. The results of this analysis found no significant effects. That is, we found no evidence that saliency and congruence (or their interaction) have any effect on people’s levels of explanation agreement.

4.3. Experiment Three

Before reporting the statistics let us summarise our main findings. We predicted that if the assumption trio is true, then some non-philosophers would judge that ‘God approves *because* good’ is true, and others would judge that ‘good *because* God approves’ is true. We found this ¹³. We also predicted that these judgements would be insensitive to which intervention is made salient, and to whether the condition is congruent. We did not find this. Instead, we found a significant main effect of congruence on peoples’ levels of truth agreement.

We also predicted that some non-philosophers would judge that an utterance of ‘good *because* God approves’ is an explanation and some would judge that ‘God approves *because* good’ is an explanation. We found this. Finally, we predicted that non-philosophers would be more likely to judge that an utterance of ‘good *because* God approves’ is an explanation in a condition in which intervening on God in order to intervene on good is made salient, than they are in the incongruent condition in which intervening on good to intervene on God is made salient, and would be more likely to judge that an utterance of ‘God approves *because* good’ is an explanation in a condition in which intervening on good in order to intervene on God is made salient, than in the incongruent condition in which intervening on God in order to intervene on good is made salient. This prediction was supported. We found a significant main effect of congruence on peoples’ levels of explanation agreement.

Table 8 below summarises the descriptive data from the experiment of participants’ levels of truth agreement. The ‘Yes’ column represents the proportion of participants who reported that what Maria said is true (5, 6, or 7). The ‘No’ column represents the proportion of participants who reported that what Maria said is false (1, 2, or 3). The ‘I’ column represents the proportion of people who reported being indifferent between these two

options (4). We also included the results of the one-sample *t*-tests for each condition. The results of these tests show mean levels of truth agreement are significantly above 4 in congruent conditions, but *not* in incongruent conditions.

Table 8. Descriptive data from all conditions of participants’ levels of truth agreement.

Condition	% Yes	% No	% I	Mean	SD	<i>t</i> -Test	<i>p</i> -Value
Good because God approves congruent (<i>n</i> = 76)	65.8	15.8	18.4	4.96	1.66	5.042	<0.001
Good because God approves incongruent (<i>n</i> = 79)	24.0	38.0	38.0	3.81	1.76	−0.958	0.341
God approves because good congruent (<i>n</i> = 80)	67.5	17.5	15.0	4.86	1.63	4.737	<0.001
God approves because good incongruent (<i>n</i> = 84)	41.7	38.1	20.2	4.35	1.97	1.604	0.113

To compare participants’ level of truth agreement across conditions, we tested level of truth agreement with a 2 (salience) × 2 (congruence) between-subjects ANOVA. The results of this analysis found a significant main effect of congruence $F(1, 315) = 17.772, p < 0.001$.

The main effect of congruence showed that levels of truth agreement were significantly higher in congruent conditions ($M = 4.91, SD = 1.76$) than in incongruent conditions ($M = 4.08, SD = 1.76$).

Table 9 below summarises the descriptive data from the experiment of participants’ levels of explanation agreement. The ‘Yes’ column represents the proportion of participants who reported that what Maria said is an explanation for her (5, 6, or 7). The ‘No’ column represents the proportion of participants who reported that what Maria said is not an explanation for her (1, 2, or 3). The ‘I’ column represents the proportion of people who reported being indifferent between these two options (4). We also included the results of the one-sample *t*-tests for each condition. The results of these tests show mean levels of truth agreement are significantly above 4 in all conditions except for the “God approves because good” incongruent condition.

Table 9. Descriptive data from all conditions of participants’ levels of explanation agreement.

Condition	% Yes	% No	% I	Mean	SD	<i>t</i> -Test	<i>p</i> -Value
Good because God approves congruent (<i>n</i> = 76)	68.4	15.8	15.8	4.87	1.58	4.799	<0.001
Good because God approves incongruent (<i>n</i> = 79)	41.7	41.8	16.5	3.97	1.99	−0.113	0.910
God approves because good congruent (<i>n</i> = 80)	73.7	16.3	10.0	4.88	1.62	4.838	<0.001
God approves because good incongruent (<i>n</i> = 84)	45.2	31.0	23.8	4.49	1.95	2.289	0.025

To compare participants’ level of explanation agreement across conditions, we tested level of truth agreement with a 2 (salience) × 2 (congruence) between-subjects ANOVA. The results of this analysis found a significant main effect of congruence $F(1, 315) = 10.081, p = 0.002$.

The main effect of congruence showed that levels of explanation agreement were significantly higher in congruent conditions ($M = 4.87, SD = 1.80$) than in incongruent conditions ($M = 4.23, SD = 1.80$).

5. Discussion

There are two notable results. First, in all three cases, we found that non-philosophers have what we can call *symmetric judgements* both about truth and explanation. That is, across all three cases they judge both that ‘ x because y ’ and ‘ y because x ’ is true, and in all

three cases judge that utterances of both 'x because y' and of 'y because x' are explanations for the relevant party.

This appears to show that the assumption trio cannot be true. If non-recherché cases of metaphysical explanation are asymmetric, and if the folk notion of metaphysical explanation is continuous with that of the philosophical notion, then non-philosophers should judge that one of 'x because y' and 'y because x' is true, and that at most an utterance of either 'x because y' or 'y because x' is an explanation, but not both. This is not what we find. So, either the folk notion of explanation is quite different from that of the philosophical notion (folk–philosophical continuity assumption is false) or many more token explanations are symmetric than philosophers had supposed (the widespread asymmetry assumption is false).

The second notable result is the effect of congruence and salience on non-philosophers' judgements. If the assumption trio is true, then we expect these factors to have no effect on people's truth judgements, leaving open that they might have an effect on people's explanation judgements. This is not what we found. In the mind/brain case, we found that the interaction of salience and congruence had a significant effect on people's truth judgements but did not have an effect on their explanation judgments. In the set/member case, we found that salience and congruence had no effect on either judgement, and in the God/good case we found that congruence had an effect on both truth and explanation judgements. If both the truth *simpliciter* assumption and the folk–philosophical continuity assumption are correct, however, this is not what we should find. Instead, we should find that people judge that a proposition is true, or false, *simpliciter*, regardless of which intervention is salient.

Jointly, these two sets of results seem to show that at least two members of the assumption trio are false.

There are, however, several strategies that defenders of the assumption trio might adopt to preserve the trio in the face of these results.

First, they might argue that it is consistent with the two notions being continuous that non-philosophers have somewhat different judgments than philosophers, and that all our studies reveal is that this is the case.

There are two ways of pressing this idea. According to the first, what should be resisted is the idea, which we presupposed at the beginning of this paper, that if non-philosophers' concept of metaphysical explanation is the same or, or very similar to, that of philosophers, then they will tend to make similar judgements so that, for instance, if philosophers tend to judge that if p metaphysically explains q, and q does not metaphysically explain p, then so will non-philosophers. The idea would be that even though philosophers and non-philosophers share a very similar concept of metaphysical explanation, for some reason this is not reflected in their judgments. Perhaps, for instance, there are certain performance errors as regards philosophers or non-philosophers or both, so that their judgments differ even though the underlying concept is very similar.

We certainly concede that this *might* be so. Nothing we can tested here, can rule out that possibility. However, if this were so, then we take it that this would be newsworthy. When philosophers say that non-philosophers' concept of metaphysical explanation is continuous with that of philosophers, they seem to at least in part have in mind the idea that both parties will make similar judgements. If this turned out not to be so, that would be interesting. However, we also think that in the absence of a clear view about *why* these judgements systematically differ if the underlying concept is very similar, this is not a very plausible view. So, while we certainly cannot rule it out, much more would need to be said to defend the assumption trio, if one took this route.

An alternative way to press this idea would be to say that continuity is a matter of degree, and so even though the judgements of philosophers and non-philosophers differ in certain ways, this need not be thought to show that the underlying concept is *significantly* different. However, it seems to us that if the folk notion is one of which folk judgements on three non-recherché cases are so very different from philosophical judgements, then

this should give us pause in making the claim that the two notions are continuous in any interesting sense. We will, however, not have more to say on this idea.

Second, defenders of the assumption trio might argue that studies such as these are poor ways of probing the folk notion. Perhaps participants were confused; perhaps they require a careful tutorial on the notion that they are supposed to be employing before they can usefully be asked to make their judgements.

We have two responses to this worry. First, we see no reason to think that participants were confused. Although some of the vignettes describe unusual circumstances, they are not difficult to understand. Second, and perhaps more importantly, if the folk require a tutorial that explains to them the notion in question before they are able to render such judgments, then this itself tends to undermine our reasons for thinking that the folk–philosophical continuity assumption is true. The claim of continuity is meant to be the claim that the notion employed by non-philosophers *in ordinary contexts* is continuous with the notion employed by philosophers. This means that in ordinary contexts non-philosophers can and do employ the folk notion without someone first giving them a tutorial to explain that notion. So, while for all we know it might be true that prior to tutoring non-philosophers on the notion they are to employ, they are unable to employ the relevant notion, that being so would itself undermine the folk–philosophical continuity assumption.

Third, defenders of the assumption trio might argue that our methodology is a poor way to test the widespread asymmetry assumption. That is because we do not ask participants *both* whether ‘*x because y*’ is true *and* whether ‘*y because x*’ is true for each vignette. Instead, each participant only sees *one* such claim, and responds to whether they think it is an explanation or not. It could be, then, that participants are, for each case, split between those who think that ‘*x because y*’ is true and ‘*y because x*’ is not, and those who think the converse. Then, even though we would find agreement that both ‘*x because y*’ and ‘*y because x*’, this would not be because any individual participant judges that *both* are true, and so this evidence would not undermine the widespread asymmetry assumption.

There are several things to say about this idea. First, this proposal cannot explain our data, given the high levels of agreement that both ‘*y because x*’ and ‘*y because x*’ are true. If this proposal were correct, then we would expect to find much lower levels of agreement that either claim is true: for in any given case of evaluating ‘*y because x*’ or ‘*x because y*’ participants who take *only* the converse to be true will judge that the claim with which they are faced is false. Since according to this proposal in *every* such case there are participants who judge that only the converse explanatory claim is true, we should expect to find much lower levels of agreement that each of the claims are true, than what we in fact find. Moreover, there is empirical reason to reject this proposal. Recent empirical work by Latham and Miler [22] presented participants with similar vignettes and asked them *both* whether ‘*y because x*’ and whether ‘*y because x*’. They found that participants judged *both* claims to be true. So, we can set aside this response.

Fourth, defenders of the assumption trio might argue that non-philosophers employ many different explanatory notions, and that they are using different such notions to evaluate ‘*y because x*’ than they are to evaluate ‘*x because y*’. In particular, it may be that they are using the notion of metaphysical explanation to evaluate ‘*x because y*’, and as a consequence are judging that the relevant claim is true, and then are using some other notion to evaluate ‘*y because x*’, and as a consequence they are judging that *that* claim is also true. But if that is so, this would do nothing to undermine the widespread asymmetry assumption.

One version of this strategy would be to suggest that some of the vignettes that we offer can be read in causal terms, and so some participants take themselves to be evaluating causal explanations instead of metaphysical explanations. For instance, one might worry that the cortisol case can be read as causal: that the presence of cortisol *causes* stress, rather than constitutes its being the case that there is stress. However, even if this were so, we would still not predict the results that we attained. Participants who read the vignette as causal would be expected to respond that cortisol levels explain stress (since the causal

arrow runs from cortisol to stress) but *not* that stress explains cortisol. So, at best we might explain why we found agreement that cortisol explains stress.

For the multiple sources of explanation hypothesis to explain our results, we also need to explain why people judge that stress explains cortisol. Suppose the remaining participants interpret the explanatory claims as metaphysical rather than causal. To explain our results, it would need to be that these participants are inclined to judge (a) that stress explains cortisol and (b) that cortisol *does not* explain stress. Why is that? Well, they must judge (a) to explain our result, namely that people do judge that stress explains cortisol. Moreover, they must also *not* judge that cortisol explains stress, since if they did then this would undermine the widespread asymmetry assumption. But it seems unlikely that participants who interpret the vignette as metaphysical are judging that stress explains cortisol and not the converse. And, if they are, then those judgments are radically different from those of philosophers, which would go towards undermining the idea that the notion they are using is continuous with philosophers. Thus, this version of the suggestion seems to us to be implausible as an attempt to rescue the assumption trio.

Similarly, one might try to argue that the set/member case might be interpreted causally by participants and that this explains our results. The idea might be that since we describe the cases in terms of a biconditional, “sets only exist when their members do, and their members only exist when the set does” people might interpret this biconditional causally. Let us suppose that is so. If we begin with the assumption that people typically take causal relations and thus explanations, to be asymmetric, this means that people will either suppose that the existence of the set causes (and hence explains) the existence of the member, or the reverse, but not both. On the face of it, it seems much more likely that people will suppose that the existence of the member, a concrete person, causes the existence of the set rather than the other way around. If, however, people take that to be true, then they will tend to judge that the set exists because the member does and will not judge that the member exists because the set does, and so we should still expect to observe an asymmetry in judgements. So, in order for this claim about causal interpretation to explain our results, it would need to be that either (a) some people interpret the vignette as one in which the existence of the set causes the existence of the member, and some interpret it the opposite way or (b) that people take causation/causal explanation to be (sometimes) symmetric, so that even though they judge that, say, the existence of the member causes the existence of the set, they also think that the existence of the set causes the existence of the member, and thus the explanation proceeds in both directions.

We take (b) to be unlikely, though of course our results do not speak to causal cases. Still, we do know from empirical work that people tend to make asymmetric causal judgements (e.g., [34]). We also take (a) to be unlikely. It seems unlikely that many participants interpret the causal claim to go in one direction, and many in the other direction too. Indeed, we take it to be unlikely that people will suppose that an abstract object, like a set, causes the existence of a person. So, while we cannot rule out that people do interpret this vignette causally, even if they do, we think that it is very unlikely that this would be a good explanation of our results.

Another version of this strategy, that the defender of the assumption trio might adopt, is to argue that there are multiple kinds of non-causal explanation, and that participants are in fact employing several notions. One suggestion, for instance, is that amongst the class of non-causal explanations there are not only metaphysical explanations, but also *rationalising* explanations and *evidential* explanations¹⁴. Rationalising explanations are explanations that give us reasons why something is the case¹⁵. Evidential explanations are explanations that present evidence or justification for the explanans.

For instance, consider the following three ‘because’ claims:

1. The sign is red because it is maroon.
2. The sign is red because the colour of the bulb above it was changed yesterday.
3. The sign is red because it looks red under ideal conditions.
4. The sign is red because it signals danger.

(1) is (arguably at least) an example of metaphysical explanation. We take it that (2) is an example of a causal explanation. (3) is an example of evidential explanation¹⁶. The fact that the sign looks red under ideal conditions is evidence that it is red. (4) is an example of a rationalising explanation. The *reason* the sign is red is that red signals danger, and the purpose of the sign is to signal danger.

Given that, it might be then argued that some participants are evaluating *metaphysical* explanations, while others are evaluating *evidential* or *rational* explanations. One reason to think that this might be occurring is because participants are motivated to try and make sense of the statement with which they are presented. For example, participants presented with 'x because y' will read that as a metaphysical explanation if they are inclined to judge that y metaphysically explains x, and they will then judge that this is so. But, if they are inclined to judge that y does not metaphysically explain x, then they will read the 'x because y' in terms of some other kind of non-causal explanation whereby y *does* explain x. As a result, they too will then judge that 'x because y' is true. But again, if this is what is happening, it does not undermine the widespread asymmetry assumption.

Could this proposal explain our data? Well, appealing to rationalising explanations is unlikely to help. That is because for an appeal to this kind of explanation to explain our data, it would need to be that where people are inclined to judge that y metaphysically explains x, they are also inclined to judge that x rationally explains y. But, in all the cases that we offer this seems implausible, because the direction of the rationalising explanation tends to run in the same direction as the metaphysical explanation. So, while participants might be employing both notions of explanation, we would not expect them (collectively) to judge both that 'x because y' and that 'y because x'.

What about evidential explanations? Again, for this proposal to work, it will need to be the case that evidential explanations tend to run in the *opposite* direction from the metaphysical ones (in our cases). Is that plausible? Perhaps it is. We regrettably do not know enough about how non-philosophers think about evidential connections to be sure that they will judge that wherever we have metaphysical explanations we have evidential connections that run in both directions. Perhaps non-philosophers do think of evidential connections in this manner. If so, then if they judge that the presence of cortisol *metaphysically* explains the presence of stress, they will also judge that the presence of either stress or cortisol is *evidence* for the other. If that is right, then defenders of the assumption trio might explain our results by appealing to people's judgements regarding both metaphysical *and* evidential explanations. They could argue that people judge that cortisol metaphysically explains stress and not the converse, in line with the widespread asymmetry assumption, and that people also judge that stress *evidentially* explains cortisol. That is why participants agree both that cortisol because stress, and that stress because cortisol: they interpret the first statement as a statement about metaphysical explanation, and the second as a statement about evidential explanation. *Mutatis mutandis* for the other two cases.

We think that there is something to this suggestion; certainly, one that would be worth following up on in future research. To succeed, it requires that people judge that evidential connections obtain in what we might pre-theoretically have thought was the 'wrong' direction. For example, it requires that people judge that the presence of a singleton set is evidence for the presence of its member, and that is why they are inclined to judge that 'member because set', is true. We are not confident that people employ this notion of evidential explanation, but nor are we confident that they do not. We think that further research would be required to show that this is in fact the correct interpretation of our results, and hence to show that the assumption trio is in good standing. However, as it stands, we think that our results do put significant pressure on the assumption trio.

Fifth, defenders of the assumption trio might argue that non-philosophers may be unable, without prior tutoring, to clearly distinguish a proposition's being true, and an utterance of that proposition being an act of explanation. Perhaps our participants are treating these as equivalent, and all of their judgements are judgements about whether

the utterance is an act of explanation. If so, the fact that in some cases non-philosophers' "truth" judgements are sensitive to salience and congruence factors merely tells us that their *explanation* judgements are sensitive in this manner, something that is consistent with the assumption trio.

This possibility also strikes us as promising. Do we have reason to suppose it is true, beyond the fact that drawing the distinction between truth and explanation does appear to be something that non-philosophers might find difficult? There is evidence that non-philosophers draw *some* distinction here. In experiment 1, we found that the interaction of salience and congruence has an effect on one set of judgements and not the other. Rather puzzling, though, they have an effect on truth judgements but not explanation judgements. So, while participants do appear, at least in this case, to be treating these as different, there could be a case made that they are not thinking of judgments of truth, and judgements of utterances as being acts of explanation, in quite the way we are imagining. Notably though, in all the remaining experiments, participants' truth and explanation judgements are aligned.

So, let us suppose that participants are simply judging, in all cases, whether the utterance is an act of explanation. If so, then the experiments we report do nothing to undermine the truth *simpliciter* assumption. Nevertheless, they still show that the assumption trio is false. For they show that at least one of the other two assumptions is false. Recall that in all three experiments, participants judged both that utterances of $\lceil x \text{ because } y \rceil$ and of $\lceil y \text{ because } x \rceil$ are acts of explanation. As noted previously, philosophers hold that an utterance can only be an act of explanation if the proposition expressed is true, though it might *fail* to be an act of explanation even if the proposition expressed is true. If non-philosophers are employing roughly this notion of an act of explanation, then it follows that they are implicitly judging that both $\lceil x \text{ because } y \rceil$ and $\lceil y \text{ because } x \rceil$ are true, contra that widespread-asymmetry assumption. By contrast, if non-philosophers suppose that a speech act can be an act of explanation *even if the proposition expressed is false*, then they surely are employing a notion of metaphysical explanation (broadly construed) that is not continuous with that of philosophers. Then, the folk-philosophical continuity assumption is false. In either case, our results undermine the assumption trio.

The defender of the assumption trio might also argue that there is no *unified* diagnosis of what has gone wrong here; instead, she might suggest that there is something peculiar, or confounding, or potentially misleading, about each case, and this is why we get the results we do. For instance, she might suggest that in the mind/brain case some people might read the case as causal, and that some people might be swayed by background knowledge of minds/brains. She might then argue that in the set case, non-philosophers who are not mathematicians, simply do not know enough about sets to have a firm view, and so perhaps take the direction of explanation to proceed in both directions because the two are necessarily associated and in a way, all that is known about singleton sets and members is that they are associated in this way. Finally, she might try to tell some different story about why we get the results we do, in the Euthyphro case. Perhaps, for instance, she could think that people take different views on this depending on their religious views. We have already argued that appealing to a causal reading of the mind/brain case is not explanatory of our results. And while we acknowledge that we certainly cannot rule out that there are different explanations of our results in each case, where these explanations are consistent with the assumption trio, we think that the fact that we find results across three very *different* cases suggests that the results are robust.

Finally, the defender of the assumption trio might suggest that what explains these results in a way that is consistent with the assumption trio, is that although non-philosophers employ a notion of metaphysical explanation which embraces the widespread asymmetry assumption, their judgements in the cases we offer do not reflect this because in these cases participants confuse modal covariation with metaphysical explanation.

The idea, we take it, is this. Just as we know that people distinguish the *notion* of correlation from causation, we also know that sometimes people mistake merely correlational

relations for causal ones. This can fairly easily be explained by the fact that mechanisms that evolved to detect causal relations do so by, in the first instance, detecting associations, and then sorting these into the causal and noncausal (see, for example, [34] and [37]). We know, however, that sometimes having detected an association, the mechanism then fails to properly sort that association into causal versus not.

Interestingly, Miller and Norton [30] suggest that the same mechanisms that evolved to detect causal relations, also generate people's judgments about metaphysical explanations. While they offer their account as a debunking account of metaphysical explanation, it could instead be offered as an account of *how* people come to detect metaphysical explanations. It might be argued that the mechanisms which allow us to detect causal relations sometimes mistake mere modal association for metaphysical explanation, particularly in cases where not only does *x* supervene on *y*, but *y* also supervenes on *x*. This, at least, could potentially explain people's judgements in the Euthyphro and Sets cases. In both cases, we have mechanisms that detect the modal association; however, the mechanism then fails to properly sort these associations into explanatory and non-explanatory, instead labelling them *both* as explanatory. So, it is not that people do not generally take metaphysical explanation to be asymmetric, it is just that in these cases people's cognitive mechanisms lead them astray.

Still, while we think that the general idea here has some plausibility, because we think that it may be that people's judgements about metaphysical explanation are at least in part informed by the functioning of something like the casual detection mechanism, in fact this explanation of our results does not sit well with our results taken as a whole. If this were the right explanation, then we would expect to find people's judgements to be much stronger in the congruent conditions than the incongruent ones. Why so? Because the causal detection mechanism functions, in part, using actual or imagined interventions, which allow us to detect *asymmetries*: that is, to detect that intervening on one thing has an effect on the other, but not the converse. When this is the case, the mechanism tends to signal that the correlation is causal, and tells us which is the cause, and which the effect [34,37]. As Haggmeyer et al. put it:

Interventions often enable us to differentiate amongst the different causal structures that are compatible with an observation. If we manipulate an event *A* and nothing happens, then *A* cannot be the cause of event *B*, but if a manipulation of event *B* leads to a change in *A*, then we know that *B* is a cause of *A*, although there might be other causes of *A* as well [37].

There is a range of empirical data that support the idea that the causal detection mechanism seeks out asymmetries via various environmental cues. So, if people were using this sort of mechanism to detect metaphysical explanations by sorting the explanatory from the non-explanatory modal associations, then we would expect that in conditions of congruence, people will tend to be much more likely agree that there is an explanation than in conditions of incongruence. After all, the former are precisely cases where a relevant intervention is made salient, and this should trigger the causal detection mechanism to categorise the association as a metaphysical explanation. However, recall that while in the God/good case we found that congruence had an effect on both truth and explanation judgements, in the set/member case we found that salience and congruence had *no* effect on people's judgements. This seems to us to militate against this explanation of our results¹⁷.

At this point, the defender of the assumption trio might argue that even if all that we say here about our results is correct, it still does not undermine the assumption trio, because it does not undermine the assumption of continuity. For it would not show, or so goes the idea, that non-philosophers do not share a notion of metaphysical explanation that is closely connected to the one employed by philosophers. Here is a case in point¹⁸. It seems *prima facie* reasonable to think that non-philosophers have a concept of acceleration that is close to the notion as it appears in Newtonian physics. But suppose, as it also plausible, non-philosophers judge that some things do not count as accelerations, which according to Newtonian physics do count as accelerations. (For instance, consider the case of a carnival

ride in which people sit in chairs connected to a central rotating column spinning at a uniform rate. Suppose that people judge that the chairs are not accelerating in that case.) Even though the judgements of Newtonian physicists regarding acceleration and those of non-philosophers differ, we would still be inclined to say that they share a concept, or at least, that their concepts are strongly continuous. If so, then it does not follow that just because certain judgements of philosophers regarding metaphysical explanation are different from those of non-philosophers, that they do not share the same, or at least a continuous, concept of metaphysical explanation.

We agree that the mere fact that two parties disagree about the extension of some concept does not show that their concepts are not continuous. However, we do think that if (a) two parties disagree on sufficiently many cases, or if (b) they disagree about certain core or essential cases, or if (c) they disagree on a range of easy or paradigm cases, then this gives us a good reason (though still defeasible) to think that their concepts are neither the same nor continuous (or at least, not strongly continuous). As such, we take it that the evidence we present goes some very small way towards showing (a) to be true and can plausibly be said to show that (c) is true too. We think that the cases we presented are easy (except perhaps sets and members), paradigmatic ones, and so if philosophers and non-philosophers disagree about them, then it gives a reason to think that their concepts are not strongly continuous. Moreover, although we only investigated three cases, and so our results certainly do not show that philosophers and non-philosophers disagree about *lots* of cases, it does show that they disagree about the only cases that have been empirically investigated. And that should give us pause to worry that they might not agree about remaining the cases either. So, while we do not think that this evidence *decisively* shows that the two concepts are not continuous, we certainly think it puts pressure on the claim that they are. Again, this is a place where follow up empirical work could profitably be undertaken to see just how widespread the disagreement lies when it comes to such cases.

What, though, should we conclude if we are inclined to say that at least one member of the assumption trio is false? We could simply embrace the idea that the notion of metaphysical explanation employed by non-philosophers is quite different from that employed by philosophers. Perhaps we philosophers have developed a specialized notion of explanation that performs certain theoretical work for us, and which is just very different from the everyday notion.

This approach leaves much of current philosophical methodology untouched; but it requires we concede that the notion of metaphysical explanation, as it is used by philosophers, does not have some long-standing pedigree: it is not a well understood, broadly used notion on which we all have a good grasp. Indeed, it leaves open the charge that philosophers have invented some new, parochial notion, which may be uninteresting, or worse, unintelligible. It also leaves us open to the charge that our theorising fails to target the very sorts of questions and concerns of non-philosophers which motivated employing the notion to begin with.

Alternatively, we might cling to the idea that the notions are indeed continuous, and engage in a project of explicating the content of those notions by appealing to empirical work such as ours. To pursue this approach would require further empirical research into the conditions under which people judge that one thing metaphysically explains another. Given the results we report here, if we pursue this approach we have reason to think that at the very least the widespread-asymmetry assumption is false: current ways of thinking about metaphysical explanation have over emphasised the role of asymmetry in our judgements. That, in turn, has significant down-stream consequences for theorising about the conditions under which propositions of the form 'x because y' are true.

There are already a great number of these accounts, ranging from those that appeal to relations of ground in some way¹⁹, through to those that do not²⁰. Many, but not all, of the accounts that appeal to ground posit a primitive, asymmetric, relation, in part motivated by the purported need for some asymmetric worldly structure that can back asymmetric metaphysical explanation. But if metaphysical explanation is not asymmetric,

then positing a relation with these features, (at least for these reasons) is clearly a mistake. Similar considerations apply to non-grounding approaches that attempt to capture the asymmetry of metaphysical explanation.

So if philosophers embrace this approach, our results suggest that some of the theorising about the properties of the worldly structures that back metaphysical explanations are mistaken: for they are premised on erroneous assumptions about the properties of metaphysical explanations.

This second approach requires significant methodological upheaval. If philosophical judgements are not a good guide to the judgements of non-philosophers (as this research suggests) then much more engagement with empirical work is required. Moreover, there are then questions about why philosophical judgements are so different from everyday judgments if they are both judgments about the same notion, or at least, two notions that are continuous, and in turn, questions about how best to systematise these quite different judgements into a cohesive account of metaphysical explanation.

Whichever approach is pursued, we take it that there is still interesting and important work to be done in investigating non-philosophers' notion of metaphysical explanation whether or not this is taken to be continuous with the notion employed by philosophers. We hope to have begun making some, small, inroads into this task.

Author Contributions: Conceptualization, A.J.L. and K.M.; methodology, A.J.L. and K.M.; software, A.J.L.; validation, A.J.L. and K.M.; formal analysis, A.J.L.; investigation, A.J.L.; resources, A.J.L. and K.M.; data curation, A.J.L.; writing—original draft preparation, A.J.L. and K.M.; writing—review and editing, A.J.L. and K.M.; visualization, A.J.L.; supervision, K.M.; project administration, K.M.; funding acquisition, K.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by The Australian Research Council DP180100105 and FT170100262.

Institutional Review Board Statement: The study was conducted according to the guidelines of the Declaration of Helsinki, and approved by the University of Sydney Human Research Ethics Committee (2019/349).

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The original data presented in the study are openly available in the OSF at <https://osf.io/sc4qf/>.

Conflicts of Interest: The authors declare no conflict of interest.

Notes

¹ We use corner quotes here to signify that 'x because y' is a kind of sentence, where x and y are variables that range over sentences. We will speak of 'an instance of 'x because y'' when we intend to talk about a particular instance of the schema. We will simply speak of 'x because y' in order to talk about all instances of the schema.

² Though there are some, such as Wilson [1] and Schaffer [2], who think that metaphysical explanation is a sort of non-diachronic causation.

³ See Correia and Schnieder [6], Jenkins [7], Raven [8,9] and Skiles [10] as endorsing this view. Notable exceptions to this trend include Thompson [11,12], Maurin [13] and Dasgupta [4].

⁴ See Schaffer [3,14], Raven [8,9], Rosen [15], Audi [16,17] and Cameron [18].

⁵ Bliss [19,20].

⁶ Sometimes these are framed as examples of symmetrical instances of grounding. Since authors often move freely between these terms, it is clear that these examples constitute examples of symmetrical metaphysical explanation.

⁷ The only work we know of is that by Latham and Miler [22], which investigates whether judgments about metaphysical explanation are context sensitive.

⁸ Gopnik [32], Steyvers, Tenenbaum, Wagenmakers & Blum [33] and Sloman [34].

⁹ Einhorn & Hogarth [35] and Murphy & Medin [28].

¹⁰ This is an abstract characterization of the sentence participants respond to, which has the form 'Fred is stressed because he has high cortisol levels.' Likewise for the other sentences in our predictions.

- 11 These are people in a large database who partake in a range of online experiments, usually in psychology, behavioral economics and sociology, for monetary compensation. While they have significant experience in completing online experiments, there is little reason to think that these people will have a particular interest in, or knowledge of, philosophy.
- 12 You might wonder whether people's background knowledge of psychology, neuroscience, and related areas, would influence their responses here. This is an interesting suggestion and would be worth investigating in future. In the current study, we think that the numbers of participants who would have specialized knowledge in these areas is very likely to be low, and so we think it unlikely that the current results would be significantly impacted.
- 13 Like with fn. 12, you might wonder whether there is an association here, between people's responses and either their religiosity (or religious belief). This too would be worth investigating in future. However, while different religions and denominations do take a clear stance regarding the correct direction of explanation in this case (i.e., Divine Command Theory) there is still persistent disagreement between them regarding *which* direction is the correct one.
- 14 With thanks to an anonymous referee for pressing us further on this issue.
- 15 For discussion, see [36].
- 16 At least, on the assumption that you do not think that what it is for something to be red *just is* for to look red under certain conditions.
- 17 One might worry that in the God/good vignette it is not specified that the prayer in fact changes God's mind. However, we take it that even if participants did not come to believe that the prayer changed God's mind, that the intervention that proceeds via God's mental states is, in this vignette, made salient as compared to an intervention on what is good. Indeed, that seems to be borne out by the results here which found an effect of congruence on peoples' judgements.
- 18 This case was suggested by an anonymous referee.
- 19 Schaffer [3], Cameron [18], Audi [16], Raven [38] and Trogdon [39].
- 20 For a variety of such approaches, see Kovacs [40], Miller & Norton [30], Baron & Norton [41], Wilson [42], Shaheen [43], Bertrand [44] and Norton & Miller [45].

References

1. Wilson, A. Metaphysical causation. *Noûs* **2018**, *52*, 723–751. [CrossRef]
2. Schaffer, J. Grounding in the image of causation. *Philos. Stud.* **2016**, *173*, 49–100. [CrossRef]
3. Schaffer, J. On What Grounds What. In *Metametaphysics: New Essays on the Foundations of Ontology*; Manley, D., Chalmers, D., Wasserman, R., Eds.; Oxford University Press: Oxford, UK, 2009.
4. Dasgupta, S. Constitutive Explanation. *Philos. Issues* **2017**, *27*, 74–97. [CrossRef]
5. Glazier, M. Explanation. In *Routledge Handbook of Metaphysical Grounding*; Raven, M.J., Ed.; Routledge: London, UK, 2020; pp. 121–132.
6. Correia, F.; Schnieder, B. *Metaphysical Grounding: Understanding the Structure of Reality*; Cambridge University Press: Cambridge, UK, 2012.
7. Jenkins, C. Is Metaphysical Dependence Irreflexive? *Monist* **2011**, *94*, 267–276. [CrossRef]
8. Raven, M.J. Is ground a strict partial order? *Am. Philos. Q.* **2013**, *50*, 193–201.
9. Raven, M.J. Ground. *Philos. Compass* **2015**, *10*, 322–333. [CrossRef]
10. Skiles, A. Against Grounding Necessitarianism. *Erkenntnis* **2015**, *80*, 717–751. [CrossRef]
11. Thompson, N. Metaphysical Interdependence. In *Reality Making*; Jago, M., Ed.; Oxford University Press: Oxford, UK, 2016.
12. Thompson, N. Questions and Answers: Metaphysical Explanation and the Structure of Reality. *J. Am. Philos. Assoc.* **2019**, *5*, 98–116. [CrossRef]
13. Maurin, A. Grounding and metaphysical explanation: Its complicated. *Philos. Stud.* **2019**, *176*, 1573–1594. [CrossRef]
14. Schaffer, J. The Internal Relatedness of All Things. *Mind* **2010**, *119*, 341–376. [CrossRef]
15. Rosen, G. Metaphysical Dependence: Grounding and Reduction. In *Modality: Metaphysics, Logic, and Epistemology*; Hale, B., Hoffmann, A., Eds.; Oxford University Press: Oxford, UK, 2010; pp. 109–136.
16. Audi, P. Grounding: Toward a Theory of the In-Virtue-Of Relation. *J. Philos.* **2012**, *109*, 685–711. [CrossRef]
17. Audi, P. A clarification and defense of the notion of grounding. In *Metaphysical Grounding: Understanding the Structure of Reality*; Correia, F., Schnieder, B., Eds.; Cambridge University Press: Cambridge, UK, 2012; pp. 101–121.
18. Cameron, R. Turtles all the way down: Regress, Priority and Fundamentality. *Philos. Q.* **2008**, *58*, 1–14. [CrossRef]
19. Bliss, R. Viciousness and the Structure of Reality. *Philos. Stud.* **2013**, *166*, 399–418. [CrossRef]
20. Bliss, R. Viciousness and Circle of Ground. *Metaphilosophy* **2014**, *45*, 245–256. [CrossRef]
21. Rodriguez-Pereyra, G. Grounding is not a Strict Order. *J. Am. Philos. Assoc.* **2015**, *1*, 517–534. [CrossRef]
22. Latham, A.J.; Miller, K. Judgements of Metaphysical Explanations are Context Sensitive. Available online: <https://philpapers.org/rec/LATJOM> (accessed on 30 May 2024).
23. Vasilyeva, N.; Wilkenfeld, D.; Lombrozo, T. Contextual utility affects the perceived quality of explanations. *Psychon. Bull. Rev.* **2017**, *24*, 1436–1450. [CrossRef] [PubMed]
24. Waskan, J.; Harmon, I.; Horne, Z.; Spino, J.; Clevenger, J. Explanatory anti-psychologism overturned by law and scientific case classifications. *Synthese* **2014**, *191*, 1013–1035. [CrossRef]

25. Wilkenfeld, D.; Lombrozo, T. Inference to the best explanation (IBE) versus explaining for the best inference. *Sci. Educ.* **2015**, *24*, 1059–1077. [[CrossRef](#)]
26. Kahneman, D.; Miller, D.T. Norm theory: Comparing reality to its alternatives. *Psychol. Rev.* **1986**, *93*, 136–153. [[CrossRef](#)]
27. Hitchcock, C.; Knobe, J. Cause and norm. *J. Philos.* **2009**, *11*, 587–612. [[CrossRef](#)]
28. Murphy, G.L.; Medin, D.L. The role of theories in conceptual coherence. *Psychol. Rev.* **1985**, *92*, 289–316. [[CrossRef](#)]
29. Chapman, L.J.; Chapman, J.P. Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *J. Abnorm. Psychol.* **1969**, *74*, 271–280. [[CrossRef](#)] [[PubMed](#)]
30. Miller, K.; Norton, J. Grounding: Its (probably) all in the head. *Philos. Stud.* **2017**, *174*, 3059–3081. [[CrossRef](#)]
31. Wilson, A. Grounding entails counterpossible non-triviality. *Philos. Phenomenol. Res.* **2018**, *96*, 716–728. [[CrossRef](#)]
32. Gopnik, A.; Glymour, C.; Sobel, D.; Schulz, L.; Kushnir, T.; Danks, D. A theory of causal learning in children: Causal maps and Bayes nets. *Psychol. Rev.* **2004**, *111*, 3–32. [[CrossRef](#)] [[PubMed](#)]
33. Steyvers, M.; Tenenbaum, J.B.; Wagenmakers, E.J.; Blum, B. Inferring causal networks from observations and interventions. *Cogn. Sci.* **2003**, *27*, 453–489. [[CrossRef](#)]
34. Sloman, S. *Causal Models: How People Think about the World and Its Alternatives*; OUP: Oxford, UK, 2005.
35. Einhorn, H.J.; Hogarth, R.M. Judging probable cause. *Psychol. Bull.* **1986**, *99*, 3–19. [[CrossRef](#)]
36. Berker, S. The Unity of Grounding. *Mind* **2018**, *127*, 729–777. [[CrossRef](#)]
37. Hagmayer, Y.; Sloman, S.; Lagnado, M.D.; Waldman, A. Causal Reasoning Through Intervention. In *Causal Learning: Psychology, Philosophy and Computation*; Gopnik, A., Schultz, L., Eds.; OUP: Oxford, UK, 2007.
38. Raven, M.J. In Defence of Ground. *Australas. J. Philos.* **2012**, *90*, 687–701. [[CrossRef](#)]
39. Trogdon, K. An introduction to grounding. In *Varieties of Dependence*; Hoeltje, M., Steinberg, B.S.A., Eds.; Philosophia Verlag: Munich, Germany, 2013; pp. 97–122.
40. Kovacs, D.M. Metaphysically Explanatory Unification. *Philos. Stud.* **2020**, *177*, 1659–1683. [[CrossRef](#)]
41. Baron, S.; Norton, J. Metaphysical Explanation: The Kitcher Picture. *Erkenntnis* **2021**, *86*, 187–207. [[CrossRef](#)]
42. Wilson, J. No Work for a Theory of Grounding. *Inquiry* **2014**, *57*, 196–215. [[CrossRef](#)]
43. Shaheen, J.L. The causal metaphor account of metaphysical explanation. *Philos. Stud.* **2017**, *174*, 553–578. [[CrossRef](#)]
44. Bertrand, M. Metaphysical Explanation by Constraint. *Erkenntnis* **2019**, *84*, 1325–1340. [[CrossRef](#)]
45. Norton, J.; Miller, K. A Psychologistic Theory of Metaphysical Explanation. *Synthese* **2017**, *196*, 2777–2802. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.