Generality explained

Øystein Linnebo

Preprint – October 2021

Abstract

What explains the truth of a universal generalization? Two types of explanation can be distinguished. While an *instance-based explanation* proceeds via some or all instances of the generalization, a *generic explanation* is independent of the instances, relying instead on completely general facts about the properties or operations involved in the generalization. This intuitive distinction is analyzed by means of a truthmaker semantics, which also sheds light on the correct logic of quantification. On the most natural version of the semantics, this analysis vindicates some claims made—without a proper defense—by Michael Dummett, Solomon Feferman, and others. Where instance-based explanations are freely available, classical logic is shown to be warranted. By contrast, intuitionistic logic (or slightly more) remains warranted regardless of what explanations are available.

Consider the following universal generalizations:

(1) Every student of mine is born on a Monday

(2) Every whale is a mammal

(3) Every natural number has a successor

Suppose each generalization is true. What explains its truth? That is, in virtue of what is the generalization true? These questions are metaphysical, not epistemic. We are not asking for *evidence* for the generalizations. Rather, we want to know what it is about reality that is responsible for the truth of the generalization.

These metaphysical questions admit of different kinds of answer. Let us say that an explanation of a universal generalization is *instance-based* insofar as it proceeds via individual instances of the generalization. For example, (1) permits an instance-based explanation. Suppose that my students are $a_1, a_2, \ldots$, and a_m. Then (1) is true because my students are $a_1, a_2, \ldots, a_n$, and because each of these objects is born on a Monday. By contrast, (2) admits of an entirely different kind of explanation. Its truth can be fully explained without mention of any particular whale, namely by citing the completely general fact that part of what it is to be a whale is to be a mammal. Likewise, the truth of (3) can be fully explained by completely general facts about the property of being a natural number and the successor operation. Let us say that an explanation of the truth of a universal generalization is *generic* insofar as it does not proceed via individual instances but is based on general facts about the properties or operations involved in the claim that is generalized. The list of thinkers who have taken an interest in non-instance-based explanations of true generalizations is long and varied; see, for example, (Kant, 1997, B3), (Russell, 1908, §II), (Weyl, 1921, p. 54) (building on Brouwer), (Hilbert, 1926, p. 194), (Carnap, 1931, §4), (Gödel, 1944, pp. 455-56), (Kaufmann, 1978) (building on Husserl), (Dummett, 1993, p. 440), (Rosen, 2010, §8), (Fine, 2014), (Wright, 2018), and (Hale, 2020).

Some clarifications are in order. First, being instance-based or generic is a matter of degree. An explanation can proceed via all instances of a generalization, via some but not all instances, or via none at all. The mentioned explanation of (1), for example, is partially but not wholly instance-based. While the explanation is directly concerned with each and every student of mine, it is not directly concerned with any particular object that is not a student of mine, although these objects too are in the range of the generalization.[1]

Second, the two forms of explanation are not in competition. There is no need to choose one kind of explanation to the exclusion of the other. For one thing, there are generalizations that naturally call for a hybrid explanation. Consider:

(4) Every object that is either a student of mine or a whale is either born on a Monday or a mammal

This odd but (we are supposing) true generalization is most naturally explained in a way that combines instance-based and generic considerations. For another, there are true generalizations that admit of both kinds of explanation. Consider (2). For simplicity, assume that the domain consists of past and present objects and that the two predicates used in the generalization have been provided with perfectly precise definitions. As observed, the resulting truth admits of a generic explanation. But it also admits of an instance-based explanation. Suppose the whales are $b_1, \ldots, b_n$. Then (2) is fully explained by the fact that these objects comprise all and only the whales and that each of them is a mammal. Of course, this instance-based explanation is far less informative than the generic explanation. But informative or not, this *is* an explanation: it specifies some features of reality that fully account for the truth of (2).

Third, I am not postulating an ambiguity in the quantifiers reflecting the kind of ex-

[1] To make this clear, we may suppose that (1) is analyzed as $\forall x(Sx \rightarrow Mx)$, with the obvious definitions of 'S' and 'M'.

planation that a generalization requires. Nothing I have said should alter your view on the univocality of the quantifiers. My claim is rather that perfectly univocal generalizations can, when true, be explained in interestingly different ways: some in a highly instance-based way, some in a purely generic way, and some in both of these ways.

My overarching aim is to understand how the different kinds of explanation work, both in their pure forms and in combination. To achieve this aim, I develop a truthmaker semantics and explore how states of the world can verify, or make true, sentences of our language—including universal generalizations.[2] This investigation is carried out using ordinary ZFC set theory in a classical logic.

I investigate two more specific questions as well. First, under what conditions are the different kinds of explanation available? A partial answer is obvious and uncontroversial: in many cases, only an instance-based explanation is available. There is, we may suppose, no uniform reason why (1) is true. It is not the case that only students born on a Monday are admitted into my classes, or that they alone (perhaps for some bizarre psychological reason) are interested in what I teach. So a generic explanation of (1) is out of the question. A less obvious—but more interesting—question concerns the availability of instance-based explanations. I show that, on a wide range of metaphysical views, there are domains that don't consist of a definite range of instances. And where there isn't a definite range of instances, instance-based explanations are not always available. Thus, there will be true generalizations that only admit of a (wholly or partially) generic explanation.

Second, what consequences does my analysis have for the correct logic of quantification? I show that, on the most natural version of the truthmaker semantics, there is an interesting connection between the availability of the two forms of explanations and the correct logic of quantification. Assume that instance-based explanations are freely available. Then the quantifiers can be shown to obey classical logic. Next, assume instead that instance-based explanations are *not* freely available. This makes the truth of a universal generalization a "bigger deal", as the generalization may need to be underwritten by a (wholly or partially) generic explanation. The truth of a universal generalization $\forall x\, \varphi(x)$ may thus be more demanding than the mere absence of counterexamples, that is, the failure of $\exists x \neg\varphi(x)$ to be true. In technical parlance, when instance-based explanations are not freely available, there is no guarantee that the universal and existential quantifiers are dual to each other: universal quantification may well be stronger than the dual of existential (namely the absence of

[2] There may well be other ways to analyze the different kinds of explanation of true generalizations. An analysis in terms of the metaphysical notion of grounding seems a promising alternative.

counterexamples.) This loss of duality undermines the validity of classical logic. I prove that *intuitionistic* logic remains valid regardless of what explanations are available. Intuitionistic logic thus provides an extremely robust fall-back, available even on metaphysical views that severely restrict the explanations that are available for true universal generalizations. This vindicates some claims famously made by Michael Dummett, and more recently also by Solomon Feferman, but which have never been properly substantiated.[3]

In another respect, however, my analysis differs starkly from Dummett's (or, for that matter, Brouwer's): my analysis completely divorces intuitionistic logic from the form of anti-realism that is so often used to support this logic and with which this logic is therefore often associated.[4] This form of anti-realism, central to the views of Brouwer and Dummett, forges a tight—and extremely controversial—link between truth and proof or evidence. By contrast, on my analysis, the validity of classical or intuitionistic logic for the quantifiers turns on the extent to which instance-based explanations are available. This is an abstract structural question that arises in connection with a wide range of metaphysical views, including some of a robust realist character, as we shall see in the next section.

1 Why generic explanations are needed

Consider a true universal generalization. An instance-based explanation requires that it be possible to consider all of its instances, provide an explanation of each instance, and then conjoin all of these individual explanations to produce the desired instance-based explanation. Thus, when there isn't a definite range of instances to consider, the instance-based approach is unavailable.

How can a domain of quantification be indefinite? One source of indefiniteness is *local*, in the sense that it concerns whether or not some given object is in the domain. Suppose we are investigating the causes of baldness and thus wish to consider the domain of all and only bald men. Since the predicate 'bald' is vague, we may have failed to determine a domain. This local form of indefiniteness will not concern us here. Indeed, we shall proceed on the assumption that all of our predicates have perfectly sharp definitions, such that for any given string of objects, each predicate either determinately applies to the objects (in the given order) or determinately does not.

Our concern is a *global* form of indefiniteness: not whether some given object belongs

[3] See e.g. (Dummett, 1991, p. 319) and (Dummett, 1993, p. 440), as well as (Feferman, 2014) and the unpublished but widely circulated (Feferman,).

[4] See (Fine, 2014) and (Rumfitt, 2012) for some earlier arguments for this divorce.

to the domain, but the extent of the domain in its entirety. Let me sketch some possible examples. To be absolutely clear: I am not endorsing the metaphysical views involved in the examples; I merely use them as examples.

First, consider the view that the future is metaphysically open. For example, there may or may not be a sea battle tomorrow; nothing currently settles the matter one way or the other. Consider

(2) Every whale is a mammal

interpreted such that the generalization includes not only past and current objects but also future ones. Suppose we have a perfectly sharp (perhaps cladistic) definition of 'whale' and 'mammal'. But if the future is open, it is currently not fully determinate which objects will exist in, say, ten years. The indeterminacy concerning the future affects not only about what will happen to objects that already exist but also what objects will come to exist. This makes it impossible to consider all the instances of the generalization: the objects with which the generalization is concerned are not all available to be considered. Thus, there cannot be a wholly instance-based explanation of (1). Any explanation would have to be at least partially generic, thus ensuring that the generalization will continue to hold whatever objects the future brings into existence.

Second, consider the ancient notion of potential infinity. According to Aristotle and a majority of mathematicians and philosophers until the Cantorian revolution of the late 19th century, this is the only legitimate notion of infinity (at least outside of theology). On this view, an infinite domain is "always unfinished". However many natural numbers have been "produced"—say by producing an associated numeral or a collection that instantiates the number in question—it is possible to "produce" more. There can be no such thing as a complete list of all natural numbers. Any list of numbers can be used to define a larger number, namely the least number larger than all the ones on the list. This means that the true generalization

(3) Every natural number has a successor

does not admit of a wholly instance-based explanation. Since the natural numbers are merely potentially infinite, they are never all available to serve in an instance-based explanation. So any explanation of (3) would have to be at least partially generic, thus ensuring that the generalization will continue to hold of any number that might be produced.[5]

[5]For details, see (Linnebo and Shapiro, 2019), especially the discussion of "strict potentialism".

There are more relaxed forms of potentialism as well. *Set-theoretic potentialism* agrees with Cantor that there is a "complete totality" of all natural numbers but denies that there is such a totality of all sets. Rather, the hierarchy of sets is open-ended and "always unfinished"; however many sets have been "formed", all these sets can be used to form yet more sets.[6] But if there is no "complete totality" of sets, it is impossible to consider all the sets, as would be required for a wholly instance-based explanation of a set-theoretic generalization such as:

(5) Every set has a power set

Any explanation of (5) would have to be at least partially generic.

Third, recall that Hilary Putnam rejects "metaphysical realism", understood as the view that there is a definite totality of absolutely all objects.[7] Rather, he claims, what objects there are depends on which concepts we bring to bear in our thoughts and theories. Can there still be true generalizations concerning absolutely all objects? Some philosophers think not.[8] Although natural, this conclusion is too quick. True, Putnam's view entails that no all-encompassing generalization can have a wholly instance-based explanation; for there is no definite range of instances to consider. But his view leaves the door open to such generalizations—provided that they admit of a (wholly or partially) generic explanation, thus ensuring that the generalization continues to be true of any objects "carved out" by concepts yet to be introduced.

Fourth, many philosophers are currently attracted to a hierarchical conception of reality, according to which objects and truths are arranged in a hierarchical manner, with each stage fully explicable in terms of the preceding ones.[9] In particular, every true universal generalization must be "located" somewhere or other in the hierarchy of truths. This gives rise to a problem. If a true generalization ranges over absolutely everytfhing, it will have instances that are not available at the stage of the hierarchy where this truth is located. Yet the generalization must be fully explicable in terms of truths that are available at preceding stages. Thus, if the hierarchical conception is to permit absolutely universal generalizations to be true at all, the needed explanation cannot be wholly instance-based but must involve some reliance on generic explanation as well.

---

[6] Set-theoretic potentialism can be traced back to Cantor; see his letters to Dedekind and Hilbert, available in (Ewald, 1996), as well as (Jané, 2010) and (Linnebo, 2013) for discussion. Arguably, this form of potentialism is just an instance of a more general phenomenon of indefinite extensibility; see (Dummett, 1991, ch. 24) and (Dummett, 1993), as well as (Linnebo, 2018a) for a recent discussion.

[7] See e.g. (Putnam, 1987).

[8] See (Parsons, 2006).

[9] See e.g. (Fine, 2012) and (Schaffer, 2009).

Finally, consider the mentioned anti-realist view that explicates truth in terms of proof. On this view, an explanation of the truth of a statement is naturally understood as a finite proof. It follows that no generalization over an infinite domain can have a wholly instance-based explanation, since the associated proof would be infinite. Notice, however, that this example of a metaphysical view that limits the availability of instance-based explanations is neither representative of such views in general nor particularly distinguished in terms of its plausibility.

All these metaphysical views share a common structure. There is a range of objects that are "available", namely by being past or present; by being produced, formed, or "carved out"; or by existing at the stage of a hierarchical reality at which an explanation is to be given. Yet there are generalizations whose domains include the available objects but exceeds these in a way that cannot be properly circumscribed. The range of available objects can always be expanded, namely by letting time pass, by actively producing or "carving out" more objects, or by considering stages that are higher up in a hierarchical reality. Even so, the available objects can never exhaust the domain of the generalizations, as it is always possible for even more objects to become available.

This shared common structure prompts a daunting challenge. We would like to explain every true generalization on the basis of the available objects. Is this possible? For existential generalizations, the answer is straightforward. The truth of the generalization is explained by appeal to a witness among the available objects—and this witness will remain available as ever more objects become available. By contrast, universal generalizations are problematic. Even though all the objects available today satisfy some generalization, a counterexample might arise tomorrow. How can we explain the truth of a universal generalization whose range includes objects yet to become available, all on the basis of material that *is* available? Generic explanations provide the answer. Such explanations are robust enough to ensure that the generalization will continue to hold no matter how the range of available objects expands.[10]

I believe non-instance-based generality has a number of other applications as well. Here I can only briefly mention some examples. First, this form of generality suggests a response to some (though not all) of (Fine, 2010)'s "puzzles of ground", namely (as in the fourth example above) to reject the idea that every universal generalization is grounded in each of its instances. Second, consider (Russell, 1908)'s famous "vicious circle principle", which

[10]In technical parlance: while universal generalizations are not always upwards absolute, those underwritten by a generic explanation are.

bans definitions that quantify over a totality to which the defined entity would belong. By quantifying over such a totality, Russell claims, a definition would illicitly presuppose what it set out to define. This claim loses its force, however, when we admit generalizations that do not presuppose each of their instances.[11] Third, non-instance-based generality enables us to retrieve a form of absolute generality on the basis of material that is accepted even by some theorists who deny the possibility of such generality. The key is to use the "systematic ambiguities" invoked by these theorists as generic explanations.[12]

2 Towards an account of intrinsic truthmaking

Having argued that there is a need for generic explanations, even on some robustly realist views, I turn to the task of explicating the distinction between such explanations and the more familiar instance-based ones. This explication will be based on a *truthmaker semantics*, which analyzes how states of the world explain the truth of a statement. The central notion is that of a state s *verifying* (or *making true*) a formula φ relative to a variable assignment σ; this relation is denoted $s \Vdash_\sigma \varphi$. For convenience, I write this as $s \Vdash \varphi(a_1, \ldots, a_n)$, where φ has free variables $v_1, \ldots, v_n$ and $\sigma(v_i) = a_i$. A *statement* will be a formula relative to a variable assignment.

My target idea is that a state s verifies a statement φ just in case material intrinsic to s suffices to explain φ, leaving no need to "look beyond" s to account for the truth of φ. The truth of φ is in this sense intrinsic to s. This is the right notion, I contend, to explicate the ideas with which we began. For example, if the future is metaphysically open, then everything that is now true must be made true by some state concerned solely with the past and the present. On this conception, it is natural to take truthmaking to be *monotonic*, in the sense that, when a state s verifies φ, so does any other state that is at least as informative as s.

For present purposes, we need not take a stand on how, exactly, the states are to be understood. The options parallel those in the metaphysics of possible worlds: the states can be parts of Lewisian worlds, (perhaps non-maximally specific) properties of the world, or (perhaps partial) stories about the world. Nor do we need to decide precisely what states there are; in particular, the generic verifiers can be taken to be essential truths in the sense of (Fine, 1994), higher-order generalizations of identity in the sense of (Rayo,

[11]For a development and defense of this idea, see (Linnebo, 2018b).

[12]See (Glanzberg, 2004) (who uses the Russellian term "typical ambiguity") and (Parsons, 2006).

2013) or (Dorr, 2016), or non-Humean laws of nature. The analysis I develop is neutral on all these questions. What matters for present purposes is how the states combine to explain the truth of various statements, and in particular, what it is for an explanation of a universal generalization to be instance-based, generic, or something in between.

The central concepts of my truthmaker semantics can be explained by means of examples. (A formal exposition can be found at the end of this section and in two appendices.) Consider a system consisting of just two balls, b_1 and b_2, and the following two states of the system:

b_1's being red

both b_1 and b_2's being red.

There is a natural ordering $\leq$ of the states in terms of what we may think of as their informational content. The monotonicity requirement, mentioned above, can now be expressed as follows:

$$\text{if } s_1 \leq s_2 \text{ and } s_1 \Vdash \varphi, \text{ then } s_2 \Vdash \varphi$$

We write $s < t$ for $s \leq t \wedge s \neq t$. For example, the first of the states mentioned above is strictly weaker than the second.

At the bottom of the ordering $\leq$ is the trivial state $\mathbb{0}$, which has no information whatsoever. At the top is the inconsistent state $\mathbb{1}$, which has "too much" information and therefore verifies every statement whatsoever. For example, $\mathbb{1}$ is identical with the state of ball 1's being all red and all green.[13] A state s is *atomic* iff every strictly weaker state is identical with the trivial state:

$$\forall t(t < s \rightarrow t = \mathbb{0})$$

A state s is *maximal consistent* iff every strictly stronger state is identical with the inconsistent state $\mathbb{1}$:

$$\forall t(s < t \rightarrow t = \mathbb{1})$$

Finally, there is a fusion operation on the states, $\sqcup$, which corresponds to combining pieces of information. This operation is always defined: the worst that can happen is that the fusion of two states is inconsistent.

Let us illustrate the concepts just introduced by means of some examples. We start

[13] Contrast the semantics of (Fine, 2014), where the state spaces must be permitted to contain an arbitrary finite number of distinct inconsistent states. In fact, on my approach even the single inconsistent state could be eliminated at the cost of adopting some slightly less elegant clauses for verification.

with some cases where the domain is fixed.

Example 1 Consider again a system consisting of two balls, b_1 and b_2, each of which can be either red or green. We thus have four atomic states, which we designate r_1, r_2, g_1, and g_2. The maximal consistent states are $r_1 \sqcup r_2$, $r_1 \sqcup g_2$, $g_1 \sqcup r_2$, and $g_1 \sqcup g_2$. Other fusions of distinct atomic states are inconsistent: $r_1 \sqcup g_1 = r_2 \sqcup g_2 = \mathbb{1}$. The state space can thus be represented as follows:

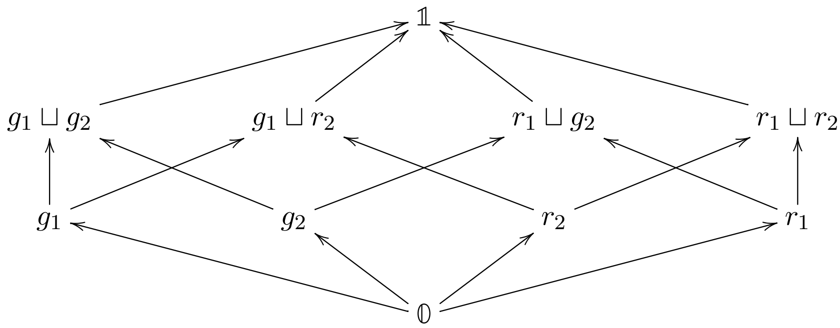

Notice also that the states support a natural notion of *aboutness*. For example, g_1 is about b_1, $g_1 \sqcup g_2$ is about both balls, and $\mathbb{0}$ is about nothing.

Example 2 Consider a system consisting of one ball, b_1, which can be green or either of two non-overlapping shades of red, namely crimson and scarlet. The atomic states are g_1 and r_1. The maximal consistent states are g_1, c_1, and s_1. We have $c_1 \sqcup r_1 = c_1$ but $c_1 \sqcup s_1 = \mathbb{1}$. This state space can be represented as follows:

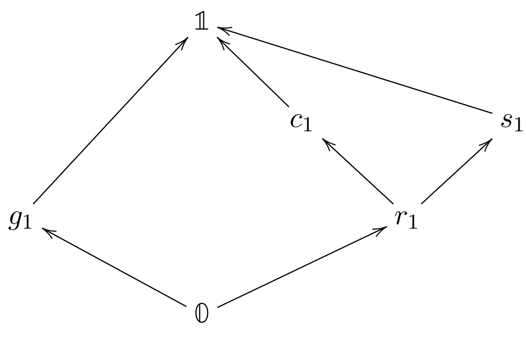

Example 3 Consider a system consisting of a countable infinity of balls, b_i for $i \in \omega$, each of which can be green or either of two non-overlapping shades of red, namely crimson and scarlet. The atomic states are g_i and r_i for each i. The maximal consistent states are obtained by choosing one of g_i, c_i, and s_i, for each i, and fusing all the chosen states.

Our next task is to use the truthmaking relation, $\Vdash$, to begin the analysis of the different forms of explanation of universal generalizations with which the article began. A more complete analysis will be provided in the next section. Again, my discussion will be informal and example-driven, although all of my claims will eventually be borne out by the formal details provided later in this section and in the appendices.

First, consider Example 1. The state of b_1 being red verifies the statement that b_1 is red, written 'Rb_1'; that is, $r_1 \Vdash Rb_1$. To explain Rb_1, there is no need to look beyond r_1. Likewise, it seems, $r_1 \sqcup r_2 \Vdash Rb_1 \wedge Rb_2$. What about the universal generalization $\forall x\, Rx$? A natural answer is:

(6) $$r_1 \sqcup r_2 \Vdash \forall x\, Rx$$

This yields a paradigm of a fully instance-based verification—under our present simplifying assumption of a fixed domain.

Some characteristics of instance-based truthmaking suggest themselves. The verifying state can be "factorized" as a fusion of simpler states, often such that each "factor" is only about a single instance of the generalization.[14] Thus, the fusion is about many instances. Furthermore, the verification has a low degree of modal robustness: were there to be more colored balls, $r_1 \sqcup r_2$ would no longer verify the generalization. (Of course, to examine this possibility properly, we would have to lift the simplifying assumption of a fixed domain.)

Next, consider Example 3. What verifies that, if an object is crimson, it is also red? This seems to require no information about the state of the system but rather to result from the natures of the properties involved. This suggests that $\mathbb{0} \Vdash Cb_i \to Rb_i$, for each i.[15] To explain each conditional, there is no need to look beyond $\mathbb{0}$. And this, in turn, suggests that $\mathbb{0}$ also verifies the corresponding universal generalization:

(7) $$\mathbb{0} \Vdash \forall x(Cx \to Rx)$$

This yields a paradigm of a generic verification.

[14] Why the qualification "often"? Consider a club whose members include both individuals and families, with no prohibition against an individual being a member of the club both directly and indirectly via her family. Say that a family is smart just in case each member of the family is smart. Consider now the generalization: every member of the club is smart. Assume the members of the club include both a family f and a member m of f. Then any state that verifies that f is smart will also verify that m is smart.

[15] Alternatively, the verifier could be a non-trivial state of crimson's being a shade of red. For present purposes, nothing hangs on this choice. What matters is that the verifier be a state that is generic in the sense that it is not about any of the balls in particular.

Again, some characteristics of generic truthmaking suggest themselves. In a case of fully generic truthmaking, the verifying state often cannot be factorized into simpler states.[16] Moreover, the verifier is not about any particular instance of the universal generalization. For this reason, the verifier enjoys a high degree of modal robustness. Even if more objects were introduced, this state would still verify the generalization.

Let us now make explicit the clauses concerning the verification of logically complex formulas on which we have tacitly relied. (These clauses are defended in the appendices.) Each of the connectives $\vee$, $\wedge$, and $\rightarrow$ is taken as a primitive; as are the quantifiers $\exists$ and $\forall$. Additionally, $\neg\varphi$ is defined as $\varphi \rightarrow \bot$, and $\varphi \leftrightarrow \psi$ as $(\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$. We relied on the following clause for conjunction:

$$(\Vdash \wedge) \qquad s \Vdash \varphi \wedge \psi \quad \text{iff} \quad s \Vdash \varphi \text{ and } s \Vdash \psi$$

Let us also adopt the analogous clause for disjunction:

$$(\Vdash \vee) \qquad s \Vdash \varphi \vee \psi \quad \text{iff} \quad s \Vdash \varphi \text{ or } s \Vdash \psi$$

Next, we adopt the following clause for the conditional.

$$(\Vdash\rightarrow) \qquad s \Vdash \varphi \rightarrow \psi \quad \text{iff} \quad \text{for each } t\text{, if } t \Vdash \varphi\text{, then } s \sqcup t \Vdash \psi$$

To motivate this clause, recall Example 1, where two balls can be either all red or all green. Now ask yourself which states verify the following formula:

$$(8) \qquad Rb_1 \rightarrow Rb_1 \wedge Rb_2$$

One minimal verifier of (8) is r_2. For b_2's being red ensures that, if b_1 too is red, then both balls are red.[17] Another minimal verifier is g_1. For b_1's being green excludes the antecedent and thus also verifies the conditional.

The treatment of negation requires some explanation. We begin with two clauses concerning the inconsistent state $\mathbb{1}$. This state verifies everything, and only this state verifies

[16] To understand the need for the qualification 'often', consider a three-membered residents' committee that decides by unanimous voting that all residents are required to wear a hat. The generic verifier of 'Every resident wears a hat' factorizes into three states concerned with each of the three votes. Thanks here to an anonymous referee.

[17] See (Yablo, 2014) and (Rumfitt, 2015).

falsum:[18]

$$(\Vdash \mathbb{1}) \qquad \mathbb{1} \Vdash \varphi \qquad \text{if } s \Vdash \bot\text{, then } s = \mathbb{1}$$

Next, notice that ($\Vdash\to$) implies:

$$(\Vdash \neg) \qquad s \Vdash \neg\varphi \quad \text{iff} \quad \forall t(t \Vdash \varphi \to s \sqcup t \Vdash \bot)$$

To make this more intuitive, say that s *excludes* φ iff there is no consistent extension of s that verifies φ. Thus, ($\Vdash \neg$) has an equivalent but more helpful formulation, namely that $s \Vdash \neg\varphi$ iff s excludes φ. As an illustration, consider Example 1 again and the formula $\neg Rb_1$. This has a single minimal verifier, namely g_1. For the minimal piece of information which ensures that b_1 is not red is that b_1 is green, which excludes its being red.

Finally, there are the two quantifiers. Let $D(s)$ be the domain associated with a state s; that is, the objects whose existence is verified by s. We adopt the following clauses:

$$(\Vdash \exists) \qquad s \Vdash \exists x\, \varphi(x) \quad \text{iff} \quad s \Vdash \varphi(a) \quad \text{for some } a \in D(s)$$

$$(\Vdash \forall) \qquad s \Vdash \forall x\, \varphi(x) \quad \text{iff} \quad s \sqcup t \Vdash \varphi(a) \quad \text{for every } t \text{ and every } a \in D(t)$$

That is, to verify an existential generalization is to verify the existence of an object a and that a is a witness to the generalization. Furthermore, to verify a universal generalization is to verify each instance *when fused with a state that verifies the existence of the object with which the instance is concerned.*[19] It would be unreasonable to require that, for a state s to verify a universal generalization $\forall x\, \varphi(x)$, the state s verify each instance $\varphi(a)$ completely on its own. Since s may "know" nothing about a, we may need information about what object a is. Such information is contained in any state t such that $t \Vdash Ea$. The reasonable requirement is therefore only that the fusion $s \sqcup t$ verify $\varphi(a)$. Intuitively, for you to have a

[18] It is noteworthy that only a single inconsistent state is needed (cf. fn. 13).

[19] Why not let the clause for the universal quantifier be the dual of that for the existential? The answer is that this would undermine the monotonicity of universal generalizations. Suppose s verifies the existence of certain objects and that each of these objects satisfies $\varphi(x)$. Then a stronger state might well verify the existence of further objects without verifying that these further objects satisfy $\varphi(x)$. A more promising alternative is to tweak ($\Vdash \exists$) in light of ($\Vdash \forall$) by letting $s \Vdash \exists x\, \varphi(x)$ iff $s \Vdash \neg\forall x \neg\varphi(x)$. Since this poses no threat to monotonicity, a semantics based on this alternative clause can be constructed and might well be useful for certain purposes. For our purposes, however, this alternative is not very natural. A state of the world up to and including a certain time verifies that there is a cat iff that chunk of the world in fact contains a cat (as per my preferred clause); it is not sufficient that that chunk exludes an entirely cat-free future (as per the alternative clause). When told that there exists a cat, it makes sense to demand to know where it is.

verifier for a universal generalization, it suffices that, for any object that I might give you, your verifier and the given object *together* verify that the corresponding instance. I cannot expect you to verify this instance until I have actually given you an object, about whose existence your verifier may be entirely “ignorant”.

With these clauses in place, it is now a routine—but useful—exercise to confirm the truthmaking claims (6) and (7) on which the preceding discussion was based.

3 Instance-based vs. generic verification

As announced, my overarching aim is to understand the intuitive distinction between instance-based and generic explanations of the truth of a universal generalization. To this end, I have developed a truthmaker semantics and used this to provide paradigm examples of the two forms of explanation. Our next task is to use the semantics to provide a systematic analysis, which subsumes and generalizes the ideas suggested by the paradigm examples of the previous section.

Recall our opening example:

(1) Every student of mine is born on a Monday

Suppose (1) is true because my students are $a_1, a_2, \ldots, a_m$ and because each of these objects is born on a Monday. Suppose s_i verifies Ma_i but doesn't verify Ma_j for any $j \neq i$. Clearly, $s_1 \sqcup \ldots \sqcup s_m$ verifies the conjunction $Ma_1 \wedge \ldots \wedge Ma_m$. However, this fused state fails to verify the generalization (1). As far as this state is concerned, I might have students distinct from $a_1, a_2, \ldots, a_m$ who are not born on a Monday. To verify (1), we additionally need to verify that the listed individuals are all of my students. Suppose there is a state t that verifies this:

$$t \Vdash \forall x(Sx \rightarrow x = a_1 \vee \ldots \vee x = a_m)$$

Then the larger fusion $s_1 \sqcup \ldots \sqcup s_m \sqcup t$, which adds t to all the instance-based information, verifies (1).

Let us say that a state is a *totality state* for the formula $\varphi(x)$ iff the state verifies that every instance of $\varphi(x)$ belongs to some fixed range of objects. (See Definition 3 in Appendix B for a formal definition.) Thus, in the above example, t is a totality state for the predicate S.

We are now ready to make some important observations. First, as mentioned, instance-based verification is a matter of degree. For example, the verifier $s_1 \sqcup \ldots \sqcup s_m \sqcup t$ of (1) is certainly instance-based, but not entirely so. In virtue of fusing each s_i, this verifier is about each student of mine. It may also be about me, my university, or some limited range of other objects. But the verifier is *not* about every object whatsoever! Because of the contribution made by the totality state t, only a limited range of objects are involved in the explanation of the truth of (1).

Second, note that our analysis incorporates *truthmaker necessitarianism*, defined as the view that truthmaking claims hold of necessity. For example, were there to be more objects than there actually are, the totality state t would still ensure that every student of mine is among the a_i's, each of whom would by the s_i's be born on a Monday.

The third observation concerns the possibility of purely instance-based explanations of universal generalizations. Is there a state s that verifies a generalization $\forall x\, \psi(x)$ without involving any form of generic verification or reliance on totality states for ψ or subformulas thereof? A negative answer follows from an interesting assumption (to which this article is not, however, committed), namely that for any state s there is an extension $s' \geq s$ with a strictly larger domain. If so, then for a state s to verify $\forall x\, \psi(x)$, it is not enough for s to verify that each of the objects that s is about satisfies $\psi(x)$; additionally, s needs to verify that any other object too satisfies $\psi(x)$. This means that s must "have an opinion" concerning objects that s is not about, which in turn means that the verification effected by s must be in part generic or involve some reliance on totality states.

Totality states are sometimes found problematic. They shouldn't be—at least not in the present context. For example, it is intrinsic to my office right now that it contains a single human being. Thus, the current state of my office contains not only the state of my being there but also the totality state of no one else being there. This totality state need not be understood as a mysterious absence of people other than myself. The state is better understood in terms of my office being circumscribed in a way that entitles us to add to a vast list of positive information concerning it that *that's it*: there is nothing more to say about my office. In fact, the same kind of completeness claim is implicit in our talk about possible worlds. Just like my office, a possible world cannot be fully described by a list of positive information but must be "closed off" by adding that there is nothing more to say about the world.

Some characteristics of instance-based and generic explanations were adumbrated in the previous section. These characteristics can now be properly stated and confirmed. First, instance-based verifiers have *low uniformity*. The verifier can be "factorized" as a fusion of totality states and particular states s_a, where each s_a is about an object a but typically not about any other member of the domain. Second, such verifiers typically have *high aboutness*: the verifier is about many objects. For example, the mentioned verifier of (1) is about each of my students. Third, instance-based verifiers have *low modal robustness.* Suppose the verifier is a fusion of some totality states t_i and some particular states s_a. Were there to be more objects—that is, given some different totality states T_i—then the fusion of the s_a's and the T_i's might no longer verify the generalization. As we observed, (1) provides an example.

Let us now consider the characteristics of generic verification. First, such verifiers have *high uniformity.* When the verifier is fully generic, it makes no reliance on totality states and cannot be "factorized" into simpler states. We saw this in the example where the trivial state $\mathbb{0}$ verifies the generalization that all crimson balls are red. Second, generic verifiers have *low aboutness.* A generic verifier need not be about any individual object, as illustrated by examples (2) and (3). Third, generic verifiers have *high modal robustness.* Since the verifier works without any reliance on totality states, it would still verify the generalization were there to be more objects. For example, the fact about color inclusion that explains the generalization that crimson balls are red works in the same way irrespective of what objects there are.

Having analyzed the difference between instance-based and generic explanation (or verification), I turn to two more specific questions that were posed in the introduction. The first of these concerns the availability of each type of explanation. Which true generalizations admit of which type of explanation? The answer suggested in Section 1 can now be confirmed and illuminated on the basis of the analysis just summarized. Each type of explanation has its distinctive characteristics, which potentially limit its scope. Let us begin with the easy case. Since generic explanations are uniform across all instances of the generalization, such explanations cannot be given for merely accidental generalizations, such as the student generalization, (1).

The scope of the more familiar instance-based type of explanation is a harder question. Such explanations require large fusions of states that are about a large number of objects; indeed, a *purely* instance-based explanation would involve a fusion of states that are about each and every member of the domain. Are such large fusions available? The answer will depend on our metaphysics. In Section 1, I described a family of metaphysical views on which the answer is negative. In a nutshell: if there isn't a definite domain of all objects, there isn't a definite range of all instances to consider for the purposes of given an entirely instance-based explanation.

Let me elaborate, using as an illustration the view that the future is metaphysically open. Suppose the domain of quantification includes not only past and present but also future objects. As time passes, more and more objects become available; but the available objects never exhaust the domain. Now, as ever more objects become available, ever more states too become available. For a state isn't available until all of the objects that the state is about are available. It follows that no state is ever about all objects; there will always be stronger states, which are additionally about objects that are not yet available. Any view with these kinds of implications will severely limit the scope for instance-based explanations. While we can fuse any states that are solely about the available objects, there is no such thing as a fusion of states that are about everything in the domain. Since many members of the domain aren't yet available, nor are any states that are about all these objects. Generic explanations, by contrast, are unaffected by these complications. So these explanations can pick up some of the slack where the instance-based explanations give out.

4 The logic of quantification

Our second more specific question concerns the logic of quantification. In the introduction, I promised a connection between the availability of the two forms of explanations and the correct logic of quantification. Let me redeem this promise.

According to Wittgenstein's *Tractatus*, a logical truth is 'sinnlos', that is, a truth with no content. In our present setting, the idea of having no content can be explicated in two different ways: either as being verified by the trivial state $\mathbb{0}$, or as being verified by every maximal consistent state. The first explication speaks for itself. If no information is required to verify φ, then φ is surely without content. For the second explication, recall that a state s is said to be maximal consistent just in case any stronger state is inconsistent: $\forall s'(s < s' \rightarrow s' = \mathbb{1})$. Intuitively, a state is maximal consistent just in case it is maximally opinionated. The maximal consistent states thus correspond, in our truthmaker semantics, to the possible worlds of the more traditional possible worlds semantics.

Our two explications of the Tractarian conception of logical truth suggests the following two logics of truthmaking.[20]

[20]To be entirely precise, the handling of free variables requires some care. Say that a state s *supports parameters in* Σ *relative to an assignment* σ iff, for every free variable v in Σ, $\sigma(v) \in D(s)$. Let '$\Sigma \models \varphi$' mean for every state space $\mathcal{S}$, every assignment σ, and every s from $\mathcal{S}$ that supports parameters in $\Sigma \cup \{\varphi\}$

Definition 1 (a) Let $\Sigma \models \varphi$ iff: for every state space $\mathcal{S}$ and every state $s \in S$, if s verifies every member of Σ, then s verifies φ as well.

(b) Let $\Sigma \models^* \varphi$ mean that, for every state space $\mathcal{S}$ and every maximal consistent $s \in S$, if s verifies each member of Σ, then s also verifies φ.

How do these two consequence relations compare with more familiar logics? In Appendix B, we prove the following proposition, which provides the answer.

Proposition 1 Consider the language of first-order logic. Let $\vdash_{\text{IL}}$ and $\vdash_{\text{CL}}$ represent deducibility in intuitionistic and classical logic, respectively. Then:

(a) $\Sigma \vdash_{\text{IL}} \varphi$ iff $\Sigma \models \varphi$. In particular, $\vdash_{\text{IL}} \varphi$ iff $\mathbb{0} \Vdash \varphi$ for every sentence φ and every state space $\mathcal{S}$.

(b) $\Sigma \vdash_{\text{CL}} \varphi$ iff $\Sigma \models^* \varphi$. In particular, $\vdash_{\text{CL}} \varphi$ iff $s \Vdash \varphi$ for every sentence φ and every state space $\mathcal{S}$ and every maximal consistent $s \in S$.

This proposition has a pleasing moral.[21] Depending on which explication of logical consequence we choose, the logical truths are either those of intuitionistic logic or those of classical logic. Thus, *both* logics have an important role to play.

The relative importance of these roles, however, will depend on our metaphysics. If our metaphysics is friendly towards maximal consistent states, it will often suffice to restrict our attention to such states—in which case we are entitled to classical logic. By contrast, if maximal consistent states are not always available, we are forced to take into account non-maximal states—in which case intuitionistic logic provides a fallback, robust enough to be available irrespective of what states there are. In fact, as we shall see in the next section, a stronger subclassical logic is often justified as well.

There is also a connection with the availability of instance-based explanations. The key is to observe that *a metaphysics is friendly towards instance-based explanations just in case it is friendly towards maximal consistent states.* Suppose the domain is definite,

relative to σ: if $s \Vdash_\sigma \theta$ for each $\theta \in \Sigma$, then $s \Vdash_\sigma \varphi$. The case of $\models^*$ is analogous. See (Troelstra and van Dalen, 1988, ch. 2, sect. 5.7) for the convention of restricting our attention to assignments that support parameters. This restricted can be dropped if we use *free* intuitionistic logic.

[21] Of course, the proposition depends on the clauses adopted in Section 2 concerning the verification of logically complex formulas. Interesting alternatives exist as well, especially on the approach that considers all states, as in the definition of $\models$. As is well known, classical logic too can be validated on this approach by replacing each formula with its double negation translation (Troelstra and van Dalen, 1988, ch.2 §3). Possibility semantics for classical logic, as in (Humberstone, 1981), can be seen as absorbing this translation into the semantic clauses. See (Rumfitt, 2015) and (Fine, 2017) for yet other options. In the appendices, I provide a preliminary defense of the naturalness of the clauses adopted here, relative to our purposes. A full defense will have to await another occasion.

such that all objects are available. Then there are no restrictions on the operation of fusing states, which in turn ensures that maximal consistent states are always available and instance-based explanations can always be given. Suppose instead that the domain is *in*definite, such that not all objects are available. Then the fusion operation must be restricted to states that are simultaneously available, which in turn restricts the pool of maximal consistent states and limits the scope for instance-based explanations.

5 Semi-intuitionistic logic

Our story could have ended here. In this optional, final section, however, I wish to press on and provide a philosophical defense of a *semi-intuitionistic logic*, which adds to intuitionistic logic a law to the effect that certain restricted generalizations behave classically and is thus intermediate in strength between intuitionistic and classical logic. Semi-intuitionistic logics are sometimes studied for their interesting technical properties,[22] but have so far lacked a proper philosophical defense.

Consider a metaphysical view on which instance-based explanations are not freely available. It might nevertheless be possible to say something informative about the extent to which such explanations are available—and a restricted form of classical quantification is thus legitimate. A good example is set-theoretic potentialism, which regards the hierarchy of sets as incompletable and for this reason denies that it is possible to give fully instance-based explanations of generalizations over all sets. But this form of potentialism regards each individual set as complete, which means that any true generalization restricted to a set admits of a fully instance-based explanation. In other words, when we restrict our attention to any given set—no matter how large—set-theoretic potentialists are in complete agreement with actualists; in particular, they agree that the logic of these restricted generalizations is classical.

Let us investigate, therefore, what it is for instance-based explanations to be available for a restricted generalization. We begin by clarifying the relation between a restricted generalization and its instances. Consider $\forall x(\varphi(x) \to \psi(x))$, which we often abbreviate as $(\forall x : \varphi(x))\psi(x)$. Assume that t is a totality state for the restricting clause $\varphi(x)$; for example, t verifies $\forall x(\varphi(x) \to \bigvee_i (x = a_i))$. Let us call the corresponding statements $\psi(a_i)$ *the critical instances* of the restricted generalization. For example, the critical instances of example (1) are the statements Ma_i to the effect that each of my students is born on

[22] See e.g. (Feferman, 2010).

a Monday. *What is the relation between verification of the restricted generalization and verification of all of its critical instances?*

As our discussion of (1) revealed, a state can verify all the critical instances without thereby verifying the restricted generalization. However, an equivalence between verification of a restricted generalization and verification of all of its critical instances can be shown to obtain *provided that the state in question extends a totality state for the restricting clause.*

Proposition 2 Suppose t is a totality state for φ, and let $s \geq t$. Then s verifies $(\forall x : \varphi(x))\psi(x)$ iff s verifies each of its critical instances. Moreover, s verifies

(9) $$(\forall x : \varphi(x))(\psi(x) \vee \neg\psi(x)) \rightarrow (\forall x : \varphi(x))\psi(x) \vee (\exists x : \varphi(x))\neg\psi(x)$$

The proposition too has an intuitive and pleasing moral. A totality state for φ ensures that all the φs are available, which in turn ensures that generalizations restricted to φ admits of instance-based verification. The availability of instance-based verification has two upshots. First, given a totality state t for φ, verification of any generalization restricted to φ comes to the same thing as verification of all of its critical instances. Second, generalizations restricted to φ behave classically, in the following sense. Suppose a state $s \geq t$ ensures that $\psi(x)$ behaves classically, that is, $\forall x(\psi(x) \vee \neg\psi(x))$. Then s ensures that the restricted generalization $(\forall x : \varphi(x))\psi(x)$ too behaves classically: either it is true or there exists a counterexample. In other words, given a totality state for φ, we know that any non-classical behavior of the restricted generalization cannot be blamed on the operation of quantification restricted to φ but must be traced back to the subformula $\psi(x)$.

In light of this discovery, it is natural to ask for which conditions φ there is a totality state. The answer will depend on one's metaphysical views. We can, however, describe an important class of conditions for which a totality state is always guaranteed, namely the condition of being one of some circumscribed objects. I shall explicate this using plural logic. So let us add plural variables (xx, yy, ...), each of which can take as its values any circumscribed objects. We also add two new logical predicate $\prec$ and E, where '$x \prec yy$' and 'Eyy' are read as "x is one of yy" and "yy exist", respectively.

What does it take for a state s to verify the existence of some objects aa? A plausible answer is that s must be *correctly and fully informed about which objects aa comprise.* We explicate this by stipulating that $s \Vdash Eaa$ iff the following conditions are satisfied:

(i) if $a \prec aa$, then $s \Vdash a \prec aa \wedge Ea$;

(ii) if $a \not\prec aa$ and $t \Vdash Ea$, then $s \sqcup t \Vdash a \not\prec aa$.

In other words, whenever a is one of aa, then s verifies that this is so and that a exists; and whenever a isn't one of aa, then s, supplemented with information about which object a is, verifies that that is so.

We now make a crucial observation. *Any state that verifies the existence of some objects aa also serves as a totality state for the condition '$x \prec aa$' of being one of these objects.* This ensures that Proposition 2 applies. Abbreviating '$\forall x(x \prec aa \rightarrow \psi(x))$' as '$(\forall x \prec aa)\psi(x)$', this proposition entails that for every state s that verifies the existence of aa, we have:

(i) s verifies $(\forall x \prec aa)\psi(x)$ iff s verifies each critical instance, that is, s verifies each $\psi(a)$ where $a \prec aa$;

(ii) s ensures that the restricted quantifier '$\forall x \prec aa$' behaves classically in the sense explained above.

Thus, quantification restricted to a plurality permits instance-based explanation and behaves classically. Quantification not so restricted need not permit instance-based explanation and can thus only be assumed to obey intuitionistic logic.

In fact, the trivial state $\mathbb{0}$ can be shown to verify all instances of (a version of) a principle known as *bounded omniscience*, namely:

(BOM) $$\forall yy\big((\forall x \prec yy)(\varphi(y) \vee \neg\varphi(y)) \rightarrow (\forall x \prec yy)\varphi(y) \vee (\exists x \prec yy)\neg\varphi(y)\big)$$

Suppose additionally that every atomic predicate behaves classically:

(At-LEM) $$\forall \bar{x}(P\bar{x} \vee \neg P\bar{x})$$

This principle can be seen as stating that the only source of non-classical behavior are the quantifiers. The logic obtained by adding these two principles to an intuitionistic logic of plurals is an instance of semi-intuitionistic logic. This is the appropriate logic for studying logicality in the sense of being verified by the trivial state $\mathbb{0}$.

All in all, the following picture emerges. When the domain is properly circumscribed, instance-based explanations are freely available, and consequently so is classically behaved quantification. Each circumscribed domain can be represented by a plurality of objects. Provided that the atomic predicates behave classically, these pluralities form classical islands in an intuitionistic ocean. In fact, this much should be granted by everyone. The proper locus of disagreement concerns the size and nature of the classical islands. Some theorists wish to drain the ocean entirely, leaving only dry land everywhere.[23] Others posit only finite classical islands in an infinite intuitionistic ocean. Set-theoretic potentialism, which I find attractive, calls for an intermediate position where the islands can be vast (namely of any set-sized cardinality) but the surrounding ocean is even greater (namely a proper class).[24, 25]

Appendices

A Sentential truthmaking

Let us adopt a language that allows (but does not require) the logic to be intuitionistic. There is a countable infinity, p_i for $i \in \omega$, of sentence letters. Additionally, we include $\bot$ as a sentential constant, which is false on any interpretation. As mentioned, the connectives $\vee$, $\wedge$, and $\rightarrow$ are taken as a primitive, whereas $\neg\varphi$ is defined as $\varphi \rightarrow \bot$, and $\varphi \leftrightarrow \psi$ as $(\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$.

We now provide a proper definition of a state space. As discussed, we have a set of states S and a partial order $\leq$. More precisely, our state spaces are upper join semi-lattices with minimal and maximal elements $\mathbb{0}$ and $\mathbb{1}$, respectively. This means that $x \sqcup y$ is the least upper bound of x and y in the partial order $\leq$. It is straightforward to prove that for every such structure we have $\mathbb{0} \sqcup s = s$, $\mathbb{1} \sqcup s = \mathbb{1}$, $s \sqcup s = s$, that $\sqcup$ is commutative and associative, and that we have the closure conditions:

- if $s_1, s_2 \leq t$, then $s_1 \sqcup s_2 \leq t$
- if $s \leq t_1$, then $s \leq t_1 \sqcup t_2$

Next, we need a function that encodes information about what each state $s \in S$ verifies. In the case of sentential logic, we do this by means of a two-place function $[\![\xi]\!]_\zeta$ such that, for each $s \in S$, $[\![\xi]\!]_s$ is a function from the set of sentence-letters to $\{0, 1\}$, subject to the

[23] More prosaically, these theorists accept the unrestricted plural comprehension scheme $\exists\varphi(x) \rightarrow \exists xx \forall y(y \prec xx \leftrightarrow \varphi(y))$, which ensures there is an all-encompassing plurality.

[24] As noted in the introduction, my metatheory throughout this article is classical ZFC. On the view in question, this can be justified by choosing—and then working within—an inaccessible-sized classical island.

[25] Thanks to Salvatore Florio, Kit Fine, Vera Flocke, Bob Hale, Jon Litland, Agustín Rayo, Sam Roberts, Ian Rumfitt, Benjamin Schnieder, Tim Williamson, as well as audiences in Bristol, Cambridge, Gothenburg, London, Oslo, Oxford, Tel Aviv, and Warsaw.

following monotonicity requirement:

$$\text{if } s \leq s' \text{ and } [\![p_i]\!]_s = 1, \text{ then } [\![p_i]\!]_{s'} = 1$$

Intuitively, $[\![p_i]\!]_s = 1$ (or 0) means that s does (or does not) verify p_i. Summing up, a state space for sentential logic is a triple $\mathcal{S} = \langle S, \leq, [\![\xi]\!]_\zeta \rangle$, subject to the mentioned requirements.

Finally, we need to define what it is for a state to *verify*, or *make true*, a logically complex formula. Since there are many options, we begin by laying down some constraints on the desired definition. Clearly, our definition must:

(i) cohere with our target idea that $s \Vdash \varphi$ iff material intrinsic to s suffices to explain φ.

In particular, this means that our definition must:

(ii) be monotonic; that is, if $s \leq s'$ and $s \Vdash \varphi$, then $s' \Vdash \varphi$.

After all, if s suffices to explain the truth of φ, then *a fortiori* any stronger state s' suffices as well.[26] Third, our definition should:

(iii) put the information intrinsic to the states to maximally good use.

This constraint will be illustrated shortly.

Fourth, our definition must:

(iv) be compatible with the standard Tarskian account of truth.

This requires some explanation. Observe that every Tarski model can be regarded as a state, namely the state of our system being just as the model describes. As we shall see, these states are highly opinionated: they settle every question concerning the objects in the domain of the model. By contrast, not every state corresponds to a Tarski model; for a state can fail to be maximally opinionated. Now, suppose that a Tarski model $\mathcal{M}$ corresponds to a state $s_\mathcal{M}$. Then our third requirement is that $s_\mathcal{M}$ verify all and only what is true in the ordinary Tarskian sense:

(T) $$s_\mathcal{M} \Vdash \varphi \quad \text{iff} \quad \mathcal{M} \models \varphi$$

Finally, our definition should:

[26] This contrasts with the semantics of (Fine, 2014), as well as most other accounts of truthmaking. However, this may reflect a difference in explanatory target more than a substantive disagreement.

(v) give rise to a workable logic of truthmaking.

This means that one or both of the natural logics of truthmaking set out in Definition 1 must be "workable". I offer no definition of this notion, and none is needed. On the analysis I develop, the correct logic of truthmaking is intuitionistic logic, whose workability is beyond doubt.[27]

The clauses governing the verification of logically complex formulas were presented in Section 2, namely:[28]

$(\Vdash \mathbb{1})$ $\qquad \mathbb{1} \Vdash \varphi \qquad$ if $s \Vdash \bot$, then $s = \mathbb{1}$

$(\Vdash \wedge)$ $\qquad s \Vdash \varphi \wedge \psi \quad \text{iff} \quad s \Vdash \varphi \text{ and } s \Vdash \psi$

$(\Vdash \vee)$ $\qquad s \Vdash \varphi \vee \psi \quad \text{iff} \quad s \Vdash \varphi \text{ or } s \Vdash \psi$

$(\Vdash\rightarrow)$ $\qquad s \Vdash \varphi \rightarrow \psi \quad \text{iff} \quad \text{for each } t, \text{ if } t \Vdash \varphi, \text{ then } s \sqcup t \Vdash \psi$

As noted, $(\Vdash\rightarrow)$ implies:

$(\Vdash \neg)$ $\qquad s \Vdash \neg\varphi \quad \text{iff} \quad \forall t(t \Vdash \varphi \rightarrow s \sqcup t \Vdash \bot)$

that is, $s \Vdash \neg\varphi$ iff s excludes φ, in the sense that there is no consistent extension of s that verifies φ.

We should accept these clauses and the analogous ones for the quantifiers, I contend, because they do well with respect to our constraints—indeed better than any alternatives. I will now use the constraints just articulated to provide at least a preliminary defense of this contention. (A fully worked-out defense would have to consider a range of different alternatives, so has to be left for future work.) Let me begin with constraint (i) of cohering with our target notion of truth making. The clause $(\Vdash \wedge)$ is unproblematic in this regard. A state s suffices for an explanation of a conjunction iff s suffices for an explanation of each conjunct. Moreover, since a disjunction is explained by explaining one of its disjuncts, $(\Vdash \vee)$ too is well motivated, given our target notion.

It might be objected that $(\Vdash\rightarrow)$ conflicts with our emphasis on truthmaking as intrinsic.[29] How can the truth of a conditional be intrinsic to a state s given that $(\Vdash\rightarrow)$ quantifies

[27] Many earlier investigations of the logic of truthmaking have either failed to offer a proper logic or offered weak logics whose workability is debatable. For example, (Restall, 1996) offers a logic close to first degree entailment of relevant logic.

[28] Because we assume monotonicity, these clauses are substantially simpler than those of (Fine, 2014).

[29] The clause for the universal quantifier adopted in the next section gives rise to an analogous complaint.

over states other than s? But this objection conflates *what* a state space represents with *how* it represents it. While $s \Vdash \varphi$ does indeed represent that φ is fully explained by material intrinsic to s, the state space often represents this fact in structural terms that involve states other than s. In this way, we gain important structural information about intrinsic truthmaking. An example just considered illustrates the point: $g_1 \Vdash \neg Rb_1$ iff g_1 has no consistent extension that verifies Rb_1. Although both sides of the biconditional are true, the left-hand truth stands entirely on its own: b_1 being green suffices, all by itself, to explain why it is not the case that b_1 is red. The right-hand claim about consistent extensions of g_1 *represents* that g_1 rules out the possibility of b_1 being red. But *what is represented* is that material intrinsic to g_1 suffices to explain $\neg Rb_1$; this is a claim about intrinsic truthmaking.

Closely related to constraint (i) is the monotonicity requirement (ii), which is precise enough to admit of a proof.

Proposition 3 (Monotonicity) If $s \Vdash \varphi$ and $s \leq t$, then $t \Vdash \varphi$.

The proof is an easy induction on syntactic complexity.

I turn now to constraint (iii), namely putting the information in a state to good use. The clause for the conditional does particularly well in this respect. To see this, suppose we had analyzed the conditional as material and thus adopted:[30]

$$s \Vdash \varphi \to \psi \quad \text{iff} \quad s \Vdash \neg\varphi \vee \psi$$

This alternative would gratuitously shift the set of verifiers of many conditionals towards the stronger end of the spectrum. Consider Example 1 of the two colored balls. On the material analysis of the conditional, the minimal verifiers of our example

(8) $$Rb_1 \to Rb_1 \wedge Rb_2$$

would be g_1, $r_1 \sqcup r_2$; in particular, r_2 would no longer be a verifier. But intuitively, the state of b_2 being red does verify the conditional. Or consider Example 2, which, we recall, involves a single ball which can be green, scarlet, or crimson. What verifies 'if the ball is crimson, then it is red'? On my analysis, $\mathbb{0}$ suffices. On the alternative analysis, the minimal verifiers would be r and g, which again would violate constraint (iii). To verify that an object is red if crimson, one need not take a stand on whether the object is red or

[30] Of course, ($\Vdash \neg$) would then have to be primitive, not derived.

green.

Turning now to constraint (iv), we begin with an important technical observation.

Proposition 4 Suppose s is maximal consistent. Then s verifies every instance of the Law of Excluded Middle.

Proof. Since every state verifies $\varphi \to \varphi$, it suffices to show that for any φ and ψ, we have $s \Vdash \varphi \to \psi$ iff $s \Vdash \neg\varphi \vee \psi$. Right-to-left is straightforward. For the other direction, assume the left-hand side. If $s \Vdash \psi$, we are done. So suppose $s \nVdash \psi$. We wish to show $s \Vdash \neg\varphi$. So assume $t \Vdash \varphi$. Since we have assumed $s \Vdash \varphi \to \psi$, we get $s \sqcup t \Vdash \psi$. Since s alone doesn't verify ψ, we must have $s < s \sqcup t$. Since s is maximal consistent, this means $s \sqcup t = \mathbb{1}$, which ensures $s \sqcup t \Vdash \bot$ and thus also $s \Vdash \neg\varphi$, as desired. $\dashv$.

On reflection, this result is unsurprising. Since a maximal consistent state s is maximally opinionated, it takes a stand on every φ: either $s \Vdash \varphi$ or $s \Vdash \neg\varphi$. Next, we need to explain what it is for a state to correspond to a Tarski model. A Tarski model for sentential logic is simply an assignment of a classical truth value to each sentence letter. A state $s_{\mathcal{M}}$ in a state space S therefore corresponds to a Tarski model $\mathcal{M}$ iff, for each $i \in \omega$, s verifies p_i or $\neg p_i$ in accordance with whether p_i is true or false in $\mathcal{M}$. Given this definition and Proposition 4, it is easy to prove that our definition of truthmaking is compatible with Tarskian truth, in accordance with requirement (iv), as explicated by equation (T).

Finally, requirement (v) is unquestionably satisfied in light of the following proposition and the next one, which are first steps towards Proposition 1. The proofs of the former propositions are accordingly subsumed under that of the latter.

Proposition 5 The sentential logic of truthmaking is intuitionistic. That is, $\Sigma \vdash_{\mathrm{IL}} \varphi$ iff $\Sigma \models \varphi$. In particular, we have $\vdash_{\mathrm{IL}} \varphi$ iff $\mathbb{0} \Vdash \varphi$ for every state space.

This is a pleasing result, since intuitionistic logic is without question a workable logic.

It would nonetheless be good to retain some role for classical logic. Proposition 4 points the way. Combined with Proposition 5, this entails that every maximal consistent state verifies every truth of classical logic. That is, the maximal consistency of a state ensures that the state behaves classically. We can now prove that the logic of maximal consistent truthmaking is classical.

Proposition 6 The sentential logic of maximal consistent truthmaking is classical. That is, $\Sigma \vdash_{\mathrm{CL}} \varphi$ iff $\Sigma \models^* \varphi$. In particular, we have $\vdash_{\mathrm{CL}} \varphi$ iff $s \Vdash \varphi$ for every state space $\mathcal{S}$ and every maximal consistent $s \in S$.

B First-order truthmaking

Let $\mathcal{L}$ be a first-order language. For simplicity, we assume that all function symbols have been eliminated in favor of predicates. We wish to extend our account of truthmaking to formulas of $\mathcal{L}$.

Our definition of a state space first needs to be adjusted to the richer language. As before, a state space $\mathcal{S}$ consists of a set of states S and a partial order $\leq$ on S. Additionally, the state space consists of a set D, which is its domain, and a two-place function $[\![\xi]\!]_\zeta$ specifying which atomic formulas are verified by each of its states. More precisely, for each $s \in S$, $[\![\xi]\!]_s$ is a function from the set of predicate letters to subsets of appropriate Cartesian products of D, such that an n-place predicate letter is mapped to a subset of D^n. As we shall see, $\langle a_1, \ldots, a_n \rangle \in [\![P]\!]_s$ represents that $s \Vdash P(a_1, \ldots, a_n)$. We impose the following monotonicity requirement:

$$\text{if } s \leq s', \text{ then } [\![P]\!]_s \subseteq [\![P]\!]_{s'}$$

for each predicate P of $\mathcal{L}$. Thus, when a state verifies an atomic statement, then any stronger state does so as well.

The identity predicate requires some comments. Let $D(s)$ be the field of $[\![=]\!]_s$; intuitively, this is the set of objects that s is *about*. Notice that the mentioned monotonicity requirement ensures the monotonicity of domains as well; that is, $s \leq s'$ implies $D(s) \subseteq D(s')$. Let 'Et' abbreviate '$t = t$'; this functions as an existence predicate. We require that D be the union of the sets $D(s)$ for every consistent $s \in S$.

Should we require $[\![P]\!]_s \subseteq D(s)^n$ for each n-place predicate P? That is, should we require that a state s only verify atomic statements involving objects that s is about? Perhaps. But a weaker, and more compelling, requirement suffices for our purposes. Suppose neither a nor a' is in $D(s)$. Then s must not discern between a and a'. Let us therefore require that, if the supposition holds, then

(Indisc) $$s \Vdash \varphi(a) \quad \Leftrightarrow \quad s \Vdash \varphi(a')$$

for any atomic φ. As is easy to see, this requirement and the definition of verification provided below entail that (Indisc) holds for every formula φ.

We now extend the clauses governing the verification of logically complex formulas, as previewed in Section 2. As before, we ignore use-mention distinctions and write $s \Vdash \varphi(a_0, \ldots, a_n)$ instead of $s \Vdash_\sigma \varphi(x_0, \ldots, x_n)$, where $\sigma(x_i) = a_i$ for each i. We have:

($\Vdash$At) $\quad s \Vdash Pa_1 \ldots a_n \quad$ iff $\quad \langle a_1, \ldots, a_n \rangle \in [\![P]\!]_s$

($\Vdash \exists$) $\quad s \Vdash \exists x\, \varphi(x) \quad$ iff $\quad s \Vdash \varphi(a) \quad$ for some $a \in D(s)$

($\Vdash \forall$) $\quad s \Vdash \forall x\, \varphi(x) \quad$ iff $\quad s \sqcup t \Vdash \varphi(a) \quad$ for every t and every $a \in D(t)$

These clauses do well with regard to our five constraints, as will now be shown.

The first clause merely explicates how we record information about the verification of atomic formulas. The clauses for the two quantifiers were explained and motivated in Section 2. It will be useful, though, to consider some examples. To make things formally precise, we need a definition that connects talk about truth in the ordinary Tarskian sense with our truthmaking semantics.

Definition 2 Let $\mathcal{M}$ be an $\mathcal{L}$-structure with domain D. Consider a state space $\mathcal{S}$ based on D. Then $s \in S$ is *an* $\mathcal{M}$*-state* iff any literal that s verifies is true in $\mathcal{M}$. Further, let us say that *an* $\mathcal{M}$*-state of* $\mathcal{S}$ *makes* φ *true* (symbolized $\mathcal{M} \Vdash \varphi$) iff there is an $\mathcal{M}$-state s in $\mathcal{S}$ such that $s \Vdash \varphi$.

Recall Example 3, which involves balls b_i for $i \in \omega$, each of which can be green or red, and, if red, either crimson or scarlet. So for each i, there are atomic states g_i and r_i. Let these states also be the minimal verifiers of Eb_i. Now, suppose one of the balls is green, say b_i. (Formally: consider a structure $\mathcal{M}$ according to which b_i is green.) Then the corresponding state g_i is a verifier for $\exists x\, Gx$ and for $\neg\forall x\, Rx$. Suppose instead that all the balls are red. (Formally: consider the associated structure $\mathcal{M}$.) If there is an infinite fusion $\bigsqcup_i r_i$, then this fusion is a verifier for $\forall x\, Rx$ and for $\neg\exists x\, Gx$. But what happens if the state space only permits finite fusions? Then the two mentioned formulas, which are true in the ordinary Tarskian sense, will not be verified by any $\mathcal{M}$-state. Indeed, even some truths of classical logic will lack a verifying $\mathcal{M}$-state; $\forall x\, Rx \vee \exists x\, \neg Rx$ provides an example. Some true universal generalizations are nevertheless verified, for example $\forall x(Cx \rightarrow Rx)$ and $\forall x(Rx \rightarrow Rx)$. Since each of their instances is made true by the trivial state $\mathbb{0}$, these generalizations too are made true by $\mathbb{0}$.

It is interesting to observe that $\mathbb{0} \Vdash \forall x\,(Rx \vee \neg Rx)$.[31] This illustrates how a state can verify a universal generalization without verifying even a single instance; for the minimal verifiers of the instance $Rb_i \vee \neg Rb_i$ are r_i and g_i. This phenomenon arises because, when a state verifies the existence of a ball, this state must take a stand on the ball's color, thus verifying the corresponding instance. By contrast, the verifier of the universal generalization (namely $\mathbb{0}$) has no opinion about the color of any individual ball.

Let us turn to the constraints that have precise technical content. As before, the monotonicity constraint can be proved to be satisfied.

Proposition 7 For any formula φ and states s and s' we have:

$$\text{if } s \leq s' \text{ and } s \Vdash \varphi, \text{ then } s' \Vdash \varphi$$

The proof is an easy induction on syntactic complexity.

Next, there is the constraint of compatibility with the Tarskian notion of truth. So let us compare this notion of truthmaking with the more familiar Tarskian notion of truth in a model. This requires a definition.

Definition 3 Consider a state s and a formula φ. Let *the extension of φ at s*, denoted $[\![\varphi]\!]_s$, be the set $\{\vec{a} \in D^n : s \Vdash \varphi(\vec{a})\}$ for appropriate n. Next, let us say that s is *a totality state for φ* iff s is consistent, $[\![\varphi]\!]_s \subseteq D(s)^n$, and $[\![\varphi]\!]_{s'} = [\![\varphi]\!]_s$ for every consistent $s' \geq s$. (Intuitively, a totality state for φ knows each member of the extension of φ and knows that these are all the members of the extension.)

It is easy to show that a state s is maximal consistent iff s is a totality state for 'Ex' and, for every atomic statement φ about the members of $D(s)$, the state s verifies either φ or $\neg\varphi$. As before, every Tarski model $\mathcal{M}$ corresponds to a maximal consistent state $s_{\mathcal{M}}$. Also as before, we can prove that $s_{\mathcal{M}}$ verifies φ iff $\mathcal{M}$ satisfies φ.

The final constraint calls for a workable logic of truthmaking. To show that this constraint is satisfied, we first recall the two logics of truthmaking introduced in Definition 1. Then, we prove that our soundness and completeness results extend to first-order logic—as already stated in Proposition 1, which I reproduce here for the reader's convenience.

Proposition 1 Consider the language of first-order logic. Then:

[31] To establish this, suppose that s verifies Eb_i. By our assumption concerning the verifiers of Eb_i, it follows that s verifies a color ascription to b_i and thus also the corresponding instance of the generalization, $Rb_i \vee \neg Rb_i$.

(a) $\Sigma \vdash_{\mathrm{IL}} \varphi$ iff $\Sigma \models \varphi$. In particular, $\vdash_{\mathrm{IL}} \varphi$ iff $\mathbb{0} \Vdash \varphi$ for every sentence φ and every state space $\mathcal{S}$.

(b) $\Sigma \vdash_{\mathrm{CL}} \varphi$ iff $\Sigma \models^* \varphi$. In particular, $\vdash_{\mathrm{CL}} \varphi$ iff $s \Vdash \varphi$ for every sentence φ and every state space $\mathcal{S}$ and every maximal consistent $s \in S$.

Proof. We begin with the soundness claim of case (a). For the sentential logic, we use the Hilbert-style axiomatization of (Moschovakis, 2015, §2.1). To prove soundness, we first prove that $\mathbb{0}$ verifies each axiom. Consider $\varphi \to (\psi \to \varphi)$. Assume $s \Vdash \varphi$. We need to show that s verifies $\psi \to \varphi$. Assume that $t \Vdash \psi$. We need to show that $s \sqcup t \Vdash \varphi$, which is straightforward. The other axioms are handled analogously.

It remains only to prove that logical consequence is closed under the rules of inference. First, assume we have truth-makers for φ and $\varphi \to \psi$. By fusing these, we obtain a truth-maker for ψ. Next, assume that IH holds for $\Sigma \vdash \varphi(x) \to \psi$, where '$x$' does not occur free in Σ or ψ. We want to show that it also holds for $\Sigma \vdash \exists x\, \varphi(x) \to \psi$. Assume $s \Vdash_\sigma \Sigma$ (i.e. $s \Vdash_\sigma \theta$ for every $\theta \in \Sigma$), where s supports parameters in $\Sigma \cup \{\exists x\, \varphi(x) \to \psi\}$ relative to σ. We want to show $s \Vdash_\sigma \exists x\, \varphi(x) \to \psi$. Assume $t \Vdash_\sigma \exists x\, \varphi(x)$. Then there is $a \in D(t)$ such that $t \Vdash_{\sigma[a/x]} \varphi(x)$, where $\sigma[a/x]$ differs from σ at most in assigning a to 'x'. Our assumption concerning 'x' ensures $s \Vdash_{\sigma[a/x]} \Sigma$. Observe that $s \sqcup t$ supports parameters in $\Sigma \cup \{\varphi(x) \to \psi\}$ relative to $\sigma[a/x]$. Thus, IH entails $s \sqcup t \Vdash_{\sigma[a/x]} \varphi(x) \to \psi$. So $s \sqcup t \Vdash_{\sigma[a/x]} \psi$ and hence also $s \sqcup t \Vdash_\sigma \psi$, as desired. The rule for universal generalization is handled analogously.

We now turn to the completeness claim of case (a). This claim follows from the completeness proofs of (Sambin, 1995), which show that every intuitionistically consistent set of formulas has a state space model, i.e. a model where $\leq$ is an upper join semi-lattice with minimal and maximal elements.

Finally, we consider case (b). This follows from case (a), using Proposition 4 and recalling that adding the Law of Excluded Middle to intuitionistic logic yields classical logic. $\dashv$

It remains only to prove *Proposition 2.* Throughout this proof, let t be a totality state for φ and $s \geq t$. First, assume $s \Vdash (\forall x : \varphi(x))\psi(x)$. We want to show that s verifies each critical instance. Assume $t \Vdash \varphi(a)$. Since t is a totality state for φ, we have $t \Vdash Ea$. Hence our main assumption ensures $s \Vdash \psi(a)$, as desired.

Next, assume that s verifies each critical instance. We want to show that s verifies the restricted generalization. Assume that $u \Vdash Ea$ and $u \Vdash \varphi(a)$. Since t is a totality state for φ, this entails $t \Vdash \varphi(a)$. Hence our main assumption ensures $s \Vdash \psi(a)$ and thus a fortiori $s \sqcup u \Vdash \psi(a)$, which proves the desired claim.

Finally, to prove $s \Vdash (9)$, assume $u \Vdash (\forall x : \varphi(x))(\psi(x) \vee \neg\psi(x))$. Then, for each a such that $s \Vdash \varphi(a)$, we have either $s \sqcup u \Vdash \psi(a)$ or $s \sqcup u \Vdash \neg\psi(a)$. Suppose that, for all of these a's, the first disjunct obtains. Then, by the first half of Proposition 2, it follows that $s \sqcup u \Vdash (\forall x : \varphi(x))\psi(x)$. Alternatively, suppose that for at least one of these a's, the second disjunct obtains. Then it is easy to prove that $s \sqcup u \Vdash (\exists x : \varphi(x))\neg\psi(x)$, as desired. $\dashv$

References

Benacerraf, P. and Putnam, H., editors (1983). *Philosophy of Mathematics: Selected Readings*, Cambridge. Cambridge University Press. Second edition.

Carnap, R. (1931). The logicist foundation of mathematics. In (Benacerraf and Putnam, 1983).

Dorr, C. (2016). To be F is to be G. *Philosophical Perspectives*, 30(1):39–134.

Dummett, M. (1991). *Frege: Philosophy of Mathematics.* Harvard University Press, Cambridge, MA.

Dummett, M. (1993). What is mathematics about? In *The Seas of Language*, pages 429–445. Oxford University Press.

Ewald, W. (1996). *From Kant to Hilbert: A Source Book in the Foundations of Mathematics*, volume 2. Oxford University Press, Oxford.

Feferman, S. Is the continuum hypothesis a definite mathematical problem? Unpublished manuscript dated 2011.

Feferman, S. (2010). On the strength of some semi-constructive theories. In Feferman, S. and Sieg, W., editors, *Proofs, Categories and Computations: Essays in Honor of Grigori Mints*, pages 109–129. College Publications, London.

Feferman, S. (2014). Logic, mathematics and conceptual structuralism. In Rush, P., editor, *The Metaphysics of Logic*, pages 72–92. Cambridge University Press, Cambridge.

Fine, K. (1994). Essence and Modality. In Tomberlin, J., editor, *Language and Logic*, volume 8 of *Philosophical Perspectives*, pages 1–16. Ridgeview, Atascadero.

Fine, K. (2010). Some puzzles of ground. *Notre Dame Journal of Formal Logic*, 51(1):97–118.

Fine, K. (2012). Guide to ground. In Correia, F. and Schnieder, B., editors, *Metaphysical Grounding*, pages 37–80. Cambridge University Press.

Fine, K. (2014). Truth-maker semantics for intuitionistic logic. *Journal of Philosophical Logic*, 43(2-3):549–577.

Fine, K. (2017). Truthmaker semantics. In Hale, B., Miller, A., and Wright, C., editors, *A Companion to the Philosophy of Language*, 2nd revised and expanded edition 22, pages 556–577. Blackwell, Oxford.

Glanzberg, M. (2004). Quantification and Realism. *Philosophy and Phenomenological Report*, 69:541–72.

Gödel, K. (1944). Russell's mathematical logic. In (Benacerraf and Putnam, 1983).

Hale, B. (2020). What makes true universal statements true? In Leech, J., editor, *Essays on Essence and Existence*, chapter 6, pages 104–123. Oxford University Press, Oxford.

Hilbert, D. (1926). Über das Unendliche. *Mathematische Annalen*, 95:161–190. Translated as "On the Infinite" in (van Heijenoort, 1967).

Humberstone, I. L. (1981). From worlds to possibilities. *Journal of Philosophical Logic*, 10(3):313–339.

Jané, I. (2010). Idealist and realist elements in Cantor's approach to set theory. *Philosophia Mathematica*, 18(2):193–226.

Kant, I. (1997). *Critique of Pure Reason*. Cambridge University Press. Transl. by Paul Guyer and Allen Wood.

Kaufmann, F. (1978). *The Infinite in Mathematics: Logico-Mathematical Writings*. Reidel, Dordrecht.

Linnebo, Ø. (2013). The potential hierarchy of sets. *Review of Symbolic Logic*, 6(2):205–228.

Linnebo, Ø. (2018a). Dummett on indefinite extensibility. *Philosophical Issues*, 28(1):196–220.

Linnebo, Ø. (2018b). On the permissibility of impredicative comprehension. In Fred-Rivera, I. and Leech, J., editors, *Being Necessary: Themes of Ontology and Modality from the Work of Bob Hale*, chapter 9, pages 170–187. Oxford University Press, Oxford.

Linnebo, Ø. and Shapiro, S. (2019). Actual and potential infinity. *Noûs*, 53(1):160–191.

Mancosu, P. (1998). *From Brouwer to Hilbert: The Debate on the Foundations of Mathematics in the 1920s*. Oxford University Press.

Moschovakis, J. (2015). Intuitionistic logic. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Spring 2015 edition.

Parsons, C. (2006). The problem of absolute universality. In Rayo, A. and Uzquiano, G., editors, *Absolute Generality*, pages 203–19. Oxford University Press, Oxford.

Putnam, H. (1987). *The Many Faces of Realism*. Open Court.

Rayo, A. (2013). *The Construction of Logical Space*. Oxford University Press, Oxford.

Restall, G. (1996). Truthmakers, entailment and necessity. *Australasian Journal of Philosophy*, 74(2):331–340.

Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In Hale, B. and Hoffmann, A., editors, *Modality: Metaphysics, Logic, and Epistemology*, pages 109–36. Oxford University Press.

Rumfitt, I. (2012). On a neglected path to intuitionism. *Topoi*, 31(1):101–109.

Rumfitt, I. (2015). *The Boundary Stones of Thought: An Essay in the Philosophy of Logic*. Oxford University Press.

Russell, B. (1908). Mathematical logic as based on a theory of types. *American Journal of Mathematics*, 30:222–262.

Sambin, G. (1995). Pretopologies and completeness proofs. *Journal of Symbolic Logic*, 60(3):861–878.

Schaffer, J. (2009). On what grounds what. In Manley, D., Chalmers, D. J., and Wasserman, R., editors, *Metametaphysics: New Essays on the Foundations of Ontology*, pages 347–383. Oxford University Press.

Troelstra, A. S. and van Dalen, D. (1988). *Constructivism in Mathematics: An Introduction*, volume 1. Elsevier, Amsterdam.

van Heijenoort, J., editor (1967). *From Frege to Gödel*, Cambridge, MA. Harvard University Press.

Weyl, H. (1921). Über die neue Grundlagenkrise der Mathematik. *Mathematische Zeitschrift*, 10(1–2):39–79. English translation in (Mancosu, 1998).

Wright, C. (2018). How high the sky? Rumfitt on the indeterminacy of the set-theoretic universe. *Philosophical Studies*, 175(8):2067–2078.

Yablo, S. (2014). *Aboutness*. Princeton University Press.