ISBN
9798191654065

Author
May, Ainsley

UNIVERSITY OF CALIFORNIA,
IRVINE

Meaning: Models, Modelling and Mathematics

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Philosophy

by

Ainsley May

Dissertation Committee:
Professor Toby Meadows, Chair
Professor Kai Wehmeier
Assistant Professor Ari Koslow

2026

DEDICATION

To the Palestinian people.

Contents

List of Figures

ACKNOWLEDGMENTS

My thanks first and foremost to my friend and advisor Toby Meadows, for seeing me through to the end. Special thanks also to Foz, Silas and collective Meadows' cats, for being a second home.

Additional thanks to the other members of my committee Kai Wehmeier and Ari Koslow. You are right about so many things and I wish it had taken me less time to learn them but I appreciate your encouragement while I tried.

The above sentiment also holds of my mother, Tanya Ringuet and of Ari Khoudary, who is a great person and knows more about modelling than I ever will.

More broadly, the LPS department tries its best to be supportive and has been subjected to far too many emails on my behalf.

Peers make or break graduate school and I was very lucky to spend time with Adam Chin, Tori Cotton, Margaret Farrell, Rena Goldstein, Rebecca Korf, Jeahyun Lee, Curtis Mason, Helen Meskhidze, Ellen Shi, Will Stafford and Darby Vickers. I appreciate all your work making our community welcoming.

My fellow logicians, Antoine Mercier, Brian Ransom and Jess Schirle.

My fellow neurodivergent organisers, Sayid Bnefsi, Dani Galvan, Stella Moon and Alex Vasquez.

The hosts of many excellent parties, most especially Bella Pepke for introducing me to lovely Arts friends.

Finally this dissertation would not have been possible without therapeutic, psychiatric and gender affirming care being accessible.

VITA

Ainsley May

Bachelor of Arts in Philosophy	**2016**
University of Queensland	*Brisbane, Australia*
Bachelor of Science in Mathematics	**2016**
University of Queensland	*Brisbane, Australia*
Teaching Assistant	**2016–2018**
University of Queensland	*Brisbane, Australia*
Master of Arts in Logic and Philosophy of Science	**2020**
Ludwig-Maximillian University	*Munich, Germany*
Teaching Assistant	**2021–2026**
University of California Irvine	*Irvine, California, USA*
Master of Arts in Philosophy	**2024**
University of California Irvine	*Irvine, California, USA*
Lead Instructor	**2025**
University of California Irvine	*Irvine, California, USA*
Doctor of Philosophy in Philosophy	**2026**
University of California Irvine	*Irvine, California, USA*

ABSTRACT OF THE DISSERTATION

Meaning: Models, Modelling and Mathematics

By

Ainsley May

Doctor of Philosophy in Philosophy

University of California, Irvine, 2026

Professor Toby Meadows, Chair

When do two theorems mean the same thing? Obviously the answer depends on how one understands meaning. In this dissertation, we approach meaning as a topic analysable into more theoretically refined concepts which serve as targets for scientific models. This "model-based" approach, which emphasises the informative role of models in enabling surrogative reasoning, is put forth in Chapter 1 and informs our assessment of accounts in the subsequent two chapters. In Chapter 2, we propose the folkloric account of meaning in mathematics, whereby the content of a mathematical statement is the class of logical models that make it true, with respect to a background class of models chosen to capture a particular mathematical setting of interest. In Chapter 3, we critique Soames' structured proposition account arguing with the use of examples that it fails to identify the differences which are most salient in mathematics.

CHAPTER 0

Introduction

Meaning is a ubiquitous and often hotly contested concept, in both our everyday lives and in academic settings. Consequently, many methodologies have been developed to help us discuss and analyse meanings, prominent among which is formal semantics. As the name suggests, formal semantics uses formal tools, such as those found in logic and mathematics, to capture linguistic features of natural and artificial languages.

This dissertation explores and applies a broadly naturalistic approach to formal semantics with a focus on meaning in mathematics, it consists of three projects each discussed in their own chapter. Chapter 1 sets the stage for the subsequent chapters by discussing the use of model-based inquiry in formal semantics. This approach where formal semantic accounts are intended as models of meaning (or its related concepts) underpins the analysis of the two accounts presented in later chapters. Chapter 2 provides a positive account of meaning in mathematics, called the folkloric account. Chapter 3 details Soames' structured proposition account and critiques its viability as an account of mathematical meaning in particular.

The original interest towards meaning in mathematics was motivated by the clear failures of possible world semantics when applied to mathematics. Infamously, such accounts, combined with the widely held view that mathematical truths are necessary, conclude that every mathematical statement is synonymous either with all tautologies or with all contradictions. Ultimately the intensionalists may have uncovered the truth about some grand metaphysical reality, but this is cold comfort for those who want to distinguish mathematical statements.

Recognising the difference between my goal which was more pragmatic and localised to mathematics in contrast to the metaphysical aims of more traditional accounts, led to isolating the broader conception of formal semantics that motivates this dissertation. At its core, this conception takes semantics to be engaged in model-based inquiry as is common in the sciences. We isolate this view of semantics, not to defend it directly (which would be a significant undertaking in its own right) but to better situate the reader with respect to our aims.

Unfortunately, although with good reason, the term 'model' has distinct referents relevant to our discussion; typically distinguished as scientific models and logical models. Scientific models consist of a carrier and an assignment. The carrier is any (physical or abstract) object or collection of objects and the assignment relates the carrier to the target which the model is intended to help the user reason about. On the other hand, a logical model consists of an interpretation over a collection of objects. If we're being strict then a logical model is always given with respect to a formal language, so the interpretation associates (collections of) objects with the symbols of this language. In this way logical models are scientific models of a formal language. Similarly, logical models can be scientific models of mathematical structures or the world. This is done by having the formal language represent parts of a mathematical, physical or conceptual domain. Often people speak directly of the functions, individuals and relations, e.g., addition, Da Vinci and believes, rather than the symbols naming them in the formal language.

Furthermore, logical models may be used as a part of a larger scientific model, so rather than looking at an individual logical model, a scientific model's carrier can contain any number of logical models. The most prominent case of this is the tradition of model-theoretic semantics, where statements are evaluated over different logical models. In modal logic, model-theoretic semantics are used to model different modalities. Most centrally to this dissertation, model-theoretic semantics are scientific models of natural languages. Rather than attempting to maintain consistent terminology through-out the entire dissertation (or specifying the kind of model every time), we use the terms which best fit the discussion at hand.[1]

The scientific use of model is most central to chapter 1 which is the only chapter to deal with modelling directly. In particular, chapter 1 outlines how model-based inquiry operates and how modelling can benefit investigations. This set-up allows us to discuss the fruitfulness of viewing the structured proposition account from (Cresswell 1985) as a model of semantic content. Much of this benefit comes from cross-model comparison, as we show how comparing Cresswell's model to other accounts facilitates learning about the concepts of interest and envisioning ways of building on his model.

Although modellers can be as dogmatic as any other kind of researcher, modelling often cultivates an attitude of open-mindedness, which manifests in a variety of ways including accepting provisional models and accepting multiple models. This openness often comes from recognising the limitations of specific models and the fundamentally constrained nature of

[1] Fear not, following which type of model is being referred to is far easier than the preceding paragraphs make it seem.

models. Models need not be perfect in order to be useful. As such provisional models are accepted as a necessary way-point; they provide us with information which helps direct further investigation and model building. Another factor contributing to this open-mindedness is that models are highly goal relative since they are evaluated with respect to their adequacy for purpose.

This open-minded attitude sets us up to embrace the folkloric account proposed in chapter 2, both in supporting a provisional model as a step in the right direction and the demand for an account better suited for mathematics specifically. On the folkloric account the content of a mathematical statement is the class of logical models that make it true, taken from a background class of models.[2] This account stays within the world-based or circumstantialist tradition which identifies meaning with unstructured sets of truth conditions, with two significant departures from the prominent intensional approach.[3] First, content is relative to a contextual base class of models (rather than a fixed set of metaphysically possible worlds) capturing the theory relativity of mathematical equivalence. Second, these models can correspond to all kinds of mathematical systems rather than just metaphysically possible worlds. This chapter applies the folkloric account to several cases and evaluates its strengths and weaknesses, ultimately concluding that it is a promising start which warrants further investigation.

Our positive evaluation of the folkloric account relies on the claim that other available accounts are inadequate. Although we briefly touch on the limitations of these accounts in chapter 2, chapter 3 allows us to explore one alternative, Soamesian structured propositions, in more depth. The account of structured propositions presented in (Soames 1987) shaped much of the subsequent literature. However, its formal details receive minimal attention as they are ancillary to the broader motivations and metaphysical project of the work. Chapter 3 goes over these details more carefully to allow clearer evaluation and future development of the account. Having presented the modelling essentials of the account, we assess its suitability as a model of mathematical meaning through several examples. I argue that Soames' account is not successful as it misidentifies the source of hyperintensionality salient to mathematics, placing too much emphasis on structure while retaining trivial truth conditions for mathematics. Furthermore, these structural distinctions produce judgements that are too fine-grained.

[2] Here, content is the theoretical refinement of meaning that the account targets.

[3] The name folkloric derives from the idea that the account is a "common sense" variation of an intensional account for specifically mathematical purposes.

So far we have spoken fairly broadly about semantics but in actuality this dissertation engages primarily with the truth conditional tradition of semantics. In order to situate the accounts discussed in this dissertation let's consider the following rough taxonomy of accounts of meaning (see figure 0.1). Already by focusing on meaning, we're omitting large swathes of linguistic investigation, such as those relating to language development. Additionally, this taxonomy is not intended to be complete and each approach contains multiple accounts which we don't name or differentiate.[4] Nonetheless it highlights key points of difference, focusing on those most relevant for distinguishing the accounts we're interested in.

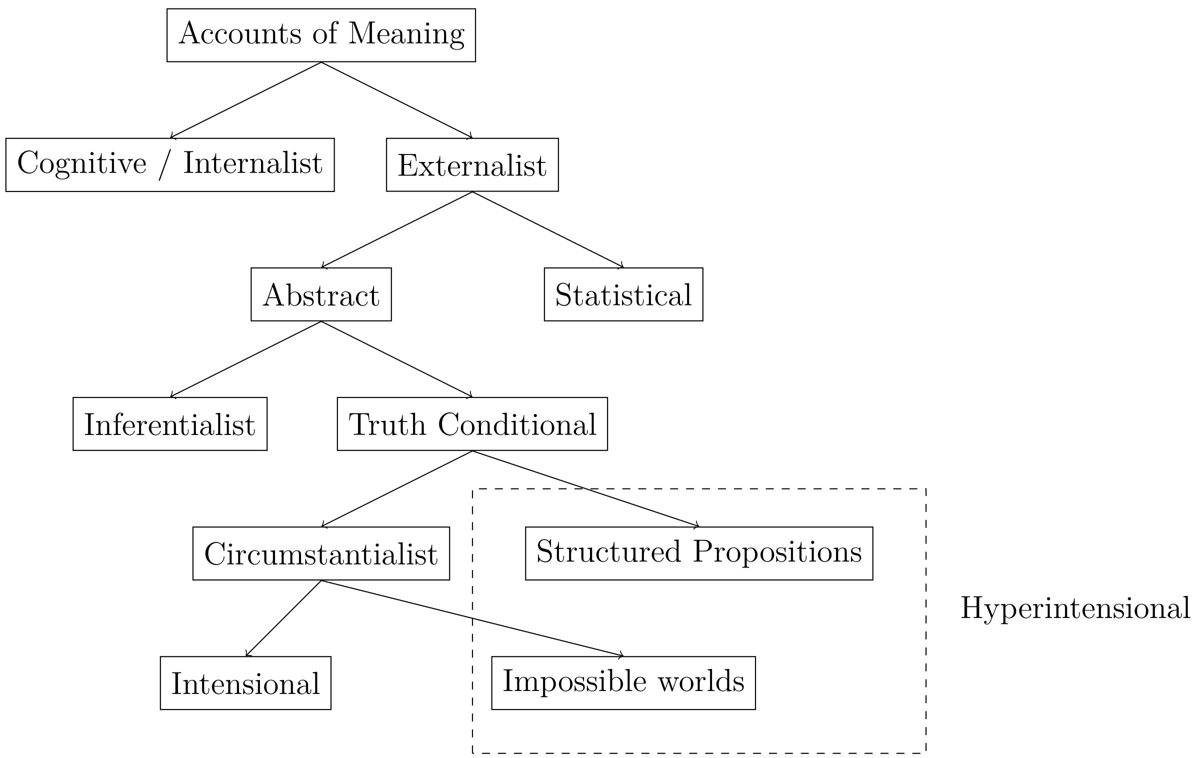

Figure 0.1: Taxonomy of different accounts of meaning

The first division is between a cognitive or internalist approach and an externalist approach. On the cognitive approach (see e.g. Jackendoff 1990, Fodor 1980, Chomsky 1980), meanings are mental representations. Alternatively for an externalist, meanings are outside the head and in the world. Within the externalist camp, some take a more statistical approach that deals with actual corpora, i.e., records of spoken or written language. This approach is more common in linguistics specifically whereas in philosophy, it is more common to posit specific abstract entities.

[4] For example, the account of (Davidson 1967) does not fit neatly into the taxonomy, being perhaps closest to a truth conditional account which is neither circumstantialist nor structured.

Truth conditional approaches and the formal implementations of inferentialism provide two kinds of proposals for these abstracta.[5] Inferentialist approaches hold that the meaning of a sentence is given by its inferential role, for example by the inferences it permits. This approach is prominently associated with (Brandom 2009) and is pursued formally in the field of proof theoretic semantics (see e.g., Schroeder-Heister 2018). Truth conditional approaches hold that the meaning of a sentence is given, in whole or in part, by the conditions under which the sentence is true. Within the truth conditional camp we can further divide into circumstantialist and structured proposition approaches.

Circumstantialist accounts (also known as model-theoretic accounts) hold that the meaning of a sentence is given wholly by the collection of circumstances in which it is true.[6] Accounts differ by which circumstances are considered, but in all cases circumstances can be formally represented by an interpreted structure, i.e., a logical model.[7] On the most prominent approach, circumstances are maximal, consistent, metaphysically possible worlds (see e.g. Carnap 1956; Lewis 1970; 1973; Kripke 1980).[8] Alternatively circumstances can be impossible worlds, which are interpreted structures but need not be maximal, consistent, nor metaphysically possible (see e.g. Nolan 1997, Berto and Jago 2019).

On the other hand, advocates of structured propositions hold that the semantic content of a sentence, sometimes called a complex, has a further internal configuration (see Soames 1987, Cresswell 1985, Moschovakis 2006 for some contemporary accounts). A complex is a recursive object that contains the semantic contents of the sentence's components (such as names, connectives, and predicates). The configuration of the complex records how the proposition's components contribute to the overall (meaning of the) proposition. Structured proposition accounts differ on their definition of semantic contents for different components, but common options include the component's intension, its extension or a propositional function.

Another important distinction is between intensional and hyperintensional truth conditional accounts. In our usage, intensional accounts are circumstantialist accounts that require the

[5]Sometimes inferentialist accounts are classified as a meta-semantic rather than semantic view. we include it here given its common role as a contrast to truth conditional approaches.

[6]We prefer the term *circumstantialist* over *model-theoretic* to avoid confusion with the mathematical discipline of model theory, which is used more extensively by the folkloric account in chapter 2. Not to mention our already extensive use of "model" terms.

[7]Some circumstantialists, identify circumstances with logical models and others, most notably Lewis, do not. Even in cases where circumstances are distinct from models, models are able to serve as a formal representation of circumstances.

[8]There are many other kinds of possible worlds, relating to different modalities, e.g., epistemic, deontic, and of course logical, but accounts of meaning primarily use metaphysical possibility.

circumstances to be metaphysically possible worlds. This includes possible world semantics and Carnap's method of intension. Then hyperintensional accounts are any truth conditional accounts which are more fine-grained than intensional accounts. This includes the folkloric account, impossible worlds and structured proposition accounts.

Meaning is too complex to be captured by a single model. However, this dissertation demonstrates the benefit of taking a modeller's perspective to semantic inquiry, especially concerning the demand for an account of mathematical meaning. This modeller's perspective allows us to compare models, in order to better understand the function and benefits of competing accounts. This in turn sets us up to build even more successful models. By accepting purpose relativity and focusing on mathematics in particular, the folkloric account is able to resolve issues with the intensional account while maintaining a circumstantialist core. This mathematical focus also leads to a rejection of Soames' structured proposition account. Yet the modeller's perspective allows us to appreciate the linguistic features captured by the Soamesian account as responding to different demands.

CHAPTER 1

Model-based inquiry in truth-conditional semantics

1.1 Introduction

(Yalcin 2018) and (Jackson 2020) both argue that semantics should be understood as model-based science. That is, that semantic accounts don't deal directly with meanings but rather provide us with a surrogate for reasoning about meaning. Although I am inclined to agree with this position, I do not attempt to defend it here. Instead the goal of this chapter is much more pragmatic, to highlight how model-based inquiry can contribute to semantic investigations.[1] In doing so, this chapter offers a foundation for exploring the neglected cornucopia of (formally and philosophically) interesting modelling work to be done in semantics, which has suffered from a disconnect between philosophical theorising and modelling in semantics.

We start by introducing model-based inquiry in semantics in section 1.2. First, we provide a basic introduction to models before describing the overall process of inquiry. Then we discuss how models can be informative in this process and what properties this requires of the models. Finally we connect our discussion back to semantics, by outlining Cresswell's structured proposition account as an instance of model-based inquiry in semantics.

Subsequently, section 1.3 demonstrates the benefits of comparison across models by contrasting Cresswell's model with other models containing intensions or structured propositions. These comparisons show how identifying points of common ground and difference can lead to better understanding of the models, their targets and the relationships between them including by directing further investigation.

[1]We prefer the term *model-based inquiry* over the more common *model-based science*, since modelling practices are not the sole purview of science. Similarly, our use of *models* refers to what are elsewhere called *scientific models* in contrast to *logical* models, such as the first order structures used in possible world accounts. Such structures often appear in semantic models, introducing further terminological strife into such discussions.

1.2 Model-based inquiry

In this section, we describe model-based inquiry and how it is used in semantics. Our account of modelling expands on the view of models as mediators, described in (Morrison and Morgan 1999), supplemented by more recent discussions of modelling and some novel terminology.[2] We start with a basic introduction to models and their role in enabling surrogative reasoning. Then in section 1.2.1, we zoom out to the broader inquiry process that modelling is embedded within. This allows us to discuss the desired properties of models, and how they contribute to inquiry, in section 1.2.2. Finally, we bring this all together in section 1.2.3 to give an instance of model-based inquiry in semantics, by discussing Cresswell's account of structured propositions.[3]

Characteristically, models provide a surrogate for part of the inquiry's subject, called the *target*.[4] This allows inferences to be drawn about the target on the basis of the behaviour of the model, through a process often called *surrogative reasoning*.[5] Sometimes this information can be read off directly from the model, for example using a world map to determine if a given country is landlocked. Other times, the models act as instruments that need to be manipulated in order to gain the desired information, such as running a computer simulation to determine if a hypothetical star system reaches a stable state.

Models are also partial, i.e., they do not capture everything about a target.[6] For example, a (standard Mercator) world map cannot be used to compare relative sizes of countries at different latitudes, at least not effectively. Thus models are assessed by their adequacy for purpose, i.e., whether they suit the particular goals of an inquiry. This partiality also leads to the possibility of multiple models for a given target, which may or may not be compatible with each other. Sometimes we can learn even more about the models, and thus their targets, by comparing them to each other, as we'll see later in this chapter.

[2](Yalcin 2018) also mentions the models as mediators account but he doesn't discuss it in much detail since he is more focused on the meta-semantic consequences of his view.

[3]Cresswell's account is chosen because it provides an interesting comparison to other accounts for our discussion in section 1.3.

[4]Cases of "targetless" models are outside the scope of this chapter. Note sometimes the term "surrogate model" refers to a specific class of models where the target is difficult or impossible to measure directly. We don't restrict ourselves to this class and instead use *surrogate* simply to name the part of the model corresponding to the target.

[5]We don't attempt to characterise the nature of the relationship between the target and its surrogate in the model. Often it is understood as representational, (Frigg and Nguyen 2020) provides a detailed account of this view. Whilst the representational view (especially one which allows for productive idealisation) can be a fruitful lens we aren't strongly committed to it.

[6]Colloquially, the map is not the territory.

Models consist of a *carrier* and an *assignment* which connects parts of the carrier to parts of the subject of inquiry.[7] For illustration, let us consider Fisher's reservoir model of a gold based economy, pictured with labels describing the assignment in figure 1.1 below.[8]

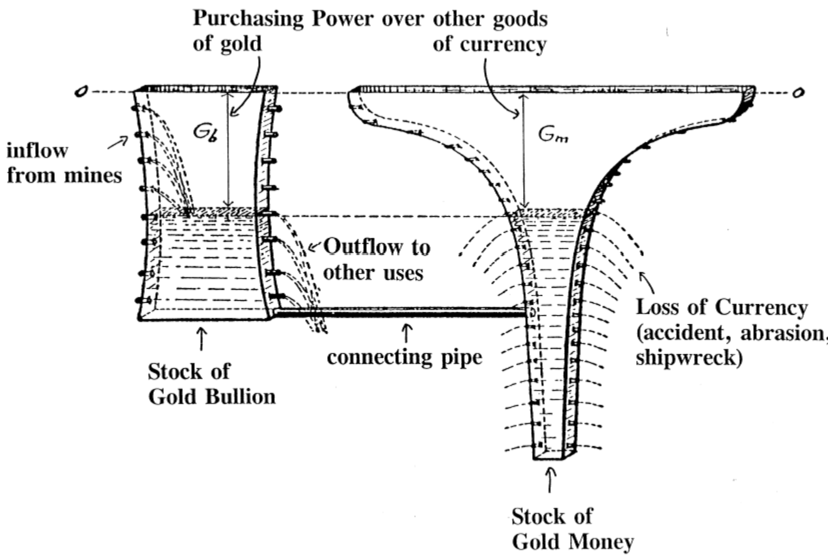

Figure 1.1: Fisher's connected reservoirs model with labels.

In this model, the carrier consists of two reservoirs filled with water connected by a pipe. One reservoir is assigned to the stock of gold bullion and the other the stock of gold currency/coin, with the connecting pipe allowing the transfer of bullion into currency and vice versa. In addition to the connecting pipe, each reservoir has holes allowing more outflow. For the bullion, this outflow is assigned to other uses of gold such as in jewellery and scientific applications. For currency, this outflow is assigned to currency removed from circulation. The distance from the top of the reservoir to the water level is assigned as the purchasing power of bullion and coin respectively and so on.

Assignments are typically neither total nor surjective, i.e., some parts of the carrier don't relate to any part of the subject and some parts of the subject don't relate to any part of the carrier. For example, the font on a rail map doesn't reflect anything about the actual railway system and the distance between stations is not captured by such a map.

[7] Other accounts of modelling attribute more to models. For example, (Weisberg 2015) requires that a carrier be equipped with a construal, which contains what he calls the model's scope and fidelity criteria in addition to the assignment. We prefer to stay minimal and associate fidelity criteria with the evaluation of the model, rather than the model itself.

[8] Taken from (Morgan 1999, Figure 12.3). This article contains detailed discussion of several of Fisher's models.

Assignments are what transform a carrier into a model, since it is what connects the carrier to the target, allowing it to act as a surrogate. This is essential since all kinds of things can serve as carriers of models, including both material and non-material/abstract objects. Mathematical objects are the most common type of non-material carriers but there are other examples, such as collections of possible worlds or even fictional objects.[9]

The models as mediators view highlights that models help us investigate theories and the world. In order to serve this mediating role, there needs to be a sufficient connection to theory and/or the phenomenon for the model to function as a surrogate yet sufficient independence that the model can still provide novel insights. Thus there is a bi-directional information flow between models, theory and the phenomenon. Roughly speaking, the theory and phenomenon provide constraints on the model and the model provides predictions about them, as depicted in figure 1.2.

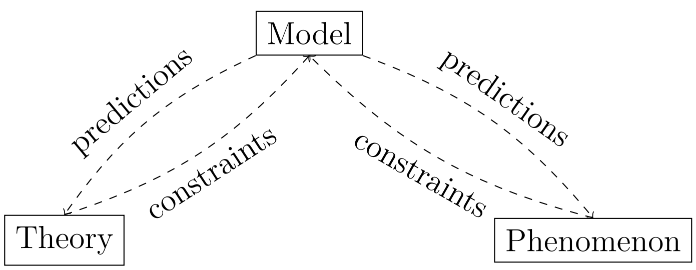

Figure 1.2: Information flow in modelling

	Y	y
Y	YY	Yy
y	Yy	yy

Figure 1.3: A Punnett square

Theory and phenomenon inform models by providing constraints used to construct and evaluate models. Often constraints are used in tandem, with information from both the world and our theoretical understanding of it being synthesised into a model. This synthesis can also draw on other sources such as narratives and of course modelling paradigms (Morrison and Morgan 1999). We can see both theory and phenomenon informing models in the case of Punnet squares, a model of Mendelian inheritance.

Punnett squares show the possible combinations of a pair of genotypes, see e.g. figure 1.3. Outside the grid are the parent's genotypes and each cell in the grid takes one allele (gene variant) from each parent, so the grid contains all the possible genotypes for the offspring. Each cell is equally likely to be the offspring's genotype so for our example given in figure 1.3, there is a 25% chance of the offspring having the YY genotype, 50% chance of it having Yy genotype and 25% chance of it having the yy genotype.

[9] For some philosophers, these kinds of non-material carriers may not be entirely distinct but that need not concern us here.

Now we need to appeal to the division between recessive and dominant alleles proposed by Mendel's theory, it tells us that any cell with a dominant allele will express the trait. This allows us to use the Punnett square to determine the predicted prevalence of a trait. Supposing the Y allele is dominant, our Punnett square captures a situation with a 75% chance of having the dominant trait and a 25% chance of having the recessive trait. Figuring out which allele is dominant is derived from observation of the phenomenon, such as noting which trait is more common.

Of course the reason we build models out of these constraints is because models can be informative. Sometimes they are informative in a direct manner, simply providing inquirers with new knowledge. Other times they are informative due to their role in evaluation, such as determining if the consequences of the theory line up with observations of the phenomenon. Models can also lead to novel hypotheses. If the model behaves in some particular way, then inquirers can check if this is borne out in the phenomenon or look for a theoretical basis of the behaviour.

To recap, models consist of some carrier object that is attributed an assignment so it can act as a surrogate for its target. This surrogate serves a mediating role whereby we use it to learn about the target and more readily relate to both theory and phenomenon.

1.2.1 Inquiry in more detail

In this section, we zoom out to consider the overall process of inquiry of which modelling is a part. In particular, we describe the different *aspects* of the subject of inquiry and how they relate to each other throughout an investigation involving concept-driven modelling.

In concept-/ theory-driven modelling the model's target is a theoretical concept, by contrast data-driven modelling takes a phenomenon directly as its target.[10] In this latter case, observations of the phenomenon produce data that forms the basis of the model and conceptual theorizing comes later to help explain the model and the phenomenon.[11] Concept driven modelling may also be prompted by some phenomenon, especially for more naturalistically inclined modellers.[12] This initial phenomenon informs the model, but is not the direct target, which only develops after some theorising about the phenomenon.

[10]This distinction between concept- and data-driven modelling is also described as being between top-down and bottom-up modelling respectively.

[11]For example, we had mathematical equations modelling the orbits of the planets, long before we had a theory of gravity to derive them from.

[12]Historically, this kind of concept driven modelling with an initial phenomenon of interest has been the paradigm case of modelling considered by philosophers of science.

Let us sketch the development of a model for inquirers engaged in concept driven modelling with an initial phenomenon of interest. The *phenomenon* is a "manifestation" of the subject of inquiry, typically some part of the world that we're interested in. Then some conceptual *theme* is abstracted from this phenomenon, this theme is often a pre-theoretic concept which serves as the starting point for further theoretical analysis. Next inquirers refine this theme to a more precise *theoretical concept* better suited for intensive investigation and theorising. Models are then constructed with this refined theoretical concept (or some part of it) as their target.

Since the target is understood in relation to the theme and phenomenon, the model also captures parts of them through "induced modelling relations." That is, the model also provides a less comprehensive surrogate for the theme and phenomenon and is correspondingly also evaluated with regard to them. Once we have a model, its properties are investigated leading to inferences about the concepts or phenomenon, with a focus on the theoretical concept as its direct target. Figure 1.4 depicts this process of inquiry instantiated for the investigation of inheritance.[13]

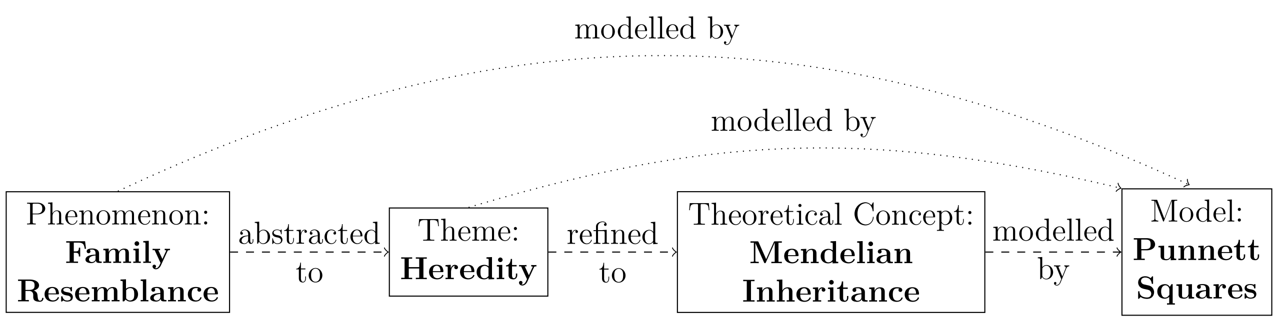

Figure 1.4: Sketch of investigation involving concept-driven modelling applied to a biological example

Let's review the investigation of inheritance as a reasonable scientific example of the process described. Inquiry starts with the phenomenon of family resemblance, i.e., that related plants and animals exhibit similarities in their traits. From this we abstract to a pre-theoretic notion of heredity, i.e., that something is passed on from parents to their offspring. This pre-theoretic notion got refined into the concept of Mendelian inheritance, i.e., that discrete genes determine the traits of individuals and that these genes are governed by the genes of the individual's ancestors. Then models such as Punnett squares and pedigree charts were developed.

[13] To preempt our later discussion, we take semantics to be interested in the pre-theoretic concept of meaning derived from language use especially synonymy judgements and refined/ theoretical concepts such as semantic content and truth conditions.

Such models are most useful when applied to individual genes and can only identify certain simple patterns of inheritance. Eventually, more complex gene interactions were discovered leading to a broader notion of genetic inheritance which subsumed Mendelian inheritance.[14]

Despite the limitations (or even inaccuracies) of models of Mendelian inheritance, they are still employed for handling simpler cases and in introducing topics of inheritance to students. This demonstrates that models are applied, not necessarily due to their correctness but due to their adequacy for purpose, such as whether the model warrants the kinds of inferences relevant to the current inquiry. Inquirers may have many interrelated purposes, including: identifying causal/explanatory structures, pedagogical goals, obtaining measurements or simulated data, testing theory consistency etc.

Adequacy often partially consists in how well the model "fits" the target but the purpose helps determine what exactly this fit consists in. For example (Weisberg 2015) identifies two classes of criteria for assessing fit, dynamical and representational fidelity criteria. Dynamical fidelity criteria govern how well the model's output should match observations of the target. Representational fidelity criteria govern how similar the model's internal structure should be to the structure of the target. Typically inquirers have both kinds of criteria but may prioritise them differently in certain contexts. Dynamical fidelity criteria are emphasised when the purpose relates more to simulating data whereas representational fidelity criteria are more important for causal/explanatory uses.

Each of the processes depicted in figure 1.4, abstracting, refining and modelling, are one to many. That is, several concepts may be abstracted from a given phenomenon, each of these concepts may be refined into many narrow topics, and each refined topic may be modelled in various ways. Investigations are generally open ended and ongoing, so inquirers may return to these processes multiple times, especially as new methodologies, concepts, modelling techniques etc. arise. Similarly, these processes are generally intertwined with each other and so need not proceed linearly in the order introduced above.[15]

Indeed, the refinement process is somewhat mutually constitutive. We start with a general investigation of some subject and as we learn we become better able to define (or dare I say, conceptualise) it. However, it is only by successfully identifying this refined understanding of our subject that we can say the learning was about *it* in particular. Much of this refinement is

[14] For simplicity, the narrative here reflects our contemporary perspective and so is somewhat anachronistic. More accurately, genetic inheritance originally referred to Mendel's theory and the distinction between them was only introduced once the limitations of Mendelian inheritance were identified.

[15] Thus the processes are actually many-to-many relations.

conceptual as connections are drawn and distinctions are made between the subject of inquiry and related concepts but it can also be informed by phenomenal evidence, with naturalistic inquiry giving more weight and importance to such evidence than non-naturalistic inquiry. For instance, the shift from Mendelian inheritance to genetic inheritance was partially due to Mendel's theory being incompatible with observations of more complex traits.

Thus the model contributes to this refinement as well because the development of post-Mendelian theories relied on having model based predictions that could be compared with the phenomenon. The inability for the theory to accommodate certain inheritance patterns was more clearly seen through failed model predictions than from the theory's conceptual commitments alone. We now turn to look more closely at how models can be informative.

1.2.2 Beneficial Properties of Models

In this section, we discuss different properties models can possess and how they assist surrogative reasoning. These properties include providing access to epistemic resources associated with the carrier, being more definite than theories, being able to be manipulated and being circumspect. In addition to their direct contribution to surrogative reasoning, these properties can also help inquirers evaluate and learn about the model itself. This is helpful, both in its own right and for allowing subsequent learning about the subject of investigation.

Providing greater epistemic access is the most obvious desideratum of models. If we were already able to do all the reasoning and have all the understanding we desire about our target then there would be no need to introduce models in the first place. By their very nature, models provide access to epistemic resources associated with the carrier, that is, access to tools for manipulating the carrier as well as existing knowledge from the carrier's field of study. The more powerful these tools, the more benefit gained from using that kind of carrier. Since mathematical resources are so strong, mathematical models are very common.[16] Alternatively, Fisher's reservoir model, mentioned earlier, makes use of fluid dynamics.

In order to apply these epistemic resources the model needs to be sufficiently definite. Again mathematical models are often the epitome for providing something definite in terms of

[16]We might characterise this as a virtuous cycle, where mathematical models are employed because they are so well understood and then because they are so ubiquitous they are studied even more.

precision, although there is likewise a lot to recommend the concreteness of material models. The specificity of the model also provides a cognitive benefit, by letting inquirers get a handle on the subject. These benefits are increased when the model is used to give specific examples or case studies. Furthermore, such examples provide a distinctively didactic advantage, for those trying to understand the model and its related subjects.

Combining the specificity of a model with the new epistemic resources allows us to make predictions with the model. These predictions are an important way that models provide a connection back to the phenomenon or the theory, as depicted in figure 1.2. This connection is essential to the models role as mediator. These predictions can contribute in various ways. They might be accepted as something we have learned about the subject of inquiry or one of its aspects, or they can be compared to other information (e.g. theoretical commitments or observations) to evaluate the success of the model.

Making predictions requires supplementing our assignment with a key. While an assignment connects components of the model to those of the target, the key allows us to convert values taken by these components.[17] For instance with a litmus test, the assignment specifies that the strip's colour represents that the pH of the tested solution falls within some range, but we need the key to tell us which range. Sometimes the key is easy to miss because it is so readily understood. Keys can be overlooked when they follow a well-known convention, like the top of a map being north (unless otherwise specified) or when it is simply identity, like for proportions of a Punnet square to a proportional likelihood.

It is important to consider the key and assignment separately, since keys need not cover every part of the assignment. For example, Fisher's reservoir gives the qualitative prediction that the system will reach equilibrium, i.e., that the purchasing power of gold bullion and gold currency will equalise, but we can't read off what price this will be at. This kind of quantitative prediction is not possible because although we have an assignment telling us that the price is given by the distance from the water level to the top of the respective reservoir (see figure 1.1), we don't have a key telling us how to convert this distance to a price.

Models can make more predictions when they are able to be manipulated. In the most limited case, manipulating a model means being able to change values taken by parts of the model or their configuration, so that we can see the consequences of these changes. These manipulations allow a model to apply to a broader range of situations. They are not

[17] The key is particularly important in the account of Frigg and Nguyen 2020.

restricted by some initial settings as a static object is. This allows models to function as instruments, allowing us to investigate a variety of situations.

In order to get the most benefit out of the manipulations, the changes need to correspond with something of interest within the target and something covered by the key. Changing the colour of the lines on the rail map doesn't produce new predictions so it isn't a fruitful manipulation.[18] Returning to Fisher's model, he also lacked a method for establishing the precise relationship between reserve stock and price which would determine the shape of the reservoir, showing that although we can manipulate this feature, doing so cannot provide us with much information (Morgan 1999).[19] This shows that the predictive power of models also relies on having the facts both to build the model with appropriate constraints and to provide it with the relevant inputs. Thus whether one has access to the information necessary for manipulating the model and obtaining predictions guides the kind of models one can construct and evaluate.

Similarly, the benefits of being manipulable are not unbounded since the model needs to conform to the constraints that allow it to act as a surrogate of its target. A model where every feature can be changed can no longer provide us with information about the target. For a less extreme case, we prefer not to be able to manipulate part of the model if it corresponds to some constant part of the target. If we cannot ensure this through model construction directly, we must stipulate that such manipulations are outside of the model's scope, i.e., its range of applicability. Other times inquirers will try to determine this scope by attempting to "break" a model. This involves testing to find the model's most extreme consequences and seeing which manipulations break this connection between model and target. Often this is done by applying the model to a problem case, such as one that the theory is known to struggle with or simply an uncommon variant.

Manipulating the model is especially helpful when it allows inquirers to identify the specific contribution of some model feature. For example, robustness checks determine whether or not the behaviour of computational models depends on the choice of parameters (the numerical values assigned to components of the model). Alternatively, we might compare more qualitatively different models, to see if both still capture an aspect of the target or if a feature possessed by only one model is somehow indispensable.[20] Especially in the case of

[18] Although recolouring the rail map may serve other purposes, like improving accessibility.

[19] The curved top of the currency reservoir in figure 1.1 reflects the hypothesis that as the purchasing power of coin decreases, the rate of this decrease also slows.

[20] For example, (Ripley 2012) argues that impossible world accounts can capture everything the structured proposition accounts can.

formal models, inquirers can make more "surgical" interventions where the minimal amount of other features are changed.

This modularity, which allows focus on specific features, also relates to the benefit models receive from being circumspect. It's okay if a model doesn't capture everything about its target or the broader subject of inquiry, so long as what it does capture is informative. Even if all things considered, an inquirer disagrees with some feature of the model (or its related theory), the model may still be the right tool for a particular job (e.g. where the opposing considerations are less salient). In such cases, the inquirer can still avail themselves of the model, in an instrumental fashion. Toy models, which employ radical simplifications to focus on a few key components, are a paragon example of this.

Similarly, inquirers can pull together multiple models to get a clearer understanding of a subject, which may not be possible through a single model alone (at least not at their current stage of investigation). Consider gas models which come in macroscopic or microscopic forms (among others), with the former capturing overall behaviour, like in the ideal gas law, and the latter focusing on particle behaviour, as in statistical mechanics. Both classes of models capture important properties and it is only by considering them together that one is able to fully understand the behaviour of gases.

Both cross-model comparison and instrumental uses of models demonstrate how models can serve as a common ground. This allows a more collaborative approach especially when combined with the circumspect stance that models are rarely definitive and instead can be continually improved and refined for different purposes. We explore this further in section 1.3 where we discuss the benefits of comparing different semantic accounts.

Models transform less tractable subjects into manipulable objects of inquiry, by providing a definite apparatus and utilising the epistemic resources of their carriers. Yet the predictive power granted by manipulability is bounded as models must remain tethered to their targets through appropriate constraints. Likewise models require relevant information to construct these constraints and specify keys for more detailed predictions. The circumspect character of individual models invites both their instrumental use in specific contexts and their combination for richer understanding.

1.2.3 An Instance of Model-based Inquiry in Semantics

In this section, we relate the preceding discussion about modelling to semantics specifically. We describe a particular research program in semantics in terms of model based inquiry, introducing Cresswell's account of structured proposition as a model produced by this program.[21] This sets us up to compare Cresswell's account with other approaches to exemplify different relationships between models in section 1.3.

There are many projects and programs in semantics, each with their own relationship to modelling. Indeed these projects span numerous disciplines including Logic, Philosophy of Language, Linguistics and often the intersections between them. For the purposes of this chapter we focus on the one summarised by figure 1.5.[22] Namely the project investigating the phenomenon of language use, which is abstracted into the theme of meaning refined to the theoretical concept of semantic content modelled by Cresswell's account of structured propositions.

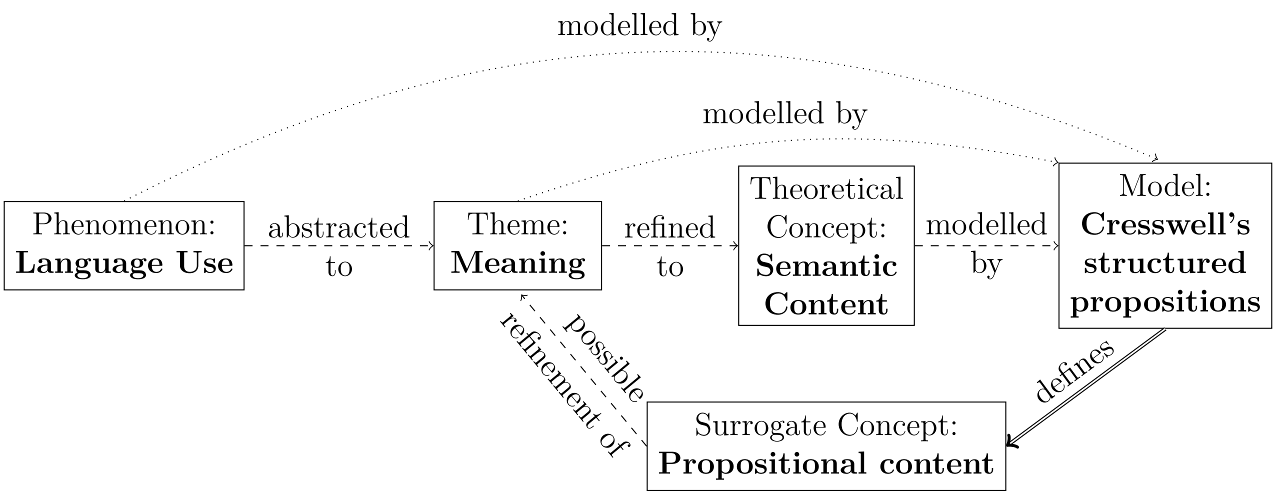

Figure 1.5: Sketch of conceptual relationships for an instance of semantic inquiry

Figure 1.5 also introduces a new aspect called the surrogate concept, which is defined as whatever concept is precisely realised by the model. In practice, this surrogate concept rarely has an independent existence, but naming it helps distinguish what the model actually defines from its target.[23] By investigating this realised concept, we learn more about what

[21] Here we're using structured proposition in the broad sense of any account where the internal composition of expressions explains differences between necessarily equivalent expressions. See below for the alternative classification of Cresswell's model as a structured intension account instead.

[22] This program is best understood as falling within philosophy of language if only because it is more logical than those in linguistics and more linguistic than those in logic.

[23] The surrogate concept is most likely to be attributed an independent existence or investigated in its own

the model is a good (in fact, perfect) model for. In turn, this can inform our theorising about the original model target, based on whether they accord with each other. For our purposes, we'll call the surrogate concept, propositional content, where this is understood to be distinct from semantic content.[24]

Recall that other than the double arrow for the model defining the surrogate concept, all the arrows are one-to-many, so as we move through the figure, there are alternative choices that could also lead to other projects. By taking meaning as the pre-theoretic concept, we distinguish this program from the purely logical uses of semantics towards the more linguistic side. However, the linguistic influence doesn't extend to interest in the "language faculty," especially in its more cognitive and developmental forms. Thus, the choice of language use as the phenomenon serves primarily as a reminder of the worldly and naturalistic aspect which initially prompted the inquiry and remains as a point of evaluation.

Let's now briefly introduce the structured proposition account of (Cresswell 1985) mentioned in the previous figure. We'll use this account in the next section to demonstrate different relationships between models, providing more details as they become relevant.[25] Recall that the intension of an expression is the function from the collection of metaphysically possible worlds to the expression's extension at that world.

Cresswell's account uses a language where expressions are sorted by types, built up from the basic types of e for entities/individuals and t for truth values.[26] So names belong to type e, nouns and intransitive verbs belong to type $e \to t$, adjectives and adverbs belong to type $e \times e \to t$, binary logical operators belong to type $t \times t \to t$ etc. These expressions are combined to build complex expressions and this process is recorded by placing the intensions of the basic expressions in a recursive configuration, i.e., the account's eponymous structure, often represented by nested tuples.

For example, we construct the semantic content of 'Lutruwita is an island and K'gari is sandy' as follows. First we combine the intensions of 'Lutruwita' and 'island' through $\overline{island}$

right in more formal or mathematical fields.

[24] In general the surrogate concept will be distinct from the model target, unless the map is in fact the territory.

[25] Those familiar with Cresswell's account may find its presentation here somewhat unusual as our discussion is organised so that the account (and possible alterations of it) is an illustrative point of comparison. Be assured that the discussion does eventually cover the essentials of the account.

[26] (Jackson 2020) calls this a Montagovian approach, due to Montague's role developing it (see e.g. (Montague 1973)). We depart from the notation used by Cresswell as we find ours more transparent, intuitive and likely to be familiar to the reader, although it is slightly less compact. Cresswell instead uses 0 for the type of entities, 1 for the type of propositions and / for $\to$.

taking in $\overline{Lutruwita}$ since 'island' is of type $e \to t$ and 'Lutruwita' is of type e. So we have $\langle \overline{island}, \langle \overline{Lutruwita} \rangle \rangle$. Similarly 'K'gari is sandy' becomes $\langle \overline{sandy}, \langle \overline{K'gari} \rangle \rangle$. Then these are combined to form their conjunction, 'and' is of type $(t \times t) \to t$ so it takes in our two sub-sentences. Thus we have $\langle \overline{and}, \langle \langle \overline{island}, \langle \overline{Lutruwita} \rangle \rangle, \langle \overline{sandy}, \langle \overline{K'gari} \rangle \rangle \rangle$ as the overall semantic content.

Most of the idealisation which occurs in this investigation is of an omissive kind, where the scope is successively restricted. For example, the abstraction from language use to meaning removes topics like the physiology of speech from the investigation. The refinement from meaning to semantic content is much more theory dependent, in our case it removes features like tone and also kinds of expressions like questions.

The formalisation which occurs at the modelling phase is much more distortive. Notably, using sets of discrete objects to capture the extension of nouns is incongruous when applied to mass nouns such as 'water.' Although it is often not necessary to literally provide extensions, such a level of detail is rarely suited to the force of the model, it is the discrete structure of such a extension itself which struggles to accommodate 'water'. We do have the option here to simply stipulate once again that such terms are out of scope, but often users accept the distortion instead.

The most significant benefit of Cresswell's model is in providing something definite to work with. It provides detail on how to understand structure, namely through grammatical categorisation, and a way to represent it. In doing so, it allows inquirers to construct explicit examples to further explore its consequences. The model also makes available mathematical resources by employing type theory and the recursive tuple structure.[27] The tuple structure is used even more effectively by another structured proposition account given in (Moschovakis 2006), which defines isomorphisms on this structure to capture more synonymous expressions.

1.3 Comparing Semantic Models

In this section, we compare Cresswell's model to the intensional model and to other structured proposition models. This comparison reflects some perspective based benefits of the modeller's outlook, for instance having circumspect goals allows us to look for (more incremental) improvements and common ground with other models. In both cases, we show how this

[27] We see more about the role type theory plays in Cresswell's account in section 1.3.2.2.

comparison leads to insight about the model itself and its relationship to these other models. In particular, both cases include friendly ways to use a model even if disagreements remain.

The comparison with the intensional accounts shows how distinguishing possible targets provides greater clarity about these concepts and how an alternative target can let other inquirers use the account. The comparison with other structured proposition accounts shows how features of a model can be repurposed to expand its scope of application and that one model can act as a surrogate for another. The benefit of these comparisons relies partially on specific features of our case, such as the relationship between truth conditions and semantic contents. Nonetheless, I suggest that model comparison will generally provide such benefits.

Before we continue, let us preepmt the objection that our discussion below is misguided because it mostly ignores the meta-semantic differences between the various accounts. This omission is intentional and captures part of the modeller's attitude. As (Yalcin 2018) points out natural language semantics is a relatively young field and as such it may be prudent to continue to let the field grow and diverse models develop before attempting to read off the structure of reality from them.

1.3.1 Structured Propositions vs Intensions

Structured propositions and intensions are often presented as fierce rivals, but in fact they share a lot of common ground. In this section, we see how comparing these models leads to greater understanding. In particular, attributing different targets to the accounts lets us identify similarities and differences both between the target concepts and the models.

1.3.1.1 Target Disagreement and Clarifying Concepts

It is relatively uncontroversial that intensional semantics provide an account, i.e., a model, of truth conditions.[28] The contentious claim, is that semantic contents also consist solely of these truth conditions. For those who support intensional accounts, intensions are a successful model of semantic contents precisely because they are a successful model of truth conditions. Others disagree, arguing that semantic contents cannot be reduced to truth conditions, usually pointing to counter-intuitive consequences of the intensional account. We might describe this as a disagreement about the target of intensional accounts.

[28]Recall, we are focused on the uses of intensional semantics most directly related to meaning. As a formal system, possible world semantics is also used in a variety of modal logics which may have different targets.

Strictly speaking then there are two models associated with the intensional account, one of which is a model solely of truth conditions and another that is a model of semantic contents. Typically when inquirers disagree about the target of a given model it is because there are competing models for (at least one of) the purported targets. In our case, structured propositions are a rival model of semantic content. So really there's two comparisons going on; comparing different models of a single target and comparing the targets to each other.

Having two models of a single target, e.g. semantic contents, gives inquirers a definitive point of comparison since the models make distinct claims about the target and its behaviour. Notably, structured proposition accounts hold that semantic contents have an internal structure. Identifying these contrast points, sets the stage for investigating deeper, so inquirers can look into which situations seem to require such structure. This benefit remains even if one ultimately rejects the alternative model, what matters is that the model provides specific claims about the target, which one can accept or deny. Likewise, many inquirers will have existing ideas about the disputed target but such ideas are improved and made more accessible when systematised into a model.

Given that one of the accounts has an alternative target, claims about how the disputed target behaves can be translated into claims about how the possible targets differ from one another, thereby clarifying them. That is, these claims help us understand what semantic content could be over and above truth conditions. This benefit is strengthened in our semantic case, by truth conditions being a relatively clear concept and having a straight forward role in intensional accounts. As such characterising intensional accounts as modelling truth conditions leads to a firm understanding of the model.[29]

This secure understanding provides a clear locus of difference by which to evaluate intensional accounts as a model of semantic contents. We must determine whether the intensional account is still promising or if its success was totally dependant on capturing truth conditions alone. Again this evaluation is supported by having a clear conception of semantic content and rival models are one way to obtain such clarity. Finally, by having an alternative target for the intensional account, we are able to keep what is appealing about that account, even if we reject it as a model of semantic contents as we explore below.

[29]It might seem counterintuitive that there is a benefit to simply identifying truth conditions as the target, when this claim is so uncontroversial. However, it is surprisingly easy to merge truth conditions and semantic contents together especially if the dominant paradigm supports this merger. But in doing so, we not only lose the possibility of exploring differences between them but also miss the stronger claim made by advocates of intensional accounts.

1.3.1.2 Friendly Uses of Competing Models I

Distinguishing different targets for the intensional account allows using it as a model of truth conditions despite significant disagreement about semantic contents. Such applications can occur either because the user is interested primarily in modelling truth conditions or because it forms a part of their distinct model for semantic contents (similarly for other targets). So by identifying truth conditions as an alternative target we retain what is fruitful in the intensional account and allow its use by more inquirers.

By restricting the target to truth conditions, even inferentialists can accept and use intensions. They simply have a larger theoretical disagreement since they hold that meaning/semantic content is best understood in terms of inference rather than truth conditions (which may ultimately limit the fruitfulness of the model to them). Although the model serves as a common ground, different inquirers may draw different conclusions based off it. Most obviously, inquirers may agree two expressions have the same truth conditions without concluding from this that they have the same semantic content. This difference highlights a benefit of separating the results of the model and the conclusions drawn from it. Thus providing a reason to more explicitly adopt a model-based perspective.

Advocates of the structured proposition approach go a step further and incorporate intensions (as a model of truth conditions) into their account of semantic content. As (Soames 1987, p. 63) puts it when discussing his own account "[Structured propositions] do not replace truth-supporting circumstances in a semantic theory; rather, they supplement them with a new kind of semantic value." Recognising this role that intensions (as sets of truth-supporting circumstances) can play in the structured proposition account, simultaneously shows a point of connection and of difference. They both accept intensions as a model of truth conditions but differ where structured propositions supplement them with structure. Distinguishing these targets lets us more readily describe this relationship between the models.

Identifying structured proposition accounts as building on intensional accounts rather than fundamentally rejecting them can also provide didactic benefits because intensions are the standard approach these days and building from this familiar territory can make learning structured propositions easier. In some cases this will also have a rhetorical benefit for those willing to weather the issues of intensional accounts, but who may become more convinced of its inadequacy once they are presented with a better approach.[30]

[30] On the other hand, the partial continuity between intensional and structured proposition accounts may be seen as a disadvantage, by those who think that the issues with intensional accounts are severe enough to warrant starting afresh. Thus making this identification instead a rhetorical tool for more radical alternatives,

1.3.2 Structured Intensions vs Structured Propositions

Cresswell's model hews more closely to the intensional account than other structured proposition models. In particular, Cresswell accepts the intensional account for statements without propositional attitudes and that each component of a structured proposition contributes its intension. In contrast, Soames' structured propositions can also contain extensions, because some lexical items, such as names, have their extension as their semantic content, not its intension. Consequently, Cresswell's account is sometimes called a structured intension account to distinguish it. In this section we relate Cresswell's model to other structured proposition accounts. First, we discuss how the common ground shared by structured intensions and structured propositions, allows the former to model the latter and the virtues that distinguish them. Second, we look at how parts of Cresswell's model might be repurposed by inquirers who disagree with his choice to limit the scope of hyperintensionality to propositional attitudes.

1.3.2.1 Friendly Uses of Competing Models II

For some purposes, structured intensions are a good surrogate for structured propositions. That is, Cresswell's account can be used as a model of other structured proposition models. In particular, if one is committed to the rigidity of names (and logical operators), then structured intension and structured proposition accounts with the same structural rules produce the same synonymy relation.[31]

If two accounts have the same structural rules, then they each give the same configuration to any expression. So they would only disagree if its components were not synonymous. However, synonymy of simple components comes down to their truth conditions and both a constant intension and an extension contribute the same truth conditions. As such if one's goal is to assess the synonymy claims of these accounts or to test the consequences of changing structural rules, then it suffices to use structured intensions.[32] In doing so, inquirers take advantage of the greater uniformity of structured intensions to learn about structured propositions (even if they ultimately prefer structured propositions).

Of course, structured intensions are also models of semantic content in their own right.

such as those proposed by inferentialists.

[31] Assuming the rigidity of logical operators may not seem substantive but it is worth noting that this need not hold in impossible worlds accounts.

[32] An example of changing the structural rules would be changing between the richer or more basic language.

Anyone with commitments against rigidity or towards Millian names will prefer structured intensions and structured propositions respectively. From an instrumental perspective though, adjudicating between structured intensions and structured propositions (given the rigidity commitments above) comes down to virtues.

The structured intension approach might be preferred for its uniformity and simplicity, either on the basis of these virtues themselves or because they lead to an easier implementation of the model. Whereas structured proposition models might be preferred for accurately representing the available degrees of freedom, i.e., only including a world parameter when it actually influences the valuation of that lexical component.[33]

Since there are multiple justifications for choosing between the models, it is helpful to communicate the actual justification, especially as justifications can relate to wholly different facets of the investigation.[34] Modellers and non-modellers alike are inclined to view the issue through their preferred lens, making it easy to misidentify the justification; either by incorrectly justifying the choice on the basis of their own preferences or through lacking familiarity with the possibilities on the "other side of the fence."[35]

Similar to the target disagreements, viewing structured intensions as a model of structured propositions highlights both the common ground and difference between the accounts. The difference really boils down to whether to treat the semantic contribution of names and logical connectives as extensional or intensional, a choice with both instrumental and ideological justifications. For certain purposes, this difference makes no impact, revealing a high level of commonality between the accounts. Nonetheless the differences between these two accounts make them appealing to different inquirers on the basis of their goals and theoretical commitments.

1.3.2.2 Scope Disagreement and Repurposing Models

Inquirers can also disagree about a model's scope of application. In semantics, these scope disagreements are often about the prevalence of a given linguistic feature since a model capturing said feature is needed whenever that feature arises. Likewise, models which don't

[33]It's not clear if this distinction between the semantic content literally being the extension or simply being best represented by its extension can be maintained, but minimally this shows different ways of expressing the judgement.

[34]Although no doubt disparate justifications for a given choice may coincide for some inquirers.

[35]Lest this diagnosis sound too critical, I admit I have been guilty of both mistakes.

capture this feature remain applicable in situations where it does not arise. Exploring both sides of a scope disagreement can be beneficial in the same way as target disagreements, such as by clarifying the stakes of the argument. It also provides an opportunity for inquirers to repurpose models or parts thereof.[36] In this section, we demonstrate what this process can look like by discussing attributes of Cresswell's account that could be repurposed for a broader class of situations and deserve further investigation.

The scope debate about hyperintensionality asks which situations require distinguishing necessarily equivalent statements. Cresswell's own position is that hyperintensionality primarily appears in relation with propositional attitudes, such as believes. Furthermore, he holds that statements without such attitudes are well treated by intensional accounts. Thus intensional models apply when statements contain no propositional attitudes but outside of this scope, a different model is needed, namely the one he proposes in (Cresswell 1985). Other advocates of hyperintensionality claim that it is much more prevalent and intensional accounts are in greater need of revision.

Recall that on Cresswell's account, each sentence is associated with a structured proposition, i.e., a configuration of nested tuples containing the intensions of the basic expressions found in the sentence.[37] The central mechanism Cresswell uses for propositional attitudes actually applies to 'that' rather than to the attitude directly.[38] 'That' is treated as a name forming operator, taking in a sentence's structured proposition and returning a name for it, i.e., a new expression of type e. The that-operator returns the same name if two sentences are built from intensionally equivalent components in the same configuration, i.e., if they have the same structured proposition. For example, since 'Tasmania' is another name for Lutruwita, $\langle \overline{island}, \langle \overline{Lutruwita} \rangle \rangle$ and $\langle \overline{island}, \langle \overline{Tasmania} \rangle \rangle$ are the same structured proposition.

The propositional attitude then is treated like any other transitive verb, mapping pairs of entities, in particular an agent and the name of a 'that'-clause, to a truth value. For an expression not containing propositional attitudes, the tuple configuration has no impact, reducing its semantic content to its intension.[39]

[36] Of course, there need not be a scope disagreement for repurposing to be beneficial. Indeed modelling techniques are often repurposed to subjects far outside their original application.

[37] Thus, Cresswell's account provides a clear example of intensions being incorporated as a model of truth conditions.

[38] Since 'that' is doing all the work, this account has the benefit of dealing with all propositional attitudes uniformly. However, this also prevents the account from applying to other *non-propositional* intensional attitudes, e.g., seeks.

[39] This is not the case in other structured proposition accounts.

Limiting the applicability of intensional accounts to situations not involving propositional attitudes is a position worth isolating in its own right, since it illustrates how a more explicitly modelling based methodology can discuss existing philosophical disagreements. Yet what is more interesting is the opportunities for those who disagree with Cresswell's position and think that hyperintensionality is more widely prevalent (or those who just prefer a uniform account).[40] Much of Cresswell's work can be repurposed to build an extended account where the configuration is always relevant to determining an expression's semantic content.

Such an extended account would retain two key differences from other structured proposition accounts (at least the Millian ones). The first difference is that each component contributes its intension, which we discuss in the following section. The second difference is the use of a richer language that reflects the typing of various grammatical categories whereas many structured proposition accounts only divide their non-logical vocabulary into individuals and predicates.[41] We discuss this difference to explore parts of Cresswell's model available for repurposing.

The simpler categorisation of terms is standard logical practice but given that an expression's configuration is so central to the structured proposition approach, it is surprising that more accounts are not interested in this further structural analysis.[42] This is especially surprising when structured proposition accounts already make important use of the lambda operator, which is a cornerstone of type theory.

To illustrate the impact of this richer language consider the statement 'Fozzy snored loudly.' Using either language, we have that 'snored loudly' is a predicate (specifically one of type $e \to t$). In the richer language, we can treat 'snored loudly' as a complex predicate obtained by applying the adverb 'loudly' (of type $(e \to t) \to (e \to t)$) to the intransitive verb 'snored' (of type $e \to t$) so the statement becomes $\langle\langle \mathsf{loudly}, \langle \mathsf{snored} \rangle\rangle, \langle \mathsf{Fozzy} \rangle\rangle$. In a more basic language 'snored loudly' has to be collapsed into a single predicate, so in the tuple notation the statement would be $\langle \mathsf{snoredloudly}, \langle \mathsf{Fozzy} \rangle\rangle$. Thus the richer language allows us to retain more of the statement's original structure.

[40] For example, mathematics is another commonly identified source of hyperintensionality outside of propositional attitudes. See chapter 3 for a critical assessment of structured propositions modelling the hyperintensionality in mathematical statements.

[41] Alternatively, (King 1995) reflects more grammatical structure by using the kind of syntactic trees found in linguistics. In this case, we still have a richer categorisation of the language but the typing is far less prominent.

[42] There are pragmatic reasons to only deal with an individual/predicate only language, such as this being more familiar to the average philosopher and the notational burden of type theories. So this decision is understandable, especially for early proposals, but I find the lack of subsequent work disappointing.

This richer language is common in some areas of linguistics, especially those in the Montagovian tradition. Indeed the choice of a richer or simpler language may partially reflect differences between more linguistic or metaphysical purposes. Inquirers with linguistic interests will obviously want to capture more detailed aspects of language, whereas those primarily engaged in a metaphysical inquiry only need a model to show that their view is coherent. Alternatively choosing the simpler language might also be based on a specific commitment, such as the primacy of properties over properties of properties (such as adverbs) or that only the logical configuration is relevant for distinguishing semantic contents. Regardless, the existence of models based on these different languages puts pressure on inquirers to justify their choices, regardless of where such justifications come from.

Cresswell's account also features his so-called Macro-structure Constraint which permits 'that' to be flexibly typed. This feature allows users to reverse the granularity introduced by the richer language. Since 'that' provides a name for the phrase it applies to, its input type is typically given by the phrase's components. With flexible typing, we can instead pay attention only to some macro-structure in the phrase's configuration, as in the following example, analysing 'Turing is an unmarried man.'

First, 'Turing' is of type e. Then as it stands, 'unmarried man' has two components, the adjective 'unmarried' of type $(e \to t) \to (e \to t)$ and the noun 'man' of type $e \to t$.[43] If we wanted to embed this within a propositional attitude, 'that' would normally take in this whole complex expression and return a singular entity, so it should be of type $(((e \to t) \to (e \to t)) \times (e \to t) \times e) \to e$. So our that-clause would be formalised as $\langle \mathsf{that}, \langle\langle \mathsf{unmarried}, \langle \mathsf{man}\rangle\rangle, \langle \mathsf{Turing}\rangle\rangle\rangle$.

Alternatively, we might consider the internal structure of 'unmarried man' to be irrelevant, and want to treat it as a simple expression. This expression is of type $e \to t$ since it is obtained by applying an adjective to a noun, i.e, an expression of type $(e \to t) \to (e \to t)$ to one of type $e \to t$. We can capture this by using a different 'that', the one of type $((e \to t) \times e) \to e$. Leading to the following formalisation $\langle \mathsf{that}^*, \langle\langle \mathsf{unmarried\text{-}man}\rangle, \langle \mathsf{Turing}\rangle\rangle\rangle$. This propositional fragment is the same as $\langle \mathsf{that}^*, \langle\langle \mathsf{bachelor}\rangle, \langle \mathsf{Turing}\rangle\rangle\rangle$, allowing us to model a situation where 'bachelor' and 'unmarried man' are synonymous.[44]

[43] We follow (Moschovakis 2006) in allowing adjectives to be type $(e \to t) \to (e \to t)$. Other accounts such as (Heim and Kratzer 1998) use predicate modification to ensure correct typing.

[44] One might be concerned that the flexible typing just results in the same formalisation as the basic language, but this is purely an artefact of using such a simple example. If instead we were dealing with 'handsome unmarried man', we could still keep 'handsome' and 'unmarried man' separate, unlike in the basic language where this would need to be a single predicate again (or rewritten as a conjunction).

Cresswell's discussion focuses on 'that' but in order to have flexible typing for 'that', the macro-structure must already be a well-typed configuration for the statement. Since Cresswell only requires hyperintensionality for propositional attitudes, this focus is understandable. For him 'that' is the gateway to hyperintensionality, allowing us to shift from (the otherwise adequate) intensional representations into more finely delineated objects. However, if one holds that the finer structure is always relevant then the need for such a gateway is obviated. Instead the tuple structure can do all of the work itself, allowing us to distinguish intensionally equivalent statements even when not embedded into a propositional attitude.

Although we lack the space here to detail such an account where the tuple configuration directly combines a richly categorised language with the flexibility of the Macrostructure Constraint, they are both interesting features which warrant further investigation. Additionally, their impact need not be limited to propositional attitudes and so there is an opportunity to repurpose these features and build on the important work done by Cresswell. This case considers just a single model among many which may provide fruitful features and tools, if only they would be reappropriated. Thus, there is no doubt much more work to be done and benefit to be gained from exploring semantic accounts through a more model-based perspective.

1.4 Conclusion

In this chapter, we explored the use of model-based inquiry in semantics. We started by introducing the fundamentals of model-based inquiry and how it is situated within larger investigations, including the benefits provided by models. This discussion allowed us to identify a research program culminating in Cresswell's model of structured intensions /propositions as a specific instance of model-based inquiry in semantics. Subsequently, we demonstrated that exploring connections between models can multiply their benefits, informing inquirers about the models and their targets. In particular Cresswell's model has several interesting connections to other models; it can serve as a competing account, a source of modelling infrastructure to repurpose and a surrogate for another account.

Treating Cresswell's account as an explicit contrast against intensional accounts revealed their common ground where intensions were retained as a model of truth conditions. Whilst also identifying structure as a potential differentiator between semantic content and truth conditions. Cresswell's account has features which may be useful outside of the scope he specifies for it, namely a rich Montagovian typed categorisation of language and the Macrostructure Constraint. These warrant further investigation in the spirit of repurposing models and their features.

In addition to providing its own account of semantic content, Cresswell's account can also act as a model of other structured proposition accounts. In doing so, inquirers dismiss ontological differences between the accounts in order to determine specific consequences common to both approaches. These comparisons borne from just a single research program are only the starting point to the benefits to be gained from model-based inquiry in semantics.

CHAPTER 2

Meaning in Mathematics: a folkloric account

Abstract The assumed necessary truth of mathematics poses well-documented issues for intensional accounts of meaning, which claim that all mathematical theorems are synonymous. Various hyperintensional accounts, such as those using impossible worlds or structured propositions, avoid this issue but the generality of such approaches also produces incredibly fine-grained accounts which make identifying the salient semantic features difficult. Between these issues, we lack an adequate account of mathematical meaning and I respond to this gap by proposing the folkloric account.

On the folkloric account the content of a mathematical theorem is the collection of (logical) models, within some reference class of models, that make the theorem true. The appeal of this account is partly that it retains central aspects of world-based accounts, such as evaluation within a model. Yet it overcomes their limitations by incorporating more models to represent different mathematical theories and structures without allowing all such structures. Furthermore, by varying the reference class of models, the folkloric account is able to capture the diversity of practice within mathematics. In this chapter, I introduce the folkloric account and use examples to identify its strengths and weaknesses including avenues for further development.

Consider any two mathematical theorems, for concreteness let's say the well-ordering principle of the reals and the Dedekind completeness of the reals. The former says there is an ordering relation on the real numbers with no infinite descending chain. The latter says that every bounded set of reals has a least upper bound. These statements seem to mean different things. However, on the most prominent account of meaning, possible worlds semantics, all mathematical theorems have the same meaning because they are necessarily true. The problem of all mathematical theorems having the same meaning is shared by other so-called 'intensional' accounts.

The judgement that all theorems have the same meaning is puzzling given how readily we distinguish between theorems like those above. Certainly any such account is unhelpful for considering mathematical meaning in particular. Trivialising mathematical meaning is a well-known issue of intensional accounts and has motivated several hyperintensional accounts. However, we'll see in section 2.1, that existing hyperintensional accounts are overly fine-grained, making distinctions that are not relevant for mathematical use.

This chapter presents a preliminary account of mathematical meaning, which I call the folkloric account. This account succeeds in the minimal desideratum of distinguishing the meanings of different theorems while still allowing distinct formulations of theorems to have the same meaning. The account combines the ideas that a statement's meaning is given by its truth conditions and that in mathematics these truth conditions are best represented by collections of interpreted structures (i.e. logical models). Roughly, then the meaning of a theorem is the collection of such structures which make the theorem true. Furthermore, the folkloric account allows relativising to different initial collections of structures that capture different informal mathematical frameworks. I suggest that this account is 'folkloric' because it adapts a minimal truth-conditional account in the obvious way to describe mathematics. However, this kind of 'folk-account' doesn't appear to have been seriously worked out in the literature. In this chapter, I give a precise account based off this folk idea and investigate the account's virtues and limitations. Overall, I argue that despite the remaining challenges, the folkloric account is promising and warrants further investigation

This evaluation proceeds from the view that formal semantics provides models (in the scientific sense) of meaning.[1] That is, formal semantics provides an idealised representation of meaning (or an aspect thereof) for various epistemic purposes, usually collectively known as 'surrogative reasoning.'[2] As such the claims in this chapter do not relate to some "thick" notion of meaning/propositions and should not be read as such. Furthermore, models are implemented and evaluated with respect to the aims of their users, which is reflected through the relativity of the folkloric account allowing different models of meaning for different mathematical contexts. This intentionality also informs my evaluation of the folkloric account especially in contrast to other existing accounts that are less attuned to mathematical meaning, and hence their inadequacy results from serving a different purpose.

This chapter is structured as follows. In section 2.1, I identify the challenges facing existing accounts of meaning, which motivate the folkloric account. In section 2.2, I introduce the folkloric account and apply it to a few cases. In section 2.3, explore and defend the relativity of the folkloric account. I reject alternatives that emphasise an intended interpretation or

[1]Note that model is only used in the scientific sense for this paragraph and later we revert to using it in the logical sense most common to formal semantics. See Yalcin 2018 and Jackson 2020 for two versions of this view.

[2]Subsequently, we use 'content' to refer to the technical notion introduced in the folkloric account, keeping meaning for the informal notion.

consolidate the approach into a single function and showing how the relativity is beneficial and even essential for reflecting mathematical practice. Section 2.4 concludes and describes future directions.

2.1 Existing Semantic Accounts

The folkloric account is motivated by the failure of existing semantic accounts when applied to mathematics. In this section, I briefly outline the main challenges with existing accounts. We start with intensional accounts before moving to impossible worlds and structured propositions as two prominent approaches to hyperintensional accounts.[3] As mentioned the intensional accounts are too coarse-grained for fruitful application to mathematics. On the other hand, existing hyperintensional accounts are too fine-grained, recording distinctions which are not typically relevant in mathematical settings.

The fact that all mathematical statements have trivial intensions is well-known and widely acknowledged by both friends (Lewis 1970, Carnap 1956) and foes (Soames 1987, Ripley 2012) of intensional accounts, so we discuss it only briefly here. This triviality follows from the fact that i) a statement's intension corresponds to the set of metaphysically possible worlds (or models thereof) which make it true and ii) mathematics is often presumed to be metaphysically necessary.[4] Consequently, true mathematical statements are necessarily true and hence their intension is the set of all metaphysically possible worlds. Similarly, the intension of all false mathematical statements is the empty set, since they are not true at any metaphysically possible world. This allows little room for the "meaningfulness" of mathematical statements.

This triviality doesn't only impact the meaning of individual statements, it can be exacerbated by conditionals. Since any conditional statement with a true consequent is true, any conditional with a theorem as a consequent is itself necessarily true, no matter what the antecedent is. Thus such accounts are forced to conclude that the following statement is true.

(1) If $\mathbb{R}$ is Dedekind complete then there exists a well-ordering of $\mathbb{R}$.

[3]Recall that all these accounts fall within the truth-conditional approach from our taxonomy in figure 0.1.

[4]As this chapter is not interested in a thick notion of meaning, my definition of intension glosses over several disagreements in the literature, such as the nature of possible worlds which allow them to play this role.

This conflicts with the fact that many mathematicians would consider (1) false.[5] This result is especially concerning given how central conditionals are to mathematics. This issue does rely on using the material conditional specifically, but this is not unreasonable given that classical logic is used extensively in mathematics.

In my view, the primary issue here rests on the combination of metaphysical necessity being the relevant modality for meaning and assumed of mathematics. If one instead finds it acceptable to consider different ways for mathematics to be and holds that this variation reflects the meaning of mathematical statements, then the triviality dissolves. The folkloric account attempts to flesh out such an alternative.

2.1.1 Hyperintensional Accounts

This triviality partially motivates hyperintensional accounts of meaning, such as impossible worlds and structured proposition accounts. However, much of the work on (and arguably much of the intuitive appeal of) hyperintensional accounts focuses on propositional attitudes, where the differences between intensionally equivalent statements becomes more apparent. The two main approaches to hyperintensionality are those based on impossible worlds and structured propositions.

Impossible worlds generalise possible worlds first by dropping the requirement that worlds be a *possible* state of affairs, as the name suggests. Additionally impossible worlds need not be maximal nor consistent.[6] A circumstance is maximal if for all sentences φ, either φ or $\neg\varphi$ holds in the circumstance.[7] A circumstance is consistent if at most one of φ or $\neg\varphi$ holds in the circumstance. By relinquishing maximality and consistency, impossible worlds are no longer closed under logical consequence. Instead they have four semantic values; the classical truth values, 1 and 0, for true and false, as well as $\emptyset, \{0, 1\}$ for neither true nor false and both true and false respectively.

Despite mathematics serving as a motivation for impossible worlds accounts, surprisingly little work has been done developing this specific application. Most work in the area focuses more on the overall conceptual justification for impossible worlds rather than building and applying specific accounts. As such our critique is more a reflection of differing goals rather than denying the viability of such an account.

[5] We discuss the reasons for (1) being false in footnote 30 and the surrounding text.

[6] In contrast, the folkloric account, motivated by the proposed solution to the primary issue above, only drops the restriction on metaphysical possibility but retains maximality and consistency.

[7] This property is sometimes also called completeness.

Indeed, the flexibility of the impossible worlds approach, which allows developing so many particular accounts, is simultaneously its greatest strength and greatest weakness. Impossible worlds are able to make incredibly subtle distinctions and consequently distinguish a wide-range of phenomena, but this granularity requires very detailed circumstances.[8] Furthermore, these circumstances and collections thereof are no longer tied to features, like metaphysical possibility and logical consequence, that allowed us to characterise them and apply the account at a larger scale.

Since even the smallest differences are now salient, it is easy to describe circumstances that disagree but much harder to classify agreement. In the most extreme case, it is unclear how we ever get synonymy, since it seems for any two sentences there is always a circumstance where one is true and the other is false. For example, on the 'modest' account of (Nolan 1997), any formula is true at some circumstance. So, if we have arbitrary sentences φ and ψ, we can construct a circumstance where $\varphi \wedge \neg\psi$ is true. If conjunction and negation behave classically in this circumstance (which we also seem to be able to specify), then φ is true and ψ is false. Taking the standard circumstantialist position that the meaning of a sentence is the set of circumstances that make it true, then these two sentences have different meanings. This construction can be repeated to make any two sentences non-synonymous.

Of course, the flexibility of impossible worlds allows one to introduce restrictions relevant to the particular situation of interest. Such constraints are a hyperintensional parallel to Carnap's meaning postulates. Meaning postulates are like axioms in that they are formulas used to restrict the collection of possible worlds to only those where they are true.[9] Examples from (Carnap 1952) include defining a bachelor as an unmarried man and requiring cooler to be an irreflexive and transitive relation.[10] We can think of the folkloric account as providing precisely such mathematically motivated constraints on impossible worlds. Whereas Carnap uses meaning postulates primarily to investigate analyticity, the folkloric account explores mathematical meaning instead.

In contrast, advocates of structured propositions suggest that truth conditions don't suffice to capture meaning and that it's necessary to reflect an expression's internal structure as well, i.e., how it is constructed from its components. Roughly speaking, a structured proposition

[8]For example, a single impossible world has to specify how the connectives work (on all four semantic values) and how sentences are related to each other, i.e., how changing the semantic value of one sentence affects the semantic values of others.

[9]In Carnap's terminology possible worlds are state descriptions.

[10]The former is axiomatised by $\forall x(B(x) \leftrightarrow \neg W(x) \wedge M(x))$, where B, W, and M are (unary) relation symbols standing for 'bachelor,' 'wed' and 'man' respectively. The latter is axiomatised by $\forall x \neg Cxx$ and $\forall xyz(Cxy \wedge Cyz \rightarrow Cxy)$ where C is the (binary) relation symbol for 'cooler.'

is a recursive object, usually represented by nested tuples, containing the semantic contents of its components, such as names, predicates and connectives.

I suggest that structured proposition accounts misidentify the source of non-synonymy most pertinent to mathematics.[11] Structured propositions make very fine-grained distinctions which are hard to motivate in mathematical applications. For example, structured proposition accounts generally claim '2 is even' and '2 is not odd' are not synonymous and neither are '2 is prime and even' and '2 is prime and 2 is even.'[12] Additionally, most structured proposition accounts still represent some expressions (such as predicates) within a structured proposition by its intension.[13] Consequently, such accounts retain the challenging restriction to metaphysically possible worlds. As such, mathematical theorems don't differ because of their truth conditions but (merely) because of their grammatical structure.

This simultaneous dissatisfaction with the coarseness of intensions and the incredible fineness of hyperintensional accounts is often called the granularity problem. (Bjerring and Schwarz 2016) identify this as an in principle unsolvable problem, arguing that no 'intermediately-grained' content can satisfy our intuitive requirements of it. Rather than taking this impossibility as terminal, I lean into the goal-dependent nature of modelling.[14] As (Bjerring and Schwarz 2016) put it, 'different phenomena and different contexts seem to call for different levels of granularity,' so we must search for the appropriate granularity in a given context. Allowing different levels of granularity lets us appreciate the fruits of existing accounts whilst still finding them inadequate for *our purposes.*

Thus, our project is to find the level (or levels) of granularity that allow us to capture distinctions found in mathematical practice. The granularity that we settle on is between possible and impossible worlds. As such the folkloric account can be viewed as either an expansion of possible worlds or a contraction of impossible worlds accounts, at least in terms of permitted circumstances.[15]

[11]We expand on this criticism in chapter 3.

[12]The former pair are non-synonymous because 'not even' is a complex predicate and the later pair are distinguished as a predicative and conjunctive sentence respectively.

[13]This is true even of Millian accounts like (Soames 1987) that take the semantic content of a name to be its extension. In fact the ability to recover intensions from structured propositions is often considered an advantage of such accounts.

[14]This also manifests as hewing closer to a linguistic sense of meaning, although, in our case, a rather coarse grained one.

[15]I encourage the reader to choose the view of the folkloric account that they consider more sympathetic. From my perspective, treating the folkloric account as expanding possible worlds or constraining impossible worlds differ primarily in their recommendation for its further development.

2.2 The Folkloric Account

In this section, I introduce the 'folkloric account' and demonstrate how it works through a series of examples. We start by discussing the central ideas that underpin the folkloric account. Next we review the necessary formal apparatus to give the central definition of the folkloric account in definition 2.2.3. Then we elaborate on a specific aspect of the account called model-class selection. Finally in section 2.2.3 we apply the account to see how it works in particular cases.

There are two key ideas motivating the folkloric account.

(a) The meaning of a sentence determines the conditions under which the sentence is true.

(b) Model theory studies and characterises the truth conditions of mathematical statements.

By combining these ideas, we get the following central claim of the folkloric account.

($\star$) The meaning of a mathematical sentence can be represented by the collection of models that make it true.

As phrased in ($\star$) the account may not seem particularly intuitive.[16] However, I call this a folkloric account because I think ($\star$) reflects a reasonable precisification of (a) for mathematics, especially given the problems facing intensional accounts. Of course by their very nature, pieces of folklore are rarely directly or explicitly addressed. As such I only give brief motivation for the account's folkloric status and thereafter help myself to the name, the 'folkloric account.'

The core sentiment of (a) is widely endorsed, even featuring as the opening words to a prominent linguistics textbook, which states 'To know the meaning of a sentence is to know its truth conditions' (Heim and Kratzer 1998, p. 1). Indeed truth-conditional accounts get their name from the fact that they all endorse (a). Circumstantialists go the furthest by directly identifying meaning with truth conditions and representing these truth conditions by sets of circumstances. Furthermore such circumstantialist accounts dominate the literature.

We combine (a) with the claim of (b), namely that model theory is the proper tool for investigating and characterising the truth conditions of mathematical statements. Model theory is the study of particular set-theoretic structures called models and their relationship to formal theories. Formal theories (hereafter theories) are collections of sentences in some

[16] If like me, you already find this position appealing have some patience while I convince others of this.

formal language. Models specify which sentences of the language are true by giving an interpretation of the language's primitive vocabulary, such as constant and predicate symbols, as we show in the formal definitions below. Model theory is both its own discipline of mathematics and a tool used by mathematicians, logicians and philosophers to investigate structures and theories. In particular, model theory is often used to characterise and study the properties of (classes of) models which satisfy, that is make true, either particular theories or classes of theories. Our use of model theory in this chapter is primarily in this latter role as a tool. [17]

In essence then the folkloric account (roughly captured by $\star$) modifies the central idea of circumstantialist accounts to accommodate the huge variety of mathematical structures. Furthermore it does so by making use of model theory, an area of mathematics which studies such structures. I expect that for many users of circumstantialist accounts, including the possible worlds semantics, the change from (a) to ($\star$) will barely register. This is because circumstances are simply (formal) objects in our semantic theory so we can adapt them to our goals, as we already do to account for different modalities. Mathematics concerns not the physical world governed by metaphysical possibility but the mathematical world, which is captured by all kinds of models some of which may not be metaphysically possible.

Before getting to the details of the folkloric account, let's clarify its goals. The folkloric account isn't intended as a complete account of mathematical meaning, rather it captures one important aspect of meaning, which I call *content*.[18] Specifically, content captures an aspect of meaning focusing on a statement's truth conditions. This focus on truth conditions makes content especially amenable to formal analysis with existing tools, and so it is an ideal starting place.

I prefer the terminology of content since the account's judgements about synonymy (sameness of content) diverge from certain intuitions we have about meaning.[19] For example, content does not capture a notion of subject matter.[20] Nonetheless, I am committed to the claim

[17] Since model theory has its own subject matter, there's a possibility of circularity arising from applying a model-theoretic account of meaning to model theory itself. It is outside the scope of this chapter to resolve this issue. However, I wish to clarify that the use of model theory is not intended to grant model theory a more fundamental status than other areas of mathematics, it is simply well suited for meta-theoretic analysis.

[18] Despite moving away from the term meaning, synonymous and related words are retained for ease of presentation.

[19] If pressed I would propose something like a pluralist theory of meaning, where our reports of meaning judgements actually refer to any one of a cluster of related notions, of which content is one, perhaps subject matter is another. As such, none of these individual notions is able to exhaust what we "mean" by meaning.

[20] For example in section 2.2.3 I defend the result that Gödel's compactness theorem and Heine-Borel compactness have the same content, despite being about theories and metric spaces respectively.

that content is best understood as a part of meaning, rather than a separate notion. More accurately, content is intended to be a scientific model (in the surrogative reasoning sense) of mathematical meaning, and hence subject to idealisation.

2.2.1 Formal Definitions

In order to give a formal presentation of the folkloric account, it is helpful to review some technical definitions from model theory.[21] First, we need a language $\mathcal{L}$, consisting of a set of relation and constant symbols, $R_0, R_1, \ldots, c_0, c_1, \ldots$ and a set of logical symbols including $\neg, \rightarrow, \exists$ and a set of variables, $v_0, v_1, \ldots$. For ease of presentation, we follow the common practice of ignoring function symbols and typed or multi-sorted languages. Since we don't include function symbols, the terms of this language are the variables and constant symbols. The atomic formulas of this language have the form $R_i t_0 \ldots t_{n-1}$ where R_i is a relation symbol of arity n and t_i is a term. If φ, ψ are formulas then so are $\neg\varphi$, $(\varphi \rightarrow \psi)$ and $\exists x \varphi$.[22] We refer to variables which fall within the scope of a quantifier, as bound, otherwise they are free. Sentences are formulas with no free variables.

We evaluate formulas within a structure called a model.[23] Roughly speaking, a model specifies which objects exist and gives the extensions of the non-logical symbols in $\mathcal{L}$.

Definition 2.2.1. A model $\mathcal{M} = (M, \cdot^{\mathcal{M}})$ for $\mathcal{L}$ consists of a set of objects M forming the domain of $\mathcal{M}$ and an interpretation function $\cdot^{\mathcal{M}}$. The interpretation function provides the extensions $(R_0^{\mathcal{M}}, R_1^{\mathcal{M}}, \ldots, c_0^{\mathcal{M}}, c_1^{\mathcal{M}})$ of the relation and constant symbols in $\mathcal{L}$.

The extension of a constant symbol is the object that the constant names, i.e., $c_i^{\mathcal{M}} \in M$. The extension of a relation symbol is the set of (n-tuples of) objects that stand in the relation, i.e., for a relation of arity $n \geq 1$, $R_i^{\mathcal{M}} \subseteq M^n$. These extensions allow us to talk about what is true in a model, which is defined formally in definition 2.2.2 below.

For simplicity, we follow common practice and assume that every object in the domain is named by some constant symbol, i.e., for all $a \in M$ there exists $\bar{a} \in \mathcal{L}$ such that $\bar{a}^{\mathcal{M}} = a$. Similarly, we give only the classical definitions for our logical symbols (thus

[21] Readers familiar with the basics of model theory can safely skip to definition 2.2.3. Conversely, a reader interested in a more detailed introduction to model theory should consult a textbook on the topic, such as (Marker 2002).

[22] As usual other connectives and the universal quantifier can be defined as abbreviations for formulas involving only $\neg, \rightarrow, \exists$.

[23] From here onward, model refers to the logical sense common in mathematics and formal semantics (rather than the scientific sense used earlier).

restricting us to classical models) but in principle, the folkloric account also allows varying the underlying logic. Finally, we only give the satisfaction definition for sentences and not for open formulas.[24]

Definition 2.2.2. We read $\mathcal{M} \models \varphi$ as φ is true in $\mathcal{M}$. It is defined as follows:

- $\mathcal{M} \models R_i c_0 \dots c_{n-1}$ if and only if $\langle c_0^{\mathcal{M}}, \dots, c_{n-1}^{\mathcal{M}} \rangle \in R_i^{\mathcal{M}}$
- $\mathcal{M} \models \neg\varphi$ if and only if $\mathcal{M} \not\models \varphi$
- $\mathcal{M} \models \varphi \rightarrow \psi$ if and only if $\mathcal{M} \not\models \varphi$ or $\mathcal{M} \models \psi$
- $\mathcal{M} \models \exists x\varphi$ if and only if there is some $a \in M$ such that $\mathcal{M} \models \varphi(x \mapsto \bar{a})$ where $\varphi(x \mapsto \bar{a})$ means that $\bar{a}$ has been substituted for every free instance of x in φ where $\bar{a}$ is the constant symbol naming a.

Now we can give the key definition of the folkloric account, that the content of a sentence with respect to a collection of models $\mathfrak{C}$ (hereafter *model-class*) is the class of models in that collection which make the sentence true.[25]

Definition 2.2.3. The *content* of a sentence φ, relative to a model-class $\mathfrak{C}$, is

$$[\![\varphi]\!]_{\mathfrak{C}} = \{\mathcal{M} \in \mathfrak{C} : \mathcal{M} \models \varphi\}.$$

For example let $\widehat{C}$ denote the class of all models of classical logic for an language $\mathcal{L}$ containing a zero place relation symbol A.[26] Then we have $[\![A]\!]_{\widehat{C}} = [\![\neg\neg A]\!]_{\widehat{C}}$ since every model in $\widehat{C}$ which satisfies A will also satisfy $\neg\neg A$. So A and $\neg\neg A$ have the same content with respect to $\widehat{C}$.

If two sentences, φ, ψ have the same content with respect to some model-class $\mathfrak{C}$, i.e., $[\![\varphi]\!]_{\mathfrak{C}} = [\![\psi]\!]_{\mathfrak{C}}$, they are called $\mathfrak{C}$-*synonymous* or simply *synonymous* when the choice of model-class is clear. We use $\sim_{\mathfrak{C}}$ to denote the *synonymy relation* which holds between two sentences if they have the same content with respect to the model-class $\mathfrak{C}$. Using our earlier example, we can write $A \sim_{\widehat{C}} \neg\neg A$.

[24]This is to reduce the number of definitions by not introducing variable assignments and because open formulas won't appear in this chapter.

[25]The meaning or content of subsentential expressions is also of interest but falls outside the scope of this account.

[26]Due to how we defined models, the class of all models of classical logic is simply the class of all models. But we use this notation since we can in principle use models of a weaker logic.

Remark 2.2.4. For any two sentences, φ, ψ and model-class $\mathfrak{C}$, the following are equivalent:

- $[\![\varphi]\!]_{\mathfrak{C}} = [\![\psi]\!]_{\mathfrak{C}}$.
- $\mathcal{M} \models \varphi \leftrightarrow \psi$ for all $\mathcal{M} \in \mathfrak{C}$.
- $\varphi \sim_{\mathfrak{C}} \psi$.

As such we use these notations interchangeably, based on which is clearest for the case in question.

One important way of defining a model-class is in terms of a theory. Formally, a theory is a collection of sentences and a model of a theory satisfies every sentence in the theory. We denote the collection of all models of a theory T by $mod(T) = \{\mathcal{M} : \forall\psi \in T, \mathcal{M} \models \psi\}$. We also write $[\![\varphi]\!]_T = \{\mathcal{M} \in mod(T) : \mathcal{M} \models \varphi\}$ when assessing content relative to the model-class given by a theory, i.e., $mod(T)$. Similarly we can refer to content relative to $mod(T)$ as T-content.

2.2.2 Selecting Model-classes

The content of a sentence is analogous to its intension, since both are the set of circumstances which make the sentence true. The first difference is that the folkloric account expands the class of circumstances to allow metaphysically impossible circumstances. However, the greater difference is allowing the class of circumstances, i.e., the model-class, to vary. By varying the model-class, the folkloric account produces distinct synonymy relations. This feature (and hopefully its benefits) will become clearer through the application and analysis of the account in this chapter. We defend this relativity in section 2.3 but for now we give some initial comments on the related question of *model-class selection*, namely which model-classes give us accurate and fruitful synonymy relations.

Various formal properties, such as definability, may impact how fruitful the content judgements of a model-class are. For example, all model-classes used in this chapter are elementary classes which means there exists a theory T such that $\mathfrak{C} = mod(T)$. At this point, we don't adopt any such formal requirements on which model-classes are permitted (although it is certainly an avenue for future work), instead we rely on mathematical practice to provide suitable candidates and inform our choice.

In general, model-classes, and the synonymy relations they generate, are intended to reflect different mathematical frameworks. Framework here is an informal notion which attempts to capture commitments shared by a group of mathematicians.[27] When we change frameworks, we may change which mathematical objects exist or what properties they have and consequently change the meaning of our statements.

Some frameworks, like Euclidean geometry, are already uniquely associated with a model-class usually through a formal theory. In other cases, we have to justify why the model-class is an appropriate formalisation of a given framework. This justification often appeals to mathematical practice. However, the theories regularly used by mathematicians will often be too strong for helpful content judgements, since these theories are chosen because they suffice to outright prove interesting theorems. Proving a theorem outright is unhelpful for investigating content because then the theorem is true in all models of the theory and hence has the same content as any other theorem proved by the theory.

For example, ZFC, the theory of Zermelo Fraenkel set theory with Choice, is the most prominent foundation for mathematics. However, ZFC proves most theorems of mathematics so these theorems all have the same ZFC-content, so this isn't a very informative aspect of meaning. Thus we often have to move to weaker theories to more accurately capture our synonymy judgements. Likewise using a single model $\mathcal{M}$ is unfruitful, even if it reflects the practice of mathematicians, because it will judge all theorems which are true in $\mathcal{M}$ to have the same content.

These choices of $mod(\mathsf{ZFC})$ and $\mathcal{M}$ show that the folkloric account cannot avoid trivial contents entirely. We still have that (1) is true in ZFC and thus has trivial content over $mod(\mathsf{ZFC})$. Unlike with the intensional account, the folkloric account gives us the ability to shift to a different model-class where (1) does not have trivial content. So rather than remaining stuck with the triviality, we instead have tools with which to respond to it. Furthermore, having triviality as a limiting case means that the folkloric account can capture intensional accounts as a special case.

Leaving the choice of model-classes open makes the folkloric account more of a methodology that needs careful application through well selected model-classes rather than an "off-the-shelf" solution. Since we are treating content as a formal surrogate for meaning, rather than a "thick" absolute notion, we allow users of the account to propose a suitable framework and corresponding model-class for a given application.

[27] A framework might alternatively be called a paradigm, but I wish to avoid definitional debates grounded in Kuhn.

2.2.3 Applying the Account

With the folkloric account in hand, let's apply it in four case studies to get a better understanding of how the account works, including concrete examples of model-class selection. Our first case demonstrates the non-triviality of the account. Our second case shows how synonymy relations can reflect shared mathematical content even if this diverges from more colloquial meaning judgements. Our third case builds on the first case to demonstrate how the relativity of the account can be beneficial. Finally, our fourth case looks closer at how content relates to existing equivalence results.

For our first and most basic case, let's consider the theorems from the start of this chapter.

`LUB` Every non-empty subset of $\mathbb{R}$ with an upper bound has a least upper bound.

`WO`$\mathbb{R}$ There exists a well-ordering of $\mathbb{R}$.

`LUB` states that $\mathbb{R}$ is *Dedekind-complete*. Roughly speaking, `LUB` attests to the completeness of the real numbers, meaning there are no 'gaps' in the real number line. `WO`$\mathbb{R}$ says there is a way of (in principle) listing every real number such that we start with some particular number and for each number there is a unique number which follows it in the listed order, and by following this listing we can reach every real number.[28]

It is reasonable to think that these theorems don't mean the same thing. For example Dedekind-completeness is often taken as a characteristic property of $\mathbb{R}$ whereas the existence of a well-ordering seems surprising. This may be due to the fact that `WO`$\mathbb{R}$ requires (a weak form of) the Axiom of Choice, a notoriously counterintuitive principle, whereas `LUB` does not.[29]

If we think that the Axiom of Choice plays a role in the difference between these sentences then `ZF` is a reasonable first choice of theory to use to evaluate contents. `ZF` is the theory of Zermelo Fraenkel set theory without Choice. This theory is commonly used to formalise mathematics, so the collection of its models is a reasonable context for evaluating contents.

[28] This description is intended to give the basic idea of a well-order and is not the standard definition. Note that the usual $\leq$ ordering is not a well-order because there is no first, i.e., $\leq$-least, element.

[29] Another distinction between these sentences is that they talk about distinct properties. Being well-orderable is not equivalent to being Dedekind-complete since we can well-order the rational numbers $\mathbb{Q}$ but they are not Dedekind-complete. For the purposes of this chapter, we only consider sentential synonymy so we don't pursue this distinction.

Assuming the consistency of ZF, there are models of ZF where $\mathtt{LUB}$ is true but $\mathtt{WO}\mathbb{R}$ is not, so there exists $\mathcal{M} \in mod(\mathsf{ZF})$ such that $\mathcal{M} \in [\![\mathtt{LUB}]\!]_{\mathsf{ZF}}$ and $\mathcal{M} \notin [\![\mathtt{WO}\mathbb{R}]\!]_{\mathsf{ZF}}$.[30] Therefore $[\![\mathtt{LUB}]\!]_{\mathsf{ZF}} \neq [\![\mathtt{WO}\mathbb{R}]\!]_{\mathsf{ZF}}$, so the folkloric account concludes that they have different contents as desired.

On the other hand, different mathematical theorems need not have different meanings. Consider the following theorem which states that $\mathbb{R}$ is *Cauchy-complete*.

$\mathtt{CC}$ Every Cauchy sequence in $\mathbb{R}$ has a limit.

Roughly speaking, $\mathtt{CC}$ says that every sequence of real numbers that get arbitrarily close together approaches some limiting value. This theorem is often taken to have the same meaning as $\mathtt{LUB}$ because they are both ways of attesting to the completeness of $\mathbb{R}$.[31] This is reflected by mathematicians regularly saying that '$\mathbb{R}$ is complete,' without specifying whether they mean '$\mathbb{R}$ is Dedekind-complete' or '$\mathbb{R}$ is Cauchy-complete' (which are rephrasings of $\mathtt{LUB}$ and $\mathtt{CC}$ respectively).[32] Since for all $\mathcal{M} \in mod(\mathsf{ZF})$, $\mathcal{M} \vDash \mathtt{LUB} \leftrightarrow \mathtt{CC}$, the folkloric account agrees that they are synonymous in ZF.

Sticking with ZF-content, let's consider our second case with the following theorems, the axiom of choice and the well-ordering principle.

$\mathtt{AC}$ If S is a set of non-empty sets, there is a function f such that for every $X \in S$, $f(X) \in X$.

$\mathtt{WO}$ Every set can be well-ordered.

Within ZF these statements are equivalent, i.e., $\mathsf{ZF} \models \mathtt{AC} \leftrightarrow \mathtt{WO}$. However, they are thought to mean very different things, as the following quip observes "the axiom of choice is obviously true, the well-ordering principle obviously false, and who can tell about Zorn's lemma."[33] Nonetheless, I suggest that ZF-content is tracking a salient aspect of mathematical meaning.

[30] In particular a model of ZF which also satisfies the Axiom of Determinacy satisfies $\mathtt{LUB}$ but makes $\mathtt{WO}\mathbb{R}$ false, see (Jech 2006, Lemma 33.1). Such a model also witnesses that (1) does not have trivial content, i.e., that $[\![(1)]\!]_{\mathsf{ZF}} \neq mod(\mathsf{ZF})$.

[31] In fact, $\mathtt{LUB}$ and $\mathtt{CC}$ were originally proposed as alternative formalisations of completeness and were only later shown to be equivalent.

[32] Alternatively, one might distinguish between Cauchy and Dedekind completeness on the basis that some theories cannot construct an isomorphism between the Dedekind real numbers and Cauchy real numbers. This is an interesting example that warrants further investigation but again it requires an account of subsentential meaning which we don't provide in this chapter.

[33] This quote is attributed to Jerry Bona in (Schechter 1997).

The following sketch of their equivalence reveals just how easy it is to define a well-order from a choice function and vice versa. So even though the existence of a well-order for any set is initially surprising, the equivalence shows that this really just expresses the same mathematical fact as the axiom of choice.

Suppose $\mathtt{AC}$ is true, then we have a choice function f on the set of all non-empty subsets of an arbitrary set X. From this we can define a well ordering on X as follows: let the first element of X be $f(X) = a$, let the second element of X be $f(X \backslash \{a\}) = b$, and so on until we have exhausted the set X. Hence we have $\mathtt{WO}$. Conversely, supposing $\mathtt{WO}$ is true gives that S (an arbitrary set of non-empty sets) can be well-ordered. So we define the choice function to return the least element of each set $X \in S$ as it appears in the ordering of S, hence we have $\mathtt{AC}$.[34]

Earlier I claimed that we should define content with respect to multiple model-classes, rather than just one. So let's move to our third case where changing the model-class produces distinct synonymy relations which both reflect frameworks that mathematicians might employ. We return to the properties of Dedekind and Cauchy completeness, but replace $\mathbb{R}$ with an arbitrary ordered field, $\mathbb{K}$.

$\mathtt{LUB}_{\mathbb{K}}$ Every non-empty subset of $\mathbb{K}$ which is bounded above has a least upper bound.

$\mathtt{CC}_{\mathbb{K}}$ Every Cauchy sequence in $\mathbb{K}$ has a limit.

The situation here is more nuanced than in the $\mathbb{R}$ case, because we can readily find different theories which disagree about whether $\mathtt{LUB}_{\mathbb{K}}$, $\mathtt{CC}_{\mathbb{K}}$ are synonymous or not. These are both statements about ordered fields, so the class of models of OF_2, the second-order theory of ordered fields, is a plausible choice. This theory defines what it means to be an ordered field, and it contains the field axioms (e.g. existence of identities and inverses etc.) and the order axioms.[35] These sentences have distinct OF_2-contents because $\mathsf{OF}_2 \models \mathtt{LUB}_{\mathbb{K}} \to \mathtt{CC}_{\mathbb{K}}$ but $\mathsf{OF}_2 \not\models \mathtt{CC}_{\mathbb{K}} \to \mathtt{LUB}_{\mathbb{K}}$.[36] Thus $[\![\mathtt{LUB}_{\mathbb{K}}]\!]_{\mathsf{OF}_2} \neq [\![\mathtt{CC}_{\mathbb{K}}]\!]_{\mathsf{OF}_2}$, so they are not synonymous from the perspective of $mod(\mathsf{OF}_2)$.

Now let's consider the second order theory of *Archimedian* ordered fields, AOF_2 which simply

[34]Not all equivalences to choice principles are as easy to see but largely they operate on similar principles, such as using the ordering to define other operations or to pick a canonical or extremal (e.g. greatest or least) representative.

[35]See (Marker 2002) for more details. We consider the second order theory which allows us to quantify over the powerset of field elements because we need to refer to e.g. "every non-empty subset."

[36](Propp 2013, p. 398).

adds the Archimedian axiom ($\forall x \exists n \in \mathbb{N}_{\mathbb{K}}, n > x$) to OF_2.[37] With respect to AOF_2, these two sentences have the same content, i.e., $[\![\mathtt{LUB}_{\mathbb{K}}]\!]_{\mathsf{AOF}_2} = [\![\mathtt{CC}_{\mathbb{K}}]\!]_{\mathsf{AOF}_2}$. This result accords with the earlier case where $\mathbb{K} = \mathbb{R}$ since $\mathbb{R}$ is an Archimedian field.

At first glance, it might seem troubling that our different model classes disagree about the synonymy of these two sentences. Instead, I argue it is beneficial that the folkloric account is able to capture both of these situations since both theories capture a framework that mathematicians might be interested in. In particular, one framework for all ordered fields and another framework specifically for Archimedian ordered fields. For mathematicians in the Archimedian framework, $\mathtt{LUB}_{\mathbb{K}}$ is synonymous with $\mathtt{CC}_{\mathbb{K}}$, whereas those in the more general framework can distinguish these theorems. As long as we use an appropriate model-class for the context, we don't have any issues.

In fact, this case is especially well-behaved because we can identify the Archimedian axiom, as a natural boundary between judging these sentences to be synonymous or non-synonymous. Thus, we can point to a specific and independently motivated mathematical property which explains why these theories give different judgements. This example is a little artificial since we can recover the equivalence within OF_2 by altering the statement to refer explicitly to Archimedian ordered fields.[38] Yet our ability to capture the behaviour of the stronger theory within the weaker one by restricting the domain just highlights how the theory we use impacts which objects we can refer to, which is an obvious source of meaning change across theories.

Our fourth case comes from reverse mathematics. Reverse mathematics studies the equivalence of theorems over weak base theories, commonly the theory RCA_0 which roughly captures "computable mathematics." The following two theorems have the same content with respect to $mod(\mathsf{RCA}_0)$.

`comp` If there is no model of a countable theory Γ then some finite subset of Γ has no model.

`cover` Every covering of $[0, 1]$ by a sequence of open sets has a finite subcovering.

[37] $\mathbb{N}_{\mathbb{K}}$ is the abelian semigroup isomorphic to $\mathbb{N}$ found in each ordered field $\mathbb{K}$. Again this is a well studied class of objects.

[38] That is, the following two sentences have the same OF_2-content. Every non-empty subset of an *Archimedian* ordered field with an upper bound has a least upper bound. Every Cauchy sequence in an *Archimedian* ordered field has a limit.

As with our second case these sentences might not seem to mean the same thing prima facie. Yet mathematicians themselves recognised the similarity in calling them both compactness theorems, namely, Gödel's and Heine-Borel's respectively. I claim that the similarity is based on an underlying structural correspondence. This is best seen by comparison to a third statement with the same RCA_0-content as both `comp` and `cover`.[39]

`WKL` If a tree $T \subseteq 2^\omega$ has no infinite path, then it is finite.

The fact underpinning these equivalences is that we can encode theories or coverings into a binary tree, and conversely encode a binary tree as either a theory or a covering. This is done by thinking of all three (trees, theories and coverings) as collections of sequences of binary decisions.

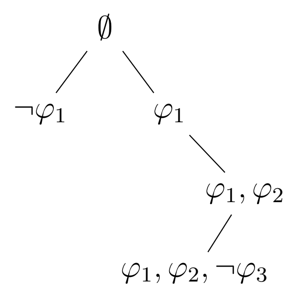

Figure 2.1: A finite fragment of a theory represented as a tree.

Rather than giving the full proofs of these equivalences, which can be found in (Simpson 2009), I only sketch the necessary encodings which allows us to focus on the use of the tree representation.[40] Figure 2.1 shows how a theory Γ can be encoded as a tree T. We let φ_i denote the ith sentence in an enumeration of the language of Γ, σ denote a binary sequence and $\sigma(i)$ denote the ith element of σ. Then we define a tree T such that $\sigma \in T$ if and only if for all i, $\sigma(i) = 1$ just in case φ_i is true in Γ. Conversely, a tree can be encoded as a theory by constructing sentences which are true if and only if there is at least one σ such that φ_i is true if $\sigma(i) = 1$ and false if $\sigma(i) = 0$.

Likewise a covering of $[0, 1]$ can be encoded by a tree by having the tree identify which intervals of length $\frac{1}{2^i}$ should be in the covering.[41] Conversely again, a tree can be encoded as a covering by having the sequences in the tree correspond to the inclusion (or exclusion) of

[39] Both `comp` and `WKL` below are better known by their contrapositive statement but this presentation is chosen to highlight their 'finite' reading to better reflect the equivalence to `cover`.

[40] We also ignore the fact that these encodings need to be done in stages, i.e., we can only construct a finite part at a time.

[41] A similar construction works for other compact metric spaces.

certain intervals in the covering. These encodings are fundamental in reverse mathematics and are also used in other subdisciplines. We can interpret these equivalences as showing that each of these theorems tells us "the same thing" about the finite versus infinite behaviour of objects which can be coded as a collection of sequences of binary decisions. As such the fact that they all share the same RCA_0-content captures a salient mathematical similarity between them.

With these examples, we have seen how to apply the folkloric account in various situations and how its judgements relate to informal discussions of meaning. We now move on to a broader assessment of the account.

2.3 Exploring Relativity

Mathematics has a reputation for absolute truth, yet the practice of mathematics reveals a huge variety of frameworks being used. Mathematicians are interested in a diverse range of objects and tools for studying them. In this section, I show how these different frameworks justify and inform the relativity permitted by the folkloric account. This relativity is the most interesting and likely controversial feature of the account, so it is important to look more closely at how it operates. First, we respond to the objection that the folkloric account doesn't capture the intended interpretation implicit in mathematical sentences. Second, we turn more directly to motivating relativity with respect to mathematical practice. Then, we defend our choice to maintain multiple distinct synonymy relations rather than subsuming them into a more general relation. Finally, we discuss some ways to limit the generation of new synonymy relations.

2.3.1 Intended Interpretation

One potential complaint about the folkloric account is that it overlooks a statement's intended interpretation. This position holds that a statement like 'Ψ' should be explicated as 'Ψ is true in the intended interpretation.' In the non-mathematical case, this intended interpretation is the actual world/ true state description, whereas in the mathematical case it would be something like the true theory of mathematics.[42] On this objection then the mathematical

[42]Note that in the mathematical case, the 'intended interpretation' is still usually a class of models, rather than a single model as in the non-mathematical case.

setting always determines a single model-class within which we should evaluate the statement (and furthermore, it will be true throughout this model-class). This issue impacts all circumstantialist accounts so I respond to it only briefly.

I agree that in practice, statements like 'Ψ' are often used to convey something about Ψ being true in the intended interpretation. Yet I consider this a feature of communication which goes beyond an account of meaning. What communication consists in and how it relates to meaning is an important question but not one that I hope to answer in this chapter.[43] The folkloric account aims to give a formal representation of content which can be fruitfully applied in mathematics, we don't intend to give an account of how expressions receive their meaning or encompass all of mathematical communication. The first project alone presents sufficient challenge.

Furthermore, circumstantialist accounts explain how the intended interpretation relates to the content. Namely, it's legitimate to communicate that 'Ψ is true in the intended interpretation' precisely when the intended interpretation is in the content of Ψ. Thus, it's correct to communicate that 'It's true that the moon shines in the actual world' just in case the actual world is a circumstance that satisfies the sentence 'The moon shines.' Similarly expressing 'The continuum hypothesis holds in the true theory of mathematics' is legitimate if and only if every model of true mathematics satisfies the continuum hypothesis. Unfortunately, we don't know what that true theory of mathematics is. But even if we did it wouldn't help us with meaning per se, it would only determine its actual truth value.

2.3.2 Motivating Relativity

Another reason for rejecting this intended interpretation objection is that it is not clear that there is a universal setting that all mathematicians work within at all times. Or as we have already discussed, such settings like ZFC, do not allow us to distinguish theorems by their truth conditions. The development of non-Euclidean geometries is a well-known historical case exemplifying the existence of multiple fruitful mathematical frameworks. In this section, we motivate the relativity of the folkloric account through an example based on this case. Furthermore, since mathematicians can work in distinct and even incompatible frameworks, there is no neutral choice of model-class available. Instead we must attempt to formalise the relevant framework by selecting an appropriate model-class for a given application.

[43] Different ways of answering this question depends on one's favourite theory of speech acts or assertions or discourse norms etc.

Non-Euclidean geometries arose when attempts to prove Euclid's fifth postulate from the prior resources of Euclid's Elements instead revealed the existence of geometries which denied this postulate. Consequently, geometry was no longer the study of a single theory, we now had Euclidean geometry which affirms the fifth postulate and hyperbolic geometry which denies it and many other geometries, like spherical and projective geometry. Although there were some attempts to preserve Euclidean geometry as the correct theory of geometry, ultimately the mathematical fruitfulness of these non-Euclidean geometries and the developments of further mathematics like manifolds, led to these geometries being embraced.[44]

Thus, we have multiple geometries, each of which are a legitimate framework to do mathematics within. Furthermore, these geometric theories are going to evaluate statements differently. These differing evaluations will be captured by distinct contents in the folkloric account. To see this let's consider the following statement.

$\mathtt{ang}$ Every triangle has angles summing to 180°.

We focus on the following three geometric theories to evaluate this statement.

- AG Absolute geometry, is given by Euclid's first four postulates and the so-called common notions,[45]
- EG Euclidean geometry, adds Euclid's fifth postulate to AG,
- HG Hyperbolic geometry, adds the negation of Euclid's fifth postulate to AG.

In Euclidean geometry, $\mathtt{ang}$ is a theorem as it is provably equivalent to Euclid's fifth postulate. In absolute geometry, $\mathtt{ang}$ cannot be proved so it is not a theorem, but its negation is also not a theorem. In hyperbolic geometry, $\mathtt{ang}$ is false, instead every triangle has angles summing to strictly less than 180°. So we have three theories which disagree about the truth, and hence the meaning, of a statement.

In terms of the folkloric account, we have that the content of $\mathtt{ang}$ in hyperbolic geometry is different from its content with respect to Euclidean geometry or absolute geometry. That is, $[\![\mathtt{ang}]\!]_{\mathsf{HG}} \neq [\![\mathtt{ang}]\!]_{\mathsf{EG}}$ and $[\![\mathtt{ang}]\!]_{\mathsf{HG}} \neq [\![\mathtt{ang}]\!]_{\mathsf{AG}}$. Since $\mathtt{ang}$ is true in all and only the models

[44] Furthermore, some of the defences of Euclidean geometry only argued that Euclidean geometry was the correct geometry for some canonical application, like that of physical space. Even granting this (now demonstrably false) claim, such an argument doesn't undermine my assertion that mathematicians work in different theories because these different theories may be for different applications.

[45] The common notions are the additional definitions required to actually set up a theory, they include basic laws such as those governing equality.

of Euclidean geometry, i.e., $mod(\mathsf{EG})$, this set of models is the content of `ang` for both AG and EG. Strictly speaking these are still distinct contents since they are given with respect to different background model-classes. We can also distinguish them because we can see immediately that `ang` has trivial(ly true) content with respect to EG, but its AG-content is not trivial.

Thus, this example shows that looking at content alone may not reveal the full picture and that in some cases we may be more interested in the synonymy relation. Supposing we did not already know that `ang` has trivial EG-content and non-trivial AG-content, we might discover this by looking at the content of something we know to be axiomatic, such as the following statement.

`tri` Every triangle has three sides.

This statement is definitional of triangles so $[\![\mathtt{tri}]\!]_{\mathsf{AG}} = mod(\mathsf{AG})$ and $[\![\mathtt{tri}]\!]_{\mathsf{EG}} = mod(\mathsf{EG})$. Consequently, $\mathtt{tri}\sim_{\mathsf{EG}}\mathtt{ang}$ but $\mathtt{tri}\not\sim_{\mathsf{AG}}\mathtt{ang}$. So we can distinguish the synonymy relations $\sim_{\mathsf{AG}}$ and $\sim_{\mathsf{EG}}$.

These examples show how using model-classes to capture different mathematical frameworks is necessary to reflect mathematical practice. Obviously this is a salient goal for the folkloric account and so justifies the choice to vary the model-class when applying the account.

Furthermore, because the model-class changes the evaluation of a statement, there is no "neutral" choice of model-class. Any model-class brings with it some commitments regarding the truth and hence meaning of various statements. These commitments often differ due to the model-classes, and the theories they capture, containing different objects. Since mathematical theories deal with abstract objects, it is reasonable to think that these theories fully determine what it is to be the kind of mathematical object(s) specified by the theory. There is no additional quality of being that kind of mathematical object which exists outside the theory. Therefore, it is not possible to formally evaluate statements or synonymy claims without first choosing a theory or model-class which determines which objects we're interested in.

My claim that mathematics requires varying the theory or model-class of evaluation seems to conflict with the fact that we don't see mathematicians specify which theory or model-class they are working in. First, we note the obvious fact that while specifying a model-class is an uncommon practice, there are situations where mathematicians explicitly do so.

This is especially common within pedagogical contexts or disciplines which more routinely vary the theory being used, such as set theory. These examples point to the second, broader reply, namely that most mathematicians work almost exclusively in a fixed context where the framework being employed is well understood. Part of gaining competence as a mathematician is becoming familiar with the theory relevant to one's specific field. Once competence is reached, reference to this theory can be elided.

Nonetheless, explicitly choosing a model-class is also somewhat of an artefact of the formal tools we use. In practice almost all mathematics is done informally, so model-classes are our formal representations of the informal frameworks being applied. This is an idealisation but a tame one that also provides access to the resources of model theory. The challenge here is primarily one of defining and choosing model-classes which successfully serve as representations for frameworks. Indeed this challenge is a central part of the problem of *model-class selection* that we have encountered throughout.[46]

Taking this formalisation as a justifiable idealisation also deals with the objection that the qualities of mathematical objects are not given by the theory but the objects that we wish to apply mathematics too, such as concrete objects or pre-theoretical concepts. Of course such objects do guide mathematical practice and there are various cases where theories are replaced for failing to capture the relevant domain. We can see this as providing "fidelity criteria" that our theories have to meet to better capture these objects. The theories still remain our best formal representation of these objects.

2.3.3 Against Unifying Relativity

We have seen that mathematicians can work in or be interested in different frameworks and in this section, I defend the choice to maintain multiple distinct synonymy relations rather than subsuming them into a more general relation. We reject each of the following two most obvious alternatives to allowing multiple synonymy relations, as they face both conceptual and formal obstacles.

i) Defining synonymy with respect to the broadest class of models and restricting to other synonymy relations, e.g., by conditionalisation.

ii) Have meaning be the function from model-classes to contents. So two sentences are synonymous if and only if $[\![\varphi]\!]_{\mathfrak{C}} = [\![\psi]\!]_{\mathfrak{C}}$ *for all* $\mathfrak{C}$.

[46]See section 2.2.2 for the start of our discussion on model-class selection.

My main objection to i) is the lack of some principled choice for this broadest model-class. As just discussed, I claim that truth and hence meaning in mathematics is essentially relative to a choice of framework and there are many equally legitimate frameworks. The most principled choice available would be the class of all logically possible worlds, but this would collapse mathematical meaning into logical equivalence. This seems prima facie undesirable and rules out synonymy of the kind between Dedekind and Cauchy completeness. The use of conditionals in i) to recover the other model classes also faces a formal issue. Namely, that we can use infinitely many axioms to characterise model-classes (e.g. the class of all models of ZFC) but in ordinary first order logic we can't conditionalise in infinitely many antecedents.

I am more sympathetic to ii) but ultimately object to it on practical grounds. I suggest that it complicates the account without providing commensurate intellectual reward. There is already much to discuss about synonymy within a fixed mathematical framework. Indeed most applications of an account of mathematical synonymy are situated within a particular framework because the salient question is whether two sentences are synonymous *for some purpose*, like usage in a proof, or a classroom discussion. Of course, it is valuable to compare contents across model-classes as in the geometric and completeness cases above. However, we lack the mathematical tools to consider all model-classes simultaneously. We face the difficulty of quantifying over such a large collection and this is compounded by the fact that even individual model-classes might not be well understood, especially those which have not been studied for some independent purpose. Working with a particular model-class already provides a fruitful analysis of mathematical synonymy with manageable complexity. As such, we prefer to have multiple distinct synonymy relations available.

2.3.4 Combating 'Content Overgeneration'

Having argued for the relativity of mathematical practice, we are faced with the fact that any statement can have many different contents. Call this the problem of *content overgeneration*. This problem arises because every model-class defines a content function and a corresponding synonymy relation. This abundance of synonymy relations can produce cases both of 'content-fracturing' where apparently equivalent statements have distinct contents and of 'content-agglomeration' where seemingly inequivalent statements have the same content. I remain committed to the relativity of the folkloric account so this problem cannot be completely avoided but in this section I argue it is less serious than it might first appear.

Obviously statements don't have different contents in a contradictory sense, Instead they only have different contents with respect to different model-classes. Since we always relativise content to a model-class, we avoid being committed to, e.g., two statements being both synonymous and not synonymous to each other. Thus, the real issue is how to adjudicate between different contents or synonymy relations.

Unfortunately, there is no uniform solution for determining the most appropriate model-class, instead this question will depend on the particular application. Mathematical practice can provide a basis for selecting a model-class in specific cases. We can assess either whether the judgements converge with those of mathematicians or provide other justification for why a model-class captures the subject matter of some framework. In the cases we've seen so far involving ordered fields and triangles, the different subject matters have been relatively straightforward to identify. Certainly when dealing with a model-class defined by a theory, most axioms will change the domain so the choice comes down to selecting the right kind of objects. We also don't need to commit to a definitive "correct" choice of model-class, as the process of exploring different model-classes can be fruitful in itself.

Furthermore, there are limits to how many synonymy relations there are and how bad they can be. First, not all model-classes define distinct synonymy relations. For example, two models are called elementarily equivalent if and only if they satisfy all the same first order sentences.[47] So adding or removing models which are elementarily equivalent to another model in the model-class doesn't change the synonymy relation since we only consider sentential meaning in this chapter.

Similarly, restricting to a smaller collection of models (which corresponds to restricting to a stronger theory) will never break synonymy between sentences. That is if $\mathfrak{C}' \subseteq \mathfrak{C}$ then $[\![\varphi]\!]_{\mathfrak{C}} = [\![\psi]\!]_{\mathfrak{C}}$ implies $[\![\varphi]\!]_{\mathfrak{C}'} = [\![\psi]\!]_{\mathfrak{C}'}$. Conversely moving to a larger model-class (or a weaker theory) will never make more sentences synonymous. That is if $\mathfrak{C}' \supseteq \mathfrak{C}$ then $[\![\varphi]\!]_{\mathfrak{C}} \neq [\![\psi]\!]_{\mathfrak{C}}$ implies $[\![\varphi]\!]_{\mathfrak{C}'} \neq [\![\psi]\!]_{\mathfrak{C}'}$. So these are both changes we can make to model-classes that do not result in new synonymy relations.

A type of synonymy relation which initially looks unappealing or unhelpful are those produced by simply adding '$\varphi \leftrightarrow \psi$' (or its negation) to an existing theory. Relative to this new theory's synonymy relation, it is not particularly interesting that φ and ψ have the same content (or different content if we added the negation). This seems like a totally arbitrary way of forcing certain content judgements which would be concerning. However, consistency and "down-

[47]Elementary equivalence is an important basic concept in model theory, see Marker 2002 for more details.

stream consequences" prevent us from using this to force every arbitrary content judgement we choose. More optimistically, such pathological cases might even provide novel insights.

By down-stream consequences, I mean the fact that adding '$\varphi \leftrightarrow \psi$' doesn't just impact the content judgement of φ and ψ, it impacts the content of everything synonymous to either sentence.[48] Adding this biconditional changes contents throughout the whole theory so it's less arbitrary than it appears. One could even consider testing the synonymy of sentences by seeing if adding the corresponding biconditional leads to any unintuitive content judgements about other statements. In particular, continuing to force other content judgements by adding more biconditionals will eventually lead to triviality.

Requiring consistency of our theories prevents unilaterally adding biconditionals to all theories because this could make the theory inconsistent. For example, adding $\varphi \leftrightarrow \psi$ to a theory that already implies $\varphi \leftrightarrow \neg\psi$ would lead to an inconsistent theory. I take it for granted that we are not interested in inconsistent theories, since classically they will judge every sentence to have the same content as any other sentence.[49] So we can only use the stronger theory, if it was possible for φ and ψ to have the same content with respect to the original theory.

2.4 Future Directions

So far we have only managed to scratch the surface of the folkloric account. Notably we need to develop more case studies to get a better understanding of how the account works and connect this (for better or worse) to the practice of mathematics. At present the account looks promising especially given the lack of adequate competing accounts. Let us end by discussing future directions which will inform how we advance the folkloric account.

Clearly the selection of model-classes plays an important role in the successful application of the folkloric account and there are several questions worth investigating about selecting model-classes. First, which model-classes produce content judgements which accord with mathematical practice? We've already seen that such model-classes are likely larger than that of the intended interpretation because we need to track distinctions which are erased by merely requiring truth in the model-class (i.e., provability).[50]

[48] This just reflects that $\sim_T$ forms an equivalence relation, which follows trivially from biconditionals being reflexive, symmetric and transitive. That is, for all $\mathcal{M}$ and for all sentences, φ, ψ, θ, we have $\mathcal{M} \models \varphi \leftrightarrow \varphi$, $\mathcal{M} \models (\varphi \leftrightarrow \psi) \rightarrow (\psi \leftrightarrow \varphi)$, and $\mathcal{M} \models ((\varphi \leftrightarrow \psi \wedge (\psi \leftrightarrow \theta)) \rightarrow (\varphi \leftrightarrow \theta)$.

[49] In principle the folkloric account could accommodate a paraconsistent logic which might make inconsistent theories more interesting but we don't investigate this further.

[50] The triviality of theorems of ZFC in $mod(\mathsf{ZFC})$ is one such case.

We might also consider restricting to model-classes which satisfy certain formal properties. For example, all model-classes in this chapter are elementary classes for theories which are recursively axiomatisable. These are selected primarily for being well-behaved and easy to describe but perhaps there are other considerations either for or against enforcing this as an required restriction to the account. Alternatively, should model-classes be closed under certain kinds of translations? Both the formal and philosophical consequences of such restrictions are of interest. On the philosophical side, we might think that these formal properties explicate what it means to form a 'coherent' collection of models rather than a collection of randomly selected models or a model-class which is purpose built to give the desired content judgements. This question of what constitutes a coherent model-class is analogous to the question of which properties a logic needs to satisfy (to be a logic). On the more formal side there is an engaging literature on the relationship between translation and theoretical equivalence, which provides its own perspective on preserving "meaning" across models.

In this chapter, we focused on making content judgements within one model-class but comparing judgements across different model-classes is also of interest. In particular, these comparisons can identify 'difference makers', which are the sentences which when added to a theory of interest make two non-synonymous statements synonymous in the stronger theory. For example, we identified the Archimedian principle as the difference maker between Dedekind and Cauchy completeness in $\mathbb{K}$, i.e., $\mathtt{LUB}_{\mathbb{K}}$ and $\mathtt{CC}_{\mathbb{K}}$. More generally, it is a real strength of the folkloric account that it relativises to different model-classes, so being able to talk more fully about the impact of these differences is potentially fruitful.

As we can see, there is much work to be done further developing the folkloric account. In the meantime, it presents a compelling step towards understanding meaning in mathematics.

CHAPTER 3
Structured Propositions: a critical analysis for modellers

3.1 Introduction

Infamously propositions are invoked to serve a variety of functions, including being: the bearers of truth values, the objects of propositional attitudes (such as belief), the referents of that-clauses, and the meanings of sentences (McGrath and Frank 2024). In this chapter, we are primarily interested in this final semantic role and we approach it from what we call a model-based or modeller's perspective.[1] In its most basic form, the model-based perspective understands accounts in formal semantics as providing models that enable surrogative reasoning about theoretical targets, such as semantic content.[2] As models, these accounts are assessed with respect to their adequacy for purpose, i.e., their utility in pursuing the goals of their user. Several factors may contribute to a user's goals, such as; the intended function of the target as above, the scope of application, the kind of information sought from the model (e.g. does it need to be computationally tractable) etc. Working within this perspective, this chapter describes and evaluates the influential structured proposition account of (Soames 1987).

First, section 3.1 briefly introduces the model-based perspective we take in this chapter and the structured proposition approach. Then section 3.2 provides a detailed presentation of Soames' account of structured propositions for a language without propositional attitudes. Finally, section 3.3 provides some challenging cases where the Soamesian account makes

[1]This perspective need not apply to all semantic investigation, it suffices that it describes some part of the practice.

[2]In this chapter, we use the term *model* in the sense of a theoretical surrogate described just now, often distinguished as "scientific" models. This is in contrast to logical models which consist of a domain and an interpretation for a formal language. The situation is further confused by the fact that models in the latter logical sense can act as or be a part of models in the former scientific sense. Indeed this chapter deals with just such a case. Consequently, we refer to the logical objects as *possible worlds*, which they are used to represent. Additionally, a collection of possible worlds or objects representing them can also be called a model, but we use the term *universe* instead.

distinctions that I argue are too fine-grained. We focus on applying Soames' account to mathematics so the cases come primarily from there. However, the structural relationships underpinning these cases also occur in non-mathematical settings and so may present concerns for the account across other areas of application.

Our presentation of Soames' system prioritises setting up its technical apparatus, so it can be applied to specific cases of interest and to be developed further (if desired). The expositional goals of this chapter, e.g., providing ample examples and meticulous definitions, were pursued with modelling in mind, but their benefits are not limited to those who endorse a modeller's perspective.

3.1.1 Model-based Inquiry in Formal Semantics

Some work in formal semantics engages in modelling, an umbrella term for a variety of practices including constructing, investigating, applying and evaluating models. Models act as surrogates for some (theoretical) target, allowing users to better understand and reason about the target by manipulating and learning about the model.

For some semanticists, modelling is central to their conception of the discipline, see e.g. (Yalcin 2018). On such a view, semantic accounts do not (and can not) give us meanings directly, instead they provide us with a definite theoretical construct which we can use to inform our understanding of some nebulous concept like meaning or linguistic competence. Regardless of whether one accepts this position, it is clear that modelling is happening. I contend that there is valuable work to be done investigating these models, including simply detailing more of their behaviour. In this chapter, we don't attempt to defend this view that semantics is best understood as a form of model-based inquiry. However, our analysis of Soames' account is underpinned by a modeller's perspective and so we briefly describe this approach here.[3]

There are many diverse semantic projects so it is necessary to describe our target, which we call *semantic content*. We focus on formal semantics as developed within the philosophical literature, especially at the intersection of logic and philosophy of language, where truth, reference and satisfaction are central. As such, we are not interested in more cognitive and/or developmental projects in linguistics.

[3]More detail on the model-based view and its application to semantics can be found in chapter 1.

A central feature of modelling activities is that they are guided by the goals of inquirers. This is most obvious when evaluating models because inquirers with different goals will assess models differently. Similarly, different models may be better suited for different purposes. This also reflects that models are partial and that it might not be possible to give a complete account of a given target. As such, it may be necessary to narrow down the scope of a model and limit oneself to more circumspect claims.

For our purposes, we emphasise the role semantic content plays in defining synonymy. That is, a pivotal function of semantic contents is to predict and explain whether statements are synonymous or not. We evaluate these predictions by comparing them to ordinary language use and linguistic judgements.[4] Such evidence is fundamentally limited since synonymy reports are rarely expressed directly and involve meta-cognitive processes. Additionally these judgements are defeasible since language users are not infallible and because they might reflect different aspects of meaning which are not our target. Nonetheless, this is the best access we have to our target. Focusing on how "meaning" shows up in the world commonly distinguishes card-carrying modellers from other inquirers who use models, such as those who rely more on a priori principles.

Additionally, we consider it a basic functional role of semantic contents that they are associated with some linguistic item, e.g., name, sentence, predicate, and that they determine the (semantic) value of that item. This is shown in figure 3.1 with $\mathcal{I}, \mathcal{I}'$ representing additional information provided by indices (such as variable assignments, possible worlds, speakers, etc.). Furthermore, semantic contents ideally should capture all and only the semantically relevant features of the linguistic item. Of course, what counts as a "truly semantic" feature is a central source of disagreement between accounts of semantic contents.

$$\mathsf{Linguistic\ Item} \times \mathcal{I} \longrightarrow \mathsf{Semantic\ Content} \times \mathcal{I}' \longrightarrow \mathsf{Value}$$

Figure 3.1: Functional Role of Semantic Contents.

For example, on the intensional account it is only the truth conditions across metaphysically possible worlds that are semantically relevant. Alternatively, structured proposition accounts hold (unsurprisingly) that semantic contents are structured in a way that mirrors syntactic features of the expression. Nonetheless they defend this structure as semantically relevant. For our archetypal modeller, the issue is less about demarcating the truly semantic and more about whether the features captured are useful for the goals of a particular application.

[4]This represents a significant departure from Soames' own approach.

Despite the philosophical prominence of structured propositions, there is surprisingly little *modelling* work done on them. Instead much of the literature focuses on metaphysical and meta-semantic debates about the essential nature of the structuring relation.[5] Whilst I have no qualms with such investigation, it is not central to the interests of myself nor our archetypal modeller. Furthermore, I think that such debates could also benefit from work by modellers, both in terms of what they can say about existing accounts and through developing models for those accounts that lack them.

3.1.2 Motivations for Structured Propositions

Structured propositions are motivated by an incredibly simple and appealing idea; roughly speaking the meaning of an expression consists of the meaning of its component words and the way in which the expression is built up from these components. Nowadays, structured proposition accounts are often motivated by the failures of intensional accounts of meaning but discussions of structured propositions predate the widespread adoption of intensional accounts. In this section, we first mention some key historical and contemporary accounts of structured propositions. Then we outline the issues faced by intensional accounts and how the structured proposition approach responds to them.

Frege's concept of sense is usually taken to be the earliest analogue of structured propositions (King 2019). However, Frege provides little information about how to determine or distinguish senses. Subsequently Russell, with whom structured propositions are most frequently associated, develops his own account(s) with (Russell 1903) providing the currently best-regarded presentation. An essential difference between Frege and Russell's views is that for Russell, objects and properties themselves (rather than their senses) can be the components of propositions. This relates to a Millian view of names whereby their semantic content is exhausted by their reference.

Later, Church produces and revises what he calls the logic of sense and denotation based on Frege's work (Church 1951; 1973; 1974; 1993). This logic characterises senses with an emphasis on their relationship to the reference (or denotation) of expressions and the ability to construct higher-order senses. The extensive use of type theory combined with the generality of this work make it a challenging read.

[5] This debate occurs both between advocates of structured propositions and their opponents (who might hold that the structural component is not semantic) and between advocates of *different* structured proposition accounts. (King 2019) discusses several positions in this debate.

Finally, we reach the contemporary accounts. There are the neo-Russelian accounts which includes the work of (Soames 1987) and (Salmón 1986a; 1992). We also have accounts that are more sympathetic to a broadly intensional account, positing that it simply needs to be supplemented with additional structure. Such accounts include that of (Cresswell 1985) where structure is only salient for sentences containing propositional attitudes. This account is also inspired by the intensional isomorphism of (Carnap 1956). Yet another approach is taken in (Moschovakis 2006), which introduces a uniquely computational perspective.

As mentioned, the challenges faced by intensional accounts provide significant motivation for the structured proposition approach. Recall that intensional accounts hold that the semantic content of some linguistic item (name, predicate, sentence etc.) is the function from metaphysically possible worlds to its extension at that world. As such intensional accounts cannot distinguish statements that are necessarily equivalent, i.e., true in all the same metaphysically possible worlds. Even if one is willing to accept this for the statements themselves, the issue is magnified when they are embedded into propositional attitudes like belief.

Posed as a response to the failures of intensional accounts, structured proposition approaches are often contrasted with impossible world accounts. Rather than introducing structure, impossible worlds produce more fine-grained meanings by considering truth conditions across a broader range of circumstances. Semantic contents are still functions from circumstances to extensions in that circumstance, but unlike possible worlds, these circumstances need not be maximal, consistent or metaphysically possible. In principle, structured propositions don't conflict with this move from possible worlds to circumstances. Instead, the tension is that adherents of structured proposition accounts hold that such a move is inadequate when accounting for differences in meaning.

Let us briefly review a few motivating examples where intensional accounts appear to be deficient. Recall such cases involve pairs of necessarily equivalent statements which appear to have different "meanings." Typically, these statements are necessary truths or falsehoods or complex statements involving such necessities. Common examples of necessary truths include identity statements and mathematical theorems. Since metaphysical necessity is weaker than logical necessity, logically equivalent statements also provide examples of necessarily equivalent statements.

The following statements, known as Frege's first puzzle, provide a traditional form of the issue for identity statements.

(1) Hesperus is Hesperus.

(2) Hesperus is Phosphorus.

Given that ‘Hesperus’ and ‘Phosphorus’ are both names for the planet Venus and identity (moreover self-identity) is a necessary property, we have that both statements are true in every possible world. Therefore these statements have the same intension. Perhaps surprisingly, structured proposition accounts do not in general resolve this issue. In fact those committed to Millianism or rigid designation are more forcefully required to accept that (1) and (2) have the same semantic content because they are committed to the names invoked as having the same semantic content.[6]

Likewise, mathematical theorems are widely held to be necessarily true and their negations necessarily false. As such all mathematical statements have trivial intensions, either they are true in every possible world or in none of them. Thus on the intensional account, all true mathematical statements (i.e., theorems) are synonymous; as are all false mathematical statements. This is clearly unsatisfactory for anyone attempting to discuss the meaning of mathematical theorems. In the final section of this chapter we assess Soames’ account with respect to its judgements about mathematical statements.

Finally, statements with trivial intensions can be combined with other statements to produce pairs of necessarily equivalent statements. For example in the following case where we add a tautology as a conjunct.[7]

(3) Harry is crying.

(4) Harry is crying and if Evie is late then Evie is late.

Since the tautological conjunct is always true, (4) is true in precisely the same worlds as (3) so they have the same intension. Such examples plus those involving logically equivalent rephrasings of contingent statements are readily handled by the structured proposition approach as we see in section 3.2.3.

[6]Options available to those endorsing a structured proposition account include paraphrasing and retreating to the analysis of propositional attitudes. As an example of the former, (Soames 1987) uses “‘Hesperus’ and ‘Phosphorus’ are co-referential” which avoids synonymy with the corresponding statement containing two instances of ‘Hesperus’ by mentioning rather than using the name. As an example of the latter, (Salmón 1986b) defends a view where the belief relation is roughly the existential generalisation of a ternary relation, among believers, propositions, and a “mode of acquaintance.”

[7]This case works best when one imagines it said by an exhausted caregiver of two young children.

These issues are all exacerbated when such necessarily equivalent statements are embedded within propositional attitudes such as believes, desires, knows. For instance, consider the following statements.

(5) Jamie believes Harry is crying.

(6) Jamie believes Harry is crying and if Evie is late then Evie is late.

Here the difference between the embedded statements is more immediately apparent since it is plausible that (5) and (6) have different truth values. Various analyses and solutions have been suggested to this issue, see e.g. (Nelson 2024) for an introduction to some of these.[8]

Since our focus is on mathematics where propositional attitudes are not necessary to motivate distinguishing more finely than intensions, we exclude propositional attitudes in what follows. This allows us to avoid certain complications involved in the analysis of propositional attitudes.

3.2 Soames' Account of Structured Propositions

In (Soames 1987), the meaning of an expression is a function from contexts and variable assignments to semantic contents and the semantic contents of sentences are structured propositions consisting of the semantic contents of its components. Rephrasing this under the modelling perspective, we might say that Soames defines structured propositions as an model of sentences' semantic contents. Following Kaplan's terminology, we'll say he also defines the *character* of an expression, the function from contexts and variable assignments to their semantic content.

In particular on Soames' account the propositional content of a formula, relative to a context and variable assignment, is called a structured proposition. A structured proposition is a tree consisting of the propositional contents of all occurrences of its semantically significant parts.[9] As the name suggests a structured proposition keeps track of structural information, typically by reflecting the syntactic structure of an expression.

[8]For example, one may deny (6) (while accepting (5)) on the basis that it is "about" Evie (and lateness) and it is possible that Jamie has no such belief.

[9]See definition 3.2.2 where we define *constituents*, our more precise terminology for "semantically significant parts." Trees are equivalent to the nested ordered tuples more commonly used by Soames and others, we prefer trees for representing the structure in a more visual manner and as a more informative name than 'complex.'

These propositional contents can then be evaluated for their value at a world, which is equal to the extension of the expression at that world on an intensional account.[10] This means we can still recover an expression's intension from its semantic content, so semantic contents supplement intensions rather than wholly replacing them.

This early work on structured propositions predates Soames' shift to a heavily Contextualist position, where pragmatic enrichments are more central.[11] On this later view, contexts play a greater role in establishing the proposition expressed by an utterance, but for this work, the context parameter isn't significantly utilised in the formal account. Combining this with the later focus of this chapter on mathematics where context relativity is minimal, we suppress the context parameter to streamline the formal presentation.[12] We also remove the belief relation due to an apparent circularity which requires additional work to eliminate.

We can now expand figure 3.1 to include characters and intensions as in figure 3.2. This preserves the most important functions along the midline. Namely $\mathfrak{p}(\cdot,\cdot)$ which maps a lexical item and variable assignment to its semantic content relative to that variable assignment, see definition 3.2.5. Then $Val^{\cdot}(\cdot)$ maps a semantic content(e.g. a proposition) and a world to the value of that content at that world, see definition 3.2.6.

$$\begin{tikzcd}
 & \mathsf{Character} \times \mathcal{V} \times \mathcal{W} \arrow[d, "\langle @(\cdot,\cdot), id_{\mathcal{W}} \rangle"] & \\
\mathsf{Linguistic\ Item} \times \mathcal{V} \times \mathcal{W} \arrow[ur, "\langle \mu(\cdot), id_{\mathcal{V}}, id_{\mathcal{W}} \rangle"] \arrow[r, "\langle \mathfrak{p}(\cdot, g), id_{\mathcal{W}} \rangle"'] & \mathsf{Semantic\ Content} \times \mathcal{W} \arrow[r, "Val^{W}(\cdot)"] \arrow[d] & \mathsf{Value} \\
 & \mathsf{Intension} \times \mathcal{W} \arrow[ur] &
\end{tikzcd}$$

Figure 3.2: Relationships between various semantic entities.

Throughout the figure, we use identity functions, $id_{\mathcal{V}}, id_{\mathcal{W}}$, to carry indices through to where they are needed. In the upper triangle, we introduce characters. In particular, we have the function $\mu(\cdot)$ which maps lexical items to their character. Then since characters are simply functions from variable assignments to propositional contents, the function $@(\cdot,\cdot)$ takes us from a character and a variable assignment to the corresponding propositional content.

[10]We call this the semantic content's value rather than its extension to help maintain the distinction between an expression and its semantic content.

[11]See (Soames 2008) for an overview of this development.

[12]It is easy to reconstruct the account relativised to contexts in addition to variable assignments by simply introducing a context parameter wherever there is a variable assignment.

It does this simply by applying (inputting) the second component (the variable assignment) to the first component (the character). So $@(\mu(\varphi), g) = \mu(\varphi)(g) = \mathfrak{p}(\varphi, g)$.

The figure's lower triangle reflects that we can map from propositional contents to intensions. Once we define $\mathfrak{p}$, it will be obvious how to build this function but it requires some tedious type lifting since Soames treats some lexical items extensionally so we leave that for the reader. Applying the usual evaluation function to an intension and a world will then lead to the same final value.

3.2.1 Syntax

In this section, we define the language, terms and formulas used by Soames (excluding propositional attitudes). Instead of constant symbols, Soames' language contains *rigid designator* symbols, this terminology is used to distinguish the fact that the interpretation of such symbols is fixed across worlds. We also follow Soames in omitting function symbols.

Definition 3.2.1. The language $\mathcal{L} = (L, ar)$ consists of a signature L and a function $ar : L \to \mathbb{N}$ assigning arities to the relation symbols. The signature L consists of the logical symbols, $\neg, \wedge, \vee, \rightarrow, \exists, \forall$ and the following:[13]

1. An infinite collection of variables: $\nu_1, \nu_2, \ldots$.
2. An infinite collection of rigid designation symbols: $d_1, d_2, \ldots$.
3. An infinite collection of relation symbols (of all arities): $R_1, R_2, \ldots$.

We also allow ourselves to use parentheses as needed for readability.

From such a language, we define its *referential terms, abstraction terms* and *formulas.* Abstraction terms, which make use of the λ symbol, are used for variable binding both in quantified formulas and building complex predicates like 'tripped and fell.'

Definition 3.2.2.
Every variable v_i and rigid designation symbol d_i is a $\mathcal{L}$-*referential term.*[14]
We use t or t_i for arbitrary referential terms.

[13]Including logical symbols in our signature simplifies some later definitions. However, λ, which we introduce later, is not an official symbol in our language (also mostly for convenience).

[14]To preempt definition 3.2.5 slightly, (directly) referential terms are terms which have their extension as their semantic content.

The $\mathcal{L}$*-formulas* and $\mathcal{L}$*-abstraction terms* are defined (double) inductively as follows:

Atomic	If R is an n-ary *relation symbol*, and $t_1, \ldots, t_n$ are *referential terms* then $R(t_1, \ldots, t_n)$ is a *formula*.
Negation	If φ is a *formula* and then $\neg\varphi$ is a *formula*.
Binary Connective	If φ, ψ are *formulas* and $\circ \in \{\wedge, \vee, \rightarrow\}$ is a binary connective then $\circ(\varphi, \psi)$ is a *formula*.[18]
Abstraction	If φ is a *formula* and x is a *variable* then $\lambda x.\varphi$ is an *abstraction term*.
Quantifier	If $\lambda x.\varphi$ is an *abstraction term* and $Q \in \{\exists, \forall\}$ is a quantifier then $Q\lambda x.\varphi$ is a *formula*.
Application	If $\lambda x.\varphi$ is an *abstraction term* and t is a *referential term* then $\lambda x.\varphi(t)$ is a *formula*.

For readability, we allow ourselves to write $\varphi \wedge \psi$ for $\wedge(\varphi, \psi)$ and $\exists x\varphi$ for $\exists\lambda x.\varphi$ but use the above official definitions to simplify subsequent definitions and ensure correct typing.

We also introduce the following terminology. A string of symbols is an *expression* (of our language) if it is a referential term, an abstraction term or a formula. A string of symbols is a *constituent of our language* if it is an abstraction term or in L, i.e., if it has the form $\lambda x.\varphi$ for some formula φ, or is a rigid designator, a variable, a relation symbol, or a logical symbol. A string of symbols is a *constituent of an expression* if it is a constituent of $\mathcal{L}$ and occurs in the expression. Constituents are the formalisation of our lexical items and we combine them to form expressions.

For example $\lambda x.\wedge(Tx, Fx)(s)$ is the formula obtained by applying the referential term s to the abstraction term $\lambda x.\wedge(Tx, Fx)$, all of which are expressions. This abstraction term is formed by prefixing λx to the formula $\wedge(Tx, Fx)$, which in turn is formed by taking the conjunction of Tx and Fx. These atomic formulas are formed by placing the unary relation symbols T and F in front of the variable x. If we translate s as Sam, T as (to have) tripped and F as (to have) fell then this formula translates to 'Sam tripped and fell.'[19] The constituents of $\lambda x.\wedge(Tx, Fx)(s)$ are $\lambda x.\wedge(Tx, Fx), s, T, F, \wedge, x$.[20]

[18]Polish notation is used to simplify later definitions. We allow ourselves to use infix notation, e.g., $\varphi \wedge \psi$ as an unofficial "abbreviation."

[19]Or more clunkily, Sam is such that they tripped and fell.

[20]Often a further distinction is drawn whereby $\lambda x.\wedge(Tx, Fx)$ and s are the *immediate* constituents of the expression. Note that $\wedge(Tx, Fx)$ is not a constituent of the expression despite being a sub-formula.

3.2.2 Semantics

In this section, we define the semantics of $\mathcal{L}$. That is, we show how to construct structured propositions for $\mathcal{L}$-formulas and how to calculate the truth value of a structured proposition at a world, W on Soames' account. In definition 3.2.5, we define the function $\mathfrak{p}(\cdot, g)$ returning the semantic contents of our expressions, relative to a variable assignment g. The function is defined first on the constituents of our language with the inductive step on complex expressions done by building a tree from the semantic contents of its constituent expressions. Then in definition 3.2.6 we define a valuation function Val^W which gives the value of a given propositional content at a world W.

First there are some supplemental definitions to set up. We define a universe, containing all the worlds we need and giving interpretations for the non-logical elements of our language (the rigid designator and relation symbols).[21] Due to the emphasis on rigid designators, the denotation function that specifies their extensions is its own component of the universe, separate from the world-relative interpretations of relation symbols.[22] Although we suppress the context-relativity in later definitions, the universe needs to contain a set of contexts so that our characters are well-defined.

Definition 3.2.3. A *universe*, $\mathcal{U} = \langle \mathcal{C}, \mathcal{W}, U, \mathfrak{D}, \{\cdot^W\}_{W\in\mathcal{W}}\rangle$, of $\mathcal{L}$ consists of:

- a set of *contexts* $\mathcal{C}$,
- a set of *worlds* $\mathcal{W}$,
- a set of objects U serving as the *domain* of $\mathcal{U}$,
- a *denotation function* $\mathfrak{D} : \{d_i\}_{i\in\omega} \to U$, which gives the world-independent extension of our rigid designators.
- a family of functions $\{\cdot^W\}_{W\in\mathcal{W}}$, where each $\cdot^W$ is the *interpretation function* of W which associates each relation symbol with their extension, at that world. That is, if R is an n-ary relation symbol then $R^W \in \mathcal{P}(U^n)$.[23]

  So the intension of R is $\overline{\mathsf{R}} : \mathcal{W} \to \mathcal{P}(U^n)$, the function mapping a world W to R^W, i.e., the extension of R at W.

[21] What we are calling a universe is often called a model, this terminology is chosen to avoid confusion with the scientific usage of 'models.'

[22] This also motivates the decision to have a fixed domain for the universe, technically it is only necessary that there be a common intersection of each world's domain which contains all the referents for the rigid designators.

[23] Equivalently $R^W : U^n \to Bool$.

We also introduce variable assignments.

Definition 3.2.4. A *variable assignment* on a universe $\mathcal{U}$ is a function $g : \{\nu_i\}_{i\in\omega} \to U$ mapping variables in our language to objects in the universe's domain. We call $g(x)$ the *extension* of x (relative to g). Let $\mathcal{V}$ denote the set of all variable assignments in $\mathcal{U}$.

Let $g\{x := \mathsf{o}\}$ denote the variable assignment which differs from g at most by assigning $\mathsf{o} \in U$ to the variable x. We call $g\{x := \mathsf{o}\}$ an *x-variant of g.*

Now, we can build structured propositions for formulas.[24] In general, a structured proposition of a formula φ is a tree labelled with the semantic contents of φ's constituents. We define such trees inductively below in definition 3.2.5 and represent them graphically in figure 3.3 for each kind of formula. Before giving the full definition let's see a rough characterisation of the semantic contents of our language's constituents.

Since Soames is committed to direct reference, the semantic content of a referential term (variables and rigid designators) is its extension (given by g or $\mathfrak{D}$ respectively).[25] The semantic content of a logical symbol are also its extension. That is for connectives, their semantic content is the corresponding truth-conditional function and for quantifiers, their semantic content is the characteristic function of non-empty sets (for $\exists$) and the entire domain U (for $\forall$).[26] The semantic content of a relation symbol R is its intension $\overline{\mathsf{R}}$. Finally, the most complex case is abstraction terms of the form $\lambda x.\varphi$. Roughly speaking, its semantic content is the propositional function mapping an object o to the structured proposition of φ where x has been replaced with o.

Definition 3.2.5 (Propositional contents)**.** Let $\mathfrak{p}(\cdot, g)$ denote the function from constituents and formulas of $\mathcal{L}$ to their semantic contents relative to a variable assignment g. Given a constituent or formula γ, $\mathfrak{p}(\gamma, g)$ is its *semantic content*, relative to g. In the case of a formula φ, $\mathfrak{p}(\varphi, g)$ is also called the *structured proposition* of φ, relative to g.

[24]Note that because semantic contents are always defined with respect to a variable assignment even open formulas are attributed a structured proposition.

[25]See (Soames 2005) for his later view regarding pragmatic enrichment of this semantic content.

[26]There are two points to note here. First, the semantic contents for logical symbols are described as if they act on the *value* of semantic contents (truth values and sets) rather than on the semantic contents themselves (propositions and propositional functions). This reflects the second point that this choice (to define the semantic contents for logical symbols in this way) isn't actually doing any formal work. It's not until we calculate the value of the structured propositions that we need to distinguish the different logical operations. So we could put any placeholder we want here (including just leaving in the symbol), as long as it is different for the different logical symbols (and we define the valuation function correctly).

We define $\mathfrak{p}(\cdot, g)$ inductively, the base case covers the different kinds of constituents of our language (excluding abstraction terms) and the inductive case covers formulas, which are strings of those constituents, and abstraction terms, which contain formulas.

Constituents

1. For a rigid designator d, $\mathfrak{p}(d, g) := \mathfrak{D}(d) \in U$
2. For a variable x, $\mathfrak{p}(x, g) := g(x) \in U$
3. For an n-ary relation symbol R, $\mathfrak{p}(R, g) := \overline{\mathsf{R}} : W \to \mathcal{P}(U^n)$ is the intension of R.[27]
4. For the negation symbol, $\mathfrak{p}(\neg, g) := \mathsf{Neg} : \mathit{Bool} \to \mathit{Bool}$ is the truth-conditional function for negation.
5. For a binary connective $\circ$, $\mathfrak{p}(\circ, g)$ is the corresponding truth-conditional function from $(\mathit{Bool} \times \mathit{Bool}) \to \mathit{Bool}$. We denote these $\mathsf{Conj}, \mathsf{Disj}, \mathsf{Imp}$ for $\wedge, \vee, \to$ respectively.
6. For a quantifier Q, $\mathfrak{p}(Q, g)$ is the corresponding function on $\mathcal{P}(U) \to \mathit{Bool}$. That is, $\mathfrak{p}(\exists, g) := \mathsf{Some}$ is the characteristic function of non-empty sets and $\mathfrak{p}(\forall, g) := \mathsf{All}$ is the function which takes value 1 on input U and 0 otherwise.

Formulas and abstraction terms

7. For a formula $\varphi = \alpha(\beta_1, \ldots \beta_n)$ where α, β_i are constituents or formulas of $\mathcal{L}$, $\mathfrak{p}(\varphi, g) = \mathfrak{p}(\alpha(\beta_1, \ldots \beta_n), g) = \langle \mathfrak{p}(\alpha, g), \langle \mathfrak{p}(\beta_1, g), \ldots, \mathfrak{p}(\beta_n, g)\rangle\rangle$.
8. For an abstraction term, $\lambda x.\psi$, $\mathfrak{p}(\lambda x.\psi, g) := \mathfrak{a}_{\psi}^{\mathsf{g,x}} : U \to \mathrm{Prop}$ is a one-place propositional function mapping $\mathsf{o} \mapsto \mathfrak{p}(\psi, g\{x := \mathsf{o}\})$. That is, $\mathfrak{a}_{\psi}^{\mathsf{g,x}}$ maps an object o to the semantic content of ψ relative to the variable assignment g' which differs from g at most by setting $g'(x) = \mathsf{o}$.

Figure 3.3 provides a more graphical depiction of structured propositions for each kind of formula, with the nested tuple notation also given in the caption.[28]

[27] This is the key location for context relativity so the full account would have $\mathfrak{p}(R, C, g) = \overline{\mathsf{R}}_C$ where $\overline{\mathsf{R}}_C$ is the intension of R in context C. For example if T stands for "is tall" then $\overline{\mathsf{T}}_C$ could be (the intension of) "is tall (for a basketballer)" or "is tall (for a 5-year-old)." We drop the context parameter moving forward, since predicates are the only lexical components they affect and because in most cases the context can be made explicit in the formalisation instead. We could use B ="is tall for a basketballer" in the original formalisation rather than translating it simply as T ="is tall."

[28] We have five kinds of formulas based on which syntactic operation was applied last in their construction, these are atomic, application, negation, binary connective and quantification. We use conjunction and the existential quantifier as representative for their formula kinds.

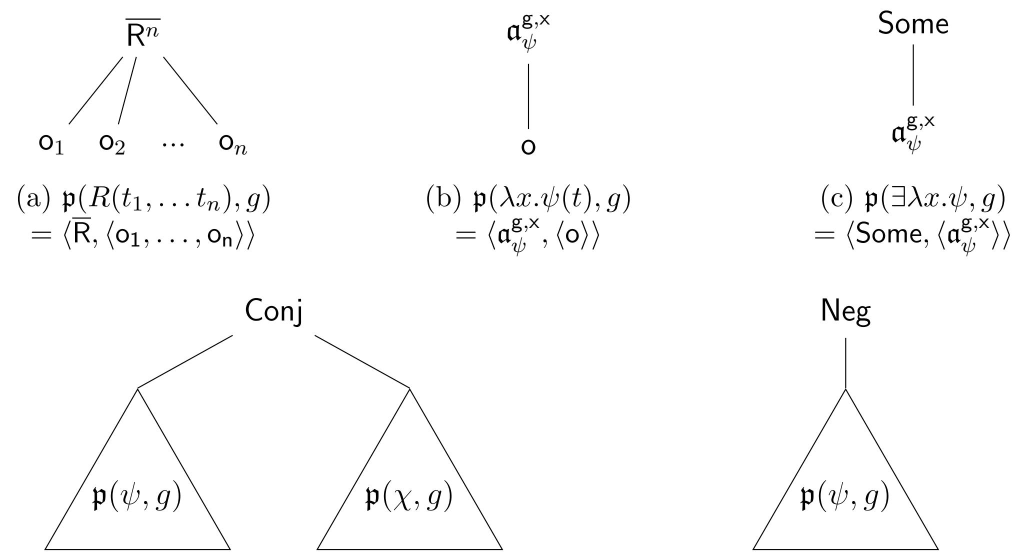

(a) $\mathfrak{p}(R(t_1, \ldots t_n), g) = \langle \overline{\mathsf{R}}, \langle \mathsf{o_1}, \ldots, \mathsf{o_n} \rangle\rangle$

(b) $\mathfrak{p}(\lambda x.\psi(t), g) = \langle \mathfrak{a}_\psi^{\mathsf{g,x}}, \langle \mathsf{o} \rangle\rangle$

(c) $\mathfrak{p}(\exists \lambda x.\psi, g) = \langle \mathsf{Some}, \langle \mathfrak{a}_\psi^{\mathsf{g,x}} \rangle\rangle$

(d) $\mathfrak{p}(\wedge(\psi, \chi), g) = \langle \mathsf{Conj}, \langle \mathfrak{p}(\psi, g), \mathfrak{p}(\chi, g) \rangle\rangle$

(e) $\mathfrak{p}(\neg\psi, g) = \langle \mathsf{Neg}, \langle \mathfrak{p}(\psi, g) \rangle\rangle$

Figure 3.3: Structured propositions for different kinds of formulas.

Now that we have the semantic content of any expression, we need to be able to determine the value of these semantic contents at a world. This value should correspond to the usual extension of an expression on an intensional account of meaning. Thus we give the following definition of Val^W.

Definition 3.2.6.

The *valuation function* Val^W of a world W is defined inductively as follows:

1. $Val^W(\langle \overline{\mathsf{R}}, \langle \mathsf{o_1}, \ldots, \mathsf{o_n} \rangle\rangle) = \begin{cases} 1 & \text{if } \langle \mathsf{o_1}, \ldots, \mathsf{o_n} \rangle \in R^W; \\ 0 & \text{otherwise.} \end{cases}$

2. $Val^W(\langle \mathsf{Neg}, \langle \mathfrak{p}(\psi, g) \rangle\rangle) = 1 - Val^W(\mathfrak{p}(\psi, g))$.

3. $Val^W(\langle \mathsf{Conj}, \langle \mathfrak{p}(\psi, g), \mathfrak{p}(\chi, g) \rangle\rangle) = min\{Val^W(\mathfrak{p}(\psi, g)), Val^W(\mathfrak{p}(\chi, g))\}$.

4. $Val^W(\langle \mathsf{Disj}, \langle \mathfrak{p}(\psi, g), \mathfrak{p}(\chi, g) \rangle\rangle) = max\{Val^W(\mathfrak{p}(\psi, g)), Val^W(\mathfrak{p}(\chi, g))\}$.

5. $Val^W(\langle \mathsf{Imp}, \langle \mathfrak{p}(\psi, g), \mathfrak{p}(\chi, g) \rangle\rangle) = max\{1 - Val^W(\mathfrak{p}(\psi, g)), Val^W(\mathfrak{p}(\chi, g))\}$.

6. $Val^W(\langle \mathfrak{a}_\psi^{\mathsf{g,x}}, \langle \mathsf{o} \rangle\rangle) = Val^W(\mathfrak{a}_\psi^{\mathsf{g,x}}(\mathsf{o}))$.
   By definition $\mathfrak{a}_\psi^{\mathsf{g,x}}(\mathsf{o}) = \mathfrak{p}(\psi, g\{x := \mathsf{o}\})$, so $Val^W(\langle \mathfrak{a}_\psi^{\mathsf{g,x}}, \langle \mathsf{o} \rangle\rangle) = Val^W(\mathfrak{p}(\psi, g\{x := \mathsf{o}\}))$.

7. $Val^W(\langle \mathsf{Some}, \langle \mathfrak{a}_\psi^{\mathsf{g,x}} \rangle\rangle) = \begin{cases} 1 & \text{if there exists } \mathsf{o} \in U \text{ such that } Val^W(\mathfrak{a}_\psi^{\mathsf{g,x}}(\mathsf{o})) = 1; \\ 0 & \text{otherwise.} \end{cases}$

Equivalently, $Val^W(\langle \mathsf{Some}, \langle \mathfrak{a}_\psi^{\mathsf{g,x}} \rangle\rangle) = max\{Val^W(\mathfrak{p}(\psi, g\{x := \mathsf{o}\}) : \mathsf{o} \in U\}$.

8. $Val^W(\langle \mathsf{All}, \langle \mathfrak{a}_\psi^{\mathsf{g,x}} \rangle\rangle) = \begin{cases} 1 & \text{if every } \mathsf{o} \in U \text{ is such that } Val^W(\mathfrak{a}_\psi^{\mathsf{g,x}}(\mathsf{o})) = 1; \\ 0 & \text{otherwise.} \end{cases}$

   Equivalently, $Val^W(\langle \mathsf{All}, \langle \mathfrak{a}_\psi^{\mathsf{g,x}} \rangle\rangle) = min\{Val^W(\mathfrak{p}(\psi, g\{x := \mathsf{o}\}) : \mathsf{o} \in U\}$.

The output of Val^W is called the *value* of that semantic content at W. We can use this to define the satisfaction relation so that $W \models \mathfrak{p}(\varphi, g)$ if and only if $Val^W(\mathfrak{p}(\varphi, g)) = 1$.

3.2.3 Non-synonymy Examples

Now we are able to define synonymy on Soames' account and provide examples where it successfully distinguishes intensionally equivalent sentences (at least by our perspective). Some readers may not agree with the judgements of Soames account but the examples still provide insight into how the account is applied.

Soames doesn't explicitly give a definition of synonymy but since "meanings" on his account are functions from contexts and variable assignments to semantic contents, it must be that two expressions are synonymous if they have the same function as their "meanings," i.e., they have the same character. More concretely, they are synonymous if they have the same semantic content for every context and variable assignment.

Definition 3.2.7. Two constituents, α, β are *synonymous* if and only if for all contexts C and variable assignments g they have the same semantic content, i.e.,

$$\forall C, g,\ \mathfrak{p}(\alpha, C, g) = \mathfrak{p}(\beta, C, g).$$

Lacking any direct information from (Soames 1987), we assume that contexts only serve to fix referents for indexicals. Similarly, the variable assignment g only serves to give referents for (free) variables. As such, closed formulas without indexical expressions don't change relative to contexts or variable assignments. Thus if φ is a closed formula without indexical expressions we refer to $\mathfrak{p}(\varphi)$ as *the* structured proposition of φ. Such formulas are the focus of this chapter.

Definition 3.2.8. Two sentences lacking contextual vocabulary φ, ψ are *synonymous* if and only if they are assigned the same structured proposition. Two structured propositions are the same if and only if they have identical tree structures and the label on each node is equal to the label on the corresponding node of the other tree.

Let's start with some basic examples from mathematics where Soames' account is successful.

(7) Six is even.

(8) Six is perfect.

(9) Eight is even.

Here, we want to distinguish (7) from (8) because they contain different relations and (7) from (9) because they contain non-coreferential names/ objects. These statements have the following structured propositions.

(7⋆) $\langle \overline{\mathsf{even}}, \langle \mathsf{6} \rangle \rangle$

(8⋆) $\langle \overline{\mathsf{perfect}}, \langle \mathsf{6} \rangle \rangle$

(9⋆) $\langle \overline{\mathsf{even}}, \langle \mathsf{8} \rangle \rangle$

Indeed, we get the desired result. We have $\langle \overline{\mathsf{even}}, \langle \mathsf{6} \rangle \rangle \neq \langle \overline{\mathsf{perfect}}, \langle \mathsf{6} \rangle \rangle$, since $\overline{\mathsf{even}} \neq \overline{\mathsf{perfect}}$ and $\langle \overline{\mathsf{even}}, \langle \mathsf{6} \rangle \rangle \neq \langle \overline{\mathsf{even}}, \langle \mathsf{8} \rangle \rangle$, since $6 \neq 8$.

Now let's return to some of our motivating examples.

(3) Harry is crying.

(4) Harry is crying and if Evie is late then Evie is late.

For clarity let's formalise these statements first into our language $\mathcal{L}$, allowing ourselves the benefit of using infix notation for the connectives.

($3^\dagger$) $crying(Harry)$

($4^\dagger$) $crying(Harry) \wedge (late(Evie) \rightarrow late(Evie))$

These sentences have the following structured propositions (see also figure 3.4).

(3⋆) $\langle \overline{\mathsf{crying}}, \langle \mathsf{Harry} \rangle \rangle$

(4⋆) $\langle \mathsf{Conj}, \langle \langle \overline{\mathsf{crying}}, \langle \mathsf{Harry} \rangle \rangle, \langle \mathsf{Imp}, \langle \langle \overline{\mathsf{late}}, \langle \mathsf{Evie} \rangle \rangle, \langle \overline{\mathsf{late}}, \langle \mathsf{Evie} \rangle \rangle \rangle \rangle \rangle \rangle$

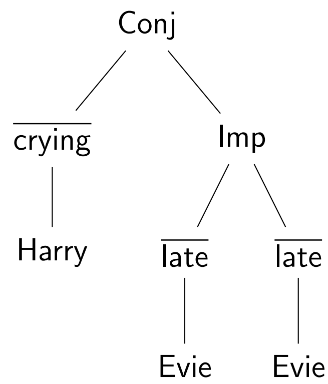

Figure 3.4: Structured proposition for (4⋆)

Clearly these two structured propositions are not equal, because they are different sizes. So these two statements are not synonymous (despite being intensionally and even logically equivalent).

Let's consider another example.

(10) Lex is funny and not funny.

(11) Lex is lonely and not lonely.

This time our statements are intensionally equivalent because they're both necessarily false, but we may want to distinguish their meanings because they talk about different properties.[29] We can use lambda abstraction to formalise 'is funny and not funny' into a single complex relation, rather than rephrasing the sentence as 'Lex is funny and Lex is not funny' as we typically do in first order logic without lambda terms. And similarly for 'is lonely and not lonely.' This produces the following formulas.

(10†) $\lambda x.(funny(x) \wedge \neg funny(x))(Lex)$

(11†) $\lambda x.(lonely(x) \wedge \neg lonely(x))(Lex)$

Applying $\mathfrak{p}$ to these formulas produces the following structured propositions.

(10⋆) $\langle \mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Fx \wedge \neg Fx}}, \langle \mathsf{Lex} \rangle \rangle$ such that $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Fx \wedge \neg Fx}}(\mathsf{o}) = \langle \mathsf{Conj}, \langle \langle \overline{\mathsf{funny}}, \langle \mathsf{o} \rangle \rangle, \langle \mathsf{Neg}, \langle \overline{\mathsf{funny}}, \langle \mathsf{o} \rangle \rangle \rangle \rangle \rangle$

[29] Of course, both complex relations have empty extensions because they require an individual to have two contradictory properties.

(11⋆) $\langle \mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Lx \wedge \neg Lx}}, \langle \mathsf{Lex} \rangle\rangle$ such that $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Lx \wedge \neg Lx}}(\mathsf{o}) = \langle \mathsf{Conj}, \langle\langle \overline{\mathsf{lonely}}, \langle \mathsf{o} \rangle\rangle, \langle \mathsf{Neg}, \langle \overline{\mathsf{lonely}}, \langle \mathsf{o} \rangle\rangle\rangle\rangle\rangle$

These structured propositions are not equal because $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Fx \wedge \neg Fx}} \neq \mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Lx \wedge \neg Lx}}$. For these propositional functions to be equal they would have to output the same structured proposition for every input. However, $\langle \mathsf{Conj}, \langle\langle \overline{\mathsf{funny}}, \langle \mathsf{o} \rangle\rangle, \langle \mathsf{Neg}, \langle \overline{\mathsf{funny}}, \langle \mathsf{o} \rangle\rangle\rangle\rangle\rangle \neq \langle \mathsf{Conj}, \langle\langle \overline{\mathsf{lonely}}, \langle \mathsf{o} \rangle\rangle, \langle \mathsf{Neg}, \langle \overline{\mathsf{lonely}}, \langle \mathsf{o} \rangle\rangle\rangle\rangle\rangle$. This is because $\overline{\mathsf{funny}} \neq \overline{\mathsf{lonely}}$, since there is some possible world W where someone is funny and not lonely so $lonely^W \neq funny^W$. Since they have distinct structured propositions, these statements are not synonymous.

Similarly we can distinguish either of those statements from the following statement, which involves a more obviously distinct relation.

(12) Lex is taller than themself.

Like (10) and (11) this statement applies a necessarily false relation to Lex, but the relation is not a complex relation although it still requires lambda abstraction in its formalisation as follows.

($12^\dagger$) $\lambda x.taller(x, x)(Lex)$

This formula is assigned the following structured proposition.

(12⋆) $\langle \mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Txx}}, \langle \mathsf{Lex} \rangle\rangle$ such that $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Txx}}(\mathsf{o}) = \langle \overline{\mathsf{taller}}, \langle \mathsf{o}, \mathsf{o} \rangle\rangle$

This structured proposition differs from the previous two because $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Txx}}$ differs from $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Lx \wedge \neg Lx}}$ and $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Fx \wedge \neg Fx}}$. In particular we can see that $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Txx}}(\mathsf{o})$ has three components whereas $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Lx \wedge \neg Lx}}(\mathsf{o})$ and $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{Fx \wedge \neg Fx}}(\mathsf{o})$ both have six.

Similarly, we can capture differences between logically equivalent statements that don't rely on necessary truths.

(13) There's someone who is tall but doesn't play basketball.

(14) It's not the case that everyone either plays basketball or isn't tall.

These statements are formalised as follows.

($13^\dagger$) $\exists \lambda x(tall(x) \wedge \neg basketballer(x))$

(14†) $\neg\forall\lambda x(basketballer(x) \vee \neg tall(x))$

Leading to these corresponding structured propositions (see also figure 3.5).

(13$\star$) $\langle \mathsf{Some}, \langle \mathfrak{a}^{\mathsf{g,x}}_{\mathsf{T\wedge\neg B}} \rangle\rangle$ where $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{T\wedge\neg B}}(\mathsf{o}) = \langle \mathsf{Conj}, \langle\langle \overline{\mathsf{tall}}, \mathsf{o}\rangle, \langle \mathsf{Neg}, \langle \overline{\mathsf{basketballer}}, \mathsf{o}\rangle\rangle\rangle\rangle$

(14$\star$) $\langle \mathsf{Neg}, \langle \mathsf{All}, \langle \mathfrak{a}^{\mathsf{g,x}}_{\mathsf{B\vee\neg T}} \rangle\rangle\rangle$ where $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{B\vee\neg T}}(\mathsf{o}) = \langle \mathsf{Disj}, \langle\langle \overline{\mathsf{basketballer}}, \langle \mathsf{o}\rangle\rangle, \langle \mathsf{Neg}, \langle \overline{\mathsf{tall}}, \langle \mathsf{o}\rangle\rangle\rangle\rangle\rangle$;

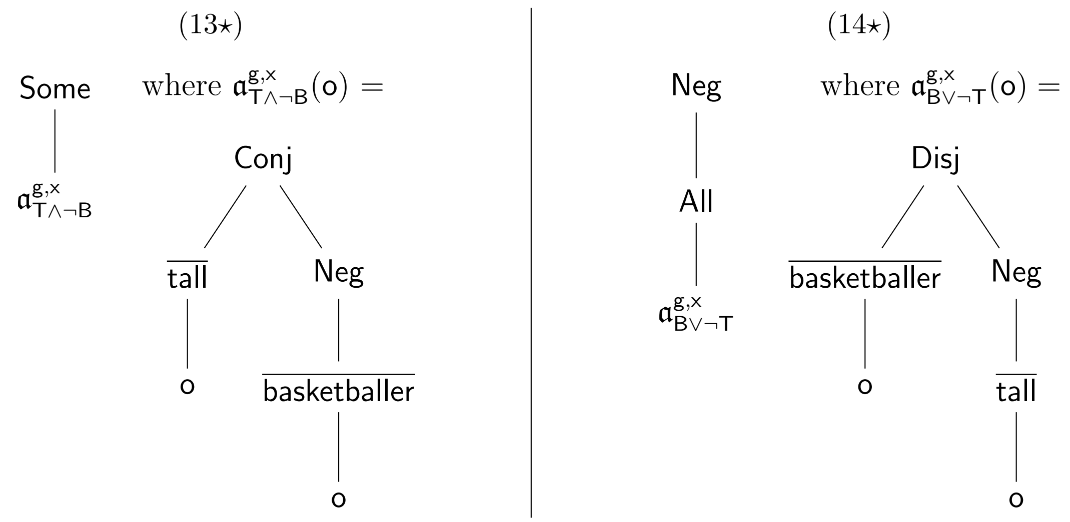

Figure 3.5: Structured propositions for (13$\star$) and (14$\star$)

These structured propositions are distinct as they have different sizes. So the corresponding statements are not synonymous. Hence, Soames' account can distinguish between logically equivalent statements as desired.

3.3 Problem Cases for Soames' Account

This section covers the more unfavourable analysis of Soames' structured proposition account, providing five cases where Soames' account makes distinctions that I suggest are overly fine-grained. The five problem cases consist of statements which have distinct structured propositions but which I suggest are mere paraphrases of each other and so should be assigned the same semantic content. Each case is representative of a specific relationship between the statements that arguably are only ever ways of paraphrasing. Namely, the cases involve i) reordered conjunctions, ii) converse predicates, iii) obverse statements, iv) sentential vs predicative forms and v) definitions.[30]

Recall that as a part of our modeller's perspective this evaluation relies more on the use of language and associated linguistic judgements than it does on commitments to a priori properties of propositions. We focus on mathematical examples for these cases, since this is an area where it's clearer that the differences between the statements are not salient for how they are used.[31] However, these cases may be concerning across a much broader swath of language.[32] If you don't find these cases concerning, they nonetheless highlight features tracked by Soames' account and show how they are treated.

Case A: Re-ordered Conjunctions

Our first problematic example is that conjunction is not commutative on Soames' account, i.e., $\varphi \wedge \psi$ and $\psi \wedge \varphi$ have different semantic contents (for $\varphi \neq \psi$). Let's start with two simple statements and consider the two orders in which we can form their conjunction.

(15) Lee is tall and Remi is happy.

(16) Remi is happy and Lee is tall.

These statements have the following structured propositions (see also figure 3.6).

(15⋆) $\langle \mathsf{Conj}, \langle\langle \overline{\mathsf{tall}}, \langle \mathsf{Lee} \rangle\rangle, \langle \overline{\mathsf{happy}}, \langle \mathsf{Remi} \rangle\rangle\rangle\rangle$

(16⋆) $\langle \mathsf{Conj}, \langle\langle \overline{\mathsf{happy}}, \langle \mathsf{Remi} \rangle\rangle, \langle \overline{\mathsf{tall}}, \langle \mathsf{Lee} \rangle\rangle\rangle\rangle$

The inequality of these two structured propositions relies solely on the fact that these are ordered trees. For example, these structured propositions disagree at the bottom left node, since it is labelled Lee in (15⋆) and Remi in (16⋆) and Remi and Lee are distinct individuals

[30]The *obverse* of a statement is obtained by replacing a relation/predicate with the negation of its complement. This term is being used in analogy to the process of obversion from syllogistic logic. Obversion retains the subject (and the quantity), but replaces the predicate with its complement and flips the quality (i.e., whether it is a positive or negative statement). Thus the primary difference from our usage is that there may not be a quantifier determining the statement's quantity.

[31]Mathematical examples are also used due to the contingent fact that this work grew out of a project on mathematical meaning.

[32]In other words, I don't think that the problem with these cases is due to some feature unique to mathematics. In particular, the relationships identified are closer than just being mathematically equivalent or mutually derivable.

in our domain. In fact these structured propositions disagree at every node except the root. So our two conjunctions (15, 16) are not synonymous.[33]

(15⋆) (16⋆)

Figure 3.6: Structured propositions for (15⋆) and (16⋆)

This result is pretty counter-intuitive, it seems like $\varphi \wedge \psi$ means the same as $\psi \wedge \varphi$. The broader phenomenon of logically equivalent statements having distinct structured propositions (e.g. 3, 4) is a virtue of hyperintensional accounts. However, it is much less compelling that merely reordering conjuncts should change the semantic content.

The cases where 'and' behaves non-commutatively tend to involve it acting as other than a logical conjunction. Many such cases involve an implicit temporal ordering like 'I unlocked the door and opened it' versus the unnatural 'I opened the door and unlocked it.' In such cases 'and' should not be translated as $\wedge$, in our door example, it should be translated by something that preserves this temporal component. If we limit ourselves to mathematical settings, distinguishing between reordered conjunctions is even less appealing.

A different structured proposition account, given in (Moschovakis 2006), solves this problem by allowing certain kinds of permutations, i.e., re-orderings of the tree elements. However, to define the synonymy preserving permutations, Moschovakis uses a more sophisticated tree structure. These permutations also allow the account to more accurately capture symmetric relationships. For instance, the account judges 'Tegan is a sibling of Sara' and 'Sara is a sibling of Tegan' as synonymous.

[33]This result relies on the accuracy of definition 3.2.7 but since Soames doesn't suggest anything to the contrary, it seems reasonable to require the identity of trees since a standard isomorphism between trees would undermine the decision to use ordered trees.

Case B: Converse Relations

Our next case involves a similar issue from (Moschovakis 2017), namely, converse relations.

(17) Seven divides ninety-one.

(18) Ninety-one is divisible by seven.

These statements are formalised as follows.

($17^{\dagger}$) $divides(7, 91)$

($18^{\dagger}$) $divisible(91, 7)$

So we have the following structured propositions.

($17\star$) $\langle \overline{\mathsf{divides}}, \langle \mathsf{7}, \mathsf{91} \rangle \rangle$

($18\star$) $\langle \overline{\mathsf{divisible}}, \langle \mathsf{91}, \mathsf{7} \rangle \rangle$

Figure 3.7: Structured propositions for ($17\star$) and ($18\star$)

Here the issue isn't just that we have reordered the components, instead we have genuinely different components, since $\overline{\mathsf{divides}} \neq \overline{\mathsf{divisible}}$. Thus these are distinct structured propositions and so the statements are not synonymous on Soames' account. Yet, to merely note that these properties are not identical is to miss how closely related they are. They are converses, so one relation holds of an ordered pair $\langle a, b \rangle$ if and only if the other relation holds of $\langle b, a \rangle$.

Arguably there is really just one mathematical fact that is being expressed by both statements.[34] Namely, the fact that 91 divided by 7 is an integer.[35] Indeed, this might help explain why

[34]Referring to a "mathematical fact" likely provokes associations to grounding and truth-maker semantics. However, in line with our modeller's approach, we don't attribute strong metaphysical or universal claims to such facts which are common in those other literatures.

[35]Unfortunately, Soames' account doesn't deal with function symbols, so it's not clear what structured

only one of these relations (divides) is given a symbol (|), since it would be superfluous to have more notation for something we can already readily symbolise.[36] If both statements express the same mathematical fact then I suggest that this gives us reason to conclude that they are synonymous, against the judgement of Soames' account.

Converse relations also appear in active and passive constructions. For example, (Frege 1879) claims that 'At Plataea the Greeks defeated the Persians' and 'At Plataea the Persians were defeated by the Greeks' have the same conceptual content. Although, Frege replaces conceptual content with his later Sinn/Bedeutung distinction, it gives a clear precedent for wanting to treat such statements as equivalent under an appropriate synonymy relation.[37]

Case C: Sentential vs Predicative forms

Since our language has lambda abstraction, we can form more complex predicates, such as ones which contain logical connectives. Let's look at the impact of this ability by considering the following two statements.

(19) The reals are complete and the reals are infinite.

(20) The reals are complete and infinite.

These two statements can be given different formalisations, corresponding to whether the conjunction applies to a sentence or appears within a predicate respectively. The former is simply a conjunction of atomic formulas. For the latter, we use lambda abstraction to form a complex predicate and perform the grammatical construction of coordination.

($19^\dagger$) $complete(\mathbb{R}) \wedge infinite(\mathbb{R})$

($20^\dagger$) $\lambda x.(complete(x) \wedge infinite(x))(\mathbb{R})$

These give the following structured propositions.

proposition to assign to $integer(91/7)$. It seems plausible to treat term construction as analogous to formula construction. Especially, since in the case of division, we avoid issues of world relativity and hence the need to identify when an expression is being used as a world relative functional term or as a rigid designator for its actual extension.

[36] Of course, mathematics is typically plagued by variant notations, so the dominance of divides is actually rather impressive.

[37] The issue of converse relations is also taken up from a more metaphysical lens in both (Williamson 1985) and (Fine 2000).

$$(19\star) \quad \langle \mathsf{Conj}, \langle \langle \overline{\mathsf{complete}}, \langle \mathbb{R} \rangle \rangle, \langle \overline{\mathsf{infinite}}, \langle \mathbb{R} \rangle \rangle \rangle \rangle$$

$$(20\star) \quad \langle \mathfrak{a}^{\mathsf{g},\mathsf{x}}, \langle \mathbb{R} \rangle \rangle \text{ such that } \mathfrak{a}^{\mathsf{g},\mathsf{x}}(\mathsf{o}) = \langle \mathsf{Conj}, \langle \langle \overline{\mathsf{complete}}, \langle \mathsf{o} \rangle \rangle, \langle \overline{\mathsf{infinite}}, \langle \mathsf{o} \rangle \rangle \rangle \rangle$$

Obviously these are distinct structured propositions because they are different sizes. Even though we can see from the definition of $\mathfrak{a}^{\mathsf{g},\mathsf{x}}$ that these two structured propositions will have the same value at every possible world, since $\mathfrak{a}^{\mathsf{g},\mathsf{x}}$ maps $\mathbb{R}$ to $(19\star)$. In type theory, $(19\star)$ is usually called the β reduct of $(20\star)$, so this shows that β reduction does not preserve synonymy on Soames' account, since (20) and (19) are not synonymous.

Again, I suggest that only one "fact" is being expressed by these statements, namely that the real numbers have the property of being complete and the property of being infinite. Although there may be a metaphysical and/or cognitive difference between whether these (or any pair of) properties are two halves of a more complex property or two distinct properties, I reject that there is a semantic difference. In particular, these statements are used interchangeably and so should be treated synonymously.

Even worse, I suggest it might actually be misleading to take the purported difference between these statements seriously in a mathematical setting. For style and convenience, mathematicians are far more likely to say 'The reals are complete and infinite' which is the predicative form. This gives the sense that there is a single complex property being 'complete and infinite' which is satisfied by the real numbers. However, it's more reasonable to think that the information being conveyed is about two separate properties, 'complete' and 'infinite', each of which hold of the real numbers. Thus mathematicians use the statement which is more literally formalised as a predicative sentence but they intend something closer to the sentential use of conjunction to combine two separate conjuncts.

Case D: Obverse Statements

Earlier, we discussed converse relations and now we turn to a similar but less common connection. Consider (7) and its obverse, the statement obtained by replacing its relation (even) with the negation of its complement (not odd).[38] In attempting to do so, we encounter a situation similar to that in the preceding section as there are two possible positions for the negation. Hence we have three statements to consider.

[38]See footnote 30 for more about why we call these obverse statements.

(7) Six is even.

(21) It's not the case that six is odd.

(22) Six is not odd.

These statements are formalised as follows.

($7^\dagger$) $even(6)$

($21^\dagger$) $\neg odd(6)$

($22^\dagger$) $\lambda x.(\neg odd(x))(6)$

These formulas have the following structured propositions.

($7\star$) $\langle \overline{\mathsf{even}}, \langle 6 \rangle\rangle$

($21\star$) $\langle \mathsf{Neg}, \langle \overline{\mathsf{odd}}, \langle 6 \rangle\rangle\rangle$

($22\star$) $\langle \mathfrak{a}^{\mathsf{g,x}}_{\mathsf{notodd}}, \langle 6 \rangle\rangle$ such that $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{notodd}}(\mathsf{o}) = \langle \mathsf{Neg}, \langle \overline{\mathsf{odd}}, \langle \mathsf{o} \rangle\rangle\rangle$

No two of these statements are synonymous with any other, since each of these structured propositions are distinct. Obviously ($21\star$) is a different size from the others and then ($7\star$) differs from ($22\star$) since the former contains an intension and the latter a propositional function.[39] Although ($21\star$) is the β-reduct of ($22\star$), which we discussed in the previous section.

However, these statements seem to again be mere linguistic paraphrases of each other. Yet, the situation is more complicated than in previous cases to argue that a single mathematical fact is being expressed. Suppose we attempt to describe the fact, it would be that some property holds of the number 6. For mathematical use, it doesn't matter whether this core property is negated or not. But neither ($21\star$) nor ($22\star$) are strictly speaking (the semantic content of) a relation holding of 6. Obviously ($21\star$) is a negation and as noted ($22\star$) contains a propositional function.

[39]We could instead translate 'even' as an abstraction term $\lambda x.(even(x))$ and assign it the semantic content $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{even}}$. However, this would still not be synonymous with (22) since $\mathfrak{a}^{\mathsf{g,x}}_{\mathsf{even}}(\mathsf{o}) \neq \mathfrak{a}^{\mathsf{g,x}}_{\mathsf{notodd}}(\mathsf{o})$ since the former ($\langle \overline{\mathsf{even}}, \langle \mathsf{o} \rangle\rangle$) has two components and the latter ($\langle \mathsf{Neg}, \langle \overline{\mathsf{odd}}, \langle \mathsf{o} \rangle\rangle\rangle$)has three.

Part of the issue here is that although we put the negation in the predicative *position*, it is still the sentential or Boolean negation. So in both cases 'not odd' is treated as the denial of a property rather than as a property in itself. In order to treat 'not odd' directly as a property, we would need a predicate negation that acts more like a complement operation and takes relations to different relations, like adverbs do.[40]

This case relies more heavily on the mathematical setting than the others, because it's only when we restrict our domain to integers that 'even' and 'not odd' have the same extension. Although the need to restrict the domain in this way weakens our claim that they are in fact the same relation, the relation itself only applies to integers. Thus, we might hypothesise that this restriction is already part of the relation since applying it to non-integers results in a linguistically unacceptable statement.[41] Although our suggestion here involves "ignoring" some structure, we don't want to collapse down to intensions. For example, we maintain that 'perfect' is also a relation on integers which is clearly distinct from 'even', so we can maintain the lack of synonymy between (7) and (8).

Case E: Definitions

Another way that we might approach the mathematical fact behind (7) is to apply the definition of 'even.' So let's consider the following pair of statements.

(7) Six is even.

(23) Six is divisible by two.

These statements have the following structured propositions.

(7⋆) $\langle \overline{\mathsf{even}}, \langle 6 \rangle \rangle$

(23⋆) $\langle \overline{\mathsf{divisible}}, \langle 6, 2 \rangle \rangle$

Again these are distinct structured propositions because they have different sizes. Additionally (17⋆) contains a referential term, namely 'two,' not found in (7⋆). Thus these sentences are

[40] Although English lacks the vocabulary to distinguish between predicative and sentential negation (we can only do so indirectly with syntactic analysis as in the previous section), languages such as Russian do have words for both (Borschev et al. 2006).

[41] One initial way to capture this behaviour might be to allow partial relations. So 'even' is true of 6, false of 117 and undefined on Julius Caesar.

not synonymous. However, (23) simply replaces the relation in (7) by its definition so it's clear that there is some kind of meaning shared between the statements.

This canonical style of case using definitions is often discussed in terms of the 'Paradox of Analysis' (Robinson 1954). This paradox is more accurately a dilemma about the inability for analyses to be both correct and informative. In order for an analysis to be correct, it must establish an equivalence between the definiendum and the definiens. But if both terms are equivalent then it appears the analysis is merely tautological and has not achieved anything. On the other hand, we expect that some work has been done in the course of the analysis. But if the analysis is substantial then the definiendum and the definiens must differ in some way, undermining their purported equivalence.

In our example, we are closer to having a stipulative definition, where the term 'even' is introduced precisely to name the property given of being divisible by two. (Russell and Whitehead 1963, p. 11) describe such cases as 'a declaration that a certain newly introduced symbol or combination of symbols is to mean the same as a certain other combination of symbols of which the meaning is already known.'[42] Such definitions are often found in mathematics, where we reach the ideal of providing necessary and sufficient conditions, removing any leeway between a predicate and its definition on which to base this lack of synonymy.

Given our interest in mathematical meaning in particular, this may suggest that mathematical predicates always need to be replaced by their definition. Although this wouldn't resolve our issue from the previous section of equivalent but not definitional properties. There might also be cases where it is difficult to assess which among a collection of equivalent properties is the official definition. I suggest that the definition of transitive sets is such a case. The property that "every element of A is also a subset of A" and that "if $x \in A$ and $y \in x$ then $y \in A$" are both given as definitions and neither is particularly more intuitive than the other (the latter makes the connection to transitivity obvious but the former tends to be more useful in applications). These statements are not synonymous on Soames' account since the latter contains a conjunction and the former does not, and this difference remains even if one insists on replacing "subset" by its definition in terms of set membership. Thus replacing mathematical predicates with their definition may not be an immediate solution to this issue.

[42] Technically Russell and Whitehead claim that all (mathematical) definitions act in this way, I follow (Robinson 1954) in restricting this claim to certain types of definition.

3.4 Conclusion

Although Soamesian structured propositions may be an appropriate model of semantic contents for some purposes, it faces significant challenges. In particular, this chapter provides five cases where the model is too fine-grained especially when applied to mathematical statements. The distinctions tracked by structured propositions focus primarily on syntactic features and do not connect with mathematically salient reasons for distinguishing statements. Nonetheless, I think the structured proposition approach warrants further investigation as many of these distinctions can be motivated, even if they are not compelling for our particular use case. Thus, I hope the presentation given here which strips Soames' account down to its formal skeleton is beneficial for further work developing this account from a more sympathetic perspective.

Bibliography

Berto, Francesco and Mark Jago (2019). *Impossible Worlds*. Oxford: Oxford University Press. 1333 pp.

Bjerring, Jens Christian and Wolfgang Schwarz (2016). “Granularity Problems”. In: *The Philosophical Quarterly* 67.266, pp. 22–37.

Borschev, Vladimir et al. (2006). “Sentential and Constituent Negation in Russian BE-sentences Revisited”. In: *Formal Approaches to Slavic Linguistics: The Princeton Meeting 2005 (FASL 14)*.

Brandom, Robert B. (2009). *Articulating Reasons: An Introduction to Inferentialism*. Harvard University Press.

Carnap, Rudolf (1952). “Meaning postulates”. In: *Philosophical Studies* 3.5, pp. 65–73.

— (1956). *Meaning and Necessity: A Study in Semantics and Modal Logic*. Second expanded edition. Chicago: University of Chicago Press.

Chomsky, Noam (1980). *Rules and representations*. Woodbridge lectures 11. New York: Columbia University Press. 299 pp.

Church, Alonzo (1951). “A Formulation of the Logic of Sense and Denotation”. In: *Structure Method and Meaning: Essays in honor of Henry M. Sheffer*. Ed. by S. K. Langer Paul Henley H. M. Kallen. New York: The Liberal Arts Press, pp. 3–24.

— (1973). “Outline of a Revised Formulation of the Logic of Sense and Denotation (Part I)”. In: *Noûs* 7.1, pp. 24–33.

— (1974). “Outline of a Revised Formulation of the Logic of Sense and Denotation (Part II)”. In: *Noûs* 8.2, pp. 135–156.

— (1993). “A Revised Formulation of the Logic of Sense and Denotation. Alternative (1)”. In: *Noûs* 27.2, pp. 141–157.

Cresswell, Maxwell John (1985). *Structured meanings: The semantics of propositional attitudes*. MIT press.

Davidson, Donald (1967). “Truth and meaning”. In: *Synthese* 17.1, pp. 304–323.

Fine, Kit (2000). “Neutral Relations”. In: *Philosophical Review* 109.1, pp. 1–33.

Fodor, Jerry A. (1980). *The language of thought*. The language and thought series. Cambridge, Mass.: Harvard University Press. 214 pp.

Frege, Gottlob (1879). *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Halle an der Saale: Lubrecht & Cramer.

Frigg, Roman and James Nguyen (2020). *Modelling Nature: An Opinionated Introduction to Scientific Representation*. Synthese Library: Studies in Epistemology, Logic, Methodology, and Philosophy of Science 427. Springer International Publishing.

Heim, Irene and Angelika Kratzer (1998). *Semantics in generative grammar*. Blackwell.

Jackendoff, Ray (1990). *Semantic structures*. Current studies in linguistics series 18. Cambridge, Mass.: MIT Press. 322 pp.

Jackson, Brendan Balcerak (2020). "Model-theoretic semantics as model-based science". In: *Synthese* 199.1–2, pp. 3061–3081.

Jech, Thomas J. (2006). *Set theory*. 3rd ed. Springer Monographs in Mathematics. Springer Berlin / Heidelberg. 1754 pp.

King, Jeffrey C. (1995). "Structured Propositions and Complex Predicates". In: *Noûs* 29.4, pp. 516–535.

— (2019). "Structured Propositions". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2019. Metaphysics Research Lab, Stanford University.

Kripke, Saul A. (1980). *Naming and necessity*. Harvard University Press, p. 172.

Lewis, David (1970). "General semantics". In: *Synthese* 22.1-2, pp. 18–67.

— (1973). *Counterfactuals*. Cambridge: Harvard University Press.

Marker, David (2002). *Model Theory. An Introduction*. Springer, p. 360.

McGrath, Matthew and Devin Frank (2024). "Propositions". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Fall 2024. Metaphysics Research Lab, Stanford University.

Montague, Richard (1973). "The Proper Treatment of Quantification in Ordinary English". In: *Approaches to Natural Language*. Springer Netherlands, pp. 221–242.

Morgan, Mary S. (1999). "Learning from models". In: *Models as Mediators*. Cambridge University Press, pp. 347–388.

Morrison, Margaret and Mary S. Morgan (1999). "Models as mediating instruments". In: *Models as Mediators*. Cambridge University Press, pp. 10–37.

Moschovakis, Yiannis N. (2006). "A Logical Calculus of Meaning and Synonymy". In: *Linguistics and Philosophy* 29.1, pp. 27–89.

— (2017). "Sense and denotation as algorithm and value". In: *Logic Colloquium '90: ASL Summer Meeting in Helsinki*. Ed. by Juha Oikkonen and Jouko Väänänen. Lecture Notes in Logic. Cambridge University Press, pp. 210–249.

Nelson, Michael (2024). “Propositional Attitude Reports”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Fall 2024. Metaphysics Research Lab, Stanford University.

Nolan, Daniel (1997). “Impossible Worlds: A Modest Approach”. In: *Notre Dame Journal of Formal Logic* 38.4.

Propp, James (2013). “Real Analysis in Reverse”. In: *The American Mathematical Monthly* 120.5, p. 392.

Ripley, David (2012). “Structures and circumstances: two ways to fine-grain propositions”. In: *Synthese* 189.1, pp. 97–118.

Robinson, Richard (1954). *Definition*. Oxford University Press.

Russell, Bertrand (1903). *Principles of Mathematics*. New York: Norton.

Russell, Bertrand and Alfred North Whitehead (1963). *Principia Mathematica volumes 1, 2, and 3*. Cambridge University.

Salmón, Nathan (1986a). *Frege's Puzzle*. Atascadero, CA: Ridgeview Publishing Company.

— (1986b). “Reflexivity”. In: *Notre Dame Journal of Formal Logic* 27.3, pp. 401–429.

— (1992). “On Content”. In: *Mind* 101.404, pp. 733–751.

Schechter, Eric (1997). “Constructivism and Choice”. In: *Handbook of Analysis and Its Foundations*. Elsevier, pp. 131–154.

Schroeder-Heister, Peter (2018). “Proof-Theoretic Semantics”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2018. Metaphysics Research Lab, Stanford University.

Simpson, Stephen G (2009). *Subsystems of second order arithmetic*. Vol. 1. Cambridge University Press.

Soames, Scott (1987). “Direct Reference, Propositional Attitudes, and Semantic Content”. In: *Philosophical Topics* 15.1, pp. 47–87.

— (2005). “Naming and Asserting”. In: *Semantics versus Pragmatics*. Oxford University Press, pp. 356–382.

— (2008). “Introduction”. In: *Philosophical Essays, Natural Language: What It Means and How We Use It*. Vol. 1. Princeton: Princeton University Press, pp. 1–20.

Weisberg, Michael (2015). *Simulation and similarity. Using models to understand the world.* Oxford studies in philosophy of science. Oxford: Oxford University Press. 190 pp.

Williamson, Timothy (1985). “Converse Relations”. In: *Philosophical Review* 94.2, pp. 249–262.

Yalcin, Seth (2018). “Semantics as Model-Based Science”. In: *The Science of Meaning*. Oxford: Oxford University Press, pp. 334–360.