The Dissertation Committee for Steven Bischoff certifies that this is the approved version of the following Dissertation:

Scientific Pluralism and Semantics

Committee:

Joshua Dever, Supervisor

Lawrence R Buchanan

David Sosa

Jeff Speaks

Scientific Pluralism and Semantics

by

Steven Bischoff

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

December 2023

Abstract

Scientific Pluralism and Semantics

Steven Bischoff, PhD

The University of Texas at Austin, 2023

Supervisor: Joshua Dever

This project is a case study in Scientific Pluralism focused on formal natural language semantics. The case study presents a novel characterization of Scientific Pluralism while arguing for a Pluralist approach to semantics.

I begin by describing the plural state of the field. I show the lack of consensus among five different types of *background constraints* that govern semantic theorizing. Semantics consists of a set of more-or-less distinct *projects* partially defined by the background constraints practitioners adhere to.

With this description in hand, I present a Pluralist view with two main parts. First, I briefly argue that semantics should maintain multiple projects and discuss consequences of doing so. Since this ground has been covered extensively in the Scientific Pluralism literature, the bulk of the discussion centers around the second part of the view. I extend typical Pluralisms by arguing that semantic background constraints should be viewed as *soft*. A soft background constraint may remain vague, and it may be overridden without being rejected when it conflicts with other background constraints.

I apply my Pluralist view to *Taxonomic Monism*, a background constraint commonly expressed in some form throughout the literature. Taxonomic Monism provides a goal for semantic theorizing, saying that semantic theories should aim to state the meaning of every expression in the language. I argue that we should view Taxonomic Monism as a soft constraint in several ways, and I discuss the consequences of doing so.

Table of Contents

Introduction

Semantics is the scientific study of linguistic meaning. Semantic theorizing generally begins modestly, focusing on small fragments of the language before gradually expanding its scope. Suppose, then, that we're studying the meaning of the English noun 'cow'. One way to proceed with this investigation is to study the meaning of sentences containing 'cow'. Perhaps we're motivated by compositionality considerations, believing that the meaning of 'cow' (henceforth ⟦cow⟧) somehow contributes to the meaning of 'cow'-sentences such as (1):

(1) Betsy is a cow.

Following Frege (1892),[1] we characterize the meaning of (1) as a truth value. Even if truth values are appropriate meanings of *sentences*, though, they certainly aren't appropriate as ⟦cow⟧ or the meaning of any other constituent of (1). While ⟦cow⟧ contributes to the truth value of (1), it isn't a truth value itself.

Nonetheless, progress has been made: ⟦cow⟧ must interact with other constituents of (1) to produce a truth value.[2] For names like 'Betsy', a referential theory seems initially plausible. According to this theory, the meaning of a name is whatever it refers to.

(a) ⟦Betsy⟧ = Betsy.

Setting aside any semantic contribution made by 'is' and 'a', ⟦cow⟧ must interact with ⟦Betsy⟧ in a way that results in a truth value.

At this point, the theory has posited two domains: the domain of entities $\mathscr{D}$ and the domain of truth values. For the sake of simplicity, suppose the domain of truth values is {TRUE, FALSE}. We can get truth values from entities if we posit the existence of functions from entities

[1] Frege, "On Sense and Reference."
[2] Ignoring issues of context for now.

to truth values. To get a truth value as the meaning of (1), we claim that ⟦cow⟧ is one such function that takes all and only cows to TRUE:

(b) ⟦cow⟧ = the function f: $\mathscr{D} \rightarrow$ {TRUE, FALSE} such that, for all $x \in \mathscr{D}$, f(x) = TRUE iff x is a cow.

Applying ⟦cow⟧ to ⟦Betsy⟧ gives us a meaning for (1):

(c) ⟦Betsy is a cow⟧ = ⟦cow⟧(⟦Betsy⟧) = TRUE iff Betsy is a cow.

Equalities like (a) and (b) are often called *lexical entries*. The idea is that speakers of a language have access to a *lexicon*, *i.e.* a set of lexical entries for basic expressions like 'Betsy' and 'cow'. The lexicon gives speakers the raw materials needed to build complex expressions like sentences. A frequent complaint about lexical entries like (b) is that they can seem extremely uninformative. They rely on disquotation; in giving the meaning of 'cow', the lexical entry (seemingly) re-uses the word 'cow'.[3]

The concern is that the heavy reliance on disquotation limits the explanatory power of the semantic theory to the theory's detriment. By itself, (b) tells us that the meaning of 'cow' is a function that takes cows to TRUE and all other entities in $\mathscr{D}$ to FALSE. This entry gives fluent *English* speakers a solid grasp on which entities 'cow' is true of, but this is only because we already know the meaning of 'cow'. If someone who doesn't know English asks what the meaning of 'cow' is, they will likely be very dissatisfied with an answer along the lines of (b).

Fortunately, there is much more that a semantic theory can say about the meaning of 'cow'. For instance, linguists and philosophers[4] have long been aware of an apparent difference in meaning between nouns like those in Group A and nouns like those in Group B:

[3] See Yalcin, "Semantics as Model-Based Science." for discussion and the view that such lexical entries don't actually reuse the object language. One can still complain about such entries even if Yalcin's view is correct.
[4] *e.g.* Quine, *Word and Object*.

Group A: ‘cow’, ‘vehicle’, ‘idea’, ‘hole’

Group B: ‘wine’, ‘mud’, ‘enthusiasm’, ‘justice’

Group A and Group B nouns seem to have different types of objects in their extensions. Group A nouns are true of discrete, countable entities such as cows and vehicles, so they are typically called count nouns. Group B nouns, on the other hand, are true of non-discrete, non-countable “masses” or “substances”[5]: quantities of wine, enthusiasm, etc. Consequently, they are typically called mass nouns.

This apparent difference in meaning can be captured by theories of the meanings of common nouns. One option posits two distinct subdomains of the domain of entities $\mathscr{D}$:[6]

$\mathscr{D}_c$ - the domain of discrete, countable entities like cows and ideas

$\mathscr{D}_m$ - the domain of masses like quantities of wine

We can use this division of the domain of entities to distinguish between count and mass noun meanings. Our previous theory claimed that ⟦cow⟧ and ⟦wine⟧ are both functions from entities to truth values, but we can now distinguish between them by positing that the functions have different domains. ⟦cow⟧ is a function from *discrete* entities to truth values and ⟦wine⟧ is a function from *masses* to truth values

(b.1) ⟦cow⟧ = the function f: $\mathscr{D}_c \rightarrow$ {TRUE, FALSE} such that, for all $x \in \mathscr{D}_c$,
f(x) = TRUE iff x is a cow.

(d) ⟦wine⟧ = the function f: $\mathscr{D}_m \rightarrow$ {TRUE, FALSE} such that, for all $x \in \mathscr{D}_m$,
f(x) = TRUE iff x is wine.

As a consequence of tying the expressions to distinct subdomains, the theory now claims that ‘cow’ is only true of discrete, countable entities and ‘wine’ is only true of masses.

[5] As shown by ‘enthusiasm’ and ‘justice’, the substance doesn’t necessarily need to be physical.
[6] Link, “The Logical Analysis of Plurals and Mass Terms.”

Unfortunately, this theory encounters some difficulties. Consider a gruesome context in which an 18-wheeler has hit a cow head-on at great speed. Someone then utters (2), referring to the aftermath of the collision:

(2) There's cow on the side of the road.

In this context, 'cow' is referring to the "cow stuff" that's the result of the collision. In other words, 'cow' is true of a mass entity, which contradicts (b.1).

The literature on count and mass nouns contains many options for dealing with such examples,[7] but for now I will only discuss two of them. First, one might think (2) is evidence that 'cow' is *ambiguous* with more than one lexical meaning. Some expressions are clearly ambiguous. 'bank' is semantically ambiguous with distinct "financial institution" and "river bank" meanings (Further, it's syntactically ambiguous between the aforementioned nouns and a transitive verb ('bank the plane')). The first option for dealing with (2), which I will call Theory 1, says that 'cow' is also non-obviously ambiguous. The lexicon actually contains *two* 'cow' entries, one with the dominant count meaning in (b.1) and another with the more rare mass meaning in (b.2):

(b.2) $\|\text{cow}_1\|$ = the function f: $\mathscr{D}_m \rightarrow \{\text{TRUE, FALSE}\}$ such that, for all $x \in \mathscr{D}_m$, f(x) = TRUE iff x is cow.

Typical uses of 'cow' involve (b.1) and are true of count objects, but the addition of (b.2) to the lexicon allows 'cow' to be true of mass objects in contexts like that of (2).

Second, one might not be overly impressed by the data provided by (2). Even with the contextual background the utterance seems a bit strange, so perhaps it can be dismissed as ungrammatical or infelicitous. And even if the utterance *is* grammatical and felicitous, the

[7] See Pelletier, "Lexical Nouns Are Both +mass and +count, but They Are Neither +mass nor +count." for a good overview.

contexts in which it's acceptable might be considered outliers, putting little pressure on the theory to explain them. Taking this option, which I will call Theory 2, keeps (b.1) in the lexicon and doesn't add any new components like entry (b.2).

Theory 1 and Theory 2 are ostensibly competing theories of the meaning of 'cow'. Theory 2 claims that the lexicon contains only one 'cow' entry, whereas Theory 1 claims that the lexicon contains two 'cow' entries with distinct meanings. Let's assume that both theories are live options, neither having been decisively refuted. This scenario is an instance of *scientific plurality*, which occurs when multiple models, theories, methods, and / or projects are maintained simultaneously within a scientific field.

It seems intuitively plausible that the existence of multiple live options for the meaning of 'cow' is in some sense a defective situation to be remedied; at least one of the theories should be decisively refuted during the search for the best theory. The general intuition, that the existence of multiple live options constitutes a defect, is not unique to semantics. Much scientific practice seems to be implicitly motivated by an impulse to establish dominant paradigms, refute competing theories, and build consensus. Reframed as a normative thesis about the goals of the sciences, this impulse is captured by Kellert, Longino, and Waters in their characterization of *Scientific Monism*:[89]

> **Scientific Monism:**
>
> **SM.1:** The ultimate aim of a science is to establish a single, complete, and comprehensive account of the natural world (or the part of the world investigated by the science) based on a single set of fundamental principles.

[8] Kellert, Longino, and Waters, "Introduction: The Pluralist Stance.", *x*.
[9] It's unclear whether Kellert et. al intend "complete" and "comprehensive" to refer to different concepts; they sometimes just talk of "single, comprehensive" theories and methods, and I will do the same.

SM.2: The nature of the world is such that it can, at least in principle, be completely described or explained by such an account.

SM.3: There exist, at least in principle, methods of inquiry that if correctly pursued will yield such an account.

SM.4: Methods of inquiry are to be accepted on the basis of whether they can yield such an account.

SM.5: Individual theories and models in science are to be evaluated in large part on the basis of whether they provide (or come close to providing) a comprehensive and complete account based on fundamental principles.

Note that Scientific Monism can be interpreted broadly, as a claim about the sciences in general, or narrowly, as a claim about specific sciences.

Suppose that the semanticists working on theories of the meanings of 'cow' and similar nouns are Scientific Monists with respect to semantics. Because of this underlying Monism, supporters of Theory 1 *disagree* with supporters of Theory 2. They are searching for a single best theory of meaning, which can contain at most one of Theories 1 and 2. As a result, the two theories are in direct conflict, competing to be part of the ultimate best theory.

In their compendium of work on the subject, Kellert *et. al* present Scientific Pluralism as an alternative to Scientific Monism. Scientific Pluralism is more a family of views than a single, unified theory. These views are united by their denial of at least one (and usually more) of (SM.1) - (SM.5). They share the attitude that plurality is at least sometimes desirable, but this leaves significant room for difference.

Pluralist views can be distinguished based on which component(s) of Scientific Monism they deny. A Pluralist might make a metaphysical argument denying (SM.2), arguing that the world (or more narrowly, the part of the world studied by a particular science) is such that it

cannot be described or explained by a single, complete, and comprehensive account.[10] A Pluralist might also make an epistemological argument denying (SM.3). They could argue that, even if (SM.2) is true broadly or narrowly, methods of inquiry that could yield such an account do not exist (or at least are beyond our grasp).[11] Other Pluralists might make a pragmatic argument denying (SM.1), (SM.4), or (SM.5). They could argue against (SM.1) by claiming that there is no ultimate aim of science (or a specific science) or that the ultimate aim is *not* to provide a single, comprehensive account. They could also argue against (SM.4) or (SM.5) by investigating the conditions under which scientific theories and methods of inquiry are to be evaluated and accepted.[12]

Pluralist views can also differ based on their breadth. Scientific Monists can interpret their thesis broadly, accepting it for science in general, or they can interpret it narrowly, accepting it on a case-by-case basis for individual sciences. Likewise, Scientific Pluralists can broadly deny Monism, claiming that it is false of science in general, or they can narrowly deny Monism on a case-by-case basis for individual sciences. For instance, one might be a Pluralist about various behavioral sciences but a Monist about physics.

Finally, Pluralist views can differ based on whether their arguments are "top-down" or "bottom-up". Top-down arguments for Scientific Pluralism depend on general claims about the nature of the world or human perception and reasoning. For instance, Giere argues that we can only observe phenomena from limited perspectives that provide a partial view, and that adopting multiple perspectives can provide a more complete view.[13] Bottom-up arguments for Pluralism

[10] Cartwright, *The Dappled World: A Study of the Boundaries of Science*.
[11] Giere, "Perspectival Pluralism."
[12] Waters, "A Pluralist Interpretation of Gene-Centered Biology."
[13] Giere, "Perspectival Pluralism."

take the form of case studies within specific sciences. Arguments have been put forth for Pluralist views in economics,[14] physics,[15] and psychology,[16] among many others.

Case studies in Scientific Pluralism have a dual purpose. First, they promote the adoption of a Pluralist perspective on the specific science being studied. Case studies examine the metaphysical underpinnings of the science's subject matter, the epistemology of what we can know about that subject matter, and/or the pragmatics of accepting the field's theories and methods, arguing that rejection of at least one of (SM.1) - (SM.5) is warranted. Second, case studies investigate the consequences that adopting a Pluralist perspective has for practice in the field. Of course, these two goals are intertwined. The consequences of adopting a Pluralist perspective can be appealing, and one strong argument for adopting a Pluralist perspective is showing the benefits of doing so.

This paper is one such case study. I examine the field of *formal natural language semantics*. I've already provided an example of what the field studies and how it does so, but I still need to do more to distinguish my target from closely related fields. First, the science I'm examining studies meaning in *natural languages*, the evolved languages that humans use every day to communicate through sounds and symbols. This contrasts with the study of meaning in the artificially-constructed languages of logic, computer science, and mathematics (among others).

Within the study of natural language we can make a further distinction between *formal* and *distributional* semantics. There are many potential ways of characterizing this distinction,

[14] See Grabner and Strunk, "Pluralism in Economics: Its Critiques and Their Lessons." for a thorough review of the active debate in economics.
[15] Dickson, "Plurality and Complementarity in Quantum Dynamics."
[16] Longino, "What Do We Measure When We Measure Aggression?"

but I will do so by contrasting the preferred mathematical tools of the two programs. The formal semantic tradition generally uses logic and set theory as the mathematical tools of choice, whereas the distributional semantic tradition generally uses linear algebra and statistics.[17] Since my focus will be on formal natural language semantics, this is henceforth what I will mean when I use 'semantics' without qualification.

In Chapter 1, I begin by discussing the current plural state of semantics. I doubt it will be surprising that semantics currently exhibits significant plurality, but I provide a novel characterization of this plurality that aims to capture more of the variation in the field. With this characterization of semantic plurality in hand, I present my version of Pluralism in Chapters 2 and 3 and give some general, top-down arguments for the view. My Pluralist view has two main components. First, I argue that semantics is better off maintaining multiple projects or schools. Second, I argue that the constraints on theorizing within a particular project should be soft, *i.e.* vague and cancelable. These arguments are mostly pragmatic, denying (SM.4) or (SM.5).

In Chapters 4 and 5, I discuss the consequences of my Pluralist view for semantics: what does a Pluralist semantics look like, and how does it differ from current practice? The consequences of maintaining multiple projects within a field have been covered extensively within the Scientific Pluralism literature, so these chapters focus on the consequences of semantic theorizing under soft constraints. My discussion is centered around a prominent, if often implicit, theoretical goal. The following passages from various authors provide glimpses of a (perhaps *the*) principal goal of semantics:

> The aim of a semantic theory is to state the meanings of expressions of the target language. So the output of [a semantic] theory should include, for every

[17] Clark, "Vector Space Models of Lexical Meaning."

> well-formed expression of the target language, a theorem which — in some sense or other — states that expression's meaning.[18]
>
> … 'a theory of meaning' for any one entire language … [is] a detailed specification of the meanings of all the words and sentence-forming operations of the language, yielding a specification of the meaning of every expression and sentence of the language.[19]
>
> My target therefore includes truth-conditions for each of infinitely many well-formed sentences of English, including the many that are never used.[20]

I will call this common goal for semantic theorizing *Taxonomic Monism*:

> **Taxonomic Monism:** A semantic theory should state the meaning of every expression in the language.

I won't deny Taxonomic Monism and argue that it sets an undesirable goal. Instead, my plan is to argue for a Pluralist interpretation of the goal and discuss the consequences of adopting this interpretation.

[18] Speaks, "What's Wrong with Semantic Theories Which Make No Use of Propositions?" 19.
[19] Dummett, "What Is a Theory of Meaning?" 1.
[20] Williams, *The Metaphysics of Representation*. 125.

Chapter 1: Semantic Plurality

The simultaneous existence of multiple models, theories, or approaches within a scientific field has been a recent focus for philosophers of science, even beyond the core Pluralism literature. Much of the work consists of attempts to justify this plurality: under what conditions is it rational to maintain multiple models, theories, or approaches, and where do these conditions hold?

One common claim is that a field can be in a plural state because of tradeoffs among *epistemic goods*. McMullin lists predictive accuracy, internal coherence, external consistency, unifying power, fertility, and simplicity as examples of epistemic goods that can be attained by a theory.[21] If tradeoffs among these goods exist, then attaining one of them comes at the expense of attaining others: for instance, modifying a theory to increase its predictive accuracy might only be possible by decreasing its simplicity.

Many philosophers have claimed that plurality can be a rational response to difficulties with (or the impossibility of) maximizing the attainment of many epistemic goods using only one theory.[22] For instance, Levins claims (specifically for population biology) that the complexity of the world and our own cognitive limitations conspire to impose plurality on us:[23]

> The multiplicity of models is imposed by the contradictory demands of a complex, heterogeneous nature and a mind that can only cope with few variables at a time; by the contradictory desiderata of generality, realism, and precision; by the need to understand and also to control; even by the opposing esthetic standards which emphasize the stark simplicity and power of a general theorem as against the richness and the diversity of living nature. These conflicts are irreconcilable.

[21] McMullin, "Values in Science."
[22] Weisberg, "Three Kinds of Idealization."
[23] Levins, "The Strategy of Model Building in Population Biology."

According to this argument, the world (or certain parts of it) is complex enough that individual simple models are insufficient: they inevitably sacrifice some epistemic goods to gain others. As model complexity increases to approach the complexity of the world, the models become less useful to us due to our cognitive limitations. As a consequence, plurality is often the best way to attain many epistemic goods using comprehensible models and theories.

Regardless of whether plurality is a rational response to tradeoffs among epistemic goods, the existence of tradeoffs can provide an explanation of the existence of plurality. A theorist's choice of which theory to adopt can be value-laden, depending on the relative importance the theorist places on individual epistemic goods. Consider two theories, Theory A and Theory B, and two epistemic goods, Good X and Good Y. Theory A does well attaining Good X and poorly attaining Good Y, while Theory B does well attaining Good Y and poorly attaining Good X. A theorist who highly values Good X will (all else being equal) prefer Theory A whereas a theorist who highly values Good Y will prefer Theory B. The apparent disagreement between these theorists is due to differences between their preferences for epistemic goods.

If we apply this idea to semantics, the claim is that plurality persists within semantics because theorists value different epistemic goods. Indeed, this can explain the apparent disagreement in my initial example. A theorist who highly values simplicity might prefer Theory 1, since it only requires one lexical entry for 'cow' and is therefore simpler than Theory 2. A theorist who highly values generality might prefer Theory 2: unlike Theory 1, it captures both the normal and abnormal uses of 'cow'.

1. Sources of Plurality

The idea of tradeoffs among epistemic goods is important and certainly explains much plurality. However, it also leaves much unexplained. In this section, I discuss five additional sources of semantic plurality. I will collectively refer to these five sources as *background constraints*.

Background Constraints:

a. Domain constraints
b. Characterization of the data
c. Second-order constraints
d. Nature of the target system
e. Relation between the model and the target

As many of the following examples will show, these background constraints are often the source of explicit disagreement and not just mere plurality. An underlying Monism often leads to conflict among theorists who adhere to different views.

To contrast them with actual *theories*, I think of background constraints as norms governing the process of theorizing. They tell semanticists which system(s) to investigate, which phenomena emerging from those systems they should be interested in, which data to gather and how to gather it, which inferences to make, and how their theories should connect to the world. As will be seen throughout the following examples, these norms aren't always made explicit; in such cases, the norms governing the theoretical process need to be inferred by observing that process.

1.a. Domain Constraints

There is widespread disagreement over what phenomena semantic theorizing should explain. The explanatory domain of semantics is obviously constrained by common sense. As Yalcin puts it, "[semantics] is aimed at modeling, in some predictive and explanatory way, a certain category of properties, namely the semantic properties of expressions".[24] No one expects semantics to explain earthquakes or planetary motion; instead, there is a vague group of phenomena such as reference, inference, anomaly, distribution, communication, etc. that we might pre-theoretically call "meaning-relevant". Semantics is expected to explain meaning-relevant phenomena by modeling the semantic properties of expressions.

One major complication arises from the fact that semantics does not have a monopoly on explanations in the domain of meaning-relevant phenomena. Semantics competes with pragmatics to provide explanations of every type of phenomenon listed above.[25] I say "competes" because there is no widely agreed-upon line dividing the explanatory domains of semantics and pragmatics. Different theorists operate (implicitly or explicitly) with different *domain constraints* on semantics. Since there is widespread disagreement on the phenomena that semantics explains, there is also widespread disagreement on what "the semantic properties" are.

1.a.i. Clear Disagreement

Sometimes this disagreement is quite explicit. As a first example, consider the domain constraints provided by Blutner.[26] He takes "compositionality, monotonicity, and the persistence of anomaly as bounding the domain of semantics proper". My current discussion will focus on

[24] Yalcin, "Semantics and Metasemantics in the Context of Generative Grammar." 45.
[25] Semantics also competes with syntax "on its other end", but in this section I will be focusing on pragmatics.
[26] Blutner, "Pragmatics and the Lexicon.", 20-21. At times, Blutner has endorsed conflicting views (*e.g.* Blutner, Hendriks, and Hoop, "A New Hypothesis on Compositionality.").

monotonicity and the persistence of anomaly, but since compositionality is such a core concept of semantics, I will introduce all three.

There is a consensus that natural language meaning is compositional in some sense, though there is significant controversy over the specifics.[27] A first broad attempt at a principle of compositionality says that the meanings of the atomic constituents of a complex expression, in combination with syntax and composition rules, are sufficient to determine the meaning of the complex expression itself (with the normal caveats given to accommodate, *e.g.*, indexicality).

Monotonicity refers to a certain type of inference. A monotonic inference is one that remains valid whenever new premises are added. An easy example of a monotonic inference is conjunction elimination; from (P & Q) we can infer P, and this inference remains valid with the addition of any premise R. Applied specifically to semantics, the monotonicity property ensures that no material provided by semantics is contextually cancelable.

The idea of the persistence of anomaly draws a distinction between pragmatic and semantic anomaly. Pragmatic anomalies are cancelable by context. Blutner uses the following pair to illustrate the idea:

(3a) #Some people are forbidden to eat cow.

(3b) Hindus are forbidden to eat cow.

In linguistic contexts such as (3a), the phrase ‘eat cow’ leads to anomaly; however, when placed in the context of Hindu religion in (3b), the phrase no longer leads to anomaly. According to Blutner, this non-persistence is evidence that the anomaly in (3a) is pragmatic rather than semantic; thus, semantics has no responsibility to account for it. Contrast this with ‘#married

[27] Szabó, “Compositionality.”

bachelor', which seems like it could never be acceptable (as long as both words are being used literally); if this anomaly is not cancelable, then semantics must explain it.

According to some theorists Blutner's domain constraints are too restrictive, while according to others they are not restrictive enough. Asher and Pustejovsky (henceforth A&P) are among the former group who assume a larger explanatory domain for semantics.[28] In particular, explaining (some) defeasible interpretations and non-persistent anomalies is a major motivation behind their theories. Consider the following:

(4a) Jane enjoyed the book.

The preferred interpretation of (4a) is that Jane enjoyed *reading* the book. According to A&P, this is because the meaning of 'book' has a *qualia structure* that includes a *telic* role of books; ‖book‖ includes the information that the purpose of books is to be read. Roughly, when 'enjoyed' composes with 'the book', 'enjoyed' accesses the telic role of 'book' to give 'enjoyed the book' the READING interpretation. This interpretation, however, is cancelable:

(4b) The goat enjoyed the book.

This sentence is much more likely to be given an EATING interpretation than a READING one. Thus, even though 'enjoyed the book' gets a READING interpretation, this interpretation is only a default. It can be canceled by 'the goat' or other elements of the (linguistic or non-linguistic) context, and thus fails to obey the monotonicity constraint.

A&P also allow the derivation of non-persistent anomaly within their semantics. Consider the following:

(5) #Bob's idea weighs five pounds.

[28] Asher and Pustejovsky, "Word Meaning and Commonsense Metaphysics."

(6) #The mountain documents the news of the day.

In normal contexts this seems anomalous. Using the subdomain strategy that was applied to count and mass anomalies, A&P hypothesize that common nouns and verbs have narrower domains than initially thought. Instead of a function from entities to truth values, ⟦weighs five pounds⟧ is a function from the set of *physical* entities to truth values. Bob's idea is an *informational* entity, which entails that it is not in the set of physical entities. Therefore, ⟦weighs five pounds⟧ fails to return a truth value when given ⟦Bob's idea⟧ as input. Similarly for (6), where it's hypothesized that the domain of ⟦documents the news of the day⟧ is the set of informational entities and ⟦the mountain⟧ is physical. These composition errors within the theory are intended to explain the intuitive anomaly of the sentences.

However, it isn't too difficult to think of contexts in which these sentences make sense. For instance, if the mountain is in the middle of a war zone and has been repeatedly hit by explosives, one could sensibly use (6) to describe how the damage to the mountain documents the progress of the fighting. The original anomaly was explained by a typing conflict between 'mountain' and 'documents'. The fact that the anomaly is cancelable (*inter alia*) leads A&P to build *type coercion* mechanisms into their account that allow (6) to (sometimes) successfully compose. Absent some sort of contextual force, though, A&P's semantics will still diagnose (6) as anomalous.

A&P don't make their domain constraints as explicit as those stated by Blutner. However, in other work Asher does provide some justification for the inclusion of non-monotonic interpretation and non-persistent anomaly as explananda of semantics.[29] He says the following:

[29] Asher, *Lexical Meaning in Context: A Web of Words*. 15.

> Is it a defeasible but *a priori* inference that … if Julie enjoyed the book, then she enjoyed reading it? Do such inferences follow solely from one's linguistic mastery of the language? … most people can distinguish between the largely automatic interpretations that these predications seem to entail and those that require more conscious effort. One might take that to be a mark of the information as being present even during predication rather than inferred afterwards using background, nonlinguistic beliefs.

Asher seems to be operating with "automaticity" as a (perhaps soft) domain constraint; if a meaning-relevant phenomenon arises automatically (*i.e.* quickly, unconsciously) then it is at least likely that it should be explained by semantics. The apparent fact that phenomena can be automatic while being non-monotonic or non-persistent leads A&P to take a broader view of the domain of semantics than Blutner.

Unlike A&P, Schiffer takes a narrower view of semantics.[30] Like other authors I've mentioned, Schiffer is interested in analyzing an innate Chomskian language faculty, which purportedly gives us the ability to understand sentences (or at least plays a crucial role in this ability). Semantic theories are therefore theories of a portion of our ability to understand sentences. Schiffer claims that "whatever constitutes the ability of a native speaker of a language to understand sentences of her language must be something shared by all native speakers of the language".

This claim about the ability provided by the language faculty leads to Schiffer's highly restrictive domain constraints. Schiffer uses the claim to exclude both entailment relations and anomaly from the explanatory domain of semantics. Presumably, the ability to recognize a specific entailment or anomaly is *not* shared by every competent user of the language. Though Schiffer doesn't provide much argumentation, the fact that these abilities are not shared to the

[30] Schiffer, "Meaning and Formal Semantics in Generative Grammar." 77.

same extent by every competent language user is attested to by, for instance, the well-known existence of disagreement over entailment and anomaly (among both theorists and experiment subjects) or various phenomena of "semantic illusion" (where well-formed sentences are judged ill-formed or vice versa). Thus, according to Schiffer, theorists such as Blutner and A&P who use entailment and semantic anomaly as data for their theories "are mistaken about the data their theories need to explain".[31]

1.a.ii. Hidden Disagreement

The preceding discussion establishes that different semantic theorists quite clearly wish to explain different phenomena with their theories. However, not all disagreement about the explanatory domain of semantics is evident on the surface. There are cases where theorists appear to agree on domain constraints, but this appearance belies an underlying disagreement.

Consider the case of truth conditions. Truth conditions are one of the most exemplary explananda for formal semantics;[32] Lewis famously pronounced that "[s]emantics with no treatment of truth conditions is not semantics".[33] However, I claim that two distinct types of "truth condition" data have been conflated in the literature.

Pustejovsky (1995) asks us to consider the following sentence:

(7a) John baked the cake.

When you encounter this sentence, you probably imagined a certain type of process: John creating the batter for the cake, pouring it into pans and heating it in the oven, then removing the pans when the sponges are cooked through. Compare (7a) to the following:

---

[31] *ibid.* 77.
[32] Though there has been recent dissent about the extent to which semantics is responsible for explaining truth conditions (*e.g.*, Pietroski, "The Character of Natural Language Semantics." and Schiffer, "Meaning and Formal Semantics in Generative Grammar.").
[33] Lewis, "General Semantics." 18.

(7b) John baked the potato.

When you encounter this sentence, you probably imagine a different (and simpler) process: John putting the potato into the oven and heating it.

Pustejovsky takes this as evidence that, roughly, 'bake' makes a different contribution to 'baked the cake' than it does to 'baked the potato'; 'baked the cake' gets a "creative" reading whereas 'baked the potato' gets a "change of state" reading. However, this is not because 'bake' is ambiguous in the same way as 'bank', associated with more than one meaning in the lexicon. Instead, Pustejovsky takes the relatedness of the two senses as evidence that they both derive from a single lexical entry for 'bake'. Glossing over many details, the idea is that ||cake|| contains in its qualia structure the information that cakes are made by baking, whereas ||potato|| does not. Through an operation of *co-composition*, 'bake' "picks up" this feature of 'cake' during composition. This causes the information that the cake was created by John's baking to be included in the semantic value of (7a), whereas the same information about the potato is not included in the semantic value of (7b). Under this account, the interpretive difference between (7a) and (7b) is entirely due to the semantic theory.

Next, consider the following situation: John purchases a cake from a bakery. He heats his oven to 350°F and inserts the cake, removing it several minutes later. Someone then utters (7a). I am disposed to judge this utterance as true, and Fodor and Lepore (henceforth F&L) are likewise disposed.[34] They take this observation as the data for an argument against Pustejovsky's theory of 'bake'. The creative sense of 'bake' does not apply to this situation. If the meaning of 'bake' combined with the meaning of 'cake' determines that the creative sense is used in the

[34] Fodor and Lepore, "The Emptiness of the Lexicon: Critical Reflections on J. Pustejovsky's the Generative Lexicon."

interpretation of (7a), then we would expect the sentence to be judged false in this situation. But it is judged true, so Pustejovsky's theory has seemingly been falsified.

Perhaps surprisingly, F&L take this as evidence that 'bake' is indeed ambiguous (in the same way as 'bank'). However, one can imagine a third theoretical option that makes a different use of F&L's data.[35] Initially, we thought there was evidence that John *creatively* baking the cake was part of the truth conditions of (7a). However, the fact that (7a) can be true in a scenario in which John *non-creatively* bakes the cake demonstrates that initial appearances were misleading; creative baking cannot be required for the truth of (7a). Under this option, 'bake' is not ambiguous (*contra* F&L) and 'bake the cake' only makes a single contribution to truth conditions that does not require the cake to have been creatively baked (*contra* Pustejovsky). The initial appearance of a "creative" sense of 'bake the cake' arises from broadly pragmatic factors that include world knowledge about how cakes are prototypically baked.

Notice, though, that the type of evidence given by Pustejovsky in favor of his account is different from the type of evidence given by F&L (and my third option) against the account. Pustejovsky gives us two sentences and asks us to imagine the situations that they describe. F&L give us a sentence *and* a description of a situation before asking us whether the sentence is true or false in that situation. I'll call Pustejovsky's task the determination of a *prospective application condition* of a sentence, and I'll call F&L's task a *situated application judgment* of a sentence.[36]

The theoretical position of the third option I just outlined, then, would be that situated application judgments take evidential priority over prospective application judgments.

---

[35] King, "W(h)Ither Semantics!(?)." seems sympathetic to this view, though it discusses different examples.
[36] I use the term 'application' instead of 'truth' to accommodate non-declarative sentences like imperatives and questions, since such sentences have been relevant in the literature.

Prospective judgments might provide some initial evidence for truth conditions, but they can also be distorted by non-semantic noise (prototypical world knowledge, in this case). Situated judgments can provide data in which the non-semantic noise is "controlled". This theoretical position is highly compatible with a *lowest common denominator* (henceforth LCD) semantics. Soames has been a recent advocate of such a theory, which defines the lexical meaning of a sentence as "that which is common to what is asserted and conveyed by [literal] utterances of it in all [normal] contexts".[37] Presumably, the lexical meaning of a subsentential expression can then be defined as that which is common to the expression's contribution to the sentential meanings of literal utterances in normal contexts. A truth-conditional LCD semantics would, at the very least, expect truth-value judgments to be consistent with the sentence valuations given by the theory.[38] More strongly, though, it would also expect truth-value judgments *in every normal context* to be consistent with the theory's sentence valuations. Clearly, such a theory would give situated judgments evidential priority over prospective judgments when they come into conflict.

So, it seems that Pustejovsky's theorizing gives prospective judgments greater evidential weight than an LCD semantics. If Pustejovsky adopts Asher's domain constraint, this is to be expected; prospective judgments certainly seem to be quick and unconscious. Of course, such a tension between Pustejovsky and the LCD theorist could be used as another example of differing explananda within semantics. However, the takeaway I wish to stress is that the prospective / situated distinction is frequently overlooked when evidence is presented for or against truth-conditional semantic theories.[39] When it's claimed that semantics must give a treatment of

[37] Soames, "Drawing the Line between Meaning and Implicature—and Relating Both to Assertion.", 462.
[38] As long as the theory also takes these sorts of judgments seriously.
[39] Though I'm not claiming that this is always to the detriment of the theorizing.

truth conditions, it's usually uncertain whether and to what extent evidence for truth conditions is provided by prospective application judgments vs. situated application judgments.

1.b. Characterization of the Data

Differing domain constraints aren't the only way that data leads to plurality within semantics. Suppose that two theorists agree that a given phenomenon (say, persistent anomaly) is to be explained by semantics. The theorists still might disagree on how to *characterize* the data falling under that phenomenon. I'll define the characterization of a particular datum as (i) the range and structure of the possible values that could be assigned to the datum, and (ii) the actual assignment of a value to the datum. Both parts of the characterization of the data are sources of plurality.

1.b.i. Possible Values of Data

When I claim that 'married bachelor' is semantically anomalous, it's natural to ask: "As opposed to what?". The simplest response, and the one most commonly found in the literature, is that I was choosing one of a set of two possible values for the sentence:

{anomalous, non-anomalous}

Indeed, many different meaning-relevant phenomena are often similarly characterized in a binary way. Common examples are shown in Figure 1.

Phenomenon	Binary Characterization
Felicity	A sentence or utterance is either acceptable or unacceptable.
Truth Values	A sentence or utterance is either true or false.
Entailment	A sentence either entails or doesn't entail another.
Contradiction	A sentence either contradicts or is compatible with another.

Figure 1: Binary Characterizations of Semantic Phenomena

While binary characterizations of data are common, the range and structure of possible values can be more complex. For instance, we might introduce an intermediate anomaly value for phrases that are sort-of-but-not-quite anomalous and place an ordering on the new set of values:

< anomalous, semi-anomalous, non-anomalous >

Taking things to one extreme, we might even claim that anomaly should be represented by a continuous rather than discrete variable that takes degrees of anomaly as values. Similar moves can be made for the other phenomena mentioned above; the range of possible truth values is a particular locus of debate.[40]

1.b.ii. Actual Values of Data

Even if theorists agree on the range and structure of the *possible* values of semantic data, they can disagree on the *actual* values of data. There are numerous examples in the literature, and I will only be able to provide a small sample. However, this type of plurality is especially relevant in semantics when compared to many other sciences. Semantic data often consist of considered judgments made by semanticists. This means that semanticists are often both the source of data and its experimental observer. Since semantic data can be "subjective" in this way, disagreements over the data can be among the most intractable.

Reconsider the example of 'married bachelor'. Earlier, I claimed that this phrase is semantically anomalous. However, not everyone would agree with this (lack of a) value-assignment.[41] It could be argued that 'married bachelor' *does* successfully compose, but its extension is necessarily empty. Therefore, attributions of 'married bachelor' do have a truth

[40] Shramko and Wansing, "Truth Values."
[41] Fodor and Lepore, "The Emptiness of the Lexicon: Critical Reflections on J. Pustejovsky's the Generative Lexicon."

value, which they wouldn't if the phrase were anomalous. The expression seems odd because that truth value is necessarily FALSE.

For another example of anomaly data causing plurality, consider the status of cross-polar sentences such as (8a) and cross-modal sentences such as (8b):

(8a) John is taller than Mark is short.

(8b) Susan is smarter than John is tall.

Some theorists consider such sentences anomalous, constructing theories that ensure they fail to compose and have a truth value.[42] Other theorists consider such sentences semantically acceptable, constructing theories that allow them to successfully compose.[43]

For an example of a disagreement over truth values, consider the following examples from Travis,[44] supposing that the only milk in the fridge is a puddle at the bottom:

(9) There's milk in the fridge.

Travis considers two contexts of utterance. In the first context, the utterance occurs in a discussion of the cleanliness of the fridge. In the second context, the utterance occurs after the hearer indicated they would like some milk to drink. Travis claims that (9) is true in the first context but not the second. He uses this disparity to attack a particular view of semantics (exactly which view isn't important) that Berg[45] is interested in defending. Acknowledging that the disagreement can turn into a "slugfest of intuitions" at this point, Berg claims that (9) is actually true in both contexts.

[42] Kennedy, "Polar Opposition and the Ontology of 'Degrees.'"
[43] Bale, "A Universal Scale of Comparison."
[44] Travis, *The Uses of Sense: Wittgenstein's Philosophy of Language*. 18-19.
[45] Berg, "Is Semantics Still Possible?"

I could cite examples of disagreement on data indefinitely, but I'll wrap up with an example from the literature on presupposition. This example concerns the sort of presupposition carried by the following sentence:

(10) Tom is proud of his car.

(10) presupposes that Tom owns or at least has possession of a car. If the presupposition fails and Tom doesn't possess a car, then utterances of (10) are infelicitous.

What happens when we replace 'Tom' with a quantifier phrase?

(11) Every German is proud of his car.

Geurts and van der Sandt claim that the default reading of (11) is roughly paraphrased by (11a):[46]

(11a) Every German who owns a car is proud of it.

This reading is an example of the supposed phenomenon of *intermediate accommodation*, which Geurts and van der Sandt argue results from a presupposition tied to 'his car'.

Compare (11a) with another proposed paraphrase of (11):

(11b) Every German owns a car and is proud of it.

Citing slightly different sentences, Beaver demonstrates a preference for the (11b) interpretation; indeed, he seems skeptical that the (11a) interpretation is available at all:[47]

> Some time ago, I had managed to convince myself that, in agreement with van der Sandt, reading [(11b)] was indeed possible, but perhaps a little marginal. But asking naive informants about other sentences of this type caused me to doubt this view: the informants could not seem to get the reading at all.

This disagreement over the data corresponds to a disagreement over theory: while Geurts and van der Sandt argue that intermediate accommodation is the result of presuppositions tied to sentence constituents, Beaver argues that it is the result of broader contextual factors.

[46] Geurts and van der Sandt, "Domain Restriction."
[47] Beaver, "Accommodating Topics."

1.c. Second-Order Constraints

Semantic theories aren't just constrained by the data and the general theoretical desiderata. They're also constrained by what I will call *second-order* constraints on the shape of semantic theories. I will roughly characterize second-order constraints as domain-specific theoretical constraints that determine the form of semantic theories and explanations. There are (very roughly) two important types of second-order constraint:

a. Constraints on the properties of the objects posited by the theory.

b. Constraints on how the objects posited by the theory relate and interact.

Even if the phenomena to be explained is fully agreed-upon and there is no tension among various epistemic goods, adherence to different second-order constraints can still cause plurality.

An exemplar of a much-disputed second-order constraint is the principle of *Reverse Compositionality*, most notably defended by Fodor (though he most explicitly defends it for concepts instead of word meanings).[48] The "forward" compositionality principle introduced in a previous section says that the meanings of complex expressions are determined by the meanings of their constituents. The reverse principle claims that the meanings of constituents are in some sense "determined" by the meaning of the complex expression.

The sense in which constituent-meanings are determined by complex-meanings is often made more precise by being tied to the concept of *understanding* a language. One version of the forward principle says that understanding the atomic constituents of a complex expression, in combination with knowledge of syntax and composition rules, is sufficient to understand the complex. Thus, we can understand 'John read two books' because we understand 'John', 'two',

[48] Fodor, "Replies to Critics."

and whatever other constituents the sentence has. The Reverse Compositionality principle says that understanding the constituents of a complex expression is *necessary* to understand the complex expression. Johnson states the principle as follows:[49]

> **Reverse Compositionality:** A speaker understands the meaning of a (non-idiomatic) complex expression of her language only if she understands the meanings of the primitive expressions and syntactic configurations that the former is built out of.

The principle is controversial, and is often a point of contention (either implicitly or explicitly) between competing theories.

For an example of an area in which the principle is an implicit source of contention, consider two competing theories of prepositions. Both theories operate within the framework of Optimality Theory. Applied to the study of meaning, Optimality Theory provides models of the acts of producing and interpreting utterances. The models determine the optimal syntactic output to express a given meaning, the optimal meaning to interpret a syntactic input, or both simultaneously. Possible meanings or expressions are evaluated relative to a set of ranked constraints. The meaning or expression that incurs the lowest "cost" by performing the best relative to the constraints is selected as the output of interpretation or production.

Zwarts[50] and Hogeweg[51] offer different theories of the production of spatial prepositions. As models of expression production, their models take a meaning that the speaker wishes to express as input and produces an output expression. Both theories assume that there is a set of semantic features (CONTAINMENT, SUPPORT, SUPERIORITY, etc.) that are expressed by utterances

[49] Johnson, "On the Nature of Reverse Compositionality." 38.
[50] Zwarts, "Priorities in the Production of Prepositions."
[51] Hogeweg, "Optimality Theoretic Lexical Semantics."

containing spatial prepositions such as ‘in’, ‘on’, ‘above’, etc. Where they differ is in the way that they associate prepositions with features in the lexicon.

For Zwarts, ‘in’ is associated only with the semantic feature CONTAINMENT (henceforth CONT) and ‘on’ is associated only with SUPPORT (henceforth SUPP) in the lexicon. Imagine a situation in which there is a tool and a box such that the tool is contained within the box and the box is supporting the tool. The speaker wishes to describe this situation. We can say that the meaning they wish to express contains both CONT and SUPP; for simplicity’s sake, we will call this input meaning [CONT, SUPP]. According to Zwarts, ‘in’ expresses CONT and ‘on’ expresses SUPP, so how does the speaker choose one of these expressions?

In Optimality Theory, the answer is provided by the relevant constraints. In this case, the relevant constraints are those of faithfulness to the input (*i.e.* FAITH). This constraint is violated every time there is a semantic feature in the input that is not expressed in the output. ‘In’ violates FAITH with respect to SUPP since the input meaning contains SUPP but ‘in’ only expresses CONT. Likewise, ‘on’ violates FAITH with respect to CONT. Counting constraint violations clearly is not sufficient to choose between ‘in’ and ‘on’, so Zwarts’ solution is to rank the constraints. This solution treats faithfulness to SUPP (*i.e.* FAITH(SUPP)) as a separate constraint from faithfulness to CONT (*i.e.* FAITH(CONT)). Zwarts’ claim is that FAITH(CONT) is ranked higher than FAITH(SUPP); in other words, a violation of FAITH(CONT) is more severe than a violation of FAITH(SUPP). Since ‘on’ violates FAITH(CONT) and ‘in’ violates FAITH(SUPP), ‘on’ incurs a greater cost than ‘in’; thus ‘in’ is the preferred word to express the intended meaning in the toolbox scenario. This clearly fits with intuition, which says that the preferred way to describe the scenario is ‘The tool is *in* the box’.

Unlike Zwarts, Hogeweg allows 'in' and 'on' to be associated with more than one semantic feature in the lexicon. According to this theory, the meaning of 'in' contains both CONT *and* SUPP. The main benefit of this theory is that it doesn't force us to include different faithfulness constraints for different features; there is just one FAITH constraint, and violating FAITH(CONT) is no more or less costly than violating FAITH(SUPP). In the scenario where we wish to express both CONT and SUPP, the use of 'in' incurs no faithfulness violations since it expresses both features, and the use of 'on' incurs one violation since it doesn't express CONT. Unlike Zwarts' theory, then, we can see that 'in' is preferred to 'on' simply by counting constraint violations.[52]

However, this benefit comes at the expense of the principle of Reverse Compositionality. Indeed, Zwarts' arguments seem implicitly motivated by this principle (or something similar).[53] Suppose that a bat is flying around inside an attic without landing on any surface. The preferred way to describe this scenario is with (12a) instead of (12b):

(12a) The bat is *in* the attic.

(12b) The bat is *on* the attic.

Thus, 'in' gets used to describe a situation in which CONT attains but not SUPP. According to Zwarts, the fact that 'in' can refer to a situation in which SUPP does not attain means that SUPP cannot be an "inherent lexical feature" of 'in'.[54]

[52] Of course, both theories become more complicated in order to accommodate more prepositions and more data, but this simple gloss is sufficient for my purposes.
[53] Zwarts only explicitly makes this type of argument for 'around', but I imagine it would also apply to 'in'.
[54] Zwarts 15.

It's easy to see how this argument is roughly one that an adherent to Reverse Compositionality would make. Suppose that we give the following as a rough first gloss of the meaning of (12a):

(d) ⟦The bat is in the attic⟧ = $\exists !x \exists !y(\text{bat}(x) \,\&\, \text{attic}(y) \,\&\, \text{CONT}(y, x))$.

According to Reverse Compositionality, knowing that (d) is the meaning of (12a) is sufficient to understand the constituents of the sentence such as 'the', 'bat', 'in', and 'attic'. Of course, we have to be careful here in order to make this claim actually plausible. The claim shouldn't be that knowing the meaning of (12a) is sufficient *tout court* for knowing that, *e.g.* $\lambda x(\text{bat}(x))$ is the meaning of 'bat'. Without further qualifications, it's possible that the interpreter understands (12a) while associating 'bat' with ⟦attic⟧ and vice versa (while making the appropriate adjustment to ⟦in⟧).

Let's assume, though, that the appropriate qualifications can be added in order to ensure that the language user who understands (12a) also knows that 'bat' contributes ⟦bat⟧, 'in' contributes ⟦in⟧, etc. In other words, the language user is able to successfully "reverse engineer" the meaning of 'in' from (d). With the appropriate qualifications, knowledge of (d) is sufficient to know that CONT is a semantic feature of 'in'; however, it is insufficient to know that SUPP is also a semantic feature of 'in' since SUPP appears nowhere in (d). Fodor and Lepore put it quite clearly (replacing their examples with my own):[55]

> [Reverse Compositionality assumes that] each constituent expression contributes the whole of its meaning to its complex hosts. If that's right, then if you understand 'building a house', it follows that you know everything you need to determine the meanings of 'build' and 'house' and '-ing'.

[55] Fodor and Lepore, "Why Compositionality Won't Go Away: Reflections on Horwich's 'Deflationary' Theory."

As (12a) shows, a competent language user can know the meaning of complex expressions containing 'in' without knowing that SUPP is a component of ⟦in⟧. According to Reverse Compositionality, this implies that SUPP is not a component of ⟦in⟧.

If every feature of a constituent meaning is inherited by the complex meanings it contributes to, context (linguistic or extra-linguistic) can only add semantic material and not delete it. Unless you're an atomist like Fodor, this strongly suggests a view under which many lexical meanings are highly *underspecified*, containing fewer semantic features than contextual meanings.[56]

In contrast, Hogeweg offers an *overspecification* view: many lexical meanings contain more semantic features than contextual meanings, and context can only delete features of constituent meanings. This view clearly fails to adhere to Reverse Compositionality: there can be semantic material in the meaning of a constituent that fails to appear in the meaning of some complex containing the constituent. Thus, someone who hears an utterance of (12a) and recovers (d) as its meaning has, according to Hogeweg, somehow canceled the SUPP information provided by 'in'. For instance, they might have used some information about bats' (or this particular bat's) propensity to fly under the relevant contextual circumstances. Since this view doesn't adhere to Reverse Compositionality, the fact that some interpretations of expressions containing 'in' fail to contain SUPP is not definitive evidence that SUPP is not part of the meaning of 'in'.

[56] Vicente, "Polysemy and Word Meaning: An Account of Lexical Meaning for Different Kinds of Content Words."

1.d. Nature of the Target System

When theorists are constructing semantic theories/models, they are typically constructing theories *of* something. In the widely-used parlance in philosophy of science, this something is the *target system* of the theorizing.[57]

Ball gives a non-exhaustive list of different ways in which theorists have characterized the target system of semantics.[58] As already mentioned, there is the Chomskian characterization of the target as a mental faculty; followers of Lewis would characterize the target as some group of social regularities, especially conventions;[59] Soames characterizes the target as "what information is encoded by sentences relative to contexts".[6061]

The choice of target system has clear effects on semantic practice. If the target is a Chomskian language faculty, semantic theories are constrained by broader psycholinguistic and even neurobiological theories. This increases the potential relevance of theories of language acquisition and brain architecture for semantics. If the target is instead a Lewisian group of conventions, semantic theories are constrained by broader sociological theories. This increases the potential relevance of theories of cultural evolution for semantics.

1.e. Relation between the Model and the Target

Semantics provides models and theories of some target system (perhaps with a "model system" acting as a hypothetical middleman between the two), with some options for that target

[57] Godfrey-Smith, "The Strategy of Model-Based Science."
[58] Ball, "Semantics as Measurement."
[59] Lewis, "Languages and Language."
[60] Soames, "Semantics and Semantic Competence." 182.
[61] Ball lists the conception of the target of semantics in Davidson, "Truth and Meaning.", separately from the others. Given the way I'm dividing things up, I'm not sure that it is different enough from Soames' conception of the target system to warrant a separate listing in this section, though there could be differences that warrant a separate listing in the next section.

system listed above. According to Godfrey-Smith, this still leaves the *construal* of the theory, or the relationship between theory and target system, underspecified.[62]

For some theorists, semantic theory is intended to directly map onto causally efficacious components of the target system. This is particularly true of Chomskians who claim that the semantic theory is somehow represented in the brain and used in utterance comprehension and production (though it should be noted that Chomsky is careful to distinguish between the speaker's "theory" of their language and the linguist's theory of their competence). For others, semantics is construed "merely as a predictive device" with limited aspirations towards modeling any actual causal relations.[63] Soames describes semantic theory as "a rational reconstruction, rather than a causal explanation, of our ability to use sentences in communicative situations".[64] Under this construal of semantics, the input-output profile of the theory (*e.g.* the pairing of sentences with meanings) must match reality but the mechanisms of the theory are only constrained by the condition of "rationality".

How the relationship between the model and the target system is construed also has an effect on semantic theorizing. Suppose the target system is a Chomskian language faculty. On one possible construal of the semantic theory, its components and mechanisms are intended to map onto causally efficacious components and mechanisms in the brain. In this case, neurobiology will be required to inform and constrain semantic theories. On another possible construal of the semantic theory, the correspondence between its components and the components of the brain need not be as strict, but the theory must still be compatible with human

[62] Godfrey-Smith, "The Strategy of Model-Based Science." 733.
[63] *ibid.*
[64] Soames, "Drawing the Line between Meaning and Implicature—and Relating Both to Assertion." 461.

development and evolutionary history. In this case, neurobiology is less relevant than studies in developmental psychology and evolutionary biology.

2. Plurality as a General Feature of Behavioral Sciences

Observing the lack of consensus in semantics could potentially lead to despair: since semanticists can't even agree on what their subject matter is, what data their theories are answerable to, and what form their theories should take, what prospects are there for the field? In order to avoid this, it will be helpful to check whether semantics is in a unique state before moving on. If we can make similar observations about other fields, then semantics is not in a uniquely "bad" situation (of course, the situation will only seem bad from a Monist perspective).[65] Fortunately, there are such examples.

Longino provides an analysis of the study of human aggression within biology, psychology, and sociology that observes similar plurality.[66] Longino isn't sure that human aggression is even a unified subject matter: "there is no guarantee that the studies [of aggression] are studies of the same phenomenon".[67] Like 'meaning', 'aggression' is a vague and ambiguous term that has been operationalized in different ways (*i.e.* there are multiple phenomena that are viewed as indicative of aggression). In various studies, aggression has been measured using conviction of violent crime; fighting; delinquency; violent rage; anger, irritability, and verbal aggression; hitting a doll; scores on psychological tests; and diagnoses of various disorders. As in semantics, there is disagreement on exactly which of these data is relevant for the study of aggression. There is also disagreement on what exactly aggression is: is it a stable individual trait

[65] At least, semantics is not in a uniquely bad situation with respect to the concerns I'm raising; there is a long history of criticizing the scientific status of semantics on other grounds.
[66] Longino, "What Do We Measure When We Measure Aggression?"
[67] *ibid.* 697.

or something that is more situational? Thus, as in semantics, there is disagreement over the nature of the target system. Finally, there is disagreement that might be considered analogous to disagreement over second-order constraints in semantics. Consider the debate over to what extent aggression is determined by genetics vs. the environment. Intrauterine effects are considered part of the environment by the genetic approach whereas they aren't considered part of the environment by the environmental approach (this discrepancy could be seen as analogous to disagreements over the explanatory roles of semantics and pragmatics).[68]

Economics is another behavioral science with similar plurality.[6970] Robbins defined economics as the study of the relationship between ends and scarce means.[71] While this domain constraint has been widely influential, it is also far from universally accepted among practitioners.[72] Some economists argue that the definition is too narrow, excluding ethical considerations and situations in which there is a *surplus* of labor or capital. Others argue that the definition is too wide, extending to situations of scarcity in domains that *prima facie* fall under the purview of other fields.

The nature of the target system is also disputed. One issue concerns the objects found in the target system. Neoclassical microeconomic theories include individual agents in their ontologies but not any qualitatively different notion of the social: the social is simply the aggregate of individual economic agents. Some argue that this is a primary source of difficulties in reducing macroeconomics to microeconomics.[7374] A broader issue is whether the target system

[68] Longino, "Foregrounding the Background." 655.
[69] Dow, "Pluralism in Economics."
[70] Sent, "Pluralisms in Economics."
[71] Robbins, *An Essay on the Nature and Significance of Economic Science*.
[72] Backhouse and Medema, "Defining Economics: The Long Road to Acceptance of the Robbins Definition."
[73] Sent, "Pluralisms in Economics."
[74] Hands, "Caveat Emptor: Economics and Contemporary Philosophy of Science."

is *closed* or *open*. In a closed system, the nature of each relevant variable and the relationships among them are fixed and knowable. An open system is more fluid, both in its boundaries and in the properties of the relevant variables.[75]

Finally, the properties of individual agents in economic models are a major source of controversy. Neoclassical theories assume that agents are rational in the broad sense that they consistently maximize some utility or profit function. Behavioral economists such as Simon[76] and Kahneman and Tversky[77] criticize these assumptions as unrealistic. They argue that real people and firms exhibit *bounded* rather than perfect rationality, are prone to systematic biases, and use heuristic methods rather than optimization to make decisions. Neoclassical economists have responded by arguing that these criticisms incorrectly construe the relationship between economic models and the target system.[78] Perhaps rationality assumptions are justified because they sufficiently approximate the behavior of real economic agents, despite their imperfections. Or perhaps they are justified because the models do a good job predicting aggregate behavior regardless of how well they capture individual decision-making.

In multiple behavioral sciences, plurality persists because of variation among most or all of the five types of background constraints. The good news for semantics is that it is not (at least qualitatively) an outlier among its behavioral science brethren.[79] But this news still might not be very encouraging; after all, those fields are also much maligned. Perhaps plurality is problematic in all of these fields.

---

[75] Chick and Dow, "Formalism, Logic, and Reality: A Keynesian Analysis."
[76] Simon, *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*.
[77] (Kahneman and Tversky, "Prospect Theory: An Analysis of Decision under Risk."
[78] Becker, "Irrational Behavior and Economic Theory."
[79] Unsurprisingly, classifying semantics among the behavioral sciences is itself controversial.

3. Semantic Projects

In order to address this issue, I need to do a bit more groundwork to establish the conception of scientific plurality that will be used in the rest of the paper. I view semantics as consisting of a set of many distinct *projects*, which are identified at least partially by background constraints. We might observe that some background constraints are frequently jointly operative. For instance, a significant portion of semantic practice might target a Chomskian language faculty, adhere to both Forward and Reverse Compositionality, and study semantic anomaly as a binary phenomenon; another significant portion of semantic practice might also target a Chomskian language faculty, adhere to Forward but not Reverse Compositionality, and deny that anomaly is a semantic phenomenon at all; a third might target a social system of conventions, adhere to Forward and Reverse Compositionality, and study semantic anomaly as a graded phenomenon. In this case, we can identify three distinct projects as shown in Figure 2.

Project	Target System	Second-Order Constraints	Domain Constraints	Characterization of Anomaly
Project 1	Language Faculty	Forward and Reverse Compositionality	Include semantic anomaly	Binary
Project 2	Language Faculty	Forward Compositionality	Exclude semantic anomaly	N/A
Project 3	Social Convention	Forward and Reverse Compositionality	Include semantic anomaly	Graded

Figure 2: Three Distinct Semantic Projects

Of course, variation among background constraints isn't the only way to describe the plurality of a scientific field. Kuhn used the idea of a *school*.[80] A school differs from the more

[80] Kuhn, *The Structure of Scientific Revolutions*.

famous *paradigm* in that a school does not approach universal acceptance by practitioners in a field. Plurality occurs when there are multiple competing schools.

I've defined scientific projects using sets of background constraints, but Kuhn defines schools and paradigms using *exemplars* of puzzle-solutions. He claims that the process of constructing solutions to new puzzles is largely case-based instead of rule-based; scientists tend to draw inspiration from exemplary solutions to similar puzzles instead of explicitly applying rules. Thus, a school or paradigm is largely characterized by a loose set of exemplary puzzle solutions, not a set of theoretical principles.

I don't wish to be seen as disagreeing with Kuhn on a question of how to best describe plurality *in general*. The choice to describe semantic plurality in terms of background constraints was a pragmatic one. It's plausible that formal and distributional semanticists draw on a significantly different set of exemplars; it's even plausible that different "schools" within formal semantics (say, Optimality Theory and dynamic semantics) also draw on significantly different sets of exemplars. However, I'm also interested in plurality among projects that are closer to each other in theoretical space. In such cases, plurality is present, yet discerning significant differences among the projects' exemplars is difficult.

I also don't wish to give background constraints an undue role in the actual process of theory construction. Sometimes background constraints are made explicit in the course of theorizing, but, as my examples should make clear, other times they are only derived *ex post facto*. Indeed, such *ex post facto* derivations might also be difficult and controversial; for instance, there might be no clear, principled, agreed-upon way of drawing a distinction between

the entailments that a project wishes to explain and those it doesn't. For Kuhn, such a situation isn't necessarily problematic.

Kuhn had strong Monist inclinations (though I doubt he'd accept the letter of Scientific Monism as stated by Kellert *et. al*). He defined a *mature* science as a science with a single dominant paradigm (except, perhaps, in periods of revolutionary science when an old paradigm is overthrown). In contrast, he pejoratively described sciences containing a set of schools as opposed to a single dominant paradigm as *immature*. Therefore, I primarily disagree with Kuhn by promoting a positive attitude towards plurality. It's unclear how real this disagreement is, however, since I characterize plurality using background constraints and Kuhn characterizes it using exemplars of puzzle-solutions.

Chapter 2: Pluralism I - Multiple Projects

I've stated my descriptive view of semantic plurality: semantics contains many projects, where projects are defined by sets of background constraints. With this view in hand, I now turn to my prescriptive view of semantic Pluralism. My Pluralist view has two components:

1. Maintain multiple projects.
2. View background constraints as soft.

In this chapter, I motivate (1) and illustrate one of its important consequences. Since (1) covers ground that has already been covered extensively by other Scientific Pluralists, many motivations and consequences can be imported with only minor modifications from their discussions. As a result, most of the rest of this paper discusses the motivations and consequences of (2) in Chapter 3.

1. Motivation

Recall Tenet 2 of Scientific Monism:

> **SM.2:** The nature of the world is such that it can, at least in principle, be completely described or explained by [a single, comprehensive] account.

This claim can be interpreted broadly as a claim about the world in its entirety (or at least the physical world), or it can be interpreted narrowly as a claim about a specific part of the world. A narrow Monism about semantics claims that *meaning in human natural language* can be captured by a single, comprehensive account.

Chapter 1's survey of semantic plurality casts doubt on this claim. Semanticists study different target systems and different phenomena that emerge from those systems; they characterize the same phenomenon in a multitude of ways, and use different properties, relations, and laws to explain the phenomena. One plausible explanation of this plurality is a lack of unity

in the subject matter. Perhaps some semantic projects are simply studying different things with different explanations that can't be unified into a single, comprehensive account.

If this is right and (SM.2) is false of meaning in natural language, it seems clear that semantics should maintain multiple projects. As argued by Kellert *et. al*, the falsity of (SM.2) undermines the rest of the tenets of Scientific Monism:[81]

> If the world cannot be accounted for by a single comprehensive account, then there cannot be methods of inquiry that if correctly pursued would yield such an account. Hence, we should not assume that [(SM.3)] is true. And if we don't know whether the world can be fully accounted for by a single, comprehensive account, then it seems unreasonable to accept or reject scientific methods according to whether they can yield such an account [(SM.4)] or to evaluate scientific theories in terms of how close they come to providing a complete and comprehensive account [(SM.5)].

And, of course, if the semantic part of the world can't be described by a single, comprehensive account, then establishing such an account is impossible. If meaning is best captured using multiple accounts, then semantics should contain multiple theories, methods of inquiry, and underlying assumptions; in other words, semanticists should maintain multiple theoretical projects.

If (SM.2) is true, it's possible to capture meaning using a single, comprehensive account. In this case, it's much more plausible that (SM.1) is true: the ultimate goal of the study of meaning is to find a Monist account. Even so, semanticists would still have good reasons to maintain multiple projects. As mentioned, many authors have argued extensively for the benefits of plurality even if the ultimate goal is Monistic. I will provide one example of such an argument by revisiting Kuhn.

[81] Kellert, Longino, and Waters, "Introduction: The Pluralist Stance." xi.

Since he was a Monist (broadly speaking), Kuhn espoused the virtues of a mature science with a dominant paradigm and the vices of an immature science with multiple competing schools. Kuhn studied the work of pre-paradigm scientists of physical optics. He observed that the work "was often directed as much to the members of other schools as it was to nature",[82] attacking the foundations of other schools rather than building on top of its own. The emergence of a paradigm frees scientists from this internecine conflict, allowing the field to build the edifice of normal science on a secure (for now) foundation.

How, then, does a scientific field reach maturity and transition from the plurality of multiple schools to the monism of a paradigm? Kuhn observed a general pattern that he illustrated using the study of electricity. In the pre-paradigm state of this field, the "fluid" theory of electricity was just one school among many; another school took attraction and frictional generation to be the primary electrical phenomena, while another school regarded attraction and repulsion as equally important. Each school had different strengths and weaknesses, accounting for some phenomena well and other phenomena poorly, if at all.

Ultimately, a paradigm did emerge from the fluid school. Directed by the assumptions and theories of their school, fluid theorists invented the Leyden jar, the best early machine for storing and discharging electricity. Explaining the behavior of the jar became an issue of utmost importance to electricians. The fluid school won out, and this explanatory success became the most persuasive part of the argument that established the school as the first electrical paradigm.

The important lesson from this example is that the fluid school emerged as the dominant paradigm *via* work within the school similar to normal science. One way for a school to try to

[82] Kuhn, *The Structure of Scientific Revolutions*. 13.

improve its status relative to its competitors is to engage in the destructive work of criticizing those other schools and attacking their foundations. However, Kuhn argues that the constructive work of solving puzzles and demonstrating the school's explanatory power is generally more effective.

Assuming his argument is just as sound for theoretical projects as it is for schools, Kuhn gives us a reason to maintain multiple projects. The argument also elucidates what it means to maintain multiple projects. If a field contains multiple projects and its ultimate goal is Monistic (which is reasonable on the assumption of (SM.2)), positive explanatory work within the projects is a better way to achieve this goal than inter-project criticism. The best way for a project to improve its standing relative to its competitors is to demonstrate its own merits.

The current argument for Pluralism remains agnostic on (SM.2). Instead of taking a metaphysical stand, it is a pragmatic argument denying tenets 4 and 5 of Scientific Monism, repeated below:

> **SM.4:** Methods of inquiry are to be accepted on the basis of whether they can yield [a single, comprehensive] account.
>
> **SM.5:** Individual theories and models in science are to be evaluated in large part on the basis of whether they provide (or come close to providing) a comprehensive and complete account based on fundamental principles.

Specifically, this argument denies that methods and theories (*i.e.* projects) are to be largely evaluated on the basis of whether they can provide a *single* account of the domain. Evaluating a project's potential to provide a single account necessarily involves comparison with other projects: is the project able to best its peers, becoming a dominant paradigm? We move away from this zero-sum thinking when we *want* to maintain multiple projects.

2. Consequences

Semantics has ample motivation to maintain multiple projects. If (SM.2) is false and meaning can't be captured using a single, comprehensive account, then it seems that maintaining multiple projects clearly follows as a rational response. And if (SM.2) is true, then (perhaps counterintuitively) maintaining multiple projects for the time being could be the best way to achieve Monist goals. From this argument, a picture of semantic theorizing from the Pluralist perspective has started to emerge.

For Kuhn, a main impediment to the progress of a plural science is that "there are always competing schools, each of which constantly questions the very foundations of the others". Pluralism alleviates this problem prior to the emergence of a paradigm (which might never occur!). Many points of contention within semantics are viewed differently from the Pluralist perspective.

Suppose we discover tension between background constraints A and B up to but not limited to outright contradiction. If one group prefers to hold on to A while another prefers to hold on to B, the two groups can find themselves going around in circles as they argue over which constraint to collectively reject. Viewed from the Pluralist perspective, though, this conflict need not have a single resolution. We can view the tension between A and B as a point where two projects should simply go their separate ways. If some plurality is desirable, then it's possible that it's desirable to maintain some projects adhering to A and others adhering to B.

Consider another fairly recent dispute over the principle of Reverse Compositionality (repeated below):[83]

[83] Johnson, "On the Nature of Reverse Compositionality."

> **Reverse Compositionality:** A speaker understands the meaning of a (non-idiomatic) complex expression of her language only if she understands the meanings of the primitive expressions and syntactic configurations that the former is built out of.

Fodor is a major proponent of this principle, using it in the foundations of his theory of meaning. He argues that, because Reverse Compositionality is true of meaning, meanings must be atomic.[84] In response, Johnson argues against Reverse Compositionality by presenting a counter-example from a well-known theory of meaning. According to this theory, 'build' expresses *completed* events of building, and this fact is reflected in the lexicon. Data to support this claim are provided by sentences pairs such as the following:

(13a) Mary built the house.

(13b) Mary finished building the house.

(13a) seems to entail (13b). In order to explain this entailment, the theory assigns 'build' a TELIC semantic feature which ensures that the verb describes completed events. This feature opposes the ATELIC feature, which is part of the meanings of verbs that don't necessarily express completed events. However, consider the following:

(14) Mary was building the house.

Unlike (13a), (14) doesn't entail that Mary *finished* building the house. In other words, the TELIC feature of 'build' has not been carried over to the meaning of (14). The theory claims that feature cancellation is possible:[85] the progressive morpheme '-ing' removes this feature upon composition. Johnson concludes that a speaker could understand (14) without understanding 'build' (namely, that it has the TELIC feature), giving us a putative counter-example to Reverse Compositionality.

---

[84] Fodor, *Concepts: Where Cognitive Science Went Wrong*.
[85] As discussed in Chapter 1 with Hogeweg's theory of 'in'.

There are two types of response available to Fodor. First, he could deny that the speaker *actually* understands (14) in this case. Perhaps understanding a complex expression requires arriving at the correct semantic representation *in the right way*. The speaker has arrived at the correct semantic representation of (14), but has done so *via* defective means. The "correct" way to derive the semantic representation of (14) involves using feature cancellation to remove the TELIC feature from 'build'. Since the speaker didn't do this, they don't understand (14).

Second, Fodor could deny that the speaker misunderstands 'build'. Perhaps making the inference from (13a) to (13b) is not a necessary condition on understanding 'build'. If this is right, then the speaker could still understand 'build' even if they fail to make the appropriate inference.[86]

Regardless, Johnson's argument uncovers a tension between two background constraints: Reverse Compositionality and the domain constraint that requires the entailment phenomena to be explained (whatever the exact contours of that constraint might be). While Johnson doesn't argue that adhering to both constraints is *impossible*, he argues that the best existing theory of the phenomena violates Reverse Compositionality. For both Johnson and Fodor, this tension needs to be resolved by rejecting one of the constraints, and the locus of their debate is which one to collectively reject.

From the Pluralist perspective, we can see another possible resolution. While Johnson rejects the principle of Reverse Compositionality, Fodor often seems to think of it as a self-evident truth. He rejects Johnson's relatively wide domain constraint, claiming that the phenomena should be explained by syntax and/or pragmatics (if at all), not semantics. Johnson

[86] I'm highly confident that this would have been Fodor's preferred response. He was no fan of "meaning-constitutive inferences", *i.e.* inferences that one must draw on pain of not knowing the meaning of an expression (see Fodor, *Concepts: Where Cognitive Science Went Wrong*.)

criticizes the principle using data that the Fodorian has no interest in explaining with a semantic theory and a theory that the Fodorian would not accept (even if it's accepted by others). Perhaps we should view this as a point of divergence between two distinct projects, only one of which adheres to Reverse Compositionality. Johnson's project takes the domain constraint to be fundamental, focusing on how the phenomena is best explained. Fodor's project takes Reverse Compositionality to be fundamental, focusing on making explanations while adhering to the principle.

Chapter 3: Pluralism II - Soft Background Constraints

One way to be a pragmatic Pluralist denying (SM.4) and (SM.5) is to reject the Monist requirement of a *single* account. This is probably what comes to mind when thinking about "Pluralism". Another way to be a pragmatic Pluralist is to reject the Monist requirement of a *comprehensive* account. This view denies (SM.4) and (SM.5) by allowing the acceptance of theories and methods that are *known* to be incomprehensive.

In fact, these ways of being a Pluralist are generally complementary. Nancy Cartwright, one of the most prominent Scientific Pluralists, has a "dappled" view of science.[87] Scientific laws have limited scope in her view, applying to only part of the world. To cover the whole, scientists must employ a loose "patchwork" of different laws. (Cartwright's view is about the *nature* of laws; as a metaphysical view, it is stronger than the pragmatic view discussed here.)

The view developed in this chapter has two components that "soften" background constraints:

(1) Background constraints may be vague.

(2) Background constraints may be overridden.

In this chapter, I describe the components of the view in detail, discuss some general motivations, and show how the view is Pluralist by showing that each component allows theorists to accept incomprehensive theories and methods. In Chapters 4 and 5, I apply the view by examining Taxonomic Monism in detail, discussing how and why to interpret it as a Pluralist background constraint.

[87] Cartwright, *The Dappled World: A Study of the Boundaries of Science*.

1. Background Constraints may be Vague

It's tempting to view vagueness in a background constraint as a deficiency to be overcome. If a constraint contains terms that lack a precisely delimited extension, theorists will eventually encounter borderline cases that the constraint doesn't provide clear guidance on. I've previously discussed how the domain constraint "explain semantic anomaly" is vague: it's often unclear whether a particular anomaly is semantic as opposed to syntactic, pragmatic, or not really an anomaly at all. This vagueness could lead to uncertainty and disagreement, so shouldn't the constraint be precisified to eliminate borderline cases and draw sharp boundaries?

1.a. Vague Laws

Recently, philosophers and legal theorists have taken up this question as it concerns the law. They've asked whether vagueness in the formulation of a legal constraint can actually be valuable, and many have answered in the affirmative.[88] Soames argues that a vague law can rationally be preferred to a more precise one if most or all of the following conditions are met:[89]

1. The vague formulation [of the law] assigns the clear, non-borderline cases of the term the legal status desired by most lawmakers.

2. The variety of borderline cases is wide, making them hard to exhaustively anticipate.

3. The lawmaking body is either divided about the borderline cases or ignorant about the likely consequences of treating such cases one way rather than another, and so is uncertain about what legal status they should have.

4. The lawmakers recognize the value of incremental, case-by-case precisification of the law.

[88] Some have questioned whether it's actually vagueness *per se* that has value as opposed to some other closely related phenomenon that's being confused with vagueness. See Asgeirsson, "On the Instrumental Value of Vagueness in the Law."
[89] Soames, "Vagueness and the Law."

Consider the following law:[90]

Vehicle Ban: No vehicles in the park.

The law was implemented in order to promote safety and peacefulness. It clearly bans prototypical vehicles such as cars and motorcycles. Since these vehicles are loud, polluting, and dangerous, banning them promotes the law's general rationale. However, 'vehicle' is a vague term that has a wide variety of borderline cases: skateboards, strollers, drones, etc. It's difficult or impossible to formulate Vehicle Ban so that it exhaustively anticipates every potential vehicle. Some of the aforementioned borderline cases might not have been widely known or even invented when the law was created. Attempts to precisify might therefore be futile, and they might also be undesirable. As Silk puts it, "[f]uture adjudicators may have additional evidence about how certain classifications may promote or hinder the law's general rationale".[91]

The general rationale behind Vehicle Ban might be something like this:

a. Vehicles are loud, polluting, and dangerous.

b. Lawmakers want the park to be relaxing, clean, and safe.

c. These properties of vehicles conflict with the properties lawmakers want the park to have.

Suppose that Vehicle Ban was created when drone technology was in its infancy. Lawmakers are uncertain what relevant properties drones will have: size, loudness, propulsion mechanism, power source, autonomy, etc. They're also uncertain how drones will typically be used. Will they

[90] This example is quite famous in legal circles and can be traced back to a 1958 dispute between Hart, "Positivism and the Separation of Law and Morals"; Fuller, "Positivism and Fidelity to Law–A Reply to Professor Hart." The example is used somewhat differently in that dispute to talk about potential tensions between the letter of a law and its motivation.

[91] Silk, "Theories of Vagueness and Theories of Law." 137.

be used for mostly harmless recreation? Will they be used for commercial purposes by delivery companies with little regard for how safe and obnoxious they are?

As a result of this uncertainty, lawmakers don't know whether drones will have the undesirable properties in claim (a). Of course, this doesn't mean they're totally ignorant about drones. Perhaps they know some properties of drones that are *generally* relevant for classifying things as vehicles. However, the important fact is that lawmakers don't have the right information to reliably determine whether drones will conflict with the desirable properties in claim (b). They should wait until they acquire the right information before determining whether drones count as vehicles.

This example highlights one reason to allow laws to remain vague: ignorance about the entities lawmakers are trying to categorize. We can also look to the example for a second reason. Precisifying vague terms in the context of the law is generally aimed at promoting some general rationale. Up to this point, I've stipulated that the lawmakers know what they're aiming at: they agree upon a fairly simple and precise general rationale behind Vehicle Ban. However, this would likely be an idealization in most cases. Different lawmakers might hope to achieve different things with the law. Some might want the park to be more family-oriented while others want it to be a place for hobbyists. Their rationale(s) might be much more complex than a simple list of desirable properties; for instance, they are probably informed by background ideological considerations. And some lawmakers might yet be uncertain what exactly the rationale behind the law should be. After all, even the idealized rationale contains several vague terms itself.

Again, there's significant uncertainty about whether drones should count as vehicles. In this case, though, the uncertainty stems from the fact that the general rationale behind Vehicle

Ban is uncertain. It's unclear what target the lawmakers should be aiming at when they precisify 'vehicle'. Until the target comes into focus, it might be best for the lawmakers to abstain from classifying drones. The example highlights a second reason to allow laws to remain vague: ignorance about the general rationale behind the law.

This argument maps nicely onto MacFarlane's *plan expressivist* model of the contents of underspecified[92] sentences.[93] On this view, contents are sets of <world, hyperplan> pairs, where a hyperplan is a maximally specific plan for both language use and non-linguistic action (for instance, if "carry an umbrella if it's raining heavily" is a plan, then its corresponding hyperplans might specify which umbrella to carry and what counts as heavy rain). Uncertainty about vague sentences thus has *two* potential sources. There's ordinary doxastic uncertainty about which world we're in, but there's also uncertainty about what our plans are.

The model captures both reasons for allowing Vehicle Ban to remain vague with respect to drones. If lawmakers are uncertain about the present or future properties of drones (i.e. which world they're in), then asserting that drones aren't vehicles commits them to the claim that drones don't or won't have the undesirable vehicle-properties underlying the law. If drones end up having these undesirable properties, this is clearly bad for the park in a direct way, but there could also be indirect consequences. Lawmakers might accidentally find themselves committed to classifying other drone-like objects as non-vehicles even if doing so is inconsistent with the rationale behind the law. Conversely, if lawmakers are uncertain about the generale rationale behind the law (i.e.their plans), then asserting that drones aren't vehicles commits them to a plan

[92] In addition to vagueness, MacFarlane's view covers other types of linguistic underspecification like context-sensitivity.
[93] MacFarlane, "Indeterminacy as Indecision, Lecture III: Indeterminacy as Indecision." which builds on Barker, "The Dynamics of Vagueness." and Gibbard, *Thinking How to Live*.

for the park that is compatible with drones being allowed. This will cause trouble if they later settle on a plan that would be disrupted by allowing drones.

1.b. Vague Background Constraints

Background constraints are analogous to laws. Like laws, we can think of many background constraints as norms governing behavior. Background constraints govern data gathering, inference, and theory construction, among other theoretical activities. Just as there are conditions that make vagueness desirable in the law, there are also conditions that make vagueness a desirable feature of background constraints. Soames' argument can be straightforwardly adapted to provide such conditions:

1. The vague formulation of the background constraint assigns clear, non-borderline cases the status desired by most theorists.

2. The variety of borderline cases is wide, making them hard to exhaustively anticipate.

3. Theorists are either divided about the borderline cases or ignorant about the likely consequences of treating such cases one way rather than another, and so are uncertain about what status they should have.

4. Theorists recognize the value of incremental, case-by-case precisification of the constraint.

The fact that conditions (1) - (3) are frequently found in semantics should have already been made clear, but I'll give an example anyway. I've already strongly hinted that the domain constraint 'explain semantic anomaly' meets (1) - (3) because 'semantic anomaly' is vague. Presumably, there are clear, non-borderline cases of semantic anomaly that most theorists within the project are interested in explaining.[94] However, the boundaries of the phenomenon are not precisely delimited and agreed-upon, and this vagueness has multiple dimensions. There are

[94] For present purposes, we can disregard theorists outside of the project who don't adhere to the domain constraint (*e.g.* Fodor, who doubts the existence of *semantic* anomaly).

numerous cases where it's unclear whether the oddness of an expression is best classified as semantic anomaly, pragmatic infelicity, ungrammaticality, or merely some species of falsity. So conditions (1) - (3) are met.

Showing when condition (4) is met is much more involved, but the Vehicle Ban case indicates what to look for. One reason to allow a law to remain vague is "probabilistic" uncertainty about the state of the world (present or future). Likewise, theorists might allow 'semantic anomaly' to remain vague because they need more information about borderline cases in order to reliably categorize them. Another reason to allow a law to remain vague is "planning" uncertainty about what the law is supposed to accomplish. Likewise, theorists might allow 'semantic anomaly' to remain vague because the motivation for explaining semantic anomaly is still unclear. Assuming semantics is in the business of explaining linguistic anomalies, why is this the case? How and why do we distinguish *semantic* anomalies from other classes of linguistic anomaly such as ungrammaticality and necessary falsehood? Chapter 1 hopefully made it clear that planning uncertainty permeates semantics.

Allowing background constraints to remain vague is a Pluralist position. The argument for this claim doesn't require any particular view of vagueness. Suppose we have a model of vagueness under which some sentences containing vague expressions are neither true nor false. On this model, a project containing vague background constraints won't produce a comprehensive theory. For instance, a vague domain constraint remains agnostic on whether the project should explain some phenomena. A vague second-order constraint remains agnostic on whether some inferences should be made. As a result, some theoretical statements will be neither true nor false, at least from the perspective of the project. If we allow background constraints to

remain vague, we deny (SM.4) and (SM.5) by allowing incomprehensive methods and theories to be accepted.

Suppose we instead have an epistemic model of vagueness under which we are merely *ignorant* of the truth value of sentences containing vague expressions.[95] On this model, a project containing vague background constraints won't produce an *assertable* comprehensive theory. For instance, a vague domain constraint makes us ignorant of whether the project should explain some phenomena. A vague second-order constraint makes us ignorant of whether some inferences should be made. As a result, we will be ignorant of the truth value of some theoretical sentences. Since some true theoretical sentences won't be *determinately* true from our perspective, they won't be assertable.[96] Under this model, allowing background constraints to remain vague denies (SM.4) and (SM.5) by allowing us to accept methods and theories that aren't assertably comprehensive.

While allowing background constraints to remain vague is a Pluralist position according to both of these models of vagueness, each of them comes with a different *flavor* of Pluralism. On the truth-value gap model, Pluralism has a metaphysical flavor: incomprehensive methods and theories may be accepted because there is no comprehensive method or theory. On the epistemic model, Pluralism has a pragmatic or epistemic flavor: incomprehensive methods and theories may be accepted despite the possible existence of comprehensive methods and theories.

2. Background Constraints may be Overridden

Here's a pattern of argument that's familiar to all philosophers of language (and philosophers in general). An author puts forth a set of theoretical principles. They argue that

[95] Williamson, *Vagueness*.
[96] Weatherson, "Vagueness as Indeterminacy."

those principles either jointly make an incorrect prediction or are inconsistent in some way. This leads to the intermediate conclusion, often mentioned but not as often argued for, that at least one of the principles must be rejected. The author then forms an argument for which of the principles should be rejected.

2.a. A Puzzle about Belief

As an example, consider the discussion of names and attitude ascriptions from Kripke.[97] He proposes three *prima facie* plausible principles, two governing belief ascriptions and one governing translation between languages:

> **Weak Disquotation:** If a normal speaker sincerely assents to '*p*' upon reflection, then they believe that *p*.
>
> **Reverse Disquotation:** If a normal speaker believes that '*p*', then if they aren't reticent, they will sincerely assent to '*p*' upon reflection.
>
> **Translation:** If a sentence of one language expresses a truth in that language, then any translation of it into any other language also expresses a truth in that other language.

These principles can be viewed as background constraints on semantic theorizing. Both disquotation principles fit best into my taxonomy of background constraints as constraints on the characterization of the data, providing a bridge between observable assent (or non-assent) and the truth (or falsity) of belief ascriptions. Translation fits best as a second-order constraint, governing the relationship between semantic theories of different languages.

Though each of them are plausible, Kripke argues that they jointly lead to contradiction. He considers the case of Pierre, a monolingual French speaker living in France. Pierre has never been to London, but from what he's heard it seems like it would be pretty, so he assents to

[97] Kripke, "A Puzzle about Belief."

‘Londres est jolie’. By Weak Disquotation, we can derive the truth (in French) of ‘Pierre croit que Londres est jolie’, which allows us to derive the truth (in English) of ‘Pierre believes that London is pretty’ by Translation.

Pierre then moves to London. While living there, he gradually becomes a competent speaker of English. However, his neighbors don’t speak French, so he remains ignorant of many English-French translations including ‘London’-‘Londres’. He rarely ventures out of his neighborhood, which is unfortunately very ugly. Upon reflection, he declines to assent to ‘London is pretty’. By Reverse Disquotation, we can therefore derive the truth of ‘Pierre does not believe that London is pretty’.

From the three principles, we’ve derived the truth of both ‘Pierre believes that London is pretty’ and ‘Pierre does not believe that London is pretty’. It appears that we’ve derived a contradiction from the principles,[98] so shouldn’t one of them be rejected?

Surprisingly, Kripke doesn’t think so. He claims that such principles are independently well-motivated, making each of them individually plausible.[99] The puzzles motivating their rejection are (like the puzzle of Pierre) in “an area where our normal practices of interpretation and attribution of belief are subjected to the greatest possible strain, perhaps to the point of breakdown”. And it’s unclear which principle is responsible for the inconsistency. For these reasons, he thinks that “it would be foolish to draw any conclusion, positive or negative, about [the principles]”.[100]

---

[98] Though some have argued that appearances are deceiving (*e.g.* Richard, *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*.). Perhaps the sentences contain some indexical element that makes their contents non-contradictory.

[99] Kripke specifically argues against rejecting the Principle of Substitutivity, but I think the argument would extend to the principles I’ve described if they were at issue.

[100] *ibid.* 451.

2.b. Alternatives to Rejection

Kripke gives reasons to take his cautious approach but doesn't say much about what the alternative(s) look like. What are the options for dealing with conflict among background constraints without rejecting one of them? To frame this discussion, I'll consider an argument in *favor* of rejecting a background constraint in response to apparent conflict:

P1. Constraints $C_0 \ldots C_n$ conflict with each other.
P2. We should resolve the conflict.
P3. To resolve the conflict, reject one of $C_0 \ldots C_n$.
C1. Therefore, we should reject one of $C_0 \ldots C_n$.

Suppose we want to resolve the apparent conflict without rejecting any of the background constraints. Premises 1-3 correspond to three different ways of accomplishing this.

The first option claims that the appearance of conflict among the constraints is deceiving. In the Kripkean case, this means denying that the theory is actually inconsistent, rejecting Premise 1. The inconsistency immediately follows from a simple view of the logical forms of the two sentences: Believes(Pierre, *p*) and ~Believes(Pierre, *p*). However, many have argued that this view is *overly* simple.[101] Perhaps the sentences contain some ambiguity or indexicality such that their logical forms are actually something like Believes(Pierre, p_1) and ~Believes(Pierre, p_2). If so, the inconsistency, and therefore the conflict among the background constraints, has disappeared.

The second option simply lives with the consequences of the conflict and learns how best to deal with them. In the Kripkean case, this means denying that the inconsistency must be resolved, thereby denying Premise 2. Inconsistency is typically viewed as a cardinal sin of scientific theories, so it seems plausible that actions that *can* be taken to resolve inconsistencies

[101] Richard, *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*.

should be taken. However, some philosophers of science disagree. Martínez-Ordaz argues that inconsistencies can be tolerated in scientific theories without threatening the rationality of theorists who use those theories.[102] Roughly, the argument says that, in the absence of a better, consistent alternative, the best response can be maintaining the inconsistent theory while figuring out epistemic strategies to contain the effects of the inconsistency. Under classical logic (and other logics), inconsistency threatens to trivialize a set of beliefs *via* the principle of explosion. The specific challenge is therefore figuring out how to use the inconsistent theory to reason without descending into triviality.[103]

The third option, the one explored in the rest of this chapter, attempts to resolve the conflict without rejecting one of the constraints. It denies Premise 3. This option treats background constraints as if they may be overridden: when conflict arises among a set of constraints, we can choose at least one constraint to temporarily ignore. In the Kripkean case, this means treating Weak Disquotation, Reverse Disquotation, and Translation as potentially defeasible inference rules. (15) - (18) are the key sentences those inference rules act on in Kripke's example.

(15) Pierre assents to 'Londres est jolie'.

(16) Pierre assents to 'London is pretty'.

(17) Pierre croit que Londres est jolie.

(18) Pierre believes that London is pretty.

---

[102] Martínez-Ordaz, "The Ignorance behind Inconsistency Toleration."

[103] Notably, Wittgenstein made similar claims about inconsistency in mathematics throughout his body of work (especially Wittgenstein, *Remarks on the Foundations of Mathematics*.). Marconi, "Wittgenstein on Contradiction and the Philosophy of Paraconsistent Logic." 334-335, summarizes his position: "Wittgenstein's main point on this issue [*i.e.* what to do when a contradiction comes to light in mathematics] was that just throwing the calculus away is not the only "rectification" available, and that some other form of rectification is *always* on hand."

Weak Disquotation licenses the inference from (15) to (17), and Translation licenses the inference from (17) to (18). Reverse Disquotation licenses the inference from not-(16) to not-(18).

Since these inferences lead to contradictory conclusions, we must override at least one of them.[104] There are many potential ways to go about this. For instance, we might posit a preference relation among the principles that orders them by how strong we think they are. Personally, I would rank them as follows:

Weak Disquotation > Translation > Reverse Disquotation

We could then apply a *Weakest Link Principle* that says to prefer the argument with the strongest weakest link.[105] By my preference ordering, the weakest link in the argument from (15) to (18) is Translation and the weakest (and only) link in the argument from not-(16) to not-(18) is Reverse Disquotation. Since Translation is stronger than Reverse Disquotation, the argument from (15) to (18) has the strongest weakest link. Therefore, the Weakest Link Principle tells us to prefer the argument from (15) to (18), concluding that Pierre believes that London is pretty. We override Reverse Disquotation in this specific instance without thereby rejecting the constraint. We're still governed by it when it doesn't conflict with other constraints.

Allowing background constraints to be overridden is also a Pluralist position. As seen in the above example, if a project includes Reverse Disquotation as a defining background constraint and the constraint is overridable, then it might not apply in every situation. In other words, theorists can accept theories and methods even if they are known to be incomprehensive: a theorist can accept Reverse Disquotation while acknowledging that it doesn't govern inference

[104] Since we aren't taking the second option and living with inconsistency.
[105] Pollock, "A Theory of Defeasible Reasoning."

in every case. Therefore, if we allow background constraints to be overridden, we again deny (SM.4) and (SM.5) by allowing incomprehensive methods and theories to be accepted.

In Chapter 1, I claimed that background constraints are often implicit in theorizing rather than explicitly stated. This raises thorny questions about when a theorist can be said to adhere to a background constraint. Making background constraints soft raises similar questions. If a theorist claims to adhere to a background constraint but is also open to allowing exceptions, is their theorizing actually governed by the constraint?

In response, I'll note that such concerns are endemic to the study of all kinds of norms. In what sense are motorists governed by the speed limit when most of them commonly violate it without consequence (or even the expectation of consequence)? When can someone be said to act according to a Kantian imperfect duty such as "contribute to the happiness of others"?[106] If I generally use expression '*e*' in accordance with a linguistic convention to use '*e*' to refer to φ's but I sometimes use '*e*' in creative metaphors referring to non-φ's, how am I governed by the convention? Questions about what it is to adhere to a soft background constraint are instances of a broader concern about adherence to soft norms.

I'll also note the variety among the different background constraints I've discussed. They govern theoretical activity of different types and at different levels of abstraction. Some background constraints broadly orient semantic theorizing towards specific systems and phenomena within those systems. Others govern actions more narrowly by, for instance, telling semanticists what data to gather and how to gather it. Still others like Kripke's Disquotation and Translation principles govern inference, telling semanticists what follows from data and theory.

[106] Kant, *Immanuel Kant: Groundwork of the Metaphysics of Morals: A German–English Edition*.

This variety worsens the prospects for a *unified* account of what it is to adhere to a soft background constraint. It seems plausible that the account will have to incorporate a variety of theories of different kinds of norms, with its unity limited by the unity of these theories.

So I don't want to *dismiss* the question of what it is to adhere to a soft background constraint, but I do think that a satisfying answer is sufficiently complex to go beyond the scope of this project. Further, a rigorous account isn't required for my purposes. My central descriptive claim about the structure of semantics is that it contains multiple projects. To establish this claim, I only need to show that some semanticists adhere to different sets of background constraints than others. This can be done even if constraint-adherence remains uncertain in many cases.

Chapter 4: Taxonomic Pluralism I - Semantics States Meanings

Softening our view of background constraints has significant consequences for semantic theorizing. A close investigation of Taxonomic Monism gives us a good idea of what these consequences are like.

> **Taxonomic Monism:** A semantic theory should state the meaning of every expression in the language.

My investigation isolates two key components of Taxonomic Monism: the claim that semantics should *state* or *identify* meanings, and the claim that semantics should state meanings *of expressions*. For each component, I discuss how it is vague and leads to conflict with other constraints. I then apply my Pluralist view, demonstrating some consequences of allowing the constraint to remain vague and overridable.

1. Vagueness

The task of stating or identifying a meaning is much less straightforward than it sounds. Meanings are elusive. We can't gather them into a lineup and indicate which one is the meaning of a particular expression. Instead, theorists identify a meaning by referring to it in a metalanguage. The following examples were given in the introduction:

(a) ⟦Betsy⟧ = Betsy.

(b) ⟦cow⟧ = the function f: $\mathscr{D} \rightarrow$ {TRUE, FALSE} such that, for all $x \in \mathscr{D}$, f(x) = TRUE iff x is a cow.

Do (a) and (b) identify the meanings of 'Betsy' and 'cow'? In general, when do metalanguage expressions like (a) and (b) successfully refer to the intended meaning?

This question can present some surprisingly tricky difficulties. For illustration, consider a recent dispute over the relationship between semantics and truth conditions. Many theorists have

taken the determination of truth conditions to be a fundamental task of semantics,[107] perhaps even going so far as to *equate* sentential meaning with truth conditions. For such theorists, 'determining truth conditions' can be roughly synonymous with 'identifying (sentential) meanings'. Recently, however, the idea that semantics determines truth conditions has come under attack.

Consider the following examples from Pietroski:[108]

(19a) Unicycles have wheels.

(19b) Jim's unicycle has wheels.

(20a) Cars have wheels.

(20b) Jim's car has wheels.

According to Pietroski, both (19a) and (20a) are true. In contrast, (19b) is false (otherwise Jim's unicycle wouldn't be a unicycle), but (20b) is true as long as Jim's car is normal. Pietroski runs through a litany of options for specifying the semantics of 'has wheels' (and 'have wheels' if they are distinct predicates), rejecting all of them as either predicting incorrect truth values in (19a) - (20b) or being *ad hoc*.

The argument then generalizes from this example and others: tricky truth-value patterns like the one found in (19a) - (20b) are ubiquitous. As with (19a) - (20b), attempts to explain these patterns will likely be either empirically inadequate or *ad hoc*. Pietroski concludes that we should not attempt to construct semantic theories that determine truth conditions.

His discussion of the following attempt to specify the meaning of 'has wheels' is particularly relevant:

[107] With caveats to accommodate indexicality and other potential forms of context-sensitivity.
[108] Pietroski, "The Character of Natural Language Semantics."

(c) ⟦has wheels⟧ = λ*x*.true iff *x* has wheels.

Pietroski rejects this analysis because it doesn't say enough about *which* function is associated with 'has wheels':[109]

> But which function is [⟦has wheels⟧]? Is it a function whose extension includes only objects with more than one wheel; or is it a function whose extension includes some objects—like unicycles—with just one wheel? Merely surrounding the English predicate 'has wheels' with formal notation doesn't answer this question. Without an answer, we don't know what the theory says about the difficulty posed by [(19a) - (20b)].

According to this argument, (c) fails to identify a function from entities to truth values because it doesn't give us enough information. It doesn't answer the question of how many wheels an object must have, so at best it identifies a family of functions.

Pietroski is ultimately arguing that semantics *doesn't* determine truth conditions, so he rejects theories of ⟦has wheels⟧ that attempt to answer this question. King, on the other hand, argues in favor of the truth-condition view.[110] He provides a theory that he believes answers Pietroski's question (he also provides a theory of 'have wheels'):

(c.1) ⟦has wheels⟧ = λ*x*.true iff *x* has at least two wheels.

King claims that (c.1) provides the correct truth conditions in (19b), (20b), and other cases.

For Pietroski, (c) is an insufficient attempt to state the meaning of 'has wheels' because it doesn't tell us how many wheels an object must have. (c.1) ostensibly corrects this deficiency. However, it fails to answer further questions about which truth-value function is the meaning of 'has wheels'. It fails to be enlightening about what counts as a wheel (*e.g.* why is 'Frisbees are wheels.' false? Must wheels be roughly circular?). It fails to tell us in what sense something can *have* wheels (*e.g.* we're likely to deny that a normal shipping container filled with wheels has

[109] Pietroski. 228.
[110] King, "W(h)Ither Semantics!(?)."

wheels in the relevant sense; why is this the case?). Does this mean that (c.1) is also an insufficient attempt to state the meaning of 'has wheels'?

Indeed, Pietroski's discussion of data that *are* to be explained by semantics also leaves some questions open. It's unclear (at least in some cases) why his preferred data are within the purview of semantics but (19a) - (20b) isn't. Consider the following passage:[111]

> 'John is eager to please' means (roughly) that John is eager that he please someone and not that John is eager for someone to please him; 'John is easy to please' means (roughly) that it is easy for someone to please John and not that it is easy for John to please someone. Or more mundanely, 'Brutus stabbed Caesar' has no reading on which Brutus is the stabbee and Caesar is the stabber. Why?

What is the difference between these examples and (19a) - (20b) such that these examples *do* fall within the purview of semantics?

The immediate purpose of these examples is to show that meaning-statements exist along a gradient of descriptiveness. Meaning-statements that are mere tags, relying entirely on disquotation (*e.g.* ⟦has wheels⟧ = HAS WHEELS), are at the minimally descriptive extreme of the gradient. Meaning-statements found in Pustejovsky's project (discussed in Chapter 1, Section 1.a.ii) are generally further along the gradient, being much more descriptive than (c.1). This gradient is one dimension of vagueness in Taxonomic Monism: it's unclear how descriptive a lexical entry must be in order to count as identifying a meaning.

While it's difficult to definitively argue that Taxonomic Monism *should* remain vague at a given point, I can at least give reasons to allow it to remain vague. The current case shares important features with a variant of the Vehicle Ban case. When lawmakers precisify Vehicle Ban, they aim to promote some general rationale behind the law. For better or worse, the general

[111] Pietroski, "The Character of Natural Language Semantics." 18.

rationale is often unclear as well, which gives lawmakers reason to allow the law to remain vague while they determine their target.

Semanticists state an expression's meaning in order to (among other things) explain the semantic phenomena involving the expression. In other words, domain constraints form a significant portion of the general rationale for any view of how descriptive a meaning-statement must be. King's move from (c) to (c.1), adding more structure to the theory of ⟦has wheels⟧, is motivated by his preferred domain constraints: he believes that semantics should explain the truth-value data of (19a) - (20b).

However, domain constraints are themselves highly uncertain and contested. One encounters many open questions just by examining the case of 'has wheels'. Since it's unclear what the theory of ⟦has wheels⟧ should explain, it's unclear when a theory will be sufficiently descriptive. This gives semanticists a local reason to allow Taxonomic Monism to remain vague, leaving the question of what counts as a sufficient statement of ⟦has wheels⟧ unsettled unless they have a clear view of the relevant domain constraints.

2. Overriding the Constraint

Sometimes it's unclear whether a meaning has been successfully stated. Other times, it's clear that a meaning *hasn't* been stated. Consider Barwise and Cooper's classic work on quantifiers.[112] Barwise and Cooper distinguish between *logical* (*e.g.* 'every', 'some') and *non-logical* (*e.g.* 'most', 'many', 'few') quantifiers. They associate the logical quantifiers with entities that can plausibly be called meanings, defining the functions picked out by these quantifiers in every model. However, for the non-logical quantifiers they merely place

[112] Barwise and Cooper, "Generalized Quantifiers and Natural Language."

constraints on the meaning without the pretense of some stronger identification. For instance, they state the monotonicity properties of non-logical quantifiers. They also state whether the quantifiers are positive strong ('Most humans are human' is tautologous, so 'most' is positive strong), negative strong ('No humans are human' is contradictory, so 'no' is negative strong), or weak (the truth of 'A few humans are human' depends on whether the model contains a few humans, so 'few' is weak). In their Appendix B, these meaning-properties are called "semantic postulates". The semantic postulates for 'most', 'many', and 'few' place constraints on the interpretation of these expressions by restricting the class of possible models. Unlike the statements of the meanings of the logical quantifiers, which provide a stronger identification, these constraints are merely necessary conditions on the interpretation of the non-logical quantifiers. If Barwise and Cooper associate the logical quantifiers with meanings, then they associate the non-logical quantifiers with *types* of meanings.

Philosophers of science call this *abstraction.*[113] An abstraction omits information, typically because that information is irrelevant or unknown. For instance, when describing the motion of an object, we can typically abstract away from its color; we don't need to describe the object's color because it's irrelevant for calculations of velocity, displacement, etc. Barwise and Cooper's theories of the non-logical quantifiers are also abstractions because they only partially characterize the meanings of those quantifiers. Abstraction is one way that a semantic theory can *describe* an expression's meaning while failing to *state* a meaning. As demonstrated by Barwise and Cooper, it often is clear when a theory includes abstraction.

[113] Godfrey-Smith, "Abstractions, Idealizations, and Evolutionary Biology."

Idealization is another way that a semantic theory can fail to state a meaning. Idealizations say something false. For instance, we might simplify calculations in physics by claiming that the force due to friction is zero when it's known that the force is very small but non-zero. Idealizations can be much more difficult to spot, especially at points of contention in a theory, since spotting them requires a firm grasp on what is true. Nevertheless, I can at least provide tentative examples of idealization in semantics.

The meaning of a predicate is often characterized as a set containing all and only objects that the predicate is true of. This characterization struggles to accommodate the phenomenon of vagueness (considered as a linguistic phenomenon to be explained, not as I was considering it earlier). Suppose that *n* is a borderline instance of vague predicate *P*. If we're characterizing predicate-meanings as sets, then it's either definitely true or definitely false that *n* is *P*. But, at least under many views of vagueness, this isn't right; it's neither definitely true nor definitely false that the *n* is *P*. Adherents of such views can only maintain a "set model" of predicate-meanings as an idealization. Idealizations become more tolerable if some theoretical good (often simplicity) is gained by including a falsehood in the theory. Even if the "set model" is false of vague predicates, it can be a useful falsehood under the right conditions, especially when vagueness isn't the phenomenon of immediate interest.

2.a. Abstraction and Partiality

Because idealization is more difficult to identify, I will focus on abstraction in semantics. The idea of abstraction offers an alternative (though not necessarily incompatible) explanation of what Glanzberg calls *partiality* in semantic theories.[114] To illustrate the idea of partiality, consider

[114] Glanzberg, "Explanation and Partiality in Semantic Theory."

theories of the semantics of gradable adjectives ('tall', 'rich', 'wet', etc.).[115] This literature broadly characterizes the meanings of gradable adjectives as functions from entities to degrees on a scale. These degrees are ordered along a particular dimension: HEIGHT for 'tall', WEALTH for 'rich', etc.

Theories of gradable adjectives give the scales additional structure in order to explain various data. For instance, '#completely tall' is semantically anomalous. To explain this datum, it's posited that 'tall' is associated with an open scale: there is no maximum point on the scale for 'completely' to pick out. 'He is as tall as he is rich' is also sometimes viewed as anomalous, and to explain this datum, it's posited that 'tall' and 'rich' are associated with incommensurable scales that don't allow such a comparison.

However, the explanations provided by such theories are ultimately limited by the persistence of disquotation in their lexical entries. After adding the aforementioned scale structure, the theory of 'tall' still relies on disquotation (or near-disquotation) when it specifies the relevant dimension as HEIGHT. For Glanzberg, this reliance on disquotation exemplifies the partiality of semantic theorizing. Semantics does offer substantive explanations of phenomena, typically through the use of mathematics. However, it also frequently relies on disquotation to complete a lexical entry (especially for the non-logical parts of the language). Glanzberg says that, at the points where they rely on disquotation, semantic theories offer insubstantial explanations. This seems plausible. If we wish to explain why (21) is true, how substantial is an explanation saying that HEIGHT(LeBron) > HEIGHT(DeVito)?[116]

(21) LeBron is taller than DeVito.

---

[115] *e.g.* Kennedy, "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives."
[116] Better yet, if we wish to explain why 'LeBron is more athletic than DeVito.' is true, how substantial is an explanation saying that ATHLETICISM(LeBron) > ATHLETICISM(DeVito)?

The persistence of disquotation gives us reason to believe that disquotation is ineliminable. Therefore, we have reason to believe that semantic theorizing is ineliminably partial: it only partially explains some phenomena that we *prima facie* might have wanted a theory of meaning to explain.

From the partiality of semantic theories, Glanzberg concludes that the subject matter is itself partial. In the Chomskyan tradition, he assumes that the subject matter of semantics is a component of the language faculty. Therefore, he concludes that the language faculty is only part of the story of meaning. The mathematical structure found in semantic theories represents structure found in the language faculty, and persistent disquotation represents *pointers* from the language faculty to other conceptual systems. The openness of the scale associated with 'tall' is intended to represent structure within the actual lexical entry for 'tall' in the language faculty, while the disquotational specification of the dimension as HEIGHT represents a pointer from the lexical entry to other systems that are responsible for conceptualizing height. A fully satisfying explanation of the truth of (21) requires an appeal to more than just theories of language *per se*.

For Glanzberg, disquotation occurs when there are no more semantic explanations left to give. Disquotation can also occur as a type of abstraction when there *are* more semantic explanations left to give. In the literature on gradable adjectives, statements of lexical entries typically contain just enough non-disquotational information to explain the relevant data (*e.g.* anomaly). The disquotational parts of these entries are plausibly abstracting away from information that would be included if the theoretical focus were different.

2.b. Why Override the Constraint?

The Pluralist conceives of Taxonomic Monism and other such background constraints as potentially desirable but not necessary features of theorizing within a project. One way to override Taxonomic Monism is to choose to abstract away from an expression's meaning. The constraint may be overridden if the costs of adherence, in the form of conflict with other desirable constraints, outweigh the benefits.

The benefits of stating a meaning are highly situation-dependent. To illustrate this claim with a simple example, consider the following from Dever,[117] which is intended as an exemplar of semantic reasoning. Using (22) as an example, he inquires into the meaning (semantic value) of 'admires':

(22) Mary admires Susan.

Operating under the assumptions that the semantic value of a proper name is an object and the semantic value of a sentence is a function from possible worlds to truth values, he runs through the following reasoning:

1. ⟦Mary admires Susan⟧ is a function from worlds to truth values. That semantic value must be either (a) the result of applying ⟦Mary⟧ to ⟦admires Susan⟧, or (b) the result of applying ⟦admires Susan⟧ to ⟦Mary⟧.

2. ⟦Mary⟧ is an object, and hence not a function. It thus cannot be applied to anything, and option (a) is ruled out.

3. So we must have ⟦Mary admires Susan⟧ = ⟦admires Susan⟧(⟦Mary⟧).

4. Since ⟦Mary⟧ is an object and ⟦Mary admires Susan⟧ is a function from worlds to truth values, ⟦admires Susan⟧ must be a function from objects to functions from worlds to truth values.

5. The semantic value of 'admires Susan' must be either (a) the result of applying ⟦Susan⟧ to ⟦admires⟧, or (b) the result of applying ⟦admires⟧ to ⟦Susan⟧.

[117] Dever, "Formal Semantics."

6. ⟦Susan⟧ is an object, and hence not a function. It thus cannot be applied to anything, and option (a) is ruled out.

7. So we must have ⟦admires Susan⟧ = ⟦admires⟧(⟦Susan⟧).

8. Since ⟦Susan⟧ is an object and ⟦admires Susan⟧ is a function from objects to functions from worlds to truth values, ⟦admires⟧ must be a function from objects to functions from objects to functions from worlds to truth values.

Focus on the claim in (2): '⟦Mary⟧ is an object'. By convention, ⟦Mary⟧ is read as "the semantic value of 'Mary'".

Semantic value claims contain implicit quantification over occurrences of the expression. Many semantic value claims are language-universal: they say something about the semantic contribution an expression makes to each complex expression of the language in which it occurs. Assuming that '⟦Mary⟧ is an object' is making such a universal claim about the meaning of the expression 'Mary',[118] we can contrast three possible ways of interpreting the claim in decreasing order of strength. Let *o* be a variable ranging over occurrences of an expression and V(*s* , *o*) be the semantic contribution that expression *s* makes to *o*:

(i) $\forall o$(V('Mary', *o*) = Mary & object(Mary)).

(ii) $\exists x \forall o$(V('Mary', *o*) = *x* & object(*x*)).

(iii) $\forall o \exists x$(V('Mary', *o*) = *x* & object(*x*)).

The first of these is the strict Taxonomic Monist interpretation: it takes the logical form of an identification of the meaning of 'Mary'. The second and third interpretations take the quantifiers that appear throughout the above reasoning more seriously. They quantify over semantic value candidates rather than directly referring to a particular one. The second interpretation is weaker

[118] Here I'm ignoring the fact that there are multiple Marys and that 'Mary' is therefore plausibly ambiguous.

than the first because it merely places a constraint on the semantic value of ‘Mary’ without identifying it. By giving the existential quantifier wide scope it claims that there is some object that ‘Mary’ contributes in all of its occurrences; in other words, it claims that ‘Mary’ is univocal without identifying that singular meaning. The third interpretation is weaker still since it gives the existential narrow scope. Therefore, it doesn’t make the univocality claim, leaving open the possibility that ‘Mary’ contributes different objects to different occurrences.

Notice, though, that the strength of (i) and (ii) is not required for the above reasoning to go through. We don’t need the object contributed by ‘Mary’ to specifically be Mary, and we don’t even need ‘Mary’ to contribute the same object in all occurrences. All we need is for ‘Mary’ to contribute some object or other to the meaning of ‘Mary admires Susan’. For the purposes of the above reasoning, there are no additional benefits provided by the stronger claims, including the strongest (i) that adheres to Taxonomic Monism.

In addition to situationally providing minimal benefits, stating meanings can be theoretically costly by conflicting with other constraints. I’ll illustrate by returning to the case of gradable adjectives. The theory I’ve been considering includes dimensions in its lexical entries, which are functions from objects to degrees along a scale. I’ve stipulated that ‘rich’ is associated with the WEALTH dimension. For a non-comparative predication such as ‘Alice is rich’, this dimension assigns Alice a degree of wealth. The semantics then compares this degree to a degree that sets the (probably contextually determined) standard for being rich: the utterance is true if and only if Alice’s degree of wealth is at least as great as the prevailing standard. For a comparative sentence like ‘Alice is richer than John’, Alice’s degree of wealth is compared to John’s.

The claim that the meaning of 'rich' includes the WEALTH dimension encounters some difficulties. Consider the following cases:

> Measurement Change
> How wealth is measured is continually changing through changes in accounting practices, the tax code, which assets and liabilities count for and against wealth, etc. Suppose that, until t_1, wealth is measured one way. At t_1, accounting practices change and wealth starts to be measured in a new way. Which of these two measures is WEALTH?
>
> Comparison
> 'Rockefeller was richer than Bezos': is this true or false? If we compare real (inflation-adjusted) wealth, then it's true. If we compare nominal wealth, then it's false. Is WEALTH real or nominal?
>
> Mars Colony
> Elon Musk founds a Mars colony. When he arrives, he discovers that Mars is inhabited by a previously-undiscovered subterranean intelligent species. He begins to live among them. Even though he's extremely rich on Earth, in that society, his assets are mostly considered worthless. Is Elon rich?

These examples cast serious doubt on the claim that the meaning of 'rich' includes a single WEALTH dimension. Even if no other facts about the individual change, there are indefinitely many ways to assign them a degree of wealth and a wide variety of factors that can influence this assignment. Some of these factors, such as time and location, are at least familiar from the literature on context sensitivity. Others, such as whether to take inflation into account or what accounting practices to use, are highly specific to the economic domain.

Two options for dealing with this issue come readily to mind. First, we could claim that 'rich' is ambiguous: the lexicon contains multiple entries for 'rich', each with a different wealth dimension. Second, we could deny that lexical entries of expressions such as 'rich' contain *just*

dimensions. Instead, perhaps they contain *functions* from contexts to dimensions.[119] Both of these options incur a cost by decreasing the parsimony of the theory.

Philosophers of science distinguish between *quantitative* and *qualitative* parsimony.[120][121] Roughly, quantitative parsimony decreases as the theory posits more entities, while qualitative parsimony decreases as the theory posits more types or properties. What does it mean for a theory to "posit" an entity or type? A first attempt is that a theory posits an entity or type just in case the theory (perhaps in conjunction with some background theories) entails the existence of the entity or type.

If this attempt works, then it seems that the above options for dealing with context-dependence wouldn't decrease the parsimony of the theory. Prior to taking either option, the theory already includes objects and degrees. In conjunction with the appropriate mathematical machinery, the theory already entails the existence of each of the numerous wealth dimensions before the first option potentially includes them in the lexicon. Therefore, the first option wouldn't decrease the quantitative parsimony of the theory by positing new entities. The theory also already includes dimensions and, presumably, contexts. Similarly, the theory plausibly already entails the existence of functions from contexts to dimensions before the second option includes them in lexical entries. Therefore, the second option wouldn't decrease the qualitative parsimony of the theory by positing a new type.

However, it seems that this first analysis of what it is for a theory to posit an entity or type leaves something out. Consider Fregean systems of semantic types, which attempt to provide a recursively-defined taxonomy of possible types of meanings. Since they rely on

[119] Pinkal, "On the Logical Structure of Comparatives."
[120] Baker, "Quantitative Parsimony and Explanatory Power."
[121] Sendłak, "On Quantitative and Qualitative Parsimony."

recursion, they stipulate some base semantic types and rules for constructing the compound types from the base cases. Heim and Kratzer provide the following initial definition of semantic types:[122]

(a) *e* and *t* are semantic types.

(b) If A and B are semantic types, then <A, B> is a semantic type.

(c) Nothing else is a semantic type.

This definition entails the existence of infinitely many semantic types. We can construct and (with some difficulty) understand artificial expressions with highly complex types.[123] However, due at least to performance issues, only a relatively small handful of the simpler types (*e.g.* the familiar <*e*, *t*> and <*e*, <*e*, *t*>>) will be found in natural language.

Suppose that two semantic theories, T_1 and T_2, explain the same phenomena and use the same recursive definition of semantic types. T_1 therefore entails the existence of just as many types as T_2 does. However, suppose that T_1 uses fewer semantic types in its explanations. In other words, according to theory T_1, the language instantiates fewer types than it does according to theory T_2. T_1 is *ceteris paribus* simpler than T_2, but since they entail the existence of the same amount of types, the initial analysis of "posits" doesn't give this result. I won't try to capture what exactly the analysis leaves out, but the quantitative and qualitative parsimony of a theory seem to depend on more than just a count of the entities and types whose existence it entails. They depend on a count of the entities and types that are included in the theory in some stronger sense. This could mean a count of *causally efficacious* entities or *instantiated* types.

[122] Heim and Kratzer, *Semantics in Generative Grammar*.
[123] Pietroski, "Semantic Typology and Composition."

Taking a more nuanced view of quantitative and qualitative parsimony, both options for dealing with the context-dependence of gradable adjectives incur a cost of greater complexity. The first "ambiguity" option decreases quantitative parsimony by including more lexical entries in the lexicon. It expands the size of the lexicon rapidly and indefinitely. Since there are many different ways of measuring wealth, there are many different wealth dimensions and therefore many different lexical entries for 'rich'. The second "context-sensitive" option decreases qualitative parsimony by including functions from contexts to dimensions as a new type of meaning in the theory.

Up until now, the goodness of quantitative and qualitative parsimony has been accepted as a brute fact. Since brute facts are often unsatisfying, it would be nice to have reasons to think that the two types of parsimony are indeed theoretical goods. Reasons to think that parsimony is a theoretical good come in two flavors: general reasons that apply across many or all scientific domains, and specific reasons that rely on idiosyncrasies of a particular domain. Semantics (and probably the study of language more generally) provides at least one domain-specific reason to think that parsimony is a theoretical good.

A language is often thought of as something that its competent speakers "know" in some sense. To the extent that this knowledge isn't innate, it must therefore be learned. Learnability is an important motivation for some of the most fundamental features of typical semantic theories; for instance, Davidson proposes that a semantic theory must contain a finite number of basic meanings and rules in order to count as part of a learnable language.[124] A finite base of the language can be seen as a necessary condition for it to be learnable by finite beings. Of course,

[124] Davidson, "Theories of Meaning and Learnable Languages."

due to the limits imposed by our cognitive ability and the time we have available, most languages that are learnable in principle will not be practically learnable by us.

Thus, in addition to learnability in principle, *ease* of learnability is also a virtue for a language. Both quantitative and qualitative parsimony make a language easier to learn. If a language is quantitatively parsimonious by containing a smaller lexicon, it contains fewer lexical entries that a speaker must learn. And if a language is qualitatively parsimonious by containing fewer types, this restricts the hypothesis space when learning the meanings of particular expressions.

Thus, stating the meaning of 'rich' runs the risk of increasing the complexity of the theory, and increased complexity is indeed to the theory's detriment. And as before, stating a meaning can have minimal benefits. Consider the phenomenon of cross-scalar anomaly, exemplified below.[125]

(23) #Susan is taller than John is rich.

Dimensions are used to explain this anomaly: 'tall' and 'rich' contribute different dimensions to the interpretation of (23), and the resulting degrees (of Susan's height and John's wealth) are incommensurate.

Again, a very weak claim can be used to explain the phenomenon. We don't need to say *which* dimensions 'tall' and 'rich' contribute to the interpretation of (23), just that the dimensions assign incommensurate degrees to Susan and John. We don't even need to say *how* the expressions contribute the dimensions that they do. The expressions could univocally contribute the same dimensions in every occurrence; they could be ambiguous, with one of their meanings

[125] Kennedy, "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives."

selected to compose with the meanings of the other expressions; they could be sensitive to a context that determines the operative dimensions.

2.c. Application: Predicativism

There are good reasons to allow semantics to abstract away from the meaning of an expression in some situations, thereby overriding Taxonomic Monism. Further, this abstraction might be persistent. The complexity of the underlying target system is vast, especially if one's target is a language module or system of conventions. Abstraction is ineliminable if the complexity can never practically be captured. Additionally, parsimony considerations can make a theory with abstraction preferable to a more complex theory that doesn't involve abstraction.

Allowing Taxonomic Monism to be overridden in this way has consequences for semantic argumentation and theory choice. Consider recent disputes over the Predicativist theory of proper names. The orthodox view of names is Referentialism, which says that a proper name is a singular term whose semantic contribution to sentential meaning is just its referent (this is the view that was assumed in the above discussion of 'Mary admires Susan.'). Predicativism, which says that a proper name is a predicate, has recently gained popularity. The Predicativist view is motivated by so-called "predicative" uses of names, among other phenomena:[126]

(24) There are relatively few Alfreds at Princeton.

The felicity of such utterances is difficult to accommodate under the view that names are uniformly referential. One possibility is to give up the claim that names are *uniformly* referential and instead hold that they are (semantically) ambiguous. This view, which I call Ambiguity Referentialism, claims that names are sometimes referential and sometimes predicates. Another

[126] Burge, "Reference and Proper Names."

possibility is to maintain the claim that they are uniformly referential, adding type-shifting mechanisms to either the semantics or pragmatics that shift referential names to something that can compose as a predicate. Both of these popular options add complexity to the theory, though for now I will focus on the ambiguity option. A primary argument used in favor of Predicativism is that it is simpler (or perhaps more "uniform") than Referentialism: all names are predicates, and no type-shifting mechanisms are required for (24).

Many Predicativists have taken the further step of trying to identify the predicate that is supposedly expressed by names. Graff Fara defends the Being-Called Condition (BCC):[127]

> **Being-Called Condition:** A proper name '*N*' is a predicate that is true of a thing just in case it is called *N*.

Under this view, 'Alfred' is a predicate that is true of a thing just in case it is called Alfred. This means that (24) is true just in case there are relatively few things-called-Alfred at Princeton, which seems intuitively plausible.

Predicativist arguments are not without their critics. Jeshion thinks that the simplicity or uniformity considerations advanced in favor of Predicativism are mistaken.[128] Reconstructing the Predicativist argument, she claims that it depends on Uniformity Principle 1 (henceforth UP1):

> **Uniformity Principle 1:** Other things equal, a theory that explains the semantics of singular unmodified occurrences of names in the same way that it explains the semantics of those (and just those) pluralized, quantified, modified and otherwise determiner-fronted occurrences of proper names true of those that have the name is superior to a theory that does not.

Most of Jeshion's criticism of the Predicativist motivation targets the idea that the BCC provides a unified explanation of the semantics of names. As one example, suppose that Joe Romanov is

[127] Fara, "Names Are Predicates."
[128] Jeshion, "Names Not Predicates."

not a member of the Romanov dynasty. The following seems intuitively true (or at least has a reading on which it is true):

(25) Joe Romanov is not a Romanov.

However, the BCC would predict that (25) is false since Joe Romanov is *called* 'Romanov'. If (25) and similar sentences are indeed counterexamples to the BCC, then it doesn't provide a unified explanation of the semantics of names. Therefore, we can't use the BCC to claim that Predicativism is superior to Ambiguity Referentialism *via* UP1.

The Pluralism advanced in this paper doesn't take the side of either the Predicativist or the Referentialist. However, it does view UP1 and principles like it with suspicion. Most obviously, a phrase like "explains the semantics" doesn't sit well in a Pluralist setting. It seems to presuppose that the various theories of names under consideration are all competing to be part of the same single unified theory. Note, though, that Jeshion provides data that are completely different in kind from the data that initially motivated Predicativism. The counterexamples to the BCC are truth-value data, and Predicativism was motivated by felicity phenomena. A Predicativist like Pietroski who doesn't include truth-value within their domain constraints wouldn't be particularly bothered by Jeshion's data.[129]

There are other potential concerns with UP1. Both sides treat it as a principle of *quantitative* uniformity, similar to principles of quantitative parsimony discussed previously. For instance, Jeshion characterizes "the main argument originally advanced in favor of predicativism — that proper names, in both apparently referential uses and predicative uses, have the *same meaning.*"[130] Interpreted this way, UP1 tells us to prefer a theory that posits one meaning to

[129] Pietroski, "Semantic Typology and Composition.". In the same vein, they wouldn't be particularly concerned with developing an account of the truth conditions of name-predicates *a la* the BCC.
[130] *ibid.* 245, emphasis hers.

explain both types of occurrences of a name over a theory that posits two meanings to do the explanatory work.

As seen in the discussion of names in the previous section, semantic theories may abstract away from the number of meanings an expression has. The weakest interpretation of the claim '⟦Mary⟧ is an object' is agnostic on whether ⟦Mary⟧ contributes the same object in every occurrence, but the interpretation is nonetheless sufficient for many purposes. By allowing such theories, my Pluralist view creates a need for parsimony and uniformity considerations that differ in kind from those provided by UP1. When comparing theories that abstract away from meanings, we might be unable to compare the numbers of meanings they posit, and thus unable to apply quantitative principles. Instead, the parsimony or uniformity considerations will be qualitative, instructing us to prefer fewer *types* of meanings.

Jeshion acknowledges the possibility of qualitative uniformity arguments based on semantic types,[131] but her criticism remains focused on quantitative uniformity arguments. Relying on qualitative rather than quantitative uniformity, we can reconstruct a Predicativist argument in a way that isn't undermined by Jeshion's counterexamples. Predicativism is motivated by specific phenomena: for instance, bare uses of names are felicitous ('Alfred is English.'), but so are uses with an overt determiner ('Every Alfred is English.', 'There are three Alfreds at Princeton.', etc.). In order to accommodate the uses with an overt determiner, Predicativists claim that names have the semantic type of predicates (say, $\langle e , t \rangle$), which allows them to easily compose with determiners. In order to accommodate the bare uses, they claim that

[131] *ibid.* 245.

bare occurrences of names are accompanied by a covert determiner that allows them to successfully compose.

The Ambiguity Referentialist doesn't associate a name with a single semantic type. Instead, they say that a name is ambiguous between a referential meaning (*e.g.* of type *e*) and a predicative meaning. In other words, whereas the Predicativist explains the phenomena using one type of entity (predicative names), the Ambiguity Referentialist explanation uses two (predicative *and* referential names). Therefore, Predicativism is *ceteris paribus* more uniform than Ambiguity Referentialism. Point to the Predicativist.[132]

While I won't explicitly formulate a qualitative counterpart to UP1, there are two important features that would differentiate such a principle. First, its uniformity considerations would be based on the number of types of entities rather than the number of entities. Second, it would promote qualitative uniformity in explanations of specific phenomena, not "the semantics". Predicativism's ability to specifically explain a variety of felicity phenomena using only one semantic type is a point in its favor. And this remains a point in the theory's favor even if it's unable to provide uniform explanations of other phenomena.[133]

Jeshion's data are immediately intended as counterexamples to the BCC. However, one can imagine them having a broader purpose. They could be used in roughly the same way that I used the examples of 'rich' in the previous section. The 'rich' examples were intended to

[132] Arguments in favor of Predicativism against other types of Referentialism will have to be different. For instance, against the type-shifting Referentialist the point in favor of Predicativism is that it does not have to posit a type-shifting mechanism that allows referential names to be used predicatively. Here, the Referentialist is likely to respond that they don't have to posit a covert determiner in the syntax of referential uses of names. Since my purpose here is primarily illustration of an argument strategy and not a defense of Predicativism *per se*, I will not wade into these murkier waters.

[133] Here is a point of strong agreement with Jeshion. On my view, qualitative uniformity considerations merely provide points in favor and not decisive arguments. Jeshion expresses the same attitude when claiming that "[w]e should not regard uniformity as the sole or decisive factor in semantic theory selection." *ibid*. 237.

highlight the wide variety of ways of measuring wealth, thereby casting doubt on the possibility of specifying a single WEALTH dimension. Similarly, Jeshion's data could be used to highlight the wide variety of ways in which names can be used predicatively, thereby casting doubt on the possibility of a unified analysis of name-predicates.

However, as noted earlier, Jeshion provides data that are different in kind from the data that initially motivated Predicativism. To explain the felicity phenomena, the Predicativist posits that names are predicates so that they may (among other things) successfully compose with quantifiers. To explain these phenomena in a Predicativist way, the theory merely needs to state that a name contributes some predicate or other to semantic composition. They don't need to say which predicate is picked out by a name, and they don't even need to suppose that a name is univocally associated with only one predicate. Predicativism can still successfully explain the felicity phenomena even if it fails to provide a unified analysis of the name-predicate.[134] For Predicativism to lose its qualitative uniformity advantage over Ambiguity Referentialism, one must show that Predicativism can't explain the felicity phenomena with just a single type of entity.

[134] I'm not claiming that it *does*, just that it can in principle.

Chapter 5: Taxonomic Pluralism II - Meanings of Expressions

I've argued that semantics is allowed to override Taxonomic Monism by failing to state the meaning of an expression. Semantics characterizes the meanings of expressions, but these characterizations may fall short of identifying the meaning. I now turn my focus to a different part of Taxonomic Monism: the claim that semantics should state the meaning of each *expression* of the language. As in the previous section, I start by discussing the vagueness in the term 'expression'. I then argue that this is another point where Taxonomic Monism should be considered soft.

1. Vagueness

Taxonomic Monism obligates semanticists to state the meaning of each expression of the language. Natural languages contain infinitely or at least indefinitely many expressions. Providing a list of every expression and their meanings might satisfy Taxonomic Monism, but this is at least practically impossible. Instead, it's commonly assumed that there is a base set of atomic expressions and some rules for generating compound expressions from the atoms.

Determining the membership of this base set requires (among many other things) a distinction between *ambiguity* and *polysemy*. 'Bank' is often used as an exemplar of ambiguity; it has both a RIVERSIDE meaning and a FINANCIAL INSTITUTION meaning. Ambiguity is typically modeled by adding more than one expression to the lexicon for a given word; thus, in the case of 'bank', the lexicon would actually contain at least two expressions '$bank_1$' and '$bank_2$' corresponding to the different meanings.

The following pair of sentences illustrates polysemy:[135]

[135] Recanati, "Compositionality, Flexibility, and Context-Dependence."

(26a) The policeman stopped the car.

(26b) The driver stopped the car.

(26a) is most naturally interpreted as the policeman signaling for the car to stop, and (26b) is most naturally interpreted as the driver applying the brakes in order to stop the car.[136] So, 'stop' has multiple senses.[137] If (26) is an example of polysemy instead of ambiguity, then the lexicon will not include distinct expressions corresponding to the different senses. Instead, there will be a single expression whose meaning composes to generate the meanings of both (26a) and (26b). The task of explaining the different senses is then typically passed over to pragmatics.

Unfortunately for attempts to precisely delimit a set of expressions of the language, the distinction between ambiguity and polysemy is a vague one. This can even be seen with 'bank', the exemplar of ambiguity. There are both commercial banks and investment banks (sometimes operated by the same corporation); is this distinction sufficient to warrant distinct expressions corresponding to COMMERCIAL BANK and INVESTMENT BANK? 'Bank' refers to both bank buildings and corporations; is this distinction sufficient to warrant distinct expressions corresponding to BANK BUILDING and BANK CORPORATION? Such examples are easy to find.

There have been several proposed methods for determining whether something is an example of ambiguity or polysemy. For instance, Recanati claims that two senses fall under the same linguistic meaning if they "are felt as related".[138] In some of the above examples the *Felt Relatedness* test works well. The two senses of 'stop' in (26) are clearly related, as are COMMERCIAL BANK and INVESTMENT BANK, so the test classifies them as cases of polysemy.

[136] Both sentences can also be coerced into the other reading if we modify the imagined context in which they are uttered.
[137] Perhaps a very large number of senses, since examples of different ways of stopping are easy to find.
[138] Recanati, "Contextualism and Polysemy." 384.

In other cases, the verdict of the test is either less clear or conflicts with other tests.[139] 'Dissertation' has at least two senses, an abstract sense for the information in the dissertation and a concrete sense for physical copies of the dissertation. Clearly these two senses are related, as physical copies of the dissertation contain its information, so Recanati's test would classify this as polysemy. However, suppose we apply the *Conjunction-Reduction* test to 'dissertation'.[140] [141] This test takes two sentences containing the potentially ambiguous word, conjoins them, and "reduces" them by dropping one of the word's occurrences. (27c) seems strange:

(27a) Judy's dissertation is thought-provoking.

(27b) Judy's dissertation is yellowed with age.

(27c) #Judy's dissertation is thought-provoking and yellowed with age.

Since the sentence that results from attempting to conjoin the two senses of 'dissertation' is zeugmatic, the Conjunction-Reduction test says that the expression is ambiguous. The conjunction-reduction procedure unsuccessfully attempts to collapse two distinct expressions into one.

The conflicting evidence poses a taxonomic issue. (27a) and (27b) contain two distinct occurrences of 'dissertation'. 'Dissertation$_i$' in (27a) has an abstract informational sense and 'dissertation$_p$' in (27b) has a physical sense. The Felt-Relatedness and Conjunction-Reduction tests are supposed to provide evidence for the classification of 'dissertation$_i$' and 'dissertation$_p$': are these distinct occurrences instances of the same expression? Since the evidence points in different directions, the answer to this question is unclear. Further, it's unclear how we might

[139] Of course, there are other methods that assist in drawing the ambiguity/polysemy distinction, but my claims also apply to many, if not all, of those.
[140] Sennet, "An Ambiguity Test for Definite Descriptions."
[141] Zwicky and Sadock, "Ambiguity Tests and How to Fail Them."

resolve the issue. Does one test take priority over the other? If so, how should we determine which test has priority?

The ambiguity / polysemy distinction is one main way in which 'expression' is vague. Many sentences that take the form of (28) aren't assertable, since their truth value is unknown.

(28) 'e_1' and 'e_2' are instances of (the same expression / different expressions).

It's unclear what counts as a genuine expression, so the boundaries of the lexicon are vague. It's also often unclear which expression an occurrence is an instance of, so the boundaries of many expressions are vague.

Unsurprisingly, similar taxonomic issues frequently arise in biology. The species category has massive importance for the field, but it is also highly contested. Dozens of concepts have been proposed as *the* species concept.[142] Biologists have grouped organisms into species based on reproductive affinity, ecological niche, common ancestry, genotype, and other biological properties. Further, many of these concepts are incompatible, entailing different species boundaries. As a result, it's sometimes unclear whether distinct organisms are instances of the same species.[143] For instance, two organisms might have distinct ecological niches yet be able to reproduce.

'Species' is vague in the same way as 'expression'. There are multiple, often incompatible, tests for species-hood. As a result, it's often unclear what counts as a genuine species or where the boundaries of a species are found.

[142] Mayden, "A Hierarchy of Species Concepts: The Denouement in the Saga of the Species Problem." lists 24 different species concepts, many of which can be further precisified in various ways.
[143] de Queiroz, "Species Concepts and Species Delimitation."

de Queiroz identifies a common thread running through the disparate species concepts.[144] He synthesizes them into a general species concept: a *species* is a segment of a separately evolving metapopulation lineage, where a *metapopulation* is a broader population composed of connected subpopulations. For de Queiroz, this is the *only* necessary condition on species-hood.

His definition of 'species' is therefore vague, most importantly with respect to identifying "separately evolving" lineages. The previous approach argued over whether reproductive affinity, ecological niche, etc. were necessary properties of species-hood. de Queiroz's preferred approach views these properties merely as evidence that lineages are separately evolving. No one condition is decisive. If every condition points in the same direction, we can confidently say whether two lineages are separately evolving and therefore distinct species. However, if some conditions point in different directions, the lineages are in a gray area where it isn't clear whether they're separately evolving.

de Queiroz argues that the vagueness of the general species definition is beneficial. The biological properties that were previously used to *define* 'species' can now be used to *precisify* it. Different precisifications are useful for different purposes:[145]

> Subcategories of the species category are important in that they are composed of those species that are relevant to addressing particular biological questions. For example, a study of reinforcement requires species that exhibit postmating reproductive incompatibilities, whereas a study that uses a species-level phylogeny to make inferences about historical biogeography might be better served using species that exhibit monophyly.

In other words, a vague definition of 'species' is beneficial because the term can be flexibly adapted to different purposes and new information. This type of argument is familiar from Chapter 3.

---

[144] *ibid.*
[145] *ibid.* 882 - 883.

'Expression' plays a similar theoretical role to 'species', and it causes similar problems.[146] Perhaps, then, 'expression' should remain vague in similar ways to de Queiroz's proposal for 'species'. Different precisifications of 'expression' don't necessarily conflict with each other; instead, they might be appropriate for different projects and theoretical goals.

Tests like Felt-Relatedness and Conjunction-Reduction give us ways to precisify 'expression': since they deal with ambiguity phenomena, they place conditions on expression identity. As discussed, the results of these tests often conflict (or at least appear to), so they give us different precisifications of 'expression'. If the tests are more appropriate for different projects and theoretical goals, however, the apparent conflict might not require resolution.

In their seminal article on ambiguity tests, Zwicky and Sadock argue against the absolute applicability of ambiguity tests.[147] They repeatedly emphasize that "grammatical tests in general, and ambiguity tests in particular, reflect the theory within which they are framed".[148] Consider the example of Transformation tests. Sentence (29a) has two understandings:

(29a) They saw her duck.

Transformation tests check whether all the understandings survive under various transformations.

(29b) Her duck was seen by them.

The passive (29b) only has the possessive reading, which *can* be taken to mean that (29a) is an instance of ambiguity. However, Zwicky and Sadock show in detail that inferences from the results of Transformation tests are theory-laden, depending on specific claims about the existence

[146] Further parallels can be drawn, especially in light of the frequent application of evolutionary concepts to the study of language change (*e.g.* Croft, "Evolutionary Linguistics.")
[147] Zwicky and Sadock, "Ambiguity Tests and How to Fail Them."
[148] *ibid.* 3.

of various syntactic transformations as well as more general claims about the nature of transformations.

The applicability of various ambiguity tests can also depend on theoretical goals. Recanati, in his promotion of the Felt-Relatedness test, is mostly concerned with modeling how polysemous expressions are stored in memory. If two distinct senses are felt as related (*e.g.* the abstract informational and physical senses of 'dissertation'), this is evidence that they are stored together in memory; in other words, that they somehow arise from the same lexical entry.

In contrast, Zwicky and Sadock are mostly concerned with sentence-level semantic representations when they promote the Conjunction-Reduction test. The test is intended to provide evidence about whether sentences like (30) have multiple possible semantic representations:

(30) Sam thought about Judy's dissertation.

If the test answers in the affirmative, we have evidence that (30) has distinct semantic representations corresponding to (30a) and (30b), where 'd_i' denotes Judy's dissertation as an abstract informational object and 'd_p' denotes a physical copy.

(30a) thought-about(Sam, d_i)

(30b) thought-about(Sam, d_p)

The results of the Felt-Relatedness and Conjunction-Reduction tests conflict with each other (or at least appear to do so), so they provide different ways to precisify 'expression'. However, the tests are useful for related but ultimately distinct theoretical purposes. If 'expression' remains vague, it has the flexibility to be adapted for these purposes.

So far I've discussed Pluralist views of theoretical terms like 'expression' and 'species' as an alternative to Monist views, but an objection highlights a third option. An Eliminativist view of such terms argues that they should be eliminated from theoretical discourse and replaced with multiple terms, each corresponding to a different precisification of the original. Ereshefsky argues for 'species'-Eliminativism:[149]

> The term 'species' is superfluous beyond the reference to a segmentation criterion; and when the term is used alone it leads to confusion. The term 'species' has outlived its usefulness and should be replaced by terms that more accurately describe the different types of lineages that biologists refer to as 'species'.

He claims that 'species' should be eliminated and replaced with at least three new terms: 'biospecies', 'ecospecies', and 'phylospecies'. Likewise, an 'expression'-Eliminativist would claim that we should eliminate 'expression' from semantic discourse and replace it with more precise terms 'expression_1', 'expression_2', and so on.

In response, I will first note that Eliminativism towards a highly important theoretical term is a rather extreme position. This is certainly the case for a global Eliminativism that would remove the term from *all* theoretical discourse within a field. At the very least, it requires a significant coordinated overhaul of the field's linguistic conventions.

Ereshefsky's 'species'-Eliminativism is made more plausible by the conditions within the field of biology from which it emerged. The literature contains extensive discussion of the massive polysemy of 'species' and the problems this causes biologists. There are also many well-developed and well-motivated species definitions that Ereshefsky is able to put forward as effective replacement terms.

---

[149] Ereshefsky, "Eliminative Pluralism." 680. While Ereshefsky calls the view "Eliminative Pluralism", I emphasize "Eliminative" to distinguish it from the Pluralist view defended in this paper.

To gain the plausibility of 'species'-Eliminativism, the 'expression'-Eliminativist must overcome significant obstacles. They must show that 'expression' is the source of significant communicative problems, which becomes more complicated in the likely scenario that the term frequently co-occurs with other problematic terms. They must also present a set of replacement terms that overcome many problems with 'expression' without causing too many new ones. Until this happens, maintaining the use of 'expression' seems like an eminently reasonable default position.

Further, many theorists resist Ereshefsky's view despite the favorable conditions within biology. For instance, Novick and Doolittle reject 'species'-Eliminativism because it "emphasizes differences to the exclusion of similarities".[150] They argue that the similarities among different 'species'-concepts allow "knowledge, research questions, and patterns of reasoning" to easily transfer from one concept to another, even if there is no single feature that unifies them. Eliminating 'species' runs the risk of inhibiting these productive transfers.[151] So the 'expression'-Eliminativist faces yet another obstacle. Even if they show that 'expression' causes significant communicative issues within semantics, they also need to show that these costs outweigh the sorts of benefits discussed by Novick and Doolittle.

There are reasons to doubt that the 'expression'-Eliminativist would be able to overcome all of these obstacles. Ball observes that semantics is able to make progress despite general disagreement on the nature of its target system.[152] Different views of the nature of the target system entail different views of what expressions are. 'Expression'-Eliminativism that defines replacement terms with respect to different target systems has some *prima facie* appeal since

[150] Novick and Doolittle, "'Species' without Species." 79.
[151] See Brigandt, "Species Pluralism Does Not Imply Species Eliminativism." for another defense of 'species'.
[152] Ball, "Semantics as Measurement."

‘expression’ is plausibly ambiguous among them. However, theorists who think that expressions are mental representations, or conventions, or some sort of abstract object are all able to have productive discussions about expressions despite those differences. Replacing ‘expression’ with a set of terms that each correspond to different target systems could prevent productive discussions from ever occurring.

Eliminativism seems most appealing as a response to *ambiguity*. The appeal is much less evident when it comes to the phenomena that motivated this section. The *vagueness* of ‘expression’ arises from smaller-scale uncertainty about the contours of the lexicon and the rules that produce complex expressions. If we start defining ‘expression’-replacements corresponding to, say, different ways of prioritizing ambiguity tests, the number of replacement terms could quickly become unwieldy.

While the Eliminativist objection fails (at least without significant work to motivate it further), it does draw out an important feature of my Pluralist view. Allowing background constraints to be soft (*i.e.* vague and cancelable) can actually serve as a way to maintain unity in a plural scientific field. The Eliminativist faces the question of how many new terms should replace the problematic original. Once they start distinguishing uses of a term, they run the risk of descending a slippery slope towards a massive proliferation of replacements that causes more communicative issues than it solves. Similarly, the Pluralist faces the question of how many distinct projects should be acknowledged. Once they start distinguishing sets of background constraints adhered to by different theorists, they run the risk of descending a slippery slope towards a massive proliferation of projects. Pluralists distinguish projects to help avoid intractable disputes in which theorists tend to talk past one another. However, theorists can still

engage in productive disputes despite relatively minor differences in goals, assumptions, definitions, *etc*. As Pluralists distinguish more projects, they increase the risk of discouraging these sorts of productive discussions.

Allowing background constraints to be soft is a way to halt a descent down the slippery slope towards massive plurality. A vague background constraint can be a point of consensus among theorists even if they disagree on its precisifications. A cancelable background constraint can be a point of consensus among theorists even if they disagree on when exactly it applies.

2. Overriding the Constraint

Background constraints may be overridden if there are tradeoffs among theoretical goods. In Chapter 4, I argued that semantics may abstract away from the meaning of an expression if identifying the meaning significantly increases the complexity of the theory with little accompanying benefit.

Background constraints may also be overridden if they conflict with other desirable background constraints. In this section, I argue that the Taxonomic Monist imperative to theorize about the meaning-properties *of expressions* can conflict with a distinct (if related) theoretical goal for semantics. To characterize this goal, I draw on Bogen and Woodward's view of scientific explanation.[153] If the pursuit of Taxonomic Monism can interfere with other worthy goals for semantic theorizing, there are limits on how far the pursuit should go.

2.a. The Purpose(s) of Semantic Data

I'll begin with a seemingly straightforward question: what is the purpose of gathering semantic data? An obvious answer, and the one required by Taxonomic Monism, is that we use

[153] Bogen and Woodward, "Saving the Phenomena."

semantic data as evidence for and against theories of the meanings of expressions. When we ask which phrases containing an expression '*e*' are anomalous, which sentences containing '*e*' are true and which are false, which entailment relations '*e*' is part of, etc., we inquire into the meaning of '*e*'.

For a concrete example, suppose we observe some occurrences of the phrase 'much cow' and notice that each is anomalous. We naturally generalize from these observations to the conclusion that 'much cow' is anomalous in *all* of its occurrences. We then set about explaining this purported fact. We hypothesize that 'much cow' is anomalous because 'cow' is a count noun without a mass component and such nouns fail to compose with 'much'. We construct a theory of a -MASS meaning-property such that 'much [-MASS +SING NP]' is always anomalous. Our explanation of the supposed universal anomaly of 'much cow' has roughly the form of the following derivation:

Explanation of Anomaly:

P1. 'much [-MASS +SING NP]' is always anomalous.
P2. 'cow' is always -MASS.
C1. 'much cow' is always anomalous.

We then consider more occurrences of 'much cow'. We consider a (gruesome) context in which an 18-wheeler has hit a cow head-on at great speed and someone utters (2a) (a variant of (2) from the introduction):

(2a) There's much cow on the side of the road.

In this context, 'much cow' no longer seems anomalous. This observation provides evidence that our proposed explanandum (C1) is false: 'much cow' *isn't* always anomalous.

If we come to believe that (C1) is false (and we don't want to idealize by making false claims), the obvious and straightforward thing to do is reject (C1). We change our proposed explanandum by modifying (C1): instead of explaining the *universal* anomaly of 'much cow', we now want to explain the purported fact that 'much cow' is *usually*, but not always, anomalous.

By modifying (C1) and proposing a weaker explanandum, we're led to a choice between modifying (P1) and modifying (P2) (or both). We can modify (P1) by weakening the relationship between 'much [-MASS +SING NP]' and anomaly.

Explanation of Anomaly (P1 Modification):

P1*. 'much [-MASS +SING NP]' is usually anomalous.
P2. 'cow' is always -MASS.
C1*. 'much cow' is usually anomalous.

Replacing (P1) with (P1*) seems to require a significant change in the type of theory we're constructing. Suppose we've been working with the previously-discussed theory that posits the existence of two subdomains of entities: the discrete, countable entities and the "masses" or substances.[154] If 'cow' contains only discrete entities in its extension, it is -MASS. If 'much' quantifies over only masses, then 'much' will simply fail to compose with 'cow', causing the anomaly.

Under such deterministic frameworks of the kind favored by formal semantics, the anomaly of 'much [-MASS +SING NP]' is *entailed* by (among other things) the theories of 'much' and the -MASS meaning-property. Replacing (P1) with (P1*) therefore seems to require modifying the theory to include non-deterministic elements such as probabilities to allow a weaker relationship between 'much [-MASS +SING NP]' and anomaly.

[154] Link, "The Logical Analysis of Plurals and Mass Terms."

Consequently, it's unsurprising that most attempts to deal with problems like the current one modify (P2).

Explanation of Anomaly (P2 Modification):

P1. 'much [-MASS +SING NP]' is always anomalous.
P2*. 'cow' is usually -MASS.
C1*. 'much cow' is usually anomalous.

Many candidate theories support (P2*). One option builds on the distinct subdomain theory, claiming that 'cow' lacks the +MASS meaning-property in the lexicon. The 'cow' lexical entry only contains discrete entities (individual cows) in its extension and is therefore -MASS. However, before 'cow' attempts to compose with 'much', a mechanism such as a "universal grinder" occasionally intervenes to construct a +MASS NP from the -MASS lexical entry.[155]

Another option claims that 'cow' is actually +MASS in the lexicon, containing both individual cows and "cow mass" in its extension.[156] The inclusion of masses in the extension gives 'much' something to quantify. According to this theory, the meanings of lexical nouns are *overspecified* and the composition process contains mechanisms to delete components of lexical meanings. 'Cow' typically receives a count interpretation because the +MASS component of ‖cow‖ has been deleted by the time a sentential interpretation is achieved. However, composition with 'much' has the less common effect of deleting the +COUNT component of this meaning.

Whether modifying (P1) or modifying (P2), both responses to the recalcitrant datum use it to inform theories of the relationship between 'cow' and the -MASS property. Or perhaps that's too strong, and the focus on specific nouns like 'cow' is typically merely for illustrative purposes. If so, both responses at least use the datum to inform general theories of the

[155] Pelletier, "Non-Singular Reference: Some Preliminaries."
[156] Pelletier, "Lexical Nouns Are Both +mass and +count, but They Are Neither +mass nor +count."

relationship between nouns and their count / mass properties. While theorizing about the relationships between expressions and meaning-properties is undoubtedly an important use of semantic data, it can also be used for a different purpose.

2.b. Data and Hooke's Law

To motivate the idea, it will be helpful to think about how data is used in other fields. Consider the case of Hooke's Law, a physical law which claims that the force *f* needed to extend or compress a spring by some distance *d* is linearly related to that distance:

Hooke's Law: $f = -k*d$

where k is a "spring constant" determined by the properties of the spring. Suppose that we want to investigate Hooke's Law, so we perform an experiment with a coiled metal object (*i.e.* a *prima facie* spring) and a set of weights. We hang each of the weights from the coiled metal object and measure the subsequent stretching of the object.

Why might we perform this experiment? One reason is an interest in the properties of the coiled metal object. We want to know whether it's a spring (or how "spring-like" it is), so we check how well the object's behavior conforms to Hooke's Law. To do this, we check for an approximately linear relationship between the mass hung from the object and its extension. In the analogy between springs and count / mass nouns, this is like observing the behavior of 'cow' in order to determine whether it is -MASS.

If we wish to, say, use the object as a component of a machine, we should be interested in its physical properties. From a broader scientific perspective, though, our main interest probably isn't in the properties of *this particular object*. Instead, our reason for performing the experiment would be more general.

Perhaps our more general purpose is Popperian testing: by observing the behavior of the *prima facie* spring, we're attempting to falsify a theory:

Spring Theory: All springs behave according to Hooke's Law.

In the analogy between springs and count / mass nouns, this is like observing the behavior of 'cow' in order to determine whether all -MASS +SING NPs are anomalous with 'much'.

Upon closer examination, however, it appears that Spring Theory is actually *un*falsifiable. As Teller puts it:[157]

> To what materials does Hooke's Law apply, even inexactly? Well, to springs. But there is no independent, general characterization of what will count as a spring: a spring is whatever Hooke's Law (approximately) applies to. It is an objective fact about nature that there is a diverse range of materials to which Hooke's Law applies…

We have an intuitive, exemplar-based concept of springs (*e.g.* "coiled metal object"), but this is distinct from the scientific SPRING property connected to Hooke's Law. The SPRING property is defined by a type of behavior. *By definition*, all springs behave according to Hooke's Law. This means that Spring Theory, the universal claim we were potentially interested in testing, is a tautology and therefore unfalsifiable.

If our experiment can't be an attempted falsification of Spring Theory, that doesn't mean it's irrelevant for Hooke's Law. In order to get a better grip on a more general purpose for the experiment, I first need to discuss a distinction between *data* and *phenomena*.

2.c. The Data-Phenomena Distinction

Bogen and Woodward (henceforth B&W) constructed the distinction in order to counter the common view that "scientific theories predict and explain facts about 'observables': objects

[157] Teller, "'Saving the Phenomena' Today." 822.

and properties which can be perceived by the senses, sometimes augmented by instruments".[158] The data-phenomena distinction isn't intended to be sharp but can still be roughly characterized. Data are individual events or their representation, with such examples as Chomsky's single falling leaf or a pattern of spots on a photographic plate. Data have complex causes that are idiosyncratic to their experimental or observational context. Phenomena are generalizations from collections of data. In contrast to data, phenomena are repeatable, generally stable, and (hopefully) the result of a manageably small number of causes. Examples of phenomena corresponding to the above data are the attraction between Earth and nearby objects or the bending of rays of light. Because of the greater tractability of phenomena, B&W claim that science typically does and should explain them instead of individual data.

To illustrate the relevance of the distinction for semantics, consider Yalcin's non-exhaustive list of 5 types of facts that semantics is supposed to explain:[159]

1. Productivity facts: Speakers of a given language can understand and produce complex expressions in that language that they have never before encountered.
2. Communication facts: Speakers of a common language can transfer an abundant range of information systematically using that language.
3. Entailment facts: Some sentences in a language entail others; some sentences are inconsistent with others. Competent speakers manifest knowledge of such facts.
4. Acceptability facts: Some sentences or discourses in a language are judged to be unacceptable, or uninterpretable, or marked, by speakers of that language, while others are not.
5. Truth/appropriateness facts: Some sentences in a language are judged to be true or appropriate by speakers relative to actual or stipulated scenarios, while others are judged false or inappropriate relative to such scenarios.

Yalcin calls all of these "data", but upon further examination I think that there's a distinction between (1) - (2) and (3) - (5).

[158] Bogen and Woodward, "Saving the Phenomena." 303.
[159] Yalcin, "Semantics and Metasemantics in the Context of Generative Grammar."

(1) and (2) are clearly phenomena: the productivity of language and its use in successful communication are generalizations from the observation of many events (though, as I've discussed, the nature of these phenomena are controversial). Classifying (3) - (5) as data or phenomena depends on how they're interpreted. If (3) - (5) are merely talking about the fact that there are judgments of entailment, acceptability, and truth, they would clearly be phenomena, and if (3) - (5) are talking about individual judgment events, they would clearly be data. If (3) - (5) are rephrased to talk about truth, acceptability, entailment, etc. *tout court* instead of *judgments* about those concepts, it's less clear. Fortunately, drawing a clear line between semantic data and phenomena is unimportant for my purposes.

My characterization of a general purpose for the hanging weights experiment starts with the fact that Hooke's Law describes a phenomenon. This phenomenon, which I'll call *springiness*, is a linear relationship between the force applied to an object and its extension or compression. No object is expected to perfectly instantiate springiness: for instance, after a sufficiently large force is applied, any spring will simply break rather than continue stretching indefinitely. Springiness can be approximately instantiated in objects of many shapes, sizes, and material compositions.

According to B&W, data provide evidence for the existence of phenomena. On this view, the purpose of the hanging weights experiment is to provide evidence for the claim that springiness is a substantial phenomenon. Put another way, the purpose of the experiment is to provide evidence for or against the inclusion of SPRING as a property in a physical theory. The experiment provides evidence that springiness is a substantial phenomenon if and only if we

observe an approximately linear relationship between the force exerted by the weights and the extension of the object.

2.d. Demonstrating the Existence of Semantic Phenomena

B&W's view reveals another purpose for semantic data: providing evidence for the existence of semantic phenomena. Applied to the case I've been considering, the relevant phenomenon can be roughly described as follows:

> **Phenomenon 1:** Expressions of the form 'much [-MASS +SING NP]' are semantically anomalous.

Semanticists consider complex expressions such as '#much cow' for evidence that this is indeed a substantial phenomenon.

Demonstrating the existence of semantic phenomena doesn't require adherence to Taxonomic Monism. To show why this is the case, I need to say more about *how* one demonstrates the existence of phenomena. Importantly, data is insufficient on its own to accomplish this task. This is because phenomena are individuated by their causes. To borrow another example from Teller:[160]

> A whistling train rushes by you, and as it passes, the pitch of the whistle suddenly drops: an instance of the Doppler effect. Think now, instead, of a train sitting by with a whistle tone that suddenly drops, exactly as much as for the moving train, but in the second case the change is due to the engineer manually adjusting the whistle pitch. The effect sounds the same, but different cause, different effect.

If nearly-identical data have different causes,[161] then those data are instances of different phenomena.

As a result, identifying phenomena requires the theoretical work of distinguishing causes. Indeed, this theoretical background is already present in my characterization of the hypothesized

[160] Teller, "'Saving the Phenomena' Today." 817-818.
[161] Setting aside potential causes of "noise".

count noun phenomenon ('much [-MASS +SING NP]' is semantically anomalous). Even though 'much cow' and 'much the King of France' are both instances of 'much [-MASS +SING NP]' that sound odd, they are instances of different phenomena because only the former is an instance of *semantic* anomaly. Semantic anomaly is only one way in which expressions can be defective: other types of defects include syntactic unacceptability and pragmatic infelicity. The distinctions among these phenomena aren't given to us pre-theoretically;[162] instead, their contours are largely the result of significant theoretical work that attributes them to different types of causes.

The semantic anomaly of 'much [-MASS +SING NP]', if it is indeed a substantial phenomenon, is therefore caused by a semantic property of the NP. The task is to construct and justify a theory of this property, and semanticists typically hypothesize a property of NPs (which I've been calling -MASS) that causes the anomaly of 'much [-MASS +SING NP]'.

In order to justify the inclusion of -MASS in the broader semantic theory, semanticists can show that -MASS is a theoretically interesting property. For comparison, consider again the case of springs. SPRING becomes a more interesting property when we show that it applies to many different objects, not just metal coils. It also can be used to describe the displacement of, say, rubber bands or steel beams under certain conditions. SPRING also becomes a more interesting property if we show that it can be used to explain many different types of behavior. Suppose that we hang a weight from a spring, pull it away from equilibrium, and then release. The spring will become a harmonic oscillator around the equilibrium point. In conjunction with $f = m*a$ (and perhaps other hypotheses about, say, the effect of friction), Hooke's Law allows us to explain the

[162] At least, not very clearly.

motion of the spring. For instance, we may use it to predict the period of time it takes for the spring to complete one oscillation.

In short, generality forms an important part of the justification for admitting a theory of a phenomenon. Indeed, one reason that B&W think that science explains phenomena rather than data is that explanations of phenomena are more general. Since data are produced by a large and idiosyncratic set of causes, explanations of data will often fail to generalize to other entities and situations that don't share the specific features of the originals. A successful explanation of a phenomenon is *ipso facto* an explanation that generalizes to a variety of entities and situations.

We can demonstrate the generality of a -MASS theory of the semantic anomaly of 'much [NP]' in a fashion similar to how we demonstrate the generality of Hooke's Law. We can demonstrate that -MASS explains the behavior of many different entities by demonstrating that it explains the behavior of many different NPs. And we can demonstrate that -MASS explains many different types of behavior by deriving other behavior predicted by the property and confirming the existence of these phenomena. For instance, if the theory of -MASS predicts that -MASS +SING NPs cannot appear as bare arguments (*e.g.* '#Sally gave book to John'), we can demonstrate the generality of the theory by demonstrating that this new predicted behavior also occurs to a significant degree.

2.e. A Methodological Alternative to Falsification

As seen in previous sections, B&W's view of how scientific theories get "tested" opposes Popper's idea of testing theories *via* attempts at falsification. Some semanticists assume the Popperian view of their methodology. This assumption can cause difficulties that would be avoided by adopting B&W's view.

For an example, I'll consider the dispute between Pustejovsky[163] and Fodor and Lepore (F&L).[164] Pustejovsky observes the following:

(31a) The woman ate her meal.

(31b) The woman ate.

(32a) The woman devoured her meal.

(32b) #The woman devoured.

(31a) - (32a) are acceptable but (32b) isn't. Pustejovsky favors a semantic explanation of this data: 'eat' "denotes an activity of unbounded duration" whereas 'devour' "denotes a transition". Later in the book, he develops a mathematical theory of event structure that formalizes these informal descriptions.[165] The details aren't important for present purposes, but the different event structures assigned to 'eat' and 'devour' allow Pustejovsky to explain the distribution data in (31a) - (32b).

F&L take issue with this explanation. They don't claim that Pustejovsky's theory is *false*; instead, they claim that the "generalization is formulated so imprecisely that *one can't tell* whether [an introduced case] is a counterexample".[166] They ask what exactly it is to "denote an activity of unbounded duration", claiming that in at least some cases it's basically impossible to tell.

[163] Pustejovsky, *The Generative Lexicon*.
[164] Fodor and Lepore, "The Emptiness of the Lexicon: Critical Reflections on J. Pustejovsky's the Generative Lexicon."
[165] Pustejovsky, *The Generative Lexicon*. Chapter 8.
[166] Fodor and Lepore, "The Emptiness of the Lexicon: Critical Reflections on J. Pustejovsky's the Generative Lexicon." 273 (emphasis theirs).

F&L don't explicitly state the generalization they're concerned about, but I think it's something like this (with a corresponding generalization about why 'devour' can't drop its object):

> **Unbounded Activity:** All expressions that denote an activity of unbounded duration are able to drop their objects.[167]

We can represent the form of this generalization as: $\forall x(\text{condition}(x) \rightarrow \text{behavior}(x))$, where 'condition($x$)' takes as its value the metalanguage restrictor 'denotes an activity of unbounded duration' and 'behavior(x)' takes 'able to drop its object'. With their talk of generalizations and counterexamples, F&L seem to want to perform a Popperian testing of this theory. Such testing tries to falsify the theory by examining entities meeting the condition. If at least one of them fails to exhibit the behavior, then the universal generalization has been falsified.

Vagueness and uncertainty jeopardize this practice. If it's vague or uncertain which entities meet the condition (*e.g.* which expressions denote activities of unbounded duration) or which entities exhibit the behavior (*e.g.* are able to drop their objects), then it's vague or uncertain whether there are entities that meet the condition but don't exhibit the behavior. If one can't tell whether a verb denotes activities of unbounded duration, one can't tell whether phrases containing the verb are potential falsifiers of Pustejovsky's theory.

Consider the following pair:

(33a) The woman rotated the jar.

(33b) #The woman rotated.

[167] In order to explain the contrast between (31b) and (32b), there would presumably also be a generalization explaining why expressions like 'devour' *can't* drop their objects.

‘Rotated’ is unable to drop its object (unless it’s given a different semantic interpretation where the woman herself is rotating). The markedness of (33b) makes ‘rotate’ a counterexample to Pustejovsky’s theory *if* it denotes an activity of unbounded duration.

Does ‘rotate’ denote an activity of unbounded duration, though? In (33a), ‘rotate’ plausibly seems to denote a “transition” or “bounded activity” in which the woman rotated the jar until a preferred orientation was achieved. However, in (34), ‘rotate’ seems to denote an unbounded activity: the machine continuously rotates the display for an indefinite period of time.

(34) The machine rotated the window display.

So there is conflicting evidence for whether ‘rotate’ denotes an activity of unbounded duration. Unless this conflict can somehow be resolved, it’s therefore uncertain whether ‘rotate’ provides a counterexample to Pustejovsky’s theory. If falsifiabilty is the goal, this is a big problem for the theory.

However, according to the “Phenomenon View”, F&L’s focus on falsification sets them down the wrong path. When F&L complain about the vagueness of the condition, they are complaining about the vagueness of an informal condition (‘denotes an activity of unbounded duration’) stated in English. But the informal generalization isn’t *actually* Pustejovsky’s theory. The real accomplishment of the theory is translating the informal condition into a precise formal language. The translated generalization, broadly characterized as $\forall x$(formal-condition(x) $\rightarrow$ behavior(x)), is something like a theorem as long as the formal condition and the relevant auxiliary hypotheses are well-designed. Given the theory’s hypotheses about the meanings of other relevant expressions, its compositional mechanisms, and the interpretation of its outputs

(*e.g.* interpreting successful composition as semantic felicity), any expression meeting the formal condition will exhibit the behavior.

The informal condition 'denotes an activity of unbounded duration' serves as a ladder that may be kicked away once the formal theory is reached.[168] It gets us to important features of the formal theory of activities of unbounded duration: for instance, the theory contains an ordered set of times without an upper bound (*i.e.* 'unbounded'). However, once it's kicked away, it plays a much more limited role in testing the theory than F&L think it's supposed to.

The prospects for attempting to falsify Pustejovsky's theory get even worse when we move from the informal generalization to the formal theory. If it's difficult to tell which verbs denote an activity of unbounded duration or a transition, it's at least as difficult to tell which verbs have the abstract mathematical properties proposed by the theory. Perhaps we allow our intuitive understanding of which verbs denote an activity of unbounded duration to guide us to the verbs that have the corresponding mathematical properties. This method can't be any better than our intuitive understanding of the condition, which (as F&L claim) is already pretty flawed.

Fortunately, B&W's alternative view of how scientific theories get "tested" allows semanticists to avoid this problem. According to this view, the challenge facing Pustejovsky's theory isn't testing a universal generalization by attempting to falsify it. Theorists aren't compelled to find verbs that denote activities of unbounded duration and check whether they can drop objects. Instead, the challenge is demonstrating the existence of the hypothesized phenomenon. As before, this at least involves demonstrating that Pustejovsky's theory is sufficiently general, explaining many behaviors of many VPs. Importantly, though, this doesn't

[168] Though perhaps not entirely; maybe it is still used to form initial hypotheses about which expressions are in the category.

require us to be able to point to *any* expression and definitively say that *it* has the property referred to by the theory. F&L's criticism has been avoided; if we only have a vague pre-theoretical understanding of which verbs denote an activity of unbounded duration, that isn't a problem.

2.f. The Spirit of Taxonomic Monism I

In sum, the idea that science explains phenomena rather than data provides semantics with the goal of demonstrating the existence of semantic phenomena and constructing theories to explain them. I certainly don't intend my brief discussion of what goes into accomplishing this goal to be exhaustive. Instead, I hope to have made it clear that accomplishing this goal does not require adherence to the "letter" of Taxonomic Monism. To justify a theory of a semantic property like -MASS, we have to show that the theory of -MASS explains a variety of behaviors of many expressions. Clearly, this doesn't require a language-universal description of which expressions are -MASS.

More interestingly, demonstrating the existence of phenomena doesn't require adhering to background constraints in the "spirit" of Taxonomic Monism. Even if semanticists don't need to state the meaning-properties of every expression, the spirit of Taxonomic Monism might require them to at least have the ability to do so:

> **Universal Ability:** Semanticists must have the ability to state the meaning-properties of every expression.

Indeed, as I will show, Universal Ability seems to appear in the literature as (at least) a tacit assumption.

Suppose we're trying to categorize nouns as either count or mass. How can we go about doing this? One option is to consider various "well-formedness" criteria.[169] For instance, the following criteria may be used when classifying nouns like 'mud' and 'deer':

> **Criterion 1:** Mass nouns, but not count nouns, occur with quantifiers like 'much' and 'more'.
> Example: 'more mud' vs. '#more deer'
>
> **Criterion 2:** Count nouns, but not mass nouns, occur with quantifiers like 'each'.
> Example: '#each mud' vs. 'each deer'
>
> **Criterion 3:** Count nouns, but not mass nouns, can be pluralized.
> Example: '#muds' vs. 'lamps'

Another option for classifying nouns as mass or count looks at the nature of their referents, or perhaps how those referents are typically conceived by us. As previously discussed, it seems plausible that count nouns describe discrete, countable objects whereas mass nouns describe non-countable substances. Recall the theory from Link,[170] which posits that the domain of entities is divided into two subdomains: the discrete objects and the "stuff" or masses. Cars are discrete and countable, so 'car' is a count noun; mud isn't, so 'mud' is a mass noun. And even though water consists of discrete water molecules, it seems plausible that we typically conceive of it as "stuff". If referent-based classification is based on our concepts, the method classifies 'water' as a mass noun.

Different criteria frequently classify the same noun in different ways. Consider the case of 'change', disambiguated to refer to collections of coins. Well-formedness criteria classify 'change' as mass: for instance, 'more change' is acceptable but '#each change' is not. Referential criteria, on the other hand, classify 'change' as count: 'change' refers to discrete, countable

[169] Pelletier and Schubert, "Mass Expressions." The authors calls these "syntactic" criteria.
[170] Link, "The Logical Analysis of Plurals and Mass Terms."

objects (*i.e.* coins), and the claim that we conceive of change as "stuff" seems much less plausible than in the case of water.

Pelletier and Schubert seem to think that such cross-classification is a fatal flaw for any proposed set of count / mass criteria. They claim that "[t]he various criteria proffered for distinguishing count and mass expressions or occurrences differ in how they classify … [t]his seems to indicate that we should choose one of the criteria as our touchstone and not try to use a variety in concert".[171] Criticizing Bunt's hybrid classification method that uses both well-formedness and referential criteria,[172] the authors rule it out because "such a method would be bound to cross-classify expressions".

Pelletier and Schubert claim that the cross-classification of some nouns is a fatal flaw of a set of count / mass classification criteria. This is only the case if we need the criteria to put each noun into exactly one count / mass category. In other words, the authors seem to assume Universal Ability by assuming that we need to have the ability to state the count / mass meaning-properties of each noun.

Demonstrating the existence of semantic phenomena doesn't require adherence to Universal Ability. If semanticists are engaged in Popperian testing of theories, then the ability to universally classify expressions is desirable. Suppose we're testing the theory "All mass nouns are anomalous with 'each'." by attempting to falsify it. By the same argument F&L used against Pustejovsky, this task is much more difficult if we lack the ability to clearly distinguish the mass nouns from the non-mass nouns: if noun '*n*' is non-anomalous with 'each' but we're unable to

[171] Pelletier and Schubert, "Mass Expressions." 18-19.
[172] Bunt, "The Formal Semantics of Mass Terms."

know whether '*n*' is a mass noun, then we're unable to know whether our theory has been falsified.

However, if semanticists are engaged in demonstrating the existence of semantic phenomena, they aren't engaged in Popperian testing of theories. Therefore, the motivation for adhering to Universal Ability has been lost. In order to demonstrate the existence of semantic phenomena, semanticists justify the inclusion of semantic properties in the theory by demonstrating that they explain many behaviors of many expressions. This task clearly doesn't require adherence to the letter of Taxonomic Monism, and it also clearly doesn't require adherence to the Universal Ability "spirit". Even if we're unable to know whether some nouns are -MASS, we may still justify the inclusion of the property in the theory by showing that it explains many behaviors of many expressions.

By dropping the Taxonomic Monist requirement of Universal Ability, we allow the cross-classification of some expressions. If we don't *need* to be able classify 'change' as either count or mass, then we may maintain both the well-formedness and referential criteria even if they disagree on the classification of 'change'. On this view, classification criteria are supposed to be reliable indicators of category membership, but they don't need to be universally applicable. For instance, if a noun refers to discrete, countable entities, this indicates that the noun is +COUNT while leaving room for the possibility that the indication is misleading (as might be the case with 'change'). If other criteria indicate that the noun is +MASS and there is no clear way to resolve the conflict, the noun may remain unclassified.

2.g. The Spirit of Taxonomic Monism II

One way to capture the spirit of Taxonomic Monism is by requiring semanticists to have the ability to state the meaning-properties of every expression. Semanticists can demonstrate the existence of semantic phenomena without adhering to this version of the spirit of Taxonomic Monism. However, perhaps the true spirit of the constraint is weaker. Surely, semanticists must at least state meaning-properties of expressions in order to adhere to the spirit of Taxonomic Monism, even if they don't seek Universal Ability. Taxonomic Monism is a claim about the *focus* of semantic theorizing, and it implores semanticists to focus on the relationship between expressions and meaning-properties.

Expression Focus: State the meaning-properties of expressions.

In this section, I argue that semanticists can demonstrate the existence of semantic phenomena without an Expression Focus. In particular, I argue that semanticists can do this without stating the meaning-properties of *lexical* expressions. To illustrate the distinction between lexical expressions and other expressions, we can once again turn to the literature on count and mass nouns. Consider Pelletier and Schubert, who claim that the discourse surrounding the issue of count / mass nouns has been hampered by failure to adequately appreciate the following point:[173]

> When one asks: what is the semantic value of 'water', the question is ambiguous between at least the following. What is the semantic value of water, considered merely as an entry in the lexicon? vs. What is the semantic value of 'water', as it occurs in this sentence?

[173] Pelletier and Schubert, "Mass Expressions." 52.

To remedy the situation, the authors suggest that there may be several distinct "levels" of expressions. They suggest the idea of a hierarchy of levels of expressions in response to some *prima facie* puzzling data.

Consider the truth conditions of the following pair:

(35a) Chocolate tastes good.

(35b) Every chocolate comes in a wrapper.

It seems that 'chocolate' has different extensions in the two sentences. (35a) is true just in case *masses* of chocolate taste good, whereas (35b) is true just in case (roughly) every *conventional portion* of chocolate comes in a wrapper. In other words, 'chocolate' seems to have a mass meaning in (35a) and a count meaning in (35b). Importantly, these examples aren't an aberration. 'Chocolate' is a *dual-life* noun, meaning that both count and mass uses are common.[174] And many (even most) common nouns are dual-life.

What, then, is the meaning of 'chocolate' and other dual-life nouns? According to Pelletier and Schubert, this question is ambiguous. They posit syntactic structures like the one in Figure 3 for sentence (35a) that contains multiple levels of expression for 'chocolate'. The structure contains a lower-level lexical expression 'chocolate_L' at N0 and a higher-level expression 'chocolate_H' at N1 that composes with 'tastes good' and plays an immediate role in determining the sentences' truth conditions. On this view, 'the meaning of 'chocolate'' is ambiguous between 'the meaning of 'chocolate_L'' and 'the meaning of 'chocolate_H''.

[174] Kiss et al., "Issues of Mass and Count: Dealing with `Dual-Life' Nouns."

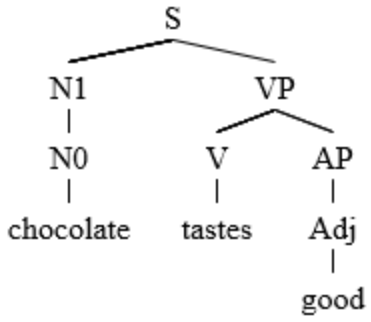

Figure 3

With this distinction between levels of expressions in mind, Pelletier and Schubert characterize what they take to be the explanatory task at hand:[175]

> The proper investigation of mass expressions will proceed by first finding a suitable type of value for the lexical entries so that, when combined to form "higher" node-values, these too will have the correct sort of values.

The project these authors are engaged in requires semanticists to state count / mass properties of expressions in the lexicon; therefore, it has a lexical Expression Focus.

2.h. DESP without Taxonomic Monism

A lexical Expression Focus requires semanticists to state meaning-properties of lexical expressions. The overarching goal of this subsection is to argue that demonstrating the existence of semantic phenomena can be done *without* a lexical Expression Focus. In other words, there's a goal for semantics that doesn't require theorists to state meaning-properties of lexical expressions.

One argument for this conclusion revisits the analogous case of Hooke's Law and springs. Demonstrating the existence of springiness is accomplished by justifying the inclusion of the SPRING property in a physical theory. To justify the inclusion of SPRING in the theory, several conditions should be met. One condition is goodness-of-fit. We need to show that SPRING

[175] Pelletier and Schubert, "Mass Expressions." 55.

captures real behavior sufficiently closely: for instance, we observe that an object's extension is approximately linearly related to the force applied over a wide-enough range of forces. Another condition is that SPRING is generally applicable. We show that it captures many other types of behavior (*e.g.* oscillation), and we show that it captures behaviors of many different objects. Therefore, in order to demonstrate the existence of SPRING phenomena we need to show that SPRING approximately captures many behaviors of many objects.

My claim is that the case of springs is analogous to the case of semantic phenomena like the count / mass phenomena I've been discussing. To demonstrate the existence of such semantic phenomena, we need to show that the relevant meaning-property approximately captures many behaviors of many expressions. We can do this without a lexical Expression Focus.

Consider how semanticists can demonstrate that a meaning-property approximately captures *one* behavior of *one* expression. In general, one can demonstrate that the meaning-property *approximately* captures the expression's behavior by demonstrating that it captures a central tendency of instances of the expression's behavior. Since there are multiple ways to measure central tendency, there are multiple potential interpretations of "approximately" capturing behavior. One option is a "mean" interpretation, in which a property approximately captures behavior just in case it generally gets close. For instance, even the most prototypical springs will always exhibit slight deviations from a linear relationship between force and extension.

More appropriate for the case of formal semantics, however, is a "median" or "mode" interpretation, in which a property approximately captures behavior just in case it perfectly captures the central or most common value. This interpretation is more appropriate because the

data in formal semantics generally takes one of a small number of values (true or false, anomalous or non-anomalous, etc.), so properties posited by formal semantic theories generally don't get merely close. They either perfectly capture behavior or miss the mark significantly. For instance, if 'much cow' is anomalous in most contexts but non-anomalous on rare occasions, -MASS perfectly captures most of the behavior of 'cow' but completely fails to capture the rare occasions. On the median or mode interpretation, -MASS approximately captures the behavior of 'cow'.

With this reduction of the problem in hand, we can see that the task of demonstrating that a meaning-property approximately captures many behaviors of many expressions can be broken down into many demonstrations that the meaning-property perfectly captures one behavior of one expression.

Suppose we observe that an occurrence of 'much cow' is anomalous. We can show that -MASS perfectly captures this instance of the behavior of the *lexical* 'cow_L' without stating meaning-properties of the lexical expression, *i.e.* without a lexical Expression Focus. Using Pelletier's distinction among different levels of expressions, we can distinguish between the lexical expression 'cow_L' and the higher-level expression 'cow_H' that attempts to compose with 'much'. To explain the observation, we posit that 'cow_H' is a [-MASS +SING NP]. Assuming our theory of the meaning of 'much' entails that 'much' is unable to compose with any [-MASS +SING NP], it follows that 'much' is unable to compose with 'cow_H'. This inability to compose results in an anomaly.

We can demonstrate that -MASS captures this instance of the behavior of the lexical expression 'cow_L' by attributing the property to the higher-level expression 'cow_H'. We don't

need to attribute -MASS to 'cow_L' or, really, state much of anything about its meaning. All we need to do is establish that this is indeed an instance of the behavior of 'cow_L' by positing a tight relationship between it and 'cow_H'. Generally, what this amounts to can be subsumed under the general claim that the meaning of 'cow_H' is a function of the meaning of 'cow_L' and the context of utterance. We can say that such a function exists without saying what the meaning of 'cow_L' is.

At this point, an advocate of Taxonomic Monism might raise an objection. I've shown we can demonstrate that a meaning-property approximately captures many behaviors of many expressions without a lexical Expression Focus. However, this intermediate conclusion doesn't automatically entail that we can *demonstrate the existence of semantic phenomena* without a lexical Expression Focus. Perhaps demonstrating that a meaning-property approximately captures many behaviors of many expressions is insufficient on its own to establish the existence of a phenomenon; further, perhaps a lexical Expression Focus is one of the necessary ingredients I've been missing.

In reply to the advocate of Taxonomic Monism, I'll provide a more direct argument that a lexical Expression Focus isn't required to demonstrate the existence of semantic phenomena. The argument depends on the observation that many semantic phenomena are multiply realizable with respect to lexical systems.

I'll illustrate the point using a different count / mass phenomenon. Consider the truth conditions of the following sentences, specifically the contribution to truth conditions made by the plural NP:

(36) Mary owns three cars.

(37) Kim hands out chocolates on Halloween.

(38) We tried many wines on our tour.

Roughly, 'cars' is true of collections of multiple cars, 'chocolates' is true of collections of multiple standard portions of chocolate, and 'wines' is true of collections of multiple varieties of wine. Each of these plural NPs is true of collections of discrete, countable entities: cars, standard portions of chocolate, and varieties of wine. In other words, each is true of collections of "count" entities. We can generalize from these examples to the claim that "[+PL NP]" is true of collections of count entities.[176]

Further, it seems that *only* the count interpretation of plural NPs is semantically well-formed. Using (39) to describe a situation in which Sam drank a large quantity of one wine is infelicitous:

(39) Sam drank many wines.

A count reading of 'many wines' on which Sam drank many varieties of wine seems to be forced. Drawing on this example and others, we can refine the generalization:

Phenomenon 2: '[+PL NP]' is true of only collections of count entities.

The count / mass literature contains many theories of the lexicon that entail the existence of Phenomenon 2. Two of these are the previously mentioned overspecification theory and a "universal packager" theory that is the converse of the universal grinder:

> **Overspecification:** The semantic value of every lexical noun contains all the values of which the noun is true. +PL deletes +MASS meanings.
>
> **Universal Packager:** Semantic values of lexical nouns contain only the "stronger" value. +PL contains a "universal packager" that creates +COUNT meanings from +MASS meanings.

[176] See Pelletier, "Lexical Nouns Are Both +mass and +count, but They Are Neither +mass nor +count." for a list of count / mass truth-conditional features.

If a language meets the conditions of either Overspecification or Universal Packager, then Phenomenon 2 is realized in that language. However, the two theories entail different theories of the lexical meanings of prototypical mass nouns like 'wine'. Overspecification entails that 'wine_L' is both +MASS and +COUNT, whereas Universal Packager entails that 'wine_L' is +MASS but -COUNT.

A commitment to the existence of Phenomenon 2 is a commitment to the existence of some underlying lexical system that can realize the phenomenon. However, this isn't a commitment to a particular theory about the lexical meanings of the relevant expressions. Indeed, one can commit to the existence of Phenomenon 2 without committing to any claims about the count properties of lexical nouns.

Perhaps the commitment to the existence of Phenomenon 2 follows from a commitment to the disjunction of Overspecification and Universal Packager. This disjunction entails that prototypical mass nouns like 'wine' are either +COUNT or -COUNT in the lexicon; in other words, it commits us to no claims whatsoever about the lexical count properties of such nouns. If belief in the disjunction is justified, then belief in the existence of Phenomenon 2 is justified.

This example shows one possible way to demonstrate the existence of semantic phenomena without a lexical Expression Focus. The existence of a phenomenon can follow from premises that are uninformative about lexical meaning.[177] If we have justified belief in the lexically uninformative premises, then we have justified belief in the existence of the phenomenon without having stated the meaning(s) of the expression(s) of interest.

[177] One might object that the premise *is* informative about the meaning of 'wine_L' since both theories entail that 'wine_L' is +MASS. In response, we can add another disjunct, an Underspecification theory on which 'wine_L' is neither +MASS nor +COUNT but only gains those properties through some contextual mechanism.

2.i. Conflict Between Taxonomic Monism and DESP

Demonstrating the existence of phenomena provides semantics with a theoretical goal that doesn't require adherence to the letter or the spirit of Taxonomic Monism. Further, as I argue in the remainder of this section, the pursuit of Taxonomic Monism can conflict with the pursuit of this alternative. As a result, a project that adheres to Taxonomic Monism may have reason to occasionally override the constraint.

In Section (2.g), I introduced Pelletier's work on count and mass nouns. Pelletier is primarily motivated by the prevalence of dual-life nouns, and his theory reflects this. He prefers Overspecification with the central claim that almost every common noun is both +COUNT and +MASS in the lexicon. This claim follows from the more general hypothesis that "the semantic value of every lexical noun contains *all* the values of which the noun is true".[178] For instance, dual-life nouns such as 'chocolate' are clearly true of both count entities (*e.g.* individual chocolates) and mass entities (*e.g.* melted chocolate), so their lexical entries are both +COUNT and +MASS.

Many occurrences of nouns have only a +COUNT or +MASS meaning. The occurrence in (37) is only +COUNT: the sentence says that Kim hands out individual pieces of chocolate on Halloween. On Pelletier's multiple-levels theory, lexical meanings need to be transformed into the meanings that immediately determine truth conditions. Since lexical meanings contain too much information for many of their uses, the theory contains rules that delete components of meanings. In order for 'chocolate' to syntactically combine with 'every' to form 'every chocolate', a semantic rule deletes the +MASS portion of the meaning of the lexical '$\text{chocolate}_{\text{L}}$'.

[178] Pelletier, "Lexical Nouns Are Both +mass and +count, but They Are Neither +mass nor +count." 19.

Only the +COUNT portion remains in 'chocolate$_H$', which allows 'every' to quantify over the correct entities.

Overspecification is designed to explain how dual-life phenomena arise from the lexicon while also meeting several desiderata. For instance, the competing Ambiguity theory posits that 'chocolate$_L$' is itself ambiguous between distinct +MASS and +COUNT lexical entries. When generalized, this theory leads to the proliferation of lexical entries, an undesirable consequence that Overspecification avoids by including +MASS and +COUNT in a single entry.

Other competing theories posit that one of +MASS or +COUNT is the basic meaning-property found in the lexicon, with the other property only appearing as the result of some rule acting on the basic property.[179] These theories face some immediate difficulties that Overspecification avoids. They require us to somehow choose one meaning-property as basic, whereas Overspecification views both properties as equally basic. They also require theorists to posit rules such as a "universal grinder" that takes in +COUNT meanings and "grinds" them into a +MASS meaning or a "universal packager" that takes in +MASS meanings and "packages" them into a +COUNT meaning. Such rules seem inevitably more complex than the simple deletion operation required by Overspecification.

There are several advantages of Overspecification, but there is also a potential disadvantage that I'd like to highlight. Recall that Pelletier based his theory on the hypothesis that the lexical meaning of a noun contains all the values that noun is true of. This hypothesis is intended to explain the behavior of dual-life nouns, which are obviously true of both count and

[179] Bale and Barner, "The Interpretation of Functional Heads: Using Comparatives to Explore the Mass/Count Distinction."

mass entities. However, it also has consequences for prototypical count nouns such as ‘cow’ and prototypical mass nouns such as ‘mud’.

I’ll return to the example that began Section 2. In normal contexts, ‘#much cow’ seems semantically anomalous, a fact that was initially explained by positing that ‘cow’ is -MASS. In some contexts, though, it seems that ‘much cow’ is in fact *non-anomalous* and true of cow stuff. If ‘cow’ can be true of masses, then it follows from Pelletier’s hypothesis that the lexical entry for ‘cow’ is +MASS. If ‘cow’ is +MASS rather than -MASS in the lexicon, then (barring some new addition(s) to the theory) the theory entails that ‘much cow’ is *always* non-anomalous.

Indeed, since it’s plausible that such “grinding” contexts can be constructed for most prototypical count nouns, it’s plausible that most prototypical count nouns are actually +MASS in the lexicon according to Overspecification. Similarly, since it’s plausible that “packaging” contexts can be constructed for most prototypical mass nouns,[180] it’s plausible that most prototypical mass nouns are actually +COUNT in the lexicon. Therefore, the theory entails that ‘much [prototypical count noun]’ will generally be non-anomalous. It also entails that similar examples of supposed mass phenomena such as ‘[prototypical mass noun] + PL’ will generally be non-anomalous.[181]

Overspecification suppresses behaviors that we *prima facie* wanted our count / mass theories to explain. It denies the existence of Phenomenon 1, the semantic anomaly of ‘much [-MASS +SING NP]’. Recall that a phenomenon is a type of behavior with a common cause. Overspecification denies the existence of [-MASS +SING NP],[182] thereby denying the existence of

[180] Suppose we visit a strange vendor that sells jars of mud; in this context, ‘I’d like three muds.’ seems acceptable.
[181] Pelletier leaves room for rare exceptions of -MASS nouns and -COUNT nouns, mentioning ‘piece’ and ‘hole’ and potential examples of both categories.
[182] Perhaps it leaves room for rare exceptions.

the phenomenon. Unless some other mechanism is proposed to explain the anomaly of 'much [prototypical count noun]', this behavior vanishes from the theory's purview.

2.j. Why is Taxonomic Monism Blameworthy?

According to Pelletier, the central task of theorizing about count / mass phenomena is describing the count / mass properties of lexical expressions and the role they play in determining the properties of "higher-level" expressions. By theorizing about the meaning-properties of expressions, Pelletier is adhering to the spirit of Taxonomic Monism. He ultimately arrives at a theory that suppresses a robust phenomenon that is *prima facie* a central characteristic of count behavior.

The example of count / mass theorizing illustrates a potential problem with strict adherence to Taxonomic Monism. Taxonomic Monism is a claim about the *focus* of semantic theorizing. It tells semanticists to inquire into the relationship between expressions and meaning-properties. This Expression Focus can bear responsibility for the suppression of phenomena and the loss of interesting and theoretically tractable behavior from the scope of the theory.

I leave room for uncertainty as to whether the current example is *actually* an instance of this type of problem. First, one might not be overly bothered by the suppression of the semantic anomaly of 'much [-MASS +SING NP]'. Perhaps the behavior *isn't* actually sufficiently robust to concern ourselves with, or it is best explained by some other type of theory (*e.g.* a pragmatic theory), or its suppression is an acceptable price to pay in order to acquire other theoretical goods. Second, if one *is* bothered by the suppression of the behavior, it's unclear how much

Taxonomic Monism *itself* is to blame. In constructing his theory, Pelletier is guided by several theoretical desiderata in addition to Taxonomic Monism; perhaps these are culpable.

Regardless, the current example illustrates how Taxonomic Monism *can* be responsible for suppressing phenomena. By imploring us to take an Expression Focus, Taxonomic Monism increases the risk of suppressing phenomena. If there is variation in an expression's behavior, then there are more viable candidate theories of the expression's meaning. When choosing among multiple candidate theories, we will often be forced to accept a tradeoff (*e.g.* losing goodness-of-fit to the data in exchange for gaining greater simplicity). One of these potential tradeoffs might be suppressing a phenomenon in exchange for other theoretical goods. Importantly, *without* the Expression Focus, suppressing the phenomenon is much less likely to occur and can indeed look like a gross overreaction.

Of these claims, the claim that variation in an expression's behavior *increases* the number of viable candidate theories is potentially the most counter-intuitive. A curve-fitting metaphor helps show why this is the case. Consider the toy example shown in Figure 4. Suppose that our initial data is represented by the blue points. With just the initial data in hand, choosing a function to represent the relationship between the x and y quantities isn't much of a choice at all: $y = 3x$ is both simple and the best fit to the data.

However, if we make some more observations (represented by the red points) and discover that the data is messier than we had initially hoped, choosing a function $y = f(x)$ becomes much more difficult. We have several decisions to make. We need to decide whether to keep the point near (5, 25) in our dataset or discard it as an outlier. We also need to decide whether to continue using a linear equation. There now appears to be convexity to the data,

especially if we keep the point near (5, 25), which makes the move to a quadratic equation more appealing. As a result of being forced into these decisions by the messier data, there are now many theoretical options available to us, whereas the uniform initial data only really presented one option.

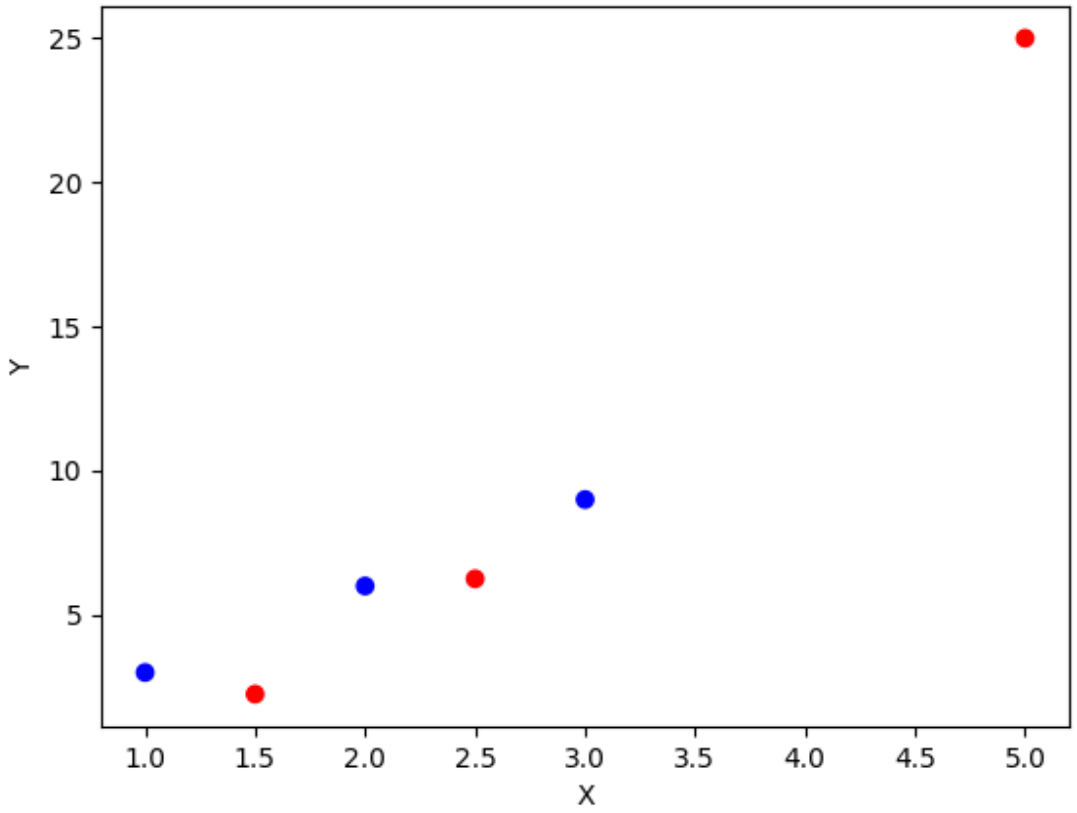

Figure 4

Likewise, the "messiness" of the count / mass behavior of expressions like 'cow' increases the number of viable theories of their meaning. Suppose that we initially consider occurrences of 'much cow' in "normal" contexts, observing that they're all anomalous. With an Expression Focus, these observations are used as evidence for theories of the meaning-properties

of expressions. A simple universal theory seems like an obvious candidate for an initial hypothesis:

Simple Theory: 'cow' is always -MASS.

We then observe more occurrences of 'much cow' in an attempt to falsify the Simple theory. Eventually we discover variation in the behavior of 'cow'. 'much cow' isn't always anomalous; in some contexts, the phrase is actually non-anomalous.

If the behavior of 'cow' was uniform, we could be content with simple universal theories. However, the observation of behavioral variation in the behavior of 'cow' disturbs this contentment. Variation opens the door to many competing theories, including several that have already been extensively discussed and are summarized in Figure 5.

Theory	Description	Benefit	Cost
Simple	'cow' is -MASS. No grinding rule	Simple, captures central tendency	Idealizes away from variation
Ambiguity	'cow_0' is -MASS. 'cow_1' is +MASS. No grinding rule	Captures central tendency and variation	Bigger lexicon adds complexity
Universal Grinder	'cow' is -MASS. Grinding rule	Captures central tendency and variation	Grinding rule adds complexity
Overspecification	'cow' is +MASS. No grinding rule	Simple, captures non-central tendency	Fails to capture central tendency

Figure 5: Benefits and Costs of count / mass theories

Each theory brings both costs and benefits. The Simple theory either idealizes away from the occasional non-anomaly of 'much cow' or brings a burden to explain it away. The Ambiguity theory, which claims that 'cow' is actually lexically ambiguous between -MASS 'cow_0' and

+MASS ‘cow_1’, can accommodate the behavioral variation but brings the cost of greater complexity in the form of an expanded lexicon. The Universal Grinder theory claims that ‘cow’ is univocally -MASS in the lexicon, adding a “grinding” rule that transforms -MASS meanings into +MASS meanings. While this option can accommodate the behavioral variation without expanding the lexicon, it also brings the cost of greater complexity by including the grinding rule. Finally (though not exhaustively), a fourth option is the one taken by Pelletier. The Overspecification theory claims that ‘cow’ is univocally +COUNT *and* +MASS in the lexicon, making a grinding rule unnecessary. While this theory has simplicity advantages and can accommodate the *non-anomalous* occurrences of ‘much cow’, it brings the cost of suppressing Phenomenon 1, the anomaly of ‘much [-MASS +SING NP]’.

Since each of these theories brings both costs and benefits, one must accept a tradeoff if choosing among them, only gaining the benefit(s) of the theory by incurring a cost. To prefer Pelletier’s theory, one must think that its simplicity advantages and accommodation of non-anomalous occurrences of ‘much cow’ outweigh its suppression of Phenomenon 1.

Taxonomic Monism contributes to suppressing the phenomenon *via* its position at the beginning of a process of theory choice. It directs us to theorize about the lexical meaning of ‘cow’; variation in the behavior of ‘cow’ opens the door to many candidate lexical theories; choosing among these theories forces us to accept tradeoffs; depending on the tradeoff we accept, our chosen theory can suppress the phenomenon.

Taxonomic Monism can be *blamed* for suppressing the phenomenon since this outcome is unlikely to occur without adherence to the constraint. In recent sections I developed an alternative view of semantic theorizing that doesn’t share Taxonomic Monism’s Expression

Focus. On this view, the role of semantic data is to provide evidence for and against the existence of semantic phenomena. As a consequence, theorizing with this view has a *Category* Focus: its task is to justify the inclusion of a semantic category in our broader theory by showing that it appears in good explanations of a wide range of behaviors.

In the current example, the category of interest is [-MASS +SING NP] and the relevant phenomenon is the semantic anomaly of 'much [-MASS +SING NP]'. With a Category Focus, observations of anomalous occurrences of 'much cow' in normal contexts are used as evidence in favor of the existence of the phenomenon. An observation of a non-anomalous occurrence of 'much cow' in an abnormal context merely *fails to confirm* the existence of the phenomenon.

Consider again the case of Hooke's Law, which defines the SPRING category. An object is a spring just in case there is a linear relationship between the force applied to it and its extension. How good this category is depends (among other things) on how well it captures a wide range of behaviors of many objects.

SPRING can only ever approximately capture object behavior. In some cases, it will capture an object's behavior well, but not perfectly, due to imperfections in the material composing the object. In other cases, it will completely fail to capture an object's behavior. For instance, an object will simply deform or break after a sufficient magnitude of force is applied rather than continue its linear extension. Of course, the fact that SPRING can only approximate behavior doesn't require us to expel SPRING from our physical theory. The category can still play a role in good explanations, including explanations of stretching and oscillation.

Observing a non-anomalous occurrence of 'much cow' in an abnormal context certainly isn't *good* news for the category [-MASS +SING NP]. Since it seems that similar observations can

be made for most (or all) *prima facie* count nouns, it seems reasonable to conclude that [-MASS +SING NP] can only ever *approximately* capture noun behavior: even if the category does a good job of capturing a noun's behavior over a wide range of cases, there will likely be some cases in which it fails.

However, as with SPRING, this conclusion doesn't require us to expel [-MASS +SING NP] from our theory. With a Category Focus, we can justify the inclusion of [-MASS +SING NP] in the theory in several ways. We can show that 'much cow' is usually anomalous, even if it isn't always; we can show that many other +SING NPs are often anomalous with 'much'; we can derive other behaviors of [-MASS +SING NP] (*e.g.* being anomalous as bare arguments) and show that these are also exhibited by many expressions.

Bibliography

Asgeirsson, Hrafn. “On the Instrumental Value of Vagueness in the Law.” *Ethics* 125, no. 2 (2015): 425–48.

Asher, Nicholas. *Lexical Meaning in Context: A Web of Words*. Cambridge University Press, 2011. https://doi.org/10.1017/CBO9780511793936.

Asher, Nicholas, and James Pustejovsky. “Word Meaning and Commonsense Metaphysics,” 2005.

Backhouse, Roger E., and Steve G. Medema. “Defining Economics: The Long Road to Acceptance of the Robbins Definition.” *Economica* 76, no. s1 (October 2009): 805–20.

Baker, Alan. “Quantitative Parsimony and Explanatory Power.” *British Journal for the Philosophy of Science* 54, no. 2 (2003): 245–59.

Bale, Alan C., and David Barner. “The Interpretation of Functional Heads: Using Comparatives to Explore the Mass/Count Distinction.” *Journal of Semantics* 26 (2009): 217–52.

Bale, Alan Clinton. “A Universal Scale of Comparison.” *Linguistics and Philosophy* 31, no. 1 (2008): 1–55.

Ball, Derek. “Semantics as Measurement.” In *The Science of Meaning: Essays on the Metatheory of Natural Language Semantics*, edited by Derek Ball and Brian Rabern. Oxford: Oxford Academic, 2018.

Barker, Chris. “The Dynamics of Vagueness.” *Linguistics and Philosophy* 25, no. 1 (2002): 1–36.

Barwise, John, and Robin Cooper. “Generalized Quantifiers and Natural Language.” *Linguistics and Philosophy* 4, no. 2 (1981): 159–219.

Beaver, David. “Accommodating Topics.” *Focus & National Language Processing* 1 (January 2004).

Becker, Gary S. “Irrational Behavior and Economic Theory.” *The Journal of Political Economy* 70, no. 1 (1962).

Berg, Jonathan. “Is Semantics Still Possible?” *Journal of Pragmatics* 34, no. 4 (2002): 349–59.

Blutner, Reinhard. “Pragmatics and the Lexicon.” In *The Handbook of Pragmatics*, edited by Laurence R. Horn and Gregory Ward. Blackwell Publishers Ltd, 2006.

Blutner, Reinhard, Petra Hendriks, and Helen de Hoop. “A New Hypothesis on Compositionality.” Sydney, 2003.

Bogen, James, and James Woodward. “Saving the Phenomena.” *Philosophical Review* 97, no. 3 (1988): 303–52.

Brigandt, Ingo. “Species Pluralism Does Not Imply Species Eliminativism.” *Philosophy of Science* 70, no. 5 (2003): 1305–16.

Bunt, H.C. “The Formal Semantics of Mass Terms.” Dissertation, University of Amsterdam, 1981.

Burge, Tyler. “Reference and Proper Names.” *Journal of Philosophy* 70, no. 14 (1973): 425–39.

Cartwright, Nancy. *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press, 1999.

Chick, Victoria, and Sheila C. Dow. “Formalism, Logic, and Reality: A Keynesian Analysis.” *Cambridge Journal of Economics* 25, no. 6 (2001): 705–21.

Clark, Stephen. “Vector Space Models of Lexical Meaning.” In *The Handbook of Contemporary Semantic Theory*, edited by Shalom Lappin and Chris Fox, 493–522. Chichester, UK: John Wiley & Sons, Ltd, 2015.

Croft, William. “Evolutionary Linguistics.” *Annual Review of Anthropology* 37 (October 2008).
Davidson, Donald. “Theories of Meaning and Learnable Languages.” In *Proceedings of the 1964 International Congress for Logic, Methodology, and Philosophy of Science*, edited by Yehoshua Bar-Hillel, 383–94. North-Holland, 1965.
———. “Truth and Meaning.” *Synthese* 17, no. 1 (1967): 304–23.
Dever, Josh. “Formal Semantics.” In *The Continuum Companion to the Philosophy of Language*, edited by Manuel García-Carpintero and Max Kölbel. Continuum International, 2012.
Dickson, Michael. “Plurality and Complementarity in Quantum Dynamics.” In *Scientific Pluralism*, edited by Stephen H. Kellert, Helen E. Longino, and C. Kenneth Waters, 42–63. Minneapolis: University of Minnesota Press, 2006.
Dow, S.C. “Pluralism in Economics.” In *The New Palgrave Dictionary of Economics*, edited by S.N. Durlauf and L.E. Blume, 2008.
Dummett, Michael A. E. “What Is a Theory of Meaning?” In *Mind and Language*, edited by Samuel Guttenplan. Oxford University Press, 1975.
Ereshefsky, Marc. “Eliminative Pluralism.” *Philosophy of Science* 59, no. 4 (1992): 671–90.
Fara, Delia Graff. “Names Are Predicates.” *Philosophical Review* 124, no. 1 (2015): 59–117.
Fodor, Jerry A. *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press, 1998.
———. “Replies to Critics.” *Mind and Language* 15, no. 2–3 (2000): 350–74.
Fodor, Jerry, and Ernie Lepore. “The Emptiness of the Lexicon: Critical Reflections on J. Pustejovsky’s the Generative Lexicon.” *Linguistic Inquiry* 29 (1998): 269–88.
———. “Why Compositionality Won’t Go Away: Reflections on Horwich’s ‘Deflationary’ Theory.” *Ratio* 14, no. 4 (2001): 350–68.
Frege, Gottlob. “On Sense and Reference.” In *Translations from the Philosophical Writings of Gottlob Frege*, translated by Peter Geach. Oxford: Blackwell, 1952.
Fuller, Lon L. “Positivism and Fidelity to Law–A Reply to Professor Hart.” *Harv. L. Rev.* 71 (1957): 630.
Geurts, Bart, and Rob van der Sandt. “Domain Restriction.” In *Focus: Linguistic, Cognitive, and Computational Perspectives*, 268–92. Cambridge University Press, 1999.
Gibbard, Allan. *Thinking How to Live*. Vol. 116. Harvard University Press, 2003.
Giere, Ronald N. “Perspectival Pluralism.” In *Scientific Pluralism*, edited by Stephen H. Kellert, Helen E. Longino, and C. Kenneth Waters, 26–41. Minneapolis: University of Minnesota Press, 2006.
Glanzberg, Michael. “Explanation and Partiality in Semantic Theory.” In *Metasemantics: New Essays on the Foundations of Meaning*, edited by Alexis Burgess and Brett Sherman, 259–92. Oxford: Oxford Academic, 2014.
Godfrey-Smith, Peter. “Abstractions, Idealizations, and Evolutionary Biology,” 266:47–56. Boston Studies in the Philosophy of Science, 2009.
———. “The Strategy of Model-Based Science.” *Biology and Philosophy* 21, no. 5 (2006): 725–40.
Grabner, Claudius, and Birte Strunk. “Pluralism in Economics: Its Critiques and Their Lessons.” *The Journal of Economic Methodology* 27, no. 4 (2020): 311–29.
Hands, D. Wade. “Caveat Emptor: Economics and Contemporary Philosophy of Science.” *Philosophy of Science* 64, no. 4 (1997): 116.
Hart, Herbert LA. “Positivism and the Separation of Law and Morals.” *L. Rev* 593 (1958): 607.
Heim, Irene, and Angelika Kratzer. *Semantics in Generative Grammar*. Blackwell, 1998.

Hogeweg, Lotte. “Optimality Theoretic Lexical Semantics,” 2010. https://api.semanticscholar.org/CorpusID:28203813.
Jeshion, Robin. “Names Not Predicates.” In *On Reference*, edited by Andrea Bianchi, 225–50. Oxford, UK: Oxford Academic, 2015.
Johnson, Kent. “On the Nature of Reverse Compositionality.” *Erkenntnis* 64, no. 1 (2006): 37–60.
Kahneman, Daniel, and Amos Tversky. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica* 47, no. 2 (1979): 263–92.
Kant, Immanuel. *Immanuel Kant: Groundwork of the Metaphysics of Morals: A German–English Edition*. Cambridge University Press, 2011.
Kellert, Stephen H., Helen E. Longino, and C. Kenneth Waters. “Introduction: The Pluralist Stance.” In *Scientific Pluralism*, edited by Stephen H. Kellert, Helen E. Longino, and C. Kenneth Waters, vii–xxix. Minneapolis: University of Minnesota Press, 2006.
Kennedy, Christopher. “Polar Opposition and the Ontology of ‘Degrees.’” *Linguistics and Philosophy* 24, no. 1 (2001): 33–70.
———. “Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives.” *Linguistics and Philosophy* 30, no. 1 (2007): 1–45.
King, Jeffrey C. “W(h)Ither Semantics!(?).” *Noûs* 52, no. 4 (2018): 772–95.
Kiss, Tibor, Francis Jeffry Pelletier, Halima Husić, and Johanna Poppek. “Issues of Mass and Count: Dealing with `Dual-Life’ Nouns.” In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, 189–98. Vancouver, Canada: Association for Computational Linguistics, 2017. https://doi.org/10.18653/v1/S17-1023.
Kripke, Saul A. “A Puzzle about Belief.” In *Meaning and Use*, edited by A. Margalit, 239--83. Reidel, 1979.
Kuhn, Thomas S. *The Structure of Scientific Revolutions*. 4th ed. Chicago: University of Chicago Press, 2012.
Levins, Richard. “The Strategy of Model Building in Population Biology.” *American Scientist* 54, no. 4 (1966): 421–31.
Lewis, David K. “General Semantics.” *Synthese* 22, no. 1–2 (1970): 18--67.
———. “Languages and Language.” In *Minnesota Studies in the Philosophy of Science*, edited by Keith Gunderson, 7:3–35. University of Minnesota Press, 1975.
Link, Godehard. “The Logical Analysis of Plurals and Mass Terms.” In *Formal Semantics*, edited by Paul Portner and Barbara Partee, 127–46. Oxford: Blackwell Publishers Ltd, 2002.
Longino, Helen. “Foregrounding the Background.” *Philosophy of Science* 83, no. 5 (2016): 647–61.
Longino, Helen E. “What Do We Measure When We Measure Aggression?” *Studies in History and Philosophy of Science* 32, no. 4 (2001): 685–704.
MacFarlane, John. “Indeterminacy as Indecision, Lecture III: Indeterminacy as Indecision.” *Journal of Philosophy* 117, no. 11/12 (2020): 643–67.
Marconi, Diego. “Wittgenstein on Contradiction and the Philosophy of Paraconsistent Logic.” *History of Philosophy Quarterly* 1, no. 3 (1984): 333–52.
Martínez-Ordaz, María del Rosario. “The Ignorance behind Inconsistency Toleration.” *Synthese* 198, no. 9 (2020): 8665–86.
Mayden, R. L. “A Hierarchy of Species Concepts: The Denouement in the Saga of the Species

Problem." In *Species: The Units of Diversity*, edited by M. F. Claridge, H. A. Dawah, and M. R. Wilson, 381–423. Chapman & Hall, 1997.
McMullin, Ernan. "Values in Science." *Zygon* 47, no. 4 (2012): 686–709.
Novick, Aaron, and W. Ford Doolittle. "'Species' without Species." *Studies in History and Philosophy of Science Part A* 87, no. C (2021): 72–80.
Pelletier, F. Jeffry. "Non-Singular Reference: Some Preliminaries." *Philosophia* 5, no. 4 (1975): 451–65.
Pelletier, Francis Jeffry. "Lexical Nouns Are Both +mass and +count, but They Are Neither +mass nor +count." In *Count and Mass Across Languages*, edited by Diane Massam, 9–26. Oxford: Oxford University Press, 2012.
Pelletier, Francis Jeffry, and L.K. Schubert. "Mass Expressions." In *Handbook of Philosophical Logic*, edited by D. Gabbay and F. Guenther, 10:1–87. Netherlands: Kluwer, 2002.
Pietroski, Paul. "Semantic Typology and Composition." In *The Science of Meaning*, edited by Derek Ball and Brian Rabern, 306–33, 2018.
Pietroski, Paul M. "The Character of Natural Language Semantics." In *Epistemology of Language*, edited by Alex Barber, 217–56. Oxford University Press, 2003.
Pinkal, Manfred. "On the Logical Structure of Comparatives." In *Natural Language and Logic*, 146–67, 1989.
Pollock, John L. "A Theory of Defeasible Reasoning." *International Journal of Intelligent Systems* 6, no. 1 (1991): 33–54. https://doi.org/10.1002/int.4550060103.
Pustejovsky, James. *The Generative Lexicon*. Cambridge, (Mass.): MIT Press, 1995.
Queiroz, Kevin de. "Species Concepts and Species Delimitation." *Systematic Biology* 56, no. 6 (2007): 879–86.
Quine, W. V. *Word and Object*. MIT: MIT Press, 1960.
Recanati, François. "Compositionality, Flexibility, and Context-Dependence." In *Oxford Handbook of Compositionality*, edited by Wolfram Hinzen, Edouard Machery, and Markus Werning, 175–91. Oxford University Press, 2012.
———. "Contextualism and Polysemy." *Dialectica* 71, no. 3 (2017): 379–97.
Richard, Mark. *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*. Cambridge University Press, 1990.
Robbins, Lionel. *An Essay on the Nature and Significance of Economic Science*. London: Macmillan, 1932.
Schiffer, Stephen. "Meaning and Formal Semantics in Generative Grammar." *Erkenntnis* 80, no. 1 (2015): 61–87.
Sendłak, Maciej. "On Quantitative and Qualitative Parsimony." *Metaphilosophy* 49, no. 1–2 (2018): 153–66.
Sennet, Adam. "An Ambiguity Test for Definite Descriptions." *Philosophical Studies* 111, no. 1 (2002): 81–95.
Sent, Esther-Mirjam. "Pluralisms in Economics." In *Minnesota Studies in the Philosophy of Science*, 80–101. 19, 2006.
Shramko, Yaroslav, and Heinrich Wansing. "Truth Values." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, n.d. https://plato.stanford.edu/archives/win2021/entries/truth-values/.
Silk, Alex. "Theories of Vagueness and Theories of Law." *Legal Theory* 25, no. 2 (2019): 132–52.

Simon, Herbert A. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*. New York, NY: Macmillan, 1947.
Soames, Scott. "Drawing the Line between Meaning and Implicature—and Relating Both to Assertion." *Noûs* 42, no. 3 (2008): 440–65.
———. "Semantics and Semantic Competence." *Philosophical Perspectives* 3 (1989): 575–96.
———. "Vagueness and the Law." In *Analytic Philosophy in America: And Other Historical and Contemporary Essays*, edited by Scott Soames, 281–98. Princeton University Press, 2014.
Speaks, Jeff. "What's Wrong with Semantic Theories Which Make No Use of Propositions?" In *New Thinking About Propositions*, edited by Jeffrey C. King, Scott Soames, and Jeff Speaks. Oxford University Press, 2014.
Szabó, Zoltán Gendler. "Compositionality." In *Stanford Encyclopedia of Philosophy*, 2008.
Teller, Paul. "'Saving the Phenomena' Today." *Philosophy of Science* 77, no. 5 (2010): 815–26.
Travis, Charles. *The Uses of Sense: Wittgenstein's Philosophy of Language*. Oxford University Press, 1989.
Vicente, Agustin. "Polysemy and Word Meaning: An Account of Lexical Meaning for Different Kinds of Content Words." *Philosophical Studies* 175, no. 4 (2018): 947–68.
Waters, C. Kenneth. "A Pluralist Interpretation of Gene-Centered Biology." In *Scientific Pluralism*, edited by Stephen H. Kellert, Helen E. Longino, and C. Kenneth Waters, 191–214. Minneapolis: University of Minnesota Press, 2006.
Weatherson, Brian. "Vagueness as Indeterminacy." In *Cuts and Clouds: Vagueness, Its Nature and Its Logic*, edited by Richard Dietz and Sebastiano Moruzzi. Oxford University Press, 2010.
Weisberg, Michael. "Three Kinds of Idealization:" *Journal of Philosophy* 104, no. 12 (2007): 639–59. https://doi.org/10.5840/jphil20071041240.
Williams, J. Robert G. *The Metaphysics of Representation*. Oxford University Press, 2019.
Williamson, Timothy. *Vagueness*. New York: Routledge, 1994.
Wittgenstein, Ludwig. *Remarks on the Foundations of Mathematics*. Vol. 7. Blackwell, 1956.
Yalcin, Seth. "Semantics and Metasemantics in the Context of Generative Grammar." In *Metasemantics: New Essays on the Foundations of Meaning*, edited by Alexis Burgess and Brett Sherman, 17–54. Oxford University Press, 2014.
———. "Semantics as Model-Based Science." In *The Science of Meaning*. Oxford: Oxford University Press, 2018.
Zwarts, Joost. "Priorities in the Production of Prepositions." In *Syntax and Semantics of Spatial P*, edited by Anna Asbury, Jakub Dotlacil, Berit Gehrke, and Rick Nouwen, 85–102. Amsterdam: John Benjamins, 2008.
Zwicky, Arnold M., and Jerrold M. Sadock. "Ambiguity Tests and How to Fail Them." In *Syntax and Semantics*, 4:1–36. Brill, 1975.