4

Negation

Contents

Our goal in this chapter is to provide a formal theory of negation *in ordinary language*, as opposed to the formal theory of negation in logic and mathematics. In order to provide for a linguistically and cognitively sound theory of negation, we argue for the introduction of a dyadic negation predicate `lack` and a force dynamic account of affirmation and negation in general. We take the linguistic horn of the dilemma first articulated by Benacerraf, 1973:

> (...) accounts of truth that treat mathematical and nonmathematical discourse in relevantly similar ways do so at the cost of leaving it unintelligible how we can have any mathematical knowledge whatsoever; whereas those which attribute to mathematical propositions the kinds of truth conditions we can clearly know to obtain, do so at the expense of failing to connect these conditions with any analysis of the sentences which shows how the assigned conditions are conditions of their *truth*.

The linguistic background is sketched in 4.1. We are equally interested in lexical semantics and the semantics of larger constructions recursively (compositionally) built from lexical elements. In 4.2, we provide a systematic survey of the negative lexical elements in `4lang`. We turn to compositional constructions in 4.3, again aiming at exhaustiveness, including many forms that involve negation only in an indirect fashion. We offer a simple, finite state formalization that embodies a more nuanced understanding of affirmation and negation, seeing these as opposing forces in the force dynamic setting (Talmy, 1988). Once the frequent cases are treated, we turn to less frequent cases that

A. Kornai, *Vector Semantics*, Cognitive Technologies, https://doi.org/10.1007/978-981-19-5607-2_4

are nevertheless often seen as diagnostic, such as double negation, discussed in 4.4, quantification and scope ambiguities in 4.5, and disjunction in 4.6.

4.1 Background

Boole, 1854, building upon thousands of years of work in the Scholastic tradition, reformulated parts of, and in important ways extended, Aristotle's logic. The structures that today bear his name, Boolean Algebras (BAs), have several features that make little sense from a linguistic standpoint, such as the commutativity of conjunction (really, *I had dinner and went home* is quite different from *I went home and had dinner*), and the basic 'Boolean' duality that stems from treating negation as a unary operation that is involutionary: $\neg\neg = id$. It is important to emphasize at the outset that what follows is a formalization of the cognitive structures underlying negation, not a critique of the standard (Boolean) negation we rely on in logic and mathematics. As we shall see, the two are very different: the economy, elegance, and tremendous usefulness of BAs came at the price of significant loss of linguistic and cognitive realism. To quote Horn, 1989:

> (...) the form and function of negative statements in ordinary language are far from simple and transparent. In particular, the absolute symmetry definable between affirmative and negative propositions in logic is not reflected by a comparable symmetry in language structure and language use. Much of the speculative, theoretical, and empirical work on negation over the last twenty-three centuries has focused on the relatively marked or complex nature of the negative statement vis-a-vis its affirmative counterpart.

In many adjectival oppositions, normally handled by some version of scalar semantics, it is very easy to pinpoint the asymmetry that Horn talks about, and assign negative value to one side of the scale unambiguously – for a summary of standard marked/unmarked diagnostic tests see Lehrer, 1985. For example, *invisible* carries overt negative marking relative to *visible*, so we conclude that conceptually it is invisible things that have no visibility, rather than visible things that lack invisibility. Yet other oppositions, such as between *full* and *empty*, offer no overt morphological cues, but are nevertheless trivial to classify, because their definition hinges on words (in this case *presence* v. *absence* of filling material) one of which is broadly synonymous to overt negatives: in this case, *absence* to *lack* or *want* (Merriam-Webster).

We will make use of the information-theoretic insight that positives, the unmarked case, are not just more frequent but, as befits a communication system, have less information content (require fewer bits). While there is no strict quantitative correspondence between frequency and the size of the code of the kind we find in artificially constructed codes (Huffman, 1952), the tendency is unmistakable in natural language and has been noted as early as Zipf, 1949.

Syntactically, the key distinction is between positive and negative polarity items (abbreviated PPI and NPI respectively), and the classification of contexts in which they

appear as positive or negative (Giannakidou, 1997). Typical NPIs are like English *ever, any, either* which occur naturally in negative contexts but not in positive ones: compare *He hasn't seen one ever* to **He has seen one ever; He hasn't seen any* to **He has seen any*; or *He hasn't seen it either* to **He has seen it either*. In other words, polarity items in English behave much like gendered items in languages like German or Polish which can appear only in the appropriately gendered context. We will not attempt to sort out the syntactic properties of polarity items here (especially as these, much like gender systems, show significant variability across languages), but will discuss their semantic import in 6.1.

4.2 Negation in the lexicon

About 12% of the 1,200 word defining vocabulary of Release V1 (S19 Appendix 4.8), 144 items altogether, involve some form of negation: *accept accident acid arrive atom bad bar behind bend black block building burn calm catch chance child clean close coal continue continuous cover curve dark dead destroy different dry eager easy elephant end fail finish firm first flat free full gas gradual green hang hard hide ill instead jump laugh leave light limit long lose mean middle must narrow natural necessary need negative new night no nothing object off offensive one only open opinion oppose out park permanent plant police practice preserve prison private protect public quiet reach remove rest right romantic rough rubber rude sad safe same send separate serious sharp short simple sincere single sleep slope smoke smooth soft solid sometimes special steady steal stiff stop straight strange stupid success sudden sure surprise take tent thick thin tie tight together twist unless waste water weak without wrong*. This list is actually a bit shorter (139 elements), because in the 144 we count with multiplicity elements that are homophonic in English, such as *thin* as in *thin paint* `hi1g liquidus rzadki 1038` `thin/1038` `flow(er_ gen)` and *thin* as in *thin reed* `ve1kony tenuis chudy 2598 gen` `thin/2598` `er_ {distance between surface}`. (Since we avoid spurious duplication of entries for metaphorical senses, treating e.g. *acid* in *vinegar is an acid* and in *an unnecessarily acid remark* by one and the same lexical item, disambiguation is rarely called for.) In Release V2 (see the Appendix) the proportion of negatives is even lower, 8%.

The list has many elements such as *water* which seem to lack any negative aspect. `water` But a closer look at the definition `vi1z aqua woda 2622 u N liquid, lack colour, lack taste, lack smell, life need` shows how negative statements enter the picture. (Recall from 1.5 that in `4lang` dyadic predicates are given in infix notation (SVO order), so `life need` means that the subject of need is life, and the object is the definiendum, whereas `lack taste` means that the object of lack is taste, and the subject is the definiendum.) The example already shows how our central innovation, the dyadic negation predicate `lack` works, and the definition `hia1nyzik` `lack` `desum brak 3306 p V =agt lack =pat` show it to be irreducible (primitive). As we shall see, `lack` is sufficient: we will not require other primitives, not even an unanalyzed unary `no`.

In many cases like *dirty* or *blind* the lexical entry carries a negative (prejudicial) sentiment, but not all of these are amenable to an analysis that contains a negative. Every analysis of blindness invokes a logical negative: 'sightless' (Merriam-Webster) 'unable to see' (Longman), etc. Within the bounds of our defining vocabulary, we can write this as `lack sight`. The critical observation here is that lack signifies the absence of a default: people (generic individuals) are sighted, which is the unmarked (default) case, but *blind* contains lexical prespecification overriding this default. Returning to *dirty*, which at first sight is defined as 'not clean'; and to *clean*, definable as 'not dirty', in terms of lack it is obviously *clean* that needs to override the default of things, in their natural state, being somewhat dirty, whereas *dirty* is definable in terms of dirt, mud, dust, soil, etc. just as *sight* is definable without recourse to negation as a form of perception that relies on eyes.

`good`
`want`
`bad`
`left`
`right`
`same`
`different`
`self`

The same treatment can be effortlessly extended to many antonym pairs, e.g. defining *good* as `gen want`, i.e. the object of `want`, where `want` is given as `=agt feel {=agt need =pat}`. Given a positive definition of `good`, we can define *bad* as `lack good`. Antonyms such as *left/right* make clear that `lack` is in some sense the dual of `has`: *left* is `side, has heart` and *right* 'dextra' is `side, lack heart`. Similarly, *same* can be defined as `lack different` but *different* need not be defined as 'lack same', we have a better definition relying on Leibniz' Principle The Identity of Indiscernibles `=pat has quality, =agt lack quality, "from _" mark_ =pat`. We also rely on this principle in the definition of the pronoun *self* which we take to be `=pat[=agt], =agt[=pat]` without any recourse to negation. In all such cases, it is really a matter of lexicographic taste whether we choose to mark antonimy on both members or just one: *invisible* means lack of visibility, and we could redundantly mark *visible* as lacking in invisibility, but we see no compelling reason to do so. Indeed, by omitting these antonymy clauses from the unmarked members of the antonymic pairs, the list we started with can be reduced considerably, a process we carried through for Release V2. Remarkably, we don't have a single example of irreducible antonymy, where both definitions would have to refer to the opposing element.

There is of course an entire class of lexical items whose primary function is to negate: the words *no, not*, the clitic *n't*, the prefixes *un-, im-, de-, non-, anti-* and the like. Ideally, we wish to represent these by a unary negation operator, provisionally written as `no`. This brings into sharp focus the issue of double negation, a matter we will discuss more fully in 4.4, but illustrate here on a contender for the title of longest English word.

Establishmentarianism is the 'movement or ideology advocating the principle of an established Church with special rights, status, and support granted by the state', an issue most people never heard of and most likely stand neutral on. *Disestablishmentarianism* is the directly opposed 'movement or ideology advocating the withdrawal of special rights, status, and support granted an established church by a state', and *antidisestablishmentarianism* is of course the movement or ideology directly opposed to this. People who prefer the status quo will likely be antidisestablishmentarian, but not establishmen-

tarian, since neither of these movements/ideologies would be content to leave things as they are.

A shorter and more common, but conceptually not any easier, case is provided by *open* versus *close* (shut). Unlike in topology, where close/open have such specialized meanings that sets can satisfy both predicates at the same time, in ordinary language no ordinary object can be clopen. Yet a third state of affairs, where the status of an object is not known, exists, just as in topology, where a set need be neither closed nor open. Tertium datur. We will denote this third state by $\circledcirc$, and use $\oplus$ and $\ominus$ to denote the positive and the negative states.

If we don't insist on lexical semantics, compositional cases, which we will treat in more detail in Section 4.3, offer much simpler examples of double negation failure. Consider *up* and *down*. Let's say we are at a construction site, perhaps standing on a ladder, and receive the instruction *move up!* which we want to defy. This can be achieved not just by moving down, but also by moving sideways, or by not moving at all. All three of these acts will conform to the negated command *don't move up*. *Don't move* or *rest* are contrary to *move*, and *move down* is contrary to *move up*, but these simply don't exhaust the entire space of possibilities, which also contains moving sideways, an action contrary to *rest, move up*, and *move down* alike. In the terms of philosophical logic, natural language negation creates *contraries* rather than *contradictionaries*. Thus, the classical Boole/De Morgan picture where negation satisfies the involution law is simply not tenable for natural language – we present our own solution in Section 4.3, and return to double negation in Section 4.4.

4.3 Negation in compositional constructions

From our perspective, the traditional Square of Opposition (Parsons, 2017) is inhomogeneous. "A" statements of the form *every s is p* are simply written `p(s)` or `s is_a p` (the two styles of writing are just syntactic variants). But a word of caution is in order: these formulas are not aimed at the logical sense of *every* ($\forall$), but rather at the everyday sense, which admits exceptions (Moltmann, 1995; Lappin, 1996). Also, such formulas typically appear in the translation of restrictive modifier clauses, where they have existential, rather than universal import.

For example, when we say in naive physics (Hayes, 1978) that *atoms* are small particles that have nuclear energy (never mind how well this definition fits modern physics, our target is ordinary language), the definiens is formulated as `small, particle, atom has nuclear(energy)`, and here `nuclear(energy)` doesn't embody the claim, not even in naive physics, that all energy is nuclear. Only the much narrower claim, that the energy that atoms have is nuclear, is part of the definition. In this respect, generic is_a is closer to "I" statements of the form *some s is p*.

Of particular interest here is the style of default inference supported: if energy is provided by atoms, that energy is nuclear, if a cane is owned by a blind person, that cane is white, and so forth. This is indeed in opposition to "E" statements *no s is p*

whose central goal is to block similar inferences: persons have organs, these organs are typically functioning, so persons can walk, talk, see, etc. – this all goes without saying. The inferences are highly automatic/preconscious, yet we rely on such inferences in the process of making sense of natural language utterances all the time.

`blind` Clearly, the raison d'être of the word *blind* is to guarantee that some of these inferences are blocked, hence our definition `lack sight`. Further, this prohibition on the inference is absolute, we treat a blind person with a black cane as unusual, exceptional, out of the ordinary, but reality overrides the default, whereas we treat a blind person that can see as paradoxical, impossible, and our best interpretation strategy upon encountering a situation like this is to say that the person was not really blind, that this has something to do with some technical definition 'legally blind' rather than the everyday meaning of blindness.

Finally, "O" statements, *some s is not p* mean lack of implication from s to p, a view equally compatible with Aristotle's original formulation *not every s is p*, which need not carry the existential implicature that many take for granted in the analysis of *some*. This becomes a bit clearer if we take into account the Aristotelian view that the predicate inheres in the subject: there is no difference, other than surface form, between *Joe is fat* and *Joe has fatness* or *Joe fat(ten)ed*. Whether the predicate is expressed adjectivally, nominally, or verbally has no bearing on its relation to the subject, which is one of subsumption. On this view, O forms are simply `s no p` which leaves it ambiguous between `s isa no p` (adjectival/nominal form using the copula) and `s (no p)` (overtly negated verb). To make the type theory work out, in 2.1 we assumed a broad type of *matter*s, which are neutral between things (ordinary nominals), action nominals, events, actions (verbal elements) and properties (adjectival elements). English verb-nouns such as *divorce* furnish a rich class of surface examples.

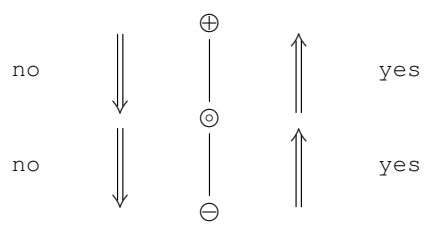

Fig. 4.1: Forces in negation and affirmation

The outstanding issue is explaining why unary `no` is absolute while binary `lack` is generic. `lack` signifies that the predicate in question does not inhere in the subject. What does `no` signify? It is at this point that the information-theoretic view comes to the fore. By the logic of compressibility, `no` must be adding some extra information, but this is not simply negating the statement, as the Boolean solution would have it, but rather *applying a force to make it negative*. As in naive physics (Hayes, 1979) we assume that

matters have three basic states, positive, zero (default, resting state), and negative: we will depict this in a three-state finite automaton arranged top to bottom as in Figure 4.1.

A word of caution is in order: while finite state automata of the sort depicted here are capable of limited counting, e.g. `no yes yes no no` would move the current state from the initial ⊚ to ⊖, this really doesn't correspond to anything in natural language. Motion, both ordinary physical motion of objects and more general 'movements' or 'processes' provide another example of the same tripartite characterization that we have seen in Fig. 4.1, this time with *start, steady*, and *stop* states.

To see how the state transitions actually work, and to refine the picture to include not just negation but also affirmation, we analyze some ordinary language expressions here. We start with imperatives, both because these are a major source of negatives and because they justify some of the key features of our model. Consider the negatives *Don't smoke!* or *No smoking* and their paired affirmatives *Smoke!* or *#Smoking*.

Normally, locations are unspecified for smoking/nonsmoking, though there are many places where the default is nonsmoking and some where the default is still smoking. A sign that simply says *No smoking* has the same force as one with an overt deontic operator *Smoking prohibited*. The opposite of this is a sign *smoking (permitted)*, and not *#smoking mandatory* which would carry a much stronger affirmation of smoking. This is not because we don't find obligatory rules, there are many from *seatbelts mandatory* to *you must agree to our privacy policy first*, but rather because we find smoking increasingly restricted to special settings like dedicated smoking rooms at airports.

Returning for a moment to our starting example, it is clear, even if we don't take overt morphological marking into consideration, that the normal (default) state of things is to be visible, and invisibility, to the extent it exists, is the marked case. The primary goal of prohibitions is to designate their object as abnormal. Consider *You shall not kill*. Biblical Hebrew (and English at the time of King James) made no distinction between imperative and future negative, the normative effect (of an ideally kept command) is that in the future there is simply no killing (*retzach*). In `4lang` we can write this as `after(gen lack kill).`

As we have briefly discussed in 3.4, we frequently encounter antonyms that fit well with the tripartite picture of Fig. 4.1: *heavy* means `weight(er_ gen)` and `light` means `weight(gen er_ )`. Since the generic will unify with the subject, the effect that (Parsons, 1970) illustrates with the example of *enormous flea*, that such a flea is still rather small, is easily explained: such a flea has size much larger than `gen`, but this automatically refers to a generic flea, not any generic object.

```
heavy
light/1381
```

Returning to our theory of *You shall not kill*, `gen` is the same proquant that we use elsewhere to denote a non-specific entity. After the utterance of the command who does no killing? Somebody. Everybody. People. Recipients of the command. It is precisely the generic nature of the subject that guarantees the universal import of the prohibition. This gives an answer to the question we raised above: we will not need a unary negation operator *no* since it can be defined as `gen lack`. From the logic point of view `gen` may appear mysterious, but from the vector semantic view it is simply the row vector

```
no
gen
```

(bra) $\langle 1/n, \ldots 1/n|$ which has the same nonnegative constant $1/n$ at each coordinate, where n is the dimension of the linguistic subspace L (see 2.3 for the notation). `gen` may be a primitive in the `4lang` dictionary, not reducible to other words, but it is not a primitive in vector semantics, in fact it is one of the elements most easily defined.

The picture of negation that emerges from these considerations is very nontraditional. From the mathematical side we have already seen that it requires an intuitionistic protologic of some sort, since it admits at least one extra value in addition to true and false. And from the linguistic side, instead of the standard, unary negation operation `no` that is analogous to Boolean $\neg$, we have a dyadic operation `lack` that signifies that its first argument does not have some defaults normally associated to it, with the second argument determining which default gets overridden. For example, persons are assumed to have fully functioning organs (in fact, this assumption is held for all living beings, and is inherited to persons via animals) so `person, lack sight` defeases an entire chain of inferences whereby `eye is_a organ` and `living_being has organ(working)` that would normally lead us to believe that persons have working eyes i.e. they are sighted. Compositional `no` is derived as `gen lack`, the unary negation operator is formed by quantifying over the first argument of the dyadic `lack`. Since `gen` is a fixed vector, and `lack` is a fixed matrix, the unary `no` operator is simply the vector obtained from postmultiplying `gen` by `lack`. In particular, this is not a matrix like $-I$ (with -1 at each diagonal element and zeros elsewhere), it's just a vector.

How the (primitive) dyadic negation operator `lack` and the (derived) unary `no` interact with auxiliaries, main verbs, adjectives, and adverbials is a complex matter. We can't possibly do justice to the syntax of negation in this book, especially as this changes from language to language. But the semantics is constant, and is simple enough to derive some major conclusions that appear to have syntactic import as well. In particular, note that `lack` is a sparse matrix with only a few dozen prespecified elements in the uroboros set given by equations such as *night* $\stackrel{d}{=}$ `period, follow sunset, sunrise follow, dark, lack sun, <sleep at>` or *opinion* $\stackrel{d}{=}$ `thought, person has, person[confident], person lack proof`. The simplest of these contain just one clause: *public* $\stackrel{d}{=}$ `lack owner` or *lose* $\stackrel{d}{=}$ `after(lack)`.

night
opinion
public
lose

We call attention to the fact that `lack` is not a pure negation operator (scalar product of a vector with -1) but rather the subtraction of something that is normally there. A person normally has intact bodily functions, including sight, so *blind man* is perfectly normal, while *#blind stone* is markedly infelicitous, just as *#sighted stone* would be. In geometric terms, negation is better explained as partial complementation, i.e. a subset of the complement, not the entire complement. We return to this matter in 7.2, where we discuss scalar semantics.

4.4 Double negation

In general, double negation is out. Negative imperatives are easy (in English, they require *do*-support, but this is exceptional), from *go!* it is easy to form *don't go!* with the intended meaning *stay!*. But double negatives *???don't don't go* are hard to produce, people tend to express the intended meaning by *don't stay*. A British National Corpus (BNC) search reveals 40 examples of *don't don't*, all in live conversation (as opposed to writing), and all with the meaning 'emphatically don't' as in *Charlotte please don't don't go noisy* or *Don't don't you think that there's a conflict of interest there*. This is from a total of 92,334 *don't*s in the corpus. The asymmetry is not restricted to imperatives: consider a grocery store with a sign *no bananas (today)*. Once the shipment arrives, they will not advertise *???no no bananas*. To quote De Mey, 1972:

> 'Natural' negation only involves objects or elements a speaker or listener is attending to . . . It makes no sense to instruct a listener to suppress a thought he is not considering or an idea he is not having.

The only standard case of double negation is when the first negative is syntactic and the second morphological: *a not unhappy person, a not unfriendly letter, . . .* (see Horn, 1989 5.1.3). What is remarkable about such cases is that they are no longer about the negation of some default: there is no assumption that people are generically happy or letters are friendly. It is the unhappiness of a person that is being negated here, an idea that we couldn't reasonably assume to have already been present in the listener's mind as a default assumption. Rather, it is the compositional meaning `person is_a unhappy` that gets negated in its entirety. We conclude that `no`, as a syntactic operator, negates the main predicate, so from aRb we obtain $a(\neg R)b$ by the corresponding compositional semantic rule. We assume, without argumentation, a rule-to-rule hypothesis (Montague, 1970; Bach, 1977; Gazdar et al., 1985) between rules of compositional syntax and semantics.

In this case, the negation of the predicate is easy: both `¬is_a` and `¬has` can simply be taken as `lack`, so we obtain `person lack unhappy`. To negate *John ate fish* we need to invoke some form of *do*-support on the syntactic side to obtain *No, John didn't eat fish*. Note that the main predicate *John ¬eat fish* is coordinated with *No*: to obtain the desired result that this is a singly negated statement about eating we take $\neg X$ to be headed by $\neg$ rather than by X. Since our meaning representations can't have nodes with multiplicity (without the use of the `other` operator), the sentence-initial `no` is unified with the `no` of `no eat`, and we obtain `John no eat fish`. Returning to `person lack unhappy`, we can accept this as is, or proceed syntactically from *not (unhappy person)* or from *(not unhappy) person*. We investigate both possibilities.

Since standard tests of constituency (Wells, 1947) support the second analysis, we start with *not unhappy* and substitute, salva veritate, the definition of unhappy, to obtain `no (gen lack happy)`. As we have seen, the syntactic negation operator affects the main predicate, in this case `lack`. A suitable candidate for `¬lack` will be `has`,

which means 'doesn't lack' after all. This way, we obtain `gen has happy` which, when applied to `person`, will yield the desired `person has happi(ness)`.

In the other analysis, we start with *unhappy person* with the semantics `person is_a unhappy`. Again substituting salva veritate, we obtain `person is_a gen lack happy`. Here `person` can unify with `gen` and to yield the more specific `person`, and similarly `is_a` can unify with `lack` to yield `lack`, so altogether we have `person lack happy`, a very reasonable semantic representation of *unhappy person*. Negating this by the syntactic `no` again amounts to negating the main predicate, so we obtain `person has happy` as before, irrespective of the constituent structure we started with.

When both `no`s in a double negation are compositional, the above analysis would yield `gen lack gen lack` which, without special pleading, will simply reduce to `gen lack` i.e. to single negation, a result we are not unhappy with, given the absence of real-life examples suggesting otherwise. For the better attested *Don't you ever NOT clean up after yourself!* we can invoke extra rules, e.g. that the contrastive stress actually keeps the second negation distinct from the first, and indeed, such sentences sound natural only with contrastive stress/intonation.

4.5 Quantifiers

Following Frege, 1879 and Russell, 1905 the treatment of a restricted class of lexical elements, quantifiers, has become virtually inseparable from the treatment of negation. In this regard, our treatment is a considered return from Montague, 1973 and subsequent work to the earlier tradition, whose last significant exponent was Peirce (Böttner, 2001). While Montague Grammar eventually treated nominals as generalized quantifiers (Gärdenfors, 2007; Badia, 2009), we move in the other direction, and treat quantifiers as nominals whose compositional behavior is largely dictated by their semantic content, rather than as special term-binding operators. In doing this, we make purposely very little distinction between an individual fox, the species *Vulpes vulpes,* the set of foxes in the world, or the class of potential foxes in all possible worlds.

That some kind of quantificational ur-element is needed is already clear from a closer look of our definition of *good* as the object of `want`. To write out the definiens in infix (SVO) order, it is not enough to write `want good`, for this would be interpreted as the definiendum filling the subject slot, saying in effect *(the) good wants (the) good*, or worse yet, *(the) good wants itself*. Since the intended meaning is that *good* is what people want (a consensus theory of value), who is the subject, one person, an exemplary and perhaps even God-like person, or just anybody? We will use the same generic `gen` that we used in 4.3 to fill the subject slot, but caution the reader that this element doesn't have universal import, it's just a placeholder that 'plugs up' the valence. The closest overt element in English with roughly the same meaning and distribution is *one* used generically, as in *One should take an umbrella if the sky is cloudy*, but we use `gen` so as to avoid confusion with numerical `one`. Unlike `one` whose semantics clearly involves

the singular, gen, being at the top of the subsumption hierarchy, will unify with any x. Whereas one, book means a single book, gen, book is simply book, and we leave it open whether this means an arbitrary book, the set (or class) of all (actual of potential) books, or some abstract notion of 'bookness' as in *the book of nature*.

Lexicalized quantifiers either in their base form *some, any, no, ...* or in a subtyped form *someone, somebody, something, somewhere, somehow, anyone, anybody, anything, anywhere, anyhow, noone/no-one, nobody, nothing, nowhere, ...* will be treated on a par with pronouns, including interrogatives, as members of a new lexical category *proquant*, whose crosslinguistic coherence (but not the name proquant) is argued for by Szabolcsi, 2015. Many, if not most, of the proquants are either lexical primitives, or have a compositional analysis that directly relies on abstract primitives such as the wh morpheme responsible for interrogatives. Here our focus is on overtly negated elements such as *nobody*, and the main question is whether these require a unary negation operator no.

One area where the standard theory appears vastly superior to the one presented here is assigning semantics to obviously compositional quantifier structures such as *at most seven, no more than ten*. But this is accomplished at the price of sweeping under the rug the fundamental problem we started out with, assigning semantics to the atomic units. What is the semantics of *seven*? The dictionary suggests 'the number 7', but this is not exactly helpful, since '7' is left undefined.

Could we actually use here the standard mathematical semantics that rests on the Peano axioms? The requisite formulas $\leqslant 7, \neg(> 10)$ seem to capture the intended meaning quite nicely, and the task of assembling them in a rule-to-rule fashion appears feasible. Yet the same approach is notoriously problematic for common 'fuzzy' cases like *at least a few, some, many/much* .... A more subtle problem is posed by overgeneration: the standard semantics smoothly extends to zero and negative integers, yet expressions like *at most minus one* are hard to interpret by ordinary speakers, and the more math we apply the clumsier the corresponding natural language expressions become. Do we have to translate *greater than i* as denoting the complex plane with the unit disk removed? If so, why don't we assign this as the meaning for *greater than 1* as well? If not, how do we account for expressions like *greater than z*, with z any complex number, which are perfectly common and ordinary in complex function theory?

Altogether, the standard logical approach is inappropriate for handling what little overlap there is between the semantics of logical and natural language expressions. It offers spurious precision, not just in the handling of 'fuzzy' quantifiers but also for any number above the magical number 7 ± 2 (Miller, 1956). Since the standard theory was developed in order to overcome the well-known limits of human numerosity (Dehaene, 1997), it is incapable, by design, of accounting for these limits. A fuller discussion would go beyond the scope of this book, but a step in the right direction is already taken in Gordon and Hobbs, 2017, who restrict Peano arithmetic to the metatheory, and concentrate on the cognitively relevant structures like 'half orders of magnitude'.

Using this notion, we can assign meaning to lexically complex quantifiers such as *somewhat* in constructions such as *It will be somewhat warm(er)* which we take to mean

'it will be perceptibly warm(er)' where *perceptibly* means 'by half order of magnitude'. Since this is arguably an adverbial meaning, we will concentrate here more on the proquants, where *some-* has a pure existential import. Deriving the lexical meaning of quantifiers is made easier by the fact that in most languages they share a sortal type with pronouns, so we will have interrogatives *who, what, where, when, ...* and follow the same typing *everyone/anyone/someone/noone, everything/anything/something/nothing, everywhere/anywhere/somewhere/nowhere, everytime/anytime/sometime/never*.

The sortal types are quite transparent: *who* requires a `person`, normally spelled out in English as *one*; *what* requires a `thing`; *where* requires a `place`, spelled in these proquants as *where* but historically *ere* (also seen in *here, there*); *when* requires a `time`; and *how* requires a proadverbial, spelled variously as *how (anyhow, somehow)* or as *way (anyway, someway, no way/nohow)*. Another suppletive form is *never*, with *no+ever* used interchangeably with *no+time*.

who
what
where
when
how
manner

As standard (Katz and Postal, 1964; Langacker, 2001), we analyze *who* as `person, wh`; *what* as `thing, wh`; *where* as `at, wh`; *when* as `time, wh`; and *how* as `manner, wh`, where `manner` is `quality, do has`. By taking *some-* to mean `exist`, arguably a primitive, we obtain for *someone* the definition `exist, person` and similarly for *something, somewhere, sometime, somehow*. We take *every-* to be synonymous with `gen`, and again use the conjunctive combinations `gen, place` to define *everywhere*; `gen, manner/1706` to define *everyway*, etc.

any

In systems of Knowledge Representation (KR) such as Cyc (Lenat and Guha, 1990) it is common to distinguish individuals, e.g. some particular poet, say Allen Ginsberg, from the class Poet, of which Ginsberg is an InstanceOf. The semantics of *any-*, however conceived, will have to express the choosing of one particular instance from a class, the central element of the meaning being that it doesn't matter which instance (Kadmon and Landman (1993) call this the 'free choice' reading of *any*). Here we take advantage of the mechanism that we have at our disposal independent of negation and quantification, thematic roles (Dowty, 1986) and the fact that we already have a fundamental `is_a` relation in the system. With this, we can define *any* as `<one>, =agt is_a`. We note that optionality (the use of defaults, see 6.4) is another feature of the system that has broad justification already on the quantifier-free fragment (Reiter and Criscuolo, 1983). When we say *any poet* this will mean any (one) x such that *x is_a poet*, and it is the same semantics that we apply to *anyone, anything, anywhere, ...*.

With the other proquantal roots out of the way, we can turn to our central subject matter here, the semantics of *noone, nothing, nowhere, ...*. This requires no special effort, in that *no-* is already defined as `gen lack` and the sortal types just unify with `gen`, leading to `person lack` for *noone*; `thing lack` for *nothing*; etc. Thus *noone slept* is simply `person lack sleep`, and the key scope effect, that this really means 'nobody *among the people relevant in this context* slept' is obtained by reading `person` in this manner. Unlike the Generative Semantics tradition, where this scope restriction is obtained via tracing the scope of (typically covert) high-level speech act operators that act indexically (Lakoff, 1970; Kaplan, 1978), here we take the genericity as basic and

find, to the very limited extent one can (Kornai, 2010b), episodic readings by special effort. In this regard, our system is closer to the database logics that rely on a locally closed world assumption (Doherty, Lukaszewicz, and Szalas, 2000) than to classic Montague Grammar.

Compare *Everyone on Cormorant Island speaks two languages* to *Two languages are spoken by everyone on Cormorant Island*. There is a sense that the active sentence does not require these to be the same two languages for everyone, whereas the passive sentence does. But how strong is this sense? Early generative theory (Katz and Postal, 1964) assumed that both readings are available for both sentences. This left explaining which reading is preferred in which context to factors that go beyond syntax and semantics such as communicative dynamism (Firbas, 1971), as there is a similarly strong sense that the active sentence is about the inhabitants of Cormorant Island while the passive is about two languages. Also, it is worth keeping in mind that the entire phenomenon is somewhat marginal. The ratio of passives to actives is somewhere between 4% and 18% depending on genre (Givón, 1979), e.g. the BNC has 662 instances of *killed by* compared to 4407 instances of *kill*. Quantifier phrases (nearly 70k examples in the BNC) will appear in the *by-* phrase only in about 1.5% of the cases.

In `4lang` the active sentence means `person in Cormorant, person speak language(two)` (recall that the two instances of `person` that appear in the linearly rendered formula are automatically unified). The passive sentence means `language(two) is_spoken_by person in Cormorant Island`. It is unclear whether these become the exact same thing as soon as we acknowledge a lexical redundancy rule (Bresnan, 1982) that relates active V to passive *is V-ed by*: there are surprisingly many design choices even within LFG where the idea that the active/passive relation is to be captured in the lexicon is taken for granted (Genabith and Crouch, 1999).

Here we consider, very briefly, the other proquants. *Anyone on Cormorant Island speaks two languages* versus *Two languages are spoken by anyone on Cormorant Island* has the same level of uncertainty in regards to judgments of grammaticality and readings as the *everyone* examples we started out with. To avoid bracketing, we will write `Cormorant_Islander` or just `C_I` for `person in Cormorant Island`. With this abbreviation the active sentence can be paraphrased as `C_I speak language(two)` and `lg(two) is_spoken_by C_I` and again the outcome depends on the status of the redundancy rule (or in other generative treatments, the transformation) that relates actives to passives. *Someone* does not bring in the same ambiguity problem, since `exist C_I speak language(two)` is implicationally equivalent to `lg(two) is_spoken_by C_I, exist C_I`, no matter how we handle active/passive.

Finally, let us consider the examples most relevant to our subject matter, negated universals or "E" statements. Clearly, *Noone on Cormorant Island speaks two languages* means `C_I lack speak language(two)` and this is subject to the downward entailment issues that smart alecs often play on: ... *but Joe here speaks seven!* More important, we see `lack` as negating a non-default proposition, as in the double negation

cases discussed in 4.4, indicating that the mechanism we proposed there is available for these cases as well.

As for "E" passives, we get `lg(two) is_spoken_by lack C_I` which says, in a somewhat clumsy fashion 'among the people who speak two languages we don't find Cormorant Islanders'. This offers the same episodic reading as the active, and is subject to the same downward entailment problem. Note, however, that the phenomenon is even more marginal: *by noone/nobody* phrases are just 0.1% of the total occurrences of *noone/nobody* in the BNC, for a total of 8 sentences among over ten million. One would really have to be superbly confident about having already captured 99.9999% of English grammar before seeing these as a descriptive challenge.

4.6 Disjunction

In BAs, De Morgan's Laws connect conjunction to disjunction in a perfectly symmetrical fashion. But in natural language semantics conjunction is the default operation: unless some other particle is present we interpret phrases and clauses conjunctively. In case of proper nouns, we treat the conjunct as a collective (Scha, 1981). Given that negation is a marked operation, there is no way to follow the BA technique and reduce disjunction to conjunction. In fact, `no (A and B)` ends up negating the head predicate, so we get `A ¬and B`. This is tantamount to the well-known deontic paradox: *No food and drink* is actually obeyed by a person who only brings food but no drink. The obverse of this, Ross's Paradox (Ross, 1941) brings in the same concerns.

It is fair to say, then, that our interest is with a positive, rather than a De Morgan-style definition of disjunction. While we take the rather unsurprising route that `or` is a primitive, not at all reducible to `and` and `no`, let alone to `and` and `lack`, there is more to disjunction than 'well, it's a primitive'. The cognitive import of `or` is clearly to keep both disjuncts open, whereas in conjunction a higher (collective) node is formed and the conjuncts themselves are no longer active. We define *or* by `"_ or _" mark_ choose`, but note that it is unconnected to the broader system: since not one `4lang` definition contains `or`, it is eliminable from the uroboros set.

`or`

A systematic study of `or` in the larger lexical domain must await later releases, but it is worth noting that almost 40% of LDOCE (Bullon, 2003) definitions uses this word. Most of these seem trivial from the `4lang` perspective: for example *abandon* (V) is defined as 'to stop having a particular idea, belief, or attitude'. Here we could simply add much-needed generality by saying 'stop having =pat', for this would add back more literal senses of *abandon* like 'to leave someone, especially someone you are responsible for'; 'to go away from a place, vehicle etc permanently, especially because the situation makes it impossible for you to stay'; and 'to stop doing something because there are too many problems and it is impossible to continue' as in *The soldiers abandoned the battlefield/their weapons*.

Manually checking over 32,800 cases (the 1st edition of LDOCE (Procter, 1978), is far more frugal in this regard) is beyond our powers, but the number of cases seems

large enough to justify a deeper study. What seems clear is that eschewing *or* leads to a less precise description of synonyms and of selectional restrictions, as in in *zonked* 'very tired or suffering from the effects of drugs or alcohol, so that you do not want to do anything'. The vectorial perspective may make this considerably easier: to the extent we define conjunction by intersection of polytopes, we may be able to define disjunction by their union, but we must leave this at the conjectural level for now. This may work well even if negation does not correspond to complementation, as we have argued above.

At sentence level, *or* signifies either a future choice to be made, or a past, unknown, choice. This makes `or` more closely related to exclusive or (xor) than to standard Boolean $\vee$, though this is often hard to discern since the alternatives are disjoint to begin with. Further, while natural language `and` must involve incrementing the time index on successive verbal conjuncts (cf. the example we started out with, *I went home and had dinner*), `or` has no temporal update associated to it, which again highlights the lack of duality between these two. Another diagnostic pointing at the same conclusion is the clear ability of *or* to introduce alternatives that are counterfactual: *It can wait, or they would have called us by now*.

There is no question that the proposal made here sacrifices quite a bit on the mathematics side: conjunction is not commutative, Boolean duality is gone, and there are many ripple effects through the entire system we haven't even discussed, e.g. that existential quantification no longer amounts to infinite disjunction. But the gains on the linguistic side are considerable: we have a formal theory of word meaning whereby we can assign semantics to morphological operations in a manner that smoothly extends to compositional semantics.

`4lang` captures well the key observation that negation is not an involution, and in general offers translations whose processing difficulty correlates inversely with their frequency. Clearly, the theory is a better fit with the classical Knowledge Representation tradition (Brachman and Levesque, 1985; Brachman and Levesque, 2004) and with database logic than with the first- and higher-order (intensional) calculi familiar from MG and related theories. We do not see this as a loss, especially not from the learnability perspective, a matter we shall return to in 5.3.

We started this chapter with Benacerraf's observation that sentences in natural language and in mathematics are different enough to merit separate semantic frameworks. Were this not so, it would actually be hard to explain why Boolean Algebra, and modern logical calculi in general, took so long to develop from Aristotle's logic. The approach presented here, in many ways a considered return to a more Aristotelian perspective, is not an attempt to 'reform' standard mathematical logic, which we consider to be the correct theory of the domain. Rather, our goal is to build, with the same care, a formal theory of natural language semantics, even at the price of finding this theory insufficient in the mathematical domain the same way as e.g. in the measurement domain discussed in 3.4.