Dyadic negation in natural language

András Kornai

SZTAKI Computer and Automation Research Institute

H-1111 Budapest Kende u 13-17, Hungary

andras@kornai.com

Abstract

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.

1 Introduction

The transformational theory of negation, starting with Klima, 1964, is intent on capturing the semantics of negation by a unary negation operator often paraphrased as *It is not the case that S*. For Klima, *neg* is an (optional) daughter of S, and the occurrence of *neg* in surface configurations is to be derived by transformations. In this theory, deep structure is conceived of as 'the language of thought' (Fodor, 1975; Fodor, 2008) composed of discrete, symbolic units arranged in a tree structure that reflects the semantics, in particular the scope relations, in ways that may not be fully transparent at surface level. In fact, Chomsky, 1966 argued that this conceptualization, including the idea that the deep structure is language-independent, goes back at least to the Port-Royal Grammar (Arnauld and Lancelot, 1975 [1660]):

> Arnauld observes (p. 208; PRL 160) that the sentence *There are few pastors nowadays ready to give their lives for their sheep*, though superficially affirmative in form, actually contains implicitly the negative sentence "Many pastors nowadays are not ready to give their lives for their sheep." In general, he points out repeatedly that what is affirmative or negative "in appearance" may or may not be in meaning, that is, in deep structure. In short, the real "logical form" of a sentence may be quite different from its surface grammatical form. The identity of deep structure underlying a variety of surface forms in different languages is frequently stressed (Chomsky 1966:87).

While the transformational analysis gradually came to include a variety of non-sentential negatives, (see Jackendoff (1969) for an intermediate, and den Dikken (2019) for a well worked out modern version), the primary focus of this theory is still on logical negation at the sentence level, at the expense of key pragmatic aspects such as presuppositions and conversational/conventional implicatures.

In this paper, we retain the transformational generative grammarian's identification of semantic representation with underlying form, a method most popular in Generative Semantics (Goldsmith and Huck, 2013), but one that actually goes back to Pāṇini. However, we avoid calling it Logical Form (LF), and speak simply of semantic representation for two reasons. First, because the term 'LF' is increasingly used in a narrow sense to mean formulas of Montague's Intensional Logic (Montague, 1973) and related calculi, whereas the work presented here relies on a different system of Relevant Logic (Meyer and Martin, 1986). Second, because the representation format we use, the Resource Description Framework (RDF, Beckett, 2004) has its roots in a different tradition, that of Knowledge Representation (King, 1979; Sowa, 2000) and conceptual graphs (Quillian, 1967; Minsky, 1975; Sowa, 2008) that computational linguistics uses extensively (see Ch. 1 of Kornai, 2023).

The shift in representation format and the underlying logic is caused by a shift in the range of data surveyed. Negation in natural language gives us a Himalayan body of phenomena, lying uncomfortably where two tectonic plates, logic and linguistics, each with its own hypotheses and methods of argumentation, come together. Our perspective is determined by the information-theoretical view (see Ch. 1.3 of Kornai, 2019) that the information content of sentences is dominated by the meaning of words, and

> logical structure accounts for no more than 12–16% of the information conveyed by a sentence, a number that actually goes down with increased sentence length (Kornai 2019:6).

Therefore we begin Section 2 by an exhaustive survey of lexical negation, and throughout the paper we view syntactic negation as a small appendage of the main body. Emphasis is redirected from the sentential (compositional) to the lexical (non-compositional) aspects of negation. Together with Klima, and most subsequent workers like Ladusaw (1980), we assume that negative polarity items can be lexically specified for *neg* not just overtly, as in *nowhere* or *never* ← *not ever*, but also covertly, as in *seldom* or *at all* where the morphology fails to show (even traces of) the negation. We aim at providing a *formal* theory of negation, but the object of our study is ordinary language, where expressions of technical English such as *It is not the case that* are absent (Kornai, 2010b), rather than the formal theory of negation in logic and mathematics. In what follows, 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*.

Our work is intended as a contribution to the linguistic tradition. 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 *neg* 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 (p xiii).

Our main goal in this paper is to pinpoint the source of this asymmetry by using a dyadic predicate LACK. This fits well with the Aristotelian notion of negation, *apophasis*, which tells (phasis) apart (apo) something from something. Just as the Aristotelian term for negation is parallel to the term for affirmation (*kataphasis*, telling something about something), LACK is parallel to our (obviously dyadic) predicates for affirmation, IS_A and HAS. It also fits well with the 'Australian Plan' of Relevant Logic (Meyer and Martin, 1986) in that we keep the law of the excluded middle (no additional truth values beyond True and False) and we do not permit any proposition to be both. A key syntactic element of the Australian approach is the dual operator '*' of which Meyer and Martin, 1986 say:

> Although it is not a particle of English, it *should* be. (p 309, emphasis in the original)

Remarkably, it is not only English that lacks a * morpheme but, to the best of our knowledge, all natural languages do, an empirical fact just as striking as the marked/unmarked asymmetry emphasized by Horn. The solution proposed here is to use an actual morpheme, LACK, as the primitive element, and derive the linguistic asymmetry from the fact that it is dyadic. As we shall see in Section 2, the analogy between logical and arithmetical negatives is quite clear as long as we restrict ourselves to the ancient Greek understanding of arithmetic, where $7 - 5$ could be easily computed as 2, but $5 - 7$ simply made no sense, as there was no concept of -2 to begin with. LACK subtracts something that is by default present, and makes no sense (is infelicitous) otherwise. If *blind* is defined as 'LACK sight', *blind person* makes eminent sense as 'person LACK sight' whereas *#blind stone* runs into the problem of subtracting a property that only animals and image-making devices enjoy by default.

We lay out a theory of negation built on 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. From this perspective, monadic negation is overanalysis in the morphological case, and in Section 3 we argue that the unary negation operator `no`, the pivotal element (written *neg* or $\neg$) in the modern theory, is actually a derived notion obtained from LACK by (generic) quantification over the first (subject) variable of LACK.

Our discussion of compositional constructions in Section 3 also aims 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). The machinery is put to use in Section 4, where we describe how some puzzles generally considered central to the semantics of negation such as double negation, compositional quantifiers, disjunction, and scope ambiguities can be handled with the dyadic system presented here.

There are several aspects of the system presented here that cannot be justified within the bounds of the paper, such as the lack of underlying ternary operators (ditransitives, see Kornai, 2012); no probabilistic or other semiring weighting in the metalanguage (Gyenis and Kornai, 2019); using grammatical functions as linkers (Kiparsky, 1987; Butt, 2006); hypernode (as opposed to hyperedge) graphs as the basic data structure (Woods, 1975); and no doubt many smaller design decisions that sometimes go against the mainstream choices. For the reader interested in a more extensive defense of these we offer a 'hook' to Kornai, 2023, where many of these issues are discussed in greater detail, and the same machinery is used not just for negation, but also for spatial and temporal semantics (Ch. 3); probabilistic reasoning (Ch. 5); modals and counterfactuality (Ch. 6); implicature and gradient adjectives (Ch. 7); proper names and the integration of real-world knowledge (Ch. 8); and some computational linguistic applications (Ch. 9). Readers who prioritize syntax over morphology, and in general syntactic phenomena over lexical ones, are particularly urged to take a look at Ch. 2.4 on linking, and readers with a more lexical/morphological set of concerns are advised to look at Ch. 2.5.

2 Negation in the lexicon

Our survey of negation in the lexicon is designed to be exhaustive, based on the entire vocabulary of English. Starting with the Collins COBUILD (Sinclair, 1987) and LDOCE (Procter, 1978) dictionaries, we found that the words defined in either of these cover well over 99% of running text once numerals, punctuation marks, and proper names are excluded. This reduced the task to inspecting the COBUILD or Longman definitions for negative elements. We search for 'negative' aspects broadly, so as to include not just those words where *no* or *not* appear in the definition, but checking also for the clitic *n't*, the prefixes *un-, in/m/r/l-, de-, dis-, mis-, non-, anti-*, and the suffixes *-less* and *-free*. For safety, we looked both at COBUILD and LDOCE, but we present only the results building on the the Longman Defining

Vocabulary (LDV), since COBUILD definitions can be replaced by Longman headwords, so a negative COBUILD definition is still captured by this test. The LDV, originally about 2,200 elements including bound morphemes, was in turn reduced to a smaller vertex cover set in the definition graph, 1,200 elements in the `4lang` version V1.0, presented in Kornai, 2019 pp 122–124. These 1,200 items were both machine- and manually inspected. Their number have been further reduced to 776 primitive senses in V2.0 (see the Appendix of Kornai, 2023), but to be on the safe side we used the larger set here. An important caveat is that primitive status is not determined uniquely:

> Another difference between the generative and the algebraic approach is that only the former implies commitment to a specific set of primitives. To the extent that work on lexical semantics often gets bogged down in a quest for the ultimate primitives, this point is worth a small illustrative example. Consider the cyclic group Z_3 on three points given by the elements e, a, b and the following multiplication table.

	e	a	b
e	e	a	b
a	a	b	e
b	b	e	a

Table 1. Multiplication in Z_3

> The unit element e is unique (being the one and only y satisfying $yx = xy = x$ for all x) but not necessarily irreducible in that if a and b are given, both ab and ba could be used to define it. Furthermore, if a is given, there is no need for b in that aa already defines this element, so the group can be presented simply as $a, aa, aaa = e$ i.e. a is the 'generator' and $a^3 = e$ is the 'defining relation' (as these terms are used in group theory). Note, however, that the exact same group is equally well presented by using b as the generator and $b^3 = e$ as the defining relation – there is no unique/distinguished primitive as such. This non-uniqueness is worth keeping in mind when we discuss possible defining vocabularies. (Kornai, 2010a)

About 12% of the defining set (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* 'liquidus' as in *thin paint* versus *thin* 'tenuis' as in *thin reed*. The technical means of disambiguating such lexical entries are irrelevant for this paper, but we note that 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, so that disambiguation is rarely called for.

The list has many elements such as *water* which seem to lack any negative aspect. But a closer look at the definition 'liquid, life NEED, has no color, has no smell, has no taste' shows how negative statements enter the picture. Many of these can be handled by our central innovation, in our case replacing the above definitions by `liquid, life NEED, LACK color, LACK taste, LACK smell`. In the formal system that our parser relies on, dyadic predicates are given in CAPS and 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. In addition to subjects and objects of dyadic predicates, denoted by '1' and '2' as in Relational Grammar, see Perlmutter, 1980, our formal system (for details see Ch. 1.3 of Kornai, 2023) also relies on an undifferentiated attribution/predication relation, denoted by '0', that subsumes both *is* and *is_a*, so we have `animal` and `clever` as conjuncts in the definition of *fox* as `animal, four-legged, hairy, red, clever` again conflating, rather than carefully separating, 'direct' and 'metaphorical' usage.

In many adjectival oppositions, normally handled by some version of scalar semantics such as (Kennedy, 2007), 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).

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 affirmatively 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.

The same treatment can be effortlessly extended to many antonym pairs, e.g. defining *good* as the object of WANT, and *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* may be `LACK different` and *different* may be `LACK same`, but only one of these terms has a positive definition: *x is the same as y* means x has all the essential properties of y and y has all the essential properties of x. Since `x IS_A y` means 'x has all the essential properties of y' (Kornai, 2010a), we can define *x same y* by `x IS_A y, y IS_A x` without any recourse to negation. In all such cases, it is really a matter of lexicographic taste whether we choose to mark antonymy 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, and only 83 elements of the original 144 remain, less than 0.7% of the defining vocabulary. 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 first illustrate 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. Conservative people who prefer the status quo will likely be antidisestablishmentarian, but not establishmentarian, 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). In topology, these predicates have such specialized meanings that sets can satisfy both at the same time (these are called *clopen* sets). In ordinary language objects cannot be clopen: a door is either closed or not, in which case it is open. Yet a third state of affairs exists where the status of the object is not known, and this differs in significant ways from graded predicates like *slightly open* or *practically closed*. In the epistemic sense, tertium datur. We will denote this third state by $\circledcirc$, and use $\oplus$ and $\ominus$ to denote the positive and the negative states, but emphasize that these are not truth values, the underlying logic is still binary. We follow Berto and Restall, 2019, who defend the 'Australian Plan' semantics for negation

> (...) based on two ideas. The first is that negation is an exclusion-expressing device: we utter negations to express incompatibilities. The second is that, because incompat*ibility* is modal, negation is a modal operator as well.

As our survey demonstrates, the dyadic negation operator LACK, even though it is limited to cases of incompatibility with lexical prespecification, is already sufficient for capturing the lexical semantics by means of RDF-style meaning postulates – the traditional *neg*, which has no such limitation, is plainly overkill for this. The modal aspect is also clear, not just for the epistemic case used for $\circledcirc$ above, but also for the deontic cases that will be central to our discussion of compositional negation in Section 3. 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. 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 3, and return to double negation in Section 4.1.

Several variants of quantum logic have resources to express the idea that the door is in some quantum superposition state. However, neither the logic used here nor natural language have such resources: what $\circledcirc$ describes is the common situation when we cannot (or just do not wish to) make a committment to $\oplus$ or $\ominus$, rather than some exotic "Schrödinger's Cat" situation.

2.1 Quantifiers

Ever sice Aristotle's *De Interpretatione*, and particularly in the hands of 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, 1987; Badia, 2009), we move in the other direction, and treat quantifiers as nominals whose compositional behavior, which we defer to Section 3, 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" (Kornai, 2018).

That some kind of quantificational ur-element is needed is already clear from a closer look at 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 a default generic, `gen` to fill the subject slot, but caution the reader that this element doesn't have universal import – for now 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. Quantifiers of a clearly compositional nature, like *at most seven, no more than ten*, are deferred to Section 4.2, but we note here that *two* is not defined as *one plus one*: 'being two' is an inherent perceptual property just like 'being blue'. This works well up to the limits of human numeracy, the magic number seven plus or minus two (Miller, 1956), and means exactly two (as opposed to the 'at least two' readings sometimes proposed as default (Horn, 1972) from which the 'exactly two' reading is derived by exhaustification, see Haida and Trinh, 2020 fn. 1). 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`.

At this point, the question may be asked: rather than using *gen* and LACK as primitives, and deriving unary *no* from these as *gen* LACK, why not use binary HAS and unary ¬ as primitives, and derive LACK as ¬ HAS? To the extent ¬ and HAS are independently motivated, this appears considerably simpler (even though *gen* is also well motivated, see Carlson and Pelletier, 1995), and certainly more in keeping with tradition. Our answer is threefold. First, such an analysis of *lack* misses the key lexical insight, that it negates something that is ordinarily present. It is true that *blind* means 'not have sight', but if this were the entire story then *#blind stone* should be just as felicitous as *blind person*. Second, the contrast in felicity smoothly extends from the lexical to the compositional domain: consider the priest at the end of a marriage ceremony uttering

(1) The bride and the groom shall HAVE each other (for the rest of their lives, in sickness and health, ...

If the paraphrase of LACK as 'not have' would be reasonable, by negating (1) (e.g. as a response to the standard callout *Speak now or forever hold your peace*) we should obtain

(1a) The bride and the groom shall not have each other (for the rest of their lives ...)

or, by the suggested analysis

(1b)# The bride and the groom shall lack each other (for the rest of their lives ...)

which is clearly infelicitous. A third point against this suggestion is that it leaves Horn's observation about the linguistic asymmetry between positive and negative statements entirely unexplained, indeed, mysterious. This will be particularly clear for double negation, a matter we will return to in 4.1.

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* (∀), 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, 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.

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 is_a no p` (adjectival/nominal form using the copula), `s (no p)` (overtly negated verb). To make the type theory work out, we will assume a broad type of *matters*, 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.

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* (Talmy, 1988). 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 1:

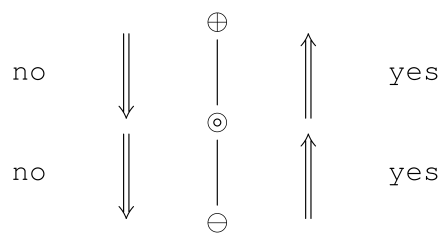

Fig. 1: Forces in negation and affirmation

A word of caution is in order: while finite state automata of the sort depicted here are capable of counting modulo the number of states, the iteration could go to any depth. For example, `no yes yes no no` would move the current state from the initial ⊚ to ⊖, but 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. 1, this time with *start, steady*, and *stop* states (and in a physical system, the kind of signed addition of forces performed by the automaton does make sense).

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 they are the dominant source of negation in primary linguistic data (kids are urged *don't touch* and *don't put it in your mouth* long before they encounter any other negatives), and because they make clear how key features of our model fit with the Australian Plan. An imperative X calls for some action that results in a state of affairs (situation, possible world) where X is fulfilled. Using the `4lang` representation system we can simply write `do X`, and if the command is kept, the meaning postulates associated to *do* will guarantee `after(X)`.

Now, to negate (defy) an imperative means doing some Y that is incompatible with the X targeted by it. Any $Y \curlywedge X$, including the important special case where X calls for some direct action we refuse to perform, is a good way for negating the force of the command, as in the construction site example above. Since phrasal verbs like *move up* are often considered lexical (semicompositional) here we will consider the negatives *Don't smoke!* or *No smoking* and their paired affirmative *Smoke!* where the negation is clearly compositional. Frege already noted that the deontic element takes wide scope over the negation. In the semantic representation, written here with unary ¬ to make the point in standard terms, we have `I order (you ¬smoke)` rather than `I ¬order (you smoke)`. As a referee noted, this observation generalizes to experiencers, which also take broad scope over the embedded proposition. Consider

(2) Diamonds are valuable to John (but Mary considers them stupid, overpriced trinkets).

In the deep structure most analyses would posit a matrix element with an experiencer subject and the embedded proposition as object:

(2a) `John CONSIDER (diamond HAS value)`
(2b) `Mary CONSIDER (diamond LACK value)`

The work to be done by the syntax, relating the matrix subject *John/Mary* to the experiencer expressed by *to* on the surface; and getting the subject *diamond* of the predication expressed by the small clause in surface subject position is the *same* work for the positive HAS and the negative LACK case. There are many syntactic theories capable of doing this, from old-style transformational grammar to minimalism, and it is not clear how on the strength of examples like (2) we could choose among them. The picture is further complicated by the interaction of tense with modals and negation: as Han, 2001 notes, in many languages (e.g., Italian, Modern Greek, and Spanish), imperatives cannot be negated, whereas in English, German, etc. they can be.

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 to a moment to our earlier 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 murder*. Biblical Hebrew (and English at the time of King James) made no distinction between imperative and imperfect, the normative effect (of an ideally kept command) is that in the future there is simply no murder (*retzach*). In our formal language of semantic definitions we can write this as `after(gen LACK murder)`.

We often see antonyms that fit well with the tripartite picture of Fig. 1: *heavy* really means 'has weight greater than gen' and *light* means 'has weight less than gen'. 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.

Returning to our theory of *You shall not murder*, `gen` is the same proquant that we use elsewhere to denote a non-specific entity. After the utterance of the command who does no murder? 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 in Section 2.1: we will not need a unary negation operator `no` since *no(P)* can be defined as `gen LACK P`.

4 Putting it all together

The picture of negation that emerges from our considerations is somewhat nontraditional: instead of the standard, unary negation operation `no` analogous to Boolean ¬, 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)` 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`.

How the (primitive) dyadic negation operator `LACK` and the (derived) unary `no` interact with auxiliaries, main verbs, adjectives, and adverbials is a complex matter. No matter how much this would add to the strength of the theory exposed here, we can't possibly do justice to the syntax of negation in this paper, 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. Before discussing these implications for issues that many semanticists consider key, such as double negation (4.1), compositional quantifiers (4.2), disjunction (4.3), and scope ambiguities (4.4), let us first consider an example that has none of these issues yet remains challenging to many systems.

We will use negative focus sentences like *Mutual fund trades don't take effect until after the market closes* where proposals differ greatly on where they draw the boundary between what is said and what is implied. For most readers of this sentence, the fact that they *do* take effect after market closure is very much part of the meaning – in fact the whole point of the utterance is to tell the impatient traders (especially those more used to stocks, where trades take effect immediately) that they have to wait till the evening. The analysis of *no X until after Y* must directly include, or at the very least must imply `after Y, X`. Whether we consider this a problem in semantics or in pragmatics is of little concern: in Kornai, 2010b we introduced a *Principle of Responsibility:*

> The semantics of any expression must be fully accounted for by the lexicon and the grammar taken together.
>
> The Principle of Responsibility is only slightly stronger than the standard Principle of Compositionality which takes the semantics of any expression to be determined by the semantics of its lexical components and by the grammatical way those are combined. The additional requirement it imposes is that the 'pragmatic wastebasket' remain empty at all times: it doesn't matter whether we call ordinary inferences grammatical, lexical, or pragmatic (and perhaps extragrammatical), the overall system needs to account for these, either in one specific component, or by means of tracing the inference process through several components.

Here the entire work is performed in the semantics, using one and the same implicational mechanism for pragmatic and 'purely semantic' inference, a mechanism that includes substitution of the definiens for the definiendum *salva veritate*. (There are other parts, such as *kal va-chomer*, see Kornai, 2019 Ch 19.4, and spreading activation (Nemeskey et al., 2013) that are not discussed here as they don't play a constitutive role in the following.) We analyze our *X* constituent, *trades take effect* as `X happen`. Since we define `happen` as `change`, we can use the definition of *change,* which is `after(=pat[different])`, and repeat the process by substitution of the definition of `different`, which includes `=pat has quality, =agt lack quality` – the reader can consult the Appendix of Kornai, 2023 to see that the definitions were not created for this particular derivation. Altogether, we obtain that `trade LACK effect before Y` becomes `trade HAS effect` after Y by elementary pieces of temporal deduction (see Ch. 3.2 of Kornai, 2023) and by the very definition of change. But there is considerable burden on the syntax, which must realize that *do*-support separates the subject *mutual fund trades* from the phrasal verb *take effect*, that *until after* is a complex temporal adverb, and on the lexicon, which must find the syntactic and semantic information associated to multi-word expressions. Whether the bulk of the syntactic work is performed by transformations or by grammar formalisms that permit discontinuous constituents is a choice we don't have to make here.

4.1 Double negation

In general, double negation is out (Collins, 2018). 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 been already 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 ¬ 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 that covers both *unhappy person* and the neutral ◎ state 'neither unhappy nor happy' both. 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. Unification of the two `no`s to yield a single negation, rather than cancellation by $\neg\neg = id$ is precisely what we would expect to be the correct semantics for emphatic (reduplicated) *don't*s. 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.

4.2 Compositional quantifiers

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 $\leq 7, \neg(> 10)$ 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 paper, 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 `matter`; *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*.

As standard (Katz and Postal, 1964; Langacker, 2001), we analyze *who* as `wh, person`; *what* as `wh, matter`; *where* as `wh, place`; *when* as `wh, time`; and `how` as `wh, way_2`, where we use the subscript to distinguish the proquantal element from `way_1` 'via'. 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, way_2` to define *everyway*, etc.

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 Allan 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 thematic role mechanism that we have at our disposal independent of negation and quantification (Dowty, 1989) 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` where the angled brackets denote optionality (default), 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*; `matter 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.

4.3 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 by means of De Morgan's laws. 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 double negative, 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. *Or* typically signifies either a future choice to be made, or a past, unknown, choice:

> An *alternative* (or *or*) proposition contains two statements, the acceptance of one of which involves the rejection of the other . . . either may be agreed to, but not both (Lakoff, 1971 p 142)

This makes `or` more closely related to exclusive or (xor) or Latin *aut* than to standard Boolean $\vee$, Latin *vel* (though this standard characterization of Latin *aut* and *vel* is disputed in Jennings, 1994). There are also linguists who dispute this view (Pelletier, 1977; Gazdar, 1979; Pullum, 2006) often on the strength of rather plain examples like *John doesn't walk or talk* whose dominant reading is clearly *John ¬walk* $\wedge$ *John ¬talk*, no matter how much normative grammars like *Fowler's Modern English Usage* object to this. Here we follow the tradition where inclusive *or* would require a separate lexical entry, one with conjunctive semantics, perhaps to be written `vel`.

We note that `vel` is distinct from our conjunctive ‘,’ which is modeled on natural language *and* and therefore involves incrementing the time index on successive verbal conjuncts (cf. the example we started out with, *I went home and had dinner*). Our primitive `or` signifying choice has no temporal update associated to it, and clearly has the ability to introduce alternatives that are counterfactual: *It can wait, or they would have called us by now*. In these respects, inclusive *or* seems to follow Scha’s collective reading: *walk or talk* appears as a single entity ‘perform basic human functions’ rather than a genuine disjunction.

4.4 Scope ambiguities

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 the semantic representation system we rely on (see Ch. 1.3 of Kornai, 2023), the active sentence means `person IN Cormorant Island, 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` for `person IN Cormorant Island`. With this abbreviation the active sentence can be paraphrased as `Cormorant_Islander speak language(two)` and `lg(two) is_spoken_by Cormorant_Islander` 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 Cormorant_Islander speak language(two)` is implicationally equivalent to `lg(two) is_spoken_by Cormorant_Islander, exist Cormorant_Islander`, 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 `Cormorant_Islander 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 4.1, 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 Cormorant_Islander` 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.

4.5 Open problems

One area of notorious complexity that we left untouched is prosody, especially contrastive stress, what Manaster Ramer, 1995 called 'the last refuge of the formal grammarian'. Comparing *It was not 'Peter and 'Kate but only 'Kate who made this mistake* to *It was not Peter 'and Kate but (')only 'Kate who made this mistake* is hard. The task of the grammarian is made even harder by a key design decision of the mainstream syntactic framework, starting with the "T model" of Chomsky's Pisa Lectures, to separate semantic and phonological interpretation early on in the derivation. In contemporary such as the 'new Minimalist Program' of Chomsky et al., 2023 this decision is preserved, so the underlying representation must contain some contrastive stress morpheme "' to track effects that impact both the phonology and the semantics. Following (Harley, 2014), the morphemes are classified as roots or as features, but any contrastive stress element "' seems to display properties of both.

5 Conclusions

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.

In regards to negation, the semantic theory proposed here and in related work (Kornai, 2010a; Kornai et al., 2015) captures well the key observation that negation is not an involution, and in general offers translations whose processing difficulty (Xiang, Grove, and Giannakidou, 2016) 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 Montague Grammar and related theories. We do not see this as a loss, especially not from the learnability perspective (Gyenis and Kornai, 2019).

We started 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. Our work, in many ways a considered return to a more Aristotelian perspective, is not an attempt to 'reform' standard mathematical logic, which we take 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.

Acknowledgments

We thank Anna Szabolcsi (NYU), András Máté (ELTE), and Hans-Martin Gärtner (RIL) for incisive comments on an earlier version of this paper. We are grateful to Dávid Nemeskey (ELTE), Gábor Recski (TU Wien), Márton Makrai (Institute of Cognitive Neuroscience and Psychology), Attila Zséder (Lensa), and all other contributors to the ongoing development effort at GitHub.

We are grateful to the anonymous reviewers for *Acta Linguistica*. The author was particularly pleased by the tough report by Referee "A", which was tremendously helpful for making this a better paper. To quote the timely Bonnici and Simske, 2023 "A truly excellent review, whether recommending acceptance or non-acceptance, motivates the authors to do more in the area their submission addresses", and this is precisely what happened here. We are also grateful to Márta Abrusán (CNRS) and Geoff Pullum (Edinburgh) for pointing the author to very relevant literature.

Kornai was supported in part by 2018-1.2.1-NKP-00008: Exploring the Mathematical Foundations of Artificial Intelligence, by the Hungarian Scientific Research Found (OTKA), contract number 120145, and by the Ministry of Innovation and Technology NRDI Office within the framework of the Artificial Intelligence National Laboratory Program, and MILAB, the Hungarian artificial intelligence national laboratory.

References

Arnauld, Antoine and Claude Lancelot (1975 [1660]). *General and Rational Grammar: The Port-Royal Grammar*. Trans. by Jacques Rieux and Bernard E. Rollin. Mouton.

Bach, Emmon (1977). *An extension of classical transformational grammar*. UMASS.

Badia, Antonio (2009). *Quantifiers in Action: Generalized Quantification in Query, Logical and Natural Languages*. Springer. ISBN: 9780387095639.

Beckett, Dave, ed. (2004). *RDF/XML Syntax Specification (Revised)*. URL: https://www.w3.org/TR/PR-rdf-syntax.

Benacerraf, Paul (1973). "Mathematical Truth". In: *The Journal of Philosophy* 70.19, pp. 661–679.

Berto, Francesco and Greg Restall (2019). "Negation on the Australian Plan". In: *Journal of Philosophical Logic*. URL: https://doi.org/10.1007/s10992-019-09510-2.

Bonnici, Alexandra and Steve Simske (2023). "Reviewer #2 Must Be Stopped!: Or, The Art of Providing Good Reviews". In: *DocEng '23: Proceedings of the ACM Symposium on Document Engineering*, pp. 1–2. DOI: https://doi.org/10.1145/3573128.3607810. URL: https://dl.acm.org/doi/abs/10.1145/3573128.3607810.

Boole, George (1854). *An Investigation of the Laws of Thought on Which are Founded the Mathematical Theories of Logic and Probabilities*. Macmillan.

Böttner, Michael (2001). "Peirce Grammar". In: *Grammars* 4.1, pp. 1–19.

Brachman, R.J. and H. Levesque (1985). *Readings in knowledge representation*. Morgan Kaufmann Publishers Inc., Los Altos, CA.

— (2004). *Knowledge Representation and reasoning*. Morgan Kaufmann Elsevier, Los Altos, CA.

Bresnan, Joan (1982). "The passive in lexical theory". In: *The mental representation of grammatical relations*. Ed. by Joan Bresnan. MIT Press, pp. 3–86.

Butt, Miriam (2006). *Theories of Case*. Cambridge University Press. DOI: 10.1017/CBO9781139164696.

Carlson, Greg and Francis J. Pelletier, eds. (1995). *The Generic Book*. University of Chicago Press.

Chomsky, Noam (1966). *Cartesian Linguistics*. Harper and Row.

Chomsky, Noam et al. (2023). *Merge and the Strong Minimalist Thesis*. Cambridge University Press. DOI: 10.1017/9781009343244.

Collins, Christopher (2018). "*NEG NEG". In: *Glossa* 3.1:64, pp. 1–8. DOI: https://doi.org/10.5334/gjgl.611.

De Mey, M. (1972). "The Psychology of Negation and Attention". In: *Logique et Analyse* 15, pp. 137–153.

Dehaene, Stanislas (1997). *The number sense*. Oxford University Press.

den Dikken, Marcel (2019). "'Not': Not Non-projecting". In: *Argumentum* 15.2, pp. 504–533.

Doherty, Patrick, Witold Lukaszewicz, and Andrzej Szalas (2000). "Efficient reasoning using the local closed world assumption". In: *Proc 9th International Conference on AI: Methodology, Systems, Applications (AIMSA 2000)*.

Dowty, David (1989). "On the semantic content of the notion thematic role". In: *Property theory, type theory and natural language semantics*. Ed. by G. Chierchia, B. Partee, and R. Turner. Dordrecht: D. Reidel.

Firbas, J. (1971). *On the Concept of Communicative Dynamism in the Theory of Functional Sentence Perspective*. Brno University.

Fodor, Jerry A. (1975). *The Language of Thought*. Harvard University Press. ISBN: 0–674-51030-5.
— (2008). *LOT 2: The Language of Thought Revisited*. Clarendon Press. ISBN: 978–0–19–954877–4.
Frege, Gottlob (1879). *Begriffsschrift: eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Halle: L. Nebert.
Gärdenfors, Peter, ed. (1987). *Generalized quantifiers*. Reidel.
Gazdar, Gerald (1979). *Pragmatics: Implicature, presupposition, and logical form*. Academic Press.
Gazdar, Gerald et al. (1985). *Generalized Phrase Structure Grammar*. Oxford: Blackwell.
Genabith, Josef Van and Richard Crouch (1999). In: *Semantics and syntax in Lexical Functional Grammar: The resource logic approach*. Ed. by Mary Dalrymple. MIT Press, pp. 209–260.
Givón, Talmy (1979). *On understanding grammar*. Academic Press.
Goldsmith, John A. and Geoffrey J. Huck (2013). *Ideology and linguistic theory: Noam Chomsky and the deep structure debates*. Routledge.
Gordon, Andrew and Jerry Hobbs (2017). *A Formal Theory of Commonsense Psychology: How People Think People Think*. Cambridge University Press.
Gyenis, Zalán and András Kornai (2019). “Naive probability”. In: *ArXiv*, p. 1905.10924.
Haida, Andreas and Tue Trinh (2020). “Zero and triviality”. In: *Glossa* 5.1. DOI: 10.5334/gjgl.955.
Han, Chung-Hye (2001). “Force, negation and imperatives”. In: *The Linguistic Review* 18, pp. 289–325.
Harley, Heidi (2014). “On the identity of roots”. In: *Theoretical Linguistics* 40.3/4, pp. 225–276.
Hayes, Patrick J. (1978). *The Naive Physics Manifesto*. Geneva: Institut Dalle Molle.
— (1979). “The naive physics manifesto”. In: *Expert Systems in the Micro-Electronic Age*. Ed. by D. Michie. Edinburgh University Press, pp. 242–270.
Horn, Larry (1972). “On the semantic properties of the logical operators in English”. PhD thesis. UCLA.
— (1989). *The Natural History of Negation*. Chicago: University of Chicago Press.
Huffman, David A. (1952). “A method for the construction of minimum redundancy codes”. In: *Proceedings of the IRE*. Vol. 40, pp. 1098–1101.
Jackendoff, Ray (1969). “An Interpretive Theory of Negation”. In: *Foundations of Language* 5, pp. 218–241.
Jennings, R. E. (1994). *The Genealogy of Disjunction*. Oxford University Press. ISBN: 0-19-507524-2.
Kadmon, Nirit and Fred Landman (1993). “Any”. In: *Linguistics and Philosophy* 16 (4), pp. 353–422.
Kaplan, David (1978). “On the logic of demonstratives”. In: *Journal of Philosophical Logic* 8, pp. 81–98.
Katz, Jerrold J. and Paul M. Postal (1964). *An Integrated Theory of Linguistic Descriptions*. Cambridge: MIT Press.
Kennedy, Christopher (2007). “Vagueness and grammar: the semantics of relative and absolute gradable adjectives”. In: *Linguistics and Philosophy*. DOI: 10.1007/s10988-006-9008-0.
King, Margaret (1979). “Knowledge representation in database systems”. In: *ISSCO Working Paper* 33.
Kiparsky, Paul (1987). *Morphosyntax*. Stanford University: ms.
Klima, E.S. (1964). “Negation in English”. In: *The structure of language*. Ed. by J. A. Fodor and J. J. Katz. Prentice-Hall, pp. 246–323.
Kornai, András (2010a). “The algebra of lexical semantics”. In: *Proceedings of the 11th Mathematics of Language Workshop*. Ed. by Christian Ebert, Gerhard Jäger, and Jens Michaelis. LNAI 6149. Springer, pp. 174–199. DOI: 10.5555/1886644.1886658.
— (2010b). “The treatment of ordinary quantification in English proper”. In: *Hungarian Review of Philosophy* 54.4, pp. 150–162.
— (2012). “Eliminating ditransitives”. In: *Revised and Selected Papers from the 15th and 16th Formal Grammar Conferences*. Ed. by Ph. de Groote and M-J Nederhof. LNCS 7395. Springer, pp. 243–261. DOI: 10.1007/978-3-642-32024-8_16.
— (2018). “Truth or *dare*”. In: *Tokens of Meaning: Papers in Honor of Lauri Karttunen*. Ed. by Cleo Condoravdi and Tracy Holloway King, pp. 511–521. URL: http://kornai.com/Drafts/dare.pdf.

Kornai, András (2019). *Semantics*. Springer Verlag. ISBN: 978-3-319-65644-1. DOI: 10.1007/978-3-319-65645-8. URL: http://kornai.com/Drafts/sem.pdf.

— (2023). *Vector semantics*. Springer Verlag. DOI: 10.1007/978-981-19-5607-2. URL: http://kornai.com/Drafts/advsem.pdf.

Kornai, András et al. (2015). "Competence in lexical semantics". In: *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics*. Denver, Colorado: Association for Computational Linguistics, pp. 165–175. DOI: 10.18653/v1/S15-1019. URL: https://www.aclweb.org/anthology/S15-1019.

Ladusaw, William A. (1980). *Polarity Sensitivity as Inherent Scope Relations*. New York: Garland Press.

Lakoff, George (1970). *Irregularity in Syntax*. Holt, Rinehart, and Winston.

— (1971). "Presupposition and Relative Well-Formedness". In: *Semantics: An Interdisciplinary Reader in Philosophy, Linguistics, and Psychology*. Ed. by D. Steinberg and L. Jakobovits. Cambridge University Press, pp. 329–340.

Langacker, Ronald (2001). "What *wh* means". In: *Conceptual and Discourse Factors in Linguistic Structure*. Ed. by Alan Cienki, Barbara Luka, and Michael B. Smith. CSLI Publications, pp. 137–152.

Lappin, Shalom (1996). "Generalized Quantifiers, Exception Phrases, and Logicality". In: *Journal of Semantics* 13, pp. 197–220.

Lehrer, Adrienne (1985). "Markedness and Antonymy". In: *Journal of Linguistics* 21.2, pp. 397–429.

Lenat, Douglas B. and R.V. Guha (1990). *Building Large Knowledge-Based Systems*. Addison-Wesley.

Manaster Ramer, Alexis (1995). "Book Reviews: The Language Complexity Game". In: *Computational Linguistics* 21.1, pp. 124–131. URL: https://aclanthology.org/J95-1010.

Meyer, Robert K. and Errol P. Martin (1986). "Logic on the Australian plan". In: *Journal of Philosophical Logic* 15, pp. 305–332.

Miller, George A. (1956). "The magical number seven, plus or minus two: some limits on our capacity for processing information". In: *Psychological Review* 63, pp. 81–97.

Minsky, Marvin (1975). "A framework for representing knowledge". In: *The Psychology of Computer Vision*. Ed. by P.H. Winston. McGraw-Hill, pp. 211–277.

Moltmann, Friederike (1995). "Exception Phrases and Polyadic Quantification". In: *Linguistics and Philosophy* 18, pp. 223–280.

Montague, Richard (1970). "Universal Grammar". In: *Theoria* 36, pp. 373–398.

— (1973). "The proper treatment of quantification in ordinary English". In: *Formal Philosophy*. Ed. by R. Thomason. Yale University Press, pp. 247–270.

Nemeskey, Dávid et al. (2013). "Spreading activation in language understanding". In: *Proceedings of the 9th International Conference on Computer Science and Information Technologies (CSIT 2013)*. Yerevan, Armenia: Springer, pp. 140–143. URL: https://hlt.bme.hu/media/pdf/nemeskey_2013.pdf.

Parsons, Terence (1970). "Some problems concerning the logic of grammatical modifiers". In: *Synthese* 21.3–4, pp. 320–334.

— (2017). "The Traditional Square of Opposition". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2017. Metaphysics Research Lab, Stanford University.

Pelletier, Francis J. (1977). "Or". In: *Theoretical Linguistics* 4, pp. 61–74.

Perlmutter, David M. (1980). "Relational grammar". In: *Current approaches to syntax*. Ed. by Wirth and Moravcsik. Academic Press, pp. 195–229.

Procter, Paul (1978). *Longman Dictionary of Contemporary English*. 1st ed. Longman.

Pullum, Geoffrey K. (2006). "Exclusive OR: free dinner and stay out of jail". In: *Language Log*. URL: https://languagelog.ldc.upenn.edu/nll/?p=46.

Quillian, M. Ross (1967). "Semantic memory". In: *Semantic information processing*. Ed. by Minsky. Cambridge: MIT Press, pp. 227–270.

Reiter, Raymond and Giovanni Criscuolo (1983). "Some representational issues in default reasoning". In: *Computers and Mathematics with Applications* 9.1, pp. 15–27.

Ross, Alf (1941). "Imperatives and Logic". In: *Theoria* 7, pp. 53–71.

Russell, Bertrand (1905). “On denoting”. In: *Mind* 14, pp. 441–478.

Scha, R. (1981). “Distributive, Collective and Cumulative Quantification”. In: *Formal Methods in the Study of Language, Part 2*. Ed. by J. A. G. Groenendijk, T. M. V. Janssen, and M. B. J. Stokhof. Mathematisch Centrum, pp. 483–512.

Sinclair, John M. (1987). *Looking up: an account of the COBUILD project in lexical computing*. Collins ELT. DOI: 10.1016/0022-5193(74)90110-6.

Sowa, John F. (2000). *Knowledge representation: logical, philosophical, and computational foundations*. MIT Press.

— (2008). “Conceptual Graphs”. In: *Handbook of Knowledge Representation*. Ed. by F. van Harmelen, V. Lifschitz, and B. Porter. Elsevier.

Szabolcsi, Anna (2015). “What do quantifier particles do?” In: *Linguistics and Philosophy* 38.2, pp. 159–204. ISSN: 1573-0549. DOI: 10.1007/s10988-015-9166-z. URL: https://doi.org/10.1007/s10988-015-9166-z.

Talmy, L. (1988). “Force dynamics in language and cognition”. In: *Cognitive science* 12.1, pp. 49–100. DOI: 10.1207/s15516709cog1201_2.

Wells, Roulon S. (1947). “Immediate constituents”. In: *Language* 23, pp. 321–343.

Woods, William A. (1975). “What’s in a link: Foundations for semantic networks”. In: *Representation and Understanding: Studies in Cognitive Science*, pp. 35–82. DOI: 10.1016/B978-0-12-108550-6.50007-0.

Xiang, Ming, Julian Grove, and Anastasia Giannakidou (2016). “Semantic and pragmatic processes in the comprehension of negation: An event related potential study of negative polarity sensitivity”. In: *Journal of Neurolinguistics* 38, pp. 71–88.

Zipf, George K. (1949). *Human Behavior and the Principle of Least Effort*. Addison-Wesley.