Elliot Murphy

Chapter 7
The nature of language and the structure of reality

Abstract: Natural language affords humans with the ability to construct an unbounded array of hierarchically structured expressions. The syntactic component involves forming binary-branching sets via the operation MERGE, taking objects from the lexicon or objects already part of the syntactic workspace. These structures then regulate linguistic meaning, which amounts to forms of conceptual instructions. A major strand within the generative enterprise further assumes that assembling linguistic structures does not involve reference to extra-mental entities (i.e., there is no 'word-world' relation). This type of semantic internalism is defended here using a broad range of case studies. The unifying theme throughout is that linguistic structure and meaning are wholly mind-internal processes exhibiting a specific computational and representational architecture. By attempting to inject models of syntax and semantics with concerns relating to cognitive constraints, I review how recent efforts in this direction may be able to unveil a means to reconcile properties of human language with endogenous properties of the brain.

1 Introduction

Natural language syntax yields an unbounded array of hierarchically structured expressions, as opposed to linear combinations of elements into string-like units. Further, linguistic rules appear to be structure-dependent (Chomsky 1975), a unique feature not shared with artificial forms of intelligence (Leivada et al. 2022). Within contemporary generative grammar and certain strains of the biolinguistic enterprise, syntactic computation is thought to adhere to "general principles that may well fall within extra-biological natural law, particularly considerations of minimal computation" (Chomsky 2011: 263), such that certain linguistic theories might be engaging with general properties of organic systems (Gallego & Martin 2018; Murphy et al. 2024). Relatedly, language is fundamentally framed as being a product of an individual's mind/brain, termed an 'I-language' (i.e., 'internal', 'individual', 'intensional')

Elliot Murphy, Vivian L. Smith Department of Neurosurgery, McGovern Medical School & Texas Institute for Restorative Neurotechnologies, UTHealth, Houston, Texas, USA,
e-mail: elliot.murphy@uth.tmc.edu

https://doi.org/10.1515/9783111293776-008

(Chomsky 1986, 2000): "It is not that speakers communicate because they have an E-language ['external'] in common; rather, where I-languages overlap sufficiently, communication is possible" (Hornstein and Anthony 2003: 9; see also Murphy 2020b).

More broadly, the language system is thought to be optimally designed for structure-meaning mapping, whereas issues pertaining to externalization (speech, sign, etc.) introduce complexities above and beyond the basic concern of constructing sets of syntactic objects (as Chomsky often puts it, language is not 'sound with meaning', but rather 'meaning with sound'). In keeping with these ideas about the fundamental nature of human language, in this chapter I will begin by reviewing the most recent generative conception of linguistic *structure* (Section 1), before moving onto considerations of linguistic *meaning* (Section 2), with a view to surveying the major architectural components of linguistic computation.

2 Structure

> Things take strange shapes in the mind. (Ishiguro 2015: 156)

Humans, according to Otto Jespersen, have "a notion of structure" that they impose on linguistic data. The syntactic component of language involves the construction of binary-branching hierarchically organized sets via the operation MERGE. This operation initiates with lexical access, manipulating objects from the lexicon *or* objects already part of the syntactic workspace.

MERGE

Given the set {X, Y}, we can either select a new lexical object and MERGE it, to form {Z, {X, Y}}, or we can select an existing object to form {X, {X, Y}}. These sets are then 'labeled' and given a syntactic identity (Chomsky 2013; Murphy 2015a, 2015b; Murphy et al. 2022b), based on which element is most structurally prominent and easiest to search for (i.e., Z in the structure {Z, {X, Y}}). There is increasing evidence that only elements in the workspace that have been labeled can be subject to movement (Bošković 2021), suggesting that labeling forms a core architectural component of regulating syntax-semantics. Natural language structure-building via MERGE satisfies Turing's original definition of computation, involving the composition of a series of elementary functions which are non-decomposable and atomic.

The earliest version of MERGE was introduced in the early 1990s ('lower case' Merge) by Chomsky and developed by a number of researchers, most of whom

would later form what David Foster Wallace (1997) called the "Militant Grammarians of Massachusetts". This early definition eliminated the distinction between the phrase structure procedure and the transformational procedure. Merge, unlike phrase structure rules, does not rewrite symbols, but merely adds information about their grouping (Chomsky 1995). Recently, ('capital') MERGE has been defined as an operation on a workspace and its objects (Chomsky et al. 2019), formalized as follows (WS = workspace; P/Q = workspace objects; X = additional elements):

WS = {P,Q, . . . }
MERGE(P,Q,WS) = WS' = {{P,Q},X_1, . . ., X_n}

Below is a basic case of traditional External (responsible for argument structure):

WS = [X, Y, Z] → MERGE(X, Y) → WS' = [{X, Y}, Z]

In contrast, we can also implement Internal Merge (movement; responsible for discourse- and information-related properties):

WS = [{X, Y}, Z] → MERGE(X, {X, Y}) → WS' = [{X, {X, Y}}, Z]

A simple case of deriving the structure 'the man sells fish' is represented below:

WS = [the, man, sells, fish] → MERGE(sells, fish)
WS' = [the, man, {sells, fish}] → MERGE(the, man)
WS'' = [{the, man}, {sells, fish}] → MERGE({the, man}, {sells, fish})
WS''' = [{{the, man}, {sells, fish}}]

Given that MERGE can target any given syntactic object in a workspace, it can combine 'the' and 'man' independently of any previous application of constructing the VP. MERGE is strictly Markovian; the present stage is independent of what was generated earlier, unlike standard recursion. Hence, the labeled status of both the NP and the VP are respected. We have no need to invoke multiple workspaces or cross-workspace comparisons: we only need recursive applications of MERGE to a set of accessible objects.

Generative grammar assumes that a core part of an individual's knowledge of language is the possession of such a computational procedure (Chomsky 1995). Other potentially innate forms of knowledge that pertain to higher cognition include: knowledge of objecthood; algorithmic representations; a type-token distinction; a capacity to represent sets and path trajectories; spatiotemporal contiguity; causality; a capacity for some form of cost-benefit analysis; 'core knowledge sys-

tems' such as intuitive mechanics, intuitive physics, intuitive biology; a self-concept and a theory of mind (Marcus 2001; Pinker 2007; Spelke 2010).

Units of computation

The atoms of linguistic computation are assumed to be bundles of lexical features or functional categories (Chomsky 1967; Ramchand & Svenonius 2014). Linguists within, for example, the Distributed Morphology tradition have long abandoned hope in a coherent notion of 'word'; there are simply groups of phonological, semantic and formal syntactic features. Some elements have [+syntax] and [+semantics] but [-phonology], like copies of movement, elided elements or 'PRO' elements, whereas other elements have [+phonology] and [+syntax] but [-semantics], like expletive elements ('It is raining'), and others have [+phonology] and [+semantics] but [-syntax], like single-unit utterances ('Hello').

Another consequence of this position is that 'words' are also built by MERGE. Consider 'institutionalize', and other morphologically complex structures, which are 'mini sentences' in their hierarchical relations and exhibit a compositional semantics. Morphemes, formal features, and perhaps also phonemes (Kayne 2019) may be elements within the workspace. Notice that non-compositional words, like 'socialism', 'university' and 'reactionary' (which is not semantically composed from the meaning of 'react', '-ion' and '-ary') may need to be stored in memory, as with idioms, perhaps in a distinct 'pragmatic lexicon' generated by concerns of performance (as opposed to competence) (Carston 2022), potentially posing some challenges to the single-lexicon, single-MERGE account.

These atoms of computation are intensional entities, mental symbols that are abstractions of brain states, connecting to conceptual systems that use them as instructions for generating particular inferences (Adger 2022a). Word meanings can act as "concept assembly instructions" (Pietroski 2018: 292). As Section 2 will explore below, these linguistic representations are not intentional (with a 't') entities: they are not 'of' or 'about' anything in the world, they are simply addresses for conceptual representations. Some representations generate a phonetic representation, while others purely play a role in syntax-semantics (e.g., C, *u*wh, T). When linguists speak of VPs and c-command, these are simply mental mechanisms that are abstracted from currently unknown neurophysiological processes (Benítez-Burraco and Murphy 2019; Murphy 2020a).

What can be said with some confidence, however, is that the meanings of words are representations providing detailed instructions to conceptual systems, with even seemingly very similar words like 'more' (a subset-subset relation) and 'most' (a subset-superset relation) yielding a number of distinct psycholinguistic

consequences (Knowlton et al. 2021). These words themselves are likely not atomic concepts, but call upon a range of primitive logical concepts like cardinality, comparison, superset and subset. Atoms of semantic instructions remain unclear, and require further theoretical and experimental research (Pietroski 2018).

Likewise, semantically very similar words can have distinct selectional requirements and permit distinct syntactic environments. Verbs like 'said' permit complementizer deletion, whereas 'whisper' and other manner of speech verbs ('murmur', 'chuckle') do not (Adger 2003), and they also do not permit extraction from their complements (Franks 2005):

(1) a. Jason whispered that the phoenix had escaped.
b. * Jason whispered the phoenix had escaped.

(2) a. What$_i$ did Saul say [$_{CP}$ (that) [$_{TP}$ he saw t_i]]?
b. * What$_i$ did Saul whisper [$_{CP}$ that [$_{TP}$ he saw t_i]]?

Natural language syntax exhibits *discrete units* which lead to a *discreteness-continuity duality*. This results in the possibility that the boundary between syntactic categories can be non-distinct. Syntax is driven by *closeness of computation* (syntactic objects X and Y form a distinct syntactic object, {X, Y}). Its objects are *bounded* (a fixed list, e.g., N, V, Adj, Adv, P, C, T, *n*, *v*, Asp, Cl, Neg, Q, Det) and their hierarchical ordering is based on a specific functional sequence such as C-T-*v*-V which imposes direct restrictions on combinatorics (Adger and Svenonius 2011). This functional sequence may be learned by infants, or may form part of innate knowledge arising from mereological relations between conceptual categories (Ramchand and Svenonius 2014).

Some other issues that emerge here include the possibility of lexicon-internal operations, such as categorization, involving MERGE-type operations within the lexicon itself, prior to workspace insertion. For example, <v^*, R> or past tense objects like 'walked' may be good candidates for lexicon-internal MERGE operations. Pursuing this architecture further, note that functional categories such as v^*, T and C typically have a specifier position, whereby Spec-Head agreement is available. However, while this is permitted in English, other languages do not exhibit such an agreement relation. Recent work (Sugimoto 2022) argues that such languages introduce functional categories as "an amalgam where two or more . . . are formed into one element before merging into narrow syntax", extending the scope of 'early' syntactic combinatorics even further. One consequence of this is that the specifier position is reduced, capturing the absence of ϕ-agreement in Japanese.

The operation of Self-Merge has also been discussed in recent literature, a candidate for another intra-lexicon operation:

Self-Merge: √man → {√man}

Recursive Self-Merge has been used to turn acategorial roots within the lexicon into interpretable syntactic objects (syntax/labeling cannot 'see' them otherwise), which are then characterized by the categorial and functional features associated with the Self-Merged object (e.g., N{√man}). This motivates Adger (2013) to dispense with postulating functional heads as the source of these syntactic features, since they can instead by built from the 'bottom-up' via lexical features and Merge.

These types of theoretical developments are part of a more global aim (emerging in the early 2000s) to explain grammatical constraints in *procedural* terms arising from the behavior of syntactic operations, rather than utilizing kinds of *declarative* constraints that simply stipulate and formalize a constraint on computation.

Labeling and minimal search

After lexical selection and MERGE, the operation of labeling takes place when conceptual systems access the structures generated by syntax. This occurs at distinct derivational punctuations based on the configuration and subcategorization demands of the lexical items involved. For example, in the case of head-complement structures this is done immediately after MERGE (Bošković 2016). Syntax is an autonomous module, with interpretability of features arising at the syntax-semantics interface at the point of labeling.

Labeling is assumed to locate the closest or least embedded head to determine the category of a given structure (Shim 2022).

(3) SO_1 = {X, YP}

(4) SO_2 = {{X, YP}, {Z, WP}}

In (3), the head X is found by the labeling algorithm since it is the least embedded head. However, (4) presents an ambiguity in search depth, and so an assumption in the literature is that if the two heads share common features then this serves as the label, as in (5) (but see Murphy and Shim 2020 for an alternative account), or one of the two heads is what is left behind after some labeling-driven movement, as in (6).

(5) {{$X_{[F]}$, YP}, {$Z_{[F]}$, WP}}, where [F] = [F] (e.g., <3SG>)

(6) {{X, YP}, ~~{Z, WP}~~}

Alongside labeling, other operations include Agree and feature inheritance (Murphy 2015a; Smith et al. 2020). Agree and labeling both involve the same *search algorithm* via minimal search (i.e., heads like X in (3) are feature-bundles, whereas sets carry no features, so heads are detected before the search procedure checks into YP) (Cao 2017). However, they differ in their *search target* and *search domain* (Ke 2022). The search target for Agree is the feature attributes of unvalued features on the head initiating the search (e.g., Person, Number). Agreement, for example, checks the [*u*ϕ] on the search-initiating probe and the [*u*Case] on the goal (Sobin 2014). The search target for labeling is categorial features on heads (Murphy and Shim 2020). Agree, resulting in feature valuation (Arregi and Nevins 2012), utilizes a generic copy operation, while labeling seems to utilize categorization and object maintenance (Murphy 2015a, 2015c) – all sufficiently generic to form strong candidates for (Marrian) implementational mapping (Murphy 2020a).

Minimal search is likely not executed through an encapsulated series of Fodorian modules with limited central access to compare across searches, but rather through a parallel breadth-first search procedure (as opposed to depth-first), allowing the syntax to search a particular domain of sets at once to find a specific feature, respecting also c-command relations across the entire structure; e.g., "derivation is in parallel" for cases such as NP-VP, for Chomsky (2024). Chow (2022) presents a promising candidate for a minimal search algorithm that exhibits certain features of *both* depth- and breadth-search, compatible with 'smuggling' derivations in that it exhaustively searches within a transparent (accessible) specifier before searching the specifier's sister. The specific implementation and scope of the memory/search architecture of syntax may constitute part of the species-specific computational apparatus of language. It may turn out that the regulation of breadth- vs. depth-search algorithms arises from the interaction of certain syntactic objects (and not others) with interface demands, and that these demands are not imposed by the conceptual systems of non-human primates onto their more restricted combinatorial systems. Following this theme, Rizzi (2016) presents a distinction between 1-Merge, 2-Merge and 3-Merge systems (respectively, unit-unit, unit-phrase, phrase-phrase). Human language is safely within the 3-Merge category. Other species have been claimed to implement recursive self-embedding, e.g., two carrion crows via the elegant design in Liao et al. (2022), which may speak to the existence of a pre-MERGE capacity for basic recursive rule following, although this data might also be accounted for via non-recursive ordinal lists, rather than a reflexive facility for center-embedding recur-

sion. Generative linguistics stresses the specific type of recursion natural language syntax exhibits; what objects it manipulates; and how structures are *mapped* to interpretations.

Surreal syntax

Other formal systems share certain properties with this model of linguistic structure-building, perhaps even properties related in cognitive basis. For instance, Chomsky has pointed out that recursive application of MERGE to a single restricted element, the empty set { }, applied to itself yields the natural numbers via the successor function.

Pursuing this line of inquiry, consider here the surreal numbers, discovered by John Conway (1976), which include both the ordinals and the real numbers. Conway (1976) extends the von Neumann ordinals to posit a kind of dual set theory to implement a parallel comparison procedure to re-define the class of all possible numbers in a non-Archimedean system. The surreal numbers can be defined as equivalence classes of cuts, i.e., an object {L | R}, where L and R are sets whereby each element of L is less than each element of R. Issues of minimality arise immediately: the structure {0 | } is the simplest number greater than 0 (i.e., 1). {1 | ω} is the simplest number greater than 1 and less than infinity (i.e., 2), and so on; the definition of zero is therefore { | }, and issues become much more complex with numbers such as the square root of ω, or ω^3, or 2ω = {ω + n | } (Conway and Guy 1996). The left/right sets constitute encapsulated planes of computation that provide instructions for inductive reasoning about the defined number. Natural language syntax utilizes this type of dual memory system – potentially across multiple workspace types, as in Adger (2017), but also across minimal search in labeling/Agree – for the purposes of assigning to structures a range of interpretations. The mechanics of Conway's number system are what we arrive at if we map a comparison search procedure onto our 'number sense' core knowledge system (with its own representations and rules of induction), rather than the lexicon.

Domains of derivation

MERGE can also derive some set-theoretic properties of linguistic relations, such as *membership*, *dominate* and *term-of*, as well as the derived relation of *c-command* (=sister of) which is relevant for interpreting hierarchical relations between linguistic elements (Epstein 1999; Chomsky 2021a). Syntactic objects can be combined in

phases (Chomsky 2008), which can be extended to form non-local dependencies (see Murphy 2016; Murphy & Leivada 2022; Murphy et al. 2024 and also Roberts 2019; Lightfoot 2020; Collins 2022 for how economy conditions may aid language acquisition by reducing the number of tree structures that the child needs to consider). Phases are essentially chunks of syntactic workspace memory accessed by the interfaces, typically associated with clausal/complementizer (CP) and verb phrase (*v*P) domains, since these seem to be the only domains in which the full range of morphophonological, syntactic and semantic effects that should be linked to intermediate movement is attested. Still, the PP and DP domains may also be phasal, despite some key effects (multiple spell-out; semantic effects) being absent (van Urk 2020). Related ideas in recent literature include the proposal that the highest head in the extended projection of a lexical head is a phase head, irrespective of the features of the head; i.e. phasehood might be a more context-dependent issue bearing a uniform configuration (Bošković 2014).

Issues of phasehood relate directly to issues of reducing computational complexity. Examining some core principles of recursion, natural language clearly exhibits minimisation (Kleene 1952) alongside a constrained set of representational categories. Binary branching of structures (Radford 2016) limits redundant computation, reducing the range of possible computations, and is also "the computationally most parsimonious combinatory principle capable of yielding an unbounded set of structures" because "it requires just two temporary buffers of operative memory in which to store the elements to be strung together" (Rizzi 2013), in contrast to unary Merge and *n*-ary Merge.

One of the implications of this framework is that the innate component of the language faculty (Universal Grammar) is likely unrestricted MERGE, but when situated within the context of an organic computational system that is constrained by economy factors ('3rd factors' in language design) it is restricted further (Collins 2022):

I-Language: unrestricted MERGE

I-Language + Economy: binary branching MERGE

The present workspace-centric definition of MERGE has a number of conceptual and empirical advantages over older versions, and permits the formulation of new efficiency criteria, such as 'Resource Restriction'. Some current topics of concern in the literature focus on pair-Merge of heads, since pair-Merged structures of heads generated by MERGE lead to rendering the notion of accessibility ill-defined (for discussion, see Shim 2022). In the remainder of this section, I will briefly discuss some of these issues.

Resource Restriction

Resource Restriction (Chomsky 2019; Chomsky et al. 2019) states that when the combinatorial operation MERGE maps workspace *n* to workspace *n*+1, the number of computationally accessible elements (syntactic objects) can only increase by one (Komachi et al. 2019). This can account for a peculiar property of natural language recursion that separates it from other forms of recursion (e.g., propositional calculus, proof theory): natural language MERGE involves recursive mapping of workspaces that removes previously manipulated objects (Chomsky 2021b). Similar observations apply to the idea that when MERGE targets objects in a workspace, non-targeted elements remain conserved and intact once the new workspace has been established (Chomsky et al. 2019), continuing the intuition of 'No Tampering'. MERGE itself exhibits the formal characteristics of a finite-state rewrite rule (Trotzke and Zwart 2014), with minimal computational complexity. Resource Restriction is sympathetic to basic observations from neurobiology and active inference (Parr et al. 2022), whereby the brain instinctively disregards the vast majority of incoming data in order to make inferences over a limited set of inputs.

A topic of recent discussion concerns how we can define workspace size. Fong et al. (2019) suggest that the size of a syntactic workspace should be considered to be the number of accessible terms plus the number of syntactic objects. This theory requires only element counting – opening up questions about whether other measures may be used to establish legal operations, such as compressibility (Murphy et al. 2024).

The examples in (7) illustrate how Resource Restriction can determine legitimate MERGE applications.

(7) a. WS = [X, Y] → MERGE(X, Y) → WS' = [{X, Y}]
b. WS = [{X, Y}] → MERGE(X, {X, Y}) → WS' = [{X, {X, Y}}]
c. WS = [X, {W, Y}] → MERGE(X, W) → WS' = [{X, W}, {W, Y}]

In (7a), the number of accessible objects is 2 (X and Y), and the resulting workspace (WS') has 3 (the elements X and Y, and the new set, {X, Y}). Hence, this is a legitimate operation, since the number of accessible objects has increased only by one. However, in (7b) the number of accessible objects in WS is 3, but in WS' it is 5. Yet, Chomsky (2021b) assumes that another criteria is relevant in determining legitimacy alongside Resource Restriction. While Resource Restriction checks the increase in accessibility, it does not make any assumptions about the nature of accessibility. A principle that does address accessibility is minimal search. In (7b), WS' includes two occurrences of X, however they differ in their structural embed-

ding. And so, the X within {X, Y} becomes inaccessible. Hence, using Resource Restriction alongside minimal search, we find that (7b) is indeed a legal operation. Lastly, we can then conclude that (7c) involves elements that can be searched for equally, and so WS' increases accessibility by more than one, hence is an illegal syntactic operation.

This architecture assumes that if the conceptual systems do not attempt to access or succeed in accessing syntactic structures, then no interpretation can be assigned to them. Chomsky (2021b, 2024) assumes that the interface systems directly access the derivational workspace for interpretation, eliminating the need for, e.g., Transfer operations or the Phase Impenetrability Condition. Language-external interfaces in turn impose their own range of constraints; for example, Truswell (2007) proposes a condition at the interpretation interfaces whereby cases of movement that cross structures describing multiple events are ruled out.

Puzzles about parsing

Lastly, there is a common distinction made between the grammar and the parser. A grammar licenses derivations that pair structures/meanings in an unbounded fashion. A parser finds a structure that can be paired with a meaning when data is encountered. One approach to this issue that many researchers have found appealing is the 'one-system' hypothesis (Lewis and Phillips 2015) that grammatical theories and language processing models describe the same cognitive system, as evidenced by the fact that grammar-parser misalignments only seem to occur as a consequence of limitations in domain-general systems such as memory access and control mechanisms, and as evidenced by the convergence between online and offline responses to grammatical anomalies (Sprouse and Almeida 2013).

The centrality of structural information in human sentence parsing is sometimes underestimated in the literature, but there are many areas where structural factors seem important; e.g., measures of syntactic surprisal can explain processing dynamics of garden path sentences independently of more commonly used measures like lexical surprisal (Arehalli et al. 2022). It has been extensively documented how syntactic structure impacts parsing: Structural distance (between nodes) is a significant predictor of reading times, while linear distance is not (Baumann 2014); violations of hierarchical sentence rules lead to slower reading times (Kush et al. 2015); and expectations of word category based on hierarchical grammars also predicts reading times (Boston et al. 2011). However, online processing measures indicate that dependency formation (e.g., for agreement and movement) take constant time, and not proportionately more time for longer search paths/spaces (Kush 2013). This suggests that online language processing

does not execute a node-by-node minimal search procedure of the kind outlined above – which only indicates that the computational account may not map directly onto algorithmic models and implementation. Complexity at one level of description does not need to match the complexity of a different level. Building connections between successful computational accounts of structure-building, and successful psycholinguistic accounts of parsing, is a major challenge for future research.

New directions, old mysteries

The architecture I have outlined here raises a number of serious questions. Issues remain about how to relate the lexicon to the workspace, and whether or not the workspace is a syntactic object. Recent controversies concern whether set theory provides the right tools for representing natural language phrase structures (Gärtner 2022), or whether other models, such as mereological systems (Adger 2022b), are more suitable. Other issues concern 'rule ordering' in syntax (Obata et al. 2015; Goto and Ishii 2020), whereby, for example, feature inheritance, agreement, labeling and movement must take place in a certain order, but without much in the way of a principled explanation as to why (though rule ordering differences can contribute to explanations of cross-linguistic variation; Sugimoto 2022). With respect to domains of computation, one major recent debate concerns the extent to which ellipsis can be used as a diagnostic for phasehood (Bošković 2014; Todorović 2020).

A much larger issue looms, as ever, in the background: How and why are certain lexical elements entered into the workspace to begin with, prior to any application of MERGE? The 'creative aspect' of language that this question calls upon remains entirely outside all current models of language (Chomsky 1995).

3 Meaning

> Nature is a temple, where, from living pillars,
> Confused words are sometimes allowed to escape;
> Here man passes, through forests of symbols,
> Which watch him with looks of recognition.
>
> (Baudelaire, *Les fleurs du mal*, 1857,
> 'Correspondances' no. 4)

Once linguistic structures are assembled via this system of lexical selection, MERGE, labeling, agreement, and interface access, they are then interpreted. We turn now to the issue of what kind of interpretations are afforded by language.

The citadel itself

Darwin wrote in his *Notebook N* in 1838, aged 29, that "[t]o study Metaphysics, as they have always been studied appears to me like puzzling at astronomy without mechanics. – Experience shows the problem of the mind cannot be solved by attacking the citadel itself. – The mind is a function of body" (cited in Gould 1975: 125). Darwin's recommendation for the naturalistic study of mind – that it should not sideline the concerns of "body" – should carry over to the investigation of language. Any account of linguistic meaning must be cognitively and psychologically plausible.

A biolinguistic approach to lexical reference, taking stock of Darwin's citadel and the endogenous properties of human psychology, assumes that "there are no independently specifiable external physical objects that correspond to the referential expressions we use" (Hinzen 2007: 21). What is known as 'semantic internalism' investigates internal states (Valente 2019); it studies the organism in abstraction from the complexities of the social, physical and cultural environment (i.e., the I-language). As such, the claim that linguistic structures like 'this water' derive their meaning internally is to say that they do *not* derive their meaning through the achievement of some internal approximation to a feature of the environment. We impose these structures on sensory epithelia and non-linguistic conceptual domains (see also Parr et al. 2022); a framework no different from the intensional definition of atoms of syntactic computation outlined above in Section 2. In this sense, natural language semantics is essentially a form of syntax – both are concerned with internal symbol manipulation.

In contrast to this biolinguistic perspective, the majority of working philosophers identify as externalists (Gertler 2012), exploring word-world and alethic constructs. Externalists assume that some, or all, expression meanings are individuated externally. They typically place emphasis on metalinguistic and communicative processes (Rast 2023). Semantic externalism maintains that word meanings are *denotational*, while sentence meanings involve *truth conditions* (some configuration of states in the world – putatively identifiable – that can be mapped onto language). The remainder of this chapter will defend the idea that internalist models of meaning deserve to be considered part of what we mean by the nature of language. Instead of focusing on trying to specify conditions under which sen-

tences can be true (externalism), I will focus on the cognitive details pertaining to the format of representations (internalism).

Reference

A common belief in philosophy of language is that words refer to entities, qualities and events in our sensorium and are directly related – via some connection – to the external world. In Kant, Bolzano and the early Wittgenstein we find the idea that a mind-internal representation can exhibit an isomorphism with the thing representated. More recently, we find the Davidsonian view that the compositionality of meaning can be accommodated only via postulating truth-conditions for sentences, and referents for words (Davidson 1984). Meanwhile, Lewis (1975) held that 'meaning' is to be found in the notion of 'conventions', or regulated forms of verbal behaviour, treating language as a form of sociolinguistics or game theory.

Yet, reference is not a pre-theoretical notion (Ludlow 2003; Leivada and Murphy 2021). It is a theoretical term that works differently in different contexts. Reference can be seen as an *action* (Strawson 1950; Chomsky 2000), but also the process through which individual lexical items can be used to fetch conceptual representations (Pietroski 2018). Roots of these ideas – of mapping linguistic forms to semantic interpretations – in modern thought can be traced to Katz and Fodor (1963) and Jackendoff (1972). The act of referring involves (at least) the speaker, the utterance, the context, and some properties of the external world (i.e., Humean sense data and phenomenology, or related formulations).

In keeping with these ideas, traditional analytical philosophy was essentially psychological in character, from the work of Frege to the later Ayer to Chrisholm (Johnston and Leslie 2012). Only relatively recently has externalism been substantively defended (Putnam 1975; Burge 1979; Kripke 1979, 1980).

There is a rich history behind these developments. Nineteenth-century philologists stressed the importance of sounds; by the 1920s, Otto Jespersen and others had shifted attention to the written word; in the 1950s, Chomsky moved the study of language towards abstract mental structures, beyond any specific form of input modality. By the 1970s, Tarskian systems of 'semantics' and truth-conditions had been developed to a sufficient degree of sophistication, but a false conclusion was made: Natural language has meaning; meaning is a system of semantics; semantics can be represented via Tarskian models; therefore, the semantics of natural language is Tarskian. While one can readily model a system of semantics via Tarski-style truth-conditions, it does not follow that the meanings we derive from

natural language expressions (and not formal language expressions) follow the same 'externalist' system of reference (Higginbotham 1990; Pietroski 2022).

The issue of reference becomes more difficult when we entertain the following cases ('#' = semantically anomalous):

(8) a. The book was interesting but in poor condition.
b. # The school caught on fire and was welcomed at the Oval Office.
b. The White House was repainted before issuing a statement.
c. # Donald Trump sued the left-wing newspaper that fell off the table.

We can access multiple senses of 'book' and metonymic organizational expressions like 'the White House' simultaneously, yet nothing in the physical world can satisfy the conditions of being, for instance, a unit of information and a collection of pages. Not even Donald Trump himself can negotiate the metaphysical presuppositions of information-instutition-physical entities. This phenomenon of jointly associating multiple senses to a single expression is termed copredication (Murphy 2012, 2021a, 2021b, 2021c, 2022). Internalists point to the centrality of human functional interests and concerns in linguistic reference, undermining the view that these nominals are a mereological sum (forming a 'complex object' in the mind-independent world) or a union of the extensions of different senses. Notice that verbs also generate copredications:

(9) Doctor Tandon writes well but illegibly.

This eventive description can simultaneously refer to a process, informational content, and a product. While copredications can arise via verbs, it seems possible (Levin and Rappaport Hovav 2013) that verbs encode only thin, schematic meanings to be filled by rich nominal representations.

Certain of these polysemous concepts exhibit a semantic structure which assumes the form of a set of clustered and overlapping interpretations. For example, a number of senses can be attributed to a single nominal in certain configurations:

(10) The school with large windows starts at 9am and has a strict headmaster and unruly students.

We can think of institutions with these complex polysemous senses in a manner that seems to be possible only due to linguistic structure.

Yet at the same time, our MERGE-based computational system yields the ability to provide a kind of skeletal structure to thoughts, often *without* the need to

do deep conceptual search. A simple sentence like 'the drink on the table' provides only the bare minimum conceptual detail which can then be filled in by other cognitive systems. Compare visual meaning, such as a painting of the Parthenon, with the linguistic meaning of the phrase 'the Parthenon'. We can easily imagine the Parthenon, but if someone asks us 'How many columns are there?', we cannot provide an answer: the vagueness of the representation is immediately exposed. We use language, very often, so that we *don't* have to think.

As soon as we engage in the seemingly harmless act of naming something, we are immediately in the game of highlighting category-defining features of the thing, and rendering less salient non-diagnostic features that are not relevant to category membership but which still pertain to the object (Lupyan 2012). What provides us with a sense of "chairs and their chairness" (Coetzee 2013: 122) is not visual perception (which contributes sortal and colour features), but rather language (labeling, in particular), which projects root/lexical concepts into 'higher' planes of computation, "boost[ing] the dimensionality of the human mind" (Hinzen 2009: 31). Linguistic meaning seems to assist us towards highlighting abstract relational commonalities between things, very regularly recruiting older cognitive systems (e.g., analogous reasoning). The "use of a linguistic label prompts a search for sameness across instances", highlighting "shared relational structures" in ways that a label from some other modality might not (Pritchard 2022: 15).

Composing meanings

> His own opinion, which he did not air, is that the origins of speech lie in song, and the origins of song in the need to fill out with sound the overlarge and rather empty human soul. (Coetzee 2000: 4).

Consider next a seemingly trivial observation: the meaning of a stop sign is different from the meaning of the MERGE-generated structure '[a [stop sign]]'. Only the latter involves combinatorial composition to derive a hierarchically structured meaning. Syntactic headedness and semantic headedness typically map biunivocally onto each other (Piatelli-Palmarini 2020), with close ties existing between syntactic structure and interpretation (Hinzen 2006). 'Red boat' is a Noun Phrase, and it is a boat that is red, not a red quality with boat-like features. There are exceptions involving subsective adjectives: with 'former lawyer' or 'alleged threat', the adjective syntactically modifies the noun, but semantically they involve modifying the adjectival phrase (a former lawyer is not a lawyer). Yet even here, a former lawyer was a lawyer at some point, and an alleged threat or a fake lion or a toy train all exhibit some resemblance – semantics cannot escape syntax.

Composing meanings is a non-trivial matter even at the single word level, before we get to phrases. Consider how any polity concept invites joint contributions from categorically distinct domains: *Liverpool* can be 'fun' and 'polluted', and can be burned down and re-built on the opposite side of the Mersey river, yet still be called *Liverpool* (Chomsky 2000). The city does not so much *have* a function as it *is* a function. Recent experimental work points to the prominent role of functionality over materiality (Murphy 2017, 2021a, 2021b, 2021c). As Hinzen (2006: 66–67) notes: "An object shows its true nature not by exploring its natural and holistic connections with other things, but by putting it in an artificial circumstance where it establishes accidental and 'unnatural' connections with other objects".

I-Language and I-Semantics

> The question young David raises is an important one. What does it mean to be an orphan? Does it simply mean that you are without visible parents? No. To be an orphan, at the deepest level, is to be alone in the world. (Coetzee 2020: 5)

To unpack the intricacies of linguistic meaning further, consider 'water'. This is loosely defined as a physical entity (H_2O) by externalists, considered to be a rigid designator referring to the same entity across all possible worlds; as such, the meaning of the word 'water' determines an extension for it, i.e. the set of water substances (H_2O plus some impurities). It seems more accurate to characterize the lexical meaning of 'water' as a functional concept, which is weighted to a certain degree (though not entirely) by its physical constitution: Some Europeans crossing the Congo in 1648 were so famished and starved that they were forced to "drink water [which] resembled horse's urine" (Braudel 1992: 227).

In opposition to this, we find the "super-externalist" (Ludlow 2003: 149) position that linguistic form can only be determined via a complete ontology (e.g., of H_2O). Externalism maintains that the meanings of words are determined by extra-mental properties. But is the externalist therefore also committed to the reverse claim (which seems to follow logically) that the extra-mental world contains entities which correspond to mind-internal lexical representations?

Other problems emerge in the following case (Chomsky 2000): A tea factory contaminates my local water supply. What comes out of my tap is very slightly contaminated with the substance from the tea leaves, but we still judge it to be water. Meanwhile, a chemically identical substance is being made a few blocks away in a neighbourhood with a clean water supply, when someone makes a cup of tea with the same concentration of tea chemicals. My substance is water, theirs is tea.

Relatedly, Pietroski (2017: 207), using US government statistics, notes how Diet Coke has a higher percentage of H_2O than "the stuff from my well". He adds that "Diet Sprite® and club soda are even more like H_2O" yet are not deemed *water* for reasons to do purely with "intended purposes". Pietroski (2018) maintains that denotational acts, of the kind involving complex polysemy, cannot be explained purely by semantic knowledge (see also Pietroski 2022). For rough approximation, Ludlow (2003) makes a distinction between I-substances (water) and P-substances (H_2O), which refer to intensions and physical entities.

This argument connects directly to another internalist position, the idea that the external world (post-Newton) is fundamentally unknowable, exhibiting properties that directly conflict with common sense reasoning, and that only scientific models can be developed and tested. Models can be intelligible, even if the world itself is not – contrary to a long tradition in philosophy under which "the structure of reality . . . is intelligible" and knowable (Cator 1907) – or so the internalist argument goes (see also Murphy 2023). While experiential content itself *simply is a property of the universe* (Strawson 2008), in virtue of existing within reality, anything that falls outside of this domain can only be approximated via the limits of our perceptual and cognitive faculties. In Conway's (2009) formulation, it is not our job to understand the universe: we cannot comprehend a physics that is no longer deterministic, and in which (quantum) particle behavior is strictly *uncaused* and *independent of past history* – disproving Leibniz's principle of sufficient reason, that everything must have a reason/cause, perhaps the most fundamental of classical scientific assumptions. We can only hope to generate intelligible models of reality, coherent to ourselves but to no other forms of organic or artificial life in existence – present or future.

Consider next the Twin Earth thought experiment (Putnam 1975). In some parallel universe, set a few hundred years before the present (i.e., before modern chemistry), it is said that water is not made of H_2O but rather some other substance, XYZ. Can the inhabitants of this Twin Earth use 'water' to refer to this substance, having the same subjective experience with it? Putnam says 'No'. The semantic internalist says 'Yes', because 'water' seems to be polysemous between some more common, function-based sense, and a more technical sense. The strict externalist reading also leads us to conclude that all substances that make use of 'twin water' (XYZ) also must be conceived differently. Ice cream and diet soda on Twin Earth must *also* be different entities to their equivalent on Earth, and therefore so must our own bodies.

Indeed, imagine some other parallel universe. Pietroski (2017) offers Fraternal-Earth, where doppelgangers of our scientists discover that what they have all been loosely calling 'mud' in fact has a deep, uniform chemical structure. All forms of 'mud' are of substance XYZ, hence they can use the concept 'mud' to

refer successfully to all physical structures of 'mud'. But it does not follow from this that the inhabitants of Fraternal-Earth could not have used 'mud' to refer to our chemically-diverse samples of mud, if they entered a black hole and ended up in the English countryside.

Putnam agrees with the internalists that words in and of themselves do not refer but maintains that our intuitions about how words can refer are constructed via person X referring to object Y using sign S (1992: 382); which presupposes some mind-independent structure for object Y, which simply begs the question of how words relate to any notion of objecthood.

Moving to another example, Kripke (1980, 2019) provides a rejoinder to Dummett's social use of natural kind terms, and also Putnam's externalism *vis-à-vis* its treatment of linguistic division of labor, whereby part of our meaning of the word 'gold' lies in our deference to experts to provide an explanation for some chemical basis. Though Kripke uses the following argument to defend a version of externalism, for our purposes it serves to highlight what the semantic internalist would see as the inherent polysemy in terms like 'gold'. Kripke notes that 'gold', like water, can be said to be composed of the chemical element Au. Imagine that all of the world's gold experts (who can distinguish gold from fool's gold) are invited to a conference in a large convention center, and are all killed in an explosion. Kripke (1980, 2019) asks: Does it follow that the extension of gold now encompasses fool's gold, since there is nobody around to tell the difference? If we assume that 'gold' has no such extensions to begin with, we have no problem. The crucial factor that Kripke does not enter into discussion here is the background knowledge of the people he is presenting this argument to. It is already too late for Kripke's audience to remove what Ludlow would term the P-substance sense from their lexical entry for 'gold'. In terms of the externalist viewpoint, there are seemingly two distinct responses: one is to say that the reference is fixed via a causal-historical chain, and the other is to say that the reference is fixed via word meaning changing over time (a view supposing that natural language is related, somehow, to 'conventions' and 'contexts' and 'circumstances').

The version of semantic internalism that I have been reviewing here assumes that material constitution may indeed factor as an essential component of our conceptions of water and gold (indeed, *some* material basis seems necessary), but our representations of these entities assume more of a complex grid-like structure, with different weights being afforded to different senses (see also Hampton et al. 2007; Martí 2015; Nichols et al. 2015; Tobia et al. 2020).

More recent frameworks can also be scrutinized in this light. The development of *two-dimensional semantics* (which distinguishes between a word's primary and epistemic intension) presupposes a word-object reference relation (Byrne and Pryor 2006: 38). Two-dimensional semantics works on the assumption

that, "[s]ince it is not a priori that water is H_2O, there is an epistemic possibility in which water is, say, XYZ, and not H_2O, even though there is no such metaphysical possibility" (Ibid. 41). But this cannot be the case, since the concept of water carries no ontological presuppositions, unlike H_2O. There are simply concepts or representations, neither of which are *of* anything except themselves ('water' and 'H2O' are objects from distinct Wittgensteinian language games), making talk of "reference via causal-historical chains" (Russell 2010: 190) unmotivated. For the semantic internalist, lexical items can be deployed as *hypotheses* about the world, but they do not host some extra-mental property that the language system independently investigates.

Moving on from water and gold, consider lastly the case of 'jade'. After initial usage framed this as being a natural kind of some type, it was subsequently revealed to be composed of two different substances (silicate minerals), jadeite and nephrite. Yet this has no implications for the validity of semantic internalism, and the lexicon treats it the same as 'gold'. By happenstance, the letter 'J' is the only letter not currently occupying the periodic table. Imagine that scientists discover in ten years from now that the two elements that jade is composed of turn out to be derived from a single, more fundamental element, J. Hence jade can once more be neatly packaged into a convenient word-object reference – but with no implications for the validity of the claim that 'jade' has no extensions.

I heard you paint houses

Recent empirical work has sought to address some of these issues concerning internalism/externalism. Haukioja et al. (2021) conducted a series of experiments revealing that people use both appearance and underlying identity when defining a selection of five natural kind terms ('water', 'lightning', 'gold', 'diamond', 'tiger'), depending on which aspect is highlighted as being novel. The authors say that this shows that both externalism and internalism face problems, since neither 'essence' nor appearance are utilized in a monolithic manner (however, the internalist accounts I have reviewed do not suppose such monolithic power for 'essence', and internalism would not have been proven right if material basis were judged to be almost irrelevant). These experimental results seem to show that different conceptual demands are at play depending on the scenario presented to participants. As the authors themselves note, the results are possibly compatible with a causal homeostasis view of natural kind concepts, whereby deep and superficial properties cannot be separated (Rehder and Kim 2010).

Aristotle pointed out in his *Metaphysics* that the form of a house (its wood and stones) differs from its function. He realised that material constitution – or

"the facts of the universe" (Carnap 1956: 11) – is only one factor we take into account when deciding objecthood. Many other factors have been entertained in recent work (Pustejovsky and Batiukova 2019), building on Aristotelian aitiational semantics, attempting to explore the explanations/causes for objecthood.

Other features beyond use and functionality, such as geometry, sometimes seem more salient. If we use Chomsky's (2000) example of painting a house red, it seems to be understood, *a priori*, that it is the exterior surface of the house that is painted, not the interior, reflecting in part an underlying CONTAINER concept.

(11) John painted the house red.

This is because "the semantic features impose an analysis in terms of specific properties of intended design and use, a designated exterior, and indeed far more intricacy . . . The exterior-interior dimension has a marked and unmarked option; if neither is indicated, the exterior is understood" (Chomsky 2000: 125). If John is stood outside, 5 metres from the house, and Mary is stood inside, 5 metres from the surface, John is 'near' the house, but Mary is not. Even if we say that John painted 'the spherical cube' red, we know the the exterior surface was painted, despite the non-existence of this impossible entity (Chomsky 2000: 35, 2007).

Colour terms like 'red' in 'red house' raise similar extensional problems. How can we specify the supposed extension of 'red'? Proffering 'the set of all red things' does not help the matter. Not all portions of a red thing are necessarily red: compare a 'red house' (exterior or interior) with 'red cave' (interior). While red ink is red through and through, red food does not need to be. A lexical item like 'red' is "placed in an abstract space . . . with its own organization that is only marginally influenced by the spatial relations in the real world" (Arsenijević 2008: 15).

In a recent cognitive neuroscience study of object conceptualization, Cheng et al. (2021) note that we think of ourselves as being 'inside' a scene but 'outside' an object, and speculate that one diagnostic feature of a scene may be concavity. This leads to the prediction that images depicting concavity will be behaviorally categorized as scenes more often than those depicting convexity. A second prediction is that the neural scene-processing network will respond more to concave images than convex. Both predictions were realized. Instead of evincing extensions of wordly entities, linguistic meaning instead seems to be tapping into robust and complex conceptual networks, with their own principles of computation and criteria for access and deployment in ordinary inference and communication.

These and many other examples suggest that language seems to provide (or at least encode lexically) function-related notions, and functional information "is

at the root of conceptual formation" (Mandler 2004: 210). We seem to have a deep, almost unshakeable intuition that concepts get their meaning from something outside the head. But this is just as baseless as our intuitions about mass, space and time, which have all been shattered by modern physics (Chomsky 2000; Conway 2009).

Overall, it seems that a 'house' is much more than its (i) *material constitution* and its (ii) *function* (as Aristotle assumed): we can also use its (iii) *origin* (Hobbes) or derive a sense of (iv) *continuity* (Locke's view of personhood, which can be shown to generalize far beyond humans) to individuate entities. In addition, we also have (v) *extra-linguistic biases* for shaping objecthood, and (vi) extra-lexical encyclopedic knowledge and prototypical instantiations, closely wedded to each lexical item. The conceptual content of words cannot be wholly isolated from their linguistic properties (sadly, for many experimental psycholinguists) since it is directly shaped by them. Lexical knowledge is not simply a 'label' for concepts; it is a unique mode of human conceptualization (Pietroski 2018; Acquaviva 2022), providing an efficient, accessible link to clusters of representations from distinct core knowledge systems.

4 Conclusions

> Wisps and flashes of uncolored cloud flashed past the window. Above and below were a different story, but there was always something disappointing about clouds when you were inside them; they ceased to be clouds at all. It just got really foggy (Wallace 2011).

With the biolinguistic enterprise, what philosophers typically define as "the problem of the external world" (Neta 2003), or somesuch formulation, becomes not a purely metaphysical question, but a cognitive and naturalistic one. The problem instead becomes how to alter one's prior representations/inferences as minimally as possible so as to match the data of sense. The external world may be an inconvenience, but it is not a "problem". The very notion of semantic externalism qua Putnam is deeply at odds with how major currents in the brain sciences view organic worldly interactions (Ramstead et al. 2020), defying Darwin's principle of centering properties of "the citadel".

Chomsky will occasionally claim – as he did in a recent conversation with Kripke (Chomsky 2021c) – that reference is only achieved if a word refers to something that a physicist could examine, which is something of an extreme caricature of externalism, but the general argument stands. And Kripke will rebut, as he did to Chomsky, that when we discuss the intricate properties of 'river' and 'house' and the many factors that seem to conspire to determine possible individuation,

that we are in fact not denying that these words have extensions, only that the extensions happen to be very "complicated". Yet this brings with it the inherently cognitively implausible feature of complex extensions for 'river' and 'house' and 'school' somehow being set during language acquisition. Chomsky (p.c.) further rebuts this specific externalist position: "What notion of extension permits an object X to be within or outside the extension of a term depending on how we think about it? I've never heard of any, and can't imagine what it would be". Entertaining highly "complicated" extensions will not suffice if the same entity can be, but also not be, a house or river depending on how we conceive of it. As Collins (2021: 401) has noted in this connection: "It is a grand philosophical achievement of Chomsky's to so much as entertain the idea that language is its own thing, an object of interest that may sustain deep inquiry regardless of its poor design relative to the ends to which formal systems are developed, and regardless, too, of whether ancient philosophical quandaries are resolved".

In this chapter, I have defended the claim that the MERGE-based mechanics of the human mind generate internal representations that signal to other cognitive systems possible interpretations. These mechanics make reference to other cognitive domains, but they do not appear to deem the external world fit for such referential treatment. The mapping of linguistic structures to their meanings – regulated by economy conditions and recruiting an array of thousands of lexical representations – is more than complicated enough without "the facts of the universe" (Carnap 1956: 11) getting in the way.

Despite some major advances in explanatory scope since the origins of modern linguistics, I think the classical, Enlightenment era picture remains correct; a picture of human language as being "a brief and rare spark in the universe, one that we may soon extinguish", constituting in some fundamental ways "what may be a true marvel of the cosmos" (Chomsky 2021b: 4), providing some appropriate context for the following reflections of anthropologist and philosopher Loren Eiseley (1959: 161–162):

> In a universe whose size is beyond human imagining, where our world floats like a dust mote in the void of night, men have grown inconceivably lonely. We scan the time scale and the mechanisms of life itself for portents and signs of the invisible. As the only thinking mammals on the planet – perhaps the only thinking animals in the entire sidereal universe – the burden of consciousness has grown heavy upon us. We watch the stars, but the signs are uncertain. We uncover the bones of the past and seek for our origins. There is a path there, but it appears to wander. The vagaries of the road may have a meaning, however; it is thus we torture ourselves.
>
> Lights come and go in the night sky. Men, troubled at last by the things they build, may toss in their sleep and dream bad dreams, or lie awake while the meteors whisper greenly overhead. But nowhere in all space or on a thousand worlds will there be men to share our loneliness. There may be wisdom; there may be power; somewhere across space great in-

struments, handled by strange, manipulative organs, may stare vainly at our floating cloud wrack, their owners yearning as we yearn. Nevertheless, in the nature of life and in the principles of evolution we have had our answer. Of men elsewhere, and beyond, there will be none forever.

References

Acquaviva, Paolo. 2022. Word meaning: A linguistic dimension of conceptualization. *Synthese* 200. https://doi.org/10.1007/s11229-022-03910-9

Adger, David. 2003. *Core Syntax: A Minimalist Approach*. Oxford: Oxford University Press.

Adger, David. 2013. *A Syntax of Substance*. Cambridge, MA: The MIT Press.

Adger, David. 2017. "A memory architecture for merge". https://lingbuzz.net/lingbuzz/003440 (accessed 23 September 2024)

Adger, David. 2022a. What are linguistic representations? *Mind & Language* 37(2). 248–260. https://doi.org/10.1111/mila.12407

Adger, David. 2022b. "Syntax is mereological"l. CRISSP Lectures Online. https://www.crissp.be/wp-content/uploads/2022/05/adger-lectures.pdf (accessed 23 September 2024)

Adger, David & Peter Svenonius. 2011. Features in minimalist syntax. In Cedric Boeckx (ed.), *The Oxford Handbook of Linguistic Minimalism*, 27–51. Oxford: Oxford University Press.

Arehalli, Suhas, Brian Dillon & Tal Linzen. 2022. Syntactic surprisal from neural models predicts, but underestimates, human processing difficulty from syntactic ambiguities. arXiv: 2210.12187. https://doi.org/10.48550/arXiv.2210.12187

Arregi, Karlos & Andrews Nevins. 2012. *Morphotactics: Basque Auxiliaries and the Structure of Spellout*. New York: Springer.

Arsenijevic, Boban. 2008. From spatial cognition to language. *Biolinguistics* 2(1). 3–23. https://doi.org/10.5964/bioling.8615

Baumann, Peter. 2014. Dependencies and hierarchical structure in sentence processing. *Proceedings of the Annual Meeting of the Cognitive Science Society* 36(36). 152–157.

Benítez-Burraco, Antonio & Elliot Murphy. 2019. Why brain oscillations are improving our understanding of language. *Frontiers in Behavioral Neuroscience* 13. https://doi.org/10.3389/fnbeh.2019.00190

Bošković, Željko. 2014. Now I'm a phase, now I'm not a phase: On the variability of phases with extraction and ellipsis. *Linguistic Inquiry* 45. 27–89. https://doi.org/10.1162/LING_a_00148

Bošković, Željko. 2021. "Merge, move, and contextuality of syntax: the role of labeling, successive-cyclicity, and EPP effects". https://boskovic-linguistics.media.uconn.edu/wp-content/uploads/sites/2801/2021/04/CambridgeHandbookShortFinal_FixedReferences.pdf (accessed 23 September 2024)

Boston, Marisa Ferrara, John T. Hale, Shravan Vasishth & Reinhold Kliegl. 2011. Parallel processing and sentence comprehension difficulty. *Language and Cognitive Processes* 26. 301–349. https://doi.org/10.1080/01690965.2010.492228

Braudel, Fernand. 1992. *Civilization and Capitalism 15th-18th Century, Volume I: The Structures of Everyday Life: The Limits of the Possible*. Berkeley/Los Angeles: University of California Press.

Burge, Tyler. 1979. Individualism and the mental. *Midwest Studies in Philosophy* 4(1). 73–122. https://doi.org/10.1515/9780748629299-015

Byrne, Alex & James Pryor. 2006. Bad intentions. In Manuel García-Carpintero & Josep Macià (eds.), *Two-Dimensional Semantics*, 38–54. Oxford: Oxford University Press.
Cao, Yu. 2017. Labeling for linearization. https://lingbuzz.net/lingbuzz/003427 (accessed 23 September 2024)
Carnap, Rudolf. 1956. *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago/London: University of Chicago Press.
Carston, Robyn. 2022. Syntactic structures and pragmatic meanings. *Synthese* 200. https://doi.org/10.1007/s11229-022-03861-1
Cator, Gerald. 1907. The structure of reality. *Mind* 16(61). 54–69.
Cheng, Annie, Dirk B. Walther, Soojin Park & Daniel D. Dilks. 2021. Concavity as a diagnostic feature of visual scenes. *NeuroImage* 232. https://doi.org/10.1016/j.neuroimage.2021.117920
Chomsky, Noam. 1967. The formal nature of language. In Eric H. Lenneberg (ed)., *Biological Foundations of Language*, 397–442. New York, NY: Wiley and Sons.
Chomsky, Noam. 1975. *Reflections on Language*. New York, NY: Pantheon Books.
Chomsky, Noam. 1986. *Knowledge of Language: Its Nature, Origin, and Use*. Westport, CT: Praeger.
Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.
Chomsky, Noam. 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
Chomsky, Noam. 2008. On phases. In Robert Freidin, David Michaels, Carlos P. Otero & Maria Luisa Zubizaretta (eds)., *Foundational Issues in Linguistic Theory: Essays in Honor of Jean-Roger Vergnaud*, 133–166. Cambridge, MA: The MIT Press.
Chomsky, Noam. 2011. Language and other cognitive systems. What is special about language? *Language, Learning and Development* 7(4). 263–278. https://doi.org/10.1080/15475441.2011.584041
Chomsky, Noam. 2013. Problems of projection. *Lingua* 130. 33–49. https://doi.org/10.1016/j.lingua.2012.12.003
Chomsky, Noam. 2019. Some puzzling foundational issues: the Reading Program. *Catalan Journal of Linguistics Special Issue*. 263–285. https://doi.org/10.5565/rev/catjl.287
Chomsky, Noam. 2021a. Simplicity and the form of grammars. *Journal of Language Modelling* 9(1). 5–15. https://doi.org/10.15398/jlm.v9i1.257
Chomsky, Noam. 2021b. Minimalism: where we are now, and where we can hope to go. *Gengo Kenkyu (Journal of the Linguistic Society of Japan)* 160. 1–41. https://doi.org/10.11435/gengo.160.0_1
Chomsky, Noam. 2021c. Invited lecture at the Rule Following and Normativity class, Spring. Saul Kripke Center. The Graduate Center, CUNY.
Chomsky, Noam. 2024. Genuine explanation. In Giuliano Bocci, Daniele Botteri, Claudia Manetti & Vincenzo Moscati (eds.), *Issues in Comparative Morpho-Syntax and Language Acquisition*, 15–44. Oxford: Oxford University Press.
Chomsky, Noam, Ángel J. Gallego & Dennis Ott. 2019. Generative grammar and the faculty of language: insights, questions, and challenges. *Catalan Journal of Linguistics Special Issue*. 226–261. https://doi.org/10.5565/rev/catjl.288
Chow, Keng Ji. 2022. A novel algorithm for minimal search. *Snippets* 42. 3–5.
Coetzee, John Maxwell. 2000. *Disgrace*. London: Vintage.
Coetzee, John Maxwell. 2020. *The Death of Jesus*. New York: Viking.
Coetzee, John Maxwell. 2013. *The Childhood of Jesus*. London: Harvill Secker.
Collins, Chris. 2022. The complexity of trees, Universal Grammar and economy conditions. *Biolinguistics* 16. https://doi.org/10.5964/bioling.9573
Conway, John Horton. 1976. *On Numbers and Games*. New York: Academic Press.

Conway, John Horton. 2009. “The free will lectures”. https://web.math.princeton.edu/facultypapers/Conway/ (accessed 23 September 2024)
Conway, John Horton & Richard K. Guy. 1996. *The Book of Numbers*. New York, NY: Springer-Verlag.
Davidson, Donald. 1984. Truth and meaning. *Truth and Interpretation*. Oxford: Oxford University Press.
Eiseley, Loren. 1959. Little men and flying saucers. *The Immense Journey*. London: Vintage.
Epstein, Samuel D. 1999. Un-principled syntax: the derivation of syntactic relations. In Samuel D. Epstein & Norbert Hornstein (eds.), *Working Minimalism*, 317–345. Cambridge, MA: The MIT Press.
Fong, Sandiway, Robert C. Berwick & Jason Ginsburg. 2019. The combinatorics of Merge and workspace right-sizing. Paper presented at Evolinguistics Workshop, Tokyo, 25–26 May, 2019.
Franks, Steven. 2005. What is that? In Yoshihisa Kitagawa & Dorian Roehrs (eds.), *Indiana University Working Papers in Linguistics* 5. 33–62.
Gärtner, Hans-Martin. 2022. Copies from “standard set theory”? A note on the foundation of minimalist syntax in reaction to Chomsky, Gallego and Ott (2019). *Journal of Logic, Language and Information* 31. 129–135. https://doi.org/10.1007/s10849-021-09342-x
Gallego, Ángel J. & Roger Martin (eds). 2018. *Language, Syntax and the Natural Sciences*. Cambridge: Cambridge University Press.
Gertler, Brie. 2012. Understanding the internalism-externalism debate: what is the boundary of the thinker? *Philosophical Perspectives* 26(1). 51–75.
Goto, Nobu & Toru Ishii. 2020. Determinacy theory of movement. *NELS* 50. 29–38.
Gould, Stephen Jay. 1975. Review of *Evolution of the Brain and Intelligence* by Harry J. Jerison. *Paleobiology* 1(1). 125–129. https://doi.org/10.1017/S0094837300002244
Hampton, James A., Zachary Estes & Sabrina Simmons. 2007. Metamorphosis: essence, appearance and behaviour in the categorization of natural kinds. *Memory & Cognition* 35. 1785–1800. https://doi.org/10.3758/BF03193510
Haukioja, Jussi, Mons Nyquist & Jussi Jylkkä. 2021. Reports from Twin Earth: both deep structure and appearance determine the reference of natural kind terms. *Mind & Language* 36(3). 377–403. https://doi.org/10.1111/mila.12278
Higginbotham, James. 1990. Contexts, models, and meanings: a note on the data of semantics. In Ruth M. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality*, 29–48. Cambridge: Cambridge University Press.
Hinzen, Wolfram. 2006. *Mind Design and Minimal Syntax*. Oxford: Oxford University Press.
Hinzen, Wolfram. 2007. *An Essay on Names and Truth*. Oxford: Oxford University Press.
Hinzen, Wolfram. 2009. The successor function + LEX = human language? In Kleanthes K. Grohmann (ed.), *InterPhases: Phase-Theoretic Investigations of Linguistic Interfaces*, 25–47. Oxford: Oxford University Press.
Hornstein, Norbert & Louise M. Anthony. 2003. Introduction. In Louise M. Anthony & Norbert Hornstein (eds.), *Chomsky and His Critics*, 1–10. Oxford: Blackwell.
Ishiguro, Kazuo. 2015. *The Buried Giant*. London: Faber & Faber.
Jackendoff, Ray. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, MA: The MIT Press.
Johnston, Mark & Sarah-Jane Leslie. 2012. Concepts, analysis, generics and the Canberra Plan. *Philosophical Perspectives* 26(1). 113–171.
Katz, Jerrold & Jerry A. Fodor. 1963. The structure of a semantic theory. *Language* 39(2). 170–210. https://doi.org/10.2307/411200
Kayne, Richard S. 2019. What is suppletion? On **goed* and on *went* in Modern English. *Transactions of the Philological Society* 117(3). 434–454. https://doi.org/10.1111/1467-968X.12173

Ke, Alan Hezao. 2022. Can agree and labeling be reduced to minimal search? *Linguistic Inquiry*. 1–22. https://doi.org/10.1162/ling_a_00481
Kleene, Stephen Cole. 1952. *Introduction to Metamathematics*. New York, NY: Van Nostrand.
Komachi, Mamoru, Hisatsugu Kitahara, Asako Uchibori & Kensuke Takita. 2019. Generative procedure revisited. *Reports of the Keio Institute of Cultural and Linguistic Studies* 50. 269–283.
Knowlton, Tyler, Tim Hunter, Darko Odic, Alexis Wellwood, Justin Halberda, Paul Pietroski & Jeffrey Lidz. 2021. Linguistic meanings as cognitive instructions. *Annals of the New York Academy of Sciences* 1500(1). 134–144. https://doi.org/10.1111/nyas.14618
Kripke, Saul A. 1979. A puzzle about belief. In Avishai Margalit (ed.), *Meaning and Use*, 239–283. Dordrecht: D. Reidel.
Kripke, Saul A. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
Kripke, Saul A. 2019. "Naming and Necessity revisited. Centre for Logic and Language Special Lecture". https://philosophy.sas.ac.uk/podcasts/naming-and-necessity-revisited-prof-saul-kripke (accessed 23 September 2024)
Kush, Dave W. 2013. *Respecting relations: Memory access and antecedent retrieval in incremental sentence processing*. Maryland: University of Maryland dissertation.
Kush, Dave, Jeffrey Lidz & Colin Phillips. 2015. Relation-sensitive retrieval: evidence from bound variable pronouns. *Journal of Memory and Language* 82. 18–40. https://doi.org/10.1016/j.jml.2015.02.003
Liao, Diana A., Katharina F. Brecht, Melissa Johnston & Andreas Nieder. 2022. Recursive sequence generation in crows. *Science Advances* 8(44). https://doi.org/10.1126/sciadv.abq3356
Leivada, Evelina & Elliot Murphy. 2021. Mind the (terminological) gap: 10 misused, ambiguous, or polysemous terms in linguistics. *Ampersand* 8. https://doi.org/10.1016/j.amper.2021.100073
Leivada, Evelina, Elliot Murphy & Gary Marcus. 2022. DALL-E 2 fails to reliably capture common syntactic processes. *Social Sciences & Humanities Open* 8(1). https://doi.org/10.1016/j.ssaho.2023.100648
Levin, Beth & Malka Rappaport Hovav. 2013. Lexicalized meaning and manner/result complementarity. In Boban Arsenijević, Berit Gehrke & Rafael Marín (eds.), *Subatomic Semantics of Event Predicates*, 49–70. Dordrecht: Springer.
Lewis, David. 1975. Languages and language. In Keith Gunderson (ed.), *Language, Mind, and Knowledge*, 163–188. Minneapolis, MN: University of Minnesota Press.
Lewis, Shevaun & Colin Phillips. 2015. Aligning grammatical theories and language processing models. *Journal of Psycholinguistic Research* 44. 27–46. https://doi.org/10.1007/s10936-014-9329-z
Lightfoot, David W. 2020. *Born to Parse: How Children Select Their Languages*. Cambridge, MA: The MIT Press.
Ludlow, Peter. 2003. Referential semantics for I-languages? In Louise M. Anthony & Norbert Hornstein (eds.), *Chomsky and His Critics*, 140–161. Oxford: Blackwell.
Lupyan, Gary. 2012. What do words do? Toward a theory of language-augmented thought. In Brian H. Ross (ed.), *The Psychology of Learning and Motivation*, 256–297. Waltham, MA: Academic Press.
Mandler, Jean Mandler. 2004. *The Foundations of Mind: Origins of Conceptual Thought*. Oxford: Oxford University Press.
Marcus, Gary Fred. 2001. *The Algebraic Mind: Integrating Connectionism and Cognitive Science*. Cambridge, MA: The MIT Press.
Marti, Genoveva. 2015. General terms, hybrid theories and ambiguity: a discussion of some experimental results. In Jussi Haukioja (ed.), *Advances in Experimental Philosophy of Language*, 157–172. London: Bloomsbury Academic.
Murphy, Elliot. 2012. *Biolinguistics and Philosophy: Insights and Obstacles*. London: Lulu.

Murphy, Elliot. 2015a. Labels, cognomes, and cyclic computation: an ethological perspective. *Frontiers in Psychology* 6. https://doi.org/10.3389/fpsyg.2015.00715
Murphy, Elliot. 2015b. Reference, phrases and individuation: topics at the labeling-interpretive interface. *Opticon1826* 17. https://doi.org/10.5334/opt.cn
Murphy, Elliot. 2015c. The brain dynamics of linguistic computation. *Frontiers in Psychology* 6. https://doi.org/10.3389/fpsyg.2015.01515
Murphy, Elliot. 2016. Phasal eliminativism, anti-lexicalism, and the status of the unarticulated. *Biolinguistics* 10. 21–50. https://doi.org/10.5964/bioling.9047
Murphy, Elliot. 2017. Acquiring the impossible: Developmental stages of copredication. *Frontiers in Psychology* 8. https://doi.org/10.3389/fpsyg.2017.01072
Murphy, Elliot. 2020a. *The Oscillatory Nature of Language*. Cambridge: Cambridge University Press.
Murphy, Elliot. 2020b. Language design and communicative competence: the minimalist perspective. *Glossa: A Journal of General Linguistics* 5(1). https://doi.org/10.5334/gjgl.1081
Murphy, Elliot. 2021a. Predicate order and coherence in copredication. *Inquiry. An Interdisciplinary Journal of Philosophy* 67(6). 1744–1780. https://doi.org/10.1080/0020174X.2021.1958054
Murphy, Elliot. 2021b. Persistence conditions of institutional entities: Investigating copredication through a forced-choice experiment. *Frontiers in Psychology* 12. https://doi.org/10.3389/fpsyg.2021.528862
Murphy, Elliot. 2021c. *Linguistic Representation and Processing of Copredication. PhD thesis*. London: University College London dissertation.
Murphy, Elliot. 2022. Copredication and complexity revisited: A reply to Löhr and Michel. *Cognitive Science* 46(10). https://doi.org/10.1111/cogs.13207
Murphy, Elliot. 2023. Animals sick with language: From syntax to socialism in Nietzsche. In Matthew McManus (ed.), *Nietzsche and the Politics of Reaction: Essays on Liberalism, Socialism, and Aristocratic Radicalism*, 263–286. London: Palgrave Macmillan.
Murphy, Elliot, Emma Holmes & Karl Friston. 2024. Natural language syntax complies with the free-energy principle. *Synthese* 203(154). https://doi.org/10.1007/s11229-024-04566-3
Murphy, Elliot & Evelina Leivada. 2022. A model for learning strings is not a model of language. *PNAS* 119(23). https://doi.org/10.1073/pnas.2201651119
Murphy, Elliot & Jae-Young Shim. 2020. Copy invisibility and (non-)categorial labeling. *Linguistic Research* 37(2). 187–215. https://doi.org/10.17250/khisli.37.2.202006.002
Murphy, Elliot, Oscar Woolnough, Patrick S. Rollo, Zachary J. Roccaforte, Katrien Segaert, Peter Hagoort & Nitin Tandon. 2022b. Minimal phrase composition revealed by intracranial recordings. *Journal of Neuroscience* 42(15). 3216–3227. https://doi.org/10.1523/JNEUROSCI.1575-21.2022
Neta, Ram. 2003. Contextualism and the problem of the external world. *Philosophy and Phenomenological Research* 66(1). 1–31. https://doi.org/10.1111/j.1933-1592.2003.tb00241.x
Nichols, Shaun, N. Ángel Pinillos & Ron Mallon. 2015. Ambiguous reference. *Mind* 125. 145–175. https://doi.org/10.1093/mind/fzv196
Obata, Miki, Samuel Epstein & Marlyse Baptista. 2015. Can crosslinguistically variant grammars be formally identical? Third factor underspecification and the possible elimination of parameters of UG. *Lingua* 156. 1–16. https://doi.org/10.1016/j.lingua.2014.12.003
Parr, Thomas, Giovanni Pezzulo & Karl J. Friston. 2022. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. Cambridge, MA: The MIT Press.
Piatelli-Palmarini, Massimo. 2020. Minds with meanings (pace Fodor and Pylyshyn). *Rivista Internazionale Di Filosofia e Psicologia* 11(1). 1–18. https://doi.org/10.4453/rifp.2020.0001

Pietroski, Paul M. 2018. *Conjoining Meanings: Semantics Without Truth Values*. Oxford: Oxford University Press.
Pietroski, Paul M. 2022. "One word, many concepts: endorsing polysemous meanings." https://www.terpconnect.umd.edu/~pietro/research/papers/OneWordManyConcepts.pdf (accessed 23 September 2024)
Pinker, Steven. 2007. *The Stuff of Thought*. NY: Viking Press.
Pritchard, Tim. 2022. Proprietary linguistic meaning. *Synthese* 200. 1–20. https://doi.org/10.1007/s11229-022-03776-x
Pustejovsky, James & Olga Batiukova. 2019. *The Lexicon*. Cambridge: Cambridge University Press.
Putnam, Hilary. 1975. The meaning of 'meaning'. In Keith Gunderson (ed.), *Language, Mind and Knowledge*, 131–193. Minneapolis, MN: University of Minneapolis Press.
Putnam, Hilary. 1992. Replies. *Philosophical Topics* 20. 347–408.
Radford, Andrew. 2016. *Analysing English Sentences*. Cambridge: Cambridge University Press.
Ramchand, Gillian & Peter Svenonius. 2014. Deriving the functional hierarchy. *Language Sciences* 46 (B). 152–174. https://doi.org/10.1016/j.langsci.2014.06.013
Ramstead, Maxwell J. D., Karl J. Friston & Inês Hipólito. 2020. Is the free-energy principle a formal theory of semantics? From variational density dynamics to neural and phenotypic representations. *Entropy* 22(8). 1–30. https://doi.org/10.3390/e22080889
Rast, Erich. 2023. Metalinguistic disputes, semantic decomposition, and externalism. *Linguistics and Philosophy* 46. 65–85. https://doi.org/10.1007/s10988-022-09357-y
Rehder, Bob & Shinwoo Kim. 2010. Causal status and coherence in causal-based categorization. *Journal of Experimental Psychology: Learning, Memory and Cognition* 36(5). 1171–1206. https://doi.org/10.1037/a0019765
Rizzi, Luigi. 2013. Introduction: Core computational principles in natural language syntax. *Lingua* 130. 1–13.
Rizzi, Luigi. 2016. Monkey morpho-syntax and merge-based systems. *Theoretical Linguistics* 42(1–2). 139–145. https://doi.org/10.1515/tl-2016-0006
Roberts, Ian G. 2019. *Parameter Hierarchies and Universal Grammar*. Oxford: Oxford University Press.
Russell, Gillian. 2010. Analyticity in externalist languages. In Sarah Sawyer (ed.), *New Waves in Philosophy of Language*, 186–205. London: Palgrave Macmillan.
Shim, Jae-Young. 2022. "MERGE and pair-Merge of heads". https://scholar.kyobobook.co.kr/article/detail/4010036888790 (accessed 23 September 2024) https://doi.org/10.14342/smog.2022.115.1
Smith, Peter W., Johannes Mursell & Katharina Hartmann (eds.). 2020. *Agree to Agree: Agreement in the Minimalist Programme*. Berlin: Language Science Press.
Sobin, Nicholas. 2014. Th/Ex, Agreement, and Case in expletive sentences. *Syntax* 17(4). 385–416. https://doi.org/10.1111/synt.12021
Spelke, Elizabeth. 2010. Innateness, choice, and language. In Julie Bricmont & Jean Franck (eds), *Chomsky Notebook*, 203–210. New York: Columbia University Press.
Sprouse, Jon & Diogo Almeida. 2013. The empirical status of data in syntax: A reply to Gibson and Fedorenko. *Language and Cognitive Processes* 28(3). 222–228. https://doi.org/10.1080/01690965.2012.703782
Strawson, Galen. 2008. *Real Materialism and Others Essays*. Oxford: Oxford University Press.
Strawson, Peter F. 1950. On referring. *Mind* 59(235). 320–344.
Sugimoto, Yushi. 2022. *Underspecification and (im)possible derivations: Toward a restrictive theory of grammar*. Ann Arbor: University of Michigan dissertation.
Tobia, Kevin P., George E. Newman & Joshua Knobe. 2020. Water is and is not H2O. *Mind & Language* 35(2). 183–208. https://doi.org/10.1111/mila.12234

Todorović, Neda. 2020. Aspect interacts with phasehood: Evidence from Serbian VP-ellipsis. In Jeroen V. Craenenbroeck, Cora Pots & Tanja Temmerman (eds.), *Recent Developments in Phase Theory*, 11–34. Berlin: De Gruyter.

Trotzke, Andreas & Jan-Wouter Zwart. 2014. The complexity of narrow syntax: minimalism, representational economy, and simplest Merge. In Frederick J. Newmeyer & Laurel B. Preston (eds.), *Measuring Grammatical Complexity*, 128–147. Oxford: Oxford University Press.

Truswell, Robert. 2007. *Locality of wh-movement and the individuation of events*. London: University College London dissertation.

van Urk, Coppe. 2020. How to detect a phase. In Jeroen V. Craenenbroeck, Cora Pots & Tanja Temmerman (eds.), *Recent Developments in Phase Theory*, 89–130. Berlin: De Gruyter.

Valente, Matheus. 2019. Communicating and disagreeing with distinct concepts: a defense of semantic internalism. *Theoria* 85. 312–336. https://doi.org/10.1111/theo.12201

Wallace, David Foster. 1997. *Infinite Jest*. London: Abacus.

Wallace, David Foster. 2011. *The Pale King*. London: Pegnuin.