



Internalist priorities in a philosophy of words

John Collins¹ 

Received: 23 November 2021 / Accepted: 2 February 2023 / Published online: 15 March 2023
© The Author(s) 2023

Abstract

Words appear to be denizens of the external world or, at any rate, not wholly mental, unlike our pains. It is the norm for philosophical accounts of words to reflect this appearance by offering various socio-cultural conditions to which an adequate account of wordhood must cleave. The paper argues, to the contrary, that an adequate account of word phenomena need avert to nothing other than individual psychology along with potential external factors that in-themselves do not count as linguistic. My principal leverage will be that, by everyone's lights, whatever words are, they are syntactically combinable and possess structural properties. But such conditions cannot be externally realised; instead, they are aspects of our internally realised cognitive capacity. It will also be argued, however, that the position is consistent with much of our common lore about words, albeit sans an externalist linguistic ontology.

Keywords Externalism · Internalism · Ontology · Syntax · Words

1 Introduction

A standard view is that language, and so words, are social in nature. Thus, Wittgenstein (1953), Quine (1948), Lewis (1969, 1975), Burge (1989), and many others echo the pre-generative structuralist tradition in linguistics exemplified by Bloomfield (1933, p. 198): 'the facts of language are facts of social, not of individual psychology'. My argument will be, to the contrary, that an adequate account of word phenomena need avert to nothing other than individual psychology along with potential external factors that in-themselves do not count as linguistic. My core claim is that, by everyone's lights, whatever one's idea of a word is, words are syntactically combinable into hierarchical structures and possess inherent structural properties; yet such conditions are not externally realised in linguistic vehicles (speech, script, hand gestures), which

✉ John Collins
John.Collins@uea.ac.uk

¹ School of Politics, Philosophy, Language & Communication, University of East Anglia, Norwich NR4 6XE, UK

are linear, save if conceived through the lens of the cognitive resources we deploy when producing or consuming the vehicles *as* linguistic.

A question arises about the appropriate modality for this claim: *could* such conditions be externally realised in some medium? That is, for explanatory purposes, is it possible for externalia to be appropriately structured in ways that cognition respects or represents rather than projects onto? I shall not attempt to answer this question, although, if I am right, no linear medium would suffice, and it is unclear, given the facts of human psychology, how a relevant medium could be constructed and produced and consumed in the way speech and sign are.

Ontologically speaking, I offer an eliminativist position, for words, as commonly conceived, are unrequired for any explanatory purpose. Still, we shall see how to preserve the truth of much—although not all—of what we want to say about words from an intuitive or common-sense perspective. The crucial rider here is that what makes true such common lore is not an ontology of words, but only the exercise of our linguistic capacity that is internally constituted. On this view, there is no strict individuation of words, but only context-sensitive decisions on when speaker-hearers' internal states count as enough alike for the purposes at hand; what are invariant are the relevant linguistic properties that enter into the specification of mental states that support linguistic competence, but such properties can only be what our best theories identify.

It bears emphasis upfront that my present concerns are narrow. Ontological worries have been raised for all sorts of 'manifest' entities, whether intuitively material, such as tables and trees, or more abstract, such as numbers, propositions, money, etc. Officially, I take no stand on these extra-linguistic ontological matters. My argument proceeds from the fact that there is an on-going scientific programme that accounts for word-phenomena in an internalist manner at apparent odds with our externalist intuitions about words. There is no corresponding programme for the other examples of manifest ontology; for instance, there is no science of statues or tables or even of numbers whose explanations require an ontology at odds with any intuitive take on the objects; indeed, in the case of numbers and money, say, it is wholly unclear what ontology, if any, the folk uphold. In this light, even if someone of an externalist bent towards words could fashion a general deflationary approach to ontology that found all kinds of social objects acceptable, such as Thomasson (2010, 2018) advances in general, and Tasker (2022) adopts for words, my case would remain unmolested.

The next section will clarify some basic notions. After that, the general problem of word individuation will be presented as it infects an externalist construal of words. This will lead onto my setting out the combinable and structural properties of words, where we shall see that they cannot be plausibly externally realised. To rub the point in, various objections to the argument will be discussed. To finish, I shall show how hypothesised internal structures can be seen to explain and so preserve much (but not all) of the 'manifest image' of words sans any ontological commitment to linguistic externalia.

2 Clarifications

Three clarifications are in order before beginning in earnest.

First, as is now known, word segmentation of an acoustic stream is a matter of construction, or of the ‘mind building words’, as Isac and Reiss (2013, p. 31) put it. This is because the ‘slices’ in the acoustic stream that mark phonemic distinctions, and so word boundaries, have no correlate in the stream itself. Put ontologically, sounds understood phonologically, as discrete units that make up words, are not sounds understood acoustically, which have no discreteness relative to a host stream. This fact would damn any externalism that identified words with sounds. I shall, however, side line this consideration, for even under a phonemic idealisation where sounds are discrete, the ideal strings would remain linear and so not apt to realise syntax in the appropriate sense; further, words can have a discrete external vehicle, as it were, albeit not acoustic (sign or text, say).

Secondly, by ‘linguistic’, I mean, in particular, syntactic, phonological, and semantic properties. Just what these properties are is, of course, disputed, but my argument will proceed by only assuming that they minimally involve the (non-linear) combinability of words and selectivity relations between words. One is free to have a richer conception of the linguistic realm, encompassing, say, orthography and speakers’ reflexive intentions (e.g., that a speaker intends an audience to recognise their audience-directed intentions). On the first count, say, many hold that semantic properties involve relations between symbols and entities in the world; so, if semantics is essential to word-phenomena, then an internalism can’t be the whole explanatory story. Still, once it is acknowledged that the structural properties advertised are both essential to words and not externally realisable, then words are not externally individuated, even if certain semantic aspects of words do turn out to be world-involving (if semantic externalism is true). On the second count, first, orthography, hand gestures, and sound waves are not essentially linguistic, for no invariance of such properties shows up in linguistic tokening, as is most easily appreciated by noting that one can rehearse a word internally with it counting as the same word one may utter or write down. I shall be at pains to show, however, that we can still treat as true a lot of common word-talk that does appeal to orthographic and other external properties; it is just that such claims are not made true by an ontology of words.

Likewise, it may be thought that certain intersubjective conditions relating to intentions to communicate are essentially linguistic. We may happily grant that speakers do have such intentions, and even that they are unavailable to non-linguistic organisms. Yet so much would only indicate that language enables certain cognitive capacities, in concert with other systems, such as theory of mind capacity. Such intentions might seem essential for the fixation of a robust same-word relation. I deny the need for any such relation, however (Sect. 7).

Thirdly, a natural objection to any species of linguistic internalism is that it presupposes a far too strict notion of what counts as external or, alternatively, just defines the external as not the internal, rendering the positive view ultimately unclear. Regarding the external as, say, what physics might be able to identify appears too strict, ruling out nations and football teams, say. Equally, that entities have mind-involving individuation conditions is insufficient to support a relevant internalism, for much of ‘ordinary’

ontology involves such conditions. As it is, for present purposes, I am happy to let externalism be characterised via a notion of mind-independence:

Externalism An entity e is external iff e necessarily involves the instantiation of a mind-independent property, i.e., one whose obtaining does not require the exercise of any human cognitive capacity.

So, atoms and molecules count as external, but so do tables and chairs, for while human cognition is required for their conception and manufacture, some raw material needs to be instantiated, which is not mind-dependent. Similarly, all kinds of cultural artefacts count as external, for they require some public medium beyond human cognition, although many problems arise here with particular cases.¹ Other humans also count as external insofar as their existence does not depend upon others' cognitive capacity. About some entities (numbers, say), we do not know if they are external or not. At any rate, as we shall see, some extant brands of linguistic externalism are happy to have linguistic entities supervene on internal states of speaker-hearers and mind-independent properties variously construed. My species of internalism says that no specific mind-independent properties are necessary for the explanation of word phenomena, although some such properties will be involved in the explanation of all kinds of complex phenomena involving language, such as communication or reading and writing.

An interesting question arises about the relation of other minds external to the mental states of a given speaker-hearer; this issues will be addressed in Sect. 7.3.

In the next section, I shall raise a general individuation problem for words, which *prima facie* affects all accounts save for the various stripes of eliminativism. Ultimately, I shall contend that the individuation problem dissolves on my internalist approach, for there are no discrete entities to individuate, even though we may still judge word talk to be true or false (Sect. 7).

3 The common-sense ontology of words and the problem of their individuation

Miller (2020a, 2020b) offers some plausible desiderata on any adequate account of the common notion of a word: words should be *expressible*, *creatable*, and *evolvable*. These conditions, without further ado, should apply to both types and their tokening. It might seem that such conditions rule out certain metaphysical options, for they appear to render words as, effectively, communicative artefacts. It might be, however, that being abstract and creatable are not inconsistent (Irmak, 2019), and a range of options are available for how tokens stand towards types; for example, types might be abstracta that tokens represent (Bromberger, 1989; Nefdt, 2019; Szabó, 1999). My present aim is not to adjudicate on these internecine metaphysical disputes. What is broadly uncontested is that all views that seek to capture the intuitive notion of a word face an individuation problem.

¹ One might imagine stories and pieces of music, say, to require an external medium, but the cases are not so straightforward. One may construct either in one's mind without putting them on paper or there being otherwise performed.

An adequate individuation of a word W will tell us, for any tokens x and y , if they count as instances of the type W . What is wanted is some feature(s) of a word, which, if shared between x and y , will make them both W -type things. The problem is that no class of features appears to suffice for all cases (cf., Feinsinger, 2021; Miller, 2019). For example, two tokens of the form /'bænk/ might count as the same word if our concern is transcription or pronunciation, but not semantically or syntactically, for the phonemic form is ambiguous between riversides and financial institutions, at least for the nominal. It will not do to say that there is the one word with distinct meanings or syntactic properties, for semantics matters to word individuation, lest a word is just an empty sign. Syntactic position also matters as a determinant of what a word might mean, a matter to which I shall return at below (Sect. 4).

Sameness of meaning will not suffice, either, of course, for then all synonymies will count as the same word, both within and across languages. A tempting move is to individuate a word in terms of both semantic and morphophonemic properties. A problem now arises for the individuation of these properties. Variation between speaker-hearers in matters of construal and pronunciation precludes any simple appeal to the relevant properties. Such variation also obtains between occasions. Is an auctioneer talking ten to the dozen or a drunk slurring his speech using the same words as the sober newsreader? In one sense, yes, for the drunk, say, might remember what he said. Still, intention to utter certain words is neither necessary (not intending to cuss in church is no excuse) nor sufficient (words can always come out wrong, such as with spoonerisms) for words to be tokened, and what was produced might be unrecognisable as linguistic (cf., Stojnić, 2021). To be sure, here as elsewhere, there are ways to idealise away from complexity, and admitting vagueness is not to banish distinctions. Still, doubt is cast on the enterprise of individuation, if the application of the relevant individuating factors are interest-relative, or supererogatory to any explanatory project.

The general point has been made forcefully by Chomsky (1986), who argues that *all* external linguistic notions, whether of words or languages, are conditioned by matters of interest and value. If so, then linguistic externalia ought not to attract serious ontological commitment, i.e., a commitment that issues from explanation of phenomena rather than colloquial ways of talking. The status of words might be similar to that of the sky. We talk of the sky, and do not deny its existence, but the sky is not a thing that calls for analytical clarification or theoretical explanation. All we want to say about the sky, if true, can be explained without appeal to the sky as something requiring individuation. In this light, the idea that words are repeatable, intersubjective entities, depending on speakers' mental states, but not identifiable with them, is erroneously to characterise the phenomenon of communication via an ontology that is unrequired for its explanation. I shall return to this point in proposing a general divorce between ontological commitment and truth (Sect. 7.2).

The individuation problem does not necessarily sink all metaphysical efforts to individuate words as externalia, but it should encourage suspicion about the goal itself. A general externalist orientation may be preserved by backing off on the demands of individuation; after all, preciously few things are individuatable to the satisfaction of all philosophers. Perhaps we just ought to live with a certain looseness, or even deflate our metaphysical ambitions (Thomasson, 2010, 2018). There is, however, a

grave problem for the very idea of words as ordinarily conceived as some kind of external entity, quite apart from individuation worries.

Whatever words are, by everyone lights, they must be syntactically combinable and possess structural properties that go to determine their potential interpretation in host structures. Yet neither of these properties appear to be externally realised. The natural way to put this claim is that words are partly constituted by how speaker-hearers represent and process them, which are matters for individual psychology. I want to press this point further and claim that all the theoretical action concerning word phenomena resides in individual psychology, and any externalia that may be said to be a word (ink marks, an utterance), is only such because of the cognitive states involved in its consumption or production. In this light, what is essential to words turns out to be phenomena that render words as reflections of cognitive capacity rather than entities meriting serious ontological commitment.

My ultimate goal is to live happily with indeterminate context-sensitive individuation of words construed as external entities while being able to explain how external media get to be conceived of as linguistic. In short, we shall retain some if not all of our linguistic manifest image while looking internally for explanation, which is exactly where the manifest image nudges us to look once we consider the structural conditions any notion of a word must satisfy, over and above the socio-cultural aspects we ordinarily impute to words.

4 Two linguistic conditions on words: combinability and selectivity

In the previous section, I entertained, after Miller (2020a, 2020b), three intuitive conditions an account of words must satisfy: words should be expressible, creatable, and evolvable. These factors reflect socio-cultural aspects of words as things populations of speakers use over time. In this section, I shall spell out two conditions on words that reflect the fact that they belong to, or issue from, a linguistic system that treats words as combinable in ways that systematically affect the semantics of the resulting structures. Words, after all, are not traffic lights or styles of dress, still less ‘natural meanings’ in Grice’s (1989) sense.

In theoretical linguistics, the notion of a word is typically suspected of yoking together distinct conditions that should receive separate explanations (Di Sciullo & Williams, 1987; Pustejovsky & Batiukova, 2019). The notion of a *lexical item* tends to substitute for a word, the former including various affixes and roots that might not occur overtly as individual expressions in morphology. Furthermore, lexical items include unpronounced covert items (e.g., PRO as the subject of infinitives). How morphology works to issue in surface form is currently controversial. In so-called distributed morphology, there is no lexicon as traditionally conceived as a list of exceptions; morphology is built up in tandem with the syntactic derivation being defined over root items (*locus classicus*, Halle & Marantz, 1993). This invites a similar attitude towards syntactic and semantic information a lexical item encodes. At one extreme, Borer (2005) conceives of a lexical item as simply an index or root that is linked to extra-linguistic information in long-term memory, containing no essential syntactic information or selectional restrictions (see below), both being captured by

general structural forms the items incarnate. The traditional view assumed far richer lexical items from which both syntactic structure and compositional interpretation are projected under general constraints.² Hale and Keyser (1993, 2003) elaborate a highly influential middle ground, where selection is configurational, but items still retain syntactic and semantic information that constrains ‘sentential’ syntax. For our purposes, we may prescind from these disputes. I shall simply assume that words or lexical items, whatever they are by everyone’s lights, are subject to syntactic conditions and structural semantic interpretation. If it turns out that the ‘best theory’ finds no place for lexical items as currently variously conceived or anything else dimly related to our intuitive notion of a word, then so be it. For my wider dialectical purposes, all I require is that word phenomena are explained by cognitive invariances internally characterised.

The first condition insists that words are open to syntactic operations:

Combinability: Whatever words are, they are essentially combinable into structural units that inherit their identities from the constituent words and how they are organised.³

Whatever a word is, it is combinable with other words to form phrases (including what we colloquially call sentences) that themselves may combine with other words and phrases. The principles governing such combinatorics we call syntax. Of course, we say much that is not syntactic/combinable, such as ‘um’s and ‘ah’s and various expletives. Indeed, normal speech is highly degraded and fragmented relative to syntactic principles (false starts, jumbled words, unfinished phrases, etc.). This does not cast doubt on syntax, for the principles enter into an explanation of consumption without the errors, and enter into an explanation of the shape of the characteristic errors. As Chomsky (1965) argued, if we want a theory of performance, we still need a theory of competence from which perspective performance is an interaction effect of dedicated linguistic systems and more general extra-linguistic systems that determine the contribution of working memory, attention, etc. to linguistic performance.

Rather than assuming a specific theory of syntax, I shall rest content with a core aspect any theory must respect: syntax determines units of combined lexical items that are not identifiable or individuated in terms of linear order or any other perceptible property associated with morphophonemic form. The fundamental reason for this is that external media (speech, sign, text) are linear and so possess an associative structure, which preserves the order of items. Thus, concatenating ‘a’ with ‘b’ to produce ‘a[∩]b’, and then right-concatenating with ‘c’ to produce ‘a[∩]b[∩]c’ does not preserve ‘a[∩]b’ as a constituent, for the exact same structure is produced by ‘b[∩]c’ being left-concatenated with ‘a’. In distinction, syntax is non-associative, and so not linear: the structure of a whole reflects the constituents that were combined to form the whole regardless of their order in the whole (Collins, 2011). Words of a sentence

² This has been the traditional view since Katz and Fodor (1963), Katz and Postal (1964) and Chomsky (1965), and has been pursued in various ways; see, for example, and Levin and Rappaport Hovav (1994). For a detailed overview and problems for the general approach, see Levin and Rappaport Hovav (2005).

³ I take this thesis to be trivially true, although there might be thought to be exceptions, such as salutations and exclamatives (Quirk et al., 1985). Whether there is covert structure in such cases is an open question (see, for example, Stainton, 2006). That issue aside, items usable as single-word utterances are combinable:(i) a They said their goodbyes.b The helloing went on for hours.c Everyone was like wow.

are not beads on a string, but relate to one another in hierarchical terms, forming constituents.

To make the point concrete, take the sentence:

The brackets mark the constituents. So, *tall man* is a constituent (a noun phrase), as is *the tall man* (a determiner phrase), and *the tall man, who likes women* (a determiner phrase including a relative clause), but *the tall* is not, and neither is *tall man who*, etc. Just what labels are the right ones is an open issue, but some labelling is required, which records the identity of the constituent phrase.

The bottom-line is that an account of words must render them as items that can be combined into phrases. As we shall see, words conceived as externalia do not meet this demand. This is not to say that hierarchical structure cannot be realised externally; I assume it is realised all over the place in the organisation of organic structure, for example. The point is only that hierarchical structure is not realised in the linear external media for language, which is essentially associative, whereas syntax is essentially non-associative.

The second condition on words that issues from the nature of language relates to semantics:

- (1) a The tall man, who likes women, likes himself
 b [TP [DP [DP The [NP tall man]][CP who [VP likes women]]][T -s [VP likes himself]]]

Selectivity Whatever words are, as combined, they are interpreted relative to being selected by other words.

On the classic way of satisfying this condition, verbs (and other categories, but we'll stick with verbs for purposes of exposition) have a certain number of 'roles' to assign, and a properly interpretable structure results iff each expression with which the verb is combined is uniquely assigned a role (see Levin and Rappaport Hovav, 2009, for overview). For example:

- (2) a *Sam sneezes Mary
 b *Sam kissed Mary the table
 c *Bill persuaded to leave

The verb *sneeze* assigns a single role—an agent of the activity—and so (2a) is unacceptable because it includes a nominal (*Mary*) without a role. We might, of course, conjure up a reading, imagining a context where people induce sneezes in others, but then we have effectively created a new meaning for the word, for on this construal, *Sam sneezes* would not be true unless two individuals were involved. Similarly, *kiss* assigns agent and patient (thing affected) roles, and so *the table* in (2b) lacks a role. In (2c), *persuade* requires a nominal specifying the patient of Bill's persuasion.

There is great controversy over how such facts should be accommodated. One clear problem is that many verbs can have variable argument assignments without any change of meaning:

- (3) a Bill kicked Sam the ball
 b Bill kicked Sam
 c Bill kicked (out)

This encourages the thought that selection is a function of configuration (a result of combination) rather than an inherent property of the putative selecting item.

For our purposes, the mere facts are what is relevant rather than any theoretical account of the phenomena. Words are interpreted in combination, and how they are combined determines interpretational options.

In the next section, I shall spell out the problems these conditions on words (whatever words are) pose to an externalist model.

5 The problems for externalism

There are many ways of accounting for words as externalia. Kaplan (1990, 2011) thinks of words in terms of a ‘common currency’, as essentially historical entities. Katz (1981; cf. Katz & Postal, 1991) considers them to be platonic abstracta, much like numbers. Nefdt (2019) considers words to be like numbers too, although on a structuralist conception of mathematics. Others consider words to be artefactual abstracta, much like pieces of music (Irmak, 2019). Devitt (2006) is more nominalist: words are tokens that satisfy certain high-level functional roles. Miller (2019) suggests that they are bundles of properties, both mind internal and mind external.⁴ Millikan (2005) and Feinsinger (2021) think of words as social entities supported by community co-ordination. Others have non-externalist positions much closer to the kind I favour. Rey (2003, 2020) argues that words and all other linguistic entities are intentional inexistents; that is, speaker-hearers have representations of words that do not represent anything extant at all, but the theorist adopts a pretence attitude towards the objects of the representations to save herself the awkwardness of referring to the speaker-hearers representations of such and such and so and so, much as the vision scientist does in referring to triangles and colours instead of representations of them. The basic point is that all explanations proceed over the representations regardless of the ontological status of the represented. Azzouni (2013) similarly but independently argues that there are no public linguistic entities, but we involuntarily perceive things to be linguistic.

I shall depart from the externalist views, but my disagreement with Rey is more a friendly amendment over how we ought to think of the nature of cognitive representation and the linguistic manifest image. Our disagreement is ideological, not ontological. To be sure, Rey’s claim may be construed very thinly, where word representations are simply a way of specifying the phenomena to be explained. I take him to mean something more robust, however, where the putative representations are on explanatory duty in entering into the specification of the internal states in the form of a ‘common coin’ that constitutes the interfaces between distinct systems that in concert explain the phenomena.

The externalist views share the thought that while the existence of words and some of their properties might depend or supervene on speaker-hearers’ internal states, the words themselves are not such states. Words here are variously conceived of

⁴ I am sympathetic to Miller’s position as a take on how we think of words, and it is not out-and-out externalist. One motivation for his bundle approach is to make sense of the fact that we speak of the same word being both thought and written (*I had just thought of that very word before you wrote it*), which is not easily accommodated on either a wholly internalist or externalist view. I think, however, as will be explained, that this consideration can be dealt with in a wholly internalist manner (see Sect. 7).

as being public, tokened in external media, possessing a history independent of the minds of speaker-hearers that employ them, and so on. Space precludes any significant discussion of the various ways just intimated as to how such an externalism might be developed and defended. Fortunately, for my purposes, there is a general problem that infects all of the views: they do not satisfy the combinability and selectivity conditions on words. As Bromberger (2011, p. 490) avers:

Any worthwhile conception would take aboard that words function as constituents of phrases and sentences. It would acknowledge that they play their defining roles merged with other terms, and thus, that—whatever their intrinsic perceptual and referential features—it is of the essence of words that they can appear in juxtapositions through which they receive and assign thematic roles, and stand in various functional relationships.

Bromberger is right, and his basic thought can be readily generalised.⁵ The form of the argument is as follows:

Argument:

- (i) Whatever words are, they can occur in phrases (both token and type).
- (ii) As so occurring, words are subject to the conditions of combinability and selectivity.
- (iii) But the conditions only apply to internal states, not to externalia.
- (iv) Words, therefore, not be externalia.

I shall take (i)–(ii) as given. The crucial premise is (iii), and the inference to (iv) might be resisted in various ways, all of which I shall seek to dispel. First, then, let me defend premise (iii).

There are two species of externalism to consider. First consider views that conceive of words not as abstracta, but as concretely realised in some way or other, or, if types, as conditions that tokens may realise. The problem, simply put, is that even if we imagine words to be concretely realised, syntax is not, and so a constitutive condition on would-be external words is not itself externally realised.

Syntax, at its most simple, is the hierarchical organisation of lexical items into further combinable units. Yet the units are in no sense realised in any external medium. As noted above, the point here is not that hierarchical organisation *tourt court* cannot be realised externally—of course it can be. The point, rather, is that any external medium for language must be linear insofar as we cannot produce or consume a number of signs simultaneously but only in some sequence. This makes any possible segmentation of the medium associative under which order is preserved. Yet syntax is non-associative.

To see the point, consider a familiar case of ambiguity (structural homophony):

- (4) a Mad dogs and Englishmen go out in the midday sun
 - b [[Mad dogs] and Englishmen]...
 - c [Mad [dogs and Englishmen]]...

The ambiguity of (4a) can be resolved in two ways, with either *mad* modifying the first conjunct (*dogs*) or the conjunctive phrase (*dogs and Englishmen*). So, the adjective

⁵ As Bromberger (2011, p. 503) notes, philosophers need to ‘turn their metaphysical binoculars around to see the facts more clearly’, i.e., away from the external or abstract to human beings’ internal states.

is either combined with the nominal, *dogs*, the result of which is then merged with the conjunction (*and Englishmen*), or the adjective is merged with the conjunctive phrase (*dogs and Englishmen*). This difference, however, is in no sense essentially marked in (4a) or otherwise constituted or registered in any external media that might be said to realise the sentence.⁶

This is not to suggest that properties of externalisation, such as intonation and accent, can't signal structural information. For some constructions, intonation marks out the scope of a focus item, such as *only*, *always*, or *just*.⁷ For other constructions, focus accent can mark modification:

- (5) a The big big car
 b The [BIG BIG]_F car
 c The [BIG]_F big car

If both occurrences of the adjective are focused, then the car is simply really big, which indicates that the structure is read flat. If the higher occurrence is focused, then the big car among the big cars (the biggest) is being specified, which has the structure:

- (6) The [big [big car]]⁸

These kinds of phenomena, however, do not suggest, let alone show, that syntactic structure is always marked externally. In fact, it is more appropriate to say that it never *really* is; rather, sometimes external features indicate what structure is being used and what discourse mechanisms are used to resolve decisions. Consider a case where there simply is no phonological guide:

- (7) a Who does Bill want to succeed?
 b (which person *x*)(Bill wants *x* to succeed)
 c (which person *x*)(Bill wants to succeed *x*)

With irrelevant details elided, the standard explanation of the ambiguity is that *who* has two potential launch sites: as the SPEC of the infinitive intransitive *to succeed* or

⁶ One might find dispositionalism tempting, according to which the structural properties are 'secondary', explained by a string's 'primary' external properties along with its social/functional role in expressing whatever meaning it has for a population (cf., Devitt, 2006). Let me register a basic problem. If syntax is a dispositional property of a string, we should be able to say what the primary properties are that have the secondary effects, but this looks to be impossible given the modality-independence of syntax, i.e., it is the same syntax across text, utterance, sign, and internal monologue. In short, even if we were to think of syntax as a 'secondary' property, it turns out that the realisation of it has precious little to do with the external object, and everything to do with us. It might be thought that the primary properties could be 'idealized' phonological or other vehicular properties (Rey, 2020), but then the syntax would be tied to a particular modality, but there are no invariant or 'normal' or even idealised external properties that serve as the vehicle for syntax.

⁷ For example:(i) a Mary only gave Bill flowers.b Mary only gave [BILL]_F flowers.c Mary only gave Bill [FLOWERS]_F (ia) is ambiguous with respect to the alternatives being excluded. Thus, (ib) means that no-one other than Bill was given flowers by Mary, although Mary might have given Bill other things too, whereas (ic) means that nothing other than flowers was given to Bill by Mary, although Mary might have given flowers to other people. See Beaver and Clark (2008) for extensive discussion.

⁸ See Cinque (2010) for discussion of the interplay of adnominal modification and relative clauses in these kinds of construction. For present purposes, the phenomena rather than their explanation is what matters.

as the complement of the transitive verb *succeed*. The lower positions are construed as variables bound by the fronted *wh* item. Given that these two positions are phonologically null in (7a), the form fails to determine which position *who* binds. This entails that *only* the two readings are available, since *who* must bind one of the positions, i.e., (7a) has no vacuous reading corresponding to (8):

(8) (which person *x*)(Bill wants *y* to succeed *z*)

This can be observed by noting that if the object position of *succeed* is overtly occupied, then *who* must bind the SPEC position of the infinitive:

(9) Who does Bill want to succeed Sam?

In other words, syntactically speaking, in (7a) *who* must move from one of the lower positions of the subordinate TP:

(10) a [CP *who* does [TP Bill want [TP **who** to succeed]]]
 b [CP *who* does [TP Bill want [TP PRO to succeed **who**]]]

Note that in these cases and innumerable others we are not concerned with any fact to do with sentences as might be taken to exist independently of the mental representations of speaker-hearers. We want to explain a *modal* fact: why can the relevant constructions *only* be interpreted in exactly two ways, not fewer or more?

It might seem that a Platonic externalism will evade these problems, for on such positions, syntax has no concrete realisation; tokens in a medium only count as linguistic because of their relation to abstract types. The move certainly does relieve the externalist of certain burdens, but it also denudes syntax of its explanatory role. The phenomena syntax explains, as indicated, are not necessary truths akin to theorems, whose truth (let's assume) are somehow independent of any empirical facts. The modality to which syntax pertains involves the possible scope of an underlying system of principles, not anything we happen to do or how things are independent of us. Thus, even if we were to suppose that syntax is externally constituted outside of space and time, let alone the minds of individual speaker-hearers, all questions would remain open as to why such abstract structures constrain linguistic cognition.

This objection is not a version of the standard epistemological complaint against Platonism: how can we know such abstracta? Platonically speaking, syntax could be very different from whatever explains the systematicity of our linguistic cognition, or not exist as a recursive combinatorial operation at all. There is enough room in Plato's heaven for clutter. The pressing questions are why the organisation of words is systematic in precisely the way it is, and why we all acquire languages with just that structure, and can't acquire other kinds of conceivable languages. This sharply distinguishes language from mathematics as an object of study, the chosen analogy for the Platonist (cf., Katz, 1981; Katz & Postal, 1991). In more general terms, mathematics, let's assume, concerns the truth about numbers, functions, etc., not mathematicians. With syntax, the questions are invariably about us.

This might be a mere appearance, if, for whatever reason, we are good at mainlining the truth out there when it comes to language (just as we are with mathematics, it seems). There is a crucial disanalogy. We don't a priori discover syntactic truths, as we do in mathematics about numbers, etc. (let's suppose). It is, for example, not a

(logically or metaphysically) necessary truth that (4a), qua linear form, is ambiguous. Suppose we were to make advances in neuroscience, and our best theory told us that (4a) is not ambiguous after all. This would be no kind of contradiction, but a puzzling empirical matter: how come speaker-hearers are so systematically misled? Yet there are clear and familiar examples where we do go wrong about syntax, such as with centre-embedded relative clauses (e.g., *Sheep dogs farmers train chase run*).

What I hope to have done so far, then, is provide a reason for why syntax as a combinatorial principle is not externally constituted, which doubles as a *prima facie* reason why words are not so constituted either, insofar as such principles essentially apply to words (whatever they are). I shall come to objections to this inference in the following section.

Much the same reasoning as applies to our combinatorial condition applies to the selectivity condition. The problem for externalism here is that the relevant semantic properties are only manifested when words are in combination, but understood as external entities, words need not be in any combination at all; that is, the relevant properties of words are structural ones that relate them to one another, and which determine certain interpretive roles the words possess within the structure. Take a simple example:

(11) Bill_{agent} kicked Mary_{patient}

In the combination, the words acquire their respective roles, either, depending on the theory, from the verb *kick* or from the structural configuration. There is, however, nothing essentially agentive about *Bill* or patient-like about *Mary*. One might think that there is something essential about *kick* as the assigner of the roles, but whatever is essential about it is dependent on the structure in which the verb can occur. We should not, therefore, semantically speaking, individuate the words in isolation, but only as items in possible structures. Again, these structures are not externally realised.

It might be thought that the relevant properties are merely relational ones, which can be externally accommodated; being an externalist about words doesn't entail that all properties must be intrinsic. Consider (natural) satellites. Being a moon of a planet, say, is a relational property, but moons as kosher external objects are not thereby impugned. The analogy, however, is not a good one. The Earth's moon as a hunk of mass need not be in any orbit, and it is a mere contingency that it is; it would not be if, say, the Earth were destroyed, or if galaxies collided. Words, as we saw, *must* be combinable, hunks of mass need not be satellites, and only count as such given certain complex facts. Words are not in all the denumerable sentences in which they can occur; nor are they abstractions from the sentences.

A better analogy, originating with Frege, is to think of words as atom-like with phrases akin to molecules. Atoms don't need to be in molecules, and are not defined by them, but molecules only emerge given the properties of the constituent atoms (number of electrons, etc.), and only certain molecules are possible. So, it is perfectly ok for words as external entities to have relational properties, but selectivity is not among them, for it depends upon syntactic organisation that itself is not external. Words must be such as to enter into syntactic combination by which they acquire semantic roles given the properties of the constituent words.

The next section will support the inference that these two conditions of combinability and selectivity ought indeed to lead us to reject words as externalia.

6 Supporting the anti-externalism inference

There are numerous ways to resist the inference from the conditions on words not being externally constituted to words themselves not being externalia. I shall here try to rebut four objections.

6.1 Projection

One might suspect that the force of the argument offered trades on an overly wooden conception of what counts as external. To be sure, the thought goes, the mind imposes structural conditions on word individuation, but words remain external. On analogy, we might think of lots of social kinds or artefacts as partly individuated by how we conceive of them or use them, but computers, musical scores, forks, etc. remain external. As flagged in Sect. 1, I want to park the general question of the mind-inflected ontology that populates our social world just to focus on the linguistic case. The point I insist upon is that when we go to explain word phenomena naturally specified in externalist terms, the phenomena will fractionate into a linguistic internal component and a non-linguistic external component without any recourse to words as some amalgam of both factors. From the optic of explanation, therefore, all the linguistic action is internal.

No-one should be interested in policing common speech and thought, denying, say, the existence of words on a page, the Manchester accent, the offensiveness of slurs, etc. The crucial point is that sanctioning our ordinary talk of words doesn't demand that we treat externalia as linguistic once we go to explain the phenomena. The general reason for this is that the linguistic component remains an internal invariance, which can have wildly different externalia associated with it, or none at all.

6.2 A difference in kind

Another natural move is to consider the structural conditions essentially associated with words as somehow different in kind from the words themselves such that the latter may remain external, whereas the former are conceded to be internal. There are different ways of pursuing this line.

First, one might adopt a broadly anti-realist position about structure, while remaining robustly realist about what incarnates the structure. Consider Quine (1972, p. 451):

I find the phrase 'logical analysis' misleading, in its suggestion that we are exposing a logical structure that lay hidden in the sentence all along. This conception I find both obscure and idle. When we move from verbal sentences to logical formulas we are merely retreating to a notation that has certain technical

advantages, algorithmic and conceptual... [D]eep structure loses its objectivity... [T]he grammarian's deep structure is similar in a way to logical structure. Both are paraphrases of sentences of ordinary language.

Quine's central point here is that the relevant structure is just a form of paraphrase, rather than a genuinely explanatory posit. If so, we could think of words as externalia, whereas standard generative syntax would just be one way among others of depicting strings of words for some clarificatory purpose. In particular, it would be illegitimate to impose a syntactic condition on words as if there were any 'objectivity' to it.

There are two things wrong with this view. Combinability and selectivity do not amount to mere conditions on paraphrase. Though they do allow for a paraphrase in many cases, in other cases they don't. For example, treating tense as the head of a phrase doesn't license any particular paraphrase.⁹ More significantly, the structural conditions *are* objective at least in that they explain the phenomena rather than offer paraphrase. As we saw above, it is a perfectly good question to ask why a given string is precisely two-ways ambiguous, as opposed to three ways, or not ambiguous at all. This kind of phenomenon is not explained by citing paraphrases, for it is precisely why a string admits and precludes certain paraphrases that calls for explanation.

Of course, to say that syntax is objective is not to say that we know the answers to our questions; it is only to say that the questions concern matters of fact, rather than a choice of notation or convenience, and positing one syntax as opposed to another can provide coherent answers to such factual questions.

Another move in the same direction, but which does not impugn the objectivity of syntax is suggested by Kaplan (2011, p. 511):

My creationism about words does not extend to sentences. The world in which sentences and other compounds live is brimming with untokened types. Put roughly, the basic elements of the language [words] are earthly creations, but the compounds generated by syntactical rules (the rules also being earthly creations and thus subject to change) are structures—types—which may or may not have tokens.

The claim here is that the sentences that include words can be considered as types living in one 'world', however characterised, whereas the words themselves can be individuated concretely or 'earthly'. Warranted objections can be made to this partition concerning the presumption that words are somehow simples with no intrinsic structure (cf., Hawthorne & Lepore, 2011; Nefdt, 2019). Indeed, my selectivity condition does not exactly entail that structural conditions apply to word individuation, but it is most often read as doing so, if lexical items are not treated as bare roots, which I presume Kaplan does not endorse.

A more straightforward objection to Kaplan's proposal, however, is that if words and sentences exist in different 'worlds', then it becomes opaque how the two can properly interact, why, in particular, sentential syntax imposes a condition on words, given that syntax is not externally realised.

⁹ Certain general conditions preclude ready paraphrases. It is broadly assumed, for example, that movement is successively cyclic, which means that any displaced item in a structure, such as *wh*-item in a question, will be copied at various intermediate locations (the SPEC positions of subordinate phrases).

Kaplan appears to evade this problem by seeing both words and rules that generate syntactic structures as being both finitely ‘earthly’, while the unbounded class of sentences is not. Yet the problem that the combinatorial condition poses does not concern a difference in cardinality and its supposed basis (a difference between the earthly and the abstract). The problem, rather, is that syntax, qua non-externally realised, applies to words, qua putatively externally realised. It is erroneous to deflate this problem as if it merely amounts to the thought that sentences, unlike the rules for their generation, are too numerous to be tokened (‘earthly’). To see this, suppose, say, that the number of possible sentences were finite and that all were tokened (repeatedly, if you want). The problem would remain that the syntax would not be tokened in any external sense, for it is not constituted by the properties of any media or any ordering of any media items (syntactically speaking, a sentence is not an n -tuple of phonemes). So, let the syntax be ‘earthly’ in the bare sense of not being Platonic. It remains opaque how putative externalia might be structured in ways not so much as registered by externally realised properties.

Nefdt (2019) may be read as offering an answer to this quandary, which constitutes my third candidate way of resisting the inference from anti-externalism about syntax to an anti-externalism about words.

6.3 Reconciling the concrete and the abstract

Commendably, Nefdt (2019) takes seriously the demands of syntax, considering it to be a desideratum of an adequate account of words that it accommodate their occurrence in syntactic structures. The essence of his view is expressed as follows:

Tokens have a non-natural ability to represent types, which is why tokens of varying phonographic profiles can represent a single type and tokens of a similar or the same profile can represent distinct types. There is a certain arbitrariness to this relation. Tokens of all sorts can perform (and fail to perform) in their representational tasks but they are generally defined by that capacity nonetheless. Following Szabó’s [(1999)] picture, word tokens and sentence tokens etc. represent types and indirectly referents. Departing from this picture (but hopefully still within its spirit), word types do not refer to or represent objects in the world directly (or abstract objects) but represent nodes in larger linguistic structures in either the places-as-objects or the places-as-offices perspective. So types are abstract places in structures and tokens are the individual sounds, inscriptions, signs we use to fill some of those places, i.e. terminals in trees (Nefdt, 2019, p. 904)

The idea is that words are ‘quasi-concrete’ in that they are abstracta that require a concrete representation for their individuation. Hence, token words count as particular words because they represent types, and it is from such types that the tokens inherit their linguistic properties. Crucially, the types are partly individuated by syntactic structure. In effect, therefore, words qua types are positions in a structure, and words qua tokens are externalia that represent such positions. This view is designed to satisfy

the desideratum that an account of words should have fidelity to linguistics, not only by factoring in syntax, but also because.

[I]inguists do not stop at pure structures in their work, they are particularly interested in those structures that are represented in real world languages. This follows from the fact that linguistics is an empirical science. Thus, the structuralist picture needs to be amended to include what Parsons calls quasi-concrete objects. These objects or positions-in-structures have a mixed ontology (Nefdt, 2019, p. 903)

Although the position is ingenious and commendable in its attempt to meet conflicting desiderata, it remains doubtful that linguists have any need for ‘quasi-concrete’ objects and our initial problem remains unresolved, for it is unclear whether the putative objects actually incarnate syntax.

On the first count, linguistics counts as an empirical science because it seeks to explain empirical phenomena, and is sensitive to all phenomena as potentially germane to the corroboration of its hypotheses (cf., Antony, 2003). None of this entails or presupposes that linguistics, qua empirical, is in any sense about linguistic externalia, which the ‘quasi-concrete’ might realise. Post-Chomsky, generative linguistics, at any rate, is not in the business of corpus analysis, and while such analysis might offer some insight into some questions, it does so by offering evidence for hypotheses that are not about corpora as such. Of course, such a position might be rejected as ultramontane, and Nefdt (2019, p. 903, n.23) follows Stainton (2011; cf., Tasker, 2022) in endorsing a ‘mixed ontology’, meaning that we can be pluralist about what we take linguistic objects to be relative to particular explanatory projects. At least for some such projects, then, Nefdt might reckon that ‘quasi-concrete’ objects must be posited. It really needs to be established, however, that *any* such pluralism is actually entailed or presupposed by the relevant explanations of the branches of linguistics.

The point here is that *all* theorists blithely speak of sentence and word tokens, but it is another matter whether the apparent ontology of such talk has any role to play in the explanations offered (cf., Collins, 2020; Rey, 2020). Historical linguistics, for instance, appears to be interested in things such as Old English, Middle English, and Modern English, and how one led to another, but nothing appears to be lost if we eschew such apparent ontology for the cognitive states of speaker-hearers of various populations. We might look to *Beowulf*, Bede, Chaucer, and Edmund Spenser for evidence, say, but the historical hypotheses are not about such texts, which came about due to independent factors (historical linguistics is not a branch of literary criticism, still less the history of publishing). This matter is somewhat by the by, anyway, for the syntactic hypotheses to which Nefdt wants his account to cleave concern structures which need not be tokened, and might not even be possible to token. Whatever is tokened, qua linguistic, is a consequence of an underlying system, and so constitutes evidence for the nature of the system, but it remains an arbitrary snapshot of no essential significance in itself.¹⁰

For some explanatory purposes, one might precisely be interested in the properties of specific speech events or what linguistic properties are tokened in a given locale,

¹⁰ For example, crucial data to decide between competing hypotheses might simply not be available by looking at corpora. This is certainly the case for acquisition, where children only have the data to which they are exposed to go on. Thus, experimentation, as in any other science, requires the manufacture of data. See Crain and Thornton (1998).

and so a corpus might be of immediate interest. None of this need involve the positing of linguistic externalia, however. Insofar as the concern is for the linguistic properties, then the interest is in the internal states of particular populations or specific persons relative to what is produced in time and space. Again, as per normal, the phenomena are fractioned into internal and external components, the former realising what is linguistic to the phenomena.

It might be thought, ‘Well, ok, strictly speaking, that is right, but people do utter sentences and so syntax must apply to words in such cases, which may thus be viewed as ‘quasi-concrete’ incarnations of the structure insofar as the structure is necessary for their individuation’. This takes us to the second and more serious complaint flagged above: it is unclear whether syntax ever does apply to concreta of any kind.

Suppose, following Nefdt, we say that words (/lexical items) are terminals on trees (trees are just a vivid representational format for depicting constituency and dependence relations). It doesn’t follow that the terminals are concreta in any sense whatsoever that ‘fill positions’; minimally, they are simply atoms relative to phrases or higher nodes on the tree, i.e., they have no constituents. What would establish externalia as the terminals would be the application of syntax to the actual concretea of the text or speech vocalisation. We have seen, already, however, that there is no such application. Syntax just doesn’t obtain externally. We speak as if it does, in treating media of various kinds as if they had linguistic properties, but the externalia themselves are not syntactically organised, but may be viewed as being so organised if we treat the externalia as the output of a certain kind of system, which thus evidences the nature of the system.

Take a toy example. Consider:

(12) aabb

As a concatenation, it has no structure beyond linearity. We can think of it as being generated in different ways, however:

(13) a [a [ab] b]
 b [[aa][bb]]
 c [a [a [b [b]]]]

The string itself remains invariant; what changes is the hypothesis about the nature of the system that generates it. The terminals, qua *positions* in a structure, therefore, are not anything that might occur externally, whether token or type, but mere elements of a system that we may variously hypothesise to be the generative system underlying the production of the string. Of course, in this toy example, the ‘a’s and ‘b’s remain invariant, which is unlike the linguistic case, where the terminals need not even be externally tokened, and any external token has a host of extra-linguistic properties. It bears emphasis, though, that even in the toy example, if we take the ‘a’s and ‘b’s to be partly individuated by their being terminals, then the property of being this or that terminal is not externally realised.

This response in turn might be challenged in a way that leads us to our third consideration of how the inference from non-external syntax to non-external words might be resisted.

6.4 Representationalism

Suppose all of the above is granted. A response remains: true, words and syntax are not externally realised, but syntax and words are represented; all you have shown is that syntax is mental content, and if words are syntactically individuated, then they must be partly mentally individuated, i.e., with specific reference to the cognitive capacity of speaker-hearers. In terms of Nefdt's position, it might be thought that the terminals or linguistic atoms are not externalia, but representations of externalia. This inverts his intent, somewhat, as on his account word tokens are supposed to represent types defined structurally. Still, it is an available position and would, if viable, preserve words as externalia while accommodating the non-external nature of syntax. The position faces a decisive objection, however.

As Rey (2003, 2020) argues, if syntax and words (what Rey calls 'standard linguistic entities') are represented by the mind of speaker-hearers, it just doesn't follow that there is anything represented out there in any way at all. To be sure, we think and talk as if what is represented is extant, yet what intuitive pull there is to this feeling might be a complex form of illusion (cf., Bromberger, 2011; Azzouni, 2010, 2013). The force of this rejoinder is due to the explanatory work we ask of representations.

First, linguistic properties are not rendered in external terms, as we have shown at length, and appear to be no worse off for that. Secondly, *prima facie*, no explanations require there to be such entities; apparently, it suffices that we think there are or perceive them to be heard or seen. This can be appreciated by considering the difference between imagining speech (rehearsing a lecture in your head, say) and hearing speech. As far as syntax and semantics is concerned, there just is no difference. One can, as it were, token syntax and words with nothing out there.¹¹ This is not so with imagining a table of a certain weight, colour, and dimensions. In so imagining, one hasn't created a photon-absorbing mass in some space-time region. The externalist about language needs to show that *at least one* explanation in linguistics entails or presupposes external linguistic properties being tokened in any way whatsoever.

Nefdt (2019, p. 903, n. 23) does briefly consider Rey's position, and says the following:

[Rey] argues against the existence of standard linguistic entities. Despite this stance, I think his idea of linguistic objects as intentional nonexistents is compatible with an eliminativist structuralist theory such as Hellman (1989) since Rey's view would still have to presuppose that these entities are designated by the rules and structures of linguistic theory.

This is odd, since Rey (2020, p. 291/332) thinks there are no relevant entities at all to be designated. To be intentionally non-existent is not to *be*, albeit in some different manner from tables and chairs, it is just *not to be at all in any way at all*, at least as far as Rey is concerned. What is true is that the vehicle for the contentful intentional

¹¹ This is not to suggest that inner speech as a quotidian feature of our mental lives typically involves a full syntactic structure; it is mostly fragmentary. Still, we *can* rehearse full syntactic structure 'in our heads', and the fragmentary inner speech amounts to fragments of a particular system, not just any old jumble. See Langland-Hassan and Vicente (2018) for controversies on the nature of inner speech.

inexistent must exist, but it will not be in a syntactic position, for syntax itself is still another species of intentional inexistence. The vehicle will be a physical state.

Much could be said here, back and forth, but space demands the point be left as a challenge for the externalist (see Collins, 2020; Rey, 2020).

So far, then, our argument is supported. Words, by everyone's lights, require syntactic and selectivity properties, but such properties are not externally realised, and so words aren't either. Various ways of resisting the inference were rejected. In the next section I shall offer a fully internalist approach to the explanation of word-phenomena without appeal to external linguistic items. The position will depart from an intentional inexistent view insofar as it doesn't take the representation of (intentionally inexistent) linguistic entities to be required for any relevant explanations. As mentioned above, my problem here is not with the inexistent bit, but the intentional bit. Further, I shall argue that the view on offer is perfectly consistent with the truth of much common sense about words.

7 The explanatory role of internal states

If, therefore, we grant the combinability and selectivity conditions on words are not externally realised, but are nevertheless essential to words on any worthwhile conception of them, we should look to internally realised capacities for the explanation of word-phenomena. It would be confusing to speak of words as internal states, but for all explanatory purposes certain cognitive invariances that serve as the atoms for syntactic computation will be essential to the capture of word phenomena, with no external properties essentially required as invariances. In this section, I shall set out the view, show that it is consistent with the truth of much what we colloquially want to say about words, and defend it against two likely objections.

7.1 Internalism *sans* intentional inexistence

I mentioned just above that we should identify certain cognitive invariances as essential to the *explanans* for word phenomena. This is demanded due to the vast range of extra-linguistic factors that can influence cognitive performance in both linguistic consumption and production. What fits the bill are entries in long-term memory that link morphophonemic instructions (how to pronounce/recognise) with semantic information and syntactic categorisation of a somewhat flexible kind, perhaps just a root label.¹² The formation of such entries is called lexicalisation, i.e., a packaging of otherwise available information (encyclopaedic) and some bespoke features. The entries can be activated and so tokened as an occasion demands for production or consumption, where they are subject to syntactic conditions in a workspace where, as we would theorise it, the entries are combined into a structure that becomes available for broader cognitive activities, such as production and parsing. The rest is a matter

¹² Stojnić (2021) proposes, similarly, to use lexical selection as a condition on *u* and *u'* being utterances of the same word along with a causal/historical chain leading back to a neologising event. While I am very sympathetic to the mentalist aspect, I ultimately see a thorough-going internalism as giving up on word individuation; a fortiori, no causal/historical chain is required (see below).

for on-going empirical inquiry. As mentioned in Sect. 4, a current dimension of much dispute is how rich and structured these memory entries are. For our purposes, the bare idea will suffice.

As mentioned in Sect. 5, there is a thin reading of ‘intentional inexistence’, under which it simply picks out word phenomena as essentially involving the cognition of speaker-hearers without any ontological commitment on our part. I assume Rey to have a more robust notion in mind where the putative representational content is explanatory. At any rate, this is the claim I am interested to deny.

On my view, it is unnecessary for any explanatory ambition to think of words as intentional inexistents or objects of pretence or the intentional objects of perceptual illusion, even if we don’t actually see words out there, any more than we actually see triangles (qua mathematical objects). This is because the perception of external media *as* linguistic is a downstream phenomenon of the memory entries that explain such features of perception as well as production. In other words, you just don’t get to perceive anything as linguistic unless the right cognitive states are activated, and it is such states that constitute the linguistic realm for theoretical purposes. That is, rather than thinking of words as entities we represent, the real invariances are not in perception but the upstream stable states that explain and constrain downstream representational states.

So, it is true that we mistake as a word an external entity what we can reidentify as the same *thing*, but the invariance here is the activation of the same long-term memory entry in a speaker being subject to the same syntactic conditions rather than any vehicular features experientially identifiable, which need not be invariant at all, such as between reading and hearing, say, the same *word*.

I take this consideration to be the beginning of an answer to what Rey (2020) and Collins and Rey (2021) call the ‘common coin’ problem, and which Rey takes to be a *prima facie* reason for intentional inexistence to be indispensable for explanation of linguistic phenomena. The problem is that since the perception/consumption of language involves the interaction of different systems, there must be interface stabilities so that, for example, syntactic information is paired with semantic information and phonological information, all of which is associated with the identification of sounds (/gestures/inscriptions) as linguistic. It cannot be adventitious that the right information is recruited on the appropriate occasion for the perception of a word. Equally, it appears that the interfaces cannot merely be causal or transductive, for linguistic items do not have causally active properties. A verb token, say, qua instance of a grammatical category, cannot cause anything; sounds and ink marks do have causal powers, but grammatical categories cannot be identified in such terms. Yet in some clear sense, we recognise externalia *as* linguistic, as having syntactic and semantic properties. How is this achievement possible unless the properties are represented?

In addressing the common coin problem, it bears emphasis that the mere notion of interfacing systems that result in linguistic identification of externalia does not entail a representation common between systems, or one system representing the content or processes of another. Chomsky (1995, p. 221) does specify the interfaces in terms of ‘legibility conditions’ (cf., Rey, 2020, pp. 284–285). Yet, minimally, what is required is just stable relations between different systems, mediated by a procedure or algorithm to translate between them. If we assume that syntax fixes a set of hierarchical

structures, legibility amounts to how different systems can use the hierarchical relations as instructions to form different kinds of structures. For example, phonology can't trade in hierarchical relations, but must, instead, linearise any hierarchical relation between x and y into one of precedence. Thus, it might be that asymmetrical c -command fixes precedence, even though the syntax doesn't represent precedence (if x c -commands y , then x precedes y) (Kayne, 1994; Nunes, 2004). Here, a linearisation algorithm will apply to a syntactic structure to determine the order between the constituent items, with such a linearised structure only then being phonologically interpreted as a form underlying a speech event. Similarly, a semantic system must identify units of structures (e.g., a relative clause inside a noun phrase), what item is the head of such a unit (e.g., *cow* in *brown cow*, making the phrase about cows, not brown), what item is an argument of what (e.g., in *wiped the table clean*, the table was wiped producing a clean table rather than a clean table being wiped), what item can bind what other item, and so on. *Ex hypothesi*, the syntax represents none of the relevant concepts necessary for semantics, but does form structural relations that can feed computational processes that do determine the relevant semantic relations (Cecchetto & Donati, 2015; Reinhart, 2006).

In short, the 'common coin' is an *effect* of the distinct systems being related by algorithmic translation devices, rather than an *explanation* of how the interfaces are arranged. For example, there is no [dog] common coin, which is invariant over syntax, semantics, and phonology; rather, [dog], as a syntactic-semantic-phonological complex is the result of a particular interface arrangement.

So, externalia can be recognised as linguistic only given a peculiar set of interfaces between systems and sensitivities to external properties, such as prosody. It is the establishment of such an arrangement that gives rise to the sense of a common coin that generalises over distinct systems, but that is just a way of specifying an arrangement of interfaces that is sufficiently stable to constitute the capacity for discriminating externalia in terms of this or that set of linguistic properties, which are constituted simply by such systems operating in concert in relation to certain stimuli. That, at least, is the shape of an answer to the problem.¹³

A further question arises as to why we should think of internal states as representational at all. From where does the putative content arise? Rey (2020, p. 375–377) has elaborated a sophisticated position on this question of representation. The gist of it is that the linguistic system as a whole can inherit an intentional character from some of its representations being stably covariant with external features, as must be the case if language is to be externalised via some medium (sound, sign, etc.). So, what I have been calling memory entries get to be *about* words (in the relevant intentional sense), because they play a role in a system that is linked up to stimuli, even though these stimuli themselves do not count as words. I cannot do justice to this position now; a worry will have to suffice.

The model would appear to be correct if all of the relevant linguistic properties were recruited in parsing, and so were required to be stably linked to stimuli that we

¹³ See Adger (2022) and Rey (2022) for discussion of whether to take linguistic items to be brain states or the non-existent objects of representations. Here I have adopted a third position in the sense that I agree with Adger that the items are not representata, but agree with Rey that they are not brain states either. The items are posits that classify unknown states for certain explanatory ends.

could recognise as linguistic. There is little reason, however, to think that *all* syntactic and semantic properties are so recruited. We are not even perfect at identifying phrase boundaries (consider so-called exceptional case marking—*Bill believed him to be honest*/**Bill believed he to be honest*) and parses familiarly break down and are ‘garden pathed’. One option, which Rey (2020, p. 343) suggests, is to ‘Ramsify’ the relevant theory so that all relevant concepts are defined structurally and acquire a content thanks to only some of them being involved in parsing, i.e., stably associated with stimuli.

It is unclear to me how the proposed inheritance of content is to work on this proposal. Consider a covert item like PRO or the various covert copies of items proposed in generative theories. They are not associated with any stimuli, and nor is it clear how to link them to items that are associated with stimuli. Abstracting over a good deal of complexity, PRO comes in three basic flavours depending on the host structure: arbitrary (PRO *running is fun*), controlled (*Jane wants PRO to leave*), and bound variable (*Everyone wants PRO to leave*). The generalisation here is that PRO is a thematic subject of non-finite verbs. So, here one does not want to Ramsify at all, for PRO (as far as the basics go), has a clear definition without appeal to a total theory, but if we don’t Ramsify, then the putative content is not a property inheritable from stimuli, for PRO can be defined without reference to any ‘observation term’. There might be some answer to this kind of problem, but I leave the issue as a challenge.

7.2 Saving the common lore

It might seem that the position on offer is an out-and-out eliminativism, for all the things we commonsensically should want to say of words externally construed appears to be unavailable or, at any rate, not true of long-term memory structures. That is correct, in one sense, for there are no words as external entities, but we can still understand much of our common linguistic lore as true, for the truth of much of our common sense in these matters just doesn’t require there to be external words, or so I shall argue. Let’s first look at the three conditions from Miller (2020a, 2020b) flagged in Sect. 3.

First, words must be expressible. The truism that we utter words need not be denied, but it doesn’t follow that this truism is *made true* by words as externalia, as a blacksmith producing horseshoes entails the independent existence of horseshoes. Instead, we may say that words are expressible insofar as long-term memory entries are activated whose instructions lead to articulation. The articulation has no *linguistic* stability, however, but we may be said to utter the same word from one occasion to another, because the activation of the memory entry counts as the same for the purposes and interests relevant (more on this anon). Given the presence of varied factors involved in articulation, great variation obtains in the sounds or signs actually produced, but this motley product is not linguistic in-itself, but is, rather, explained by the invariant involvement of certain dedicated capacities whose peculiar properties marks it as linguistic. Thus, the wind crying ‘Mary’ or a raven quothing ‘Nevermore’ does not count as linguistic, even if the product is indistinguishable from you or I making the relevant utterances.

Even granting that the same sound or gesture can occur linguistically and non-linguistically (via the wind or a bird), a more serious concern arises with two or more

people using the ‘same word’, which appears essential to communication (cf. Kaplan, 1990; Feinsinger, 2021), which I shall address shortly.

Similarly, words are creatable, with new ones coming along virtually each day. *Chumocracy* is a relatively new coinage that describes a certain form of social elite, not established by lineage or achievement, but lose social connections of friendship, same alma mater, etc., all of which predicts various advantages, such as promotion, being awarded contracts, etc. To say *chumocracy* is created, in our terms, is just to say that lexicalisation occurs in one person and then produces similar lexicalisation throughout a population. What this process involves is presently as unclear as is acquisition of a lexicon in infancy. Still, it might be thought that, however it occurs, it establishes a ‘same word’ relation that holds of a population, or across a chain of individuals, and this same-word phenomenon can’t be a matter of individual psychology. This conclusion is at best moot, though, for it remains as yet unclear why there must be something that is the same beyond sufficiently similar internal states of speaker-hearers, i.e., similar enough to support communication and identification of the ‘same word’ for the purposes at hand, notwithstanding potential variation along all linguistic dimensions. A question arises of the dimensionality of the sameness of internal states, which I shall address below.

Similar remarks apply to the intuitive evolvability condition on words. Kaplan (1990, p. 101) suggests that rather than thinking of words as replacing one another, as in a relay race, we might think of them as changing their properties, as regards spelling and pronunciation, say. From the current perspective, the truth of this thought, as far as it goes, is based on variation within the parameters of the lexical memory entries realised in the minds of the relevant speakers. There is, however, no independent truth of the matter when such evolvability is a change of properties or a change of word (a relay race), for there is no empirical difference between thinking of a word with variable properties or thinking of a word just as a collection of variable properties. The only difference is that we intuitively rank some properties as more significant (syntactic and semantic) than others (pronunciation). This is what is predicted on the present model. Speakers’ lexical memory entries vary over both time and individuals. There is no whistle that blows when the shifting pattern constitutes a change of word or not for a population.

The above responses raise two questions: what is the right response to the individuation problem? What counts as the same lexicalisation? The two questions are related.

Each individual has their own peculiar lexicon with many features shared between speaker-hearers who acquired their language together. Each item in the lexicon is a long-term memory entry encoding specific roles and instructions in individuals, which explain how they produce and consume language. Thus, there will be variation among a relevant population. In this light, all questions of the sameness of a word token relative to a type are answered in a more-or-less way depending on the interests behind the question. Words are sharable or repeatable or public to the extent that the relevant population have the same lexicon, being individual cognitive states of speaker-hearers, but the states will share many properties given broad similarities amongst members of a population that acquired their language and so lexicalised together. In some cases, we might want to say that a person has a different word, if they employ it in

different ways, but we are mostly relaxed about individual variation when it comes to encyclopaedic information. Some words will vary greatly in semantic/encyclopaedic features (common abstract nouns, say), while other will remain highly stable, such as determiners, whose semantic significance is wholly structural.

Feinsinger (2021, p. 327) thinks that we should be decidedly less relaxed: ‘without a reliable mechanism for matching content, these [generative] models leave communication unexplained, offering no intersubjective same-word relation’. It is true that no ‘intersubjective same-word relation’ is sanctioned on the kind of model proposed, but it doesn’t follow that any phenomena are thereby unexplained. The question here is whether communication actually requires any same-word relation to explain its success (or failure). In broad terms, as intimated, sameness will be a more-or-less matter conditioned by the context in which the question of the sameness of a word arises. Sometimes precision is called for, other times not. Indeed, on relevance models of communication, there is no presumption of a same-word relation; communication operates at the level of the quality of information relative to cost an interlocuter expects from an utterance (Sperber & Wilson, 1986; Carston, 2002). Relevance, in this sense, can be calculated, without a shared code.

A simple Platonic (*Parmenides*) point also bears emphasis. Suppose two speaker-hearers share the same word in the sense that the two individual memory entries share all relevant information. It doesn’t follow that there is the one word save as a way of abstracting from the states of the individuals, and it certainly doesn’t follow that there is a third external thing to which the speakers relate. So, one may employ a same-word relation in this etiolated sense, but nothing of explanatory interest turns on it being realised, and our communicative exchanges proceed smoothly with partial mappings. In short, sameness of word is not an explanatory notion, but just a way of generalising over what lexical features are relevant in some particular context.

Along the same lines, we may deflate the question of when two lexical entries count as the same. The lexical entries that enter into the explanation of the relevant phenomena are theoretical posits, so are individuated relative to whatever information the best theory specifies. There is disagreement about this (Sect. 3). Assume, however, the traditional model under which an entry contains syntactic information, category of the item and argument structure, if any; theta information, related to argument structure; relations to world/encyclopaedic information, if relevant (for nouns, not for determiners or prepositions); and instructions for articulation (phonetic or gestural). As characterised, it might well be that no two speakers share lexical entries, at least for open-class items, such as verbs and nouns, for there will almost certainly be variation in related world/encyclopaedic information and a good deal of toleration over the other factors. What the entries have in common (i.e., what makes them lexical) is the packaging of the same *kind* of information: syntactic, morphological, semantic, etc. And crucially, the entries feed into syntactic structures that underlie our competence with sentences. In short, human minds share the kind of properties and processes that enter into lexical and so linguistic competence, but there is no necessity that minds will organise themselves to have exactly the same lexical entries.

My proposal is a radical departure from common sense, but this is no mark against the proposal, of course. Still, it might appear to be burdened with the onerous task

of offering an internalist paraphrase of our externalist talk lest we adopt a thoroughgoing and wholly implausible error theory for all word talk (Tasker, 2022). Eschewing, however, an externalist model of ontology for a wholly internalist model of explanation concerning word-phenomena is *not* inconsistent with the truth of much of what we say about words even without exploring the periphrastic option.

One approach here is to distinguish between truth conditions and what makes a claim true (Azzouni, 2010; Collins, 2021). So, we may reckon the world to make true what we say, but how the world is under such circumstances is not recorded in what we have literally said. For example, *Holmes is smart* is made true by stories, but no such stories are part of what we say or refer to in saying Holmes is smart, and so not part of the truth conditions that constitute the meaning of the sentence. Thus, we speak as if words were public entities out there. If we take that at face value, then we renounce any paraphrase or analysis of our word-talk that shifts the truth conditions to long-term memory entries. That is to say, the truth conditions of our common lore about words mostly do *not* advert to internal states, and certainly not the kind of states hypothesised in generative theories, just as our talk of fiction does not advert to stories, most of which we might not even know. Still, we can see how such common-sense claims might be true by explaining the phenomena of which the lore speaks. The truth conditions remain the same, but we do not take them to be a route to what there is. We especially don't when all ontological and explanatory interests are satisfied by not positing an ontology that precisely answers to the truth conditions. In a similar way, we might explain the truth of the Sun rises, even though it doesn't, or the sky being blue, even though just what counts as the sky is unclear. Sometimes a paraphrase might be available, as with the sunrise, which we might take to be the appearance of the solar disc on the horizon, treating the Earth as our inertial frame of reference, but it is wholly unclear how to paraphrase away fiction or, indeed, word talk in internalist terms. But that only indicates the great variability of conditions that may make such talk true. A common factor, but not a sufficient one, in the case of words, will be the internal lexical entries in long-term memory of speaker-hearers, and it will be this factor that carries the explanatory load.

All that said, some of the lore must be reckoned false. For example, the marriage of a word with the properties of a particular medium are very strong. We naturally speak of words having a specific spelling or a pronunciation, which are externally constituted. On reflection, however, we may see that this bit of common sense is confused, or at least in tension with other bits of lore. Illiterate people are rightly said to share words with the literate; but the former can't read or write, which most humans who have ever lived have not been able to do. There are conventions for spelling, which are not explicable in terms of internal cognitive states, but such conventions pertain to a medium, not to language as such, and are highly contingent upon social variable factors. A similar conclusion holds for conventions for pronunciation. After all, even if there were no conventions for such matters, language would remain.

Another phenomenon is common misunderstandings, say construing being livid as being red, or being disinterested as being uninterested. Suppose a person uses *disinterested* to mean uninterested. We would naturally say that they are misusing the word, not that they simply have *their* meaning, we have *ours*. That is the right intuition, but the phenomenon here is not linguistic, but more a case of social alignment.

The person who ‘misuses’ the word does not have a semantic defect, but is merely misaligned with her peers. Apart from potential miscommunication, nothing is amiss with the person. Indeed, with time, the pattern of alignment might change, as has happened with *livid* (its ‘angry red’ meaning now being acceptable) and *refute*, which now means *strongly disagree* in many mouths. The question of whether different words are being used or the one word is being misused has no linguistic answer, but can only be answered in terms of how strongly one feels a norm of alignment should be in play.

7.3 Three objections

To finish, I shall consider three objections.

The first objection is that I am up to my ears in a use/mention confusion. It would be such a confusion to claim that words *just are* long-term memory structures that represent words. It is (*inter alia*) to avoid this kind of confusion that Rey (2020) adopts his intentional nonexistent model, under which words and other linguistic entities are represented, but no explanations turn on the entities being extant in any representation-independent way; theorists simply ‘pretend’ that there is the represented ontology for various explanatory purposes (it is less awkward to talk of ‘p’ as opposed to ‘S’s representation of p’). Although I am sympathetic to Rey’s position, there is no use/mention spectre lurking over my account. The hypothesised memory entries that constitute the atoms (or *ur*-elements) of syntactic combination are what enter into an explanation of word phenomena, but they should not be understood to be words themselves or the periphrastic base for word talk. The entries are defined in terms of their being open to syntactic processes and representing information concerning interpretation (selectivity conditions) and articulation, which downstream produces what we commonly think of as words in the consumption and production of various external media *as* linguistic. Hence, the truth of sundry claims about words as written on a page or heard at a lecture is explicable in terms of the relevant speaker-hearers recruiting certain kinds of cognitive capacities in their treatment of the relevant external media *as* linguistic.

Pressing the use/mention objection, however, one might suspect that, at some level, a use/mention error must be committed, for if I am not conflating words with representations of words, I am conflating some linguistic properties (being a verb, say) with their representation, for verbs just ain’t in the head (as it were), only states that represent verbs are. The answer to this worry is that, in some sense, the use/mention distinction is collapsed, but not guilelessly. We theorise the mind in abstract/computational terms, as trading in notions of verbs and nouns and syntactic structures, but here we intend to be specifying a system of possible internal states of the speaker that we have no way of identifying save for via our abstract categories and the theory of how they relate to one another in something like a deductive structure (Collins, 2004, 2014).

The case is analogous to mathematics. We have no access to numbers and their relations save via a notational system (or maybe diagrams), but we don’t take the properties of the system simply to be properties of the numbers. Likewise, the only current access we have to the internal states of the linguistic mind/brain is via our theories, but we don’t take properties of the theories to be the properties of the mind/brain itself beyond

certain information, as a long-term memory entry, being accessible to the mind of the speaker-hearer (Collins, 2004, 2014; Johnson, 2015). Insofar as such information is necessary to explain the relevant phenomena, we must credit it to the speaker-hearer, but such information does not characterise any system other than the linguistic one we are aiming to explain. The categories we hypothesise have no life other than to capture the cognitive phenomena the theories target (we could take the categories to describe a Platonic realm independent of the mind, but the posited abstracta would remain linguistically *sui generis*).

A posited memory entry involving a verb, as might be, is thus not true of anything beyond us, not even an abstract object, for a verb is *just* what our theory of such internal states says. The case *appears* to be different, for example, from the one where we credit people with representations of Euclidian solids or numbers. Such abstracta have a range of properties determined by mathematics, and so it makes sense to think there are two things here—triangles, say, *and* their representation—which ought not to be conflated. Yet many of the mathematical properties have no role to play in the explanations afforded by the hypothesis that the visual system deals with representations of Euclidean solids, say. Pythagoras' theorem, for example, holds for triangles, but is, as far as we know, irrelevant to visual perception. More generally, every line of a triangle contains nondenumerably many points, each being a limit of the preceding sequence of points on the line. Such a property is not an aspect of the visual triangle.¹⁴ It might be, of course, that the best theory of vision will turn out to employ the full battery of mathematical properties of what it 'represents' (the visual system *might* compute the square of the hypotenuse in relation to its other two sides for some reason or other). There is little reason, however, to think that that is true. Mathematically, shapes are defined by invariances of scale, rotation, and translation; effectively, a shape is a group. Vision, however, is perspectival: shapes as perceived are not group-theoretic invariances. The outstanding question is whether the invariances play an upstream role in the production of visual phenomenology. In our state of ignorance, the strongest claim we are entitled to is that triangle representations attributed to the visual system, as might be, are not true or false of the mathematical objects, but are simply theoretical specifications of states in terms borrowed from the properties of the mathematical objects. So, to think of an internal system as representing a triangle or a verb phrase is not to be committed to a representation and its represented, however conceived, but only to an internal system of states specified in the relevant abstract/mathematical terms, much as a physical dynamic stem requires calculus. What appears to be the represented contents of the speaker's internal states to us as theorists is just the effect of our theories ineliminably appealing to abstract categories or mathematical structures to account for the internal states.

The second objection concerns my attempt to preserve the truth of the common lore while eschewing the externalist ontology to which it is apparently committed. The problem is that the explanation of word-phenomena on offer might be seen simply

¹⁴ Marr (1982) proposed that the visual system effectively computes the Laplace-Gauss to identify a line from a distribution of points. Note, though, that here, as far as the system is concerned, a line becomes a property of a distribution, via the function, rather than the line being represented, as in real analysis, with non-denumerably points. In short, the psychological explanation doesn't involve the mathematical object, line.

to render words as illusory and so the common lore only apparently true. Consider a standard visual illusion, such as the Müller-Lyer illusion, whereby a presentation of lines of the same length appears as a presentation of lines of different length. Suppose it is explained how the illusion is brought about (perhaps due to indicators of depth, but the correct explanation is not currently important). It is not thereby shown that the illusory content is really true, that the lengths of the lines actually do differ. What is explained is only that typical perceivers are in error. Just so: if it is explained why we reckon there to be words beyond us, the truth of our common lore about words is not thereby preserved. To the contrary, just as in the illusion case, what is explained is why we think/see p , despite p being false.

The objection is well-taken, but can be accommodated. It is perfectly true that explaining a phenomenon does not always involve sanctioning the apparent ontology of the phenomenon, but I am not committed to the contrary. I have no interest in preserving the common lore as a species of unreflective metaphysics. If a philosopher says, ‘Words are socially individuated intersubjective communicative devices’, meaning to precisify and endorse the ways the folk normally speak, then she has said something false. All I am interested in preserving is some first-order talk about words absent any metaphysical generality as to why such talk is true, i.e., a general story involving word externalia. So, we can speak of words on a page, the words one hears, that word being the same as this one, etc. All claims of this kind can be shown to be true (or false), within the parameters discussed, by appeal to cognitive states; that is, it is the involvement of the same or different long-term memory entries in the identification of the external media as linguistic that makes the claims true or false. It is a further step to insist that what such colloquial talk *really* expresses is an ontological commitment to externalia. In fact, if queried on such matters, I would suggest that ordinary people would just be baffled. For example, normal people readily think of numerals and numbers to be the same (they unreflectively talk of numbers as being on a whiteboard, say), but no-one thinks that dividing a number involves erasing some part of a numeral. The normal person does not reason that numbers are therefore Platonic. They simply don’t know what to think about the metaphysical questions but their mathematical understanding doesn’t totter. The normal person’s otology, if you will, is very thin in this regard, and a proper semantics reflects that, not burdening linguistic competence with intuition about matters of fact, arcane metaphysics, or unreflective prejudice.

The same reasoning applies to rainbows. Some of what we want to say about rainbows will be true (*Indigo is in the rainbow, but brown isn’t*), whereas other things will not be (*Rainbows are semi-circular*). There need be no rainbow as such out there for a partition of truth and falsehood in our rainbow talk.

One concessive move here would be to endorse Miller’s (2019) bundle account, where a word is individuated in terms of a bundle of properties, both mind internal and mind external, which need not all be instantiated for a token of the type to be realised. In particular, Miller rightly points out that we surely intend to speak of the very same word as both mental and external (*The very words I was thinking were the ones you wrote*). I take this to be a bit of word lore that should be preserved. The general internalist reasoning so far rehearsed applies here. One doesn’t get to write a word down, say, unless it is a result of processes that involve the relevant activation of a long-term memory entry. What is thought and written down, therefore, might count

as the same for some particular purpose, not because, in metaphysical generality, a word is a bundle that has both internal and external properties, but because (more or less) the same mental states, for the purposes at hand, are involved in both cases, i.e., entertaining words cognitively and recognising words beyond one can call upon a lexical entry that counts as the same for the purpose at hand. After all, the governing intuition fades away in a case where someone says, ‘I had the very same word in mind as the wind just etched in the sand’. In such a case it is clear that a word wasn’t produced by the wind, but only something onto which we may project linguistic properties. What the wind lacks is the linguistic competence to produce a word. The determining factor is the presence or absence of internal states of the producer. The same moral applies in the communication case. Speaker A can be said to utter the same words (repeat) as speaker B, but since the sounds don’t suffice, which might issue from a parrot or the wind, we only take the words to be the same thanks to the activation of what we presume to be internal states sufficiently alike. Of course, given all that has been said, there is no context-invariant account of what counts as lexical entries to be ‘sufficiently alike’; for some purposes, subtle semantic differences matter, for others not; likewise for syntactic and phonological differences relative to different contexts.

Thirdly, one might think that while a wordly externalism might succumb to the arguments offered, the existence of other linguistic minds must be granted. That is, linguistic explanation does involve the minds of a population of speakers-hearers, and so doesn’t devolve onto individual mental states. The issues here are tricky, but, first off, note that the mere existence of a population of speakers is not germane. Humans share visual and digestive systems, but vision or digestion are not external social facts. Language might appear different insofar as features of it essentially relate to other minds as with discourse effects (focus, topic/comment), interrogatives, gender, etc. It remains unclear, however, whether such features are constitutive of language or simply effects that the language system allows for given the fact that language is used for communication. The internalist bet, as it were, is that the latter scenario is correct. There are some reasons for this view. First, gender, say, is variable in how it is represented in a language, highly marginal in English, but more extensive in morphologically richer languages. Secondly, where syntax is recruited to achieve discourse effects, as with interrogatives, the kind of syntax is otherwise available. For example, *wh*-movement occurs in interrogatives and relative clauses, but only the former plausibly presuppose other minds. Thirdly, focus clearly depends upon extralinguistic capacities (stress and accent), and tracks syntactic options. In sum, that there is a population of language users is a contingency as far as linguistic explanation goes.

8 Conclusion

There is a clear sense in which it is right to view linguistic entities as abstract, i.e., they are, as we would say, not in space–time and are causally inefficacious. The interesting question is whether we are obliged to indulge in the reification in the first place. It seems not, for all the explanatory ends linguistic entities serve are wholly satisfied by appeal to the character of the internal states of speaker-hearers. The temptation to think otherwise is subreptive in Kant’s sense: we mistake conditions on our cognition of a

‘thing’ for conditions the thing must meet anyway; because there *is* something outside of us, such as sounds and ink marks, to which we unreflectively impute properties that are not outside of us, we ineluctably feel that something beyond us must be accounted for. At the very least, I hope to have shown that we can overcome this turn of mind without loss to truisms about language and with gain to our theoretical understanding.¹⁵

Declarations

Conflict of interest There is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adger, D. (2022). What are linguistic representations? *Mind and Language*, 37, 248–260.
- Antony, L. (2003). Rabbit-pots and supernovas: on the relevance of psychological data to linguistic theory. In A. Barber (Ed.), *Epistemology of language* (pp. 47–68). Oxford University Press.
- Azzouni, J. (2013). *Semantic perception: how the illusion of a common language arises and persists*. Oxford University Press.
- Beaver, D., & Clark, B. (2008). *Sense and sensitivity: how focus determines meaning*. Oxford University Press.
- Bloomfield, L. (1933). *Language*. Henry Holt.
- Borer, H. (2005). *Structuring sense, volume I: in name only*. Oxford University Press.
- Bromberger, S. (1989). Types and tokens in linguistics. In A. George (Ed.), *Reflections on Chomsky* (pp. 58–89). Basil Blackwell.
- Bromberger, S. (2011). What are words? Comments on Kaplan (1990), on Hawthorne and Lepore, and on the issue. *Journal of Philosophy*, 108, 406–503.
- Burge, T. (1989). Wherein is language social? In A. George (Ed.), *Reflections on Chomsky* (pp. 175–191). Basil Blackwell.
- Carston, R. 2002: *Thoughts and utterances: the pragmatics of explicit communication*. Wiley
- Cecchetto, C., & Donati, C. (2015). *(Re)Labelling*. MIT Press.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. MIT Press.
- Chomsky, N. (1986). *Knowledge of language: Its nature, origin, and use*. Praeger.
- Cinque, G. (2010). *Syntax of adjectives*. MIT Press.
- Collins, J. (2004). Faculty disputes. *Mind and Language*, 19, 300–333.
- Collins, J. (2014). Representations without representa: Content and illusion in linguistic theory. In P. Stalmaszczyk (Ed.), *Semantics and beyond: philosophical and linguistic inquiries* (pp. 27–64). De Gruyter.
- Collins, J. (2020). Invariance as the mark of the psychological reality of language. In A. Bianchi (Ed.), *Language and reality from a naturalistic perspective: Themes from Michael Devitt* (pp. 7–44). Springer.
- Collins, J. (2021). Copredication and the diversity of fiction: An accommodation problem. *Erkenntnis*, 86, 1197–1223.

¹⁵ My thanks go to two external reviewers for extensive constructive criticism, Georges Rey for many discussions, and James Miller, Emma Borg, Nick Tasker, Gabe Dupre, and Elliot Michaelson.

- Collins, J., & Rey, G. (2021). Chomsky and intentionality. In N. Allott, T. Lohndal, & G. Rey (Eds.), *A Companion to Chomsky* (pp. 488–502). Wiley Blackwell.
- Crain, S., & Thornton, R. (1998). *Investigations in universal grammar: A guide to experiments on the acquisition of syntax and semantics*. MIT Press.
- Devitt, M. (2006). *Ignorance of language*. Oxford University Press.
- Di Sciullo, M., & Williams, E. (1987). *On the definition of word*. MIT Press.
- Feinsinger, A. (2021). The variation problem. *Philosophical Studies*, 178, 317–338.
- Hale, K., & Keyser, S. J. (1993). On argument structure and the lexical expression of syntactic relations. In K. Hale & S. J. Keyser (Eds.), *The view from building 20: Essays in linguistics in Honor of Sylvain Bromberger* (pp. 53–109). MIT Press.
- Hale, K., & Keyser, S. J. (2003). *Prolegomenon to argument structure*. MIT Press.
- Halle, M., & Marantz, A. (1993). Distributed morphology and the pieces of inflection. In K. Hale & S. J. Keyser (Eds.), *The view from building 20: essays in Honor of Sylvain Bromberger* (pp. 111–176). MIT Press.
- Hawthorne, J., & Lepore, E. (2011). On words. *Journal of Philosophy*, 108, 447–485.
- Hellman, G. (1989). *Mathematics without numbers: Towards a modal-structural interpretation*. Clarendon Press.
- Irmak, N. (2019). An ontology of words. *Erkenntnis*, 84, 1139–1158.
- Isac, D., & Reiss, C. (2013). *I-language: An introduction to linguistics as cognitive science*. Oxford University Press.
- Johnson, K. (2015). Notational variants and invariance in linguistics. *Mind and Language*, 30, 162–186.
- Kaplan, D. (1990). Words. *Proceedings of the Aristotelian Society*, 64 (Supp.): 93–120.
- Kaplan, D. (2011). Words on words. *Journal of Philosophy*, 108, 504–529.
- Katz, J. (1981). *Language and other abstract objects*. Totowa, NJ: Rowman and Littlefield.
- Katz, J., & Fodor, J. (1963). The structure of a semantic theory. *Language*, 39, 170–210.
- Katz, J., & Postal, P. (1964). *An integrated theory of linguistic descriptions*. MIT Press.
- Katz, J., & Postal, P. (1991). Realism vs. conceptualism in linguistics. *Linguistics and Philosophy*, 14, 515–554.
- Kayne, R. (1994). *The antisymmetry of syntax*. MIT Press.
- Langland-Hassan, P., & Vicente, A. (Eds.). (2018). *Inner speech: New voices*. Oxford University Press.
- Levin, B., & Rappaport Hovav, M. (1994). *Unaccusativity: At the syntax-lexical semantics interface*. MIT Press.
- Levin, B., & Rappaport Hovav, M. (2005). *Argument realization*. Cambridge University Press.
- Lewis, D. (1969). *Convention: A philosophical study*. Harvard University Press.
- Lewis, D. (1975). Languages and language. In K. Gunderson (Ed.), *Minnesota studies in the philosophy of science* (Vol. VII, pp. 3–35). University of Minnesota Press.
- Marr, A. (1982). *Vision*. MIT Press.
- Miller, J. (2019). A bundle theory of words. *Synthese*. <https://doi.org/10.1007/s11229-019-02430-3>
- Miller, J. (2020a). The ontology of words: Realism, nominalism, and eliminativism. *Philosophy Compass*. <https://doi.org/10.1111/phc3.12691>
- Miller, J. (2020b). On the individuation of words. *Inquiry*, 63, 875–884.
- Millikan, R. G. (2005). *Language: A biological model*. Oxford University Press.
- Nefdt, R. (2019). The ontology of words: A structuralist approach. *Inquiry*, 62, 877–911.
- Nunes, J. (2004). *Linearization of chains and sideward movement*. MIT Press.
- Pustejovsky, J., & Batiukova, O. (2019). *The Lexicon*. Cambridge University Press.
- Quine, W. V. O. (1948). On what there is. *Review of Metaphysics*, 2, 21–38.
- Quine, W. V. O. (1972). Methodological reflections on current linguistic theory. In D. Davidson & G. Harman (Eds.), *Semantics of natural language* (pp. 442–454). Reidel.
- Quirk, R., Greenbaum, S., Leech, G., & Svartvik, J. (1985). *A comprehensive grammar of the English language*. Longman.
- Reinhart, T. (2006). *Interface strategies: Optimal and costly computations*. MIT Press.
- Rey, G. (2003). Chomsky, intentionality, and CRTT. In N. Hornstein & L. Antony (Eds.), *Chomsky and his critics* (pp. 105–139). Blackwell.
- Rey, G. (2020). *Representation of language: Philosophical issues in a Chomskyan linguistics*. Oxford University Press.
- Rey, G. (2022). The innocuousness of folieism and the need of intentionality where transduction fails: Replies to Adger and to Stainton & Viger. *Mind and Language*, 37, 274–282.

- Sperber, D., & Wilson, D. (1986). *Relevance: communication and cognition*. Basil Blackwell.
- Stainton, R. (2006). *Words and thoughts: Subsentences, ellipsis, and philosophy of language*. Oxford University Press.
- Stainton, R. (2011). In defence of public languages. *Linguistics and Philosophy*, 34, 479–488.
- Stojnić, U. (2021). Just words: Intentions, tolerance, and lexical selection. *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12781>
- Szabó, Z. (1999). Expressions and their representations. *Philosophical Quarterly*, 49, 145–163.
- Tasker, N. (2022). A case for eliminativism about words. *Synthese*, 200, 1–23.
- Thomasson, A. (2010). *Ordinary objects*. Oxford University Press.
- Thomasson, A. (2018). *Ontology made easy*. Oxford University Press.
- Wittgenstein, L. (1953). *Philosophical investigations*. Basil Blackwell.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.