

Freistaat Flaschenhals

or

How the language acquisition bottleneck
shaped the lexicon-syntax interface

Matthijs Westera

October 27, 2010

Supervisor: Dr. Marijana Marelj
Credits: 30 ECTS

Contents

1	Introduction	1
1.1	Introduction	1
1.2	The lexicon-syntax interface	1
1.3	The Theta System	4
1.3.1	Clusters	5
1.3.2	Linking	7
1.3.3	Operations	8
1.4	Background assumptions	11
1.4.1	Nativism	11
1.4.2	What is in the lexicon	14
1.5	The goal of this thesis	16
1.5.1	Preview	17
1.6	Relevance to my field of study	18
2	An acquired active lexicon	20
2.1	Introduction	20
2.1.1	Three views of an active lexicon	21
2.2	The origins of lexicon operations	23
2.3	The acquisition of lexicon operations	25
2.4	Acquisition in a syntax-driven approach	27
2.5	Conclusion	29
3	The origins of verb frame alternations	30
3.1	Introduction	30
3.1.1	An intuitive (but wrong) explanation	31
3.2	Iterated learning	32
3.3	Universality	36
3.4	The Lexicon-Syntax Parameter	38
3.5	Frozen forms	40
3.6	Conclusion	43
4	The linking procedure	44
4.1	Introduction	44
4.1.1	Communication, thought and the linking procedure	45

<i>CONTENTS</i>	iii
4.2 The linking procedure and semantic bootstrapping	46
4.3 The origins of linking regularities	48
4.4 Thematic distinctness	52
4.4.1 Parsons (1990): event semantics	53
4.4.2 L-thematic distinctness through iterated learning	55
4.4.3 Carlson (1998): event individuation	56
4.4.4 I-thematic distinctness through iterated learning	59
4.5 Conclusion	59
5 Conclusion	61
References	64

Acknowledgements

Many, many thanks to Marijana Marelj for the pains she went through in order to get me to focus, clarify, focus, make my assumptions explicit, focus, make my assumptions explicit, clarify, and ultimately finish my thesis. I am unsure why progress was so slow in this respect. I appreciated our intense discussions (by and large) and I hope Anna did too.

Thanks to Hanna de Vries for helpful initial discussions and comments on an early draft.

Many thanks to Eric Reuland and Albert Visser, whose critical yet encouraging Sunday-afternoon comments I wish I had gotten earlier (the blame is on me). Processing them has not only improved my thesis, it has also taught me that all one really needs is a ten-minute powernap every twenty-four hours and a steady supply of slow sugars and tea.

Well, and fruit and biscuits. Thank you Mieke Lam, for feeding me throughout the final showdown.

Chapter 1

Introduction

1.1 Introduction

Because this introductory chapter is of considerable length (indeed, it is the longest chapter in the entire thesis), an introduction to the introduction is in place. My goal in this thesis is to approach a theoretical framework from the field of generative linguistics from a new angle, tentatively dropping some of the background assumptions to see how much of the framework can be maintained. The framework under consideration is called the *Theta System*, and it is a theory of, mainly, how the meaning of a word determines its syntactic behaviour. The link between word meaning and syntax exists at what is called the *lexicon-syntax interface*, which is introduced in section 1.2. Section 1.3 gives an overview of the Theta System. Section 1.4 discusses the background assumptions that I wish to get rid of, at least partially. Section 1.5 provides a more detailed description of the goal of this thesis and an overview of its chapters. Section 1.6 explains why this goal is relevant to my field of study, Cognitive Artificial Intelligence.

1.2 The lexicon-syntax interface

Human languages consist of a collection of lexical elements (words and elements that make up words, but also fixed expressions) and ways to recursively combine these lexical elements into an infinity of larger expressions (ultimately sentences). By assumption, lexical elements are stored in the human mind in the *lexicon* and the combination of lexical elements takes place in the *syntax*. From the syntax, syntactic structures are passed on to the *sensori-motor system*, where they are turned into speech and gestures, and to the *inference system*, where they are interpreted and inference take place.¹ Each of these systems,

¹I adopt the terminology of the literature on the Theta System, which I discuss below (e.g. Reinhart, 2002). In the literature on the Theta System, the lexicon is equated with the ‘system of concepts’. The inference system is also sometimes called ‘interpretation system’ (e.g. Reuland, 2010) or the conceptual-intentional system (e.g. Hauser, Chomsky, & Fitch,

perhaps along with some others, is contained in what Hauser et al. (2002) call the *faculty of language in the broad sense* (FLB).

One may wonder to what extent FLB can really be split up into those components, i.e. to what extent they are informationally encapsulated *modules* in Fodor's (1983, 2000) sense. I do not wish to go into this discussion, and in this thesis I simply assume that they are modules (an intuitive argument for a syntax module is that we can understand semantically strange, implausible or peculiar sentences, like 'cats chase dogs' (Foster, 1979)). At the interfaces between modules, information is filtered and perhaps preprocessed/postprocessed. Syntax, together with its interfaces to other cognitive systems, is what Hauser et al. (2002) call the *faculty of language in the narrow sense* (FLN). Because FLN seems to be uniquely human (as opposed to the other components of FLB), Hauser et al. tentatively propose the hypothesis that FLN is the only part of FLB that is an evolutionary adaptation exclusively for language (but for discussion, see Pinker & Jackendoff, 2005; Fitch, Hauser, & Chomsky, 2005; Jackendoff & Pinker, 2005).

This thesis concerns the lexicon and its interface with syntax. Lexical entries may impose requirements on the structures in which they occur (at least in one branch of approaches, as I discuss further below). Lexical entries that take arguments, like verbs, impose three kinds of requirements on their arguments: selection for particular semantic types ('s-selection'), selection for particular syntactic categories ('c-selection') and assignments of arguments to syntactic positions like Subject and Object ('linking') (Pesetsky, 1995). Throughout this thesis I will mark s-selection violations by '#', c-selection and linking violations with '*', and semi-violations of any kind by '?'.

- (1) S-selection:
 - a. Alan/the heat defrosted the ice
 - b. Alan/#his binoculars saw the mountain
 - c. Bertrand drank soy milk/#a veggie burger.
- (2) C-selection:
 - a. John asked the time.
 - b. *John wondered the time.
 (Pesetsky, 1995, p.3)
- (3) Linking:
 - a. Bill threw the ball.
 - b. *The ball threw Bill. (with Bill as Agent)
 (Pesetsky, 1995, p.3)

In addition, some lexical entries may occur in more than one *thematic form*, yielding regular *verb frame alternations* such as the causative-inchoative alternation:

2002).

- (4) a. Alan/the heat defrosted the ice. (causative)
 b. The ice defrosted. (inchoative)
 c. Alan/#his binoculars saw the mountain.
 d. *The mountain saw. (with the mountain as the thing being seen)

One of the aims of linguistics is to relate the syntactic behaviour of a lexical entry to its meaning, i.e. to reduce c-selection, linking and verb frame alternations to s-selection. For instance, the contrast in (3) may be due to the fact that Agents must always merge in Subject position rather than Object position, and the contrast in (4) may be due to the fact that *see*, and not *defrost*, requires an animate argument. The main motivation for grounding generalisations over syntactic behaviour in generalisations over meaning is that it reduces the number of entities that have to be postulated in a linguistic theory. Pesetsky (1995) motivates this reduction in terms of lexicon acquisition:

Since s-selection itself is an aspect of lexical semantics, we will have the best theory of the lexicon we can hope for: children learn pairings of sound and meaning, UG [Universal Grammar] does the rest. (Pesetsky, 1995, p.4)

The links between s-selection and c-selection and linking facilitate lexicon acquisition regardless of whether the links are innate, as Pesetsky proposes (Universal Grammar), following mainstream generative linguistics, or acquired, as e.g. Pinker proposes (cf. section 1.4).

Two branches of approaches exist that aim to capture the part of a lexical entry's meaning (s-selection) relevant for argument projection (c-selection, linking, verb frame alternations). Following the terminology of (Levin & Rappaport-Hovav, 1995), in *predicate decomposition* approaches, a predicate's meaning is decomposed into a nested structure of primitive predicates, some of which take one or several of the original predicate's arguments (e.g. Jackendoff, 1987; Levin & Rappaport-Hovav, 1995; Pustejovsky, 1995 could be grouped with these). In *role-list* approaches, a predicate's meaning is decomposed into the properties of its arguments, captured in a list of *thematic roles* or *theta roles* (Dowty, 1991; Tenny, 1994; Reinhart, 1991 *et seq.*). But as long as one does not assume that the parts of meaning affecting argument projection are represented separately from the rest of a lexical entry's meaning (as e.g. Jackendoff assumes), role-list approaches and predicate decomposition approaches are but different descriptions of the same thing, and they are representationally equally powerful (though each probably has different methodological advantages and disadvantages). Hence, the majority of the points made in this thesis with regard to the Theta System, a typical role-list approach, will equally apply to predicate decomposition approaches.

Another apparent distinction commonly found in the literature is that between *thematic approaches* and *event-structure-based approaches* (Arad, 1998; Benua & Borer, 1996; Borer, 2004). But this distinction cuts no ice. Typical 'thematic' approaches such as (Dowty, 1991; Reinhart, 2002) often invoke typical 'event-related' entailments like causation and change of state; Pykkänen

(1997) argues that the typically ‘thematic’ notion of experience be included in event structure by means of a perceive-predicate; Grimshaw (1990) argues that thematic and aspectual properties together determine argument projection; Tenny (1994) invokes a set of thematic roles defined in terms of event structure (Tenny calls these ‘aspectual roles’); and Borer (2005) likewise defines roles in terms of an event structure such as ‘originator’ and ‘subject-of-change’. Hence, the distinction should be made not between approaches, but between different parts of meaning: the thematic component and the event-related (or aspectual) component. If anything, the literature shows convergence on the idea that both meaning components (if they can be separated at all) affect argument projection.

A more essential distinction, and one that divides the linguistic field, is that between *lexicon-driven* and *syntax-driven* approaches (Borer (2005) uses the terms ‘endo-skeletal’ and ‘exo-skeletal’ to refer to the same distinction). The Theta System, introduced below, is a lexicon-driven approach: syntax projects from the information contained in a lexical entry. In syntax-driven approaches, lexical entries contain only the root component of meaning, while the rest of the meaning is contributed by one of several syntactic templates that exist independently of lexical entries (e.g. Hale & Keyser, 2002; Arad, 1998; Borer, 2005). In other words, the relation between semantics and syntax (i.e. between s-selection and c-selection, linking and verb frame alternations in lexicon-driven approaches) is moved from the lexicon-syntax interface to the syntax-inference interface. The main source of disagreement between lexicon-driven and syntax-driven approaches is whether the variable behaviour of words in a language is best captured by means of a set of syntactic templates in which each word may appear, or by means of an *active lexicon* in which one form is derived from the other. A considerable portion of this thesis is devoted to arguing that the lexicon is active. Since the Theta System is currently the most elaborate lexicon-driven approach, this conclusion will retroactively motivate my choice to devote my thesis to the Theta System.

1.3 The Theta System

The Theta System is the result of a lexicon-driven role-list approach that has been under development since the nineties (Reinhart, 1991 *et seq.*, for an overview see Reinhart, 2002). The Theta System enables the interface between the lexicon and syntax and, via syntax, the inference system. It encodes lexical concepts in a way that is legible to syntax (linking instructions) and inference (sets of features). In addition to enabling the interface, the Theta System was developed to capture the relatedness between various thematic forms of a given verb. It consists of three components:

- A formal encoding of predicates in terms of a theta grid, in which the thematic roles of each argument are represented as clusters of semantic features. The features are legible to the inference system as entailments.

- A linking procedure that derives linking instructions from the theta grid. The output of the linking procedure is legible to syntax.
- A set of operations on theta grids that derive one thematic form from another.

1.3.1 Clusters

In the Theta System, theta roles are clusters composed of two binary features $[\pm c]$ and $[\pm m]$ that specify the causal involvement and mental involvement of the role bearer, respectively. The decomposition of thematic roles into features was first proposed by Dowty (1991) as a solution to the ill-definedness of thematic roles at the time. Reinhart (2002) defines causal involvement as a perception-driven property, that holds of an argument if it is perceived to be a *necessary and sufficient condition* for the event to take place.² A reliable way to test the causal involvement of an argument is to paraphrase the sentence with ‘cause to’ or ‘make’, as in (5)

- (5) a. Alan broke the vase / Alan caused the vase to break
 b. Alan spun the wheel / Alan made the wheel spin

The feature $[\pm m]$ denotes not whether an event participant is an animate being, but whether it is mentally involved in a way relevant to the event being described. Hence, $[+m]$ entails animacy, but not vice versa.

Each feature can be specified (plus or minus) or unspecified, giving rise to nine possible clusters:

- (6) All-plus clusters:
 a. $[+c+m]$ (Agent)
 b. $[+c]$ (Cause)
 c. $[+m]$ (Sentient)
- (7) All-minus clusters:
 a. $[-c-m]$ (Theme)
 b. $[-c]$ (Recipient Goal/Benefactor)
 c. $[-m]$ (Subject Matter/Source)
- (8) Mixed clusters:
 a. $[+c-m]$ (Instrument)
 b. $[-c+m]$ (Experiencer)
- (9) Empty cluster:
 a. $[\]$ (?)

²Here the ‘perceived to be’ does the trick. After all, logically, if all participants of the event described by ‘The rock broke the window’ are necessary, how can any one of them be sufficient, i.e. a cause?

As indicated in parentheses, eight of these clusters can be roughly paired with the more traditional role labels (only roughly, for the traditional labels are all but clear).³ The four (non-empty) underspecified clusters are ambiguous. For example, while *eat* only has an agentive reading (10), *break*, taking an underspecified [+c] role, has both a non-agentive and an agentive reading (11):⁴

- (10) *eat*([+c+m][-c-m])
- a. Alan ate three slices of bread.
 - b. #The spoon ate three slices of bread.
- (11) *break*([+c][-c-m])
- a. The rock (#volitionally) broke the window.
 - b. Alan (volitionally) broke the window.

In (Dimitriadis, 2004; Reinhart & Siloni, 2003; Marelj, 2004), Parsons' (1990) neo-Davidsonian event semantics is used as a model for information at the level of inference. The neo-Davidsonian event semantics analysis, which incorporates thematic roles into the event semantics of Davidson (1967), is believed to be superior to other accounts of sentence meaning because it captures the entailment relations that exist between propositions. For example, (12a) entails (12b), (12c) and (12d) (but not vice versa), and (12b) and (12c) entail (12d) (but not vice versa). These entailments can be accounted for by assuming the event semantic representation for (12a) in (13) (and similar for the other sentences).

- (12) a. Brutus stabbed Caesar in the back with a knife.
 b. Brutus stabbed Caesar in the back.
 c. Brutus stabbed Caesar with a knife.
 d. Brutus stabbed Caesar.
- (13) $(\exists e)[\text{STABBING}(e)\&\text{Agent}(e, \text{Brutus})\&\text{Theme}(e, \text{Caesar})\&\text{In}(e, \text{back})\&\text{With}(e, \text{knife})]$

The feature clusters of the Theta System can be straightforwardly used to spell out the entailments denoted by the labels Agent, Theme, etcetera, in the event semantic representation.

Two constraints on theta clusters exist, formulated in (Marelj, 2002), of which one will play a marginal role in the present thesis. For completeness sake, I will mention them both here. First is the Full Interpretation of Thematic Roles (14), which captures the fact that the expansion of underspecified clusters, e.g. of [+c] as either [-c-m] or [+c+m], is mandatory for interpretative purposes.

³The function of the empty cluster (both feature values unspecified) is irrelevant in this thesis, but see (Marelj, 2004) for a possible use of the empty cluster as denoting the arbitrary human event participant of middles.

⁴I follow the Theta System's conventional notation for lexical entries and their theta grids, like *break*([+c][-c-m]) in the example.

(14) **Full Interpretation of Thematic Roles (FITR):**

For the purposes of interpretation, all clusters must be fully specified. (Marelj, 2002, p.3)

Second is the Identity Constraint (15), which is widely assumed in the literature, though under different names (e.g. Parsons, 1990; Carlson, 1998).

(15) **Identity Constraint (IC):**

Two identical theta-roles cannot realize on the same grid. (Marelj, 2002, p.2)

The Identity Constraint captures the observation that there are no verbs like *skick*, meaning KICK but with two Agents, as in ‘Alan skicked the ball Bertrand’ (Carlson, 1984). Marelj assumes that the Identity Constraint holds at the interface between the lexicon and syntax as well as at the interface between syntax and the inference system.

1.3.2 Linking

Following the prevailing view in the literature (e.g Grimshaw, 1981; Chomsky, 1986; Pesetsky, 1995), the Theta System contains a procedure for predicting a verb’s linking from its s-selectional restrictions (i.e. its theta grid). The linking procedure consists of two steps. The first step marks each theta role of a verb with an index depending on the polarity of its features. The second step takes a marked lexical entry and, based on the indices, merges the arguments internally (Object position) or externally (Subject position). Formally, both steps happen as follows:

(16) Marking: given an n -place verb-entry, $n > 1$

- a. Mark clusters an all-minus cluster ($[-c]$, $[-m]$, $[-c-m]$) with index 2
- b. Mark clusters an all-plus cluster ($[+c]$, $[+m]$, $[+c+m]$) with index 1
- c. If the entry includes an all-plus cluster ($[+c]$, $[+m]$, $[+c+m]$) and in addition $[-c+m]$ or $[-c-m]$, mark the verb with the ACC feature (i.e. the entry assigns accusative case).

(17) Merging:

- a. When nothing rules this out, merge externally
- b. The bearer of a role marked 2 merges internally
- c. The bearer of a role marked 1 merges externally

Note that the linking instructions, i.e. the indices 1 and 2 and the ACC feature, are the only aspects of a lexical entry that are legible to syntax. The content of the theta clusters is only passed on through syntax to the inference system. It is assumed in the Theta System that the linking instructions are derived every time a lexical entry is used. That is, they are not stored on lexical entries.

A transitive verb like *see* yields the derivation in (18), and an *unergative* intransitive verb like *buzz* yields the derivation in (19). So-called *unaccusative*

intransitive verbs seem to violate the linking procedure, in that its single argument merges in Object position (afterwards moving to Subject position, leaving a trace) (20). The Theta System explains the behaviour of such verbs by allowing changes to the theta grid to occur in between marking and merging, according to a set of operations that I discuss below. Unaccusative intransitive verbs are assumed to derive from two-place entries that have lost their [+c] argument due to *decausativisation*, leaving an argument marked ‘2’ that merges in Object position despite the verb being intransitive.

- (18) Transitive derivations
- a. Max saw him
 - b. Base entry: *see*([+m],[−c−m])
 - c. Marking: *see*_{ACC}([+m]₁,[−c−m]₂)
 - d. Merging: *Max* ([+m]₁) merges externally, *he* ([−c−m]₂) merges internally and checks accusative case (*him*)
- (19) Unergative derivations
- a. The bell buzzed
 - b. Base entry: *buzz*([−c−m])
 - c. Marking: *buzz*([−c−m]) (base entry is intransitive, hence no marking)
 - d. Merging: *the bell* ([−c−m]) merges externally
- (Marelj, 2004, p.26)
- (20) Unaccusative derivations
- a. The vase_{*i*} broke ___{*i*}
 - b. Base entry: *break*([+c],[−c−m])
 - c. Marking: *break*_{ACC}([+c]₁,[−c−m]₂)
 - d. Decausativisation operation: *break*([−c−m]₂)
 - e. Merging: *the vase* ([−c−m]₂) merges internally
- (Marelj, 2004, p.26)

These are just three examples of the enormous empirical coverage of the Theta System that is due largely to this simple linking procedure in combination with a set of operations, to be treated next.

1.3.3 Operations

As sketched in 1.2, many verbs can appear in various thematic forms. For instance, the causative-inchoative alternation is available cross-linguistically for verbs that assign a [+c]-role, like *defrost*:

- (21) a. Alan/the heat defrosted the ice. (causative)
 b. The ice defrosted. (inchoative)

These verbs are not only phonologically identical (and in most languages they share a morphological root), the causative logically *entails* the inchoative. The Theta System captures such relations between verb frames by assuming that for each verb the lexicon contains only one basic verb entry, with one thematic structure, from which the other thematic forms can be derived via *lexicon operations*. This is formulated in the Lexicon Uniformity Hypothesis, which lies at the heart of the Theta System:

(22) **Lexicon Uniformity Hypothesis**

Each verb-concept corresponds to one lexical entry with one thematic structure. The various thematic forms of a given verb are derived by lexicon operations from one thematic structure. (From Reinhart, 2000, p.20)

A similar strategy for explaining verb frame alternations is pursued in e.g. (Chierchia, 1989; Levin & Rappaport-Hovav, 1995; Pesetsky, 1995; Pustejovsky, 1995). The Theta System contains at least the following operations (Reinhart, 2002):

(23) **Saturation** existentially closes an argument, which is then available for interpretation but normally not realised in the syntax. When applied to the external argument, this operation forms passives:

- Alan_[+c] hit Bertrand_[-c-m] /
Bertrand_[-c-m] was hit (by something_[+c])

Saturation also forms middles and impersonals (Chierchia, 1995; Marelj, 2004).

(24) **Decausativisation** (also: expletivisation or external reduction) applies to a polyadic entry that assigns a [+c] role, and removes this role from the theta grid. Decausativisation applies after the marking procedure, leaving the indices of the remaining roles intact (but removing the ACC feature from the verb entry). Decausativisation derives unaccusative verbs like intransitive *break*, but also intransitive experiencer verbs like *worry*. The incompatibility of an instrument with decausativised entries shows that the role is removed entirely.

- Alan_[+c] broke the window_[-c-m] (with a rock) /
The window_[-c-m] broke __ (*with a rock)
- Alan_[+c] worried Bertrand_[-c+m] / Bertrand_[-c+m] worried

(25) **Reflexivisation** (also: bundling) identifies two roles, creating a verb with reduced arity in which a single argument fulfills two thematic roles. Reflexivisation applies to (a subset of) agentive verbs, i.e. verbs that assign a [+c+m] role. Like decausativisation, it removes the ACC feature from the verb entry.

- Alan_[+c+m] washed himself_[-c-m] / Alan_{[+c+m]&[-c-m]} washed

Reciprocalisation, an operation reminiscent of reflexivisation, forms reciprocal meanings instead.

- (26) **Agentivisation** (also: lexical causativisation) adds a [+c+m] role. If the original verb entry already contained a cluster with a [+c] feature, the feature of that cluster is inverted to [-c]. Agentivisation drastically changes the theta grid, hence the marking procedure must apply afterwards.

- The horse_[+c+m] galloped / Alan_[+c+m] galloped the horse_[-c+m]

Considerable cross-linguistic variation is explained in the Theta System by assuming that operations may take place either in the lexicon or in the syntax, depending on the language's setting of the *Lexicon-Syntax Parameter*.

- (27) **The Lexicon-Syntax Parameter:**

UG [universal grammar] allows thematic arity operations to apply in the lexicon or in syntax. (Reinhart & Siloni, 2003)

- (28) Some languages (Reinhart & Siloni, 2003, p.19):

- a. Lexicon setting: Hebrew, Dutch, English, Russian, Hungarian
- b. Syntax setting: Romance, German, Serbo-Croatian, Czech, Greek

One of the many diagnostics given in (Reinhart & Siloni, 2003) to determine a language's parameter setting is the unavailability of reflexive agent nominals in syntax languages. Since the nominalisation of verbs (e.g. deriving *dresser* from the verb *dress*) takes place in the lexicon, a reflexive verb is able to nominalise only if it can be reflexivised already in the lexicon, which is possible in English but not in French:

- (29) a. She dresses slowly because she is an elegant dresser
 b. Jean est un excellent habilleur/ maquilleur
 Jean is an excellent dresser/ make-up-er (of others only - no reflexive reading)

While all operations can apply in the lexicon, not all operations can apply in the syntax, according to the Lexicon Interface Guideline (see Marelj, 2004 for a more parsimonious explanation in terms of inclusiveness in syntactic derivations):

- (30) **The Lexicon Interface Guideline:**

The syntactic component cannot change theta grids: Elimination and modification of a theta role as well as addition of a role to the theta grid are illicit in syntax. (Siloni, 2002, p.2)

It follows that decausativisation and agentivisation can only apply in the lexicon. Indeed, no cross-linguistic variation of the sort reflexivisation is subject to is attested with respect to those operations.

1.4 Background assumptions

The Theta System is not just an concise description of linguistic data. It comes with a number of background assumptions, some firmly rooted in Chomskyan generative linguistics, other more theory-internal. In what follows I make explicit the somewhat unorthodox position that I take in this thesis with regard to the Theta System, and explain the reason why.

1.4.1 Nativism

I-language, E-language and Universal Grammar

Chomsky (1997) distinguishes *I-language*, where ‘I’ stands for individual and internalised, from *E-language*, where ‘E’ stands for extensional and externalised. An I-language is a state of the language faculty, i.e. it is a computational system implemented on the human brain that performs the kind of processing required for creating and understanding linguistic utterances. To study I-language is to study the language faculty in the brain.

The term ‘E-language’ is harder to define; an E-language could be a collection of utterances used in a particular linguistic community or collected in a particular corpus, for instance the corpus of all English newspapers. But this would only raise additional questions: what is a linguistic community, what is English? Speakers of what one would intuitively consider to be the same language inevitably differ in how they interpret certain words and in the constructions they find grammatical. The sets of utterances used in any group of individuals, i.e. candidate referents for the word ‘E-language’, are bound to lack the homogeneity required of phenomena suitable for scientific study of what it means to be a human language. Hence, according to Chomsky, the object of study for linguists is I-language: language as a computational system represented in the brain. In my thesis, too, I-language is the primary object of study.

Closely related to the study of I-language is the study of Universal Grammar, which is defined as the state of the language faculty prior to acquiring a specific language like English or Chinese. It is what allows humans, contrary to other animals, to acquire a human language, and to acquire such a language despite an *acquisition bottleneck*, or Chomsky’s (1965) ‘poverty of the stimulus’. This has also been called the *logical problem of language acquisition*: How can infants acquire one out of an infinity of grammars capable of producing an infinity of sentences, through exposure to only a finite sample of those sentences (Chomsky, 1965)?

Universal Grammar, though a necessary solution to the logical problem of language acquisition, does not need to explain all linguistic universals. For instance, the fact that in most languages a sound like *mama* denotes MOTHER is an obvious consequence of the fact that *mama* is one of the earliest sounds uttered by an infant; the universality of a regular signal-meaning mapping can perhaps partly be explained as a result of gradually accumulating language change

(Kirby & Hurford, 2002, see section 3.2 below); basic-level categories in the lexicon/vocabulary (e.g. CHAIR as opposed to FURNITURE or ROCKING CHAIR) are universal because they are the most useful for communication (Rosch, 1978, see section 3.3 below).

Why the origins of language universals are relevant

It is important that we know which universals are due to Universal Grammar, and which universals have other causes. Aside from the fact that Universal Grammar is an interesting entity in its own right, such an understanding is necessary for a deeper understanding of I-language, which is the ultimate goal. If we could freely postulate anything in Universal Grammar, then *any* mechanism in I-language would be equally plausible. For all we know, UG may contain all sentences (down to a certain depth of recursion) in a simple (but huge) look-up dictionary, reducing I-language to a mechanism for looking up the right sentence - but obviously this does not qualify as a proper scientific explanation of the linguistic data. In Chomsky's (1995) Minimalist Program, the parsimony of generative linguistic theory is increased by reducing Universal Grammar to two simple operations, Merge and Agree, together with some properties that follow from general computation (e.g. cognitive economy).⁵

Why is this relevant for my account of the Theta System? **We can only evaluate the claim that I-language contains an active lexicon if we can explain how it got there.** The Theta System explains the existence of an active lexicon in I-language by postulating that an active lexicon is innate, i.e. part of Universal Grammar. But innate lexicon operations may be unnecessary for the logical problem of language acquisition - at least, I am unaware of an argument to the contrary. And if innate lexicon operations are not necessary for the logical problem of language acquisition, then their postulation is only supported by empirical data. Perhaps the same data could be explained without postulating innate lexicon operations. What I wish to stress is that the innateness of lexicon operations, though a fine working hypothesis, need not be taken for granted.

Concept nativism

Related to the logical problem of language acquisition is the problem of word learning. How do language learners learn the amount of words they do, at the speed they do? A possible answer is that concepts are innate, while only the labels for concepts are acquired (Fodor, 1983, but see Fodor, 2000 for non-nativist view). Chomsky (1997) remarks:

⁵Although a promising step towards a more explanatory linguistic theory, the economy-driven Minimalist Program may be misguided from an evolutionary perspective (Parker, 2006). Though I do not wish to dwell on this, her main argument is that language shows signs not of perfection, but of complexity, redundancy and adaptation, containing many extraneous operations and features, as witnessed also, she argues, in explicitly minimalist accounts.

This [the ease by means of which infants seem to acquire new words] can only mean that the concepts are already available, with all or much of their intricacy and structure predetermined, and the child's task is to assign labels to concepts, as might be done with very simple evidence. (Chomsky, 1997, p.29)

Horvath and Siloni (2005) call what is acquired the *vocabulary* of a language, which is a list of labels, as opposed to the *lexicon*, which is the innate system of concepts. A possible, intuitive rationale for concept nativism is that some kinds of concepts just happen to come naturally, given the architecture of our brains, while others do not.

Concept nativism is a background assumption in generative linguistics and, hence, the mainstream Theta System. But in the Theta System a particular kind of concept nativism is advocated, according to which some thematic forms of a verb concept are innate while others are not 'just innate', but derived from other thematic forms via a set of innate operations. For instance, there is a group of intransitive verbs of which the denotations are derived from a causative concept. Apparently there is something about the denotations of such intransitive verbs due to which they do not 'come naturally', or at least not as naturally as the denotations of other intransitive verbs. There seem to be concepts that do not come naturally to us, such as categories of four-dimensional objects, so it is possible in principle. However, the great mystery of the set of intransitive verbs under consideration is that *it cannot be characterised semantically*. There is nothing about the denotations of such verbs that could possibly explain why they are 'less innate' than the denotations of other intransitive verbs. The only difference between these and other intransitive verbs is that their denotations are derived from those of causative verbs, but that is precisely the difference that I would like to see explained. In the same vein, there is no semantic difference that can explain why the causative alternates of some intransitive verbs are innate while the causative alternates of other intransitive verbs are not.

Perhaps we must take for granted that the innate system of concepts comes with a large class of exceptions, i.e. a semantically uncharacterisable set of intransitive verb concepts that must be derived. Perhaps the intuitive rationale for concept nativism, that some concepts happen to 'come naturally', is completely wrong. Perhaps I have overlooked something, as has often been the case during my adventures in the field of generative linguistics. But in any case, it will not hurt to explore what other options are available:

- There are no lexicon operations; or
- Concepts are not innate.

Neither option seems particularly easy to reconcile with the mainstream Theta System and with generative linguistics in general. Nevertheless, in this thesis I shall assume, only as a working hypothesis, the second option: concepts are not innate. If we ignored the acquisition problem that motivates concept nativism, hoping that sooner or later a solution to it will be discovered, where would that

lead us? What would be the consequences for the Theta System? Would its case in favor of an active lexicon be strengthened, or weakened? Would its empirical coverage remain intact? In this thesis I pursue such questions. I will give the long and short of my approach in section 1.5.

Under the assumption that concepts are acquired, the difference between lexicon and vocabulary fades and they coalesce into one storage facility, which I will simply call ‘lexicon’ - accepting the risk of confusion in order to avoid ‘lexabulary’ or ‘voxicon’.

E-language/primary linguistic data

Because my working hypothesis is that concepts are not innate but acquired, it will be necessary to occasionally refer to what many linguists call *primary linguistic data*. But I will use the term ‘E-language’, properly defined below, rather than ‘primary linguistic data’, because (i) the word ‘linguistic’ in ‘primary linguistic data’ is not very well-defined (it cannot mean *I-linguistic*, after all), (ii) the term ‘E-language’ is not uncommon in this sense in language evolution research (e.g. Kirby & Hurford, 2002; Smith, Kirby, & Brighton, 2003), upon which some aspects of my account of the Theta System will be based, and (iii) ‘E-language’ and ‘I-language’ sound very harmonious together.

I propose the following definition of E-language:

(31) **E-language:**

An E-language consists of all the data used for acquiring an I-language.

Figure 1.1 schematically depicts the E-languages, defined as such, of two speakers, x and y , who each construct their own I-language from it. This simple definition is very broad (‘all the data’ must include not only tons and tons of utterances, but also the contexts in which they were used and by whom they were uttered), but it is clearly delineated where it matters. By defining E-language in terms of I-language (a computational system in the brain of an individual), I avoid the problem of having to define what constitutes a suitable corpus, a linguistic community or a language like English. I believe this definition captures the sense in which ‘E-language’ has been used in language evolution research, although it has never been made explicit.⁶

1.4.2 What is in the lexicon

The mainstream Theta System makes two rather strong claims about what is *not* in the lexicon. First, the first clause of the Lexicon Uniformity Hypothesis demands that only one thematic form of each verb is stored in the lexicon. Second, the linking instructions (the indices and the case feature) of a lexical entry are not stored in the lexicon. Both claims strike me as peculiar for three reasons. The first two are orthogonal to the concept nativism debate, the third

⁶I have seen language evolutionists equate E-language with Chomsky’s notion of *performance* once, though I am positive that it was my definition they actually intended to convey.

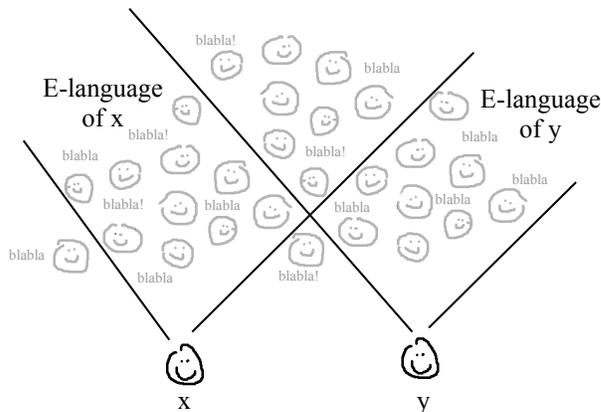


Figure 1.1: A schematical depiction of E-language, defined as the data used in constructing an I-language.

has a concept-non-nativist version and a concept-nativist version. My first reason, more an intuition than a proper reason, is as follows. Both claims concern the non-storage of something plainly visible in the input to a learner: argument structure. While one cannot rule out that our system of concepts has evolutionarily developed a language-specific trait to ignore anything of linguistic interest in the input, it does sound a little strange.

Second, the speaker would have to perform the same linking procedure and the same lexicon operation every time a particular concept must be expressed, which seems terribly inefficient. The learner may as well store them for later usage - long-term storage capacity never seems to be a problem for the human brain. The capacity of working memory, on the other hand, presumably the locus of temporary storage of derived thematic forms, is limited.

Third (concept-non-nativist version), if concepts are acquired, both claims make it impossible to acquire *exceptions* to the verb frame alternations and exceptions to the linking generalisations. How could a learner possibly know that *hit* does not decausativise, if it has not stored any of the thematic forms in which it was encountered? And how could a learner possibly know that *break* is perfectly regular, if (s)he does not have any record of its past usage? Because a learner does not know beforehand whether something will turn out to be an exception or not, it must keep track of the thematic forms and linking instructions of every verb.

Third (concept-nativist version), since vocabulary items are acquired, the link between vocabulary items and thematic forms of concepts must be acquired and stored. For instance, the infant must learn that *die* denotes inchoative DIE while *kill* denotes causative DIE. How can the link between *die* and inchoative DIE be stored if the lexicon only contains causative DIE? If inchoatives are not stored in the lexicon, then the vocabulary itself must contain some (abstract? simplified?) representation of inchoative DIE to which *die* can be linked. It

seems easier if inchoative forms are just stored in the lexicon, like all other concepts.

While I do not believe these arguments can be decisive against the two claims, they at least outnumber the zero arguments in favour of them. Therefore, I will assume in this thesis the following: **all thematic forms of a verb concept are stored in the lexicon**, and **lexical entries contain their linking instructions**. The resulting view of the lexicon is similar to the one assumed in Lexical Functional Grammar, in which all thematic forms are stored as separate entries in the lexicon, related by ‘lexical redundancy rules’ (Kaplan & Bresnan, 1982). A considerable part of my account of the Theta System will hinge on this decision, thereby retroactively supporting it.

1.5 The goal of this thesis

The Theta System is a concise system with enormous empirical coverage, but its background assumptions - the combination of concept nativism with lexicon operations, as well as the stripped-down view of the lexicon - are perhaps not unproblematic. Even when staying within the framework of generative linguistics, the Theta System’s reliance on Universal Grammar (innate lexicon operations, UG parameters, an innate linking procedure⁷) for the explanation of universals is not particularly ‘minimalist’. Recall, from the section on concept nativism, the two options I considered as a working hypothesis:

- There are no lexicon operations; or
- Concepts are not innate.

In this thesis I will pursue the second option. While this leaves open either innate or acquired lexicon operations, and either an innate or acquired linking procedure, I will maximally pursue non-nativism and reduce as many of the Theta System’s postulates to acquisition as I can. I will not claim that Universal Grammar is empty - indeed, *something* must be innate in order to solve the logical problem of language acquisition. I will only explore the idea that, be in Universal Grammar what may, perhaps it is not the Theta System. I will call my account of the Theta System the ‘non-nativist account’ or, more often, ‘my account’, as opposed to the ‘mainstream Theta System’.

I will approach the Theta System from a fresh angle, centered on language acquisition and language change rather than nativism, and reinterpret/reconstruct its postulates from the very basics. Obviously this goal does not translate into a single research question. But I will suggest answers to at least the following questions. Is the lexicon really active? What are the origins of lexicon operations? Why is there a Lexicon-Syntax parameter? What are the origins of the linking procedure? Why do verbs never assign two identical roles?

⁷Though I admit it is not trivial why a linking procedure could be un-innate, given that language must be used for thought as well as language. I will get back to this much later in the thesis.

The core mechanism that I will invoke for answering these questions is *iterated learning* (e.g. Kirby & Hurford, 2002). Summarizing, iterated learning is the process whereby each generation acquires an I-language from the E-language produced by the previous generation. Iterated learning of lexical concepts or, for the concept nativist, vocabulary labels, is subject to an *acquisition bottleneck*: only some proportion of the words in a language is acquired entirely from direct evidence. Small changes to a language due to imperfect learning may accumulate over time, resulting in universal patterns; universal patterns that, I will attempt to show in this thesis, may explain the postulates of the Theta System. Hence the title of my thesis: **How the language acquisition bottleneck shaped the lexicon-syntax interface.**⁸

In chapter 2 I argue, from the perspective of ‘my account’ of the Theta System, in favor of an acquired active lexicon. In chapter 3 I explain the origins of verb frame alternations as a result of accumulating language change due to the acquisition bottleneck. Together, both chapters form a strong argument in favor of the active lexicon and against syntax-driven approaches. Chapter 4 explains the origins of the linking procedure and targets some open questions raised by the preceding chapters on the interaction between lexicon operations and the linking procedure. Chapter 5 concludes this thesis.

1.5.1 Preview

Despite the radical differences between mine and the mainstream Theta System’s core assumptions, I will show that the Theta System in my account can remain largely intact. Its postulates can be explained without nativism, in terms of language acquisition and language change, and, as far as I can discuss this within the scope of the thesis, its empirical coverage stays almost the same.

Surprisingly, perhaps, there happens to be a fairly direct mapping between my account of the Theta System, which tentatively rejects concept nativism, and the other option, which maintains concept nativism but rejects lexicon operations. The core reasoning underlying this alternative branch, which I will call the ‘concept nativist’s account’, is as follows:

- (32) a. All thematic forms are equally innate, there are no lexicon operations (this idea I motivated in section 1.4.1);
- b. Ergo, more concepts are innate than end up in the vocabulary, for instance causative alternates of unergative intransitives never have a label in the vocabulary;
- c. Ergo, the constraints on verb frame alternations must be due to constraints in the vocabulary;

⁸Freistaat Flaschenhals (‘Free State Bottleneck’) was a tiny free state the shape of a bottleneck, that existed in Germany from 1919 to 1923. Its existence, at the interface between the French and the American zones of occupation, was due to a simple measuring error. The incoming roads sealed off, the civilians of Freistaat Flaschenhals declared its independence and designed stamps and money, occasionally plundering a passing train.

- d. A possible explanation is that there are no lexicon operations, but *vocabulary operations*.

Crucially, because vocabulary operations can be acquired (after all, the entire vocabulary is acquired), my account of acquired lexicon operations can be made *concept-nativist-compatible* by replacing ‘lexicon’ with ‘vocabulary’ in the appropriate places. The resulting Theta System will be one without an active lexicon, but with an active vocabulary, which can account for all the relevant data. A slightly more complex reasoning will show that an acquired linking procedure, too, is compatible with concept nativism - but I do not wish to run ahead of things. In the introduction to each of the main chapters, I will discuss the relation between the non-nativist account (or ‘my account’), the concept nativist account and the mainstream Theta System.

I hope that advocates of the mainstream Theta System will see the merits of at least some of my proposals, perhaps maintaining concept nativism, but giving the ‘active vocabulary’ a serious chance. I also hope that opponents of the Theta System (most notably proponents of the syntax-driven approach) will see that an active lexicon is not ‘just another way of describing language’, but an undeniable, inevitable aspect of the human brain, combined with accumulating language change - at least in the two of three radically different ways of looking at it.

1.6 Relevance to my field of study

As an interdisciplinary programme, Cognitive Artificial Intelligence (CAI) at Utrecht University is the interface between cognitive science, linguistics, philosophy and computer science. It involves different methodologies, different background assumptions, different bodies of knowledge, and different kinds of people. That is often difficult, but also interesting and, sometimes, insightful. Likewise, my thesis combines two very different methodologies, with different bodies of knowledge, background assumptions and, perhaps, people. The Theta System is fundamentally rooted in Chomskyan generative linguistics, which is typically associated with nativism and the idea that language is primarily a means for thought, as opposed to communication. The framework based on which I give a new account of this Theta System is fundamentally anti-Chomskyan, in their focus on emergence in populations of speakers, as opposed to nativism. Interfacing these two methodologies has occasionally been very difficult, but (nearly) always interesting.

More down-to-Earth, this thesis is about language and, according to Chomskyan linguists, thought. Human-like language and thought are regarded as the ultimate goals of the field of artificial intelligence. In order to reach this goal, a deep understanding of human language and human thought is believed to be necessary. At least, this seems to be the consensus in the philosophy-based CAI (as opposed to computer-science-based AI at many other universities in The Netherlands and the rest of the world).

Throughout my BSc and MSc in CAI, language has been my main interest. In my BSc thesis I approach language from a cognitive science perspective. My MSc internship, at the Language Evolution and Computation Research Unit of the University of Edinburgh, revolved around the emergence of linguistic structure through language use. In this thesis I take a theory from generative linguistics and investigate, in a way, how applicable the Edinburgh methodology is.

Chapter 2

An acquired active lexicon

2.1 Introduction

The Theta System advocates an architecture of grammar that contains an *active lexicon*, a component of grammar that not only stores lexical concepts, but also applies operations to them prior to syntactic processing. For instance, the causative verb *break* can undergo a decausativisation operation, such that it enters syntactic processing as an inchoative verb. This notion of an active lexicon is not unique to the Theta System (e.g. Chierchia, 1989, 1995; Pesetsky, 1995; Pustejovsky, 1995; Levin & Rappaport-Hovav, 1995) but it is not universally accepted either. The active lexicon is rejected in particular in syntax-driven approaches, which seek to eliminate the role of the lexicon in grammar (e.g. Arad, 1998; Borer, 2005). In this chapter and the next I show why, in my account of the Theta system, the lexicon must be active.

Although I have chosen to tentatively abandon concept nativism, I will define the term ‘active lexicon’ in a way that should be meaningful to both concept nativists and non-concept-nativists. Siloni (2002) defines an active lexicon as follows:

An active lexicon is an inventory of coded concepts, a subset of which intrinsically take participants (bear theta-roles) and can undergo arity operations as specified by UG. (p.2)

I do not adopt this definition because the requirement that the set of arity operations is universal (i.e. ‘specified by Universal Grammar’) is unnecessary. After all, if two languages have completely different sets of arity operations, this would certainly not make the lexicon passive. I also wish to avoid the strong association Universal Grammar (especially when capitalised) has with innateness. Lexicon operations may be innate, as seems to be Siloni’s proposal, or acquired, as I will propose in this chapter, but as long as they exist the lexicon should be considered active. Instead, I propose the definition in (33).

(33) **Active lexicon:**

A lexicon is active if and only if the pre-syntactic availability of one or several thematic forms of a verb concept systematically leads to the pre-syntactic availability of new thematic forms for the same verb concept.

This definition relaxes a number of assumptions made in the Theta System about the active lexicon, in addition to Siloni's UG-requirement. First, it is neutral with respect to the first clause of the Lexicon Uniformity Hypothesis; new thematic forms may well be stored in the lexicon in addition to the base form. Second, it relaxes the assumption that some kind of operations apply immediately before a concept enters syntactic processing; they may as well apply during language acquisition. Third, the manner by which one thematic form is derived from another is unspecified. That is, as long as some thematic forms are derived from others, in whatever way, and as long as this happens systematically, i.e. according to a certain set of patterns, the lexicon is active. The somewhat tedious term 'pre-syntactic availability' should be general enough so as to not distinguish between the availability of labels (in the vocabulary) or thematic forms of concepts (in the lexicon). Although the definition I propose is weaker than the one advocated by Siloni, it is still a meaningful notion and its existence would directly oppose syntax-driven approaches.

2.1.1 Three views of an active lexicon

There are three views that are compatible with this notion of active lexicon, corresponding to the mainstream Theta System, the concept nativist's account, and my account.

According to the mainstream Theta System's Lexicon Uniformity Hypothesis, the lexicon (i.e. the innate system of concepts) contains only one thematic form of each verb concept, deriving the others via lexicon operations. This situation is depicted in figure 2.1. The acquired vocabulary entries, i.e. phonological labels, are attached to these concepts in a largely predictable manner (at least in English), the exception that proves the rule being *kill* and *die*.

The evidence the mainstream Theta System tries to account for, by postulating an innate active lexicon, is a set of universal regularities in the *vocabularies* of languages: for instance, cross-linguistically, a polyadic verb that assigns a [+c] role tends to be available also as a verb without the [+c] role, one entailing the other. An alternative explanation for such regularities is the concept nativist's view: an innate, passive lexicon with an active vocabulary, as depicted in figure 2.2.¹ According to this view, we are innately endowed with many thematic forms for each verb, including, for instance, a causative alternate of an unergative, intransitive verb. But only some of these thematic forms, like inchoative and causative forms, have a corresponding entry in the vocabulary and, hence, have been the subject of the linguists' inquiry. The language learner (like the linguist) acquires the regularities in the vocabulary and may productively apply

¹So, in fact not the lexicon but the vocabulary is active - but the lexicon and vocabulary combined match my definition of 'active lexicon'.

lexicon	vocabulary
BREAK([+c][-c-m]) <i>break</i>
BREAK([-c-m]) <i>break</i>
SPIN([+c][-c-m]) <i>spin</i>
SPIN([-c-m]) <i>spin</i>
DIE([+c][-c+m]) <i>kill</i>
DIE([-c+m]) <i>die</i>

Figure 2.1: An innate, active lexicon, as postulated in the mainstream Theta System.

lexicon	vocabulary
BREAK([+c][-c-m]) <i>break</i>
BREAK([-c-m]) <i>break</i>
BREAK([+c+m][-c-m]) X
...	
SPIN([+c][-c-m]) <i>spin</i>
SPIN([-c-m]) <i>spin</i>
SPIN([+c+m][-c-m]) X
...	
DIE([+c][-c+m]) <i>kill</i>
DIE([-c+m]) <i>die</i>
DIE([+c+m][-c+m]) X
...	

Figure 2.2: An innate, passive lexicon, in which many thematic forms are equally innate, linked to an active vocabulary. The constraints on verb frame alternations are captured in the vocabulary.

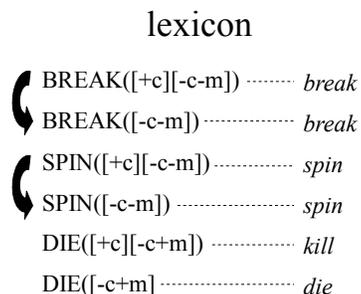


Figure 2.3: If concept nativism is tentatively rejected, lexicon and vocabulary coalesce into a single storage facility: an acquired, active lexicon.

them, i.e. predicting the inchoative form from a newly acquired causative verb. An advantage of this view is that a vocabulary must already be active anyway, in order to explain language-specific morphological regularities. And perhaps this view is more compatible with concept nativism than the mainstream view, as I have suggested in chapter 1. But only the next forty-or-so pages will tell us whether the concept nativist's view of an active lexicon has any merit beyond that.

My non-nativist account of an active lexicon, which coalesces lexicon and vocabulary, is that of an *acquired active lexicon*, depicted in figure 2.3. This will be, primarily, the kind of active lexicon that I will argue in favour of in this chapter. However, as announced in chapter 1, the lion's share of this thesis can be translated fairly directly into the concept nativist's view, by substituting 'vocabulary' for 'lexicon' in strategic places. Occasionally throughout this thesis I will guide this mental exercise on the reader's part with short in-line remarks or footnotes.

In order to support the acquired active lexicon view (and the nativist's acquired active vocabulary view), in section 2.2 I consider briefly what I believe are advantages of an acquisition view in general. In section 2.3 I argue that the acquisition of an active lexicon in the sense of (33) is inevitable if a language exhibits verb frame alternations. In other words, an active lexicon in I-language is inevitable whenever the E-language contains verb frame alternations. In section 2.4 I turn this language acquisition perspective into a strong case against syntax-driven approaches. I give a brief conclusion in section 2.5.

2.2 The origins of lexicon operations

A convincing account of the origins of lexicon operations would strongly support the hypothesis that the lexicon is active. There are two possibilities: either lexicon operations are innate, or lexicon operations are acquired. The innateness view of lexicon operations - like any innateness view - raises the following

question. What *caused* lexicon operations to be innate? (Chater, Reali, & Christiansen, 2009) argue that the only linguistic properties that could, in principle, end up in the human genome are those properties that remain constant over long periods of time. Such properties may be the asymmetry of merger or the fact that linguistic signs *denote*, and perhaps the decausativisation operation. But an operation like lexical reflexivisation does not exist in all languages (this is captured by the Lexicon-Syntax Parameter), so how could it have ended up in our genome? An alternative explanation for the innateness of lexicon operations is that they reflect deeper, language-independent cognitive mechanisms for conceptualising the world. Perhaps a decausativisation operation happens to come naturally with any neural processor. But this is so speculative that I do not believe it could yield a strong argument in favor of the innateness view. (Note that under discussion at this point is the innateness of lexicon operations, under the assumption that concepts are not innate).

The acquisition view of lexicon operations, on the other hand, seems to raise the following chicken-and-egg problem: do speakers acquire lexicon operations because languages contain certain verb frame alternations, or do languages contain such alternations because speakers have acquired lexicon operations? A possible way out of this apparent circularity may be found in the domain of gradually accumulating language change. For instance, there may be some process of language change through which every transitive verb with a [+c] role gradually comes to have an inchoative alternate, or vice versa, similar to the asymmetrical processes underlying grammaticalisation (Heine & Kuteva, 2002) and phonological change (e.g. Kirchner, 1996, from an Optimality Theory perspective). Alternatively, studies in the field of language evolution show that universal linguistic regularities can emerge due to the language acquisition bottleneck, through a process called *iterated learning* (Smith et al., 2003; Kirby, Cornish, & Smith, 2008).² In any case I am more optimistic about solving the chicken-and-egg problem faced by the acquisition view, than about solving the origins problem faced by the innateness view.

Although the language acquisition data available at present are not decisive in deciding between a lexicon-driven approach and a syntax-driven approach (Sorace, 2004), they do suggest that if grammar is lexicon-driven, then the operations are most probably acquired. Pinker (1989) brings in a lot of evidence that suggests that infants go through an initial stage in which no lexicon operation is applied (i.e. the infant uses only those thematic forms of verbs which (s)he has acquired from direct evidence), followed by a stage of overapplication of lexicon operations (i.e. overgeneralisation). This overgeneralisation initially happens freely in all directions, though in some directions more likely than in others (Sorace, 2004). That is, infants both transitive intransitive verbs and detransitivise transitive verbs, leading to such errors as (34).³

²For the concept nativist, patterns in the *vocabulary* may be the result of grammaticalisation or iterated learning, regardless of whether the underlying concepts are innate or acquired.

³For the concept nativist's account, these data are not decisive for an active lexicon. But they do, I believe, suggest an active vocabulary, which is what matters.

- (34) Come back on the light. [= make it come back on]
 Mommy, can you stay this open? [= make it stay open]
 I'm gonna just fall this on her. [= make it fall]
 And the doggie had a head. And somebody fell it off. [= make it fall]
 No, Mommy, don't eat her yet, she's smelly! [= feed]
 He's gonna die you, David. [= kill]
 ...
 (Bowerman, 1982)

I will therefore assume that lexicon operations are acquired from linguistic data. In the next section I address how lexicon operations are acquired. In the next *chapter* I provide an account of the origins and universality of verb frame alternations, i.e. a way out of the chicken-and-egg problem.

2.3 The acquisition of lexicon operations

What is necessary for the acquisition of lexicon operations? In order to be able to discover patterns in the linguistic data, the lexicon must contain all thematic forms of a verb as separate lexical entries. For instance, the decausativisation operation could not be discovered unless the causative and inchoative form of each verb are both stored as lexical entries (an assumption that I motivate in section 1.4).⁴ Consequently, under the acquisition view, lexicon operations are not mechanisms that apply every time a verb concept is used right before insertion in the syntax. Rather, lexicon operations are generalisations that help a language learner to fill in the gaps of incompletely acquired lexical entries. For instance, after acquiring the decausativisation operation and after encountering the novel verb *twuck* in sentence (35a) and (35b), the learner will infer that there must also be a lexical entry for inchoative *twuck*, as used in (35c), and the learner will store this derived verb in his/her lexicon.

- (35) a. The wind twucked the beehive
 b. Alan twucked the beehive
 c. The beehive twucked

Pinker (1989) calls this process 'productive generalisation' and the patterns that govern it 'lexical rules'. However, because my goal is to explain the entities postulated in the Theta System I will keep using the term 'lexicon operation', even though it seems to denote something more powerful than required. Nevertheless, a lexicon with lexicon operations in the sense of productive generalisation is still an active lexicon, according to my definition of the term given in (33).⁵

⁴This condition on the acquisition of lexicon operations must be restated in concept nativist terms: the language learner learns a vocabulary operation by looking only at the lexical entries that have a corresponding label in the vocabulary. This makes sense: the language learner generalises only over the lexical entries that are available in the vocabulary.

⁵And so is the combination of an innate, passive lexicon with an active vocabulary, with acquired vocabulary operations, in the concept nativist's parallel representation of the discourse.

Productive generalisation may come for free in neural networks. In *artificial neural networks*, the machine learning technique that most closely resembles the human brain, the distinction between storage and generalisation is blurred (Plate, 2002). An artificial neural network cannot store data without generalising; the act of storing itself introduces generalisations into the network. Hence, if a language contains verb frame alternations, it is impossible for an artificial neural network to store the language's lexical entries *without* also acquiring lexicon operations. Generalisation is a non-linguistic, domain-general mechanism; even if the lexicon (or vocabulary) is a language-specific module, the least we can expect it to do is generalise. If a language exhibits verb frame alternations, an active lexicon may be inevitable.

Although the acquisition of lexicon operations is straightforward *in principle*, this does not mean that acquiring the particular set of lexicon operations that correspond to verb frame alternations *in natural language* is a trivial task. The verb frame alternations found in natural language are typically constrained in a variety of ways involving often subtle semantic or phonological distinctions (Pinker, 1989):

- (36) a. John gave a dish to Sam. (also pass, tell, etc.)
 b. John gave Sam a dish.
 c. John donated a painting to the museum.
 d. *John donated the museum a painting.
- (37) a. John touched Fred.
 b. Fred was touched by John. (also hit, see, like, kick, etc.)
 c. John resembled Fred.
 d. *Fred was resembled by John.
- (38) a. Irv loaded eggs into the basket.
 b. Irv loaded the basket with eggs. (also spray, cram, splash, stuff, etc.)
 c. Irv poured water into the glass.
 d. *Irv poured the glass with water.

(Pinker, 1989, p.8)

The problem for the language learner is to find out that some verbs do not behave according to an otherwise productive generalisation, for no apparent reason and without (systematic) negative evidence. The problem for the linguist, in turn, is to explain why such constrained generalisations are nevertheless successfully acquired, generation after generation. Pinker calls this problem 'Baker's paradox', after (C. L. Baker, 1979). Pinker's solution, to his interpretation of Baker's paradox at least, is to show that many constraints are not arbitrary at all, and that in principle productive generalisation could pick up the relevant classes of lexical entries.⁶ I will put my fate in the hands of 'the linguists' to

⁶The same problem would arise, and the same solution would be available, for the concept nativist. After all, according to the active vocabulary view, the learner must acquire generali-

find out whether this pursuit is tenable, and assume in the remainder of this thesis that lexicon operations can be acquired (in other words, and somewhat paradoxically, I will be assuming in the remainder of this thesis that a child can discover in a matter of years what linguists achieve in, perhaps, decades).

2.4 Acquisition in a syntax-driven approach

The most serious opponents of an active lexicon are the syntax-driven approaches, mentioned only briefly in chapter 1. In syntax-driven approaches, recall, there exists a set of syntactic event structure templates into which category-neutral roots, so-called *listemes*, may be inserted. The following sentences, which are all grammatical (though not always semantically coherent), illustrate the core idea fairly well.

- (39) a. The dog boated three sinks.
 b. The three sinks boated some dogs.
 c. The sinks dogged the boat.
 d. The boats dogged the sink.
 e. The three dogs sank the boat.
 f. The boat sank the dogs.

(Borer, 2005, p.13)

Syntactic templates serve to explain the variable behaviour of words, including verb frame alternations, giving them roughly the same function as lexicon operations in a lexicon-driven approach. But while lexicon operations, being semantic operations, are inherently semantically constrained, syntactic templates lack such inherent constraints. Hence, overgeneralisation and Baker's paradox are an even bigger problem for syntax-driven approaches.

In this section I discuss in particular Borer's (2005) solution to overgeneralisation, but I believe that any solution to overgeneralisation in syntax-driven approaches will contain roughly the same ingredients. Borer invokes *idioms*, which are listemes enriched with additional syntactic, semantic or phonological information that constrain their usage. The following are typical examples of idioms.

- (40) a. trousers (cf. shirts)
 b. cross (that) bridge when (we) come to (it) (cf. cross that bridge when we come to it, literally)
 c. depend on (cf. sit on)
 d. kick the bucket (cf. kick the bucket, literally)

(Borer, 2005, p.25)

sations over vocabulary labels which can be conditioned on the denoted concepts, no different from generalisations over lexical entries in Pinker's view and mine.

But ‘idiom’ in Borer’s approach has a broader meaning. The inventory of idioms contains regularities that reflect constraints on verb frame alternations. For instance, all achievement verbs in English are idioms that obligatorily take an object (Borer, 2004), all verbs with a polysyllabic root in English are idioms that cannot be used in the dative construction (Oehrle, 1976), and nearly all verbs in Hebrew are idioms (Borer, 2005). Obviously, then, Borer’s solution to overgeneralisation does not solve Baker’s paradox - rather, it gives rise to it. But the postulation of idioms has an additional consequence, as I show in the remainder of this section: the mechanism required for the acquisition of listemes and idioms is identical to the mechanism required for the acquisition of an active lexicon. Therefore, the acquisition of an active lexicon is inevitable *even in syntax-driven approaches*. In light of this paradoxical result, I will reject syntax-driven approaches.

Suppose that the acquisition of listemes and idioms is a two-stage process, as depicted in (41).

(41) E-language \longrightarrow proto-storage \longrightarrow storage of listemes/idioms

In the first stage, the linguistic data is preprocessed and stored in a hypothetical *proto-storage*, in such a way that listemes and idioms can be reliably acquired and stored. I invoke this hypothetical entity because I believe it facilitates a theory-neutral analysis of the acquisition of listemes and idioms.

What should the proto-storage be like in order for the acquisition of listemes and idioms to be possible? For the acquisition of listemes it would be sufficient to store only the root meaning in the proto-storage. However, since any thematic form of a verb may turn out to be an idiom - the language learner cannot possibly foresee such idiosyncracies - not only its root meaning but every thematic form it is witnessed in must be stored in the proto-storage. It is only after storing several verb frames that the language learner can guess with some certainty whether a verb occurs in all or in only a few verb frames, i.e. whether it is a listeme or an idiom. Even if the learner’s default guess had been to store verbs as listemes - despite some languages relying almost completely on idioms, such as Hebrew - there has to be a backup mechanism for verifying and adjusting this choice if necessary. If the learner stores only the root meaning of each verb, (s)he will never be able to even check whether a root meaning has ever occurred in a particular verb frame. Hence, the proto-storage must not *a priori* ignore some of the data, because any semantic, syntactic or phonological feature could turn out to be relevant at a later stage.

In addition to this, because the learner cannot possibly acquire all features of all listemes and idioms from direct evidence (Chomsky’s (1965) ‘poverty of the stimulus’), (s)he must generalise over the data in the proto-storage. This is how the learner may discover, for instance, the idiosyncratic fact that English causative verbs with a polysyllabic root happen to be idioms that disallow dativisation, or the fact that all achievement verbs in English are idioms with an obligatory direct object. Experiments on the acquisition of novel verbs (e.g. Gleitman, 1990) confirm that these generalisations exist and remain active throughout a speaker’s lifetime (and why wouldn’t they?), suggesting that the

proto-storage is not used for some initial stage in acquisition and then discarded, but maintained, shadowing the storage of listemes and idioms throughout a speaker's life.

In sum, the proto-storage must store each thematic form of a verb separately, including semantic, phonological and syntactic information of its usage, and the proto-storage must generalise over these data and productively apply relevant patterns. It turns out that *the proto-storage is no different from an active lexicon*. That is, if a language exhibits verb frame alternations, the proto-storage will contain lexicon operations. Hence, the label 'proto-storage' in (41) can be replaced by 'active lexicon' as in (42).⁷

(42) E-language \longrightarrow active lexicon \longrightarrow storage of listemes/idioms

In light of (42), the proposal of syntax-driven approaches seems very strange indeed. The proposal requires that the language learner first constructs an active lexicon, and then constructs from it an impoverished storage of listemes and idioms that is stripped of most thematic forms and that ignores lexicon operations, postulating instead a set of syntactic templates the origins of which remain unexplained. If this is not a violation of Ockham's razor, nothing is. A lexicon-driven approach like the Theta System, on the other hand, can dispense with the storage of listemes/idioms and the syntactic templates (43).

(43) E-language \longrightarrow active lexicon

The lexicon-driven approach looks reassuringly elegant, and this result retroactively motivates the focus of my thesis: the Theta System.

2.5 Conclusion

In my account of the Theta System, an acquired active lexicon in I-language is inevitable if E-language contains verb frame alternations. This requires no additional assumptions, as the acquisition of lexicon operations *in principle* is very straightforward (productive generalisation). The acquisition of the constrained lexicon operations required for the verb frame alternations in natural language is slightly more difficult (Baker's paradox), but it is no more problematic for lexicon-driven approaches than the acquisition of idioms is for syntax-driven approaches. Indeed, syntax-driven approaches implicitly require for the acquisition of listemes and idioms precisely the thing they most strongly oppose: an active lexicon. I conclude that a lexicon-driven approach, with an acquired active lexicon, is a more parsimonious approach. Furthermore, in the introduction and by means of short remarks throughout this chapter, I have tried to show that an acquired, active vocabulary is a possible substitute for innate lexicon operations. A pressing question remains, both for my account of the Theta System and for the concept nativist's parallel account: what are the origins of verb frame alternations? This question is dealt with in the next chapter.

⁷Indeed, this seems to be an active lexicon proper, rather than merely an active vocabulary, even for the concept nativist, because its entries contain syntactic and semantic information.

Chapter 3

The origins of verb frame alternations

3.1 Introduction

In chapter 2 I proposed that lexicon operations are acquired, and that such acquisition is inevitable if there are regular verb frame alternations. This finding leads to the chicken-and-egg problem mentioned earlier: do speakers acquire lexicon operations because their languages contain regular verb-frame alternations, or do languages contain such verb-frame alternations because speakers have acquired lexicon operations? A possible way out of this circularity is to assume that verb frame alternations are the result of gradually accumulating language change. In this chapter I work out this possibility in considerable detail. Because language does not fossilise, solidifying how exactly verb frame alternations have emerged is difficult or impossible.¹ Nevertheless, I can show that even in a worst-case scenario, without invoking any special-purpose mechanisms, verb frame alternations inevitably emerge.

In this chapter I will not include as many footnotes for the concept nativist as in the previous chapter, for, having explained the basics, the substitutions required for the concept nativist (e.g. ‘vocabulary operation’ for ‘lexicon operation’) are straightforward. In assuming that all thematic forms of a verb are equally innate and that not the lexicon but the vocabulary contains operations, the concept nativist faces the exact same chicken-and-egg problem.

In section 3.2 I introduce the core mechanism responsible for the emergence of verb frame alternations: *iterated learning*. In section 3.3 I detail why verb frame alternations, most notably the causative-inchoative alternation, are universal; for instance, why does every causative verb have an intransitive alternate, but not vice versa? In section 3.4 I explain why some verb frame alternations are encoded in the lexicon in some languages but in the syntax in others, as

¹Not that it would have been easy if language had fossilized, I am aware after reading the chapter on paleontology in Bill Bryson’s ‘A Short History of Nearly Everything’.

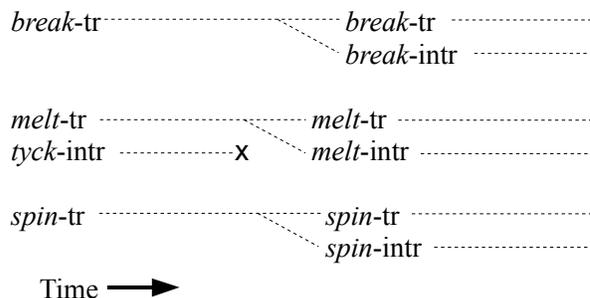


Figure 3.1: An intuitive explanation of the origins of verb frame alternations: every causative verb gradually comes to have an inchoative alternate, possibly replacing existing verbs like *tyck* (fictitious) in the process.

captured by the Lexicon-Syntax Parameter. In section 3.5 I reconcile Horvath and Siloni’s (2005) idea of ‘frozen forms’ - words that are ‘in the lexicon’ but not ‘in the vocabulary of a particular language’ - with my account of the Theta System. Section 3.6 contains a conclusion.

But first, in the short section that follows, I discuss and reject what I believe would be a more intuitive line of inquiry. This serves two goals. First, I believe that contrasting the two explanations is insightful. Second, I fear that if I do not prove this intuition wrong right away, it will block a full appreciation of the proper explanation.

3.1.1 An intuitive (but wrong) explanation

An intuitive explanation of the origins of verb frame alternations would be that there is some process of language change through which causative verbs have gradually gained an additional usage as inchoative verbs, as depicted in figure 3.1. What makes this explanation intuitive, I believe, is that the asymmetry of the hypothesised process of language change would directly parallel the asymmetry of lexicon operations. That is, we have a lexicon operation that derives inchoatives from causatives, rather than the other way around, *because* there’s a mechanism of language change that derives inchoatives from causatives. A major challenge would be the fact that, to my awareness, no process is known whereby causative verbs come to have an inchoative alternate.² But in addition this explanation is faced by the following empirical problem.

Processes of language change may involve morphological markers. If morphological markers are the result of grammaticalisation, as is commonly assumed (for an overview, see Heine & Kuteva, 2002), a concept that ends up with a

²Of course, in today’s languages, there is such a process: because the speakers of a language possess a decausativisation operation, any causative verb will instantaneously become available as an inchoative verb. But this process is only available *after* the verb frame alternation has already emerged, so it cannot be used in support of the intuitive explanation.

morphological marker is typically a concept that could initially be conveyed only through syntactic composition, in lack of a word denoting it. In this case, if the intuitive explanation presently under discussion is correct, one would expect the inchoative to be morphologically marked, because the inchoative was initially unavailable and had to be conveyed periphrastically. But this prediction is not borne out for the causative-inchoative alternation. In some languages the causative form is marked (44), while in other languages the inchoative form is marked (45) (all examples taken from Marelj's (2010) course notes):

- (44) Marking on the causative:
- a. Georgian (Haspelmath, 1993):
duɣ-s ('cook-intr') / *a-duɣ-eps* ('cook-tr')
 - b. Khalka Mongolian (Piñon, 2001):
ongoj-x ('open-intr') / *ongoi-lg-ox* ('open-tr')
 - c. Japanese (Hasegawa, 2001):
kawak-u ('dry-intr') / *kawak-as-u* ('dry-tr')
ugok-u ('move-intr') / *ugok-as-u* ('move-tr')
tob-u ('fly-intr') / *tob-as-u* ('fly-tr')
- (45) Marking on the inchoative:
- a. Russian (Haspelmath, 1993):
katat'-sja ('roll-intr') / *katat'* ('roll-tr')
 - b. Polish (Piñon, 2001):
złamać-się ('break-intr') / *złamać* ('break-tr')
 - c. Serbo-Croat:
otvoriti se ('open-intr') / *otvoriti* ('open-tr')
 - d. French:
se casser ('break-intr') / *casser* ('break-tr')

These data are problematic for any explanation based on an asymmetrical process of language change, including the intuitive explanation presently under discussion. While in the mainstream Theta System the conclusion was drawn that morphological marking cannot tell you anything about the direction of derivation (a main argument against Distributed Morphology), I conclude that apparently, *there is no direction of derivation* as far as the causative-inchoative alternation is concerned. What is required instead is a process that yields two members of a verb frame alternation in no particular order, allowing the morphological marker to arise on either side. Such a process is discussed next.

3.2 Iterated learning

Iterated learning is schematically depicted in figure 3.2. Each generation constructs I-languages based on the E-languages generated from the I-languages of the previous generation. Because the acquisition of an I-language is subject to an acquisition bottleneck, each learner will have to fill in the gaps in the

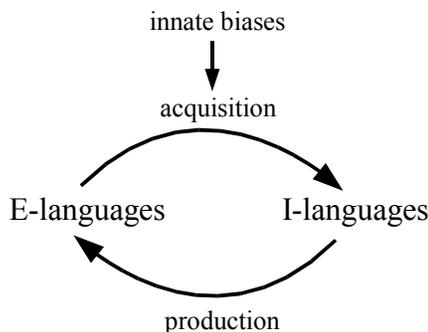


Figure 3.2: The iterated learning cycle. Roughly based on several illustrations in (Kirby and Hurford, 2001; Smith, Kirby and Brighton, 2003).

acquired data according to his/her *inductive biases*. Because generation after generation uses the same inductive biases to fill in the gaps, languages will eventually come to reflect such biases (Griffiths & Kalish, 2007; Rafferty, Griffiths, & Klein, 2009).

Universal Grammar, in generative linguistics, is an inductive bias: it allows learners to acquire a language despite Chomsky’s (1965) poverty of the stimulus. Universal Grammar is a very rigid kind of inductive bias; it strictly separates the finite set of possible grammars from the infinite set of impossible grammars, thereby solving the logical problem of language learning (cf. chapter 1). As a result, all languages on Earth reflect Universal Grammar.

But weaker, negotiable inductive biases may also lead to language universals. Consider, for instance, the learner’s expectation that similar signals have similar meanings and different signals have different meanings (the second conjunct corresponds to Pinker’s (1984) ‘Unique Entry Principle’); I will call this the *regularity bias*. The exaggeration of a weak bias through iterated learning has been attested in many computer simulations (e.g. Kirby & Hurford, 2002; Smith, 2003). The regularity bias drives generalisation; indeed, the same process through which lexicon operations can be acquired. Computer simulations confirm what one would expect by common sense: if generation after generation fills in the gaps of an acquired lexicon/vocabulary by means of generalisation, the signal-meaning mapping will become more and more regular.

As an example, consider briefly Kirby et al.’s (2008) iterated learning experiment with human participants.³ In the experiment, human participants are asked to learn and then reproduce the names of visual stimuli of coloured, moving shapes. Learning is subject to a bottleneck (each learner observes only part of the lexicon) and the output of one learner was given as input to the

³Because the experiment involves human participants, the underlying mechanics are not as transparent as those in computer simulations. However, the results are virtually identical, hence I can safely use them to illustrate what it means for a lexicon to be ‘regularised’ through iterated learning.

----->	n-ere-ki	l-ere-ki	renana	□
	n-ehe-ki	l-aho-ki	r-ene-ki	○
	n-eke-ki	l-ake-ki	r-ahē-ki	△
~~~~~>	n-ere-plo	l-ane-plo	r-e-plo	□
	n-eho-plo	l-aho-plo	r-eho-plo	○
	n-eki-plo	l-aki-plo	r-aho-plo	△
⦿	n-e-pilu	l-ane-pilu	r-e-pilu	□
	n-eho-pilu	l-aho-pilu	r-eho-pilu	○
	n-eki-pilu	l-aki-pilu	r-aho-pilu	△

Figure 3.3: A signal-meaning mapping regularised through ten generations of iterated learning. Dashes are included for clarity, they were not there in the experiment. Adopted from (Kirby, Cornish and Smith 2008).

next learner, i.e. the language underwent iterated learning. Even though the first learner was presented with a lexicon in which the stimuli’s names are completely arbitrary, after only ten generations the signal-meaning mapping exhibited a great degree of regularity, as witnessed in figure 3.3.

In this chapter I propose that verb frame alternations are nothing but regularities of the form-meaning mapping due to iterated learning. I will not discuss here whether Universal Grammar can be reduced to a set of weak biases like the regularity bias altogether, as suggested by Zuidema’s (2002) ‘how the poverty of the stimulus solves the poverty of the stimulus’ (see Lidz, 2010 for some possible counterevidence). My aim is only to show that verb frame alternations, but one aspect of grammar, can be explained as the result of the regularity bias through iterated learning. This is my aim because I wish to support my account of the Theta System, and the regularity bias is much more parsimonious than a set of innate lexicon operations (this holds regardless of whether one assumes concept nativism). It is especially parsimonious because the bias can be independently motivated: it is not only necessary for the acquisition of morpho-phonological regularities and, I have proposed, for the acquisition of lexicon operations; it is also a very useful bias *in general* in a world that is quite predictable. Similar trees yield similar fruits (different trees yield different fruits), similar actions lead to similar reactions, and similar events have similar causes. Being the bias that drives generalisation, it may well be a fundamental property of neural processing.

The emergence of verb frame alternations through iterated learning is remarkably straightforward. Suppose that in the lexicon of a language all thematic forms of a verb have completely arbitrary, unrelated phonological labels - similar to the initial lexicon in Kirby et al.’s experiment. Through iterated learning and the regularity bias, phonologically unrelated lexical entries that share a large meaning component, such as causative and inchoative forms of the same core meaning, will come to share a large component of their signal. As a result, in most languages causative and inchoative verbs share a phonological stem; in English causative and inchoative forms are even completely indistinguishable

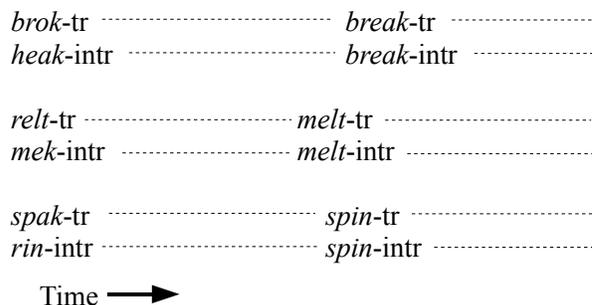


Figure 3.4: A schematical depiction of verb frame alternations due to gradually accumulating language change. From initially unrelated causative and inchoative verbs, the pairs with similar core meanings converge phonologically due to iterated learning. The initial verb forms are fictitious.

(except by the number of arguments they take). (This view of emerging verb frame alternations is depicted very schematically in figure 3.4, though mainly for comparison with the intuitive explanation I rejected earlier (figure 3.1).)

How does iterated learning explain the morphological markers in (44) and (45), which proved problematic for the intuitive explanation? As I announced, the causative-inchoative morphology is not consistent cross-linguistically because there is no ‘derivation’ going on - causatives and inchoatives are equally basic (or equally derived).⁴ Morphological marking may emerge on either the causative or the inchoative, via the following process. Differences in thematic form, such as the causative-inchoative distinction, are semantic differences like any other. As a result, the mapping between signal and causativity/inchoativity may become regular in precisely the same way in which the mapping between signal and shape/colour/motion became regular in Kirby et al.’s experiment. Iterated learning is neutral with respect to which form is marked. If some salient causatives happen to share a syllable, this syllable may spread over to the other causatives and become a causativity marker as in (44). If some salient inchoatives happen to share a syllable, this syllable may spread over to the other inchoatives and become an inchoativity marker as in (45). In sum, the data that posed a problem for the intuitive explanation are not problematic at all for an iterated learning account.

Based on this explanation of morphological marking through iterated learning, one would predict that other verb frame alternations show a similar level of cross-linguistic variation in marking. However, this prediction does not seem to

⁴I mentioned earlier that this is different from the conclusion drawn from the data in the mainstream Theta System, i.e. that morphological marking does not tell you anything about the direction of derivation. Instead, I maintain that morphological marking may tell you something about the direction of derivation - indeed, it is quite likely given how grammaticalisation works - but only if there is in fact derivation going on, which is simply not the case here. I propose that cross-linguistic variation in morphological marking can be used as a diagnostic for an origin in iterated learning.

be borne out for the parameterisable operations, i.e. operations that can take place ‘in the syntax’ such as reflexivisation. The answer to this question involves the Lexicon-Syntax Parameter, hence I will postpone it to section 3.4.

### 3.3 Universality

Why are verb frame alternations universal? If verb frame alternations derive from pairs of unrelated verbs that only accidentally share their root meaning, their universality must be explained by means of the universality of such concepts.⁵ Universally, it is useful to have a word for WASH-SELF, while a word for SELF-SEEING would not be useful. Universally, it is useful to have a word for FIGHT-EACHOTHER, while a word for BREAK-EACHOTHER would not be useful. Cross-linguistically the most typical reflexives and reciprocals denote concepts that seem quite useful as reflexives. Typical reflexives are ‘grooming’ verbs such as *wash*, *shave*, *equip* and *arm* (Reinhart & Siloni, 2005). Typical reciprocals are verbs of competition (e.g. *fight*), joint action (e.g. *communicate*), connecting (e.g. *combine*) and dividing (e.g. *separate*) (Haspelmath, 2007). But what does ‘useful’ mean?

A lexical entry is useful for expressing the concept it conveys, if it is easier to retrieve the lexical entry from the lexicon and express it than to express the same concept periphrastically, i.e. compositionally in the syntax. For example, even though DISARM-SELF is used less frequently than SEE-SELF, DISARM-SELF may still exist as a lexical entry (like Dutch *ontwapenen*) while SEE-SELF does not, because relatively, compared to their non-reflexive alternates, DISARM-SELF may be more frequent than SEE-SELF. It is of course an empirical issue to see whether this sketchy account amounts to anything. An analysis of the relative frequencies of causatives, lexical reflexives and periphrastic reflexives in a corpus could support or disprove this, and perhaps reveal an adequate usefulness measure. One may even be able to estimate from the frequencies of lexical reflexives and periphrastic reflexives the computational cost of an additional thematic role assignment. I believe that something along these lines may be the case for reflexives and reciprocals, and a similar explanation may be available for agentive verbs (‘walk the dog’). But I will leave an empirical validation of this idea for future research, and instead support the usefulness-based explanation of verb frame alternations by extending it to the causative-inchoative alternation.

Contrary to the relative usefulness of reflexives, reciprocals and agentive verbs, which is really a matter of world structure, the usefulness of inchoative verbs compared to their causative alternates can be captured in a well-defined principle, as I will show. It has proven to be difficult to discover the *directionality* of the underlying operation: is the causative derived from the inchoative (Dowty, 1979; Pesetsky, 1995) or is the inchoative derived from the causative (Chierchia, 1989; Levin & Rappaport-Hovav, 1995; Reinhart, 2000)? In the Theta System, the latter is assumed, enabling it to explain, among other things, the unaccusative behaviour of inchoatives. In the short remainder of this section

⁵Or, for the nativist, by means of the universality of vocabulary items for such concepts.

I employ Rosch's (1978) take on concept usefulness to explain the origins of the causative-inchoative alternation and the directionality of the resulting lexicon operation.

Rosch shows that all basic-level concepts (such as CHAIR, as opposed to the subordinate ARMCHAIR or superordinate FURNITURE) are concepts that optimise the trade-off between within-category similarity (e.g. the similarity between rocking chairs, office chairs, lazy chairs and stools) and between-category similarity (e.g. the similarity between chairs, trees, houses and dogs). For that reason, basic-level concepts are more useful in everyday contexts (e.g. not a furniture shop) than more specific or more general concepts. Consequently, the following principle holds:

(46) **Basic-level-preference principle:**

If a language contains a lexical entry denoting a subordinate or superordinate concept, it will also contain a lexical entry denoting the corresponding basic-level concept.

Here 'preference' does not equal temporal precedence; it denotes only a synchronic, logical relationship between two kinds of lexical entries.

In a similar fashion, I propose that the pattern in (47) is responsible for the universality of the causative-inchoative alternation:

(47) **Inchoative-preference principle:**

If a language contains a verb denoting a concept  $X([+c],\theta,\dots)$ , it also contains a verb denoting  $X(\theta,\dots)$ .

Why should this be so? A concept  $X([+c],\theta,\dots)$  always entails a corresponding concept  $X(\theta,\dots)$ , but not vice versa. Hence,  $X(\theta,\dots)$  has a smaller reference and, therefore, a higher within-concept similarity. Intuitively, an additional [+c] role adds *more variation* to the concept. In addition, a concept  $X(\theta,\dots)$  has a lower between-concept similarity: an arbitrary instance of BREAK([-c-m]) and an arbitrary instance of BEND([-c-m]) are, on average, less similar than arbitrary instances of their causative alternates. This is so because the [+c] roles of BREAK and BEND are identical, while the [-c-m] roles of, BREAK and BEND, may differ in many s-selectional restrictions, for instance in their rigidity. In other words, the addition of a [+c] role to a concept will lower the concept's within-concept similarity and increase the concept's between-concept similarity, and hence make the concept less useful. Inchoatives are to causatives what basic level concepts are to superordinate concepts: more useful. Hence, inchoative-preference is to be expected just like basic-level-preference.

The inchoative-preference principle supports from a completely new angle the Theta System's assumption that there is a decausativisation operation, rather than a causativisation operation, and, for the first time, *explains* why this should be so.

### 3.4 The Lexicon-Syntax Parameter

Some operations, recall from chapter 1, may take place either in the lexicon or in the syntax, depending on the language, as captured by the Theta System's Lexicon-Syntax Parameter (48).

- (48) **The Lexicon-Syntax Parameter:**  
 UG allows thematic arity operations to apply in the lexicon or in the syntax
- (49) Some languages (Reinhart & Siloni, 2003, p.19):
- a. Lexicon setting: Hebrew, Dutch, English, Russian, Hungarian
  - b. Syntax setting: Romance, German, Serbo-Croatian, Czech, Greek

Recall that the usefulness of a lexical entry depends on the availability of an alternative (e.g. periphrastic or otherwise syntactic) way to express its meaning. If a language has an efficient syntactic way for expressing e.g. reflexives, it is redundant to in addition have lexical reflexives. Romance languages such as French and Italian have a *se/si* clitic which can 'absorb' the remaining thematic role of a transitive verb, yielding a reflexive reading (in (Reinhart & Siloni, 2005) this is formulated in terms of checking thematic accusative case). Because Romance languages have this clitic, a reflexive meaning can be expressed efficiently in the syntax. As a result, lexical reflexives become redundant and will not survive the acquisition bottleneck.⁶

Cross-linguistically, two kinds of reflexive pronoun can be distinguished: SELF-anaphora (*himself*, Dutch *zichzelf*) and SE-anaphora (Dutch *zich*) (Reinhart & Reuland, 1993). In Dutch, a lexicon language, *zich* marks lexical reflexives. But, in addition, *zichzelf* can be used for expressing reflexives syntactically:

- (50) a. Hij waste zich.  
 He washed SE  
 'He washed'
- b. Hij waste zichzelf.  
 He washed himself

Dutch seems to be a language that has lexical reflexives *despite* the availability of a productive, syntactic way of expressing reflexives. Is that already a counterexample, even despite our vague definition of 'efficient'? I believe not, for syntax languages, too, may have both a SE-anaphor and a SELF-anaphor. Consider German, for instance:

- (51) a. Er hat sich rasiert.  
 He has SE shaved 'He shaved'
- b. Er hat sichselbst rasiert.  
 He has SELF shaved 'He shaved himself'

⁶Of course, the term 'efficiency', like 'usefulness', must be defined. I assume the working hypothesis that a suitable definition exists, though we have yet to find it in future research.

Apparently, independently of the Lexicon-Syntax Parameter, it is useful to have both kinds of anaphor. Geurts (2004) suggests that this is due to a difference in their ability to bear stress (a difference that is more pronounced in Dutch than in German):

- (52) Unstressed:
- a. Hij zag zich/zichzelf in de spiegel.
  - b. Er sah sich/sichselbst im Spiegel.  
he saw SE/SELF in the mirror
- (53) Stressed:
- a. In de spiegel zag hij *zich/zichzelf.
  - b. Im Spiegel sah er ?sich/sichselbst.  
In the mirror saw he SE/SELF

By contrast, the SE-anaphor used in Romance (syntax) languages to express reflexivity is an inflectional clitic, rather than a nominal like *sich*, and hence it cannot bear stress at all - clitics have to *lean on* (Greek *klinein*) another word (Gerlach, 2002). Hence, a similar apparent redundancy does not exist in Romance.

Reinhart and Siloni (2005) argue that all parameterisable operations share a single Lexicon-Syntax Parameter. That is, either they all take place in the lexicon, or they all take place in the syntax. This is to be expected if the same syntactic tool can be used for all operations. Indeed, Reinhart and Siloni show that the same clitic appears with reciprocals, unaccusatives, subject-Experiencer verbs, middles, impersonals and passives (illustrated in this order for Italian):

- (54)
- a. Giovanni e Maria si sono abbracciati.  
Giovanni and Maria SI are hugged  
'Giovanni and Maria hugged each other'
  - b. La porta si chiusa.  
the door SI is closed  
'The door closed'
  - c. Giovanni si preoccupa di questo.  
Giovanni SI worries of this  
'Giovanni worries about this'
  - d. Questi vestiti si lavano facilmente.  
These suits SI wash easily  
'These suits wash easily'
  - e. Si mangia le mele.  
SI eats the apples  
'One eats the apples'
  - f. Si mangiano le mele.  
SI eat the apples  
'The apples are (being) eaten'

(pp.2-3)

The Lexicon-Syntax Parameter is not a stable parameter. Even closely related languages (like German and Dutch) vary in their setting of the Lexicon-Syntax Parameter. Because we do not have either only syntax languages or only lexicon languages, change must be able to occur in both directions. Since verb frame alternations are encoded lexically *unless* they can be encoded syntactically, the trigger for such change must be syntactic in nature. The availability of a new syntactic tool may switch the parameter from lexicon to syntax, and the disappearance of a syntactic tool may switch the parameter from syntax to lexicon. A process leading to the availability of a new syntactic tool may be grammaticalisation. Cross-linguistically, words denoting BODY tend to evolve into reflexive markers, which may subsequently evolve into anticausative, middle, passive and reciprocal markers (Heine & Kuteva, 2002).

In section 3.2 I left unexplained the fact that parameterisable operations such as reflexivisation appear more often to have cross-linguistically consistent derivational morphology. This fact is surprising, recall, because iterated learning does not give rise to consistent derivational morphology - after all, initially there is no derivation going on (this explained the cross-linguistic variation in morphological marking of the causative-inchoative alternation in (45) and (44)). The derivational morphology on parameterisable verb frame alternations is, I believe, a result of the instability of the Lexicon-Syntax Parameter. When a language becomes a syntax language, which I have argued is bound to happen, the lexical entries for reflexives and other parameterisable alternations are forgotten, because from that moment their meanings are expressed syntactically. Expressing those meanings in the syntax, as opposed to in the lexicon, requires a derivation proper. As a result, the verb frame alternations will gain consistent derivational morphology. When a syntax language then later changes back to a lexicon language, the syntactically introduced derivational morphology can be maintained, incorporated in the separate lexical entries for reflexives, reciprocals, etcetera.

### 3.5 Frozen forms

Sometimes an unaccusative entry is available in a language while a corresponding causative entry is not, as illustrated in the following examples (all from Horvath & Siloni, 2005):

- (55) COLLAPSE:  
 a. Hebrew: unaccusative *hitmotet*; transitive *motet*  
 b. Hungarian: unaccusative *sszeesik*; no transitive alternate
- (56) WILT:  
 a. Hebrew: unaccusative *naval*; no transitive alternate  
 b. Hungarian: unaccusative *elhervad*; transitive *elhervaszt*
- (57) VANISH:

- a. English: unaccusative *vanish*; no transitive alternate
- b. Hebrew: unaccusative *ne'elam*; transitive *he'elim*

Because in the mainstream Theta System, unaccusative intransitive verbs are all derived from a causative alternate, the absence of a causative alternate is problematic. This leads Horvath and Siloni (2005) to postulate in such cases a *frozen entry*, a concept that ‘exists in the lexicon but cannot be inserted into syntactic derivations, and hence is not part of the actual vocabulary of the language’.

Horvath and Siloni formulate the following generalisation. Middles and adjectival passives do not exist if the transitive alternate is frozen, i.e. unavailable in the vocabulary, while inchoatives, reflexives and reciprocals do. For example, Hebrew *hitkale'ax* (SHOWER-REFL) does have a frozen causative alternate, as well as the reciprocal *hitvake'ax* (ARGUE-REC). In order to account for this, they propose the following account of a frozen entry:

- (58) **Frozen entry:** An entry is frozen if one of its theta roles is inert, that is, inaccessible outside the lexicon.

Because the [+c] role of frozen causatives is inert, they hypothesise, only operations that remove this role may yield an entry that can be used in the syntax. Indeed, decausativisation removes the inert [+c] role entirely while the operations deriving middles and adjectival passives do not. Middles and adjectival passives, though not syntactically realising the [+c] role, do assign it in the semantics (Marelj, 2004; Reinhart & Siloni, 2005). The presence or absence of an Agent (a possible instantiation of the [+c] role) in the semantics licences an Instrument role or the adverb ‘intentionally’, as in the following examples:

- (59) a. Cause role present in syntax and semantics:  
Causative: Alan broke the window (intentionally) (with a rock)
- b. Cause role present only in semantics:  
Middle: Potatoes peel easily (with a knife)  
Adjectival passive: The (intentionally) frozen pond
- c. Cause role absent entirely:  
Inchoative: The window broke (*intentionally) (*with a rock)

Furthermore, reciprocalisation and reflexivisation are assumed to replace the external role by a new role entirely by bundling it with another role in the theta grid, thereby fixing its inertness. It follows that only inchoatives, reflexives and reciprocals exist with a frozen transitive alternate.

I believe that my account of the Theta System can explain the data without invoking such theory-internal entities as frozen entries and inert roles. In my account there is nothing strange in principle about inchoatives without a causative alternate. The inchoative-preference principle works only one-way, and through iterated learning, the inchoative verb is not ‘derived’ from the causative verb; both verbs are equally basic. But what about the generalisation that only inchoatives, reflexives and reciprocals exist without a fluid transitive alternate? I

believe this can be explained by pertaining to the usefulness of lexical entries again, as in section 3.3.

I propose that, if it is useful to talk about ‘arbitrary humans’ performing an action (as in middles), or about ‘someone’ performing an action (as in passives), then it is also useful to be able to talk about a particular human, say Alan or Bertrand, performing the same action. More specific lexical entries are typically more useful, but only if the kind of events they describe are still sufficiently frequent. For instance, a concept SEE-AND-HEAR would be very specific, which is good, but it would also be very infrequent compared to both SEE, and HEAR, which is bad. Causative lexical entries are potentially much more specific than middles or passives, depending on the arguments they take, and, contrary to SEE-AND-HEAR, the events described by a causative are no more infrequent than the events described by a passive or middle. Indeed, every event that can be described with a passive or a middle can also be described with a causative verb. Intuitively: SOMEONE and ARBITRARY HUMAN are not unlike superordinate concepts compared to a specific person. It follows that causative lexical entries are more useful to have as lexical entries than the corresponding passives or middles. As a consequence, if a language contains a passive or a middle entry, the language will also contain the corresponding causative alternate.

Earlier I explained the relation between causative and inchoative verbs: inchoatives are always more useful than causatives. How about reflexives and reciprocals? Unlike passives and middles, reflexives and reciprocals are sometimes *more* specific than their causative alternates, because some lexical reflexives and reciprocals have idiomatic meanings. For instance, a reciprocal situation of two boys kissing each other’s arms cannot be described with a lexical reciprocal as in ‘the boys kissed’, only periphrastically as in ‘the boys kissed each other’. And if they lack such an idiomatic meaning, at least they are not *less* specific than their causative alternates. Hence, whether reflexives and reciprocals are more or less useful than their causative alternates, depends on the frequency of the events they can describe, relative to the frequency of events which only their causative alternates can describe. Obviously, unlike the usefulness of passives and middles compared to causatives, the usefulness of reflexives and reciprocals cannot be stated as a general rule. Therefore, one would expect that, contrary to passives and middles, reflexives and reciprocals do occasionally occur without a causative alternate, especially if the situations they apply to are very frequent. Indeed, Hebrew *hitkale’ax* (SHOWER-REFL) and reciprocal *hitvake’ax* (ARGUE-REC), used by Horvath and Siloni (2005) to exemplify reflexives/reciprocals with a frozen causative alternate, denote such typically reflexive and reciprocal events.

In sum, ‘derived’ forms may exist without a causative alternate because, in my iterated learning account, they are not in fact derived. Passives and middles never exist without a causative alternate, because their causative alternate is always more useful. Inchoatives, reflexives and reciprocals do exist without a causative alternate, because they can be more useful than their causative alternates. Hence, in my account of the Theta System, the machinery of frozen forms and inert theta roles can be dispensed with.

### 3.6 Conclusion

Approaching the Theta System from a fresh angle and tentatively dropping some of its background assumptions has made available a parsimonious explanation of some of the postulates of the Theta System. The acquisition of lexicon operations is inevitable due to the emergence of verb frame alternations, and the emergence of verb frame alternations is inevitable due to the language acquisition bottleneck. Therefore the lexicon is inevitably active. The universality of acquired lexicon operations, the existence of decausativisation rather than causativisation, and cross-linguistic variation as well as universal properties of morphological marking can all be off-loaded to world structure, the language acquisition bottleneck and iterated learning. The Lexicon-Syntax Parameter follows from the redundancy of lexical entries in the presence of efficient syntactic tools, and changes in its setting can be tied to grammaticalisation. Finally, iterated learning predicts the very patterns that motivated the mainstream Theta System's postulation of frozen forms.

These explanations are available for the concept nativist, too. So far, the relevant data can be accounted for by assuming an acquired active vocabulary instead of an innate active lexicon. The combination of an innate, passive lexicon and acquired vocabulary operations may be able to replace the mainstream Theta System's innate active lexicon. The result would be an interesting and, I believe, more parsimonious take on the mainstream Theta System.

Nevertheless, some important questions remain unanswered. Now that lexicon operations have been demoted, in a way, to patterns of productive generalisation over lexical entries (or vocabulary entries), what can be maintained of the Theta System's interaction between lexicon operations and the linking procedure? What does it mean for a lexicon operation to apply either 'before' or 'after' the linking procedure (a difference that was invoked to explain the large class of intransitive unaccusative verbs)? And what are the origins of this linking procedure anyway? These questions are the topic of the next chapter.

## Chapter 4

# The linking procedure

### 4.1 Introduction

In this chapter I propose that the linking procedure, just like lexicon operations, may be acquired, and I explore the consequences of this claim for the postulates of the Theta System. Acquiring a linking procedure does not seem problematic. The language learner will acquire lexical entries including their phonology, their semantics and their occurrence in various syntactic contexts, i.e. their linking instructions. The learner can then generalise over thematic roles of lexical entries (e.g. Agents tend to be Subjects) and over lexical entries as a whole (e.g. if nothing rules this out, arguments merge externally), and use these generalisations to fill in the gaps in the acquired lexical entries. These may be generalisations over thematic roles (e.g. Agent to Subject, Theme to Object) or over lexical entries (e.g. argument of intransitive verb to Subject *unless* if there is a causative alternate). Indeed, the acquisition of linking generalisations *in principle* appears no less straightforward than the acquisition of lexicon operations. In this chapter I will assume, as a working hypothesis, that the linking procedure can be acquired.

However, a potential argument against an acquired linking procedure exists. Pinker (1989) proposes that the linking procedure must be innate in order to *bootstrap* the acquisition of syntax (contrasting with his earlier work, in which only some innate biases are deemed necessary (Pinker, 1984, 1987)). Therefore, in section 4.2 I evaluate Pinker's argument. I will conclude that the linking procedure need not be innate, and, hence, that I can maintain my working hypothesis. In section 4.3 I explain, assuming that the linking procedure is acquired, how iterated learning can account for split intransitivity. The subsection on split intransitivity will serve to reconcile the mainstream Theta System's account of unaccusative intransitive verbs with my account (an issue left open in the previous chapter). In section 4.4 I propose an explanation, in terms of iterated learning and linking regularities, for the fact that no two arguments of a verb ever bear the same thematic role (thematic distinctness). Section 4.5

contains a conclusion.

Note that the idea of an acquired linking procedure is orthogonal to the concept nativism debate - provided that, at least, a lexical entry's merging instructions are *not* innately specified, but added by the linking procedure later. Indeed, this was one of my assumptions from the start. Hence, the points I make in this chapter will be compatible with the concept nativist's view. An acquired linking procedure is not orthogonal, however, to the question of whether language is, perhaps primarily, a means for thought rather than communication. Though an acquired linking procedure is, I believe, compatible with a language primarily used for thought, establishing this requires the following brief section, in which I distinguish between two kinds of linking procedures.

#### 4.1.1 Communication, thought and the linking procedure

Bickerton (1990) argues that language is primarily a means for thought, communication being a side-effect that turned out to be useful as well. Although this idea is by no means universally accepted (e.g. Deacon, 1997 for an opposite view), generative linguistics in general is not unsympathetic to the idea (e.g. Chomsky, 1986). Therefore I will briefly explain why I believe an acquired linking procedure is compatible with the language-for-thought view. After all, this is not trivial: how could language be primarily a means for thought if one of its fundamental building blocks, the link between lexicon and syntax, must be *acquired* first from utterances that are the result of communication? Well, perhaps there is an innate linking procedure *as well as* an acquired linking procedure, however redundantly this may sound.

I propose that linking *regularities*, while necessary for communication, are not necessary for thought. In using language for thought, the merging order of a lexical entry's arguments does not matter (much) for the resulting event semantics.¹ As long as the right thematic roles are assigned to the right arguments, it does not matter so much where in the syntax these arguments end up. For communication, however, it is important that the resulting syntactic structure is not only interpretable to the speaker, but also, in spoken form, comprehensible to the hearer. For communication to be reliable, the thematic role assignment of a verb to its arguments must be predictable from the sentence structure; i.e. there must be a regular linking procedure that all speakers of a language know and use.

Perhaps a pre-linguistic mind contains an innate linking procedure that is not conditioned on thematic roles, but on entirely different semantic criteria that are more relevant for thought. But when communication kicks in, this linking procedure is overruled by (or perhaps coexists with) an acquired, thematic-role-based, culturally shared, linking procedure required for communication; i.e. a linking procedure as transparent as the Theta System's. In sum, the hypothesis

---

¹I am ignoring relevant differences pertaining to e.g. scope, control or pronoun resolution that may exist between Subject and Object. I speculate that in the absence of a regular linking procedure, these would have been the primary criteria determining the order of merger.

that language is primarily a means for thought is compatible with an acquired linking procedure.²

## 4.2 The linking procedure and semantic bootstrapping

The bootstrapping problem in language acquisition is as follows (Pinker, 1984, 1987). Syntax is a symbolic system, in which the symbols are defined only with respect to other symbols in the same system. In (Pinker, 1984) such symbols are syntactic categories (noun, verb, etc.), grammatical functions (Subject, Object, etc.), cases (Nominative, Accusative, etc.), grammatical features (Tense, Number, etc.) and tree configurations (sister-of, etc.) (p.41). Although in current linguistic theories the inventory is somewhat reduced, the problem is the same. Each of these symbols can only be defined with respect to other symbols in the same system (for instance, the Subject is marked with Nominative case in the majority of languages). The problem for the language learner (and, as Pinker points out, for the field-linguist) is to break into this system of circular definitions.

The *Semantic Bootstrapping Hypothesis* advocated by Pinker claims that the metaphorical bootstraps by means of which the language learner pulls himself/herself out of this impasse are semantic in nature. It claims that the acquisition of syntax is at the onset guided by prototypical syntax-semantics correspondences like ‘Subject is Agent’ (Pinker, 1984, 1987) or by the linking procedure itself (Pinker, 1989). Of course, for the typical correspondences or the linking procedure to function as bootstraps, they cannot themselves be acquired and hence must be innate. Once the language learner has broken into the symbolic system, (s)he can then use syntactic distributional analysis to acquire the language-specific properties of e.g. Subject and Object. There is no consensus with respect to the validity of this solution to the bootstrapping problem, or even with respect to the reality of the bootstrapping problem (e.g. Gleitman, 1990; Rispoli, 1987). But if we suppose that semantic bootstrapping is indeed the mechanism that enables an infant to acquire syntax, I will give two arguments against Pinker’s (1989) claim that the linking procedure is innate (I will call this ‘deterministic semantic bootstrapping’), and indirectly in favor of Pinker’s (1984, 1987) idea that innate biases must suffice (I will call this view ‘probabilistic semantic bootstrapping’).³

---

²A very interesting line of argument could be that a linking procedure like the Theta System’s is so patently optimised for communication and not for thought, that its innateness would be incompatible with the idea that language is primarily a means for thought. But I cannot foresee whether this claim could be solidified, and in any case doing so would not be of particular relevance for the aims of my thesis.

³A speculative but interesting, third kind of bootstrapping hypothesis becomes available in light of the previous subsection: perhaps we bootstrap our way into syntax using our innate linking procedure for thought, after which this linking procedure is replaced with or accompanied by our acquired linking procedure for communication (which may share some rudimentary properties with the innate linking procedure, in order for semantic bootstrapping

#### 4.2. THE LINKING PROCEDURE AND SEMANTIC BOOTSTRAPPING⁴⁷

First, for deterministic semantic bootstrapping to work, Pinker needs to assume that in addition to conceptual structure there is a separate level of semantic representation. Pinker assumes that, unlike conceptual structure, this autonomous level of semantic representation is sufficiently abstract and ‘language-particular’ for there to be a deterministic mapping to syntax. I have three objections to this separate semantic representation. First, it presupposes that there *is* a deterministic mapping between the symbols of syntax and non-syntactic entities, in this case semantic predicates - precisely the opposite of what gave rise to the bootstrapping problem in the first place.

My second objection concerns Pinker’s main argument in favor of an autonomous semantic representation. He observes, correctly I believe, that the world is not carved up along the same lines by cognition as it is by language. This is illustrated by the following examples:

- (60) a. The three-way distinction between *handing*, *carrying* and *taking* is unlikely to be as cognitively salient as the distinction between, say, *throwing*, *kicking*, and *rolling*, yet as far as the dative is concerned, each of the first three belongs to a different class while the latter three belong to the same class.
- b. The English language, but not its speakers when they are not speaking, must consider *telling* to be a different kind of activity than *saying*, *shouting*, *talking* or *speaking*, but the same kind of activity as *quoting*, *leaking*, *asking*, *posing*, or *writing*. Conversely, *shouting* is no more similar to *yelling* or *screaming* than it is to *whispering* and *murmuring*.
- c. *Baking a cake* has to be construed as similar to *building a house* and *writing a letter of recommendation* but as dissimilar to *warming a cake*, *burning a cake*, or *reheating a cake*.
- d. *Betting* has to be represented as being like *envying*, *sparing*, and *begrudging* but unlike *selling*, *paying*, or *trading*.

(Pinker, 1989, p.359)

If semantic representation and conceptual structure had been the same system, Pinker argues, surely the same conceptual distinctions would have been salient for both language and cognition. But while the examples are valid, Pinker seems to miss the point that language has adapted not only to *reflect* concepts as closely as possible (which increases language learnability) but also to *convey* concepts as reliably as possible (which increases language usability).⁴ In order to convey concepts, obvious distinctions can be safely ignored while the more subtle distinctions must be explicitly encoded. For instance, there is no reason to linguistically encode (in a way other than the phonology of the words themselves) the difference between *betting* and *envying/sparing/begrudging* because

---

to work).

⁴Recall that arguments based on learnability and usability, while irrelevant in the mainstream Theta System, are relevant in both my account and the concept nativist’s account.

there is no danger of confusing these concepts. Same for *whispering*, *murmuring*, *yelling*, *shouting* and *screaming*, which are all safely onomatopoeic. The difference between *betting* and *selling/paying/trading*, on the other hand, is so subtle that it must be explicitly encoded lest it be missed. And similarly the difference between *saying*, *shouting*, *talking* and *speaking*. Language does not exaggerate small differences between words because such differences happen to be important at some autonomous level of semantic structure, but because those differences are too small to leave implicit.

Third, Pinker acknowledges, a problem for the semantic bootstrapping hypothesis (both deterministic and probabilistic) are syntactically ergative languages. In syntactically accusative languages, which form a large majority, the argument of an intransitive verb patterns with the Subject of a transitive verb. In contrast, in syntactically ergative languages, no more than 5% of the world's languages (Marantz, 1984), the argument of an intransitive verb is treated like the Object of a transitive verb and distinctly from the Subject of a transitive verb. While bootstrapping the acquisition of ergative syntax seems very problematic for a learner who is innately endowed with a linking procedure attuned to accusative syntax, as deterministic semantic bootstrapping seems to require, it seems a lot less hopeless for probabilistic semantic bootstrapping. Indeed, Pinker himself proposes that syntactically ergative languages are the result of conflicting constraints/biases.

In light of these objections to Pinker's (1989) deterministic semantic bootstrapping, probabilistic semantic bootstrapping is the most promising solution to the bootstrapping problem, and the linking procedure can be independently acquired.

### 4.3 The origins of linking regularities

If the linking procedure is acquired, this raises the same chicken-and-egg problem as the acquisition of lexicon operations: do speakers acquire a linking procedure because the mapping from meaning to syntax is regular, or does the language exhibit such regularities because speakers have a linking procedure? As in chapter 3, I propose that gradually accumulating language change is to blame. Acquiring lexical entries with their linking instructions is subject to an acquisition bottleneck. Whenever a verb's linking instructions are not completely acquired, the language learner will fill in the gaps according to his/her inductive biases, such as the regularisation bias introduced in chapter 3. Iterated learning exaggerates the regularisation bias, leading to a regular linking procedure.

While this explains why each language has linking regularities, it does not yet explain why these should be *the same* linking regularities in every language. Because iterated learning makes languages converge on the language learners' inductive biases, giving an account of the near-universality of the linking regularities is a matter of identifying which biases are involved. In what follows, I will propose an explanation of the universality of split intransitivity in terms of

inductive biases.

*Split intransitivity* is the property of natural language that intransitive verbs display two kinds of behaviour. Some intransitive verbs trigger an *unergative* derivation: the single argument merges externally, i.e. in Subject position. Other intransitives trigger an *unaccusative* derivation: the single argument merges internally, i.e. in Direct Object position. There are a number of diagnostics for unaccusativity. In Dutch, selection of auxiliary *zijn* ('be') and an available impersonal passive are two such diagnostics:

- (61) Unaccusative:
- a. De gevallen/gestorven/gebroken/verdwenen man.  
The fallen/died/broken/disappeared man
  - b. De man is/*heeft gevallen/gestorven/gebroken/verdwenen.  
The man is/has fallen/died/broken/disappeared
- (62) Unergative:
- a. De *gekomen/*gedwaalde/*gebloosde man.  
The came/wandered/blushed man
  - b. De man *is/has gekomen/gedwaald/gebloed.  
The man is/has came/wandered/blushed

Many of the syntactic differences between unergative and unaccusative verbs can be explained by assuming that the argument of an unaccusative is, underlyingly, an Object (Perlmutter, 1978). In other words, the differences between unergative verbs and unaccusative verbs are the result of differences in linking.

Reinhart (2000) shows that these differences in linking cannot be, as would be ideal, accounted for in terms of differences in s-selection. For instance, a generalisation in terms of thematic roles would not be able to explain the existence of Theme unergatives in addition to unaccusatives, which also take a Theme:

- (63) a. Agent unergatives: *walk, run, eat, etc.*  
 b. Theme unergatives: *sweat, glitter, tremble, rattle, etc.*  
 c. Unaccusatives: *break, slide, defrost, bend, etc.*

In order to account for the differences in linking, the mainstream Theta System's linking procedure contains two main deviations from what would be the most straightforward, i.e. pure Agent-Subject, Theme-Object alignment. First, the argument of an intransitive verb without a causative alternate is 'left unmarked' and merges externally, regardless of its role, triggering an unergative derivation. This is an instance of the generally accepted *merge-over-move* principle motivated by cognitive economy. Second, the argument of what I have so far called an *inchoative verb* - an intransitive verb with a causative alternate - appears to inherit its linking instructions from the causative alternate.⁵ Although this

⁵While this is not a component of the marking or merging procedure (see (16) and (17) in chapter 1), I treat this interaction with lexicon operations as part of the Theta System's 'linking procedure in the broad sense', because, indeed, it affects linking.

inheritance is only visible if the argument merges internally, triggering an unaccusative derivation (64a), it makes sense to assume that a similar inheritance occurs for verbs like *worry*, which also have a causative alternate but trigger an unergative derivation (64b).

- (64) a. *break*([+c]₁[-c-m]₂): Alan broke the window  
       *break*([-c-m]₂): The window_{*i*} broke ___{*i*}  
 b. *worry*([+c]₁[-c+m]): The war worried Bertrand  
       *worry*([-c+m]): Bertrand worried

In the mainstream Theta System this inheritance is explained, recall from chapter 1, by assuming that Decausativisation does not create a new concept, and hence applies ‘after the marking procedure’, while, for instance, Agentivization does create a new concept, and hence applies ‘before the marking procedure’ (Reinhart, 2002). But in the account of the Theta System that I explore, this explanation cannot be maintained. First, the difference between concept-creating and non-concept-creating operations does not translate to a view of lexicon operations as productive generalisation. Second, if both the linking procedure and lexicon operations are productive generalisation, there is no one applying before the other, at least not systematically for a whole class of verbs.

Instead, because I assume that the linking procedure is acquired, I have no choice but to acknowledge that the generalisation ‘intransitive verbs with a causative alternate inherit their linking instructions from it’ is as much part of the linking procedure as the generalisation ‘if nothing rules this out, merge externally’ is - both are generalisations, hence both can be acquired and productively applied. While this may appear to be begging the question, this change in perspective allows us to move the burden of explanation from the individual’s grammar to inductive biases and iterated learning; a strategy that seemed to work well in the previous chapter.

First, let me introduce some terminology. Dowty (1989) formally defines a thematic role as a set of entailments of a group of predicates with respect to one of the arguments of each. If the group of predicates is relevant for argument projection, for instance all the verbs that project the role bearer as a subject, Dowty calls the resulting thematic roles *L-thematic roles*. Presumably, the theta clusters all denote L-thematic roles. If the ‘group’ of predicates contains only one predicate, Dowty, following (Marantz, 1984; Riemsdijk & E, 1986), calls the resulting thematic roles *individual thematic roles*, which I will abbreviate as ‘I-thematic roles’. For example, the verb *break* has the I-thematic roles Breaker and Breakee, which logically contain the L-thematic roles Cause and Theme.

Based on this distinction, I introduce the following terminology. The *L-thematic alignment bias* is the expectation that identical L-thematic roles merge in identical syntactic positions.⁶ The *I-thematic alignment bias* is the expectation that identical I-thematic roles merge in identical syntactic positions. Both biases are instantiations of the more general expectation that similar forms

⁶This bias corresponds to M. C. Baker’s (1988) Uniformity of Theta Assignment Hypothesis (UTAH). However, it is only a bias, an expectation, not a rigid linking procedure.

have similar meanings (cf. section 3.2); assigning them different names is only for clarity in what follows. The I-thematic alignment bias is of course more powerful than the L-thematic alignment bias; for instance, the Breakee arguments of intransitive *break* and transitive *break* are much more similar than the Theme arguments of *glitter* and *break*. In addition, there is an *economy bias*: it is more economical to merge the argument of an intransitive verb in Subject position immediately, rather than via the Object position (as in an unaccusative derivation).

The table in (65) summarises how these biases interact for each kind of intransitive verb. For Agent unergatives, I-thematic alignment is irrelevant (there is no causative alternate with the same I-thematic role), and economy (merge over move) and L-thematic alignment (Agents of transitive verbs are also Subjects) are both satisfied. For Theme unergatives, I-thematic alignment is again irrelevant, economy is satisfied while L-thematic alignment is violated (Themes of transitive verbs are Objects, not Subjects). The existence of Theme unergatives in English suggests that the economy bias is stronger than L-thematic alignment. However, this is a close and perhaps accidental victory, as the group of Theme unergatives exhibits considerable cross-linguistic variation in their syntactic behaviour as well as within-language ‘fluidity’ (Sorace, 2004; Levin & Rappaport-Hovav, 1995). For unaccusatives, I-thematic alignment is relevant (there is a causative alternate for such verbs) and it overrules the economy bias violation.

		I-thematic alignment	Economy	L-thematic alignment
(65)	Agent/Cause unergatives		+	+
	Theme unergatives		+	-
	Unaccusatives	+	-	+

Each kind of intransitive verb is a best fit to these three competing biases. But, crucially, it is not the individual language learner who optimises his/her lexicon in this fashion. The individual learner will do so only to fill in the gaps that are due to the acquisition bottleneck. The biases are exaggerated by iterated learning, leading ultimately to the three-way partition of verbs that is captured in the Theta System’s linking procedure. This account of linking regularities does not require that a decausativisation operation applies after the marking procedure, or that the individual speaker makes the merge-over-move decision every time an unergative verb is used. Rather, these are patterns that have emerged incrementally over many generations, acquired and applied by the individual language learner through productive generalisation.

The translation of the mainstream Theta System’s linking procedure to my account of the Theta System has been remarkably easy - so easy, in fact, that one might wonder what is gained by making this translation. Foremost, it shows that the Theta System’s empirical coverage can be maintained, without postulating an innate linking procedure and innate lexicon operations, without requiring that some lexicon operations apply after and others before linking, without requiring that only one thematic form of each verb concept is lexically stored, and without

assuming that lexical entries do not contain their own linking instructions. And all this is compatible, by the appropriate substitutions, with what I have been calling ‘my account’ as well as the concept nativist’s account of the Lexicon.

## 4.4 Thematic distinctness

There is a close relation between the linking procedure, which links a verb’s arguments to syntactic positions based on their thematic roles, and the fact that verbs never assign the same role to more than one argument. I will call the latter universal property of verbs *thematic distinctness* (Dowty’s (1989) ‘distinctness’, Parsons’ (1990) ‘uniqueness’).⁷ In the Theta System, thematic distinctness is formulated as the Identity Constraint (IC) (66). Marelj (2002) assumes, in order to account for the behaviour of double Object constructions (which I will not treat in this thesis), that the IC holds at the interface between the lexicon and syntax as well as at the interface between syntax and the inference system.

(66) **Identity Constraint (IC):**

Two identical theta-roles cannot realize on the same grid. (Marelj, 2002, p.2)

Thematic distinctness or the IC implies that there are no verbs such as hypothetical *skick*, which is like *kick* but with two Agents, such that ‘John skicked the ball Bill’ means that John and Bill both kicked a ball (example from Carlson, 1984). But while thematic distinctness is assumed in most accounts of linking, its origins and universality are only rarely the topic of inquiry (Carlson, 1998).

It is useful to subdivide the notion ‘thematic distinctness’ according to the division drawn by Dowty (1989) between L-thematic roles and I-thematic roles, explained in section 4.3. Summarising, an I-thematic role is the set of entailments of a verb with respect to one of its arguments. For instance, *break* has I-thematic roles Breaker and Breakee. L-thematic roles like Agent and Theme contain only the entailments relevant for linking (c and m, in the Theta System). Following this division, I distinguish *L-thematic distinctness* from *I-thematic distinctness*:

(67) a. **L-thematic distinctness:**

No verb has two arguments that bear the same L-thematic role

b. **I-thematic distinctness:**

No verb has two arguments that bear the same I-thematic role

L-thematic distinctness is equivalent to the IC, and it is the strongest property; it entails I-thematic distinctness. If a verb assigns distinct L-thematic roles

⁷Carlson seems to mistake the first clause of Kaplan and Bresnan’s (1982) Uniqueness, which states that a lexicon-syntax mapping must assign a unique grammatical function (Subject, Object, etc.) to each of a predicate’s arguments, with thematic distinctness as he intends it. Thematic distinctness is a property of verbs while Kaplan and Bresnan’s Uniqueness is a property of the linking procedure, which follows straightforwardly from the fact that clauses in natural language happen to have only one of each syntactic position, i.e. one Subject, one Object and one Indirect Object.

(i.e. roles distinct with respect to one or more features relevant for argument projection) it automatically assigns distinct I-thematic roles (i.e. roles distinct with respect to any features). In the literature, the discussion of thematic distinctness has been limited to L-thematic distinctness.

In this section I try to explain why all languages exhibit L-thematic distinctness. The structure of this section is as follows. In section 4.4.1 I review and reject an argument raised in (Parsons, 1990). In section 4.4.2 I propose a new explanation of L-thematic distinctness in terms of iterated learning. It will turn out that, although L-thematic distinctness entails I-thematic distinctness, iterated learning can explain L-thematic distinctness only if I-thematic distinctness is independently explained. In section 4.4.3, in search of an independent explanation of I-thematic distinctness, I review Carlson's (1998) proposal that thematic roles serve the individuation of events. Rejecting it as an explanation of I-thematic distinctness, I propose a new explanation in section 4.4.4 based, again, on iterated learning.

Both explanations in terms of iterated learning will turn out to be remarkably straightforward, so why bother rejecting Parsons' and Carlson's (1998) accounts first? Because those two accounts seem to support that the IC holds not only as a property of lexical entries, but also *at the level of inference*; and I believe that that is not the case.

#### 4.4.1 Parsons (1990): event semantics

In her discussion of the Identity Constraint, Marelj (2004) quotes the following footnote from (Parsons, 1990), here with a little context:

[...] thus, each event possesses at most one Agent, at most one Experiencer, and so on.⁵ (p.74)

---

5. This requirement of uniqueness is practically forced on us by the type of logical forms employed. [...] Suppose we were to label as themes both the direct and the indirect objects. Then, the logical form of a sentence containing both such items would be logically equivalent to the sentence with the direct and indirect object interchanged. If you gave a fish to Mary, you would thereby give Mary to a fish. (p.293)

The term 'event' in Parsons' work denotes the event variable in event semantics; it does not denote event in any language-independent sense, e.g. events in our categorisation of the world. This contrast is important to keep in mind, because in section 4.4.3 I review Carlson's (1998) argument, in which 'event' seems to denote something in our categorisation of the world.

There are two readings. First, the 'suppose'-sentence in the above quote can mean the following: suppose we were to label as Themes both the direct and the indirect objects *of a verb like* give. I will call this the particular reading because it involves a particular (hypothetical) lexical entry. The second reading is more general: suppose the grammatical positions Direct Object and Indirect Object

are associated with the same L-thematic role in general, independently of the verb *give*. Although I believe that the particular reason is as Parsons intends it, I will treat both readings to show that, either way, his argument does not work.

### The particular reading

Suppose we were to label as Themes both the direct and the indirect object of a verb like *give*. In Parsons's event semantics, this would yield the following logical forms for each 'fish-giving' sentence:

- (68) a. Alan gave a fish to Mary  
 $\exists e \text{ GIVE}(e) \wedge \text{Agent}(\text{alan}, e) \wedge \text{Theme}(\text{mary}, e) \wedge \text{Theme}(\text{the fish}, e)$   
 b. Alan gave Mary to a fish  
 $\exists e \text{ GIVE}(e) \wedge \text{Agent}(\text{alan}, e) \wedge \text{Theme}(\text{the fish}, e) \wedge \text{Theme}(\text{mary}, e)$

(68a) and (68b) are logically equivalent; one would entail the other. But as Parsons acknowledges in the first sentence of his footnote, the fish-example goes wrong due to the type of logical forms employed. In particular, it goes wrong due to Parsons' assumption that the logical form of a sentence contains only the L-thematic roles. As I explained in chapter 1, L-thematic roles need not be considered in any way cognitively real. The fact that the linking procedure is sensitive only to the semantic features *c* and *m* does not mean that only these features are assigned to a verb's arguments in the syntax and passed on to logical form: after all, all the other (I-thematic role) entailments are still there at the level of interpretation. If both the direct and indirect object of *give* were Themes, as Parsons asks us to imagine, this would not mean that both arguments would be indistinguishable at the level of interpretation; it would only mean that the Given and the Givee are indistinguishable *for the linking procedure*. Indeed, a more adequate representation of the logical form might be one in terms of I-thematic roles:⁸

- (69) a. Alan gave a fish to Mary  
 $\exists e \text{ GIVE}(e) \wedge \text{Giver}(\text{alan}, e) \wedge \text{Givee}(\text{mary}, e) \wedge \text{Given}(\text{the fish}, e)$   
 b. Alan gave Mary to a fish  
 $\exists e \text{ GIVE}(e) \wedge \text{Giver}(\text{alan}, e) \wedge \text{Givee}(\text{the fish}, e) \wedge \text{Given}(\text{mary}, e)$

Obviously, (69a) and (69b) are logically distinct. Even if Givee and Given would somehow both fall under the same L-thematic role Theme, one sentence would still not entail the other. Hence, although Parsons's footnote may support L-thematic distinctness in his particular implementation of event semantics, this is only a theory-internal explanation of a theory-internal kind of thematic distinctness.

⁸This may appear redundant, but (68a) and (68b) may contain redundancy too: it seems impossible to conceptualise the meaning of the conjunct *give(e)* without conceptualising an Agent doing the giving.

### The general reading

According to the general reading, Parsons asks us to suppose that Direct Object and Indirect Object are associated with the same L-thematic role, i.e. Theme, independent of lexical entries. It is unclear to me how this could yield an argument in favor of L-thematic distinctness as I have defined it, as a property of lexical entries, but I will evaluate the argument nevertheless.⁹ Supposing that Direct Object and Indirect Object are both associated with Theme would mean that the linking procedure does not distinguish between Direct Object and Indirect Object, and that an argument labeled Theme could end up in either position. Similarly to my objection against the particular reading, this does not mean that the Direct Object and Indirect Object of a given verb are therefore indistinguishable at the level of interpretation - that is an artefact only of Parsons' event semantics.

#### 4.4.2 L-thematic distinctness through iterated learning (almost)

We need an alternative explanation for L-thematic distinctness. Labeling two arguments as Theme, as Parsons (1990) asks us to imagine, does not matter for the level of interpretation except in his implementation of event semantics. It does matter, however, for the linking procedure. If both the Giver and the Givee were Themes (supposing, for the sake of this example, that the linking procedure would not be sensitive to the Giver-Givee distinction), language learners would not be able to decide which argument goes where, except when acquiring the verb in its entirety from direct evidence. Through iterated learning, lexical entries that cannot be reliably acquired except from direct evidence tend to undergo change or disappear from a language. This is the reason we observe L-thematic distinctness in today's languages.¹⁰

However simple, this explanation has so far been overlooked. For instance, Tenny (1994) argues that Dowty's (1991) approach based on the arguments' resemblance to prototypical roles will not give rise to L-thematic distinctness, and regards this as an advantage of her own Aspectual Roles approach. But it is easy to see that Dowty's (1991) approach, too, gives rise to L-thematic distinctness. After all, a lexical entry that takes two arguments with the same number of Proto-Agent and Proto-Patient entailments cannot be reliably acquired except from direct evidence. Turning Tenny's argument around, it may even be an

⁹Instead, if valid, it would be an argument in favor of the converse of M. C. Baker's (1988) Uniformity of Theta Assignment Hypothesis (UTAH). The UTAH states that identical thematic relationships between items are represented by identical structural relationships between those items at the level of deep structure. Hence, the converse of the UTAH is that different structural relationships are associated with different thematic relationships. Parsons asks us to imagine that this converse does not hold because Direct Object as well as Indirect Object are associated with Theme, and he then shows that this goes wrong.

¹⁰To view this explanation from a different perspective: a hypothetical 'proto-language' that lacked linking regularities may not have had L-thematic distinctness either. Similarly, if what I hypothesised at the beginning of this chapter is correct, language as used for thought need not exhibit L-thematic distinctness either.

advantage of Dowty’s approach that, contrary to the theory of Aspectual Roles, it does not have thematic distinctness built-in from the start.

Note, however, the ‘almost’ in the title of this section. If two arguments bear the same I-thematic role, it *does not matter* which argument goes where, so they may as well bear the same L-thematic role. Even if we would assume like Parsons (1990) that it is L-thematic roles, not I-thematic roles, that are assigned to arguments and end up in logical form, the logical indistinguishability of two arguments would only be a problem if both arguments bear different I-thematic roles, as in the case of *give* (Givee is not identical to Given). But consider again the hypothetical verb *skick*, meaning KICK with two Agents doing the kicking, both bearing the same I-thematic role:

- (70) a. Alan skicked the ball Bill  
 $\exists e \text{KICK}(e) \wedge \text{Agent}(\text{alan}, e) \wedge \text{Theme}(\text{the ball}, e) \wedge \text{Agent}(\text{bill}, e)$   
 b. Bill skicked the ball Alan  
 $\exists e \text{KICK}(e) \wedge \text{Agent}(\text{bill}, e) \wedge \text{Theme}(\text{the ball}, e) \wedge \text{Agent}(\text{alan}, e)$

Obviously, (70a) and (70b) are logically equivalent. But the difference with *give* is that here it does not matter, because the sentences do in fact mean the same. Even replacing the L-thematic roles in the logical form by I-thematic roles, as I have proposed in my discussion of Parsons’ argument, would not make a difference.

In sum, iterated learning leaves open the possibility for verbs that violate L-thematic distinctness to survive the bottleneck, provided they assign identical I-thematic roles; such verbs need not be reliably acquired in order to be usable. For my explanation of L-thematic distinctness to be complete (and similar for Parsons’ (1990) explanation, if it had been valid), I-thematic distinctness must be independently explained (paradoxically, despite L-thematic distinctness *entailing* I-thematic distinctness). In what follows I reject a rather complex explanation of I-thematic distinctness and give a new explanation in terms of iterated learning.

### 4.4.3 Carlson (1998): event individuation

Carlson (1998) does not distinguish explicitly between I-thematic roles and L-thematic roles, although his examples seem to be aimed primarily at L-thematic roles. Since L-thematic distinctness implies I-thematic distinctness, either version of Carlson’s argument would, if valid, end our search for the origins of I-thematic distinctness.

Carlson’s (1998) explanation of thematic distinctness is as follows. By assumption, we organise the world around us by means of events as we do with things. Just like we can count the number of things in a room, we can count the number of events in a room.¹¹ Carlson assumes that we count the number

¹¹As mentioned earlier, this notion of ‘event’ is different from Parson’s (1990). Events for Carlson are non-linguistic, conceptual entities used for carving up the world in meaningful pieces. Events for Parsons are formal, linguistic entities: variables in the event semantic representation of a logical form.

of events in a room by counting the number of event participants that bear the same thematic role, i.e. we *individuate events* by identifying the roles of event participants. If the existence of e.g. two Agents or two Patients is asserted, we can infer that there must be two separate events:

- (71) Two Agents, two events:  
 a. Alan walked through the park  
 b. Bertrand walked through the park
- (72) Two Patients, two events:  
 a. Alan kissed Bertrand  
 b. Alan kissed Mary

As a further illustration, adopted from Carlson, the sentences in (73), which share an Instrument participant, can jointly describe a single event, while the sentences in (74) can only describe two different events because they have different Instruments:

- (73) One Instrument, one event:  
 a. John illuminated the room with the light  
 b. John alerted the burglar with the light
- (74) Two Instruments, two events:  
 a. John flipped the light switch with his finger  
 b. John illuminated the room with the light

If we assume the strong, L-thematic role version of Carlson's argument, his examples support the argument only when the set of L-thematic roles is as he assumes. For instance, the sentences in (75) can jointly describe the same event, because, Carlson argues, the light switch, the room and the burglar bear different L-thematic roles in the event, i.e. Patient, Location and Experiencer, respectively.

- (75) a. John flipped the light switch  
 b. John illuminated the room  
 c. John alerted the burglar  
 (from (Lombard, 1985))

But when recast in terms of feature clusters, this example does not support the argument anymore. The light switch and the room would bear the same theta cluster: [-c-m]. Neither participant is a sufficient condition for the event to take place (-c) and their mental states (if present at all) would be irrelevant for the event (-m).¹² The Theta System is not the only approach under which the light switch and the room bear the same thematic role. Under Dowty's (1991)

¹²The verb *illuminate* also has an 'experiencer'-reading, so its lexical specification may contain the underspecified [-c] rather than [-c-m], allowing for a [-c+m] expansion of the cluster. But even if this is the case, it would not contradict the point at hand, which concerns thematic roles at the level of interpretation rather than in the lexical entry.

approach both participants have the same Proto-Patient properties (change of state and causal affectedness) and under Tenny's (1994) classification both participants bear the MEASURE role.¹³

In an attempt to rescue this example we could instead assume the weak, I-thematic role version of Carlson's claim. That is, one interprets the sentences in (75) as referring to the same event, because each argument bears a distinct I-thematic role: Flippee, Illuminatee and Alertee. However, if I-thematic roles individuate events, the pair of sentences in (76) would be expected to jointly refer to the same event; after all, Alan and Bertrand bear different I-thematic roles (Flipper and Alerter, respectively):

- (76) a. John flipped the light switch with his left hand.  
b. Alan alerted the burglar with the light.

Carlson's examples are not very promising for an account of thematic distinctness in terms of event individuation, but they are not decisively against it either.

Instead, consider the formal structure of Carlson's argument. There are two premises. First, any event in the real world has at most one entity playing a given thematic role (such that thematic roles can be used for event individuation). Second, each verb expresses only one event. It follows logically from these premises that each verb has at most one entity playing a given thematic role, i.e. that thematic distinctness holds. The second premise is well-motivated, but the first premise is more speculative, and I believe it is false or, at least, unmotivated.

According to the first premise, we are unable to conceptualise events in a way that involves multiple participants with the same role. For instance, we are unable to conceptualise a collision of two objects, a shaking-hands of two people, a marrying of two people, or a cooperative piano-carrying as involving multiple participants with the same role. The claim is not that such events in the real world are never purely symmetrical, but that even if they were truly symmetrical (for instance in a computer simulation), we would still be unable to conceptualise them as such. This is a strong claim on 'something the mind cannot do'. Although I cannot prove it wrong, I can question Carlson's (1998)

---

¹³Even when using the set of L-thematic roles assumed by Carlson, some variations on his examples do not seem to match his predictions. In (1) two different participants, John's left hand and the light, bear the Instrument role, yet as far as I can tell both sentences could refer to the same event of John flipping the light switch, thereby alerting the burglar.

- (1) Two Instruments, one event:  
a. John flipped the light switch with his left hand  
b. John alerted the burglar with the light
- (2) Two Instruments, one event:  
a. John searched the room with the light  
b. John carried the light with his left hand

But this only shows that, as Carlson himself acknowledges, evaluating such data is an empirical matter and not one of intuition. My counterexample does by no means prove that L-thematic roles do not individuate events.

motivation. The only motivation Carlson (1998) provides for this cognitive inability is linguistic in nature. For instance, Carlson argues, correctly I believe, that only one thematic role is being assigned to the Subject in ‘the boys carried the piano up the stairs’ (i.e. the boys collectively are a Piano-Carrier) and to the Subject in ‘Alan and Bertrand shook hands’ (i.e. Alan and Bertrand collectively are a Mutual-Hands-Shaker). But events are encoded like this linguistically not because we cannot conceptualise events as having multiple participants with the same role, but because we happen to lack verbs that assign the same thematic role to multiple participants. In other words, Carlson’s argument is circular, because his motivation for the first premise assumes precisely the property he is trying to explain: thematic distinctness.

In conclusion, I reject Carlson’s explanation of thematic distinctness in terms of event individuation. This does not mean that I deny that thematic roles are used as a heuristic in deciding whether two sentences can describe the same event or not - indeed, it would be surprising if we did not use such salient meaning components for this purpose. But Carlson’s slogan ‘two Agents, two events’ is only a heuristic. The mapping between events and event participants is insufficiently rigid to explain the universality of L-thematic distinctness or, weaker but sufficient for our purposes, I-thematic distinctness.

#### 4.4.4 I-thematic distinctness through iterated learning

Instead of relying on event individuation, I propose that I-thematic distinctness is the result of a much more fundamental property of the mind: the regularity bias, i.e. the expectation that different forms have different meanings. Indeed, this is the same bias that guides productive generalisation in the acquisition of lexicon operations and the linking procedure. Because clauses have only one of each syntactic position (Subject, Direct Object, Indirect Object), two arguments of a verb will inevitably be linked to different syntactic positions (this is the first clause of Kaplan and Bresnan’s (1982) ‘Uniqueness’). The language learner is biased to expect that a difference in syntactic position correspond to difference in semantics. Because iterated learning exaggerates biases, the regularity bias leads to I-thematic distinctness (but not necessarily to L-thematic distinctness, which I argued earlier is a separate result of iterated learning).

## 4.5 Conclusion

An explanation of the origins of the linking procedure in terms of acquisition and iterated learning, has enabled my account of the Theta System to explain and in essence maintain some of the mainstream Theta System’s core mechanisms. Like lexicon operations, the linking procedure may be but an instance of productive generalisation over patterns in E-language. The mainstream Theta System’s solution for split intransitivity in terms of an innate linking procedure and its interaction with lexicon operations, translates remarkably well into an explanation of split intransitivity through iterated learning. It is compatible

also with the concept nativist's account.

Furthermore, in the presence of linking regularities (be they innate or acquired), the Theta System's Identity Constraint, or thematic distinctness, is a consequence of iterated learning. However, the Identity Constraint that emerges through iterated learning is a property only of lexical entries; it does not hold at the interface between syntax and inference, as has been assumed wrongly (or without motivation) by Parsons (1990) and Carlson (1998). Marelj's reason for postulating the the IC holds also at the interface to the interference system, is that it enables here to account for the behaviour of double Object constructions in the Theta System. But an evaluation of that argument and an new account of double Object constructions without pertaining to IC at the level of inference, is, I feel, beyond the scope of my thesis. (For a possible mode of explanation that I have pursued in some detail, and that I believe may lead towards a partial answer, see (Arad, 1998)).

## Chapter 5

# Conclusion

The exploration in this thesis has been motivated by three aspects of the mainstream Theta System that I found unsatisfying: the combination of concept nativism with lexicon operations, the stripped-down lexicon, and its reliance on Universal Grammar for the explanation of language universals. With regard to the first aspect, I considered two possible alternatives:

- There are no lexicon operations; or
- Concepts are not innate.

I adopted, as a working hypothesis, the second option, and explored whether the postulates of the Theta System could be explained as patterns due to an acquisition bottleneck. From this new perspective I have provided the following explanations for the core postulates of the Theta System:

- **Verb frame alternations** emerge through iterated learning due to the regularity bias. Ultimately, verb frame alternations seem to reflect world structure, which I formulated in terms of the usefulness of lexical entries. The causative-inchoative alternation, unlike the causative-reflexive and causative-reciprocal alternation, can be characterised by conditioning the usefulness of the inchoative entry on that of the causative entry, in a way similar to basic-level concepts and superordinate concepts. This yields an a-priori argument why decausativisation, and not causativisation, should exist as a lexicon operation.
- The **Lexicon-Syntax Parameter** seems to follow straightforwardly from the redundancy of lexical entries in the presence of efficient syntactic tools, where I left ‘efficient’ to be defined, and changes in its setting can be tied to grammaticalisation. I have suggested that relative frequencies of thematic forms may be used to estimate the computational cost of thematic role assignment (*zich* vs. *zichzelf*).
- Instead of postulating **frozen forms** and **inert roles**, I have shown that concept usefulness, though only qualitatively and partially defined, can be

used to predict that passives and middles never appear without a causative alternate, while reflexives, reciprocals and inchoatives sometimes do.

- **Morphological markers** may appear on either side of a verb frame alternation, not because the concept of derivational morphology is wrong per se, but because, in the case of verb frame alternations through iterated learning, there is no derivation proper. A more uniform preference of morphological markers is predicted for parameterisable operations, in which case the markers are the result of grammaticalisation.
- The **linking procedure** and its interaction with **lexicon operations** translate remarkably well into an explanation of split intransitivity through iterated learning, in terms I-thematic role alignment and L-thematic role alignment, both the regularity bias in disguise, and an economy bias. Tentatively, I have suggested that the linking procedure for **thought** and the linking procedure for **communication** may be sensitive to an entirely different range of properties, and the first may be innate while the second is acquired.
- The **Identity Constraint** emerges through iterated learning, though in two parts: I-thematic distinctness emerges through iterated learning due to the regularity bias, while L-thematic distinctness emerges as a necessity for reliable communication. The Identity Constraint through iterated learning is a property of lexical entries; I have shown that two arguments based on/in favor of an Identity Constraint at the interface between syntax and inference, are flawed.

Summarizing, in my account none of the core postulates of the Theta System I have investigated require innateness. Two innate biases suffice: economy (for explaining Theme unergatives) and regularity (for almost anything). In my account, these biases, the language acquisition bottleneck and world structure together shaped the lexicon-syntax interface.

But ‘my account’ is not the only one I have investigated. In between the lines, I have tried to show that the same set of explanations is available not only for what I have called ‘my account’, but also for yet another take of the Theta System, which maintains not lexicon operations but concept nativism. The concept nativist’s account places the burden of constrained verb frame alternations not on the lexicon but on the vocabulary. Recall that this not an a priori strange assumption to make; after all, only those innate concepts with a corresponding vocabulary entry are visible to linguists - for all we know there could be many more which the speaker never utters. The view that this assumption leads to is one of a passive, innate lexicon with an active, acquired vocabulary. The lexicon and the vocabulary together match my definition of an active lexicon, and from this perspective, too, the postulates of the Theta System I have investigated can all be explained as a consequence of two biases, the acquisition bottleneck and world structure.

The concept nativist’s account and my account are equivalent except for the concept nativism. I could have written this thesis instead on the concept

nativist's account entirely, demoting what I have been calling 'my account' to a loose collection of footnotes. Perhaps that would have been a more transparent way of presenting the same ideas, but it would have required an explicit distinction between vocabulary and lexicon everywhere throughout the thesis. Frankly, I was expecting until fairly recently that the concept nativist's account would fail when the linking procedure came into play - I had not foreseen this happy ending in which both accounts are consistent.

On the other hand, perhaps this tells us something. Perhaps the fact that two accounts can be virtually equivalent, and internally consistent, except for one fundamental assumption, suggests that this *fundamental assumption* could be taken more light-heartedly. Perhaps all lexical entries are innate to a certain degree and acquired to a certain degree, and there is no clear border between an active lexicon and an active vocabulary. When does 'acquiring concepts by means of innate constraints' change into concept nativism proper? In light of this gradient, does it really matter for theories of language acquisition whether what one acquires is a lexical entry, or a link between an innate lexical entry and a vocabulary entry? Let us call this concept-semi-nativist account of the Theta System the 'opportunistic account', and its active/passive lexicon/vocabulary a 'lazy lexabulary'.

To me - but I have had considerable discussion about this - it appears that the differences in background assumptions between the opportunistic account and the mainstream Theta System only concern (i) the separate storage of all thematic forms, and (ii) the lexical storage of linking instructions, both of which I motivated in the introduction. If, as a consequence, the opportunistic account need not postulate anything except an acquisition bottleneck on lexabulary entries, a regularity bias and an economy bias, it appears to me that this is the more parsimonious approach. However, I am at the same time aware that many of the explanations I have proposed in this thesis lack the specificity required for a fair comparison. Ultimately, a comparison of the opportunistic account and the mainstream Theta System must involve an empirical evaluation. It will be harder for the opportunistic account, with all its gradient, to yield black-or-white predictions. I suspect it will be harder for the Theta System, with all its innate rigidity, to account for gradient phenomena, such as semantic drift or fluid intransitivity.

More generally, I hope that this thesis, despite my initially radical non-nativist position and despite my aim being much more modest, has contributed to an understanding of the modes of explanation available in between the two traditional bridgeheads, nativism and non-nativism.

# References

- Arad, M. (1998). *Ucl working papers in linguistics 10 (1998)*.
- Baker, C. L. (1979). Syntactic theory and the projection problem. *Linguistic Inquiry*, 10, 533–581.
- Baker, M. C. (1988). *Incorporation: A theory of grammatical function changing*. Chicago: University of Chicago Press.
- Benua, L., & Borer, H. (1996). The passive/anti-passive alternation. paper presented at glow, athens.
- Bickerton, D. (1990). *Language and species*. University of Chicago Press.
- Borer, H. (2004). The grammar machine. In A. Alexiadou, E. Anagnostopoulou, & M. Everaert (Eds.), *The unaccusativity puzzle: Studies on the syntax-lexicon interface*. Oxford University Press.
- Borer, H. (2005). *Structuring sense - the normal course of events*. New York: Oxford University Press, Inc.
- Bowerman, M. (1982). Children's mental representation of events: some clues to structures and categories from recurrent speech errors. *Stanford Psychology Department Colloquium*.
- Carlson, G. N. (1984). On the role of thematic roles in linguistic theory. *Linguistics*, 22, 259-279.
- Carlson, G. N. (1998). Thematic roles and the individuation of events. In S. Rothstein (Ed.), *Events and grammar* (p. 35-51). Kluwer.
- Chater, N., Reali, F., & Christiansen, M. (2009). Restrictions on biological adaptation in language evolution. *Proceedings of the National Academy of Sciences*, 106, 1015-1020.
- Chierchia, G. (1989). *A semantics for unaccusatives and its syntactic consequences*. Ms. Cornell University.
- Chierchia, G. (1995). The variability of impersonal subjects. In E. Bach, E. Jelinek, A. Kratzer, & B. H. Partee (Eds.), *Quantification in natural languages* (Vol. 1, p. 107-143). Dordrecht: Kluwer Academic Publishers.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1986). *Knowledge of language: its nature, origin and use*. New York: Praeger.
- Chomsky, N. (1995). *Minimalist program*. Cambridge, MA: MIT Press.
- Chomsky, N. (1997). Language and problems of knowledge. *Teorema*, XVI/2, 5-33.

- Davidson, D. (1967). The logical form of action sentences. In N. Rescher (Ed.), *The logic of decisions and action* (p. 81-95). Pittsburgh, PA: University of Pittsburgh Press.
- Deacon, T. W. (1997). *The symbolic species: The co-evolution of language and the brain*. W.W. Norton.
- Dimitriadis, A. (2004). *An event semantics for the theta system*. Ms., Utrecht Institute of Linguistics OTS.
- Dowty, D. (1979). Word meaning and montague grammar. *Studies in Linguistics and Philosophy*, 7.
- Dowty, D. (1989). On the semantic content of the notion ‘thematic role’. In G. Chierchia, B. Partee, & R. Turner (Eds.), *Property theory, type theory and natural language semantics*. Dordrecht: Reidel.
- Dowty, D. (1991). Thematic proto-roles and argument structure. *Language*, 67, 547-619.
- Fitch, W. T., Hauser, M. D., & Chomsky, N. (2005). The evolution of the language faculty: clarifications and implications. *Cognition*, 97, 179-210.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: Bradford Books, MIT Press.
- Fodor, J. A. (2000). *The mind doesn't work that way: the scope and limits of computational psychology*. Cambridge, MA: MIT Press.
- Foster, K. I. (1979). Levels of processing and the structure of the language processor. In W. Cooper & E. Walker (Eds.), *Sentence processing: Psycholinguistic essays presented to merrill garrett*. Hillsdale, N.J.: Erlbaum.
- Gerlach, B. (2002). *Clitics between syntax and lexicon*. John Benjamins Publishing Company.
- Geurts, B. (2004). Weak and strong reflexives in dutch. In P. Schlenker & E. Keenan (Eds.), *Proceedings of the esslli workshop on semantic approaches to binding theory, nancy, france*.
- Gleitman, L. R. (1990). The structural sources of verb meaning. *Language acquisition*, 1, 3-55.
- Griffiths, T. L., & Kalish, M. L. (2007). Language evolution by iterated learning with bayesian agents. *Cognitive Science*, 31, 441-480.
- Grimshaw, J. (1981). Form, function, and the language acquisition device. In C. Baker & J. McCarthy (Eds.), *The logical problem of language acquisition* (p. 165-182). Cambridge, MA: MIT Press.
- Grimshaw, J. (1990). *Argument structure*. Cambridge, MA: MIT Press.
- Hale, K., & Keyser, S. (2002). *Prolegomenon to a theory of argument structure*. Cambridge, MA: MIT Press.
- Hasegawa, N. (2001). Causatives and the role of v: Agent, causer, and experiencer. In K. Inoue & N. Hasegawa (Eds.), *Linguistics and interdisciplinary research*. Kanda University of International Studies.
- Haspelmath, M. (1993). More on the typology of inchoative/causative verb alternations. In B. Comrie & M. Polinsky (Eds.), *Causatives and transitivity (studies in language companion series, 23)* (p. 87-120). Amsterdam: Benjamins.
- Haspelmath, M. (2007). Further remarks on reciprocal constructions. In

- V. P. Nedjalkov (Ed.), *Reciprocal constructions. 5 vols.* (p. 2087-2115). Amsterdam: Benjamins.
- Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002, 11). The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298, 1569-1579.
- Heine, B., & Kuteva, T. (2002). *World lexicon of grammaticalization*. Cambridge University Press.
- Horvath, J., & Siloni, T. (2005). *Active lexicon: Adjectival passives*. Paper presented at the Semitic Workshop, Glow 28, University of Geneva.
- Jackendoff, R. S. (1987). The status of thematic relations in linguistic theory. *Linguistic Inquiry*, 18, 369-411.
- Jackendoff, R. S., & Pinker, S. (2005). The faculty of language: What's special about it? *Cognition*, 95(2), 201-236.
- Kaplan, R. M., & Bresnan, J. (1982). Lexical-functional grammar: A formal system for grammatical representation. In J. Bresnan (Ed.), *The mental representation of grammatical relations* (p. 173-281). Cambridge, MA: MIT Press.
- Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: an experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, 105.
- Kirby, S., & Hurford, J. (2002). The emergence of linguistic structure: An overview of the iterated learning model. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (p. 121-148). Springer Verlag, London.
- Kirchner, R. (1996). Synchronic chain shifts in optimality theory. *Linguistic Inquiry*, 27:2, 341-350.
- Levin, B., & Rappaport-Hovav, M. (1995). *Unaccusativity: At the syntax-lexical semantics interface / linguistic inquiry monograph 26*. Cambridge, MA: MIT Press.
- Lidz, J. (2010). Language learning and language universals. *Biolinguistics*, 4:2, 201-217.
- Lombard, L. (1985). How not to flip the prowler: Transitive verbs of action and the identity of actions. In E. LePore & B. McLaughlin (Eds.), *Truth and interpretation*. Blackwell, Oxford.
- Marantz, A. (1984). *On the nature of grammatical relations*. Cambridge, MA: MIT Press.
- Marelj, M. (2002). Rules that govern the cooccurrences of theta-clusters in the theta system. *Theoretical linguistics*, 28.3, 357-374.
- Marelj, M. (2004). *Middles and argument structure across languages*. Unpublished doctoral dissertation.
- Oehrle, (1976). *The grammatical status of the english dative alternations*. Unpublished doctoral dissertation.
- Parker, A. R. (2006). *Evolution as a constraint on theories of syntax: The case against minimalism*. Unpublished doctoral dissertation, Theoretical and Applied Linguistics, The University of Edinburgh.
- Parsons, T. (1990). *Events in the semantics of english: A study in sub-atomic*

- semantics / current studies in linguistic series 21*. Cambridge, MA: MIT Press.
- Perlmutter, D. (1978). Impersonal passives and the unaccusative hypothesis. In *Proceedings of the fourth annual meeting of the Berkeley Linguistics Society* (p. 157-189). Berkeley Linguistic Society, University of California, Berkeley.
- Pesetsky, D. (1995). *Zero syntax: experiencers and cascades*. Cambridge, MA: MIT Press.
- Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.
- Pinker, S. (1987). The bootstrapping problem in language acquisition. In B. MacWhinney (Ed.), *Mechanisms of language acquisition*. Hillsdale, NJ: Erlbaum.
- Pinker, S. (1989). *Learnability and cognition: the acquisition of argument structure*. Cambridge, MA: MIT Press.
- Pinker, S., & Jackendoff, R. S. (2005). The faculty of language: what's special about it? *Cognition*, 95(2), 201-236.
- Piñon, C. (2001). Modelling the causative-inchoative alternation. *Linguistische Arbeitsberichte*, 76, 273-293.
- Plate, T. (2002). Distributed representations. In *Encyclopedia of cognitive science*. Macmillan Reference Ltd.
- Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.
- Pylkkänen, L. (1997). The linking of event structure and grammatical functions in Finnish. In M. Butt & T. Holloway-King (Eds.), *Proceedings of the IFL97 conference*. San Diego, CA: University of San Diego.
- Rafferty, A. N., Griffiths, T. L., & Klein, D. (2009). Convergence bounds for language evolution by iterated learning. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*.
- Reinhart, T. (1991). *Lexical properties of ergativity*. Lecture presented at the conference on lexical structure, Utrecht.
- Reinhart, T. (2000). *The theta system: syntactic realization of verbal concepts*. OTS working papers in linguistics, Utrecht University.
- Reinhart, T. (2002). The theta system: an overview. *Theoretical Linguistics*, 28(3), 229-290.
- Reinhart, T., & Reuland, E. (1993). Reflexivity. *Linguistic Inquiry*, 24, 657-720.
- Reinhart, T., & Siloni, T. (2003). *Thematic arity operations and parametric variations*. OTS working papers in linguistics, Utrecht University.
- Reinhart, T., & Siloni, T. (2005). The lexicon-syntax parameter: reflexivization and other arity operations. *Linguistic Inquiry*, 36.3, 389-436.
- Reuland, E. (2010). Imagination, planning, and working memory: The emergence of language. *Current Anthropology*.
- Riemsdijk, H. van, & E. W. (1986). *Introduction to the theory of grammar*. Cambridge, MA: MIT Press.
- Rispoli, M. (1987). The acquisition of transitive and intransitive action verb categories in Japanese. *First Language*, 7, 183-200.

- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (p. 27-48). Hillsdale: Lawrence Erlbaum Associates.
- Siloni, T. (2002). Active lexicon. *Theoretical Linguistics*, 28:3, 383-400.
- Smith, K. (2003). *The transmission of language: models of biological and cultural evolution*. Unpublished doctoral dissertation, Theoretical and Applied Linguistics, The University of Edinburgh.
- Smith, K., Kirby, S., & Brighton, H. (2003). Iterated learning: a framework for the emergence of language. *Artificial Life*, 9, 371-386.
- Sorace, A. (2004). Gradience at the lexicon-syntax interface: Evidence from auxiliary selection and implications for unaccusativity. In A. Alexiadou, E. Anagnostopoulou, & M. Everaert (Eds.), *The unaccusativity puzzle: explorations of the syntax-lexicon interface* (p. 243-68). Oxford: Oxford University Press.
- Tenny, C. (1994). *Aspectual roles and the syntax-semantics interface*. Dordrecht: Kluwer Academic Publishers.
- Zuidema, W. (2002). How the poverty of the stimulus solves the poverty of the stimulus. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems 15*. Cambridge, MA: MIT Press.