1 Introduction

1.1 Constructions and language change

This book outlines an approach to language change from the perspective of Construction Grammar. The idea of 'constructional change' will be developed as a new way of looking at diachronic processes in language. The following chapters present several case studies from the history of English in order to illustrate what processes of change come together in the development of grammatical constructions, and how the adoption of a constructional perspective can deepen our understanding of grammatical change. Naturally, an undertaking of this kind has to justify why a new term – constructional change – is introduced to group together formal and functional processes of change that affect constructions. As will be discussed, many of these are in fact quite well understood individually, and some of them are already grouped together under the heading of grammaticalization (Hopper and Traugott 2003). In what cases should we speak of constructional change rather than use established terminology? This book is an attempt to formulate a comprehensive answer to this question and to offer example analyses in which the theoretical concept of a changing, developing construction yields a shift in perspective, thus revealing issues that were previously not amenable to analysis.

The general idea of constructional change is of course not new. In fact, it has gained fairly wide currency over the past years and work on its theoretical status, similar in spirit to the present study, is underway (Bybee 2010, Trousdale 2010). Much recent work on language change, especially research with a functional orientation, focuses on the level of grammatical constructions (Bergs and Diewald 2008, Traugott and Trousdale 2010). As a theoretical framework, these works adopt some version of what is known as Construction Grammar (Fillmore *et al.* 1988, Croft 2001, Langacker 2005, Goldberg 2006). This theoretical alignment is interesting, and perhaps a little puzzling, because Construction Grammar represents a linguistic tradition that takes a synchronic and mentalist perspective. Its main objective is to find out what speakers know when they know a language and to describe this knowledge as accurately as possible (Goldberg 2003: 219). Work in language change necessarily goes beyond descriptions of a steady state, and while

1

2 Introduction

there is a great interest in cognitive aspects of language change (Hawkins 2004), many processes that are studied vastly exceed the lifetime of a single speaker, and hence go beyond the confines of a single human mind. It thus needs to be explained why and how diachronic work benefits from adopting a constructional perspective.

In synchrony, constructions are studied because of the insight that grammatical idiosyncrasies are ubiquitous: Not all transitive clauses can be passivized, not all verbs of communication can occur in a ditransitive construction, and not all monosyllabic adjectives exclusively form the morphological comparative. Exceptions, it turns out, are the rule, so that any attempt to reduce grammar to a single underlying set of rules, variable or not, is bound to fail. The slogan "Grammars contract as texts expand" (Hopper 2010) captures the idea that with every new text genre and every additional speaker, more variation enters the picture, and less of an invariant core system remains. Constructions vary across speakers and across varieties, and importantly, constructions also display category-internal variation. Speakers are assumed to keep track of this variation by retaining their experience with a given construction as a memorized collection of exemplars (Bybee 2010: 14). This means that speakers will differ in their representations of constructions. In order to describe grammar and grammatical variation adequately, it is thus necessary to determine partial generalizations, that is, generalizations that hold true for subsets of speakers in subsets of situations. To some, this emphasis on contingency may suggest that linguistic descriptions of this kind are necessarily impressionistic and vague, perhaps even unscientific. This book will try to make the case that partial generalizations can in fact be studied with quantitative rigor, by paying close attention to language use.

Partial generalizations are reflected in the usage patterns of individual constructions; taken together, these patterns form a loose, but structured, network. The concept of language as a combination of a mental grammar, containing the rules, and a mental lexicon, containing a long list with exceptions, is thus abandoned in favor of an integrated model. This is the idea of a "structured inventory" (Langacker 1987: 73), a hierarchical network that redundantly represents knowledge of language at multiple levels of abstraction, from concrete, exemplar-like representations of individual words to abstract schemas that represent syntactic relations. Each node in such a network, from the very concrete to the very abstract, qualifies as a construction and, to the extent that it displays non-predictable characteristics or other signs of cognitive entrenchment, each of these nodes merits individual description.

1.1.1 Construction Grammar, applied to diachrony

Why now has Construction Grammar become such a popular notion in the historical arena? One likely reason is that analysts of language change tend to

be very aware of the fact that the developments of many grammatical forms are highly unpredictable and idiosyncratic. To illustrate, in comparison with phonological chain shifts, which lend themselves to systemic, rule-based descriptions, the semantic developments of grammatical constructions such as modal auxiliaries, derivational morphemes, or discourse markers appear a lot more haphazard. There are of course pervasive cross-linguistic regularities (Bybee *et al.* 1994, Heine and Kuteva 2002), but once the developments of grammatical forms are studied in detail, it soon emerges that each construction has its own history. Even cognate forms that come from closely related languages and that have undergone very similar developments, as for instance English *shall* and Swedish *ska*, display fundamental differences in their respective histories and current usage (Hilpert 2008b). In the case of *shall* and *ska*, both forms have evolved to express futurity, but at the same time both forms express a range of other meanings, and the functional overlap of the two forms is fairly limited.

What unites constructional approaches to diachrony is an understanding of grammatical change as the sum of many individual metamorphoses of symbolic units. This contrasts with a perspective in which grammatical change is catastrophic and has systemic consequences (Lightfoot 1979, 1999). While the idea of wholesale grammatical restructuring thus runs counter to a constructional perspective, the assumption is often made that different constructions will change in mutually related ways. Drawing on the metaphorical conception of a grammatical space, it is often described how one construction encroaches on the territory of another one (Bergen and Plauché 2001, Rissanen 2007), or how one construction falls out of use, leaving a functional gap that is subsequently filled by other constructions (Petré and Cuyckens 2006, Petré 2010). Alternatively, the grammatical space of a language may have regions that are sparsely populated, so that a construction from another language can be borrowed into this area in order to fill that space (Heine and Kuteva 2005).

The metaphor of grammatical space is very powerful, as are the attendant assumptions that constructions compete for territory and thus strive to be dispersed evenly across grammatical space. In fact, this figure of thought underlies so much linguistic work that it is rarely made explicit, let alone challenged. Importantly though, it may not always reflect the realities of grammar. Especially explanations of grammatical developments that appeal to functional need have been severely criticized (Hopper and Traugott 2003, Tallerman 2005), as some very basic traits of grammar seem to clash with the idea of functional need. For instance, many grammatical domains are characterized by abundant layering (Hopper 1991). In English, there are quite a few constructions for the expression of future time reference (*will*, *shall*, *be going to*, *be about to*), concessivity (*while*, *although*, *despite*, *even if*, *granting*), or a skeptical stance in discourse (*actually*, *well*, *on the other hand*, *however*). These forms convey different shades of meaning, but they do not

4 Introduction

seem to be engaged in a struggle for existence. Quite to the contrary, if a
grammar has a well-populated grammatical paradigm of modal auxiliaries,
chances are that new marginal members of that paradigm will grammaticalize
as time goes on (cf. Krug 2000 on emerging modals). Further, if one form
develops a greater range of usage patterns, this does not necessarily lead to
the demise of another one. It just could be that the paradigm of constructions
as a whole is expanding. In defense of this point, it will be argued throughout
this book that *grammatical change is not a zero-sum game*.

In an extreme version then, a constructional view of language change
would hypothesize that each construction develops on its own terms and that
the change of one construction does not trigger changes in other construc-
tions. Stated in such a strong and radical way, the constructional change
hypothesis is definitely wrong. Speakers are aware of paradigmatic alterna-
tives in grammar, and if one member of a paradigm undergoes change, this
might well affect how the other members are used. Consequences may arise
in terms of both assimilation and dissimilation. The prime example of one
constructional change leading to assimilation in other constructions is per-
haps analogical leveling across paradigms, as for instance the regularization of
irregular verbs (Campbell 2004: 106). Dissimilation in constructional change
could be observed when a new construction establishes itself in a grammati-
cal domain, and an older construction from that domain changes in meaning
as a consequence, as in the scenario of encroaching constructions that was
mentioned above. Old constructions may thus exhibit characteristics of a
"donut gram" (Dahl 2000: 11), that is, they retain many of their peripheral
meanings whereas their erstwhile central meaning has been taken over by
a newly emerged construction. Haspelmath (1998) discusses the semantic
spectrum of old present tense constructions as a typical example of this phe-
nomenon. In both analogical leveling and the formation of donut grams, one
constructional change brings about, or at least facilitates, change in another
construction. However, care should be taken not to assume a causal relation
between any two constructional developments unless there are good indepen-
dent reasons for doing so. In the case of analogical leveling, these reasons are
overwhelming. In the case of donut gram formation, the evidence supports
a causal interpretation as well. In other cases, the evidence may not be quite
as clear. In the empirical studies of this book, the independence of a given
constructional change is initially taken as a null hypothesis. Alternatives are
considered on the basis of strong evidence that suggests a relation between
two or more constructional developments. Such evidence could for instance
be that the usage of two forms is subject to identical constraints that change
in similar ways over time.

1.1.2 Studying constructions in language change

As the notion of constructional change presupposes the idea of a construction,
that term deserves a proper introduction. Constructions can be informally

characterized as the 'building blocks' of grammar, as they can be combined to form phrases, sentences, and larger stretches of discourse. Most commonly, constructions are understood as signs, that is, symbolic pairings of a form and a meaning that display structural idiosyncrasies or a high level of entrenchment (Croft 2001: 16, Langacker 2005: 105). For the present purposes, the following definition from Goldberg (2006: 5) is adopted:

> Any linguistic pattern is recognized as a construction as long as some aspect of its form or function is not strictly predictable from its component parts or from other constructions recognized to exist. In addition, patterns are stored as constructions even if they are fully predictable as long as they occur with sufficient frequency.

This definition is deliberately very broad, including abstract syntactic patterns such as pseudo-clefts and the ditransitive construction, idiomatic expressions such as *blow the whistle on NP*, and concrete lexical items such as the English noun *dog*. While these forms differ in terms of how abstract and schematic they are, each of them connects a particular form with a particular meaning, thus forming a symbolic unit. A broad definition of the term 'construction' is needed because, as was indicated above, linguistic knowledge is taken to be a structured inventory of symbolic units in a network, and nothing else in addition.

Traugott (2008: 236) proposes the terms 'macro-,' 'meso-,' and 'micro-construction' to distinguish between symbolic units at different levels of schematicity. Macro-constructions are understood as constructions at a fairly abstract level; Trousdale (2008a: 170) offers the illustrating example of the English determiner construction, in which a nominal structure co-occurs with a second, determining element, such as a demonstrative, a possessive pronoun, a prepositional phrase with *of*, or a determiner proper. Macro-constructions can hence be viewed as umbrella constructions that comprise several different lower-level schemas. The next lower level is the one of the meso-construction, which is also a conglomerate of several lower-level constructions. An example of a meso-construction would be the English possessive construction (Trousdale 2008a: 170) that unites the two English genitives, the *s*-genitive and the *of*-genitive. Trousdale (2008a: 169) proposes that meso-constructions may be situated at different levels of abstraction, so that one meso-construction may actually comprise several other, less schematic meso-constructions. A micro-construction would be a formally specified pattern such as *proper N's N*, which is then instantiated in language use by constructs such as *Frank's jacket*. In the present work, the terms 'macro-,' 'meso-,' and 'micro-construction' are used heuristically to talk about constructions at different levels of schematicity. The use of these terms does not imply that the constructional network is viewed as discretely divisible into, say, three or four levels of abstraction. Rather, it is assumed that constructions are mentally represented along a continuum of schematicity.

6 Introduction

To speak of macro- and meso-constructions thus contrasts higher- and lower-level generalizations somewhere along this cline.

Diachronically, both form and meaning of a construction are subject to variation and change. Changes in form pertain to the phonology or morphosyntax of a construction, whereas meaning changes affect its semantic and discourse-pragmatic characteristics. Each of these types of change is amenable to quantitative corpus analysis, which is the main analytical approach that is taken in the present work. The defining characteristic of quantitative corpus linguistics is that questions are formulated in such a way that frequencies, as gathered from textual data, can provide answers (Stefanowitsch 2005). Frequency values that are of interest include, amongst others, the text frequency of a construction, its type frequency, its frequency of co-occurrence with other linguistic elements, and the relative frequencies of different structural and functional variants of the construction. Which one of these frequency types is to be studied depends on the kind of question one would like to answer.

The research questions of quantitative corpus linguistics are not in themselves specific to the discipline. On the contrary, the aim is to address questions of as general linguistic interest as possible, and to make the greatest use of frequency data for this purpose. Changes in form and meaning of a construction can be studied through frequency measurements of its variants: An important concept in the present work is thus the idea that constructions are not fixed, but flexible, displaying formal and functional variation. The idea that constructions are mentally represented as clouds of exemplars (Bybee 2010: 14) is helpful in this regard. The instances of an exemplar cloud will vary along multiple dimensions. This variation can be assessed on the basis of corpus data, and changes in this variation can be tracked diachronically. One variant of a construction, that is, one subtype of the construction that has a certain configuration of features, may become more frequent over time, as other variants become less frequent. Entirely new variants may develop. One constructional subtype may develop new extensions, other subtypes may converge on a single form or function.

As constructions can vary in multiple respects, constructional change rarely proceeds just along a single dimension. Rather, many formal and functional aspects of a construction change simultaneously. Not all changes will maintain the same direction over time, and as constructional features can change in mutually dependent ways, it quickly becomes fairly difficult in practice to comprehend a given constructional change just by manually inspecting examples from different historical periods. Keeping track of multiple features and the frequencies of their values over time requires some statistical machinery. The present work showcases some of the methods that are currently available and that can aid the study of grammatical change. Importantly, analyses of this kind do not obviate the need for close philological analysis. In practice, the process of gathering and coding corpus data still

involves long hours of manual inspection, which remains the most important part of the work. Looking at raw corpus data will enable the analyst to formulate meaningful hypotheses about the change, which can then be investigated quantitatively, and hopefully interpreted in a way that relates to issues of linguistic theory.

A few words need to be added on the distinction of change in linguistic structure as opposed to change in frequency. Diachronic corpus data are probably most useful for the documentation of structural changes over time. If, for instance, one is interested in the development of a construction such as the English perfect, corpus data can be used to document when the past participle started to occur next to the verb *have*, or when it ceased to show morphological agreement with the object (Traugott 1972). But even in the absence of structural change, a construction may undergo changes in absolute frequency, relative frequency, or type frequency. Absolute frequency (or token frequency) measures how often a construction occurs within a fixed amount of running text. Trivially, a construction may become more or less frequent over time. Relative frequency measures how often a construction occurs in comparison to some alternative construction. A change in relative frequency could be that a construction surpasses another one in frequency or becomes surpassed itself. Type frequency (or dictionary frequency) measures the number of different lexical elements that enter a constructional schema. For instance, the type frequency of the English derivational suffix *-ism* can be determined by consulting a reverse dictionary and extracting the number of different words with that ending. A continuing increase in type frequency over time, which can be investigated on the basis of diachronic corpus data or a dictionary such as the OED, is evidence for the productivity of a construction. A construction may also spawn fewer and fewer new types until its productivity comes to a complete halt.

It is argued here that these frequency changes, and others that could be imagined, are no less indicative of constructional change than are developments in meaning or the phonological and morphosyntactic substance of a construction. However, much current research on constructional change relegates issues of frequency to a marginal position or limits itself to reporting basic descriptive statistics. Far and away the most common use of frequency in studies of constructional change is the measure of absolute frequency over time, often normalized to a measure of tokens per million words or token per 10,000 words (cf. the contributions in Lindquist and Mair 2004, Lenker and Meurman-Solin 2007). Such values are reported to document the rise or fall of a particular construction, and usually some qualitative discussion is offered concerning irregularities in the particular shape of the curve. While there is nothing to be said against such practice, it stands to reason that, potentially, much more could be done with frequency data. The use of selected exploratory and inferential statistics can yield answers to very basic questions concerning developments of absolute frequency. One

8   Introduction

such question is whether an observed trend is strong enough to constitute
reliable evidence. Are we looking at a chance fluctuation in the data or at a
statistically significant change? Also, manual inspection of a frequency curve
often suggests different developmental phases as the slope changes. It would
be useful to have an objective way to distinguish between time periods that
are characterized by different frequency developments. Methods to address
these and other questions concerning diachronic changes are in fact available,
so that constructional change can be studied quantitatively on the basis of
many different frequency measures, not just absolute frequencies. Some of
these methods will be further discussed and exemplified in Chapter 2.

In this context, it is important to reiterate that frequency measurements,
either concerning the text frequency, type frequency, or the relative fre-
quency of structural variants of a construction, are not merely used here to
establish that a construction has changed over time. That would be trivial,
and in most cases, frequencies are not even needed for that. Frequency val-
ues should be used to support qualitative accounts of how a constructional
change unfolded over time, not only with regard to overall text frequency,
but also with an eye to all relevant grammatical traits of the construction
under investigation. A constructional approach to language involves the
basic commitment to recognize the idiomatic and non-predictable aspects
of constructions (Fillmore *et al.* 1988), and the same scrutiny should be
employed when analyzing constructional change. The art of approaching
constructional change with a corpus-based methodology is thus to oper-
ationalize theoretical questions about language change in such a way that
frequencies drawn from corpus data can provide evidence for or against a
given hypothesis, or facilitate the exploratory description of a constructional
change.

1.2 Defining constructional change

Before going further, it is necessary to define the notion of constructional
change in more detail. If the definition of a construction that was offered
above were to denote only grammatical constructions, the term 'construc-
tional change' would be largely synonymous with grammaticalization, that
is, "the change whereby lexical terms and constructions come in certain
linguistic contexts to serve grammatical functions, and, once grammat-
icalized, continue to develop new grammatical functions" (Hopper and
Traugott 2003: xv). Grammaticalization can be understood as a bundle of
processes that act in concert on both form and meaning of a construction.
Phonological reduction, reanalysis, univerbation, semantic bleaching, and
pragmatic strengthening are five particularly prominent developments that
grammaticalizing constructions typically undergo, either in combination or
individually. Constructional change is more encompassing than the changes
that characterize grammaticalization. Specifically, it includes processes of

lexicalization, processes of syntactic change that do not instantiate gram-
maticalization, processes within derivational morphology, and processes of
frequency change that are unrelated to grammaticalization. Each of these
process types is briefly discussed below.

1.2.1 Where constructional change goes beyond grammaticalization

As outlined above, the definition of constructions that is adopted here explic-
itly includes lexical items, such that all manners of lexical semantic changes
fall under the purview of constructional change (Noël 2007, Trousdale
2008b). Some trajectories of semantic change are strongly restricted to lexis.
For instance, semantic narrowing from a class to a subtype, as in English
hound, is never observed with grammaticalizing constructions. Grammatical
constructions may cease to express some meanings from a given spectrum,
but systematic semantic narrowing does not occur (see Campbell 2004: 256
for a different view). Another example is the development of evaluative
meaning, which is observed with lexical but not with grammatical elements.
The processes of amelioration, as in *nice* (from 'foolish' to 'pleasant'), and
pejorization, as in *silly* (from 'deserving pity' to 'foolish'), primarily apply to
the meaning of lexical items. Constructional change thus includes processes
that remain squarely within the field of lexis. Changes in lexical meaning
constitute an interesting domain of constructional change, but as the focus
of this book is on grammatical changes, they will not be discussed at great
length here.

Narrowing the focus to the domain of grammar and grammatical change,
it can still be argued that constructional change is a broader concept than
grammaticalization. Not everyone takes this stance. Wiemer and Bisang
(2004: 4) call for a re-definition of grammaticalization that makes it into an
all-encompassing notion of grammatical change:

> If grammar is understood as a system of more or less stable, regular
> and productive form-function-mappings, the field of grammaticaliza-
> tion . . . is to be extended to all the processes involved in the diachronic
> change and the emergence of such systems.

The present work shares Wiemer and Bisang's point of view that a tra-
ditionally defined concept of grammaticalization does not account for all
phenomena that are of interest in the development of grammatical forms
and systems. However, the position taken here is that rather than further
diluting a definition of grammaticalization that already faces the criticism
of not being concise enough (Newmeyer 1998, Campbell 2001), a notion of
constructional change should be developed in contrast to other, either more
general or more specific, processes of language change. In order to make
this argument, it first needs to be clarified why the notion of constructional
change casts a wider net than the term grammaticalization.

10   Introduction

Many definitions of grammaticalization exclude word order changes
(Hopper and Traugott 2003: 24), some of which would instantiate con-
structional change. To give a first example, a phenomenon not commonly
seen as an instance of grammaticalization is the loss of the verb-second con-
straint (V2) in the history of English (Kemenade 1987, Kroch and Taylor
1997, Roberts 2007, Los 2009). Indeed, none of the usual criteria for gram-
maticalization (Hopper 1991, Lehmann 1995) apply in this case. From the
perspective of Construction Grammar, the loss of V2 could be seen as the
deterioration of a general syntactic construction. Over time, V2 lost its gen-
eral applicability, and in Present-Day English it survives only in a handful of
stylistically marked 'V2-residual' constructions (Vikner 1995) such as loca-
tive inversion (*Into the room walked Noam Chomsky*) or negative inversion
(*Never would I leave you*). There are thus syntactic changes that are not
straightforwardly accommodated in the framework of grammaticalization
but arguably still instantiate constructional change.

Second, a linguistic domain that straddles the distinction of grammati-
calization and lexicalization, fitting only awkwardly into either of the two,
is the field of derivational morphology (Himmelmann 2004). Many English
derivational morphemes, such as *-hood*, *-dom*, or *-ship* demonstrably go back
to autonomous lexical items (Trips 2009) and thus motivate a discussion of
their history in terms of grammaticalization. But at the same time, since the
main purpose of derivational suffixes is the creation of new lexical items,
their role in lexicalization seems at least equally central (Lehmann 1989).
A view from Construction Grammar can provide a unifying perspective in
this situation (Booij 2010). If the pattern of a noun followed by *-hood* is
viewed as a construction, then the birth and development of the *N-hood*
construction can be described as one process of constructional change. This
process would include the grammaticalization of Old English *had* 'state,
condition' into a suffix, the formation of lexical items with that suffix, and
subsequent changes in the productivity of the suffix. Crucially, these changes
might include decreases in productivity. The process of host-class expan-
sion (Himmelmann 2004), which is typically exhibited by grammaticalizing
forms, only captures increases in productivity. A broader perspective, as
envisioned by Wiemer and Bisang (2004: 4), but not under the heading of
grammaticalization, can thus act as an umbrella for changes that transcend
the boundaries of pre-established fields of investigation.

1.2.2   Where grammaticalization goes beyond constructional change

So does constructional change simply subsume grammaticalization? Is it a
cover term that includes everything that is known as grammaticalization
and beyond that, a more or less random collection of lexical, syntactic, and
morphological changes? The short answer to these questions is *no*. Despite