1 *Introduction*

Nominalizations are complex nouns that are derived from verbs, adjectives, and other nouns. As described in Bauer et al. (2013), English has many ways of deriving complex nouns: affixes such as *-er* (*writer*), *-ant* (*accountant*), *-ist* (*accordianist*), *-ee* (*employee*), *-ster* (*hipster*), *-eer* (*conventioneer*), and *-meister* (*trashmeister*), whose main use is to derive personal or participant nouns; affixes such as *-ing* (*writing*), *-ation* (*destruction*), *-ment* (*amusement*), *-al* (*recital*), and *-ure* (*closure*), whose main use is to derive nouns that denote events and results; affixes such as *-age* (*assemblage*) and *-ery* (*pottery*) or *-ity* (*purity*) and *-ness* (*happiness*), whose main use is to derive collective or abstract nouns. English also has a productive process of conversion from verb to noun, the results of which can be used with a wide range of readings: *cook* – an agent noun, *attack* – an event or result noun, *wrap* – an instrument, *nosh* – an inanimate patient, *dump* – a location noun, and so on.

English nominalizations have been extensively discussed by both morphologists and syntacticians. In the generative tradition, the study of event and result nominalizations (henceforth E/R nominalizations) has been important from the start. Beginning with Lees (1960) and Chomsky (1970), an enormous literature devoted to the syntactic analysis of E/R nominalizations has accumulated within the tradition of mainstream generative grammar, including Lebeaux (1986), Grimshaw (1990, 2011), Roeper (1993), Snyder (1998), Alexiadou (2001, 2011), Newmeyer (2009), Harley (2009), Sichel (2010), Roy and Soare (2011), Fabregas (2012), Borer (2013), and Bruening (2013) among many others. Pustejovsky (1995, 1998) looks at E/R nominalizations through the lens of computational lexical semantics. The subject of E/R nominalizations has also been an important one for morphologists as well, as works by Bierwisch (1990/1991), Lieber and Baayen (1999), Melloni (2007, 2011), and Fradin (2011), among others

show. Although there is somewhat less literature by syntacticians on personal nominalizations, the work of Levin and Rappaport Hovav (1988), Rappaport Hovav and Levin (1992), Alexiadou and Schäfer (2010), Bowers (2010), and Borer (2013) comes to mind. Among morphologists, personal nominalizations have been discussed in the work of Booij (1986), Bauer (1987, 1993), Ryder (1999), Heyvaert (2001), Booij and Lieber (2004), and Lieber (2004), among others, but there is less theoretical discussion of personal affixes such as *-ant* or *-ist* and virtually none of exotic nominalizations in *-eer, -ster*, or *-meister*. Abstract nominalizations have not figured prominently in the work of syntacticians (although see Roy 2010 and van Hout et al. 2013) but have been of interest among morphologists from the start (for example, Aronoff 1976, Anshen & Aronoff 1981, Baeskow 2012, Arndt-Lappe 2014 on *-ness* and *-ity*). Lieber (2004) gives a brief analysis of collectives such as *-ery* and *-age*. Trips (2009), Lieber (2010a), Aronoff and Cho (2001), and Baeskow (2010) look at denominal complex nouns in *-dom, -ship*, and *-hood*. The only recent work that takes on the entire range of nominalizations in English is Bauer et al. (2013), which is largely a descriptive work that begins to reveal the issues I will raise here but does not attack them in theoretical terms. Thus far, no one has taken on the task of analyzing the full range of nominal derivation and trying to account for the complex relationship between form and meaning that we find in that domain.

Not surprisingly, although both syntacticians and morphologists have been interested in nominalization, they have not always asked the same questions. Morphologists have been primarily interested in forms: what are the affixes used to derive nominalizations of various sorts; what are the rules that govern them and how productive are they; in what way do they compete with each other? Morphologists, myself included, have also been concerned with affixal polysemy but seem to have concentrated primarily on the polysemy of personal affixes such as *-er* and *-ee* (Booij 1986, Booij & Lieber 2004, Lieber 2004). Less attention has been paid to the ambiguities displayed by affixes such as *-ation*, *-ment*, *-al*, and *-ing*, or by conversion (but see Melloni 2011 for discussion of comparable processes in Italian).

Syntacticians have been less concerned with the formal details of derived nouns (productivity, competition among affixes) and more interested in the relationship between sentences and the noun phrases (or determiner phrases) in which nominalizations occur: what arguments can occur or must occur to get what reading? and what verbs are allowed in

one configuration or the other? In recent years, much of the syntactic literature has concentrated on the ambiguity that E/R nominals show between an eventive reading (*the instructor's examination of the students*) and so-called result readings (*the examination was difficult/three pages long*). In many of these analyses, different readings are associated with different argument structures (Grimshaw 1990) or different underlying syntactic structures and derivations (Alexiadou 2001, Harley 2009, Borer 2013). Neither morphologists nor syntacticians have studied the full range of data pertaining to nominalizations or the intricacies of polysemy that nominalizations display.

My overall goal in this book is to rectify this state of affairs. Specifically, I intend to consider the full range of nominalizations, including the commonly discussed E/R and personal nominalizations, as well as collective and abstract nominalizations and a few seldom-discussed areas of nominalization that I will add as we go along. Second, I will try to establish on the basis of corpus data the full range of readings available to various kinds of nominalizations in various syntactic contexts. Finally, I will try to model within the lexical semantic framework of Lieber (2004, 2006, 2009, 2010b, 2015), henceforth LSF, the ways in which speakers arrive at or build those readings.[1] I will try to show how tangled the interrelationships are among various types of nominalizations, and how complex and labile the readings are that are available to them, and yet how simple the mechanisms might be that give rise to this wide range of readings.

Let me illustrate what I mean briefly by the range of readings that are frequently available for complex nouns. At first, it seems possible to start with three broad categories of nominalizations – E/R, personal/participant, and abstract/collective – with distinct sets of affixes occurring under each rubric. Of course, it is well known that nominalizations derived with suffixes such as *-ation* are systematically ambiguous between an event reading (*The professor's examination of the student was thorough*) and a so-called result reading (*The examination was two pages long*), and much has been written trying to explain that ambiguity. But it seems, as argued by Melloni (2011), that there is not a single "result" reading for nouns in *-ation, -ment, -ing*, or nouns derived by conversion. Rather, non-eventive readings can include products (*construction*), locations (*reservation*), measures (*pinch*),

[1] I have resisted in previous publications giving my theoretical framework a name and an acronym, but it is clumsy to keep referring to "the lexical semantic framework of Lieber (etc.)." So I take a plunge here in giving the framework a name – the Lexical Semantic Framework – and the accompanying acronym LSF.

paths (*descent*), and even agents (*administration, cook*) or instruments (*clip, fastening*), in addition to results. Indeed, individual E/R nouns like *construction* can have three or four possible readings, depending on syntactic context and other factors. And it has long been noted in the literature that personal nouns derived with the affixes *-er, -ant*, and *-ee* can also have a variety of readings and that those readings overlap in complex ways. Of course, they can not only have agent and instrument interpretations (*writer, printer, accountant, accelerant, attendee*) but also patient interpretations (*employee, loaner*), not to mention measure (*fiver*), means (*stroller*), and location (*diner*) interpretations. Indeed, the same *-er* noun can be used as an agent (*shooter* = a person who shoots), an instrument (*shooter* = a gun), or a patient (*shooter* = the thing which is shot).[2] Nouns in *-ee* typically denote patients but can sometimes denote agents (*attendee, escapee*). Derived nouns that typically express collectives (*pottery, acreage*) can be found in contexts where they have E/R readings (*the media's coverage of the tragedy, the mayor's bribery of the officer*) or location readings (*orphanage, nunnery*), among others. Such overlap cannot be dismissed as rare, exceptional, or even as the random effect of lexicalization – this sort of chameleon-like behavior is both productive and pervasive, as I will try to show. And given its pervasiveness, it raises many questions.

As this book progresses, I will begin to articulate the many difficulties this pervasive polysemy presents for morphological and syntactic theory. The copious syntactic literature on E/R nominalizations has led to a dizzying array of claims concerning what formal means of nominalization are attested with what kind of complements and modifiers with what kinds of interpretations. Some claims can be traced through the literature from its beginnings in Chomsky (1970) and Grimshaw (1990) to more contemporary work on nominalization such as Borer (2013). Other claims appear sporadically in one work or another. Counterexamples to previous claims crop up here and there (see especially Newmeyer 2009), oftentimes not noticed or attended to in subsequent work. Added to the problem is the apparent tendency of syntacticians to read and respond primarily to literature by other syntacticians and of morphologists to read and respond primarily to the literature of other morphologists. And with few exceptions, most of the claims are based on native speaker acceptability judgments, as has been the tradition for decades among generative linguists.

[2] See Section 4.2 for corpus examples to illustrate this point.

My contribution to this debate will be to use the tools of corpus linguistics, as exemplified in Bauer et al. (2013), to probe various claims, and thereby to try to put the study of nominalizations on a sounder empirical basis. Not surprisingly, it will turn out that intuitions about the forms and possible readings of complex nouns are often surprisingly unreliable: morphological and syntactic configurations that theorists, on the basis of intuitions, have deemed unacceptable often turn out to be easy to find and quite unproblematic in ordinary contexts. To the extent that many patterns that have been claimed to be unacceptable can be found in corpora, the theories that have been built on that data are undermined. So one goal of this book is to take a broader look at the data and try to establish what we need to explain.

My contribution will be theoretical as well as empirical. In the latter half of the book, I will argue that with a number of small modifications to be introduced in Chapters 5–8, LSF will allow us to model the way in which this complex web of polysemy arises. Briefly, LSF is a framework in which the lexical semantic representation of both simplex morphemes and affixes consists of two parts. The first is the skeleton which is made up of semantic functions and their arguments that are hierarchically arranged. Functions consist of a highly constrained set of features that allow us to characterize those aspects of lexical meaning that are relevant to the syntax; these features are simple, primitive, chosen from a universal pool of semantic features, and are such that they may be used to cross-classify lexical categories (nouns, verbs, adjectives, and so on). Skeletons are stable from one speaker to the next, but critically for what I will argue here, they may be underspecified in a number of ways. The second part of the lexical semantic representation is the body, which consists of two parts. One part comprises random bits of encyclopedic information that may vary from one speaker to the next. The other is more systematic and consists of features that might be syntactically active (and therefore part of the skeleton) in some languages, but not in the language at hand. Affixes have skeletons, just as simplex morphemes do, although they may lack semantic bodies. Derivational affixation involves the subordination of the skeleton of a base to that of an affix, with subsequent referential integration that is effected by the Principle of Coindexation.

My analysis of nominalization will begin with the theory of lexical semantic representations of Lieber (2004) but will extend that framework to look at the interplay between semantic underspecification of complex words and the resolution of that underspecification in syntactic context. My leading idea is similar to one voiced in Hanks (2013: 65). Hanks raises the question of whether words actually have determinate meanings and answers in the negative:

"The proposal here is that, strictly speaking, words in isolation have meaning potential rather than meaning, and that actual meanings are best seen as events, only coming into existence when people use words, putting them together in clauses and sentences."[3] My argument will be that LSF is well suited to operationalizing this idea. In subsequent chapters, I will work out in some detail the structures that provide the semantic potential for complex words and the mechanisms, which I will call Feature Value Matching and Contextual Coercion, by which specific readings of those words are realized in specific syntactic contexts.[4] In effect, what I will be arguing is that nominalizations do not have fixed meanings, but that they can take on a variety of readings by virtue of their sparse lexical semantics and the filling in of their representations in contexts.

Taking a panoramic view of nominalizations in terms of the range of forms, the range of readings, and both the morphological and syntactic contexts in which they occur will allow me to argue that the range of interpretation available to one kind of nominalization is inevitably influenced and shaped by the range of other nominalizations that are available to speakers of a language as well as by the contexts in which those nominalizations are deployed. In other words, one of the central claims of this book will be that nominalizations exist within a kind of derivational ecosystem where everything bears a relation to everything else.[5]

Let me be clear at the outset that the "derivational ecosystem" is a metaphor. Metaphors, as Ricoeur (1975/1977: 87) argues, are not merely verbal ornamentation but neither are they scientific models.[6] Rather, a good metaphor draws us to see something in terms of something else that is superficially unlike it; the former is what I.A. Richards (1936) calls the "tenor" and the latter the "vehicle." The tenor (for us derived nouns) and the vehicle (for us, the notion of an ecosystem) intersect in terms of some features and not others, and the metaphor serves as a filter that

[3] A precursor of this idea might be seen in some remarks of Benveniste (1966: 39), who sees the value of words as signs as being only a part of what they become in the syntagmatic context.

[4] The mechanisms by which skeletal underspecification is resolved might be seen as similar in spirit to mechanisms made available within the Generative Lexicon framework of Pustejovsky (1995, 1998), specifically what Pustejovsky (2011: 1411) calls type matching and accommodation subtyping.

[5] A disclaimer: what follows in this book is not meant in any way to be related to the branch of linguistics that is known as "ecolinguistics" as presented in Steffensen and Fill (2014) and the references cited therein.

[6] Here, Ricoeur echoes Max Black (1962).

allows us to see derived nouns in a different and potentially new light. By seeing derived nouns in a different way, we are then led to analyze them in a different way.

In what follows, I hope that the metaphor of the derivational ecosystem will draw attention to the ways in which the readings of complex nouns adapt to and are shaped by the semantic contexts in which those nouns are deployed and by the other nouns that are available to express a needed reading. I will argue that my theoretical treatment using LSF allows for this adaptability and indeed predicts that it should exist, but it should be kept in mind that LSF is nevertheless a formal theory within the general rubric of generative grammar. It is not an "ecological" theory in any sense, whatever that may mean. The metaphor simply helps us to see what the theory needs to do, but the metaphor is not itself a theory of nominal meaning. Briefly, the metaphor helps us to see two different facets of nominalization in English.

First, if we think of areas of meaning (agentives, collectives, and so on) as habitats and morphological types (particular affixes, conversion) as the organisms that exploit (or express) them, we are led to think about the ways in which forms compete in certain semantic domains and the ways in which certain semantic domains are better served by the formal morphological resources of a language than are others.[7] Some semantic habitats have several morphological types that "inhabit" them. We have, for example, lots of ways of deriving agent nouns. Interestingly, other semantic habitats are barely populated at all; there is no particular affix, for example, that forms nouns that mean "thing or stuff that has been verb-ed." One of the questions I raise in this book is what happens when there is an area of meaning which is largely uninhabited – that is, where there are no morphological types whose primary function is to express that meaning. My answer is that morphological types often expand their territories, and that different morphological types may be deployed to cover those underexploited semantic habitats under different conditions. We can take this first interpretation of the metaphor as focusing on a paradigmatic dimension: how do particular forms fit into the semantic niches that need to be expressed? Of the available means of derivation that we have, which do we choose to express a particular reading? This aspect of the metaphor will be highlighted in Chapter 4 when we look in detail at the various referential readings that derived nouns are subject to.

[7] Think of cattle and antelopes competing for grassland.

Taking the ecological metaphor in a slightly different direction, we are led to think about the way in which the syntactic and even discourse context in which a nominalization occurs pushes us toward one reading or another out of whole range of potential readings. In this sense, the syntactic or discourse context is like the ambient environment, and the complex noun like an organism that adapts to and is shaped by that environment.[8] This interpretation of the ecological metaphor takes a more syntagmatic perspective, encouraging us to see the shaping of the meaning of complex nouns in their larger syntactic and discourse contexts. This aspect of the metaphor comes to the fore in Chapters 5–7, where we look at the actual formal representations of derived nouns in LSF and at the ways in which context allows us to fix aspects of their meaning that are left lexically underspecified.

Neither interpretation of the metaphor is a perfect fit, of course; that's the nature of metaphor. In the end, affixes are not organisms like cows, antelopes, or finches; semantic categories of affixes (agent nouns, patient nouns) are not habitats like islands or grasslands; syntactic contexts are not the Galapagos Islands. But to the extent that the metaphor allows us to see that nominalizations do not have rigid denotations and to model how we arrive at their highly flexible meanings, I hope that it proves useful. For readers that are bothered by the ways in which the ecological metaphor does not work, I think that the analysis that I offer in this book nevertheless has merit.

The metaphor of a derivational ecosystem has linguistic precursors, both in the Saussurean tradition (Saussure 1916/1983) and in semantic field theory (Lehrer 1974, 1993, Kittay 1992). The notion of derivational ecosystem has its roots in the Saussurean notion of "value." What Saussure means by "value" is, roughly, that the sign is not simply a function of the signifier and the signified but is characterized as well (or as some would have it, exclusively) by its position with respect to other signs. This means that the value of a sign is not fixed but may shift, depending on where it finds itself in relation to fellow signs. Semantic field theory applies the Saussurean notion of "value" to lexical domains – color terms, verbs of motion, words for utensils, and so on. The words that occupy a lexical domain can be seen as deriving their meanings, at least in part, from their relationship to other items in the same domain. Adding or subtracting a new word in a particular domain requires a concomitant shrinking or expansion of the meanings of words already in that domain. My notion of the derivational ecosystem is related to that of the semantic field, although the domain in question is not a simple lexical domain,

[8] Think of Darwin's finches in the Galapagos Islands.

but rather an entire derivational domain – the overall system for deriving complex nouns in a language.

One thing that I will not try to do in this book is to argue against various syntactic accounts of nominalizations. Much recent work on nominalizations has been situated in syntactic frameworks such as Distributed Morphology (Alexiadou 2001, 2011, Harley 2009, Sichel 2010, Bruening 2013) or the Exo-skeletal model of Borer (2013), which claim that morphology is syntax, that vocabulary items (not morphemes) have only encyclopedic properties but no category or morphosyntactic properties, and that the properties of nominalizations can be accounted for by a series of functional projections that host affixes and trigger movement of various sorts. There are two reasons why I will not confront these accounts directly. The first is that such accounts have tended to concentrate solely on the analysis of E/R nominalizations to the exclusion of personal, participant, collective, abstract, and other nominalizations. Interesting though they are, they fail to look at the big picture. Second, and more importantly, to the extent that the data I present undermine the empirical claims on which these theories are based, they do not present viable alternatives to the lexical semantic account I give here. This is not to say that such accounts could not be modified to account for the full set of data that I will set out in what follows, just that there are no extant syntactic accounts that do so.

I have made the choice in this book to concentrate exclusively on data from English. There are a number of reasons for this. First is that many of the claims about properties of nominalizations have originated on the basis of data from English, although of course those claims have been extended and explored for many other languages. Because the origins of claims about nominalizations lie in works like Lees (1960), Chomsky (1970), and Grimshaw (1990), I would need to discuss English in detail in any case. Second, the territory I hope to cover is large. It would be difficult to cover the entire range of nominalizations in two or more languages with any degree of thoroughness. However, the main reason for concentrating on English is that arguments will frequently hinge on fine nuances of meaning: what reading(s) can a given nominalization take on in what contexts? I do not mean to claim that such nuances of meaning are only available to native speakers. Rather, I believe that I don't have good enough command of any other language to be sensitive to such nuances in a language other than my own. The arguments will also hinge on finding "live" examples of rather specific sorts for which a very large corpus will be necessary. For this, I rely largely on the Corpus of Contemporary American English, which I refer to

henceforth as COCA (Davies 2008). I have consciously tried not to rely on my intuitions about grammaticality or acceptability for reasons that I will elaborate on in Chapter 2 but take the attestation of a pattern to indicate acceptability. I do not rule out the use of intuition entirely, however, as of course I must rely on my intuitions to know what nuance of meaning is intended by any specific textual example.

In Chapter 2, I will go into some depth about the terminology I will use and my methodology in gathering examples. I will try to be clear from the outset about the terms I will use for morphological forms, for different potential readings of forms, for syntactic structures in which those can be found, and for syntactic diagnostics for various readings. Part of the difficulty in absorbing the literature on nominalization is the proliferation of terms that can be found, with a variety of terms frequently being deployed for what appear to be the same concept. With respect to methodology, I will describe how I make use of corpora, how I search for relevant examples, and how I view rare examples. My methodology is largely that used in Bauer et al. (2013) and is extensively justified in that work as well.

Chapters 3 and 4 of this book are devoted to data. Chapter 3 will look in detail at the claims that have been made in the syntactic literature regarding E/R nominalizations and consider the extent to which those claims can be supported by examples extracted from corpora. If claims are accurate, we would expect to find corpus examples of those patterns. We would similarly not expect to find examples of patterns that have been judged to be unacceptable. Of course, since it is impossible to search corpora absolutely systematically and exhaustively, we can never be certain that there are "no examples" of a phenomenon. Inability to find a pattern may be suggestive of its ungrammaticality but does not guarantee that a pattern is ungrammatical. What is more important for my purposes then is finding attested examples of patterns that have been claimed to be unacceptable. To the extent that I do find such examples, this truly changes the landscape that theorists will need to account for. I will try to show in Chapter 3 that the landscape really is quite different than we have thought for some time.

Chapter 4 will broaden the discussion beyond E/R nominalizations to a very wide range of other nominalizations, again concentrating on the formal means available and the range of readings that can be expressed. I will begin by surveying the majority of the nominalizing affixes in English and illustrating by means of corpus examples the various readings they can convey. What will emerge is that the relationship between form and meaning/reading is very far from one-to-one: morphological forms or types may not only have primary