

# Chapter 1

## Introduction

This dissertation is devoted to the language learnability problem known as *Baker's paradox* (Pinker 1989, in reference to Baker 1979). It comprises three premises:

1. Productivity
2. Arbitrariness
3. No negative evidence

To put it briefly, the paradox arises if one assumes that there are arbitrary exceptions (*arbitrariness*) to productive generalizations (*productivity*), and that negative evidence is not available to the language learner (*no negative evidence*). The purpose of this dissertation is to examine the logical and the empirical foundations of Baker's Paradox, by analyzing its logical structure and implications for learning, and questioning its empirical assumptions about the linguistic knowledge to be acquired. The central claim is a denial of the second premise: arbitrariness.

The arbitrariness premise states that there are arbitrary exceptions to productive patterns, presupposing the presence of *productivity*. A pattern is *productive* if an unbounded number of items can, in principle, instantiate the pattern. For example,

an unbounded number of verbs can be used in the *double object* (or *ditransitive*) construction (e.g. *I gave John a book*). One can show that a given pattern is productive by showing that a newly coined word can be used to instantiate the pattern. For example, when the verb *text* was coined, meaning *to send a text message*, it was immediately available for use in the double object construction (e.g. *text me your address*). Less recently, Wasow (1981) pointed out that if the verb *satellite*, meaning “to transmit messages via satellite” were invented, then it would be usable as a double object form. This example was “prophetic” (Pinker 1989:17), because such a verb was subsequently invented, with the expected syntactic behavior.<sup>1,2</sup>

The ditransitive construction, like many productive generalizations, is subject to certain restrictions. For example, verbs such as *donate* and *explain* are awkward as ditransitive verbs, despite being eligible for use in a dative construction with the preposition *to*:

- (1) a. \*The man donated the church money.  
       b. The man donated money to the church.
- (2) a. \*The doctor explained me the situation.  
       b. The doctor explained the situation to me.

Are these restrictions simply idiosyncratic facts about *donate* and *explain*, which have to be learned individually? If so, then they are *arbitrary exceptions*. Under the premise of arbitrariness, there are arbitrary restrictions on productive generalizations that must be stipulated for individual words. The alternative view, which I am

---

<sup>1</sup>The productivity of the process by which novel verbs are coined from nouns, in addition to pragmatic limitations on it, is illustrated in detail by Clark and Clark (1979).

<sup>2</sup>Zwicky (1971:232) provides a memorable illustration of what it means to be productive: Imagine a new communication verb *greem*; then “it will be possible to greem for someone to get you a glass of water, to greem to your sister about the price of doughnuts, to greem ‘Ecch’ at your enemies, to have your greem frighten the baby, to greem to me that my examples are absurd, and to give a greem when you see the explanation.”

advocating in general in this dissertation, is that verbs like *donate* and *explain* are subject to general criteria governing the productivity of the ditransitive construction.

In the domain of ditransitivity, there is an etymological generalization regarding the set of verbs that fail to show the syntactic property in question: Verbs deriving from Latin tend not to be ditransitive. Suppose that this generalization were correct, and that there were no other explanation for why *donate* is unacceptable as a ditransitive verb. Under this scenario, it could be argued that *donate* is not an “arbitrary exception,” because its behavior falls under the scope of a larger generalization. However, this sense of “arbitrariness” is not the sense that is relevant for questions about learning. Because the learner of English presumably has no way (implicitly or explicitly) to acquire knowledge of the etymology of particular words, an etymological generalization is insufficient to render these restrictions non-arbitrary from a learning perspective. If etymological origin is the only factor determining whether or not a verb can function as a ditransitive, and there is no tangible reflex of etymology that a learner could use to distinguish ditransitive verbs from non-ditransitive verbs, then there are arbitrary exceptions to be acquired in this domain (see Pinker 1989 for further discussion).

Under the third premise of Baker’s Paradox, “no negative evidence,” learners do not use negative evidence to acquire their native language. *Negative evidence* can be defined as evidence against the grammaticality of some sentence type, that is, evidence that it is ungrammatical. Evidence that *donate* is ungrammatical as a ditransitive verb would thus be an example of negative evidence.

Is it truly a “paradox”? The reasoning by which the conjunction of the three premises leads to a contradiction is laid out by Schütze (1997:122) as follows:

If subcategorization is unpredictable [arbitrariness], then it must be learned for each verb individually. With no negative evidence available [no negative evidence], the child would have to limit herself to repeating

subcategorization frames perceived in parental speech. But this contradicts [productivity] (i.e. the child will use verbs with unattested subcategorization frames).

Although there are some small differences between Schütze's assumptions and the ones advanced in this dissertation, this quotation articulates one way to derive a contradiction from the simultaneous assumption of all three premises: The conjunction of the second and third premises (arbitrariness and lack of negative evidence) implies non-productivity, which of course, contradicts the first premise (productivity). One minor difficulty with this particular way of explaining how Baker's Paradox is a paradox is that it is cast in terms of subcategorization, even though the realm of Baker's Paradox extends beyond just subcategorization phenomena. A somewhat more important difficulty stems from the fact that the idea of arbitrariness depends on the idea of productivity; the premise of arbitrariness asserts that there are arbitrary exceptions to productive patterns. Hence it is impossible to imagine arbitrariness in the absence of productivity.

Here is another way to derive a contradiction from the conjunction of the three premises: The conjunction of productivity and absence of negative evidence implies the negation of arbitrariness. Why is this so? If a pattern is productive, then there is a process (in an abstract sense) that generates instances of that pattern, so long as certain criteria are satisfied. Arbitrariness means that there are items that satisfy the criteria, yet do not undergo the process. Negative evidence can be used to place limitations on the items that undergo a process. If there is no negative evidence, then there is no source of limitations on the process other than the criteria governing the productivity of the pattern,<sup>3</sup> which implies that there can be no arbitrary exceptions. With productivity, but without negative evidence, the learner is stuck with an overly general grammar.

---

<sup>3</sup>Negative evidence is a conceivable source of evidence regarding the criteria of the pattern, so an even stronger statement is possible here.

The process corresponding to productivity can be analogized to the flow of substance through a filter: Any particle below a certain size can fit through the filter. Arbitrary exceptions can be seen as particles that are small enough to fit through the filter, which, for some reason, still fail to pass through it. Because the filter cannot restrict the passage of these items, there must be some additional mechanism to keep them back. If there were no mechanism for keeping particles back other than the filter, then every particle small enough to fit through the filter would pass through it. This outcome corresponds to what would happen in the presence of productivity, in the absence of negative evidence: an overly general grammar, and no arbitrary exceptions.

The preceding paragraphs illustrate ways of showing that the three premises cannot be simultaneously maintained, and that they do, indeed, form a paradox. How should this paradox be resolved? Which of the premises should be denied? There is plenty of evidence against the “no negative evidence” premise. However, as I will argue, this strategy for resolving Baker’s Paradox is compatible with another strategy, namely, denial of arbitrariness. I will argue further in favor of a joint solution, denying both.

Although it is useful to understand why it is a paradox and how it can be resolved, there is a more important underlying question, which is how the phenomena that appear to instantiate Baker’s Paradox in a given language – restrictions on productive patterns – are acquired by speakers of that language. In order to answer this question, it is necessary to go beyond merely resolving the paradox and to determine how many of the premises of the paradox are true, and which ones.

## 1.1 The phenomena of interest

Baker’s Paradox can be instantiated in a variety of domains, but only certain exceptions to productive rules give rise to the learning problem: those that Baker (1979) called “embarrassing,” and not those that Baker called “benign exceptions.” Although this distinction is useful, I propose an alternative way of conceptualizing it. While Baker characterizes the difference between “benign” and “embarrassing” exceptions based on optional vs. obligatory transformational rules, I suggest that the two classes should be distinguished based on whether they comprise restrictions that can be learned through categorical preemption.

### 1.1.1 “Benign” exceptions

The “arbitrariness” premise of Baker’s Paradox states that there are arbitrary exceptions to productive patterns. However, some arbitrary exceptions to productive patterns are consistent with this premise, namely, those that Baker (1979) refers to as “benign exceptions.” Baker’s Paradox – and arbitrariness in particular – arises only with “embarrassing” exceptions.

Baker characterizes the “benign” class as exceptions to obligatory transformational rules. Baker’s example of a “benign” exception is *how come*, a *wh*- phrase which exceptionally fails to trigger subject-auxiliary inversion, an obligatory rule. Subject-auxiliary inversion after *how come* is ungrammatical, as indicated by the asterisk in (3b):

- (3) a. How come he’s late?  
 b. \*How come is he late?

This exception to the rule that *wh*- phrases (*who*, *what*, *where*, etc.) trigger subject-auxiliary inversion is “benign” because it can be learned through observation of sentences in which inversion does not take place after *how come*.<sup>4</sup>

The learner can use the following logic to acquire obligatory rules: Either the input structure (e.g. an uninverted question) or the output structure (e.g. an inverted question) is grammatical, but not both. Given this exclusive disjunction between the options, it suffices to learn that one option is grammatical, in order to learn that the other option is ungrammatical. I suggest that this exclusivity is the crucial feature of “obligatory rules” that makes them “benign,” rather than the presence of obligatory rules per se.

Mutual exclusivity between two alternatives makes it possible for one form to *preempt* the other. The type of preemption in which the grammaticality of one form implies the ungrammaticality of another can be called *categorical preemption*. Categorical preemption stands in contrast to *statistical preemption*, which Goldberg (2006:93) defines in terms of “repeated witnessing” of competing forms. In categorical preemption, unlike statistical preemption, witnessing the preempting form only once could, in principle, be sufficient to learn that the preempted form is ungrammatical as an expression of the relevant meaning. For example, a learner with an excellent memory who is very astute could learn that *\*goed* is ungrammatical as the past tense of *go*, simply by hearing *went* once, as long as the learner knows that the two have (or would have) the same meaning. Statistical preemption, by contrast, is a potential mechanism for acquiring contrasts in grammaticality between forms that are not necessarily mutually exclusive.

---

<sup>4</sup>An alternative analysis of this case is that *how come* is not a *wh*- phrase in the relevant sense, but a complementizer, so what has to be learned is not that *how come* exceptionally fails to trigger inversion, but the category of *how come* (Ginzburg and Sag 2000). Baker’s logic still stands, however, assuming that the learner has somehow acquired the exclusive disjunction between the grammaticality of the inverted and non-inverted forms.

Categorical preemption typically occurs in morphological paradigms. In a morphological paradigm, there is a matrix of slots, and each slot contains exactly one form. For example, the present tense paradigm for the verb *go* consists of a matrix of slots corresponding to the singular and plural forms for first, second, and third person subjects. The grammaticality of *goes* as a third person singular present tense form is evidence against the grammaticality of, for example, *go* as a third person singular present form, because only one form can fill this slot. In other words, *goes* preempts *go* as a third person singular present form. Another example comes from number inflection on nouns: For every noun, there is a unique singular form and a unique plural form. This means that the grammaticality of a given plural form gives evidence against the grammaticality of another potential form for the plural. Thus, the grammaticality of *feet* gives evidence against the grammaticality of *\*foots*.

As DiSciullo and Williams (1987:10-14) point out, inflectional paradigms provide the “clearest cases” of this type of phenomenon (which they refer to as “blocking” rather than “preemption”).<sup>5</sup> However, Di Sciullo and Williams also argue that such exceptions are not limited to the domain of morphology, and list several examples of preemption in syntax. A relatively straightforward example of preemption in syntax comes, I believe, from the grammar of temporal expressions, which are analyzed by Fillmore (2002). Fillmore points out that *last* and *next* can be combined with *year*, *month*, and *week* (e.g. *last year*, *next year*, etc.), but not *day*; English speakers use *yesterday* and *tomorrow* instead of adverbial *\*last day* and *\*next day*.<sup>6</sup> *Yesterday* and *\*last day* can be seen as competing for the same slot in a paradigm.

Competition between forms for a single slot in a paradigm gives rise to exclusivity between those forms. As discussed above, exclusivity, coupled with evidence *for* the

---

<sup>5</sup>The choice of the term “preemption” over “blocking” here is motivated by the connotation of the term “blocking” as a restriction on forms, rather than form/meaning pairs, whereas “preemption” relates to form/meaning pairs (Clark and Clark 1979; Clark 1987).

<sup>6</sup>It seems that *yesterday* preempts *\*last day* only when the noun *day* is overt: *the day before last* and *the day before yesterday* both seem acceptable.

grammaticality of a given means of expressing a given function, gives rise to evidence *against* the grammaticality of another form. The logic can be schematized as follows (exclusive disjunction is notated with  $\oplus$ ): if  $A \oplus B$ , and  $A$ , then  $\neg B$ . It is worth noting that because this reasoning results in the conclusion that a given form is ungrammatical, it can be seen as indirectly providing a type of negative evidence, through positive evidence.<sup>7</sup> This does not seem to be the type of negative evidence whose existence (or use) is denied by the “no negative evidence” premise, however.

To summarize, the class of exceptions that Baker called “benign” can be recast as the set of restrictions that arise through categorical preemption, rather than as exceptions to obligatory rules. Categorical preemption relies on mutual exclusivity between forms, coupled with positive evidence for one form.

### 1.1.2 “Embarrassing” exceptions

According to Baker (1979), “the benign exceptions invariably involve rules that are obligatory. By contrast, the embarrassing exceptions are in connection with optional rules” (p. 546). Such optional transformational rules would rewrite (4a) as (4b), or (5a) as (5b), or (6a) as (6b).

- (4) a. I gave a book to John.  
b. I gave John a book.
- (5) a. The child seems to be happy.  
b. The child seems happy.
- (6) a. It is likely that Robin will win.  
b. Robin is likely to win.

---

<sup>7</sup>This point echoes a point embodied in Marcotte’s (2005) slogan, “positive evidence as negative evidence,” which highlights the false dichotomy between positive and negative evidence.

These rules are “optional” in the sense that they can transform the (a) structure into the (b) structure, but are not required to operate, so the (a) structure could also surface. If these rules were allowed to operate indiscriminately, then they would falsely allow unacceptable dative constructions like (7b), or the use of *happen* without *to be* as in (8b), or subject-raising with *probable* as in (9b).

- (7) a. I donated a book to the library.  
 b. \*I donated the library a book.
- (8) a. The child happens to be happy.  
 b. \*The child happens happy.
- (9) a. It is probable that Robin will win.  
 b. \*Robin is probable to win.

The words *donate*, *happen*, and *probable* can be seen as exceptions to the optional transformational rules that convert the (a) sentences into the (b) sentences.

With optional rules such as these, the grammaticality of the (a) sentence does not imply the ungrammaticality of the (b) sentence; sentences of the (a) type and the (b) type can peacefully coexist. Consider the inability of *donate* to undergo the dative alternation, for example. The dative alternation “rule” is “optional” in the sense that some verbs appear in both the double object construction and the prepositional object construction (e.g. *give*). Thus, the use of a verb in the prepositional object construction does not tell the learner that use of it in the double object construction is ungrammatical, unlike in the case of an “obligatory rule.” When learners hear the verb *donate* with a prepositional object, therefore, they do not acquire a basis for concluding that *donate* in the double object construction is ungrammatical.

I suggest that exclusivity between forms is the defining feature of the relevant class of phenomena (Baker’s “embarrassing exceptions”) rather than optional rules.

This class of phenomena can be cast as the class of restrictions on productive patterns that do not arise through categorical preemption. For these restrictions, there is no conventional near-synonymous alternative form that outcompetes the ungrammatical form in question for the same slot, and yet, the form is ungrammatical.<sup>8</sup>

This recharacterization of the set of relevant phenomena brings into the scope of the discussion constructions that are unlikely to be related to other constructions by transformational rules even under transformational assumptions, such as the attributive (or “prenominal”) adjective construction as in (10b).<sup>9</sup>

- (10) a. This man is tall.  
       b. Here is a tall man.

Adjectives that fail to appear attributively such as *aghast* (*\*an aghast man*) can be considered forms that are not categorically preempted, and are therefore among the phenomena of interest. (Being of interest, examples like *aghast* are discussed in Chapter 5.)

In summary, Baker’s Paradox arises only with a certain type of phenomenon: restrictions on productive patterns that do not arise through categorical preemption. These are the cases for which there is no conventional near-synonymous preempting form.

---

<sup>8</sup>I use “restriction” instead of “exception” because the latter connotes idiosyncrasy, begging the central question.

<sup>9</sup>In fact, it was proposed by McCawley (1964) that these constructions were related by a transformational rule (the “Adjective Shift Rule”), which rewrites a noun phrase with a relative clause containing an adjective as an attributive adjective, but this analysis has long been discredited due to the non-synonymy of predicative and prenominal forms (e.g. Bolinger 1967, Levi 1978).

## 1.2 The solution space

As Baker’s Paradox arises from the simultaneous assumption of three premises, there are three possible solutions to it: Denying the first, second, or third premise. However, the space of possible answers to the question of how restrictions on productivity are acquired has a slightly different structure. I will argue that there are two orthogonal dichotomies, one between arbitrariness and criteria-governed productivity, and one between conservatism and negative evidence. These theoretical choices are independent, so it is not necessary to choose between the “arbitrariness” and “no negative evidence” premises; in fact, I propose rejecting both.

### 1.2.1 Arbitrariness vs. criteria-governed productivity

The answer to the question of how restrictions on productivity are acquired depends in part on whether or not those restrictions are arbitrary, that is, whether or not the premise of arbitrariness holds. Denying arbitrariness can be called the *criteria-governed productivity* approach to Baker’s Paradox (Pinker 1984, 1989), because it explains restrictions on productive patterns by specifying criteria limiting the productivity of the pattern.

The criteria-governed productivity approach to Baker’s Paradox involves arguing against the existence of arbitrary exceptions by showing, for a given putative arbitrary exception to a productive pattern, that the item fails to meet the criteria governing the productivity of the pattern. For example, if one of the criteria for functioning as a ditransitive verb is for the verb to describe a transfer of possession, then any verb that does not describe a transfer of possession would fail to meet the criteria governing the productivity of the pattern. An exception to a productive pattern is only arbitrary if it fails to instantiate the pattern despite meeting its criteria.

The criteria-governed productivity approach is the one I advocate in this dissertation. My central thesis is that arbitrary exceptions (of the non-categorical preemption, or “embarrassing” variety) do not exist; rather, the restrictions on productive patterns that have been observed follow from criteria governing the productivity of the pattern.

### 1.2.2 Conservatism vs. negative evidence

To argue against arbitrariness is not to argue for the “no negative evidence” premise. Cross-cutting the dichotomy between arbitrariness and criteria-governed productivity is a dichotomy between negative evidence, on one hand, and conservatism, on the other. I suggest that the choice between conservatism and negative evidence is determined by whether or not there is a stage of overgeneralization. Either children do overgeneralize at some stage, but then cut back, or they are conservative, and never overgeneralize. If they do overgeneralize at first, they need negative evidence in order to change their grammar, regardless of whether the grammars they end up with contain arbitrary exceptions. Although the choice between these theoretical options is open in principle, there is strong evidence for the presence of an overgeneralization, hence, for the role of negative evidence.

There is a great deal of evidence that children are indeed conservative at a very young age (Akhtar 1999; Akhtar and Tomasello 1997; Berman 1978; deVilliers 1985; Dromi 1987; Olguin and Tomasello 1993; Tomasello 1992; Tomasello et al. 1997; Lieven et al. 1997; Ingram and Thompson 1996). However, there is also a great deal of evidence that children begin to generalize – and overgeneralize – from age 2 on. There are plenty of attested examples in spontaneous child speech (e.g. Bowerman 1988), and there is experimental evidence in support of this as well (Maratsos et al. 1987; Braine et al. 1990; Pinker et al. 1987; Gropen et al. 1989, 1991, 1996; Brooks and Tomasello 1999; Hochberg 1986). The notion that children go through an overgeneralization

stage is relatively uncontroversial (although according to Pinker’s (1989) “minimalist hypothesis,” there is in fact no overgeneralization stage, and children’s errors are either “one-shot innovations” or result from incorrect verb meanings). Under the assumption that overgeneralization does take place, children need negative evidence in some form in order to develop adult-like grammars.<sup>10,11</sup>

Other arguments have been made against the absence of negative evidence as well. Although there are some findings suggesting that negative evidence is not available to children (Brown and Hanlon 1970, Braine 1971, Hirsch-Pasek et al. 1984, Bowerman 1988), there is strong support for the existence of negative evidence of various types (see Chouinard and Clark 2003 for evidence and a review).

Based on these considerations, I conclude that the “no negative evidence” premise cannot be maintained and that some form of negative evidence is at work. Negative evidence need not take the form of explicit corrections; it could be obtained either *directly* through corrective feedback, or *indirectly* through statistical computations over patterns of speech (Schütze 1997, Rohde and Plaut 1999, i.a.). But negative evidence in some form – evidence against the grammaticality of certain sentence types – seems to be a logical consequence of the existence of an overgeneralization stage, and empirically supported as well.

---

<sup>10</sup>The only other solution to what Babyonyshev et al. (2001) call “the problem of late knowledge” is to hypothesize innate grammatical principles which take time to mature (Babyonyshev et al. 2001; Borer and Wexler 1987). I will ignore this possibility in the absence of strong evidence for it, as I do not take it to be the null hypothesis.

<sup>11</sup>Marcotte (2005:4) seems to see negative evidence as opposed to innateness: “if negative evidence is unavailable, then the child must cut down on overgeneralization through some method that requires no evidence, usually taken to be innately specified.” He attributes such a theory to Pinker (1989), who posited certain innate properties, but I do not believe this attribution is accurate, because Pinker’s innatist assumptions were not intended as an explanation for why children recover from errors. As far as I can tell, the only innatist explanation for error recovery is the idea that grammatical principles take time to mature.

### 1.2.3 Arbitrariness and negative evidence

Because each of the premises correspond to different ways of resolving the paradox, one might take them to be mutually exclusive, but they are not. Denying arbitrariness does not force one to adopt or deny the idea of “no negative evidence.” Negative evidence could be used to acquire general constraints on productive patterns after a stage of overgeneralization. On the other hand, the use of negative evidence also does not follow logically from the non-existence of arbitrary exceptions: Supposing that overgeneralization never occurred, it would be conceivable that the learner was conservative in such a way as to keep productivity restricted on the basis of general criteria, rather than individual words.

However, arbitrariness and negative evidence are not completely unrelated ideas. The question of arbitrariness affects the *type* of negative evidence that it would be necessary for learners to use. If there are arbitrary exceptions, then the type of negative evidence that is needed concerns specific words. In that case, the negative evidence in question would have to be construed as evidence against the grammaticality of the particular combination of the word with the target construction. If arbitrary exceptions do not exist, then a different type of negative evidence is necessary for learning the appropriate restrictions. General criteria governing the productivity of a pattern could be learned through negative evidence construed as pertaining to entire classes of words, whether directly through feedback, or indirectly. For a description of such an approach, Mazurkewich and White (1984) suggest that “children may initially overgeneralize, but they eventually identify the criteria that define the lexical class appropriate to the rule. When they do, they limit productive use of the rule to lexical items of the right class, and errors cease” (Bowerman 1988:82–83). Under this scenario, non-arbitrariness and negative evidence co-exist.

To summarize, the implications of the existence of arbitrary exceptions do not pertain to *whether* negative evidence is used in language learning, but *what type*

of negative evidence is being used. If arbitrary exceptions do exist, then negative evidence must be construed as pertaining to the use of a particular word in a particular construction. If arbitrary exceptions do not exist, it may be that negative evidence (implicit or explicit) is taken as relevant to some semantic class.

#### 1.2.4 Arbitrariness and conservatism

Just as the issue of arbitrariness affects the type of negative evidence a learner would need to use, it also affects the type of conservatism that would be involved in learning. Supposing that arbitrary exceptions do exist, and that these restrictions are learned through conservatism on the part of the learner (without negative evidence), so-called *strict lexical conservatism* would be required (Pinker 1989). According to this idea, learners never use a word in a construction unless they have witnessed the word in the construction. This scenario is rather impossible to imagine, however, because strict lexical conservatism would contradict the premise of productivity (i.e. that the pattern is productive), which arbitrariness relies on.

Supposing that arbitrary exceptions do not exist, and maintaining the assumption that restrictions on productivity are learned through conservatism, a less strict variety of conservatism would be at play. According to this idea, learners might not use a word in a construction unless it has a feature of words that have been witnessed in the construction.

#### 1.2.5 Conservatism and Attentiveness

To complete the discussion of the space of theoretical possibilities regarding how restrictions on productive patterns are learned, let us consider Culicover's (1999) idea of the *Conservative Attentive Learner*. Culicover argues that arbitrary lexical idiosyncrasies abound in English and concludes, based on the existence of such arbitrariness,

that the learner must be *conservative* and *attentive*. A *conservative* learner, broadly speaking, is one who does not over-generalize; an *attentive* learner is one who pays attention to the use of individual words in individual constructions.

The relationship between arbitrariness and conservatism has already been discussed. A corollary of the fact that the “arbitrariness” and “no negative evidence” premises are orthogonal is that the ideas of arbitrariness and conservatism are not so intricately linked as Culicover (1999) makes out, because conservatism is the flip-side of negative evidence. The nature of the relationship between arbitrariness and conservatism is the same as the nature of the relationship between arbitrariness and negative evidence. In particular, the relationship is logically orthogonal.

To the extent that arbitrariness and conservatism are related issues, it is far from being the case that arbitrariness implies conservatism, as Culicover (1999) suggests. Rather, arbitrariness seems inconsistent with conservatism, because, as discussed above, the type of conservatism that would be required to learn arbitrary exceptions is strict lexical conservatism. Strict lexical conservatism is inconsistent with productivity, and therefore, arbitrariness. Conservatism is certainly not a consequence of arbitrariness.<sup>12</sup>

The idea of attentiveness to individual lexical items is more intrinsically related to the question of arbitrariness than either conservatism or negative evidence. If there are arbitrary exceptions, then the learner must be attentive to the use of individual words in individual constructions, and furthermore encode and store such information. To put it simply, arbitrariness implies attentiveness. This means that lack of

---

<sup>12</sup>By the same token, an implicational relationship could be argued to exist between arbitrariness and negative evidence. If there were arbitrary exceptions and they were learned without any negative evidence, restrictions would be acquired solely through strict lexical conservatism. Strict lexical conservatism would contradict the “Productivity” (i.e., that the pattern is productive), which the “Arbitrariness” premise relies on. In any case, both the existence and non-existence of arbitrary exceptions are at least consistent with a view of learning that involves negative evidence.

attentiveness on the part of the learner is a possible explanation for lack of arbitrariness. However, there are other possible explanations, which will be discussed in the final chapter.

### 1.3 Roadmap

To support the criteria-governed productivity approach to Baker's Paradox, I survey a range of phenomena that have been adduced in support of the premise of arbitrariness. For each one, I aim to show that the restrictions that have been observed do not manifest arbitrary exceptions, by identifying the criteria governing the patterns involved, and showing that the items exhibiting these restrictions fail to meet the criteria governing the pattern.

The first case study focuses on the English causative alternation, relating transitive *break*, for example, to intransitive *break* (Chapter 2). Although semantic criteria have been proposed for predicting whether or not a given verb undergoes the causative alternation (Pinker 1989; Levin and Rappaport Hovav 1995), it has been claimed that these criteria do not fully capture the set of alternating verbs, and that the learner must acquire knowledge of certain arbitrary lexical exceptions (e.g. Bowerman and Croft 2008). I argue that Levin and Rappaport Hovav's (1995) semantic criteria for participating in the causative alternation in fact capture the set of alternating verbs well, correctly excluding the putative arbitrary exceptions.

In Chapter 3, I investigate constraints on the productivity of ditransitivity, as in, for example, *give John the book*. In this domain, there are some semantic restrictions, but it also appears that there are restrictions based on the form of the verb. Verbs of Latinate origin such as *donate* tend to be unacceptable as ditransitives, and it has been suggested that this etymological generalization could be learned based on

aspects of the form of the verb that would be observable to the learner: In particular, Latinate verbs tend to be *longer* than verbs of native origin, as measured by either prosodic weight (Grimshaw and Prince 1986; Grimshaw 2005) or morphological complexity (Harley 2006). The problem in this domain, therefore, lends itself to experiments involving nonce words in which meaning is held constant and the form of the word is manipulated. Building on Gropen et al.'s (1989) finding that monosyllabic nonce words like *moop* are more acceptable as ditransitives than trisyllabic ones like *dorfinize*, I present two nonce word experiments with adult native English speakers, to investigate the roles of prosodic weight, morphological complexity, perceived etymology, and formality in governing the productivity of ditransitivity.

Chapter 4 focusses on the behavior of prepositions. This investigation targets one of the cases discussed in Culicover's (1999) book *Syntactic Nuts*, in which he argues for the existence of a range of idiosyncrasies in English. He claims in particular that prepositions can differ arbitrarily in their ability to strand (as in *Who did you speak to?*) or pied-pipe (as in *To whom did you speak?*), but I show in Chapter 4 that the behavior of the prepositions he discusses can be understood on the basis of general principles of English grammar.

In Chapter 5, I investigate a putative case of arbitrariness involving adjectives. This case is brought up by Goldberg (2006), who argues that adjectives differ arbitrarily in their ability to function either prenominal (e.g. *That is a mere/\*asleep child*) or predicatively (e.g. *That child is asleep/\*mere*). Again, I argue that the behavior of these words can be understood on the basis of general principles.

I conclude that there is no evidence from any of these domains for the necessity of learning arbitrary exceptions to productive patterns. This conclusion about the linguistic situation leads to a view of the language learner as one who prefers explanatory generalizations over individual stipulations: the *explanation-seeking learner*.