Yearbook of the Poznań Linguistic Meeting 8 (2025), pp. 41–65
DOI: 10.14746/yplm-2025-0003

Possible and probable errors in child language

Sarah Brogden Payne
Stony Brook University
sarah.payne@stonybrook.edu

Abstract

This paper investigates the implications of Blind Alley Developments (BADs, Dressler et al. 2020) – ephemeral, systematic errors in early child productions – for a theory of language acquisition. I argue that a sufficient theory of acquisition should not simply *allow* for BADs, but rather should precisely account for the timeline and content of these errors in terms of the learning mechanisms that the child employs; in other words, it should be a *learning theory.* I present such an account, which draws on the complementary approaches of the Tolerance Principle (Payne & Yang 2023) and Natural Morphology (Dressler et al. 2020). This account provides both a formalization of the timeline of BADs – in terms of a calibration of productivity over the learner's internal vocabulary – and predictions regarding the expected types of BAD constructions. I provide applications of this account to both *weak BADs* – which involve the use of an incorrect form that nevertheless marks the correct target function – and *strong BADs* – which involve the use of a form never used for the target function in the learner's input.

Keywords: morphological acquisition; Natural Morphology; Tolerance Principle; Blind Alley Developments

1. Introduction: Blind Alley Developments

Blind Alley Developments (BADs, Dressler et al. 2020, 2023; Payne & Yang 2023; Dressler & Payne 2025) refer to systematic, ephemeral deviations of child productions from their input. There are two types of BADs as defined by Dressler et al (2020): *weak BADs* occur when children wrongly extend and overapply a form that is present in their input but not productive in the adult grammar, while *strong BADs* occur when nothing in the child's input can directly explain the form that they construct. Payne & Yang (2023) discuss these types of BADs in terms of *form* (i.e., the specific morphophonological process

being employed) and *function* (i.e., the thing that is being marked by the process). Under this view, weak BADs involve the use of an incorrect form that nevertheless marks the correct target function (albeit unproductively in the adult grammar), while strong BADs involve the use of a form that is *never* used to mark the given function in the input. We now turn to concrete examples of each type of BAD.

1.1. Weak Blind Alley Developments

Weak Blind Alley Developments occur when children overextend and incorrectly apply a form that nevertheless marks the correct target function in their input. Weak BADs can be contrasted with the far more common error of *over-regularization* (e.g., overapplication of the past tense *-ed* in English-learning children) in that weak BADs involve the overapplication of forms which are *unproductive* in the adult grammar, while over-regularization errors correspond to overapplications of forms which are productive in the adult grammar (Palermo & Ebhart 1968; Clahsen et al. 1992; Marcus et al. 1992; Clahsen & Rothweiler 1993; Xu & Pinker 1995; Allen 1996; Ravid & Farah 1999; Marasotos 2000; Clahsen et al. 2002; Demuth 2003; Maslen et al. 2004; Deen 2005; Mayol 2007; Caprin & Guasti 2009; Kauschke et al. 2011; Lustigman 2013; Bat-El 2014; i.a.). We consider two examples of weak BADs here: the interfixation of *-e-* in German noun-noun compounds (Korecky-Kröll et al. 2017; Dressler et al. 2020; Dressler & Payne 2025) and the formation of the English past tense with ɪ → æ/_ŋ#, as in *sing-sang* (Xu & Pinker 1995; Payne & Yang 2023; Dressler & Payne 2025).

In Standard German – including Austrian German – interfixes are sometimes added after the first constituent of a noun-noun compound, though the majority of compounds are formed without interfixes (Libben et al. 2009; Dressler et al. 2020). Specifically, *-n-* is productively added if the first constituent ends in an unstressed *-e* (e.g., *Garage-n+eigentümer* 'owner of a garage'), and other interfixes are inserted in other conditions (e.g., *Kind-er+garten* 'kindergarten'; *Land-es+kunde* 'Geography'; Libben et al. 2009). In some compounds, *-e-* is inserted (e.g., *Hund-e+würstel* 'dog poop'); importantly, however, *-e-* is infrequent and unproductive in the adult grammar (Libben et al. 2009). However, Korecky-Kröll et al. (2017) find that all three of the Austrian German-learning children they study overapply *-e-* in noun-noun compounds (e.g., **Luft-e+bon-e* for *Luft+ballon-e*, 'air ballons';

**(Mine)ral-e+wasser* for *Mineral+wasser*, 'mineral water'). This constitutes a weak BAD because the children over-apply an unproductive form (namely, *-e-*) to mark a function that it was observed marking in the input (namely, noun-noun compound interfixation). In other words, *-e-* does indeed mark noun-noun compounds in the child's input, but it does not do so productively in the adult grammar. Under the present account, *-e-* enjoys a transient stage of productivity for the children studied by Korecky-Kröll et al. (2017), leading to its observed overapplication.

Similarly, English past tense formation via the *sing-sang* pattern (i.e., ɪ → æ/_ŋ#) is unproductive in the adult grammar, as evidenced by its lack of application to novel -ɪŋ# verbs entering English: both *Bing* (the Microsoft search engine) and *bling* take *-ed* in the past tense. However, Xu & Pinker (1995) find that some children overapply this pattern: in the 4-million-word corpus of child English extracted from the CHILDES project (MacWhinney 2000), the past tense of *bring* is produced as *bringed* 32 times, *brang* 6 times, and *brung* 5 times. Though *brang* is an acceptable past tense form for adult speakers of some dialects of North American English, there is little doubt that some children do overapply the ɪ → æ/_ŋ# pattern to form *brang*: further evidence comes from the presence of *swing-swang* and *fling-flang* in the same CHILDES corpus, as neither of these past tense forms is acceptable for adult speakers. This overapplication again constitutes a weak BAD: the form (namely, ɪ → æ/_ŋ#) is used to mark the correct function (past tense inflection), despite this mapping being unproductive in the adult grammar. Notably, the overapplication of ɪ → æ/_ŋ# is the *only* error found by Xu & Pinker that appears analogical in nature. Nevertheless, some scholars have accounted for this error by adding an analogical mechanism to their framework (e.g., Bybee 1985, Rumelhart & McClelland 1986, Pinker & Prince 1988, Albright & Hayes 2003). The model presented here, however, accounts for this error not in terms of analogy but in terms of an ɪ → æ/_ŋ# rule that enjoys a brief stage of productivity in the child's developing grammar; analogy has no place in the account put forth here.

1.2. Strong Blind Alley Developments

Strong Blind Alley Developments occur when a child uses a form that is never attested with the given function in the input. We again consider two examples: Russian total root reduplication to express iterativity, durativity or

imperfectivity (Dressler et al. 2020; Dressler & Payne 2025), and two strong BADs created by a Greek child learning the subjunctive (Dressler et al. 2020; Dressler & Payne 2025).

In Russian, iterativity is expressed via imperfective verbs or secondary imperfectivization devices, but Dressler et al. report several cases of children instead making use of total root reduplication (e.g., *njam-njam* used to mean 'I am eating'). Although reduplication as a formal pattern does exist in Russian (Israeli 1997), the function of this reduplication is not to express iterativity or ongoingness, but rather the express intensification, among other things. Thus, the children making use of total root reduplication to express iterativity would have *never* seen reduplication being used for this function in their input, and we thus have a case of a strong BAD.

An even more striking example of a strong BAD comes from a Greek boy documented by Dressler et al. (2020). This boy entered two strong BADs, one after the other, in acquiring the Greek subjunctive. In adult Greek, the subjunctive is marked by an unstressed proclitic (e.g., *káni* '(s)he makes' → *na+kánki* 'let him/her make'). The child, however, first marks the Greek subjunctive by lengthening the verb root (e.g., *ká:ni*). After two weeks, he gave up this BAD and instead employed one of initial partial reduplication (e.g., *ka+káni*). Notably, Modern Greek has neither phonemic vowel length nor non-emphatic vowel lengthening. Thus, the formal pattern of vowel lengthening is not only unattested marking the subjunctive in the child's input, but is not used for any inflectional purpose. What's more, Modern Greek has no inflectional reduplication, though *total* reduplication is used to express a range of other semantic and pragmatic information (Kallergi 2015). Again, then, the formal pattern of reduplication is not attested marking the subjunctive in the child's input, though it may be attested marking other functions. Importantly, however, the reduplication that is likely attested in the child's input is *total*, while the child's construction is one of *partial* reduplication. For both vowel lengthening and reduplication, then, the child is making use of a form that is *never* employed to mark the subjunctive in his input, and we thus have two more cases of strong BADs.

1.3. Implications of BADs for theories of acquisition

Though Blind Alley Developments are relatively rare and do not affect every child, their importance for theories of language acquisition should not be

dismissed. In particular, it is not sufficient for a theory to simply *allow* for BADs: a theory that allows for BADs may well overpredict unattested error patterns in child language. This is what Dressler et al. (2020) refer to as the Feyerabendian problem of "everything goes." It is also not sufficient for a theory to explain in broad strokes how BADs occur without providing *causal mechanisms* that make precise predictions about the nature and timeline of BADs. Rather, an adequate theory of BADs is a *learning theory* that precisely accounts for the timeline and form of BADs as consequences of the learning mechanisms that the child employs. To further formalize this, I introduce two desiderata stemming from two key facets of BADs. An adequate theory of acquisition should provide a mechanistic explanation for:

(1) The *timeline* of BADs: what causes children to enter the Blind Alley exactly when they do, and why do they escape from them precisely when they do? Further, why do children escape from BADs so soon after entering them?

(2) The *content* of BADs: what error patterns do we expect the child to produce during their BADs? In the case of weak BADs, what forms in the input do we expect to be overextended? In the case of strong BADs, which unattested forms do we expect the child to use to mark a given function? Similarly, what constructions do we *not* expect the child to produce, and why? Dressler et al. (2020) pose this desideratum in terms Feyerabend's epistemological scandal of "everything goes": a theory that accounts for strong BADs by simply allowing the learner to construct *any* pattern does nothing to explain *why* we observe some Blind Alleys and not others in child productions.

In this paper, I will argue that Desideratum (1) can be accounted for under a learning-theoretic account such as that proposed by Payne & Yang (2023). Specifically, the mechanistic approach that Payne & Yang propose accounts precisely for the timeline of BADs in a way that previous accounts have failed to. By contrast, while the account of Payne & Yang determines which weak BAD constructions are *theoretically possible,* it does little to explain why the child may be attracted to some of these theoretically possible BADs over

others, and has less to say about the constructions selected during strong BADs. The latter question is the focus of Dressler et al. (2020), who develop a theory – couched in Natural Morphology – that largely satisfies Desideratum (2). At the same time, this account struggles to satisfy Desideratum (1) because it does not provide a precise mechanism by which children may escape their BADs, but only references the opposing input. As such, by marrying the complementary strengths of the accounts proposed by Dressler et al. (2020) and Payne & Yang (2023), we are able to develop a theory of acquisition that not only allows for BADs, but provides a mechanistic explanation of *why* they have the timeline and content that they do.

2. Previous accounts of BADs

Though some previous work has noted the presence of Blind Alley-like constructions in children's productions (Munson & Ingram 1985, Bittner et al. 2003, Mariscal 2009) and perception (Pater et al. 2004), Dressler et al. (2020) constitutes the first systematic, cross-linguistic study of morphological BADs. Dressler et al. provide data on the trajectories and content of nine strong and weak BADs in German, Greek, French, and Russian, and present an account of these BADs in terms of self-organization and Natural Morphology. Specifically, the authors argue that children *construct* their own course of acquisition based on interactions between their stages of representation and processing and their linguistic input, following the constructivist view of self-organization (cf. Karpf 1990; Dressler & Karpf 1995). During this interaction, children select linguistic forms from their input based on the principles such as saliency and frequency, as well as the cognitively-based preferences of Natural Morphology (cf. Dressler 1995; Kilani-Schoch & Dressler 2005; Dressler et al. 2014; Dressler & Kilani-Schoch 2016, i.a.). Such preferences include iconicity between form and meaning, morphosemantic and morphosyntactic transparency, optimal size of words according to their morphological status, and (bi)-uniqueness; particular attention is given to iconicity, which can account for all of the BADs presented by Dressler et al.. Under the approach put forth by these authors, BADs are a consequence of the preferences and preference rankings of Natural Morphology. For example, the use of reduplication to express iterativity in Russian can be explained as a consequence of the iconicity preference.

This approach has particular strengths when accounting for strong BADs, since the constructions made during these BADs can be explained not directly by the input but rather by a cognitively-based universal preference for certain forms, such as iconicity for Russian reduplication. Since these preferred constructions are expected to appear in other languages, Dressler et al. thus predict that the forms children produce during strong BADs should occur in languages other than the input language: that is, they should be typologically present. Indeed, this is how Dressler et al. address the Feyerabendian question of whether "everything goes" and satisfy Desideratum (2).

Dressler et al. argue that their theory of acquisition can account for BADs in a way that other theories – specifically, usage-based and nativist approaches – cannot. As such, they claim that BADs provide the strongest piece of evidence support their view of a constructivist approach to self-organization. It is easy to see the problem that BADs pose for usage-based accounts (e.g., Tomasello 2000; Lieven 2008; Bybee 2010; Engelmann et al. 2019): as both Dressler et al. (2020) and Payne & Yang (2023) note, these accounts overstate the extent to which the child's grammar directly mirrors the input (e.g., Yang 2013; Yang 2016: Section 2.1). Though usage-based approaches can account for omissions and analogical errors (e.g., Lieven 2008), they cannot currently account for the use of a form *never* present in the input to mark a given function. Since BADs – particularly strong BADs – involve systematic deviations from the input, such theories thus do not predict them. At the same time, Dressler et al. (2020) argue that existing nativist theories of acquisition cannot explain BADs either, since there is nothing in Universal Grammar to predict BADs or constrain possible BAD developments. In particular, they argue that the learning-theoretic approach given by the Tolerance Principle (Yang 2016) cannot account for BADs because it would rely on the presence of *sufficient* support from the input in order for a BAD to emerge, and thus a BAD should (1) never emerge to begin with, and (2) be immediately knocked out by the Tolerance Principle if it were to emerge.

As Payne & Yang (2023) argue, however, that the Tolerance Principle is able to account for BADs because it calibrates productivity over the learner's *internal* vocabulary. The Tolerance Principle (Yang 2016) is based on the intuition that a linguistic process must *earn* productivity by being applicable to a sufficiently large number of candidates, where sufficiency is calibrated over the learner's internal vocabulary. Specifically, the learner calculates two values for a rule *R*: *N*, the number of items in the learner's vocabulary which fit *R*'s

description, and *e*, the number of these items to which *R* does not apply. Under the Tolerance Principle, the rule *R* is productive only if:

(3) $$e \leq \theta_N = \frac{N}{\ln N}$$

I refer the reader to Yang (2005, 2016) for a full derivation of Equation (3), but note that it is a consequence of the hypothesis that children will generalize a rule when it is more efficient – in terms of lexical access time – to do so. To illustrate the application of the Tolerance Principle, consider a toy example in which an English-learning child knows 10 verbs and 3 do not take *-ed* in the past tense. In this case, we have $N = 10$ and $e = 3$, and since

$$3 \leq \frac{10}{\ln 10} \approx 4.3,$$

the *-ed* suffix will thus become productive for marking the past tense.

However, some verbs may never be attested in the past tense in the learner's input. To address this, Yang (2016: 177) introduces a corollary, the Sufficiency Principle, which specifies how generalizations are formed when only a subset of the *N* items have been seen in the necessary context to determine if the rule applies to them. The Sufficiency Principle posits that if *M* items have been seen with the rule applying, then the learner generalizes if $M > N - \theta_N$, a super-majority of *N*. Returning to our example, if we know the learner has seen 7 of its 10 verbs taking *-ed* in the past tense, then $N = 10$ and $M = 7$. Since

$$7 > 10 - \frac{10}{\ln 10} \approx 10 - 4.3 \approx 5.7,$$

the past tense *-ed* rule will be learned as productive *regardless* of whether the remaining 3 verbs are *unattested* in the past tense or are known *exceptions* to the *-ed* rule. It is easy to see that $N - M$ under the Sufficiency Principle is equivalent to *e* in Equation (3): both quantities must be less than θ_N for a generalization to be made. As such, the Sufficiency Principle can be thought of as a "worst case" application of the Tolerance Principle: we assume that all of the unseen items may turn out to be exceptions. The Tolerance and Sufficiency Principles have received extensive experimental support from artificial language studies with precisely controlled conditions (e.g., Schuler 2017, Emond

& Shi 2021); I refer the reader to Payne (2022) for an in-depth review of these studies. Throughout this paper, I will refer to the Tolerance and Sufficiency Principles collectively as the Tolerance Principle.

While the Tolerance Principle provides a way to evaluate a hypothesized rule for productivity, we must first hypothesize such a rule to evaluate. Previous work using the Tolerance Principle has developed learning models that make use of recursive learning with greedy subdivision: Payne (2022, 2023a) learn which inflectional categories are marked in a language by tracking "collisions" between multiple inflected forms, and Belth et al. (2021) learn inflectional processes via greedy, frequency-based subdivision. The Tolerance Principle has been applied to a wide range of morphological phenomena, including gender assignment in Icelandic (Björnsdóttir 2021), noun diminutive suffixation in Dutch (van Tuijl & Coopmans 2021), verbal inflection variation and change in Frisian (Merkuur 2021), possessive suffixation in Northern East Cree (Henke 2022), and past participle inflection in Latin (Kodner 2022). Though I focus on morphology in the current work, similar Tolerance Principle-based learning models have been proposed for phenomena in phonology (Dresher & Lahiri 2022; Belth 2023; Payne 2023b, i.a.) and syntax (Pearl & Sprouse 2021; Liang et al. 2022; Li & Schuler 2023, i.a.). Though these approaches generally make use of quantitative greedy heuristics such as frequency, one could conceive that subdivision in these learners is instead driven by the preferences and preference rankings of Natural Morphology, so that, for example, the most iconic form will be prioritized. This foreshadows the model we will develop in Section 3.

Payne & Yang (2023) argue that the missing piece in current theories of BADs – including the account put forth by Dressler et al. (2020) – is the presence of a precise, mechanistic *learning theory.* In other words, though Dressler et al. account for observed BAD patterns in terms of Natural Morphology's preferences, they do not specify the precise mechanism(s) by which the morphological system is built, or why children enter and exit BADs precisely when they do, and their account thus largely fails to satisfy Desideratum (1). The Tolerance Principle, on the other hand, provides precisely such a learning theory: it gives a "tipping point" where we can expect a form to become productive or to lose productivity, as calibrated over the learner's internal vocabulary. Payne & Yang posit that for both types of BADs, the error can only emerge once the child has learned that the relevant category must be marked in their language, and can persist either until it loses productivity – in the case of a weak BAD – or until another form becomes productive – in the case of a strong

BAD. Since the Tolerance Principle makes precise predictions about when each of these events will occur, Payne & Yang's account satisfies Desideratum (1).

For weak BADs, Payne & Yang account for the form of the BAD in terms of productivity: the given form must be productive over the learner's *internal* vocabulary in order to become a weak BAD. Because the Tolerance Principle again gives precise predictions about the conditions under which a form will be productive, this account of weak BADs also satisfies Desideratum (2). In accounting for strong BADs, however, Payne & Yang argue that these errors do not constitute "inventions out of thin air" but are rather grounded somewhere in the input data. For example, reduplication appears as a formal pattern in both Russian and Greek (Israeli 1997; Kallergi 2015). As such, Payne & Yang answer the Feyerabendian question of whether "everything goes" not in terms of *cross-linguistic* attestation, but rather attestation of a given form in the child's input: the set of possible strong BADs predicted by Payne & Yang is thus a strict *subset* of those predicted by Dressler et al.. Of course, it remains an open question whether all strong BADs employ forms that are present somewhere in the child's input – and indeed, the Greek vowel lengthening BAD discussed in Section 1.2 provides a potential counterexample – or whether the forms are more generally present in *some* language as predicted by Dressler et al. In any case, it is not enough for a given strong BAD form to be attested either in the child's input or cross-linguistically; we also want to understand *why* the child selects this form over other possible forms when they enter the strong BAD. Here, the Tolerance Principle has little to say, and the Payne & Yang account thus struggles to satisfy Desideratum (2) in the case of strong BADs.

3. Proposal: A mechanistic account of Blind Alley Developments

As discussed above, the proposals of Dressler et al. (2020) and Payne & Yang (2023) have complementary strengths; the current proposal essentially constitutes a marriage of these strengths. In terms of Desideratum (1), the account given by Payne & Yang places precise, mechanistic bounds on the timeline of BADs grounded in a formal, quantitative account of language acquisition given by the Tolerance Principle. The Tolerance Principle also provides a

measure of the critical mass of opposing input that will lead to the demise of a BAD. By contrast, the account of Dressler et al. simply states that children escape from BADs as a result of the opposing input. By definition, however, the child's input will always be opposing during a BAD, and Dressler et al. do not provide an account of why children escape from their Blind Alleys *precisely when they do*. As such, the account of Payne & Yang is better able to satisfy Desideratum (1).

In terms of Desideratum (2), however, the account of Payne & Yang has less to say. For weak BADs, this model determines which of the forms in the learner's input are theoretically possible BADs based on a precise calibration of productivity over the learner's internal vocabulary. However, it has little to say about which of the possible weak BADs a learner may choose to pursue first. While it is not strictly necessary to order the possible BADs being considered for the Payne & Yang account to succeed, it is unlikely that the child is simultaneously evaluating every possible form in their input for productivity. What's more, for strong BADs, the Payne & Yang account simply predicts that the given BAD pattern must appear somewhere in the child's input but does not make predictions regarding which forms in the child's input are likely to be posited as strong BADs. By contrast, the preferences and preference rankings of Natural Morphology allow the account of Dressler et al. to make predictions about which weak BADs the child may pursue first and which possible form the child may select for a strong BAD, although it does so without a precise formalization of this process. As such, the account of Dressler et al. is better able to satisfy Desideratum (2).

How, then, can we marry these accounts into a unified theory of BADs that satisfies both Desiderata from Section 1.3? I propose an account in which the Tolerance Principle sets strict, quantitative boundaries on the timeline of strong and weak BADs and determines which potential weak BADs can be generalized, while Natural Morphology's preferences guide the learner to consider some possible weak BADs before others and dictate what patterns are employed in strong BADs. Under this proposal, the start of both strong and weak BADs occurs sometime after the child learns that the given function must be marked in their language, following Payne (2022, 2023a) and Payne & Yang 2023. Specifically, the weak BAD emerges when the child knows that the category in question is marked and some form in their input is temporarily productive over their *internal* vocabulary, and a strong BAD will emerge when no form is productive over the learner's internal vocabulary, but the learner knows that the category must be marked. The weak BAD's demise will occur

when its productivity can no longer be supported over the learner's internal vocabulary as determined by the Tolerance Principle, and the strong BAD's demise will occur when some other process – either a weak BAD or the adult-like form – becomes productive over their internal vocabulary. As such, the Tolerance Principle sets the boundary conditions on the timeline of the BADs in a way that can be precisely calculated, and Desideratum (1) is satisfied.

Under the current proposal, once the child has learned that a category must be marked in their language, the preferences of Natural Morphology will guide them to consider some possible forms to mark this function over others. For example, if there are 5 possible forms in the input that may be used to mark a given function, but only one of these satisfies the preference for iconicity, then the child will consider this one first as a possible form. This can be seen as an adaptation of the greedy search heuristics used in models such as Belth et al. (2021): instead of determining which form we consider based on measures such as type frequency, we instead make use of cognitively-based preferences to order the hypotheses we consider. The possible forms will be evaluated against the Tolerance Principle in the order dictated by the preferences of Natural Morphology, and if a form is productive over the learner's internal vocabulary, then a weak BAD emerges. Because we consider first the forms that are favored by Natural Morphology's preferences, our proposal accounts for Dressler et al. (2020)'s observation that the attested weak BAD patterns obey preferences such as iconicity, while at the same time providing a mechanistic account – via the Tolerance Principle – of why these forms and not others enjoy a temporary period of productivity. If none of the forms being considered are productive under the Tolerance Principle, we can expect the learner to construct a strong BAD: they know that the given function must be marked, and thus invent a way to mark it. Here, Natural Morphology's preferences once more come into play: the preference for iconicity, for example, can explain the use of reduplication in the BADs discussed in Section 1.2. While it remains an open question whether the form used in the strong BAD must appear *somewhere* in the child's input (cf. Dressler et al. 2020, Payne & Yang 2023), the preferences of Natural Morphology can be employed in either case to determine the constructions that the child will make during the strong BAD. At the same time, the learning-theoretic account given by the Tolerance Principle predicts the conditions under which we expect strong BADs to occur and when we expect the child to escape from them. In this way, the combination of the Tolerance Principle and Natural Morphology satisfies Desideratum (2).

Having outlined the proposed account of BADs, we now explore its implications for weak and strong BADs, using the examples of the English to ɪ → æ/_ŋ# BAD and the Greek subjunctive BAD.

4. Weak BADs under our proposal

The model of Blind Alley Developments introduced above straightforwardly accounts for weak BADs while satisfying the desiderata for an adequate theory of BADs outlined in Section 1.3. Desideratum (1) is satisfied by the learning-theoretic portion of our model: the Tolerance Principle sets precise, mechanistic bounds on the start and end time of the BAD. Specifically, the weak BAD will emerge when the learner knows that the category must be marked – following Payne (2022) – and will end when the BAD process is no longer productive over the learner's internal lexicon. Desideratum (2) is satisfied by a combination of the learning-theoretic portion of our model and Natural Morphology's preferences. The Tolerance Principle sets mechanistic bounds on the types of weak BADs we may theoretically encounter: a BAD must be sufficiently dominant over the learner's internal vocabulary to be overextended. At the same time, the preferences and preference rankings of Natural Morphology can be seen as guiding the children towards some of these possible weak BADs over others. I illustrate these points with the concrete example of the English past tense ɪ → æ/_ŋ# weak BAD. I choose this example because the acquisition of the English past tense has been widely studied (e.g., Marcus et al. 1992, Xu & Pinker 1995), and we thus have access to the accurate estimates of vocabulary size necessary to employ the Tolerance Principle.

4.1. Modelling ɪ → æ/_ŋ#

To model the progression of the ɪ → æ/_ŋ# BAD, we must first establish a plausible model of the growth of the learner's vocabulary of English verbs over time. Previous work has demonstrated a significant correlation between token frequency and order of acquisition (e.g., Goodman et al. 2008): more frequent words are likely to be learned earlier. Thus, to model the acquisition trajectory of a "typical" English-learning child, it is sufficient to use a strictly frequency-based ordering, where token frequencies are calculated by aggregation over all North American CHILDES corpora (MacWhinney 2000). Note

that while we are modeling only one acquisition trajectory here, there is room for variation in the ɪ → æ/_ŋ# BAD resulting from differing acquisition trajectories. Indeed, as I will demonstrate below, entering the ɪ → æ/_ŋ# BAD requires learning the three most frequent -ɪŋ# verbs before learning enough regular verbs such that *-ed* becomes productive, and before learning too many -ɪŋ# verbs that do not obey ɪ → æ/_ŋ#. Thus, our model accounts for individual variation as a consequence of differences in order of acquisition: since the Tolerance Principle calculations are made over the learner's *internal* vocabulary, the state of their vocabulary at any time determines whether a given process – in this case a BAD – will be productive. Having established our frequency-based model of order of acquisition, we turn to our model's predictions for the timeline and content of the ɪ → æ/_ŋ# BAD. We begin by examining the model's predictions regarding the timeline of the BAD and then turn to the content of this weak BAD, asking why -æŋ# BADs but not -ɔt# BADs (e.g., *stink*-stought from *think-thought*) are attested.

4.1.1. Timeline

Our model asserts that the learner will know that the past tense *must* be marked in English upon entering the ɪ → æ/_ŋ# BAD. Under the learning-theoretic model developed by Payne (2022), the learner will know that English past tense marking is obligatory by the time their vocabulary contains approximately 112 verbs. Of the 200 most frequent verbs in North American CHILDES, 76 are irregular. Under the Tolerance Principle, this means that the default *-ed* rule will not be learned: $\theta_{200} = 37$, and the number of irregular verbs is thus too large to be tolerated as exceptions to the *-ed* rule. Following Belth et al. (2021), the learner then subdivides the input in search of a productive process. In doing so, the learner will encounter three verbs ending in -ɪŋ#: *bring-brought, sing-sang,* and *ring-rang* are contained in the 200 most frequent verbs being considered. Under the Tolerance Principle, 2/3 is sufficient for generalization, and ɪ → æ/_ŋ# is productive over the learner's *internal* vocabulary which contains only these three -ɪŋ# verbs. As such, the learning-theoretic portion of our model accounts neatly for the emergence of the ɪ → æ/_ŋ# BAD.

However, the productivity of ɪ → æ/_ŋ# will be short lived: of the most frequent 800 verbs in North American CHILDES, we have 8 ending in -ɪŋ#:

bring-brought, sing-sang, ring-rang, fling-flung, spring-sprang, sting-stung, swing-swung, and *wing-winged.* Of these 8 verbs, only 3 obey ɪ → æ/_ŋ#, and thus there are 5 exceptions to this rule. Under the Tolerance Principle, however, we can only tolerate three exceptions when *N*=8; that is, $\theta_8 = 3 < 5$. As such, when the early vocabulary contains 800 verbs, the ɪ → æ/_ŋ# BAD is no longer supported by the learner's internal vocabulary and must be abandoned. The learning-theoretic portion of our model thus also accounts neatly for the subsequent demise of the ɪ → æ/_ŋ# BAD and its ephemeral nature, satisfying Desideratum (1). Again, recall that the Tolerance Principle allows for individual variation in the timeline of BADs based on variation in the order of acquisition: if a child were to instead reach a productive *-ed* rule sooner, or learn e.g., *bring*, *fling*, and *wing* as their first three -ɪŋ# verbs, then the Tolerance Principle makes the testable prediction that they would not enter this BAD at all. Our model can thus also account for the fact that this BAD does not affect every child learning English, as a consequence of differences in vocabulary acquisition. A similar approach can be applied to language change: Ringe & Yang (2022) demonstrate that the Tolerance Principle can account for the historical productivity of ɪ → æ/_ŋ#. Their approach illustrates the delicate nature of productivity: changes in only a few verbs can push the productivity of ɪ → æ/_ŋ# over the "tipping point" defined by the Tolerance Principle.

4.1.2. Does everything go?

Why are ɪ → æ/_ŋ# BADs widely attested and -ɔt# (e.g., *stink*-stought from *think-thought*) BADs unattested? One part of the explanation comes from the learning-theoretic portion of our proposal, which predicts that even if an -ɔt# BAD is hypothesized by the learner, it will never be able to reach productivity over the learner's internal vocabulary. The first challenge arises when we attempt to define a context in which the -ɔt# BAD would apply: while defining the context for the -æŋ# BAD is straightforward – namely, verbs that end in -ɪŋ# in the present – this is not the case of the -ɔt# BAD. The two most frequent -ɔt# past tense verbs – *bring-brought* and *think-thought* – may lead us to hypothesize a rule of the form ŋ(k) → ɔt/ _ #, where (k) indicates optionality of the voiceless velar stop in defining the context of this rule. However, it is clear that such a rule has no hope of reaching productivity: its description encompasses not only all verbs obeying ɪ → æ/_ŋ# (e.g. *sing-sang, ring-rang*), but

also all those obeying ɪ → æ/_ŋk# (*drink-drank, sink-sank*). Indeed, of the most frequent 250 verbs in North American CHILDES, there are 6 ending in -ɪŋ(k)#: *think-thought, bring-brought, drink-drank, ring-rang, sing-sang,* and *sting-stung.* Of these, only 2 – *think* and *bring* – obey ɪŋ(k) → ɔt/ _ #, leaving the remaining 4 verbs as exceptions to this potential rule. Under the Tolerance Principle, we have that $\theta_6 = 3 < 4$, and thus ɪŋ(k) → ɔt/ _ # cannot reach productivity.

One might hypothesize that increasing the learner's vocabulary would introduce the possibility of a productive -ɔt# BAD, but reality is quite the opposite. The next most frequent -ɔt# verbs are *catch-caught* and *teach-taught*, leaving us entirely without a viable context for a unified -ɔt# rule. What's more, while the presence of *catch* might tempt us to subdivide and posit a rule along the lines of æ$\widehat{\text{tʃ}}$ → ɔt/_#, this rule will also not be able to reach productivity: *scratch-scratched, hatch-hatched,* and *attach-attached* are all similarly frequent and $\theta_4 = 2 < 3$, so æ$\widehat{\text{tʃ}}$ → ɔt/_# can never reach productivity. Similarly, introducing *teach* may lead us to hypothesize a rule of the form i$\widehat{\text{tʃ}}$ → ɔt/_#, but the similarly frequent *reach* would render such a rule unproductive. As such, the learning-theoretic portion of our model predicts that an -ɔt# BAD should never reach productivity.

While the Tolerance Principle tells us that an -ɔt# BAD, *if posited*, will not reach productivity, the preferences of Natural Morphology allow us to go a step further and posit that such a BAD may not be hypothesized in the first place. Since a weak BAD consists of the overapplication of a form that is present in the input, this entails selecting *which* form the learner will attempt to overextend. In this instance, the Natural Morphology preference for *bi-uniqueness* may steer the learner towards the -æŋ# BAD over the -ɔt# one. Specifically, most English verbs that take -æŋ# in the simple past take -ʌŋ# in the past participle (e.g., *I sing/I sang/I had sung*). This contrasts with -ɔt# past tense forms, which generally have a syncretism between the simple past and past participle (e.g., *I think/I thought/I had thought*). The fact that -æŋ# past tense verbs explicitly distinguish between the simple past and past participle may make the ɪ → æ/_ŋ# BAD a more appealing one for the learner to pursue, and thus the -ɔt# BAD may simply never be posited. The role of Natural Morphology's preferences and preference rankings is thus to guide the learner towards pursuing some possible weak BADs over others, while the Tolerance Principle sets quantitative bounds on which of these weak BADs may enjoy a transient stage of productivity. Thus, we satisfy Desideratum (2).

5. Strong BADs under our proposal

As discussed above, for weak BADs, it is the case that *some* form in the learner's input is temporarily productive over their internal vocabulary: the presence of *bring, sing*, and *ring*, for example, allow for the development of an ɪ → æ/_ŋ# weak BAD. However, in the case of strong BADs, no form in the learner's input can directly explain their constructions. The timeline of strong BADs can be explained once again by the learning-theoretic portion of our model: as for weak BADs, strong BADs emerge when the learner knows that a given category must be marked, but not yet how the category is marked. In contrast to weak BADs, strong BADs in our model occur because *no* form in the learner's input reaches productivity. The strong BAD can continue only until some process in the input becomes sufficiently dominant over the learner's internal vocabulary, either causing the learner to enter a weak BAD or to acquire the adult inflectional process. Thus, our model satisfies Desideratum (1) from Section 1.3 for strong BADs as well. To satisfy Desideratum (2), however, we must turn to the preferences of Natural Morphology. By hypothesis, the learner knows that the relevant category must be marked, but as of yet has no productive process with which to mark it. Natural Morphology then predicts which types of constructions children are likely to build: the *iconicity* preference, for example, predicts the use of reduplication in the strong BADs discussed in Section 1.2.

5.1. Modeling the Greek Subjunctive BAD

It is difficult to provide a precise Tolerance Principle analysis the strong BADs presented by Dressler et al. (2020) because doing so requires very accurate vocabulary measures, as discussed in Section 4.1, and these measures are not readily available for the children in these studies – even full transcripts only capture a subset of the children's vocabulary. While approximating vocabulary growth via frequency works for the well-studied problem of the ɪ → æ/_ŋ# weak BAD, it is more challenging to access frequency counts for the strong BADs discussed in Section 1.2, and the children who followed these strong BADs likely did not have a vocabulary acquisition trajectory determined solely by frequency. Here I will thus sketch how our approach would apply to these strong BADs were such quantitative vocabulary measures to become available, though I also refer the reader to Payne (2023c) for a quantitative

treatment of the well-studied Root Infinitive, which may be analyzed as a strong BAD (Payne & Yang 2023), in terms of the account I present here.

Our model predicts that the Greek-learning child studied by Dressler et al. (2020) will only begin producing the subjunctive BAD once he has learned that the subjunctive *must* be marked in Greek. Though the model of Payne (2022) has yet to be applied to Greek verbs, the logic is the same: when a sufficient number of verbs in the learner's internal vocabulary appear in the subjunctive and in a *different* form in a different inflectional class, then they learn that the subjunctive must be marked in Greek. Following the same procedure as discussed in Section 4.1.2, the learner will then consider the processes that have been observed marking the subjunctive in the input in an order determined by the preferences and preference rankings of Natural Morphology, evaluating each over their internal vocabulary for productivity. Our account predicts that the learner enters the strong BAD because *none* of the processes they consider are productive over their internal vocabulary. In other words, even if the proclitic *na-* is present in the learner's input, it is not sufficiently dominant over the verbs that the learner has lexicalized in the subjunctive. Natural Morphology argues (c.f., Dressler et al. 2020) that proclitics are challenging for the child to take from their input into the uptake; it is thus possible that the child may have seen enough verbs taking *na-* in the subjunctive in their input but failed to lexicalize enough of these to learn the Greek subjunctive process. As such, they know that they must systematically differentiate the subjunctive but have no productive process with which to do so.

It is here that Natural Morphology takes center stage: as argued by Dressler et al. (2020), the learner must construct *some* process to systematically differentiate the marked category, and the preferences of Natural Morphology can be seen as guiding them towards some constructions over others. Specifically, both the vowel lengthening BAD and the reduplication BAD produced by the Greek child can be predicted by the *iconicity* preference under Natural Morphology, with the latter being more iconic than the former within this framework. Payne & Yang (2023) argue that the process the child selects must be attested *somewhere* in their input as a formal pattern, and that strong BADs can thus be construed as a mismatch between form and function. While it is true that reduplication occurs as a formal pattern in Greek (Kallergi 2015), it is less clear whether this is a possibility for the vowel lengthening BAD, since Greek has only emphatic vowel lengthening. Future research should further investigate whether strong BAD constructions by children correspond to forms that are attested *cross-linguistically* – as predicted by Natural Morphology

(Dressler et al. 2020) – or *within* the language the child is learning – as predicted by Payne & Yang (2023). Even if the latter, more constrained, hypothesis is confirmed, Natural Morphology still plays a crucial role in determining *which* of the patterns attested in the language the child will choose for their strong BAD, and thus the iconicity preference can still be used to explain the two BADs the Greek-learning child constructs.

Under our model, the strong BAD will persist only until the child learns a productive process to mark the Greek subjunctive. That is, once the *na-* proclitic is sufficiently dominant over the learner's internal vocabulary, the strong BAD will be abandoned and the learner will produce the adult-like subjunctive. Any gradience in the abandonment of the BAD can be accounted for in terms of a stage of memorization before the learner generalizes the *na-* rule. For example, if a child knows 10 verbs and knows that 4 of them take *na-*, they will not be able to generalize *na-* under the Tolerance Principle, but they may still use this process rather than the strong BAD when inflecting these specific verbs, leading to an apparent gradience in their production of this strong BAD. Note that it is also possible under our model that the process that becomes productive over the learner's internal vocabulary at the end of the strong BAD is not in fact the adult-like process, in which case we expect the learner to move from a strong BAD into a weak BAD before acquiring the adult process.

6. Discussion and future directions

In this paper, I have argued that a sufficient theory of Blind Alley Developments in acquisition should not simply *allow* for BADs, but should rather provide causal mechanisms that make precise predictions about (1) the timeline and (2) the content of these errors. Specifically, I claim that an adequate theory of BADs in acquisition must be a *learning theory* that precisely accounts for the timeline and content of BADs in terms of the learning mechanisms that the child employs during acquisition. I have presented such an account by marrying the complementary approaches of the Tolerance Principle (Payne & Yang 2023) and Natural Linguistics (Dressler et al. 2020). This account provides both a formalization of the timeline of BADs – in terms of calibration of productivity over the learner's internal vocabulary – and predictions about what types of BAD patterns we expect to find cross-linguistically. In particular, we account for the timeline of weak and strong BADs in terms of the

Tolerance Principle, which provides precise quantitative predictions regarding the productivity of a given process over the learner's internal vocabulary, allowing for variation in acquisition timelines as a result of variation in vocabulary growth. We account for the content of weak BADs in terms of a combination of the Tolerance Principle and Natural Morphology: while the Tolerance Principle places precise bounds on which forms are *eligible* to be productive over the learner's vocabulary, the preferences and preference rankings of Natural Morphology can be seen as guiding the learner to some of these potential weak BADs over others. In the case of strong BADs, Natural Morphology takes center stage in explaining the patterns that children construct, since the patterns they choose are *never* attested marking the relevant categories in the input.

While the account presented here satisfies the desiderata outlined in Section 1.3, it also raises several questions for future research. Firstly, an important divergence between the accounts of strong BADs put forth by Dressler et al. (2020) and Payne & Yang (2023) is that the former predicts that the forms children use during strong BADs must be attested somewhere *cross-linguistically,* while the latter predicts that the forms must be attested *within* the language that the child is learning. As such, the potential strong BAD patterns predicted by Payne & Yang are a strict subset of those predicted by Dressler et al., since Payne & Yang argue that the form used by the child must have been seen somewhere in the input, albeit not marking the target function. The documented cases of strong BADs (Dressler et al. 2020) are as of yet inconclusive: since reduplication occurs as a formal pattern in both Russian and Greek, and since Greek does allow emphatic vowel lengthening, one can argue that the relevant forms are attested in the child's input. However, Greek does not employ vowel lengthening as a part of its morphological system, and the reduplication that does occur in Greek is total rather than partial (Kallergi 2015); one could thus argue that the relevant forms are not attested exactly in the child's input. Future explorations of strong BADs should further investigate the role of attestation of a given form in the child's strong BAD constructions.

Secondly, our account of strong BADs rests on the assumption that the child will exit the strong BAD once they learn the adult-like productive inflectional process. However, it is not always the case that the adult-like inflectional process is productive: gaps, defectivity, and other forms of ineffability arise when no process is productive (e.g., Yang 2016: Ch. 5; Gorman & Yang 2019). If the child is learning a language with such a non-productive inflection, an open question is how they escape from the strong BAD into the adult-like

process: is it simply by memorizing the relevant forms as more are seen, or is there something else that triggers their abandonment of the strong BAD? Future work should investigate the possible presence of strong BADs in the development of children acquiring inflectional processes that will not be productive in their adult grammar and build on the account presented here to account for how children escape the strong BADs in this case.

Acknowledgements

I am grateful to Charles Yang and Wolfgang U. Dressler, with whom I developed various parts of the account I present here, as well as Mark Aronoff, Patricia Donegan, Katarzyna Dziubalska-Kołaczyk, and attendees of the 2023 Poznań Linguistics Meeting, for thoughtful feedback which greatly improved this account. This work was supported by the Institute for Advanced Computational Science Graduate Research Fellowship and the National Science Foundation Graduate Research Fellowship Program under NSF Grant No. 2234683.

References

Albright, Adam & Bruce Hayes. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* 90(2). 119–161.

Allen, Shanley. 1996. *Aspects of argument structure acquisition in Inuktitut*. Amsterdam: John Benjamins.

Bat-El, Outi. 2014. *The Acquisition of Hebrew Phonology and Morphology*. Leiden: Brill.

Belth, Caleb. 2023. *Towards an Algorithmic Account of Phonological Rules and Representations.* Doctoral Dissertation, University of Michigan.

Belth, Caleb, Sarah Payne, Deniz Beser, Jordan Kodner & Charles Yang. 2021. The greedy and recursive search for mophological productivity. *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society* 42. 2869–2875.

Bittner, Dagmar, Wolfgang U. Dressler & Marianne Kilani-Schoch (eds.). 2003. *Development of Verb Inflection in First Language Acquisition: A Cross-Linguistic Perspective*. Berlin: de Gruyter Mouton.

Björnsdóttir, Sigríur Mjöll. 2021. Productivity and the acquisition of gender. *Journal of Child Language* 48. 1209–1234.

Bybee, Joan. 1985. *Morphology: a study of the relation between form and meaning.* Amsterdam: John Benjamins.

Bybee, Joan. 2010. *Language usage and cognition*. Cambridge: Cambridge University Press.
Caprin, Claudia, and Maria Teresa Guasti. 2009. The acquisition of morphosyntax in Italian: A cross-sectional study. *Applied Psycholinguistics* 30(1). 23–52.
Clahsen, Harald, Monika Rothweiler, Andreas Woest, and Gary F. Marcus. 1992. Regular and irregular inflection in the acquisition of German noun plurals. *Cognition* 45(3). 225–255.
Clahsen, Harald, and Monika Rothweiler. 1993. Inflectional rules in children's grammars: Evidence from German participles. *Yearbook of Morphology 1992*. 1–34.
Clahsen, Harald, Fraibet Aveledo, and Iggy Roca. 2002. The development of regular and irregular verb inflection in Spanish child language. *Journal of child language* 29 (3). 591–622.
Deen, Kamil. 2005. *The acquisition of Swahili*. Amsterdam: John Benjamins.
Demuth, Katherine. 2003. The acquisition of Bantu languages. *The Bantu Languages*. 209–222. Taylor and Francis.
Dresher, B. Elan & Aditi Lahiri. 2022. The foot in the history of English: Challenges to metrical coherence. In Los, Bettelou, Claire Cowie, Patrick Honeybone, & Graeme Trousdale (eds.) *English historical linguistics:*
Change in structure and meaning. Amsterdam: John Benjamins. 42–59.
Dressler, Wolfgang U. 1995. Interactions between iconicity and other semiotic parameters in language. In Simone, Raffaele (ed.), *Iconicity in Language*. Amsterdam: John Benjamins. 21–37.
Dressler, Wolfgang U. & Karpf, Annemarie. 1995. The theoretical relevance of pre- and protomorphology in language acquisition. *Yearbook of Morphology 1994*. 99–122.
Dressler, Wolfgang U., Gary Libben & Katharina Korecky-Kröll. 2014. Conflicting vs. convergent vs. interdependent motivations in morphology. In MacWhinney, Brian, Andrej Malchukov, & Edith Moravcsik (eds.), *Competing Motivations in Grammar and Usage*. Oxford: Oxford University Press. 181–196.
Dressler, Wolfgang U. & Marianne Kilani-Schoch. 2016. Natural Morphology. In Hippisley, Andrew & Gregory Stump (eds.), *The Cambridge Handbook of Morphology*. Cambridge: Cambridge University Press. 356–389
Dressler, Wolfgang U., Anastasia Christofidou, Natalia Gagarina, Katharina Korecky-Kröll & Marianne Kilani-Schoch. 2020. Morphological blind-alley developments as a theoretical challenge to both usage-based and nativist acquisition models. *Italian Journal of Linguistics* 31(2). 107–140.
Dressler, Wolfgang U., Anastasia Christofidou, Natalia Gagarina, Katharina Korecky-Kröll & Marianne Kilani-Schoch. 2023. Blind Alley Developments (BADs): in defense of our approach. *Italian Journal of Linguistics*. 35(2).
Dressler, Wolfgang U. and Sarah Payne. 2023. Self-Organization in Acquisition. *Cambridge Handbook of Natural Linguistics.*
Emond, Emeryse, & Rushen Shi. 2021. Infants' rule generalization is governed by the Tolerance Principle. In Dionne, Danielle & Lee-Ann Vidal Covas (eds.): *Proceedings of the 45nd annual Boston University Conference on Language Development.* 191–204.

Engelmann, Felix, Sonia Granlund, Joanna Kolak, Marta Szreder, Ben Ambridge, Julian Pine, Anna Theakston & Elena Lieven. 2019. How the input shapes the acquisition of verb morphology: Elicited production and computational modelling in two highly inflected languages. *Cognitive Psychology* 110. 30–69.

Goodman, Judith, Philip Dale & Ping Li. 2008. Does frequency count? Parental input and the acquisition of vocabulary. *Journal of Child Language* 35. 515–531.

Gorman, Kyle & Charles Yang. 2019. When nobody wins. *Competition in inflection and word-formation*. 169–193.

Henke, Ryan E. 2022. Rules and exceptions: A Tolerance Principle account of the possessive suffix in Northern East Cree. *Journal of Child Language* 50(5). 1–36.

Israeli, Alina. 1997. Syntactic reduplication in Russian: A cooperative principle device in dialogues. *Journal of Pragmatics* 27. 587–609.

Kallergi, Haritini. 2015. *Reduplication at the Word Level*. The Hague: De Gruyter Mouton.

Karpf, Annemarie. 1990. *Selbstorganisationsprozesse in der sprachlichen Ontogenese: Erst- und Fremdsprache(n).* Tübingen: Narr.

Kauschke, Christina, Anna Kurth, and Ulrike Domahs. 2011. Acquisition of German noun plurals in typically developing children and children with specific language impairment. *Child Development Research* 2011. 1–17.

Kilani-Schoch, Marianne & Wolfgang U. Dressler. 2005. *Morphologie naturelle et flexion du verbe français.* Tübingen: Narr.

Kodner, Jordan. 2022. Language acquisition guiding theory and diachrony: A case study from Latin morphology. *Natural Language and Linguistic Theory* 41. 1–60.

Korecky-Kröll, Katharina, Sabine Sommer-Lolei & Wolfgang U. Dressler. 2017. Emergence and early development of German compounds. In Wolfgang U. Dressler, Nihan Ketrez & Marianne Kilani-Schoch (eds.), *Nominal Compound Acquisition.* Amsterdam: Benjamins. 19–37.

Li, Daoxin & Kathryn D. Schuler. 2023. Acquiring recursive structures through distributional learning. *Language Acquisition*. 1–14.

Liang, Kevin, Diana Marsala & Charles Yang. 2022. Distributional learning of syntactic categories. *Proceedings of the 46th Boston University Conference on Language Development.* 442–455.

Libben, Gary, Monika Boniecki, Marlies Martha, Karin Mittermann, Katharina Korecky-Kroll & Wolfgang U. Dressler. 2009. Interfixation in German compounds: What factors govern acceptability judgements? *Italian Journal of Linguistics* 21(1). 149–180.

Lieven, Elena. 2008. Learning the English auxiliary: A usage-based approach. In Behrens, Heike (ed.), *Corpora in. Language Acquisition Research: History, Methods, Perspectives.* Amsterdam: John Benjamins Publishing. 25–60.

Lustigman, Lyle. 2013. Developing structural specification: Productivity in early Hebrew verb usage. *First Languag*e 33(1). 47–67.

MacWhinney, Brian. 2000. *The CHILDES project: Tools for analyzing talk*. Mahwah, NJ: Psychology Press.

Maratsos, Michael. 2000. More overregularizations after all: new data and discussion on Marcus, Pinker, Ullman, Hollander, Rosen & Xu. *Journal of Child Language* 27(1). 183–212.

Marcus, Gary, Steven Pinker, Michael Ullman, Michelle Hollander, John Rosen & Fei Xu. 1992. Over-regularization in language acquisition. *Monographs of the Society for Research in Child Development* 57. Chicago: University of Chicago Press. 1–182.

Mariscal, Sonia. 2009. Early acquisition of gender agreement in the Spanish noun phrase: starting small. *Journal of Child Language* 36. 143–171.

Maslen, Robert, Anna Theakston, Elena Lieven & Michael Tomasello. A dense corpus study of past tense and plural overregularization in English. *Journal of Speech, Language, and Hearing Research* 47. 1319–1333.

Mayol, Laia. 2007. Acquisition of irregular patterns in Spanish verbal morphology. *Proceedings of the 2007 European Summer School in Logic, Language, and Information*: 185–196.

Merkuur, Anne. 2021. *Changes in modern Frisian verbal inflection*. Doctoral Dissertation, University of Amsterdam.

Munson, Joy & Ingram, David. 1985. Morphology before syntax: A case study from language acquisition. *Journal of Child Language* 12. 681–684.

Palermo, David, and V. Lynn Eberhart. 1968. On the learning of morphological rules: An experimental analogy. *Journal of Verbal Learning and Verbal Behavior* 7(2). 337–344.

Pater, Joe, Christine Stager & Janet Werker. 2004. The perceptual acquisition of phonological contrasts. *Language* 80. 38–402.

Payne, Sarah. 2022. *When collisions are a good thing: The acquisition of morphological marking*. Bachelor's thesis, University of Pennsylvania.

Payne, Sarah. 2023. Contrast, Sufficiency, and the Acquisition of Morphological Marking. *Proceedings of the 47th Boston University Conference on Language Development* 47(1). 604–617.

Payne, Sarah. 2023. Marginal Sequences are Licit but Unproductive. *2023 Annual Meeting of Phonology.* Poster.

Payne, Sarah. 2023. Root Infinitives and the Acquisition of Morphological Marking. *Penn Working Papers in Linguistics* 30(1). In press.

Payne, Sarah & Charles Yang. 2023. Making Good on BADs. *Italian Journal of Linguistics*. 35(2). In press.

Pearl, Lisa, and Jon Sprouse. 2021. The acquisition of linking theories: A Tolerance and Sufficiency Principle approach to deriving UTAH and rUTAH. *Language Acquisition* 28. 294–325.

Pinker, Steven & Alan Prince. 1988. On language and connectionism: analysis of a parallel distributed processing model of language acquisition. *Cognition* 28. 73–193.

Ravid, Dorit & Rola Farah. 1999. Learning about noun plurals in early Palestinian Arabic. *First Language* 19(56). 187–206.

Ringe, Don & Charles Yang. 2022. The threshold of productivity and the irregularization of verbs in Early Modern English. In Los, Bettelou, Claire Cowie, Patrick Honeybone & Graeme Trousdale (eds.): *English historical linguistics: Change in structure and meaning.* Amsterdam: Benjamins. 91–111.

Rumelhart, David & James McClelland. 1986. On learning the past tenses of English verbs. In Rumelhart, David, James McClelland, & The PDP Research Group (eds),

Parallel distributed processing: explorations in the microstructure of cognition. Cambridge, MA: MIT Press. 216–271.
Schuler, Kathryn. 2017. *The acquisition of productive rules in child and adult language learners*. Doctoral Dissertation, Georgetown University.
Tomasello, Michael. 2000. Do young children have adult syntactic competence? *Cognition* 74. 209–253.
van Tuijl, Rosita & Peter Coopmans. 2021. The productivity of Dutch diminutives. *Linguistics in the Netherlands* 38. 128–143.
Yang, Charles. 2013. Ontogeny and phylogeny of language. *Proceedings of the National Academy of Sciences* 110. 6324–6327.
Yang, Charles. 2005. On productivity. *Linguistic Variation Yearbook* 5. 333–370.
Yang, Charles. 2016. *The Price of Linguistic Productivity: How Children Learn to Break the Rules of Language.* Cambridge, MA: MIT Press.
Xu, Fei & Steven Pinker. 1995. Weird past tense forms. *Journal of Child Language* 22(3). 531–556.

Corresponding author:
Sarah Brogden Payne
Social and Behavioral Sciences
Stony Brook University
100 Nicolls Rd
Stony Brook
NY11794
USA