

Making good on BADS

Sarah Payne,^a Charles Yang^b

^a Department of Linguistics and Institute for Advanced Computational Science, Stony Brook University, <sarah.payne@stonybrook.edu>

^b Departments of Linguistics and Computer Science, University of Pennsylvania <charles.yang@ling.upenn.edu>

KEYWORDS: morphological acquisition, Natural Morphology, Tolerance Principle, formal and substantive universals, Root Infinitives, reduplication.

1. BADS and their implications

Dressler *et al.* (2019 – henceforth **D&al**) report several cases of rather unusual patterns in children’s morphological acquisition. These patterns are dubbed Blind-Alley Developments (**BADs**) because of their striking deviation from the input as well as their transitory nature, eventually to be replaced with adult-like forms.

D&al discuss two kinds of BADs: we interpret both as misalignments of form and function in children’s morphology. Weak BADs involve the use of an incorrect form that nevertheless marks the correct target function. D&al give an example of the German interfix *-e-*: in adult German, *-e-* is used in some noun-noun compounds, but is infrequent and unproductive. However, some children overgeneralize it, forming compounds such as **Luft+e+bon-e* (instead of *Luft+bon-e*, ‘air balloons’). Strong BADs, by contrast, get the function wrong altogether. For example, iterativity and ongoingness are expressed by imperfectivity in Russian. Some children, however, make use of total reduplication as in *njam-njam* (‘I’m eating’) and *pik-pik* (repeated jumping). Although uncommon and not affecting every language-learning child, the importance of BADs should be recognized. Rare patterns still require explanation, and they set up boundary conditions on any theory of language and language acquisition. And we suspect that if researchers were to look for BADs with a very fine-toothed comb like D&al, more instances would turn up.

D&al note that BADs pose challenges to both usage-based and universalist theories of language acquisition. We share their skepticism toward usage-based accounts. While children do learn language from the input, the extent to which children’s grammar mirrors the input (Tomasello 2000) has been overstated by usage-based theorists (see Yang 2013, Yang 2016, Section 2.1 for discussion). At the same time, an initial stage of lexically specific grammar is hardly a unique claim of usage-based theories (e.g. Diessel 2013) but has always been acknowledged by supposedly opposing generative accounts (e.g. Marcus *et al.* 1992). A further difficulty, though not one unique to usage-based theories, is the absence of a credible learning theory that explains the child’s transition from a lexically restrictive stage to an adult-like productive and abstract grammar (Yang 2015, 2017). Indeed, recent usage-based accounts of morphological acquisition (Engelmann *et al.* 2019) overproduce patterns with highest token frequency, thus mirroring similar failings of previous models (Rumelhart & McClelland 1986). As is well known, however, minority patterns can be productive and are robustly overused by young children during language acquisition (Clahsen 1999, Elsen 2002, Yang 2016).

We are less certain about D&al’s critique of universalist theories. It may be a terminological matter. There are two senses in which the term ‘universal’ can be understood, and they are both part of the theory known as Universal Grammar (Chomsky 1965: Ch. 1, §5). In the domain of morphology, formal universals would characterize general properties of all morphological systems: words have constitutive parts (morphemes), morphemes are composed into structured units, the process of composition may be productive or unproductive, etc. By

contrast, substantive universals would describe the totality of morphological variations: particular types of morphological units (e.g. roots, infixes, templates, reduplication), the principles by which they combine (e.g. syntax-like locality conditions in Distributed Morphology, Halle & Marantz 1993, the structure and constraints on paradigms, Stump 2001), and their forms and functions (e.g. inflection vs derivation, the range of semantic relations encoded by morphology). We do not believe that anyone has seriously espoused a substantive theory of universal morphology – or that it is wise to do so. This is probably because Bromberger & Halle’s (1989) observation about phonology holds equally for morphology: both are products of history and other social/cultural contingencies, and it is difficult to imagine a theory that includes all conceivable morphological patterns in the world – and BADs produced by children. Perhaps D&al are objecting to even the possibility of such a theory; if so, then we agree. At the same time, it seems that almost everyone subscribes to a formal theory of universal morphology, e.g. words have internal structures. Admittedly, such a theory is not particularly helpful when it comes to the language-specific particulars: we need another theory – for instance, a learning theory – one which turns the formal into the substance across languages.

D&al’s own account of BADs is couched in the Natural Morphology framework (Dressler *et al.* 1987). The child constructs their morphology making use of a broad range of linguistic and non-linguistic strategies, especially a set of cognitively based universal preferences such as iconicity, transparency, etc. We are sympathetic to this view. Learning morphology clearly requires the intersecting forces of language and cognition: think the role of animacy, shape, size, etc. in noun class(ifier) systems, some of which are highly culturally specific and must be built upon the child’s conceptual and social development. But a constructivist needs to specify the precise mechanism by which the morphological system is built and subsequently refined: How do tendencies and possibilities become actualities? In particular, what exactly leads children astray into BADs – and why not all children? – and how do they manage to escape?

It seems to us, then, that all three theoretical approaches have a central missing piece: a learning-theoretic account of how children learn morphology. We now turn to such a learning-theoretic account, which in our view provides a more complete interpretation of BADs.

2. The Tolerance Principle

Learning a language requires discovering rules that generalize beyond a finite sample of data. The Tolerance Principle (**TP**) is a theory of how such generalizations are formed. Specifically,

(1) Let a rule R be defined over a set of N items. R is productive if and only if e , the number of items not supporting R , does not exceed θ_N :

$$e \leq \theta_N = \frac{N}{\ln N}$$

If e exceeds θ_N , then the learner will ‘lexicalize’ only these N items and not generalize beyond them: that is, R is unproductive. Yang (2016: 177) introduces a corollary, the Sufficiency Principle, which specifies how generalizations are formed when the ‘exceptions’ are not attested but cannot be regarded as impossible, on the grounds that evidence of absence is not absence of evidence. But it is clear that both principles generalize based on the number of positive examples: specifically, $N - \theta_N$, a super majority of N . Throughout this paper, we will refer to these variants collectively as the Tolerance Principle (TP).

The TP builds on the broadly shared intuition that a rule must earn productivity by the virtue of being applicable to a sufficiently large number of the candidates it is eligible for. For example, if there are 10 examples and all but one (9/10) support a rule, generalization ought to take place.

But no one in their right mind would extend a rule on the basis of 2/10: the learner should just memorize the two supporting examples. Productivity is a calibration of regularities and exceptions – crucially with respect to word types rather than tokens. For background and motivation for the TP, we refer the reader to Yang (2005, 2016) and much subsequent work. Recently, the TP has been implemented as the central component of a computational model of rule learning that proposes, and evaluates, rules over a corpus of data (Belth *et al.* 2021), in a step toward a discovery procedure envisioned at the beginning of generative grammar (Chomsky 1957).

By hypothesis, the productivity of a rule is determined by two integer values (N and e) defined over a vocabulary: these are obviously matters of individual vocabulary variation. Thus the TP allows room for variation in the transient stages of language acquisition as well as in the stable grammars of individual speakers. The relationship between N and e , which may change during the course of language acquisition, determines the productivity status of the rule. If e is very low as a proportion of N , then children may rapidly conclude that a rule is productive. Otherwise, a protracted stage of conservatism may ensue, which may be followed by the sudden onset of productivity, as can be famously seen in the over-regularization of irregular verbs in English (Marcus *et al.* 1992, Yang 2002). It is also possible that no rule ever reaches the productivity threshold; gaps and other phenomena of ineffability arise (Yang 2016: Ch. 5; Gorman & Yang 2019). Finally, the stochastic nature of child vocabulary acquisition may result in the productivity of a rule that is unproductive in the adult language, which ultimately corrects itself when a larger and more representative vocabulary is established.

Before seeing how the TP can help account for children's BADs, we correct a misunderstanding on D&al's behalf. They claim that "if a BAD, without support in the input, should have miraculously emerged, it should be immediately knocked out by these input-based principles, which is not the case in our example" (p. 134). This is not so. The calibration of productivity under the TP takes place in the child's internal vocabulary, via the values of N and e , rather than those in the external input (e.g. a corpus or a dictionary at the linguist's disposal). Plainly, the child does not, and can not, learn everything they hear in a corpus: their vocabulary actually grows quite slowly. An English learning child, for example, will have a vocabulary of at most 500 words at age two and at most 1,000 words at age three (Fenson *et al.* 1994, Hart & Risley 1995). Cross-linguistically, children in even the most favorable language-learning environments only manage a vocabulary of about 1,000 words by three (Hart & Risley 1995, Szagun *et al.* 2006, Bornstein *et al.* 2004), so the threat of instant knockout that D&al are concerned with will not materialize. However, BADs may eventually be knocked out once the child's vocabulary reaches a certain size such that the transitory productivity of BADs fails: the TP would thus provide an explanation for their demise.

3. BADS under the Tolerance Principle

To provide a direct TP analysis of D&al's own BAD examples would be difficult: As noted above, doing so requires very accurate vocabulary measures of the children under study. Such individual-level analyses under the TP can and have been done. For instance, in artificial language studies with children (Schuler 2017, Emond & Shi 2021), the number of (nonce) words and their morphological forms can be precisely manipulated. The outcome of learning – whether a rule generalizes or not – can be tested with additional nonce words not used in training, akin to a Wug test, and then compared against the predictions of the TP based on individual learner's vocabulary counts (N and e); the findings of both Schuler (2017) and Emond & Shi (2021) confirm the TP predictions. Unfortunately, the vocabulary measures of children in D&al's study are not available: even full transcripts can only capture a subset of the children's vocabulary. For present purposes, then, we will illustrate the TP approach to BADs with examples from cases for which

we do have accessible vocabulary data (or at least their proxies). We nevertheless hope to convey the general method that can be extended to case studies should precise vocabulary measures be available.

3.1. Weak BADS

Weak BADs have always been in the center of attention in language acquisition research albeit under different terminologies such as ‘analogy’. We agree with D&al: despite frequent allusions to analogical forces in language acquisition (e.g. Bybee 1985, Pinker & Prince 1988), analogical errors are really BADs in the sense that they are very uncommon. For instance, while over-regularization errors (e.g. *go* ~ *goed*, *hold* ~ *holded*) in child English are quite common, analogical errors such as *bite* ~ *bote*, *wipe* ~ *wope*, and *think* ~ *thunk* are vanishingly rare. The most comprehensive empirical study of analogical errors (Xu & Pinker 1995), in fact, dub these ‘weird past tense errors’ on the basis of their rarity. Xu and Pinker examined over 20,000 past-tense tokens produced by nine children: only forty weird errors (0.02%) were identified, and some of these may well be the result of phonological lenition such as *slep*, which may be the correct form *slept* undergoing the familiar process of word-final *t/d* deletion. Indeed, a sharp contrast between productive and unproductive processes can be observed in numerous acquisition studies across many languages (Lignos & Yang 2016), which partly motivated the categorical conception of productivity under the TP.

Xu & Pinker (1995)’s study does reveal a single systematic error of over-irregularization along the line of *sing* ~ *sang* / *sung* in child English, which would correspond to a weak form of BADs. The verb *bring* is over-regularized 32 times in the 4 million word corpus of child English extracted from the CHILDES project (MacWhinney 2000); over-application of productive forms is expected. But *brang* appears 6 times and *brung* 5 times in the same dataset. It is difficult to know if all of these are definitively examples of BADs: some dialects of American English do use *brang* and/or *brung* and a few instances can be found in the CHILDES input. Nevertheless, there is little doubt that some children do misgeneralize the pattern of *sing* and *ring* to *bring* to form *brang* and *brung*, as well as *swing* ~ *swang*.

The *bring* ~ *brang* error is a BAD in the weak sense: the child gets the form wrong but the function is correct (i.e. marking past tense). The TP provides a straightforward account for this, as discussed in many case studies of phonological, morphological, and syntactic acquisition (Yang 2016). Verbs such as *bring*, *sing*, and *ring* are very frequent: in fact, they are among the most frequent 200 verbs in child-directed English. As is well documented (e.g. Goodman *et al.* 2008), frequency plays an important role in vocabulary acquisition: more frequent words are more likely to be learned early.

Consider the hypothetical case where vocabulary acquisition proceeds in a manner determined by frequency: children always learn a more frequent verb before a less frequent one. By the time that the child learns 200 verbs, they will have learned 76 irregular verbs (based on the input frequencies from CHILDES): English irregular verbs are notoriously frequent. Note that 76 irregulars are far too many to support the productivity of *-ed* rule despite its coverage of a majority of 127 verbs. According to the TP, a rule can generalize over a vocabulary of 200 items only if the number of exceptions falls below $\frac{200}{\ln 200} \approx 37$. However, the learner, as in the computational model of Belth *et al.* (2021), may find a productive rule defined over a subset of the verbs when the entire vocabulary fails to yield a productive rule. Specifically, the rule $\text{ɪ} \rightarrow \text{æ} / __ \eta$ reaches productivity because it works for two (*sing* and *ring*), and the one exception of *bring* can be tolerated. In fact, the rule can enjoy an extended period of productivity. Among the top 500 verbs, the two exceptions (*bring* and *swing*) cause no difficulty for the remaining three. A case of BAD.

But the productivity of $i \rightarrow \text{æ} / _ \eta$ won't last too long. Consider the relevant words when the vocabulary reaches 800:

- (2) bring, *fling*, **ring**, **sing**, **spring**, *sting*, *swing*, wing

There are now 8 verbs that end in /ɪŋ/ but their past tense forms are scattered all over the place: the three in boldface change vowel to /æ/, the three in italics change the vowel to /ʌ/, one is completely idiosyncratic (*bring*), and the last one is regular (*wing*). Everyone loses because no one is numerically dominant enough to tolerate the rest. At this point, children will stop the generalization and memorize the past tense of irregular /ɪŋ/ verbs by rote. The absence of productivity for this class of verbs accounts for the overall rarity of 'weird past tense errors': D&al's BADs. Importantly, adults also treat this class as unproductive and do not generalize the irregular patterns. The past few decades have seen verbs such as *bing* (Microsoft search engine) and *bling* enter into the English vocabulary: they are uniformly inflected with the regular *-ed*. Of course no child learns words on a strictly frequency-determinant fashion. The actual vocabulary is surely a stochastic sample of input words: the fortuitous condition for the rise of *bring* ~ *brang* may not be met for every learner, for it requires learning only the three most frequent /ɪŋ/ verbs before learning enough regular verbs such that *-ed* becomes productive. But there is nothing unprincipled about our account: every child has a vocabulary, and the only limiting factor is our ability to measure it precisely. Once we have *N*'s and *e*'s, TP predictions are completely mechanical and unambiguous. BADs therefore receive an explanation from start to finish: it is simply the rise and fall of productivity stemmed from the child's changing vocabulary. The same approach holds for language change: see Ringe & Yang (2022) for how the TP accounts for the historical productivity of the /ɪŋ/ class and how it was lost over time.

3.2. Strong BADs

A few qualifying words on strong BADs first. In our understanding, the strong cases of BADs in D&al's study are not children's inventions out of thin air but are in fact grounded somewhere in the input data. Consider the case of reduplication in child Russian and Greek. While reduplication is not used to express, respectively, iterativity/ongoingness and the subjunctive in these languages, reduplication, as a formal pattern, does appear in both languages. It is used in Russian to express, among others, intensification (Israeli 1997), and is used in Modern Greek to express a complex range of semantic and pragmatic information (Kallergi 2015). Children's BADs are clearly different from these, but they are appropriately described as misalignment between an extant morphological form and its function.

Viewed under these lights, then, strong BADs may not be as rare or atypical as D&al suggest. Quite a lot of research in language acquisition, including morphological acquisition, has identified such misalignments of form and function as an especially revealing type of errors. One prominent example is the extensively studied phenomenon of Root Infinitives (**RI**) whereby children frequently use non-finite verbal forms in matrix sentences that require a finite form (Rizzi 1993, Hoekstra & Hyams 1998, Wexler 1998, Legate & Yang 2007). In French, for instance, non-finite matrix clauses are restricted to particular discourse situations that have a time reference as shown in the following dialogue (3) and questions (4), drawn from Rasetti (2003):

- (3) a. *Que veux-tu faire?*
what want-you do.INF?
'What do you want to do?'
- b. *Partir*
leave.INF
'(I want) to leave.'

- (4) a. *Comment lui dire cela?*
 how him.DAT say.INF this?
 'How to say this to him?'
- b. *Moi faire ceci? Jamais!*
 me do.INF that never
 'Me do that? Never!'

French-learning children, however, use non-finite matrix verbs in broader contexts and not infrequently – and crucially, in declarative sentences, over 10% of the time for the three children in Rasetti's study. Children's use of RIs falls into two broad classes in French: it may have a modal interpretation (5a) or a descriptive one (5b).

- (5) a. *manger # maman* [=reaching for chocolate]
 eat.INF mommy
- b. *caler la crayon # voilà!*
 hide.INF the pencil there

The RI phenomena, then, appears to be a strong case of BAD: the child erroneously maps an input form (non-finite verbs) to a non-target function (matrix usage in stand-alone declarative sentences). But unlike the examples of BADs in D&al's paper, which struck only select children, RI has been observed in all children learning tense-marking languages with quantitative cross-linguistic differences regarding the extent and duration of non-target form use.

RIs, too, can be accounted for under the TP. Before learning the correct mappings from form to function, children must first learn which functions are realized – in this case, which morphological features are marked – in their language. Chinese, for example, does not mark tense, while English does. English, on the other hand, does not contrast first and second person, while Spanish and French do. Which subsets of morphological features are marked, and are thus eligible to be mapped to form, vary across languages and must be learned from the input. Some proposals regarding the RI phenomenon have suggested that it results from under-specified features in the Tense node (e.g. Hoekstra & Hyams 1998). Our proposal, which draws on Payne (2022), is in a similar spirit: we suggest that the RI emerges before the child has seen sufficient evidence for tense marking; at this stage, tense is essentially under-specified, and children may supply verbal forms inappropriate for the function, i.e. non-finite forms for finite context.

The TP can be straightforwardly applied to the problem of learning which features are marked in a language. Much work on morphological acquisition (e.g. Marquis & Shi 2012, Kim & Sundara 2021) has demonstrated that children's ability to segment inflectional morphology and relate inflected forms to their stems emerges early in acquisition – as early as 8 months in English. At the same time, the Principle of Contrast (Clark 1987) provides a (perhaps innate) hypothesis that different forms will correspond to different meanings, or, in the language of the present discussion, different functions. We suggest that the child may combine their early ability to segment and relate inflected forms with the Principle of Contrast in order to detect and make use of 'collisions' – a single stem appearing in multiple inflected forms (e.g. *walk* ~ *walked*). Such collisions may then be used to determine which features are marked – and thus which functions may be mapped to form – in the language (e.g. *walk* ~ *walked* provides evidence for \pm PAST).

Yet a single collision is likely not sufficient evidence that a language realizes the corresponding features: English, for example, distinguishes first and second person in *I am* ~ *you are* but this distinction does not generalize (cf. *I walk* ~ *you walk*). Here, we may straightforwardly apply the TP: once the child learns which function is realized by individual collisions (e.g. \pm PAST above), they may employ the TP to determine whether there is sufficient evidence for the realization of that function across the language. More formally, if we encounter

a collision between inflected form A and inflected form B, where A is less frequent, then we determine the morphological features that are realized by this collision. We then search the lexicon for other words that have appeared in inflected form A. If enough of these also have a collision that realizes the same features (where enough is defined by the TP), then we learn that these features are marked in the language. We focus on the less frequent of the two inflected forms because the number of collisions between A and B will be higher relative to the total count of A than the total count of B, giving it the best chance of passing the TP. We refer the reader to Payne (2022) for the details and computational implementation of this approach.

Learning under this proposal is recursive: the child may learn, for example, that \pm PARTICIPLE is marked in English, and then that \pm 3SG is marked for $-$ PARTICIPLE forms, and so on. The order in which feature marking is learned, which is dependent on the vocabulary the child has acquired incrementally, can be translated into a feature hierarchy representation familiar in morphological theories. For languages where agreement emerges before tense, this necessarily means that tense will emerge separately for each agreement. While such an approach may at first seem counter-intuitive, it may actually provide a direct explanation for cross-linguistic differences in RI phenomena. Several accounts of RI (e.g. Phillips 1995, Guasti 2004) have suggested that the length and frequency of the RI stage is inversely correlated with the morphological ‘richness’ of the language: languages with richer inflectional morphology yield shorter RI stages. While the idea of ‘richness’ itself is not a precise or sufficient explanation, Legate & Yang (2007) build on these intuitions by suggesting that it is not the ‘richness’ of the paradigm that matters, but rather the relative amount of forms the child encounters that reward a tense-marking grammar. In the case of Spanish, English, and French, they show that the child will encounter tense-rewarding forms most in Spanish and least in English, which matches with the relative length of the RI period (longest in English and shortest in Spanish).

Our proposal offers an alternate explanation: In English, French, and Spanish, agreement generally emerges before tense (Brown 1973, Berko 1958, Prévost 2009, Montrul 2004), which means that our model will subdivide the input based on agreement and learn tense marking within each agreement node separately. At the same time, the TP tolerates relatively more exceptions for smaller values of N : when $N = 10$, $\theta_N = 4$, or 40% of N . By contrast, when $N = 100$, $\theta_N = 21$, or about 21% of N . Thus, it is easier to learn tense marking over smaller N , or equivalently, tense will emerge more quickly in languages that subdivide more based on agreement. This is indeed the case: English marks only one agreement (\pm 3SG), while Spanish marks 5 or 6, depending on the dialect. French falls between, with the first and second plural being the only agreements that are realized phonologically (Legate & Yang 2007). Our proposal straightforwardly predicts that tense-marking will emerge most quickly in Spanish, followed by French and then English, based solely on properties of subdivision under the TP. These predictions match with developmental findings regarding both age of acquisition and length of the RI stage.

It is important to note that learning what functions are realized – or equivalently, what features are marked – in a language does not mean that the child has yet learned how to map these functions to form. At the stage when the child only knows agreement marking, it is unlikely that any morphological process will be productive since the processes differ based on tense. At this point, for example, the Spanish-learning child would group *amas*, *amabas*, *amaras*, *amaste*, etc. because all forms mark the second singular, but no single morphological process can derive all four forms. This explains why the child may produce the RI, as well as its gradient nature: if there is not yet a productive rule to inflect a given agreement form, the child must rely entirely on lexicalized inflected forms with the desired agreement. In the absence of such forms, and the absence of a productive rule for inflection, it is reasonable to expect that the child may resort to an un-tensed form as long as it marks agreement correctly. Once tense and other markings do emerge, the child will then be able to map these functions to form, and weak BADS may emerge.

The TP can thus provide a direct account for the emergence, gradience, and subsequent demise of the RI, as well as cross-linguistic differences in this case of a strong BAD.

4. Does everything go?

We offer the TP as a formal universal of language (including morphology). It is possible that the TP is rooted in the non-linguistic component of cognition: a general principle of learning and generalization, although empirical work in other domains will be necessary.

But where does the substance of morphology come from? Presumably this is D&al's Feyerabendian question 'does everything go?' Natural Morphology likely constitutes part of the answer. The fact that children, unprompted, use reduplication to express iterativity seems to be grounded in cognitive factors such as iconicity: repetition is repetition. But we doubt a complete theory of substance is ever forthcoming. The usual suspects – sounds and meanings – no doubt participate in the construction of morphology, but D&al make it clear that all sorts of individualistic factors come into play as well. Children learn different words, they latch onto different aspects of the world around them, and every little perturbation, such as a morpheme segmentation error (Peters 1983), can give rise to BADs.

Whatever rules children come up with, however, are to be subjected to a TP validation on the basis of two cardinality values. Some of these will survive in the long run so we have a relatively stable grammar across individuals; BADs are those missteps that do not last.

Bibliographical References

- Belth, Caleb; Payne, Sarah; Beser, Deniz; Kodner, Jordan & Yang, Charles 2021. The greedy and recursive search for morphological productivity. *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society* 42. 2869-2875.
- Berko, Jean 1958. The child's learning of English morphology. *Word* 14. 150-177.
- Bornstein, Marc H.; Cote, Linda R.; Maital, Sharone; Painter, Kathleen; Park, Sung-Yun; Pascual, Liliana; Pêcheux, Marie-Germaine; Ruel, Josette; Venuti, Paola & Vyt, Andre 2004. Cross-linguistics analysis of vocabulary in young children: Spanish, Dutch, French, Hebrew, Italian, Korean, and American English. *Child development* 75. 1115-1139.
- Bromberger, Sylvain & Halle, Morris 1989. Why phonology is different. *Linguistic Inquiry* 20. 51-70.
- Brown, Roger 1973. *A First Language: The Early Stages*. Cambridge, MA: Harvard University Press.
- Bybee, Joan 1985. *Morphology: A Study of the Relation Between Meaning and Form*. Philadelphia, PA: Benjamins.
- Chomsky, Noam 1957. *Syntactic Structures*. The Hague: De Gruyter Mouton.
- Chomsky, Noam 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Clahsen, Harald 1999. Lexical entries and rules of language: A multidisciplinary study of German inflection. *Behavioral and Brain Sciences* 22. 991-1069.
- Clark, Eve 1987. The principle of contrast: A constraint on language acquisition. In MacWhinney, Brian (ed.), *Mechanisms of Language Acquisition*. Hillsdale, NJ: Erlbaum. 1-33.
- Diessel, Holger 2013. Construction grammar and first language acquisition. In Hoffman, Thomas & Trousdale, Graeme (eds.), *The Oxford Handbook of Construction Grammar*. Oxford: Oxford University Press. 347-364.
- Dressler, Wolfgang U.; Christofidou, Anastasia; Gagarina, Natalia; Korecky-Kröll, Katharina & Kilani-Schoch, Marianne 2019. Morphological blind-alley developments as a theoretical challenge to both usage based and nativist acquisition models. *Italian Journal of Linguistics* 31. 107-140.
- Elsen, Hilke 2002. The acquisition of German plurals. In *Morphology 2000: Selected Papers from the 9th Morphology Meeting, Vienna, 25-27 February 2000*. Amsterdam: Benjamins. 117.

- Emond, Emeryse & Shi, Rushen 2021. Infants' rule generalization is governed by the Tolerance Principle. *45th Annual Boston University Conference on Language Development*. 191-204.
- Engelmann, Felix; Granlund, Sonia; Kolak, Joanna; Szreder, Marta; Ambridge, Ben; Pine, Julian; Theakston, Anna & Lieven, Elena 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.
- Fenson, Larry; Dale, Philip S.; Reznick, J. Steven; Bates, Elizabeth; Thal, Donna J.; Pethick, Stephen J.; Tomasello, Michael; Mervis, Carolyn B. & Stiles, Joan 1994. Variability in early communicative development. *Monographs of the Society for Research in Child Development*. i-185.
- Goodman, Judith C.; Dale, Philip S. & Li, Ping 2008. Does frequency count? Parental input and the acquisition of vocabulary. *Journal of Child Language* 35. 515-531.
- Gorman, Kyle & Yang, Charles 2019. When nobody wins. In Rainer, Franz; Gardani, Francesco; Luschützky, Hans Christian & Dressler, Wolfgang U. (eds.), *Competition in Inflection and Word Formation*. Berlin: Springer. 169-193.
- Guasti, Maria Teresa 2004. *Language Acquisition: The Growth of Grammar*. Cambridge, MA: MIT Press.
- Halle, Morris & Marantz, Alec 1993. Distributed morphology and the pieces of inflection. In Hale, Kenneth & Keyser, Jay S. (eds.), *The view from Building 20: Essays in Linguistics in Honor of Sylvain Bromberger*. Cambridge, MA: MIT Press. 111-176.
- Hart, Betty & Risley, Todd R. 1995. *Meaningful Differences in the Everyday Experience of Young American Children*. Baltimore, MD: Paul H. Brookes Publishing.
- Hoekstra, Teun & Hyams, Nina 1998. Aspects of root infinitives. *Lingua* 106. 81-112.
- 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.
- Kim, Yun Jung & Sundara, Megha 2021. 6-month-olds are sensitive to English morphology. *Developmental Science* 24. e13089.
- Legate, Julie Anne & Yang, Charles 2007. Morphosyntactic learning and the development of tense. *Language Acquisition* 14. 315-344.
- Lignos, Constantine & Yang, Charles 2016. Morphology and language acquisition. In Hippisley, Andrew R. & Stump, Gregory (eds.), *The Cambridge Handbook of Morphology*. Cambridge: Cambridge University Press. 765-791.
- MacWhinney, Brian 2000. *The CHILDES project: Tools for analyzing talk*. 3rd edition. Mahwah, NJ: Lawrence Erlbaum Publishing.
- Marcus, Gary F.; Pinker, Steven; Ullman, Michael T.; Hollander, Michelle; Rosen, John T. & Xu, Fei 1992. Over-regularization in language acquisition. *Monographs of the Society for Research in Child Development* 57. Chicago: University of Chicago Press. 1-182.
- Marquis, Alexandra & Shi, Rushen 2012. Initial morphological learning in preverbal infants. *Cognition* 122. 61-66.
- Montrul, Silvina 2004. *The Acquisition of Spanish: Morphosyntactic Development in Monolingual and Bilingual L1 Acquisition and Adult L2 Acquisition*. Amsterdam: Benjamins.
- Payne, Sarah 2022. *When collisions are a good thing: the acquisition of morphological marking*. Bachelor's thesis, University of Pennsylvania, Philadelphia, PA.
- Peters, Ann M. 1983. *The Units of Language Acquisition*. Cambridge: Cambridge University Press.
- Phillips, Colin 1995. Syntax at age two: Cross-linguistic differences. *MIT Working Papers in Linguistics* 26. 325-382.
- Pinker, Steven & Prince, Alan 1988. On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition* 28. 73-193.
- Prévost, Philippe 2009. *The acquisition of French: The Development of Inflectional Morphology and Syntax in L1 Acquisition, Bilingualism, and L2 Acquisition*. Amsterdam: Benjamins.
- Rasetti, Lucienne 2003. *Optional categories in early French syntax: A developmental study of root infinitives and null arguments*. PhD dissertation, University of Geneva, Geneva, Switzerland.
- Ringe, Don & Yang, Charles 2022. The threshold of productivity and the irregularization of verbs in Early Modern English. In Los, Bettelou; Cowie, Claire; Honeybone, Patrick & Trousdale, Graeme (eds.), *English historical linguistics: Change in structure and meaning*. Amsterdam: Benjamins. 91-111.

- Rizzi, Luigi 1993. Some notes on linguistic theory and language development: The case of root infinitives. *Language acquisition* 3. 371-393.
- Rumelhart, David E. & McClelland, James L. 1986. On learning the past tenses of English verbs. In McClelland, James L.; Rumelhart, David E. & the A.T.P.R. Group (eds.), *Parallel distributed processing: Explorations into the microstructure of cognition*. Vol 2. Cambridge, MA: MIT Press. 216-271.
- Schuler, Kathryn 2017. *The acquisition of productive rules in child and adult language learners*. PhD dissertation, Georgetown University, Washington, D.C.
- Stump, Gregory 2001. *Inflectional Morphology: A Theory of Paradigm Structure*. Cambridge: Cambridge University Press.
- Szagun, Gisela; Steinbrink, Claudia; Franik, Melanie & Stumper, Barbara 2006. Development of vocabulary and grammar in young German-speaking children assessed with a German language development inventory. *First Language* 26. 259-280.
- Tomasello, Michael 2000. Do young children have adult syntactic competence? *Cognition* 74. 209-253.
- Wexler, Kenneth 1998. Very early parameter setting and the unique checking constraint: A new explanation of the optional infinitive stage. *Lingua* 106. 23-79.
- Xu, Fei & Pinker, Steven 1995. Weird past tense forms. *Journal of Child Language* 22. 531-556.
- Yang, Charles 2002. *Knowledge and Learning in Natural Language*. Oxford: Oxford University Press.
- Yang, Charles 2005. On productivity. *Linguistic Variation Yearbook* 5. 333-370.
- Yang, Charles 2013. Ontogeny and phylogeny of language. *Proceedings of the National Academy of Sciences* 110. 6324-6327.
- Yang, Charles 2015. Negative knowledge from positive evidence. *Language* 91. 938-953.
- Yang, Charles 2016. *The Price of Linguistic Productivity: How Children Learn to Break Rules of Language*. Cambridge, MA: MIT Press.
- Yang, Charles 2017. Rage against the machine: Evaluation metrics in the 21st century. *Language Acquisition* 24. 100-125.