Dependency length minimization and its limits: a possible role for a probabilistic version of the
Final-Over-Final Condition (to appear in *Language*)

Yingqi Jing^{1, 4, 5}, Damián E. Blasi^{2, 6, 7}, and Balthasar Bickel³

¹University of Zurich; yingqi.jing@uzh.ch

²Harvard University; dblasi@fas.harvard.edu

³University of Zurich; balthasar.bickel@uzh.ch

⁴Uppsala University

⁵University of Turku

⁶Max Planck Institute for the Science of Human History

⁷National Research University Higher School of Economics

A prominent principle in explaining a range of word order regularities is dependency locality, which minimizes the linear distances (dependency lengths) between the head and its dependents. However, it remains unclear to what extent language users in fact observe locality when producing sentences under diverse conditions of cross-categorical harmony (such as the placement of verbal and nominal heads on the same vs different sides of their dependents), dependency direction (head-final vs head-initial) and parallel vs. hierarchical dependency structures (e.g. multiple adjectives dependent on the same head vs nested genitive dependents). Using 45 dependency-annotated corpora of diverse languages, we find that after controlling for harmony and conditioning on dependency types, dependency length minimization (DLM) is inversely correlated with the overall presence of head-final dependencies. This anti-DLM effect in sentences with more head-final dependencies is specifically associated with an accumulation of dependents in parallel structures and with disharmonic orders in hierarchical structures. We propose a detailed interpretation of these results and tentatively suggest a role for a probabilistic principle that favors embedding head-initial (e.g. VO) structures inside equally head-initial and thereby length-minimizing structures (e.g. relative clauses after the head noun) while head-final (OV) structures have a less pronounced preference for harmony and DLM. This is in line with earlier findings in research on the Greenbergian word order universals and with a probabilistic version of what has been suggested as the Final-Over-Final Condition more recently.*

Keywords: dependency length minimization, word order harmony, head-finality, Final-Over-Final Condition, dependency treebanks

*This research was supported in part by the Ph.D grant from the China Scholarship Council (201606320224) to Yingqi Jing. Damián E. Blasi acknowledges funding from the Branco Weiss Fellowship, administered by the ETH Zurich. Damián E. Blasi's research was also executed within the framework of the HSE University Basic Research Program and funded by the Russian Academic Excellence Project '5-100'. We are grateful to the many colleagues who helped us, in particular Lena A. Jäger, Paola Merlo, Paul Widmer, Carlos Gómez-Rodríguez, Kellen Parker van Dam, Yoko Yamazaki, Arezoo Zandy, and to the associate editors (Shravan Vasishth and Titus von der Malsburg) and the reviewers for helpful comments. Author contributions: JY, BD, BB designed the research; JY and BB undertook the theoretical framing; JY and BD produced the random baselines and performed the statistical analysis; and all authors discussed the results and contributed to the final version of the paper.

1. INTRODUCTION. A growing body of cross-linguistic evidence suggests that human languages prefer shorter linear distances between a syntactic head and its dependents. This is illustrated by the data from English in 1. The addition of the prepositional phrase *from the clinic* in 1b induces a longer distance between the subject *the nurse* and the verb *supervised* (Grodner & Gibson 2005):

Experimental evidence shows that sentences with longer dependencies increase various processing demands, such as parsing time, working memory load, retrieval failures, and phrase type recognition (Gibson 1998, 2000, Grodner, Gibson, & Watson 2005, Lewis & Vasishth 2005). This suggests that longer dependencies are disfavored in use (Gibson 1998, Hawkins 2004, 2014, Liu 2008, Temperley 2008, Futrell, Mahowald, & Gibson 2015). In line with this, early corpus-based studies showed a preference for minimal phrasal domains (Hawkins 1983), and more recent studies revealed systematically shorter average dependency lengths than what one would expect under random baselines across many languages (Ferrer i Cancho 2004, Liu 2008, Gildea & Temperley 2010, Futrell, Mahowald, & Gibson 2015, Gulordava 2018, Futrell, Levy, & Gibson 2020). All these findings converge to a single principle that has come to be known as “dependency locality” or “dependency length minimization (DLM)”. In what follows, we use “dependency locality” to refer to statements of principle or theory and DLM for its measurable effect in corpus data.

Dependency locality has been posited as a probabilistic principle of language (Ferrer i Cancho 2004, Liu 2008, Gildea & Temperley 2010, Futrell, Mahowald, & Gibson 2015). The principle has been argued to predict various universal tendencies in grammars (see Futrell, Levy, & Gibson 2020, for a recent review). A particularly prominent prediction is that dependency locality (or some variant of it) explains universal trends towards WORD ORDER HARMONY (Temperley 2008, Hawkins 1983, 2014, Futrell, Mahowald, & Gibson 2015, Temperley & Gildea 2018). Word order harmony describes a statistical tendency of languages to prefer ordering heads and their dependents in a consistent way (Greenberg 1963, Lehmann 1973, Vennemann 1974, Dryer 1992). For example, languages tend to order the heads of VPs (or verb-headed dependencies) on the same side as the heads of NPs (noun-headed dependencies). The association between dependency locality and word order harmony is illustrated by 2.

(2) DLM and Harmony

In this example, minimizing the distance between the topmost head *consider* and its immediate dependent *the impact* leads to a harmonic, consistently right-branching structure (2a). The alternative disharmonic structure in 2b exhibits a longer distance because of the intervening sub-dependent *the paper's*.

However, as noted in previous work (Temperley 2008, Temperley & Gildea 2018, Futrell, Mahowald, & Gibson 2015, Yadav et al. 2020, Futrell, Levy, & Gibson 2020), the idea that harmony is associated with – or even driven by – dependency locality is challenged by two observations, one on dependency structures with parallel dependents, and one on asymmetries between head-initial and head-final structures.

The first challenge comes from the placement of parallel dependents. When there are multiple parallel dependents of the same head, dependency locality can lead to disharmonic, rather than harmonic structures (Temperley 2008, Temperley & Gildea 2018). For example, when multiple adjective phrases (APs) depend on the same head, dependency locality favors distributing them on opposite sides, as in 3a; harmonically piling APs up on the same side, as in 3b, increases one of the distances to the head:

(3) DLM in parallel dependents

Thus, dependency locality favors putting short and long dependents on either side of the head, disrespecting harmony (Dryer 1992, Hawkins 2014). Consistent with this, randomized linearization experiments show that consistent head direction alone results in longer dependency lengths than what one observes, and that random baselines with fixed head position can also minimize lengths by placing the shorter parallel dependents closer to the head without harmonizing head positions across categories (Temperley 2008, Futrell, Mahowald, & Gibson 2015).

The second challenge comes from evidence that the association between dependency locality and harmony can be modulated by the direction of a dependency (i.e. whether the head is initial

or final). In head-initial structures, DLM arises independently of harmony because of what is known as BEHAGHEL’S LAW (“Gesetz der wachsenden Glieder”, the law of growing parts) (Behaghel 1909, MacDonald 2013), which puts shorter elements before longer elements (see example 4a). In head-final structures, however, the same principle could increase dependency length, but notably, languages with more head-final dependencies are found not to follow Behaghel’s Law and prefer long-before-short orders instead (see the Japanese example 4b), albeit with varying degrees (Yamashita & Chang 2001, Hawkins 2004, Choi 2007, Temperley 2008, Ros et al. 2015, Temperley & Gildea 2018). This observation has led to postulating an alternative principle of NESTED ORDERING, which places shorter dependents closer to the head on either side (Temperley 2008, Temperley & Gildea 2018). This principle predicts DLM without appealing to harmony.

(4) Behaghel’s Law and Nested Ordering

Another way in which dependency direction interferes with the association of locality and harmony stems from observations of word order asymmetries. For example, VO languages are known to far more strongly prefer harmonic relative clauses (hence, post-nominal: cf. 5a and 5b) than OV languages, which have no strong preference for harmonic relative clauses (5c and 5d) (Greenberg 1963, cf. the data in the World Atlas of Language Structures (WALS): Dryer & Gensler 2005, Dryer 2013a). As a result, languages that prefer VO orders are expected to show DLM far more often and more strongly than OV languages.¹

(5) Asymmetries of relative clauses in VO and OV languages

- a. VO & NRel (63.82% genera in WALS) with an English example:

- b. VO & RelN (0.78% genera in WALS) with a Chinese example:

¹Even though postnominal relative clauses can be sometimes extraposed to the other side of the V in the matrix clause, it will still lead to a disharmonic word order combination, not optimal for DLM (Hawkins 2014).

c. OV & NRel (16.15% genera in WALS) with a Persian example:

d. OV & RelN (19.25% genera in WALS) with a Japanese example:

In recent work, these word order asymmetries have been summarized in the FINAL-OVER-FINAL CONDITION (FOFC; Biberauer, Holmberg, & Roberts 2014, Biberauer et al. 2017a), which bans head-initial orders directly dominated by head-final orders (and so ‘final’ must be ‘over final’).² The FOFC has been mostly discussed as an absolute universal that would strictly prohibit head-initial phrases like [*the Queen* [*of England*]] embedded in a head-final phrase, as in [[*the Queen* [*of England*]]’s opinion] (or the relative clause 5b). While such unexpected structures are sometimes explained away with additional analytical machinery (Biberauer, Holmberg, & Roberts 2014), we recast the FOFC here as a probabilistic tendency (see Bickel 2015 for further discussion):

(6) Final-Over-Final-Condition

- a. Preferred: A $\overrightarrow{\quad}$ B $\overrightarrow{\quad}$ C
- b. Preferred: C $\overleftarrow{\quad}$ B $\overleftarrow{\quad}$ A
- c. Preferred: A $\overleftarrow{\quad}$ C $\overleftarrow{\quad}$ B

²Related notions are Williams’s Head-Final Filter or Osborne’s Left-Elbow-Constraint, which hypothesize that a prenominal modifier may not have a post-dependent (Williams 1982, Osborne 2003). See Biberauer et al. (2017b) and Sheehan (2017b) for discussion.

d. Dispreferred:

Both Behaghel’s Law and the probabilistic version of the FOFC predict that the association between locality and harmony can differ between head-initial and head-final structures and between parallel and hierarchical dependencies. Tentative evidence for the overall difference between head-initial and head-final conditions comes from cross-linguistic corpus studies, showing stronger DLM in head-initial than in head-final languages (Park & Levy 2009, Futrell, Mahowald, & Gibson 2015, Futrell, Levy, & Gibson 2020, Yadav et al. 2020). Consistent with this, experimental research on head-final structures shows that informationally rich and/or unambiguous dependents facilitate the processing of the upcoming head by increasing its predictability (Konieczny 2000, Vasishth & Lewis 2006, Levy 2008, Levy & Keller 2013, Husain, Vasishth, & Srinivasan 2014) and they also facilitate sentence planning by keeping plans distinct from alternatives (Sauppe et al. 2021). These facilitatory effects tend to increase with information content in the dependent, and this in turn leads to longer distances between it and the head, an effect known as ANTI-LOCALITY or “anti-DLM”. However, as far as we can tell, neither corpus nor experimental studies have systematically controlled for the difference between parallel and hierarchical dependencies (nor their interactions with head order).

Here we explore the presence and extent of DLM under various conditions of harmony and dependency direction, and across parallel and hierarchical dependencies. We develop and fit our models on 45 dependency-annotated corpora of diverse languages from the Universal Dependencies (UD) database, version 2.5 (Nivre et al. 2019). We base our model on chains of two adjacent dependencies, so as to fully capture the local combinations of both hierarchically connected and parallel dependents in a sentence (Osborne, Putnam, & Groß 2012), as discussed below in Section 2.2.

2. DATA AND METHODS.

2.1. DEPENDENCY ANNOTATED CORPORA. The UD corpora encode syntactic structure in terms of Dependency Grammar, where each pair of words is linked via a directed arc that indicates the relations of the words, including the head directions (Hudson 1984, Meřčuk 1988, Liu 2009). Figure 1 illustrates this for an English sentence. In our study we focus exclusively on major lexical categories (blue), such as V(erb), N(oun), Adj(ective), and Adv(erb), avoiding ambiguity in, say, whether an expression like [*the article*] is headed by the determiner or the noun, or whether adpositions are different from cases in all languages of the UD database (Futrell, Levy, & Gibson 2020). We also exclude from our analysis proper names, since they can be annotated differently across languages. For example, *United States* is coded as a proper name in English but as an

adjective-noun sequence in Czech.³ Furthermore, we exclude from our analysis multiword and coordinate structures, where dependencies are unclear, and sentences with crossing dependencies, which are rare in the UD data (Ferrer i Cancho, Gómez-Rodríguez, & Esteban 2018).

FIGURE 1. Dependency graph of an English sentence (from a political debate about a law): each terminal is annotated by its parts-of-speech tag (e.g. PRON for ‘pronoun’), and each arc for its dependency relation (e.g. ‘nsubj’ for ‘nominal subject’). The directions of the arcs show the head directions, for example, the adverbial modifier *certainly* forms a head-final dependency with the head *fit*, and the word *parts* forms a head-initial dependency with the dependent *article*. Blue coloring indicates the lexical dependencies that we exclusively focus on.

In the figure we can see that the root verb (*fit*) precedes the ‘obl’ (oblique) object (*into certain parts*), which includes a genitive construction (*of the article*) as a further ‘nmod’ (nominal) modifier. This represents a consistently head-initial hierarchical structure. By contrast, the object *into certain parts* contains an ‘amod’ (adjectival) modifier to the left of the head. This is not in harmony with the rest (although expected, given the results of Dryer 1992).

Assessing the relative effect of harmony and dependency direction on DLM is only possible in languages where there is enough variation in head order. In strict head-final or strict head-initial languages, all dependencies will trivially be harmonic. Because of this, we selected only languages from the UD database with a global head-final proportion between 0.2 and 0.8, excluding the languages Arabic, Irish, Scottish Gaelic, Galician, Thai, Hebrew, Turkish, Urdu, Hindi, Amharic, Kazakh, Uyghur, Telugu, Korean and Japanese. Within this interval, there is sufficient variation in both harmony and dependency direction so that they can be analyzed in a meaningful way (see Figure S1 in the Supplementary Information). We further exclude those small datasets with less than 800 sentences (Armenian, Assyrian, Belarusian, Breton, Buryat, Cantonese, Coptic, Karelian, Komi Zyrian, Kurmanji, Lithuanian, Marathi, Naija, Sanskrit, Swedish Sign Language, Tagalog, Tamil, Upper Sorbian, Warlpiri, Welsh, Yoruba). This results in the data summarized in Table 1.

³<https://universaldependencies.org/u/pos/PROPEN.html>

TABLE 1. Overview of 45 dependency treebanks used in the study. IE: Indo-European; AA: Austro-Asiatic; AN: Austronesian; NC: Niger-Congo. We simplify the geographical distribution as Europe vs. not, since most languages are from Europe. The column on lexical dependencies summarizes the proportion of the dependencies that we target for our measurements of dependency length among all dependencies in the corpus, the column on head-finality reports the proportion of head-final dependencies, and the column on harmony the proportion of harmonically connected two-dependency catenae. The proportion of lexical dependencies is largely independent of head-finality and harmony. See Figure S2 and statistical results in the Supporting Information Section S3.2.

Language	Family/Subfamily	Europe	Sentence	Lexical dependency	Head-finality	Harmony
Afrikaans	IE, Germanic	F	1496	0.29	0.6	0.54
Ancient Greek	IE, Greek	T	8714	0.31	0.44	0.56
Bambara	Mande, East Manding	F	827	0.21	0.45	0.48
Basque	Basque, Basque	T	5768	0.37	0.74	0.58
Bulgarian	IE, Slavic	T	9944	0.34	0.44	0.47
Catalan	IE, Romance	T	14657	0.27	0.24	0.62
Chinese	Sino-Tibetan, Sinitic	F	4893	0.3	0.74	0.54
Classical Chinese	Sino-Tibetan, Sinitic	F	8204	0.39	0.68	0.47
Croatian	IE, Slavic	T	7957	0.37	0.46	0.47
Czech	IE, Slavic	T	68317	0.39	0.5	0.47
Danish	IE, Germanic	T	3897	0.33	0.42	0.52
Dutch	IE, Germanic	T	10600	0.28	0.59	0.53
English	IE, Germanic	T	12406	0.28	0.45	0.53
Erzya	Uralic, Mordvin	T	1234	0.32	0.61	0.49
Estonian	Uralic, Finnic	T	23646	0.41	0.61	0.54
Faroese	IE, Germanic	T	1021	0.26	0.55	0.45
Finnish	Uralic, Finnic	T	14842	0.37	0.51	0.51
French	IE, Romance	T	14836	0.26	0.26	0.63
German	IE, Germanic	T	12874	0.26	0.61	0.55
Gothic	IE, Germanic	T	3516	0.38	0.43	0.56
Greek	IE, Greek	T	2128	0.34	0.36	0.54
Hungarian	Uralic, Ugric	T	1293	0.42	0.78	0.62
Indonesian	AN, Malayo-Sumbawan	F	4987	0.3	0.25	0.65
Italian	IE, Romance	T	12694	0.29	0.28	0.6
Latin	IE, Latin	T	10918	0.39	0.48	0.6
Latvian	IE, Baltic	T	6349	0.37	0.59	0.49
Maltese	Afro-Asiatic, Semitic	F	1785	0.31	0.23	0.64
North Sami	Uralic, Sami	T	2542	0.36	0.5	0.46
Norwegian	IE, Germanic	T	16300	0.29	0.42	0.51
Old Church Slavonic	IE, Slavic	T	4190	0.38	0.4	0.56
Old French	IE, Romance	T	13162	0.33	0.52	0.5

Old Russian	IE, Slavic	T	12079	0.34	0.44	0.55
Persian	IE, Iranian	F	5480	0.43	0.53	0.58
Polish	IE, Slavic	T	14614	0.3	0.42	0.45
Portuguese	IE, Romance	T	11103	0.26	0.28	0.59
Romanian	IE, Romance	T	8354	0.35	0.22	0.65
Russian	IE, Slavic	T	52603	0.37	0.48	0.47
Serbian	IE, Slavic	T	3691	0.36	0.43	0.49
Slovak	IE, Slavic	T	8557	0.32	0.49	0.43
Slovenian	IE, Slavic	T	6550	0.36	0.52	0.46
Spanish	IE, Romance	T	15443	0.28	0.26	0.6
Swedish	IE, Germanic	T	5490	0.36	0.47	0.48
Ukrainian	IE, Slavic	T	4982	0.37	0.47	0.47
Vietnamese	AA, Viet-Muong	F	2856	0.44	0.28	0.57
Wolof	NC, Northern Atlantic	F	1924	0.24	0.21	0.67

2.2. MEASURING DEPENDENCY LENGTH, HARMONY AND DIRECTION. Our unit of observation for quantifying dependency length, harmony and dependency direction is the CATENA. A catena is defined as a chain of words or combinations of words that is continuous along the dependency tree (Osborne 2006, Osborne, Putnam, & Groß 2012). While catenae can be used to capture dependency combinations of any length and complexity in UD, we only look at the direction and distance of two continuous dependencies in each dependency tree. This is consistent with the pairwise comparisons of harmonic word orders in typological work and guarantees sufficient numbers of tokens in each corpus.

Figure 1 contains four two-dependencies catenae, listed in Table 2. Among these, two are hierarchically connected (*fit*→*parts*→*certain*, *fit*→*parts*→*article*), and two are constructed by parallel dependencies on a single head each (*certainly*←*fit*→*parts*, *certain*←*parts*→*article*). In what follows, we call catena of the first type HIERARCHICAL CATENA, and those of the second type PARALLEL CATENAE. Note that the same pairs of words can occur in both hierarchical and parallel catenae, for instance, the dependency *parts*→*certain* occurs in both the hierarchical catena *fit*→*parts*→*certain* and the parallel catena (*certain*←*parts*→*article*).

TABLE 2. Hierarchical and parallel catenae in the English sentence from Figure 1. The numbers on the arcs indicate the length of each dependency in orthographic word units.

Hierarchical catena	
Parallel catena	

We measure dependency length as the distance between each pair of a head and its dependents in numbers of space-separated orthographic units (‘words’). For this purpose we count all units, including the function words and proper names that fall outside the range of dependencies that we target in our analysis (cf. Section 2.1). For instance, in the hierarchical catenae of the sentence in Table 2, there are three head/dependent pairs: *fit*→*parts* spans 3 units (*into*, *certain*, and *parts*), *parts*→*article* spans again 3 units (*of*, *the*, and *article*), and *parts*→*certain* spans 1 unit (*certain*). In the parallel catenae, there are four pairs: *fit*→*certainly* spans 1 unit (*certainly*), *fit*→*parts* spans 3 units (*into*, *certain*, *parts*), *parts*→*certain* spans 1 unit (*certain*), and *parts*→*article* spans 3 units (*of*, *the*, *article*). We aggregate these counts by summing them per catena type and sentence, resulting in a total length = 3+3+1 = 7 in the hierarchical and total length = 1+3+1+3 = 8 in the parallel catenae of the sentence.

We quantify harmony by the proportion of two-dependencies catenae that are harmonic in a sentence, and dependency direction by the proportion of head/dependent pairs that are head-final in a sentence, separately for hierarchical and parallel structures. For example, in Table 2 half of the hierarchical catenae (*fit*→*parts*→*article*) and none of the parallel catenae are harmonic, resulting in harmony values of 1/2 and 0, respectively. In the same example, 1 out 3 hierarchical head/dependent pairs is head-final (*parts*→*certain*) and 2 out of 4 parallel head/dependent pairs (*fit*→*certainly* and *parts*→*certain*) are head-final, resulting in head-final values of 1/3 and 1/2, respectively.

In order to study the incidence of BEHAGHEL’S LAW or the NESTED ORDERING PRINCIPLE, we also measure how many two-dependencies catenae fall into each of the following three types: short before long, long before short, or equally-sized. We determine the size of each dependent by the total number of descendants dominated by the dependent, including only lexical elements to avoid the ambiguous dependency status of function words (Section 2.1). This is roughly equivalent to the constituent size of a phrase. To further assess the impact of differently sized parallel dependents, we also examine how they are distributed on the two sides of the head. If the dependents were randomly arranged, a sentence with more parallel dependents of the same side (e.g. most head-

initial) will have a lower chance to obey nested ordering. This in turn increases the probability of having more parallel dependents of the same size on the same side of the head. For example, a parallel structure with 3 head-initial but differently sized dependents has a 16.67% chance ($1/3!$) for nested ordering, while with 4 such dependents, the chance drops to 4.17% chance ($1/4!$). We quantify the spread of the frequency distribution in terms of the Shannon entropy across differently sized dependents. A higher entropy indicates that there are more cases with more parallel dependents on the same side, for example, just as many with 2 as with 4 dependents; a lower entropy means that the distribution is skewed, for example, many more cases with 2 than with 4.

2.3. BASELINES. DLM is defined relative to what is expected under a null model or ‘baseline’. Previous work has chiefly relied on baselines derived from randomly permuting orders (Ferrer i Cancho 2004, Temperley 2008, Futrell, Mahowald, & Gibson 2015) or dependency links (Liu & Hu 2008). However, some of these baselines can lead to mixed branching directions since the order of head and dependent is randomly assigned, with few restrictions. These disharmonies risk causing artificially high dependency lengths in the baseline and, as a consequence, inflated DLM estimates in the observed data.

To guard against this risk, we introduce a baseline that is more conservative in that it preserves more information about the corpora to be studied than previous baselines. We seek to assess whether our DLM estimates go above and beyond what one would expect from the total proportions of dependency directions in any given category and language, for example, the proportion of a dependency like *verb* → ‘nobj’ → *noun* in English. This baseline can also be interpreted psycholinguistically as the simplest model of how speakers produce word order when they are not biased towards DLM, viz. simply in proportion to the dependency directions that they have experienced and learned for the same category from other speakers in their language.

More concretely, our baseline model is generated by the following steps: (1) For each type of lexical dependency (‘nsubj’, ‘obj’, ‘nmod’, etc.) in each language, we summarize the proportion of head-final dependencies, for instance, English has a proportion of .06 head-final for nmod dependents (cf. the dependency of *article* and *part* in Figure 1). These summarized proportions of each lexical dependency will serve as the base probabilities for generating random orders in a sentence. (2) With these probabilities, we then randomly re-assign directions to all dependencies in each sentence. For example, given the .06 head-final proportion of the nmod dependency, we randomly re-assign dependency directions in such a way that there is .06 probability for them to be head-final. (3) We linearize the sentence with the re-assigned dependency directions, starting from the root of a dependency tree and collecting its dependents recursively. Concretely, in Figure 1, we start from the root word *fit*, and linearize each head/dependent pair; for example, the dependency

of *fit*→*parts* will be linearized as [*fit parts*] if a head-initial order is assigned in Step 2. We continue to linearize the word *parts* and its dependencies; for example, the dependency of *parts*→*certain* will be linearized as [*certain parts*] if a head-final order is chosen in Step 2. These two linearizations then result in the sequence [*fit certain parts*]. These procedures produce a fully projective (nested) tree, and dependencies like *fit*→*parts*→*certain* cannot be converted into [*certain fit parts*] with crossing dependencies. The whole procedure is conducted recursively, until all dependencies in the sentence have been processed. When several parallel dependents are possible, we randomly choose one of them to be the first one for linearization.

Throughout this procedure we keep all elements that fall outside our purview (i.e. function words and proper names; cf. Section 2.1) unchanged. That is, they are retained in the randomizations and linearizations of each sentence. As a result, the randomized sentences will have exactly the same elements (with the same orthographic conventions) as the original sentences, making them fully comparable. Figure 2a provides one possible randomization and linearization of the sample English sentence. For each sentence, we generate 100 baseline variants.

To make our results comparable to previous studies of DLM, we furthermore adopt two alternative baselines, a free-word-order baseline and a fixed-head position baseline (Ferrer i Cancho 2004, Temperley 2008, Futrell, Mahowald, & Gibson 2015, Futrell, Levy, & Gibson 2020) which are also both fully projective and dependency-preserving. The free-word-order baseline is generated by randomly assigning an order (head-initial or head-final) for each head/dependent pair (but we again limit this to lexical dependencies), and linearizing the whole sentence accordingly. In Figure 2b, the free-word-order baseline randomizes the orders of all dependencies with equal probabilities, not reflecting usage frequencies (unlike in our baseline). Compared with our baseline, this means that the free-word-order baseline has a higher probability to move the whole oblique object in Figure 1 before the verb, causing long dependencies, as shown in Figure 2b. This lowers the bar for detecting DLM in the observed data, making the baseline the least conservative.

The fixed-head-position baseline is generated by keeping the placements of all heads and dependents unchanged, while only allowing parallel dependents at the same side of the head to switch their positions. One random linearization in Figure 2c is that all dependencies will keep the same order as the original sentence, while allowing some parallel dependents, for example, a structure like [*certain interesting parts*] will only allow *certain* and *interesting* to switch their positions before their head *parts*. The fixed-head-position baseline is the most conservative, since it holds all head directions constant. The most important difference between our baseline and the fixed-head-position baseline is that our baseline allows flexibility for single dependent and multiple parallel dependents to move to the other side of the head. For example, the order of *certainly* and *fit* can be changed in our baseline, whereas it will be kept constant in the fixed-head-position baseline.

(a) our random baseline

(b) free-word-order baseline

(c) fixed-head-position baseline

FIGURE 2. One possible randomization and linearization of the English sentence in Figure 1 by (a) our baseline, (b) the free-word-order baseline and (c) the fixed-head-position baseline

2.4. MODELS. Our dependent variable is DLM. We operationalize this by the proportion of two-dependencies catenae in the baseline randomization compared to which the observed two-dependency catena is strictly shorter in its length. For instance, in Figure 1 we observed a total length = 7 in hierarchical two-dependencies catenae. In the one randomization shown in Figure 2a, the hierarchical two-dependencies catenae have a total length = 12 (*fit*→*parts*: 7, *parts*→*certain*: 4, *parts*→*article*: 1). If the dependency length of an observed sentence is shorter than this in, say, 65 out of 100 randomizations and linearizations of the baseline model, we take this as $DLM = .65$.

We choose this metric over comparing observed and randomized sentences in counts of orthographic units (as is usually done; Ferrer i Cancho 2004, Temperley 2008, Futrell, Levy, & Gibson 2020) because it is unclear whether orthographic units correspond to (psycho-)linguistically relevant planning, comprehension and learning units (Bickel & Zúñiga 2017) and because they differ strongly across languages as to whether they include elements like clitics and particles, or the way compounds are structured and written. For example, German orthographic conventions write particles together with their verbs when they occur immediately before the verb (*wegrennen* ‘away-run’) but not elsewhere (*sie rannte weg* ‘she ran away’). The resulting absolute dependency lengths are not comparable between languages (nor indeed within German alone). By comparing observed and randomized lengths within each sentence, these lengths are taken relative to whatever convention there might hold in that sentence. This makes them fully comparable across languages and suitable data in a cross-linguistic statistical model.

We predict DLM by fitting GENERALIZED ADDITIVE MIXED MODELS with the degree of harmony (in proportions of harmonically connected lexical two-dependency catenae) and the degree of head-final (in proportions of lexical dependencies) per sentence as fixed factors. We fit random slopes by languages and random intercepts by language, language family, and geographical region (Europe vs. not Europe) (see Supporting Information for details). We choose additive models because we are not aware of any theoretical grounds that would a priori constrain effects to be strictly linear. The fixed factors and their interaction are fitted with thin-plate regression splines as a base function.⁴ Given the substantial differences in corpora size across languages (see Table 1), we adopt a quasi-binomial distribution for the response variable with varying weights per observation, such that all of the observations of each language receive the same total weight.⁵ We

⁴The regression model is formulated in the `mgcv` syntax style (Wood 2017), as

$$\begin{aligned} \text{DLM} \sim & \text{ti}(\text{harmony}) + \text{ti}(\text{head-finality}) + \text{ti}(\text{harmony}, \text{head-finality}) \\ & + \text{s}(\text{lang}, \text{harmony}, \text{bs=re}) + \text{s}(\text{lang}, \text{head-finality}, \text{bs=re}) + (\text{lang}, \text{bs=re}) \\ & + (\text{family}, \text{bs=re}) + (\text{subfamily}, \text{bs=re}) + (\text{region}, \text{bs=re}) \end{aligned}$$

where the *ti* terms specify the tensor product smooth and constructs the model with main effects and interactions, and *bs=re* specifies random intercepts or slopes by certain grouping factors.

⁵A standard binomial model would assume the weight of the observation to be given by the corpus size in each language. This would favor evidence from some languages more than from others and thereby skew the effects of

perform all model fitting with the `bam` function in R’s `mgcv` package (Wood 2017).

We fit a set of nested models and assess the interaction between the effect of harmony and head-final proportions. We deploy marginal models separately for parallel and hierarchical catenae (rather than including this distinction as a predictor) because the two types share many dependencies in each sentence, making their estimates depend on each other. For example, in Table 2 *parts*→*certain* is both part of the hierarchical catena *fit*→*parts*→*certain* and the parallel catena *certain*←*parts*→*article*.

We evaluate the joint effect of harmony and head-finality by performing *F*-tests with the `stats` package and also by comparing mean absolute errors (MAE) in the predictions of models with vs. without an interaction term through 10-fold cross-validation in R (R Core Team R Core Team 2020). We use maximum likelihood (ML) to estimate and compare the smoothing parameters between different models, and (for efficiency) fast restricted maximum likelihood (fREML) estimation for cross-validation.

Code and data summaries are available in the Supporting Information; executable scripts and input data are accessible in an Open Science Foundation repository.⁶

3. RESULTS. Figure 3 shows the degree of DLM for the raw data, and the fitted values of a model that allows an interaction between harmony and head-finality at the sentence level. Models without an interaction fit the data slightly worse (parallel catenae: $\Delta MAE = .001$, $F = 228.62$, $df_1 = 16$, $df_2 = 325828$, $p < .001$; hierarchical catenae: $\Delta MAE = .003$, $F = 240.18$, $df_1 = 16$, $df_2 = 320344$, $p < .001$), especially in regions of low head-finality and high degrees of harmony ($\Delta MAE \approx [.01, .15]$, Figure S15 and S31), although they provide a roughly comparable picture. For both parallel and hierarchical catenae, harmony and head-finality have significant main effects on DLM ($\Delta MAE \approx [.001, .014]$, all $p < .001$, see Tables S2 and S4 in Supporting Information).

The plots in Figure 3 suggest that overall, head-initial structures (i.e. the left side of each plot) show considerably more DLM than head-final structures. This is consistent with previous corpus-based research (Park & Levy 2009, Futrell, Mahowald, & Gibson 2015, Ferrer i Cancho & Gómez-Rodríguez 2019, Futrell, Levy, & Gibson 2020, Yadav et al. 2020). For parallel catenae, we observe an anti-DLM effect in head-final structures (i.e. longer dependencies than in the baseline). This effect is especially pronounced in more harmonic structures. For hierarchical catenae, by contrast, the degree of DLM appears to be linked more to harmony than to dependency direction (head-finality). In addition, our results show that a moderate amount of harmony (around 0.8) can achieve higher DLM (darker red in the plots) than completely harmonic ordering, in both parallel

specific dependency directions and harmony degrees present in the languages with most data. The disharmonic conditions of Persian for example would have less weight than the harmonic conditions of English (cf. example 5). See section S4 in the Supporting Information for further discussion.

⁶<https://osf.io/a2jvw/>

FIGURE 3. Plots of empirical distributions of DLM (upper row) and contour plots of the fitted DLM (lower row) from models allowing interaction effects of head-finality and harmony, and random slopes by languages in parallel (a) and hierarchical catenae (b). Red colors indicate that observed dependencies are more often shorter than expected under the baseline (null) model (i.e. they suggest effects from locality), while blue colors indicate that they are more often longer than expected (i.e. they suggest effects from anti-DLM). Gray areas indicate areas that lack data.⁷ For models without an interaction term, see Figures S8 and S24 and for models based on alternative baselines see Figures S10 and S26.

and hierarchical catenae.

Very similar patterns emerge under the free-word-order and fixed-head-position baselines (see Figures S10 and S26), despite the concerns we raised in Section 2.2. This suggests that DLM effects are remarkably robust against different ways of conceptualizing the null hypothesis, consistent with earlier work that compared results under different baselines (Ferrer i Cancho 2004, Liu 2008, Temperley 2008, Futrell, Mahowald, & Gibson 2015, Ferrer i Cancho & Gómez-Rodríguez 2019, Futrell, Levy, & Gibson 2020).

The lower panels of Figure 3 show the mean effects after controlling for individual variation between languages (through random slope effects), though we also notice considerable variation in the regions where anti-DLM is observed across languages (see Figures S11 and S27). For example, Danish, English and Estonian tend to show very limited amounts of anti-DLM, while all ancient languages, like Ancient Greek, Latin and Classical Chinese, and a few modern languages show stronger anti-DLM. We include more detailed results at the level of individual languages in the Supporting Information (Figures S16-S18 and S32-S34, and Section S4.4) but our main interest is in general effects that hold after language-specific variation has been factored out.

In order to assess the predictions from BEHAGHEL’S LAW and NESTED ORDERING on dependency lengths, we further modelled the ordering of parallel dependents by partitioning the sentences according to different thresholds of head-final proportions, language family, and geographical region as random intercepts (see Section S5.2.1 in the Supporting Information for further details).⁷ While short-before-long is the dominant pattern in head-initial sentences, head-final sentences are dominated by equally-sized and long-before-short dependencies. Also, the proportions of equally-sized dependents in head-final structures are always higher than those in head-initial structures. This difference is larger in more harmonic than in less harmonic sentences. These results are not fully compatible with either BEHAGHEL’S LAW, which predicts short before long orders across all conditions, or with NESTED ORDERING, which predicts short before long orders in head-initial and long before short orders in head-final structures, regardless of harmony.

A methodological problem noted in Section 2.2 is that nested ordering of parallel dependents is easier to achieve with fewer dependents than with more dependents. To capture this potential variation across languages, we examine how parallel dependents are distributed before vs after the head. For this, we computed the (Shannon) entropy of the dependents’ frequency distribution. Higher entropies indicate that it is more likely for multiple dependents to be placed on the same side, while lower entropies mean that it is less likely for multiple dependents to be placed on the same side. Figure 5 summarizes the distribution of parallel dependents on the same side of a

⁷The quasi-binomial regression model is formulated as: $\text{proportion} \sim \text{order} + \text{type} + \text{order:type} + (\text{family, bs=re}) + (\text{subfamily, bs=re}) + (\text{region, bs=re})$, where the *proportion* is the aggregated ordering of parallel catenae by language, *order* indicates the specific orders (short-before-long, long-before-short or equally-sized) of parallel dependents, and *type* is the head-initial or head-final sentences determined by the thresholds.

FIGURE 4. Fitted orderings of parallel dependents between head-initial (HI) and head-final (HF) sentences

head in terms of entropy, subsampled and aggregated by language, and its relationship with head-finality (see Methods and Supporting Materials, S5.2.3). For higher degrees of harmony, entropy is positively correlated with head-finality (Generalized Additive Model with cubic regression splines as smoothers, and random smooths by different harmony thresholds, language family, and geographical region as random intercepts, $\chi^2 = 70.69$, $df = 3$, $p < .001$, $\Delta AIC = -169.17$, see S5.2.3 in the Supporting Materials).⁸ This suggests that parallel dependents tend to pile up more before the head (head-final) than after the head, and the correlation becomes stronger in more harmonic sentences.⁹

4. DISCUSSION. Our results reveal a consistent negative relationship between the proportion of head-final structures (“head-finality” in the figures) and DLM. In addition, we observe that the association between harmony and DLM is much stronger in hierarchical than in parallel catenae. Together, these two correlations lead to anti-DLM in catenae with high proportions of harmony and head-final structures, especially when these are parallel catenae. This means that in these regions, dependencies are longer than in most of the baseline randomizations.¹⁰

These results are robust against different baselines. This includes baselines from previous research (Ferrer i Cancho 2004, Temperley 2008, Futrell, Mahowald, & Gibson 2015, Futrell, Levy, &

⁸The regression model is: $\text{entropy} \sim s(\text{head-finality}) + s(\text{head-finality, threshold, bs=fs}) + (\text{family, bs=re}) + (\text{subfamily, bs=re}) + (\text{region, bs=re})$, where $bs=fs$ indicates the factor smooths by different thresholds, and $bs=re$ specifies random intercepts by certain grouping factors.

⁹The outlier in the lower right panel is Persian, which has many more pile-ups of parallel dependents before the head than other languages.

¹⁰Anti-DLM is also found in the lower-left corner of hierarchical catenae (Figure 3b), but this effect is driven by only a few languages and data points. We return to this at the end of the section.

FIGURE 5. Relationship between head-finality and the frequency distribution of parallel dependents on the same side of a head, measured by the Shannon entropy H . Higher entropies mean that counts of catenae with many dependents are nearly as likely as counts of catenae with fewer dependents. Lower entropies mean that most catenae have few dependents and only few have many. Each point indicates the estimated entropy values in a language within the thresholds of harmony. The grey-shaded ribbons indicate 95% confidence intervals around the generalized additive model smooth.

Gibson 2020), which begs the question why the specific anti-DLM patterns we observe here escaped attention in previous work except for one recent study (Ferrer i Cancho & Gómez-Rodríguez 2019), and occasional observations of reduced DLM effects in head-final languages (Park & Levy 2009, Futrell, Mahowald, & Gibson 2015, Futrell, Levy, & Gibson 2020, Yadav et al. 2020). We suspect that the main reason is that our models are more robust against cross-linguistic variation. Specifically, our measure of DLM is insensitive to orthography conventions and sidesteps the head-dependent ambiguity of functional elements by limiting the analysis to lexical items only. Also, we focus on general effects that hold after accounting for the variation by language. In what follows we attempt to assess the underlying principles behind our observations.

4.1. PARALLEL CATENAE. We first consider anti-DLM in parallel catenae (top right corner of Figure 3a). As noted in the introduction, one possible explanation invokes BEHAGHEL’s LAW which favors short before long orders independently of head order and harmony. In head-final structures, this will put longer dependents closer to the head, increasing dependency lengths and thereby accounting for the anti-DLM effect (see the examples in 4). However, closer inspection shows that head-final catenae do not in fact prefer short before long orders. Figure 4 shows that this is in fact the least common order when sentences are mostly head-final. Short before long orders make up a sizable proportion of all dependencies only when sentences with lower head-final proportions are included in the analysis (left and middle panels in Figure 4).

The distribution in Figure 4 is also inconsistent with the NESTED ORDERING PRINCIPLE, which predicts a preference for long before short orders in head-final structures. We find equally sized orders to be nearly as likely as long before short orders. More importantly, neither BEHAGHEL’s LAW nor the NESTED ORDERING PRINCIPLE correctly predicts the effect observed in Figure 5, which shows that parallel dependents tend to pile up more before the head than after it, and that this trend increases the more harmonic the dependencies are.

This pile-up is consistent with the anti-DLM effect (Figure 3a) and the increased amount of equally-sized dependents (Figure 4) before the head (i.e. in head-final structures): The pile-up decreases the probability of putting the shortest dependents closest to the head. But what principle might this pile-up and its anti-DLM effect reflect?

A possible candidate principle is the probabilistic version of the Final-Over-Final Condition (FOFC) that we discussed in the introduction. Such a FOFC favors harmonic dependencies inside head-final more than inside head-initial structures (see example 6d). This results in piling up parallel dependents in head-final structures (i.e. before the head) more frequently than in head-initial structures (i.e. after the head), consistent with our results in Figure 5. This harmonic pile-up in head-final structures in turn increases the chances of more equally-sized dependents, and it decreases the probability of nested ordering. Together these effects lead to the anti-DLM effect we

observe in the top right corner of Figure 3a.

TABLE 3. Schematic representations of FOFC effects on parallel catenae. The graphs in each row are ordered according to their distributions across sentences that are dominated by head-initial or head-final dependencies, roughly corresponding to the different regions in Figure 3a. For example, a sentence with [V [Adv V Adv]] order is more likely to occur in sentences with more head-initial dependencies (see Figure S39 in the Supporting Materials). The red arcs indicate FOFC violations. These offending structures are predicted to occur with decreased probabilities, thereby increasing the probability of alternative structures. The net result is that sentences with more head-final structures will show anti-DLM while more head-initial sentences allow more mixed structures with DLM. This is in line with the results in Figure 3a.

type	Head-initial	Mixed	Head-final	
adverbs	VERB VERB ADV ADV	VERB ADV VERB ADV	ADV VERB ADV VERB	ADV ADV VERB VERB
adjectives	NOUN ADJ ADJ ADV	ADJ NOUN ADJ ADV	ADJ ADV ADJ NOUN	ADV ADJ ADJ NOUN

For example, the FOFC prefers adverbs to precede the verb they modify in head-final structures, causing more catenae like [[ADV ADV] V], while it allows them to either precede or follow the verb in head-initial structures like [[ADV] V [ADV]] (see Table 3 for the placement of adverbs in parallel catenae, and Figure S39 for cross-linguistic distribution). This asymmetry of adverbs can arguably be attributed to the effects of FOFC, which restricts the possibility of *[V ADV V], leading to more pile-ups of adverbs before the verb, as shown in Table 3, but allows the disharmonic catena [V [ADV V]], leading to more balanced placement of adverbs in head-initial structures (Ernst 2003, 2004, Sheehan 2017a).

Similarly, the FOFC makes it more likely for two or more adjectives to pile up before the noun, while they can appear at different sides of the noun in head-initial structures (see Table 3 for adjectives in parallel catenae, and Figure S40 for the asymmetric distributions of adjectives cross-linguistically). This is consistent with the typological observations that prenominal adjectives are more constrained than postnominal adjectives (Greenberg 1963, Williams 1982, Croft 2002, Cysouw 2010, Dryer 2018). For example, French does not allow a prenominal adjective with a post-dependent, and will extrapose the whole AP rightwards, whereas German can prepose the (adpositional) dependent of an adjective to avoid FOFC violations (see Siewierska & Uhlířová 1998, Abeillé & Godard 2000, Osborne 2003, Sheehan 2017b). The different strategies adopted to avoid the FOFC can cause more pile-ups of adjectives before the noun than after the noun.

Other parallel catenae, such as $[[S\ O]\ V]$, $[[ADV\ N]\ V]$, and $[[IO\ O]\ V]$, also show positive correlations between head-finality and harmony (Figure S38). While this can also be captured by the FOFC, we cannot rule out effects from alternative principles. For instance, the harmony of $[[S\ O]\ V]$ might reflect the S before O preference (Greenberg 1963, Hawkins 1983, Tomlin 1986, Dryer 2013c, 2013b, Napoli & Sutton-Spence 2014, Riesberg, Malcher, & Himmelmann 2019). The preference for pile-ups like $[[ADV\ N]\ V]$ can be captured by Behaghel’s Law, which predicts that the shorter adverbs are preferentially placed before the longer oblique objects. For $[[IO\ O]\ V]$, it might reflect the effect of animacy on word order, for instance, the animate indirect object comes before the inanimate direct object.

The analysis outlined here suggests that the FOFC might indeed play a role in accounting for the variance of DLM in parallel catenae. More specifically, the FOFC makes adverbs in verb-headed catenae and adjectives in noun-headed catenae more harmonic before than after the head. These asymmetric distributions of adverbials and adjective modifiers capture the higher amount of equally-sized parallel dependents in head-final structures, leading to the anti-DLM effects observed in Figure 3a.

Crucially, dependency locality has been proposed as a potential explanation for the limited amount of disharmonic word order, especially for short parallel dependents, such as adjectives in English and Romance languages (Temperley 2008, Temperley & Gildea 2018, Gulordava 2018), but dependency locality fails to predict the observed asymmetry that the short dependents, such as adjectives and adverbs in parallel structures, are more likely to pile up before than after the head. If the limited disharmonic placements of adjectives and adverbs are driven by dependency locality, we might also expect that certain amounts of parallel adjectives and adverbs are preferentially disharmonic in head-final catenae. But we find the opposite pattern, viz. more pile-ups of adjectives and adverbs in head-final than head-initial structures. These asymmetries of adjectives and adverbs in parallel catenae thus might indeed reflect a (probabilistic) FOFC principle, rather than the dependency locality principle.

4.2. HIERARCHICAL CATENAE. Like parallel catenae, hierarchical catenae also show anti-DLM patterns dependent on head direction: Figure 3b shows that sentences with higher proportions of head-final dependencies show anti-DLM (i.e. longer than expected dependencies (blue)) while sentences with more head-initial dependencies show DLM (i.e. shorter than expected dependencies (red)). Unlike in parallel catenae, however, this effect is orthogonal to the overall positive correlation between DLM and harmony, which is consistent with earlier findings (Temperley 2008, Hawkins 2014, Futrell, Mahowald, & Gibson 2015). To some extent, these patterns can again be captured by the probabilistic FOFC. Evidence for this can best be explored by assessing what happens when head-initial vs head-final dependencies are embedded in sentences. More specifically, we examine

the relationship between dependency lengths and harmony when a dependency is embedded under the next higher node in its sentence. For example, given a genitive construction (*the office of the president*, head-initial vs. *the president's office*, head-final) and its dependency length, we quantify the probability of forming a harmonic catena inside a verb phrase, such as [*visit [the office of the president]*] or [*[the president's office] visit*], in our sample languages. Figure 6 summarizes the results across all catenae and parts of speech. The results suggest that head-initial dependencies have a higher probability to be embedded in the next higher head-initial dependency (entailing a high degree harmony), whereas a head-final dependency can be dominated just as easily by a head-initial or by a head-final dependency (entailing a lower degree of harmony). This asymmetry becomes stronger for longer dependencies ($\chi^2 = 786.82$, $df = 4$, $p < .001$, $\Delta AIC = -1563.2$). This observation is fully in line with the predictions from the FOFC, leading to more harmony and DLM with head-initial than head-final structures.

FIGURE 6. Interaction of dependency length and direction of the bottom dependency on harmony in hierarchical catenae. The colored ribbons indicate 95% confidence intervals around the generalized additive model smooth.

Further support for this comes from comparing the harmony of major hierarchical catenae in head-initial and head-final structures. For example, hierarchical catenae like relative clauses ($N \rightarrow V \rightarrow O$), complement clauses ($V \rightarrow V \rightarrow O$), and genitive constructions ($V \rightarrow N \rightarrow Gen$) show more harmony in head-initial than in head-final languages (see Table 4 and Figure S43). These asymmetries are consistent with the effects of a probabilistic FOFC. The FOFC disprefers head-initial dependencies embedded before their heads (see example 6d), that is, catenae like $*[[V\ O]\ N]$, $*[[V\ O]\ V]$ and $*[[N\ N]\ V]$ are disfavored. Embeddings after the head are more likely, and these can be – and in fact mostly are – disharmonic, as in $[N\ [O\ V]]$, $[V\ [O\ V]]$ and $[V\ [N\ N]]$ in Table 4, leading to the anti-DLM effect that we observe in head-final structures in Fig 3b.

As noted in the introduction, these FOFC predictions are fully in line with various typological observations, such as the observation that VO languages strongly favor harmonic, post-nominal relative clauses (i.e. [N [V O]] vs dispreferred *[[V O] N]), while OV languages have both pre-nominal and post-nominal relative clauses (i.e. either [[V O] N] or [[O V] N]) (Greenberg 1963, Lehmann 1984, Hawkins 1994).

TABLE 4. Schematic representations of how the FOFC impacts hierarchical catenae. The red arcs indicate the dependencies that violate the FOFC and are thus dispreferred, making room for increased probabilities of alternative catenae. This in turn increases the probability of embedding head-initial dependencies inside harmonic, DLM-favoring catenae, while head-final dependencies are embedded as easily in harmonic, DLM-favoring catenae as in disharmonic, DLM-disfavoring catenae.

type	Head-initial	Mixed	Head-final
relative clause			
complement clause			
genitive			

These observations notwithstanding, we also note an anti-DLM pattern in the bottom-left corner in Figure 3b. This pattern is not accounted for by a probabilistic FOFC, nor by any other principle. To better understand the pattern, we extracted the data in this region and found that the pattern is largely caused by specific hierarchical catenae in Romance languages, such as [[S ADJ] V] and [[S Gen] V]. All these hierarchical catenae involve lexical subjects followed by adjectives and genitival expressions before the verb. The anti-DLM effect in the bottom-left corner is driven by these data, and they are relatively sparse (1% on average, see also Figure S30).

Finally, in both parallel and hierarchical catenae in Figure 3 moderate harmony (around .8) yields higher DLM (darker red) than complete harmony. This is consistent with earlier findings that a certain amount of disharmony in parallel catenae can facilitate DLM (Temperley 2008, Temperley & Gildea 2018). But it seems that this effect generalizes to hierarchical catenae. Of course, this might be confounded by the shared dependencies between parallel and hierarchical catenae, and more research is needed to disentangle these.

5. CONCLUSIONS. Our results reveal surprisingly strong anti-DLM effects (i.e. longer than expected dependencies) under some conditions of harmony and head position once language-specific

variation is factored out. The effects differ between parallel and hierarchical catenae but both suggest a possible role for a probabilistic version of the FOFC, which disprefers head-initial dependencies directly dominated by head-final dependencies. In parallel catenae, the FOFC favors more harmony in head-final structures, especially for adjective and adverbial modifiers, resulting in pile-ups of parallel dependents before the head and thereby causing longer dependencies. In hierarchical catenae, the FOFC leads to disharmony, and therefore longer dependencies, since head-final dependencies can be as easily embedded inside head-initial and head-final dependencies, while head-initial dependencies tend to be embedded inside head-initial dependencies only. However, at least for some of these effects, we note that alternative accounts cannot be ruled out. For example, the general preference for subject-before-object orders induces longer dependencies in verb-final sentences, and no appeal might be needed to a FOFC principle.

Progress here will need a systematic comparison of what each preference and principle is ultimately based on and of where its predictions might differ. In order to move there, more extensive modeling is needed that (i) overcomes the European bias of the UD dataset, (ii) controls for the presence of (various kinds of) case and agreement morphology (Ros et al. 2015), and (iii) assesses the extent to which the FOFC can be grounded in mechanisms of comprehension, production, or learning. At present, the FOFC is a descriptive generalization, not an explanation.

REFERENCES

- ABEILLÉ, ANNE, and GODARD, DANIELE. 2000. French word order and lexical weight. *The nature and function of Syntactic Categories*, ed. by Borsley Robert, 325–360. San Diego: Academic Press.
- BEHAGHEL, OTTO. 1909. Beziehungen zwischen Umfang und Reihenfolge von Satzgliedern. *Indogermanische Forschungen* 25.110–142. DOI: 10.1515/9783110242652.110.
- BIBERAUER, THERESA; HOLMBERG, ANDERS; and ROBERTS, IAN. 2014. A syntactic universal and its consequences. *Linguistic Inquiry* 45(2).169–225. DOI: 10.1162/LING_a_00153.
- BIBERAUER, THERESA; HOLMBERG, ANDERS; ROBERTS, IAN; and SHEEHAN, MICHELLE. 2017a. Empirical Evidence for the Final-over-Final Condition. *The Final-Over-Final Condition: A Syntactic Universal*, ed. by Michelle Sheehan, Theresa Biberauer, Ian Roberts, and Anders Holmberg, 11–26. Cambridge, Massachusetts: MIT Press.
- BIBERAUER, THERESA; HOLMBERG, ANDERS; ROBERTS, IAN; and SHEEHAN, MICHELLE. 2017b. Harmony, Symmetry, and Dominance in Word Order Universals. *The Final-Over-Final Condition: A Syntactic Universal*, ed. by Michelle Sheehan, Theresa Biberauer, Ian Roberts, and Anders Holmberg, 27–42. Cambridge, Massachusetts: MIT Press.
- BICKEL, BALTHASAR. 2015. Distributional typology: statistical inquiries into the dynamics of linguistic diversity. *The Oxford Handbook of Linguistic Analysis, 2nd edition*, ed. by Bernd Heine and Heiko Narrog, 901–923. Oxford: Oxford University Press.
- BICKEL, BALTHASAR, and ZÚÑIGA, FERNANDO. 2017. The ‘word’ in polysynthetic languages: phonological and syntactic challenges. *The Oxford Handbook of Polysynthesis*, ed. by Michael D. Fortescue, Marianne Mithun, and Nicholas Evans, 158–185. Oxford: Oxford University Press.
- CHOI, HYE-WON. 2007. Length and order: a corpus study of Korean dative-accusative constructions. *Discourse and Cognition* 14.207–227. DOI: 10.15718/DISCOG.2007.14.3.207.
- CROFT, WILLIAM. 2002. *Typology and universals*. Cambridge: Cambridge University Press.
- CYSOUW, MICHAEL. 2010. Dealing with diversity: Towards an explanation of NP-internal word order frequencies. *Linguistic Typology* 14(2-3).253–286. DOI: 10.1515/lity.2010.010.
- DRYER, MATTHEW S. 1992. The Greenbergian word order correlations. *Language* 68(1).81–138. DOI: 10.2307/416370.
- DRYER, MATTHEW S. 2013a. Relationship between the Order of Object and Verb and the Order of Relative Clause and Noun. *The World Atlas of Language Structures*, ed. by Matthew S. Dryer and Martin Haspelmath. Leipzig: Max Planck Institute for Evolutionary Anthropology. Online: <https://wals.info/chapter/96>.
- DRYER, MATTHEW S. 2013b. Order of Subject, Object and Verb. *The World Atlas of Language Structures Online*, ed. by Matthew S. Dryer and Martin Haspelmath. Leipzig: Max Planck Institute for Evolutionary Anthropology. Online: <https://wals.info/chapter/81>.

- DRYER, MATTHEW S. 2013c. On the six-way word order typology, again. *Studies in Language* 37(2).267–301. DOI: 10.1075/sl.37.2.02dry.
- DRYER, MATTHEW S. 2018. On the order of demonstrative, numeral, adjective, and noun. *Language* 94(4).798–833. DOI: 10.1353/lan.0.0232.
- DRYER, MATTHEW S., and GENSLER, ORIN D. 2005. Order of Object, Oblique, and Verb. *The World Atlas of Language Structures Online*, ed. by Matthew S. Dryer and Martin Haspelmath. Leipzig: Max Planck Institute for Evolutionary Anthropology. Online: <https://wals.info/chapter/84>.
- ERNST, THOMAS. 2003. Adjuncts and word order asymmetries. *Asymmetry in Grammar*, ed. by Anna Maria Di Sciullo, 187–208. Amsterdam: John Benjamins.
- ERNST, THOMAS. 2004. Principles of adverbial distribution in the lower clause. *Lingua* 114(6).755–777. DOI: 10.1016/S0024-3841(03)00050-0.
- FERRER I CANCHO, RAMON. 2004. Euclidean distance between syntactically linked words. *Physical Review E* 70(5).056135. DOI: 10.1103/PhysRevE.70.056135.
- FERRER I CANCHO, RAMON, and GÓMEZ-RODRÍGUEZ, CARLOS. 2019. Anti dependency distance minimization in short sequences. A graph theoretic approach. *Journal of Quantitative Linguistics*.1–27. DOI: 10.1080/09296174.2019.1645547.
- FERRER I CANCHO, RAMON; GÓMEZ-RODRÍGUEZ, CARLOS; and ESTEBAN, JUAN LUIS. 2018. Are crossing dependencies really scarce? *Physica A: Statistical Mechanics and its Applications* 493.311–329. DOI: 10.1016/j.physa.2017.10.048.
- FUTRELL, RICHARD; LEVY, ROGER; and GIBSON, EDWARD. 2020. Dependency locality as an explanatory principle for word order. *Language* 96(2).371–412. DOI: 10.1353/lan.2020.0024.
- FUTRELL, RICHARD; MAHOWALD, KYLE; and GIBSON, EDWARD. 2015. Large-scale evidence of dependency length minimization in 37 languages. *Proceedings of the National Academy of Sciences* 112(33).10336–10341. DOI: 10.1073/pnas.1502134112.
- GIBSON, EDWARD. 1998. Linguistic complexity: Locality of syntactic dependencies. *Cognition* 68(1).1–76. DOI: 10.1016/S0010-0277(98)00034-1.
- GIBSON, EDWARD. 2000. The dependency locality theory: A distance-based theory of linguistic complexity. *Image, language, brain: Papers from the first mind articulation project symposium*, ed. by Alec Marantz, Yasushi Miyashita, and Wayne O’Neil, 95–126. Cambridge, Massachusetts: MIT Press.
- GILDEA, DANIEL, and TEMPERLEY, DAVID. 2010. Do grammars minimize dependency length? *Cognitive Science* 34(2).286–310. DOI: 10.1111/j.1551-6709.2009.01073.x.
- GREENBERG, JOSEPH H. 1963. Some universals of grammar with particular reference to the order of meaningful elements. *Universals of Language*, ed. by Joseph H. Greenberg, 58–60. Cambridge, Massachusetts: MIT Press.

- GRODNER, DANIEL, and GIBSON, EDWARD. 2005. Consequences of the serial nature of linguistic input for sentential complexity. *Cognitive Science* 29(2).261–290. DOI: 10.1207/s15516709cog0000_7.
- GRODNER, DANIEL; GIBSON, EDWARD; and WATSON, DUANE. 2005. The influence of contextual contrast on syntactic processing: evidence for strong-interaction in sentence comprehension. *Cognition* 95(3).275–296. DOI: 10.1016/j.cognition.2004.01.007.
- GULORDAVA, KRISTINA. 2018. *Word order variation and dependency length minimisation: a cross-linguistic computational approach*. Geneva: University of Geneva, Geneva dissertation. DOI: 10.13097/archive-ouverte/unige:106855.
- HAWKINS, JOHN A. 1983. *Word order universals*. New York: Academic Press.
- HAWKINS, JOHN A. 1994. *A performance theory of order and constituency*. Cambridge: Cambridge University Press.
- HAWKINS, JOHN A. 2004. *Efficiency and complexity in grammars*. Oxford: Oxford University Press.
- HAWKINS, JOHN A. 2014. *Cross-linguistic variation and efficiency*. Oxford: Oxford University Press.
- HUDSON, RICHARD A. 1984. *Word grammar*. New York: Blackwell Oxford.
- HUSAIN, SAMAR; VASISHTH, SHRAVAN; and SRINIVASAN, NARAYANAN. 2014. Strong expectations cancel locality effects: Evidence from Hindi. *PloS ONE* 9(7).e100986. DOI: 10.1371/journal.pone.0100986.
- KONIECZNY, LARS. 2000. Locality and parsing complexity. *Journal of Psycholinguistic Research* 29(6).627–645. DOI: 10.1023/A:1026528912821.
- LEHMANN, CHRISTIAN. 1984. *Der Relativsatz*. Tübingen: Gunter Narr.
- LEHMANN, WINFRED P. 1973. A structural principle of language and its implications. *Language* 49(1).47–66. DOI: 10.2307/412102.
- LEVY, ROGER. 2008. Expectation-based syntactic comprehension. *Cognition* 106(3).1126–1177. DOI: 10.1016/j.cognition.2007.05.006.
- LEVY, ROGER, and KELLER, FRANK. 2013. Expectation and locality effects in German verb-final structures. *Journal of Memory and Language* 68(2).199–222. DOI: 10.1016/j.jml.2012.02.005.
- LEWIS, RICHARD L., and VASISHTH, SHRAVAN. 2005. An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science* 29(3).375–419. DOI: 10.1207/s15516709cog0000_25.
- LIU, HAITAO. 2008. Dependency distance as a metric of language comprehension difficulty. *Journal of Cognitive Science* 9(2).159–191. DOI: 10.17791/JCS.2008.9.2.159.
- LIU, HAITAO. 2009. *Dependency Grammar from theory to practice*. Beijing: Science Press.
- LIU, HAITAO, and HU, FENGGUO. 2008. What role does syntax play in a language network? *EPL (Europhysics Letters)* 83(1).18002. DOI: 10.1209/0295-5075/83/18002.
- MACDONALD, MARYELLEN C. 2013. How language production shapes language form and comprehension. *Frontiers in Psychology* 4.226. DOI: 10.3389/fpsyg.2013.00226.

- MEL'ČUK, IGOR. 1988. *Dependency syntax: Theory and practice*. Albany: State University Press of New York.
- NAPOLI, DONNA JO, and SUTTON-SPENCE, RACHEL. 2014. Order of the major constituents in sign languages: implications for all language. *Frontiers in Psychology* 5.376. DOI: 10.3389/fpsyg.2014.00376.
- NIVRE, JOAKIM; ABRAMS, MITCHELL; AGIĆ, ŽELJKO; AHRENBURG, LARS; ALEKSANDRAVIČIŪTĖ, GABRIELĖ; ANTONSEN, LENE; APLONOVA, KATYA; ARANZABE, MARIA JESUS; ARUTIE, GASHAW; and ASAHARA, MASAYUKI. et. al. 2019. *Universal Dependencies 2.5*. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics, Faculty of Mathematics and Physics, Charles University. Online: <http://hdl.handle.net/11234/1-3105>.
- OSBORNE, TIMOTHY. 2003. The left elbow constraint. *Studia Linguistica* 57(3).233–257. DOI: 10.1111/j.0039-3193.2003.00105.x.
- OSBORNE, TIMOTHY. 2006. Beyond the Constituent-A Dependency Grammar Analysis of Chains. *Folia Linguistica* 39(3-4).251–297. DOI: 10.1515/flin.2006.39.3-4.251.
- OSBORNE, TIMOTHY; PUTNAM, MICHAEL; and GROß, THOMAS. 2012. Catenae: Introducing a novel unit of syntactic analysis. *Syntax* 15(4).354–396. DOI: 10.1111/j.1467-9612.2012.00172.x.
- PARK, Y ALBERT, and LEVY, ROGER. 2009. Minimal-length linearizations for mildly context-sensitive dependency trees. *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*.335–343.
- R CORE TEAM. 2020. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Online: <https://www.R-project.org/>.
- RIESBERG, SONJA; MALCHER, KURT; and HIMMELMANN, NIKOLAUS P. 2019. How universal is agent-first? Evidence from symmetrical voice languages. *Language* 95(3).523–561. DOI: 10.1353/lan.2019.0055.
- ROS, IDOIA; SANTESTEBAN, MIKEL; FUKUMURA, KUMIKO; and LAKA, ITZIAR. 2015. Aiming at shorter dependencies: the role of agreement morphology. *Language, Cognition, and Neuroscience* 30.1156–1174. DOI: 10.1080/23273798.2014.994009.
- SAUPPE, SEBASTIAN; CHOUDHARY, KAMAL K; GIROUD, NATHALIE; BLASI, DAMIÁN E; NORCLIFFE, ELISABETH; BHATTAMISHRA, SHIKHA; GULATI, MAHIMA; EGURTZEGI, AITOR; BORNKESSEL-SCHLESEWSKY, INA; MEYER, MARTIN; and BICKEL, BALTHASAR. 2021. Neural signatures of syntactic variation in speech planning. *PLoS Biology* 19(1).e3001038. DOI: 10.1371/journal.pbio.3001038.
- SHEEHAN, MICHELLE. 2017. The Final-over-Final Condition and Adverbs. *The Final-Over-Final Condition: A Syntactic Universal*, ed. by Michelle Sheehan, Theresa Biberauer, Ian Roberts, and Anders Holmberg, 97–120. Cambridge, Massachusetts: MIT Press.

- SHEEHAN, MICHELLE. 2017. The Final-over-Final Condition and the Head-Final Filter. *The Final-Over-Final Condition: A Syntactic Universal*, ed. by Michelle Sheehan, Theresa Biberauer, Ian Roberts, and Anders Holmberg, 121–150. Cambridge, Massachusetts: MIT Press.
- SIEWIERSKA, ANNA, and UHLÍŘOVÁ, LUDMILA. 1998. An overview of word order in Slavic languages. *Constituent Order in the Languages of Europe*, ed. by Anna Siewierska, 105–149. Berlin: Mouton de Gruyter.
- TEMPERLEY, DAVID. 2008. Dependency-length minimization in natural and artificial languages. *Journal of Quantitative Linguistics* 15(3).256–282. DOI: 10.1080/09296170802159512.
- TEMPERLEY, DAVID, and GILDEA, DANIEL. 2018. Minimizing Syntactic Dependency Lengths: Typological/Cognitive Universal? *Annual Review of Linguistics* 4(1).67–80. DOI: 10.1146/annurev-linguistics-011817-045617.
- TOMLIN, RUSSELL S. 1986. *Basic Word Order: Functional Principles*. London: Croom Helm.
- VASISHTH, SHRAVAN, and LEWIS, RICHARD L. 2006. Argument-head distance and processing complexity: Explaining both locality and antilocality effects. *Language* 767–794. DOI: 10.1353/lan.2006.0236.
- VENNEMANN, THEO. 1974. Topics, subjects, and word order: from SXV to SVX via TVX. *Historical Linguistics*, ed. by Anderson John and Jones Charles, 339–76. Amsterdam: North Holland.
- WILLIAMS, EDWIN. 1982. Another argument that passive is transformational. *Linguistic Inquiry* 13(1).160–163. DOI: 10.2307/4178269.
- WOOD, SIMON N. 2017. *Generalized additive models: an introduction with R*. Philadelphia: Chapman and Hall/CRC.
- YADAV, HIMANSHU; VAIDYA, ASHWINI; SHUKLA, VISHAKHA; and HUSAIN, SAMAR. 2020. Word Order Typology Interacts With Linguistic Complexity: A Cross-Linguistic Corpus Study. *Cognitive Science* 44(4).e12822. DOI: 10.1111/cogs.12822.
- YAMASHITA, HIROKO, and CHANG, FRANKLIN. 2001. “Long before short” preference in the production of a head-final language. *Cognition* 81.B45–B55. DOI: 10.1016/S0010-0277(01)00121-4.