

Concord begets concord: A Bayesian model of nominal concord typology

Kyle Mahowald, Dan Jurafsky & Mark Norris*

Abstract. Nominal concord is a phenomenon whereby nominal modifiers (e.g., adjectives, demonstratives, numerals) agree with their nominals along various dimensions (e.g., gender, number, case, definiteness). Here, drawing on a rich and typologically diverse database of nominal concord (Norris 2020), we build a Bayesian mixed effect model of nominal concord. Specifically, we consider two competing hypotheses regarding the statistical relationship between different types of concord within a language: (1) *concord begets concord*: the presence of some type of concord in a language makes it more likely that it has other types of concord vs. (2) *a little concord goes a long way*: if a language has some kind of concord, it is less likely to have other types of concord. We present evidence strongly in favor of the first hypothesis, that concord begets concord. Languages with nominal concord tend to have concord in more than one place and of more than one type. Using posterior draws from our model, we also provide quantitative evidence for a number of the tendencies described by Norris (2019a). Future work will build on this model to understand the functional role of nominal concord in language systems, how it evolves, and how it co-evolves with other typological features.

Keywords. nominal concord; Bayesian modeling; typology; probabilistic models; agreement; quantitative methods

1. Introduction. Nominal concord is a phenomenon whereby nominal modifiers (e.g., adjectives, demonstratives, numerals) agree with their nominals along various dimensions (e.g., gender, number, case, definiteness); see Norris (2014, 2017, 2019a) for detail and review. Nominal concord has major typological variation: some languages have rich nominal concord (Latvian, Greek), while others have none (Thai, Turkish). Functionally, concord can be a source of linguistic information in that agreement between a noun and its modifier reduces referential uncertainty. The extra information comes at the cost of additional complexity in the morphological system.

We can understand the functional role of concord by exploring the statistical structure of concord within languages and how concord covaries with other typological features. For instance, it has been suggested that some linguistic properties (e.g., free word order) covary with the complexity of morphological agreement systems (McFadden 2003; Futrell et al. 2015; Kiparsky 1997). If such global linguistic features affect nominal concord, we might expect that languages with gender systems, case systems, or free word order have more nominal concord across the board. On the other hand, if the complexity cost of concord is high, then we might expect that the existence of, say, gender concord in a language makes that language less likely to also have number concord or case concord. Does having one type of nominal concord make other types more likely or less likely? That is, is concord “sticky” or “repellent”? How does having a gender or case system affect the probability of concord?

* We thank the audience at our talk at LSA 2021. Authors: Kyle Mahowald, University of California Santa Barbara (mahowald@ucsb.edu) & Dan Jurafsky, Stanford University (jurafsky@stanford.edu) & Mark Norris (mjnorris@gmail.com).

Family	Yes	%	No	%	Region
Sino-Tibetan	0	0%	7	100%	Eur
Austronesian	2	25%	6	75%	Pap/Afr
Austro-Asiatic	2	33.3%	4	66.7%	Eur
Trans-New Guinea	3	37.5%	5	62.5%	Pap
Niger-Congo	4	50%	4	50%	Afr
Nakh-Daghestanian	3	75%	1	25%	Eur
Indo-European	8	88.9%	1	11.1%	Eur
Uto-Aztecan	4	100%	0	0%	NAM

Table 2. Concord in families with ≥ 4 genera in the sample

case. As examples, we provide demonstrative gender concord in Yuchi (yuc) in (3) and demonstrative case concord in Nez Perce (nez) in (4). Demonstrative number concord was exemplified in (1).³

(3) le-'e yasthæde-'e
that-III bridge-III
'that bridge' Yuchi (Linn 2001; 399)

(4) kon-ya haama-na
that-OBJ man-OBJ
'that man' Nez Perce (Deal 2010; 51)

According to Linn's (2001; 362) description, Yuchi has 6 genders for nouns, and these genders are (or can be) marked on the demonstrative in addition to the noun. In (3), we see a demonstrative *le-'e* marked for the gender we gloss as III. As for Nez Perce, according to Deal's (2010) description, demonstratives are (or can be) marked for case (and number, actually) in addition to the noun. We see this for objective case on *kon-ya* 'that-OBJ' in (4).

In addition to seeing a variety of features in concord, languages vary in terms of which lexical categories participate. Following Norris (2019a), we discuss only demonstratives, cardinal numerals (non-1 numerals), and adjectives. Demonstrative concord has been exemplified already; we show adjective concord in Ngiti (nyi) in (5) and numeral concord in Estonian (ekk) in (6).

(5) á̀ǹǹí nzo
small.PL child.PL
'small children' Ngiti (Kutsch Lojenga 1994; 345)

(6) kahe-d kangastelje-d
two-PL loom-PL
'two looms' Estonian (Norris 2018)

In the Ngiti example in (5), we see the suppletive plural form *á̀ǹǹí* 'small.PL' modifying the plural noun *nzo* 'children' (which also happens to be suppletive). In Estonian, numerals have both singular and plural forms, and plural numerals like *kahe-d* 'two-PL' are used (among other reasons) to count pluralia tantum nouns like *kangasteljed* 'loom(s)' in (6).

³ In the Yuchi example, we use III to refer to the class that Linn (2001) calls *lying* (as in lying down).

Finally, we must comment on the kinds of examples that “count” as showing concord for Norris’s (2019b) data set and thus, for us. First, the element showing concord must modify an overt noun. This is because in some languages, free-standing elements have different inflectional possibilities than those that modify an overt noun (see, e.g., Diessel (2013)). Second, the element showing concord and the overt noun must form a continuous phrase. This is because in some languages, inflectional possibilities in continuous phrases are different from those in discontinuous phrases (Clem & Dawson 2018). Since such languages are counted as not showing concord for us, this could have an overall effect of reducing the number of languages with concord (though we do not have numerical data regarding the presence of such languages in the data set).

In contrast, one area in which Norris (2019b) is generous in counting examples of concord is with regard to optionality or variation. If there are examples of concord that fit the profile for inclusion as described above, the language is coded as showing concord, even if the author described concord as optional to some degree (including “infrequent”). This may have an overall effect of increasing the number of languages showing concord, though again, we do not have numerical data regarding the frequency of such languages.

2.2. NORRIS’S (2019A) CONCORD TENDENCIES. Based on proportional results, Norris (2019a) proposes four “concord tendencies,” given below:

- (7) **Concord Tendency 1:** If a language has concord, it likely has number concord.
- (8) **Concord Tendency 2:** If a language has a grammatical gender system, it likely has gender concord.
- (9) **Concord Tendency 3:** If a language or lexical category within a language has case concord, it will likely have number concord (and if the language has gender, gender concord).
- (10) **Concord Tendency 4:** If a language has concord, it is more likely that it will have concord on both adjectives and demonstratives than on just one or the other.

We use our model to investigate these tendencies in Section 4.3.

2.3. CONCORD AND NP-INTERNAL WORD ORDER. We also consider possible effects of word order. A natural hypothesis, given the functional considerations discussed in the introduction, would be that languages with more word order freedom would have more concord. Unfortunately, we do not have adequate data on word order flexibility for the languages in our sample and so leave that for future work. As far as we know, the closest anyone has come to connecting concord and NP-internal word order is Greenberg (1963; 95) in his Universal 40. Among languages with postnominal adjectives, Universal 40 predicts two kinds of languages and precludes one type.

- (11) **Universal 40:** When the adjective follows the noun, the adjective expresses all the inflectional categories of the noun. In such cases the noun may lack overt expression of one or all of these categories (Greenberg 1963).
 - a. Predicted: N-INFL Adj-INFL (Concord)
 - b. Predicted: N-∅ Adj-INFL (Not concord)
 - c. Precluded: N-INFL Adj-∅ (Not concord)

Because Universal 40 predicts languages with concord and languages without (or more precisely, languages which do not definitely have concord, in our view), it is not specifically about concord. However, it is suggestive of a connection between concord and headedness in NPs, which made us curious about the relationship between concord and NP-internal word order.

Before continuing, we observe for the record that Farsi is potentially a language of the precluded type, as these examples demonstrate:

(12) ketab-e xub
book-EZ good
'a/the good book' Farsi (Mace 2003; 48)

(13) ketab-ha-ye xub
book-PL-EZ good
'the good books' Farsi (Mace 2003; 48)

In (13), we see a postnominal adjective, but plural *ha* appears only on the noun. We say Farsi is “potentially” a counterexample, because these are only two examples, and a more thorough investigation may reveal a more complex distribution than what we’ve presented here.

Returning to the main point, although Greenberg is talking specifically about adjectives and his conditional universal applies only to languages with postnominal adjectives, we can use our method to ask a related question: does the ordering of modifiers in the nominal phrase affect the likelihood of concord? Per Greenberg, we might expect that (at least for adjectives) noun-first languages are more likely to have concord, because Universal 40 predicts that adjectives in noun-first languages will bear the inflectional categories of the noun.⁴ We of course recognize that this hypothesis does require a small leap from Greenberg’s formulation of Universal 40, but we consider the hypothesis that concord is more common on postnominal modifiers in Section 4.4.

3. Model. For a quantitative model of nominal concord, we want something which flexibly handles non-independence in the data (language, family, and areal effects) and which also lets us simulate *new data* from a generative model (a model which allows for the generation of new simulated data) in order to explore the statistical relationship across language’s concord systems. Thus, we choose a hierarchical Bayesian model predicting concord based on the type and locus of concord, a series of language features from WALS, and random effects for language, family, and area.

A Bayesian model starts with a set of priors, which we choose here to be weakly informative (Gelman et al. 2004). Then, through a model fitting process (here, Hamiltonian Monte Carlo using the Bayesian programming language Stan (Carpenter et al. 2017)), we learn from the data in order to find a set of parameters that best fit our data.

We formulate the problem as a binary prediction task: for a given language (nested in a given family and area) and for a given type of concord (gender, number, case, definiteness) and locus (noun, adjective, demonstrative), does that language have that type of concord at

⁴ The second clause of Universal 40 presents a small complication to this prediction. Imagine a noun-first language where the adjective expresses number (for example) and in so doing, suppresses the number expression on the noun. In our view, this is not obviously concord, because it is not obviously agreement. It could be that the position of the number morpheme is determined syntactically; the reason that the adjective appears to express number is because it is the closest word to the number morpheme. Since we do not actually investigate Universal 40 in what follows (but simply draw inspiration from it), we do not discuss this complication further.

that locus? Specifically, we predict concord for a particular data point, where a data point is a binary variable stating whether a particular language has concord of a particular type at a particular locus (e.g., gender concord on adjectives in Icelandic). We make these predictions based on a set of fixed and random effects.

Crucially, our Bayesian model is *generative* in the sense that it lets us simulate data based on the learned model parameters. Intuitively, we can think of this as a more complicated version of a model that lets us simulate a coin flip. Imagine flipping a coin that you know to be fair (Heads 50% of the time and Tails 50% of the time). If you flip it 5 times, you might get the pattern HHTHT. If you flip it another 5 times, you will likely see something else: maybe TTTTH. When we sample from our generative model, we are doing something similar except in our case we are predicting not Heads or Tails but whether a given language has Concord or No Concord of a specific type at a specific locus. And, instead of having just one parameter as in the coin example, we have many parameters which reflect relevant facts like what type of concord we are dealing with, what the locus of the concord is, what kinds of concord are common in that language family and area, and so on. When we simulate in this way using our model, which has been fit to our data set, we get random draws that reflect the model’s “best guess” about the underlying structure of nominal concord in our data.

Our method here is to do thousands of these kinds of simulations: in effect, generating new random languages with a concord structure that reflects what the model has learned. Through this method, we can ask arbitrary questions about our model—which we believe is a useful tool for evaluating the kind of statistical relationships and conditional universals that characterize this kind of typological data. For instance, in all of the random languages that we generate, how many have gender concord? And of those languages which have gender concord, how many also have number concord? Because conditional claims of this nature are common in linguistic typology, we think this kind of model could be broadly useful for evaluating them.

Before turning to our results, we discuss the predictors that go into our model. Our fixed effect predictors are:

- **ConcordType**: sum-coded, one of Gender, Number, Case, Definiteness
- **ConcordLocus**: sum-coded, one of Adjective, Numeral, Demonstrative
- **ConcordType:ConcordLocus**: Interaction of type and locus
- **HasCaseMarking**: binarized WALs feature for if the language has case (from Iggesen 2013)⁵
- **NounFirst**: binarized WALs feature indicating if the noun precedes the modifier (depends on the modifier being predicted, from Dryer 2013b,c,d)
- **GenderStatus**: binarized WALs feature indicating if the language has gender (from Corbett 2013)
- **NumberStatus**: binarized WALs feature indicating if the language has number marking (from Dryer 2013a)

⁵ We should note that when a language in Norris’s data set was not part of a particular WALs data set, we either excluded it from the calculation or coded it following the criteria from the WALs chapter as we understood them.

There are random intercepts for Language, Family, and Area, with random slopes for ConcordType and ConcordLocus for each, such that each language, family, and area has its own matrix of coefficients. For a ConcordType t , a ConcordLocus o , a Language l , a Family f , and an area a , the model makes a binary prediction as to whether $y_{tol,fa}$ has concord.

We use semi-informative priors (Gelman & Hill 2006): a normal distribution with mean 0 and standard deviation of 1 for each of the fixed effect and random effect coefficients. We use a default LKJ(2) prior for the random effect correlation matrix—a standard design choice in many hierarchical Bayesian models (Vasishth et al. 2018).

By sampling from the posterior of the fit model, we obtain several kinds of results that let us quantitatively evaluate our hypotheses of interest. First, we can compare posterior estimates for different parameters (e.g., type of concord, WALs features, and area) in order to characterize languages’ overall propensity for concord, as well as to obtain model estimates for how various linguistic features affect the overall likelihood of concord. Second, we look at correlations among parameters in the posterior samples. This lets us ask questions like: “how does the probability of gender concord covary with the probability of adjective concord?” We also use the posterior draws to assess conditional claims, like Norris’s (2019a) tendencies.

4. Results.

4.1. DESCRIPTIVE RESULTS AND AREAL RESULTS. We assessed model fit by fitting the model to out-of-sample data using 5-fold cross-validation. The logic of this approach is that, by holding out some of our data and then testing on it, we can find out how good our model is. We did this in two ways: first by treating each individual data point as independent, even within a language, such that the model can, e.g., have access to information about gender concord in Icelandic when predicting whether Icelandic has number concord. Under this approach, the model accurately predicts concord for 92% of out-of-sample data points.

We also used a second k-fold validation approach in which a language was placed either entirely in the training set or in the test set. In this approach, the model is asked to predict, e.g., all Icelandic data points without having seen any Icelandic data points in its training set. This is, of course, a more difficult task since it has no within-language information (but can access language family and areal information). The model got 87% of data points correct in that case, which is worse than the classifier that treats all data points as independent. This suggests that having information about some part of the concord paradigm of a language improves classification in other parts of the paradigm. But performance is still substantially better than a simple majority classifier, suggesting that language data (e.g., presence of a gender system) and the family-specific and area-specific coefficients contribute to model performance.

Having assessed overall model fit, we can now turn to more specific research questions. To assess the overall likelihood of concord within a language, we simulate hypothetical languages with no language-specific, family-specific, or area-specific effects. We then consider how often that language has any type of concord at any locus. In Table 2, we show the probability of concord (of any type) by area. Concord is most likely in Australia and least likely in South America. In Table 3, we show the probability of concord by type and locus of concord.

Area	Probability of Any Concord
Australia	0.71
Africa	0.66
North America	0.63
Eurasia	0.58
Papunesia	0.51
South America	0.49

Table 3. Model probability (for an “average” language) of having any type of concord, by world area. Relative to the raw data, the model estimates shrink extreme areas towards the global mean.

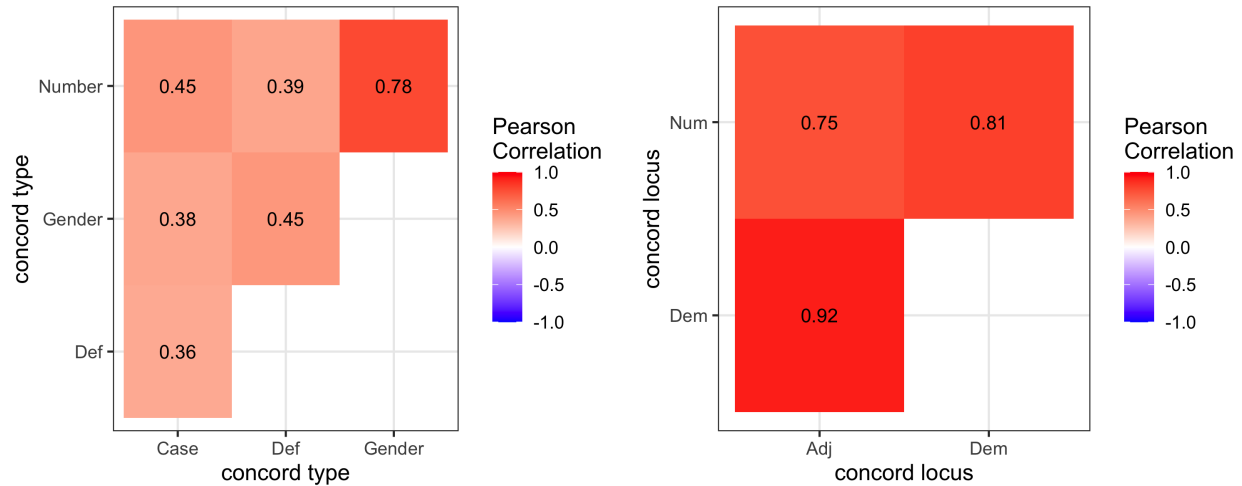
Agreement	Adj	Dem	Num
Case	0.14	0.15	0.10
Def	0.05	0.02	0.02
Gender	0.32	0.38	0.22
Number	0.47	0.50	0.17

Table 4. Model probability (for an “average” language) of having concord of the type specified. The most likely type of concord according to the model is number concord on demonstratives (.50 probability).

These results are broadly consistent with the findings in Norris (2019a). Among the fixed effect predictors, we found that languages with gender systems were significantly more likely to have concord than languages without ($\beta = -2.13$, 95% credible interval [-2.82, -1.48]). Having case and number were also both associated with having more concord, although the credible intervals on these coefficients included 0, which means that they are consistent with a null or small effect. Similarly, the credible interval for the noun-first interval included 0 ($\beta = .44$ with a 95% credible interval of [-0.04, 0.93]), suggesting that concord is more likely when the noun comes last. But, because the credible interval includes 0, we do not find strong support for a clear effect.

4.2. CORRELATIONS WITHIN A LANGUAGE. In addition to considering the likelihood of various types of concord, we want to know how different kinds of concord co-vary. Using the same method described in Section 4.1, we simulate languages and, across those languages, explore the covariance structure that emerges. That is, we ask how the presence of gender concord in a language correlates with the presence of other kinds of concord.

We found that, within a language, any kind of concord makes any other kind of concord more likely; in other words, concord begets concord. This is clear from the broadly positive correlation matrices in the figure to the right (cf. Norris 2019a). In Figure 1b, the correlation is highest between gender and number concord ($r=.78$), suggesting that the presence of gender concord makes number concord more likely and vice versa. Similarly, having concord on any modifier makes a language much more likely to have concord on other modifiers. That is, the strong correlation ($r=.92$) between adjective and demonstrative concord suggests that it is rare for a language to have adjective concord but not demonstrative concord: it typically has both or neither.



(a) Correlation between types of concord

(b) Correlation between locus of concord

Figure 1. Correlation matrices for concord type (a) and locus (b)

4.3. NORRIS’S CONCORD TENDENCIES. Broadly, our model provides robust quantitative evidence for Norris’s (2019a) Concord Tendencies, repeated below for convenience.

- (14) **Concord Tendency 1:** If a language has concord, it likely has number concord.
- (15) **Concord Tendency 2:** If a language has a grammatical gender system, it likely has gender concord.
- (16) **Concord Tendency 3:** If a language or lexical category within a language has case concord, it will likely have number concord (and if the language has gender, gender concord).
- (17) **Concord Tendency 4:** If a language has concord, it is more likely that it will have concord on both adjectives and demonstratives than on just one or the other.

First, we consider Tendency 1 (14). We simulate from our model and then condition on only simulated languages that have at least some kind of concord. If a language has concord of some kind, there is an 89.7% model probability it has number concord, 67.4% chance it has gender concord, 30.0% chance it has case concord, and a 10.7% chance it has definiteness concord. From this, we conclude, that if a language has concord it is likely to have gender concord and it is even more likely to have number concord.

Tendency 2 (15) connects gender concord and gender systems. Figure 2 shows the model probability of concord of each type as a function of whether the language has gender. It is clear that the presence of a gender system makes gender concord extremely likely, as Norris (2019a) observes. More strongly, it is also evident that the presence of a gender system makes every kind of concord more likely.

To assess Tendency 3 (16), we again simulated languages and this time include only those languages which have case concord at some locus. We then look at the probability of having other kinds of concord. We plot the probabilities in Figure 3. The figure shows that, indeed, if a language has case concord, it is more likely to have number, gender, and definiteness con-

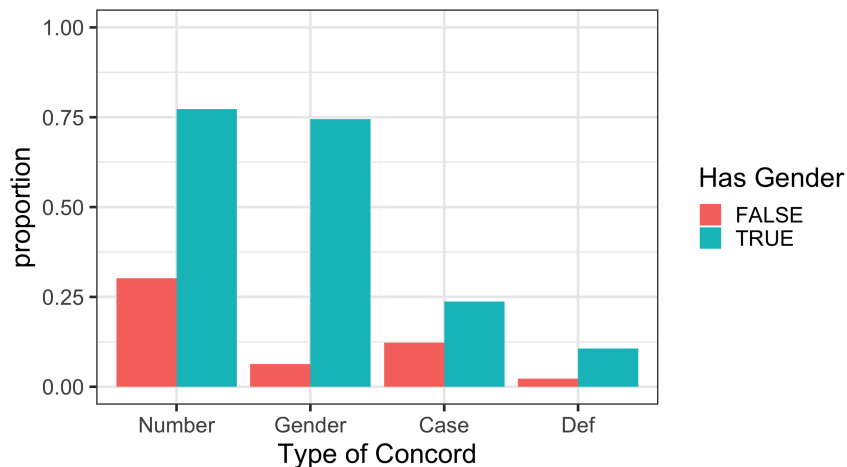


Figure 2. Probability of the given type of concord, with and without a gender system.

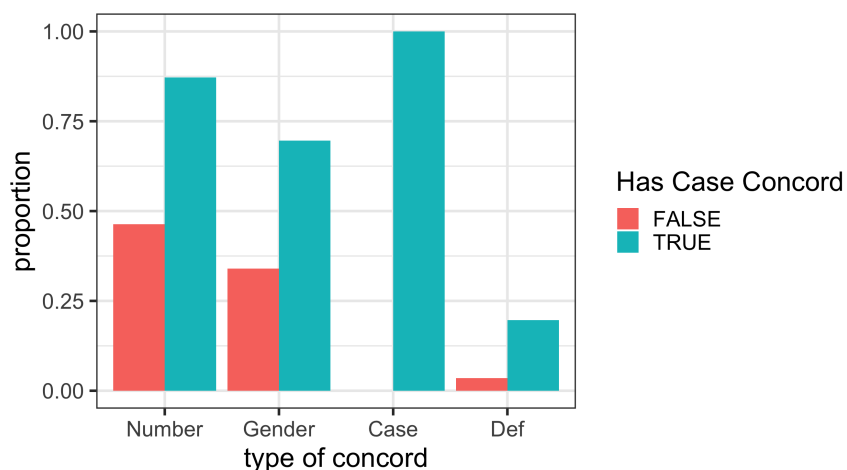


Figure 3. Probability of the given type of concord, conditioned on having case concord.

cord as well.

Finally, we investigate Tendency 4 (17), which concerns concord loci. We again simulate languages, conditioning on only those that have some kind of concord. For simulated languages which have some kind of concord, 73.5% have it on both the demonstrative and the adjective. 15.8% have it on only the demonstrative, 8.9% on only the adjective and only 1.7% on neither.

4.4. CONCORD AND NP-INTERNAL WORD ORDER. As we mentioned before, we were inspired by Greenberg’s (1963) Universal 40 to consider the relationship between concord and NP-internal word order. Specifically, we used our method of simulating languages based on the model, this time varying whether the noun came before or after its modifier, in order to ask whether we were more likely to find concord on a particular modifier if that modifier is postnominal. We found a slightly increased likelihood of concord when the for postnominal modifiers compared to prenominal, as shown in Table 5. Because the effect is small (and be-

cause, as discussed above, the noun-first parameter was not significantly different from 0 in our model), it seems like the probability of concord is largely insensitive to whether the noun comes before or after the modifier.

Locus	noun first	noun last
Adj	0.77	0.83
Dem	0.83	0.88
Num	0.50	0.59

Table 5. Probability of having concord based on modifier, as a function of whether the noun comes before or after the given modifier.

5. Conclusion.

5.1. OVERALL SUMMARY. Using a typologically diverse data set (Norris 2020), we fit a hierarchical Bayesian model to explore the likelihood of concord given various constraints. We found that languages with nominal concord tend to have concord in more than one place and of more than one type. That is, a language with case concord is also more likely to have gender concord. A language with gender concord is more likely to have number concord. A language with concord on the demonstrative is also more likely to have concord on the adjective and numeral. In other words, concord begets concord. While there were some areal and family-specific effects, these results seem robust across these sub-categories.

We also found support for the four concord tendencies proposed by Norris (2019a). We also uncovered additional generalizations. For example, a language with a gender system is likely to have gender concord (as Norris claimed), but it is also likely to have number concord (and more likely to have other forms of concord than languages without gender).

We then turned to the connection between concord and NP-internal word order. Perhaps surprisingly, we did not find strong evidence that certain word orders make nominal concord more or less likely. While this does not directly address the possible connection between word order freedom and concord, it does suggest that concord may not have much to do with word order at all.

5.2. METHODOLOGICAL CONCLUSION. There is a growing body of work seeking to use computational methods to infer how typological features covary (Bjerva et al. 2020, 2019; Daumé III & Campbell 2009). In a review of where typology is going in the twenty-first century, Bickel (2007) suggests that this kind of large-scale typological work could be fruitfully combined with the study of more fine-grained typological variation. We concur and, in particular, think there is much to be gained by the application of statistical methodology of the sort used here to specific linguistic features like concord.

In particular, fitting a Bayesian model allows us to, in part, explore the effects of languages and language families. And by simulating from the model, we are able to ask arbitrary questions about the statistical relationships present in our data. See Vasissth et al. (2018) and Schad et al. (2020) for further discussion of how to fit and use Bayesian models in linguistics and cognitive science.

Of course, any model is only as good as the data that goes into it. While our model can, in part, control for family and areal effects, we are still limited by the fact that even a typolog-

ically diverse sample covers only a small subset of human language. And it is impossible to rule out the fact that trends could be driven by non-independence of data (e.g., language borrowings and relatedness across different world regions).

Finally, we note that this work was originally made possible because the data in Norris (2019a,b, 2020) was publicly shared and available. Open, reproducible linguistic data (Forkel et al. 2018; Berez-Kroeker et al. 2018) is crucial for quantitative typological work. We strongly encourage typologists to make their data sets available to others, and we encourage researchers with strong statistical and computational training to build collaborations with typologists.

5.3. FUTURE DIRECTIONS. Building on the results discussed here, we plan to continue to explore the connection of concord to other linguistic properties. First, since concord is largely contained within noun phrases, we will investigate the extent to which concord correlates with aspects of clausal syntax (e.g., order of object and verb). Second, given that concord introduces morphological complexity, we are curious to know to what extent languages might reduce complexity in other ways. To wit, preliminary investigations suggest a higher incidence of cumulative exponence for case and number in languages that have case and number concord than for languages at large (on this property, see Bickel & Nichols 2013). On a broader level, future work will build on this model to understand the functional role of nominal concord in language systems, how it evolves, and how it co-evolves with other typological features.

References

- Berez-Kroeker, Andrea L., Lauren Gawne, Susan Smythe Kung, Barbara F. Kelly, Tyler Heston, Gary Holton, Peter Pulsifer, David I. Beaver, Shobhana Chelliah, Stanley Dubinsky, et al. 2018. Reproducible research in linguistics: A position statement on data citation and attribution in our field. *Linguistics* 56(1). 1–18. <https://doi.org/10.1515/ling-2017-0032>.
- Bickel, Balthasar. 2007. Typology in the 21st century: Major current developments. *Linguistic Typology* 11(1). 239–251. <https://doi.org/10.1515/LINGTY.2007.018>.
- Bickel, Balthasar & Johanna Nichols. 2013. Exponence of selected inflectional formatives. In Matthew S. Dryer & Martin Haspelmath (eds.), *The world atlas of language structures online*. Leipzig: Max Planck Institute for Evolutionary Anthropology. <https://wals.info/chapter/21>.
- Bjerva, Johannes, Yova Kementchedjhieva, Ryan Cotterell & Isabelle Augenstein. 2019. A probabilistic generative model of linguistic typology. *arXiv preprint arXiv:1903.10950*.
- Bjerva, Johannes, Elizabeth Salesky, Sabrina J Mielke, Aditi Chaudhary, Giuseppe GA Celano, Edoardo M Ponti, Ekaterina Vylomova, Ryan Cotterell & Isabelle Augenstein. 2020. Sig-tyt 2020 Shared task: Prediction of typological features. *arXiv preprint arXiv:2010.08246*.
- Carpenter, Bob, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus A. Brubaker, Jiqiang Guo, Peter Li & Allen Riddell. 2017. Stan: a probabilistic programming language. *Journal of Statistical Software* 76(1). 1–32. <https://doi.org/10.18637/jss.v076.i01>.
- Clem, Emily & Virginia Dawson. 2018. The emergence of case matching in discontinuous DPs. Manuscript, University of California, San Diego and Western Washington University.
- Corbett, Greville G. 2013. Number of genders. In Matthew S. Dryer & Martin Haspelmath (eds.), *The world atlas of language structures online*. Leipzig: Max Planck Institute for Evolutionary Anthropology. <https://wals.info/chapter/30>.

- Daumé III, Hal & Lyle Campbell. 2009. A Bayesian model for discovering typological implications. *arXiv preprint arXiv:0907.0785*.
- Deal, Amy Rose. 2010. *Topics in the Nez Perce verb*. Amherst, MA: University of Massachusetts dissertation.
- Diessel, Holger. 2013. Pronominal and adnominal demonstratives. In Matthew S. Dryer & Martin Haspelmath (eds.), *The world atlas of language structures online*, Leipzig: Max Planck Institute for Evolutionary Anthropology. <https://wals.info/chapter/42>.
- Dryer, Matthew S. 2013a. Coding of nominal plurality. In Matthew S. Dryer & Martin Haspelmath (eds.), *The world atlas of language structures online*, Leipzig: Max Planck Institute for Evolutionary Anthropology. <https://wals.info/chapter/33>.
- Dryer, Matthew S. 2013b. Order of adjective and noun. In Matthew S. Dryer & Martin Haspelmath (eds.), *The world atlas of language structures online*, Leipzig: Max Planck Institute for Evolutionary Anthropology. <https://wals.info/chapter/87>.
- Dryer, Matthew S. 2013c. Order of demonstrative and noun. In Matthew S. Dryer & Martin Haspelmath (eds.), *The world atlas of language structures online*, Leipzig: Max Planck Institute for Evolutionary Anthropology. <https://wals.info/chapter/88>.
- Dryer, Matthew S. 2013d. Order of numeral and noun. In Matthew S. Dryer & Martin Haspelmath (eds.), *The world atlas of language structures online*, Leipzig: Max Planck Institute for Evolutionary Anthropology. <https://wals.info/chapter/89>.
- Dryer, Matthew S. & Martin Haspelmath (eds.). 2011. *The world atlas of language structures online*. Munich: Max Planck Digital Library 2011 edn. <http://wals.info/>.
- Dryer, Matthew S. & Martin Haspelmath. 2013. *The world atlas of language structures online*.
- Forkel, Robert, Johann-Mattis List, Simon J. Greenhill, Christoph Rzymiski, Sebastian Bank, Michael Cysouw, Harald Hammarström, Martin Haspelmath, Gereon A Kaiping & Russell D. Gray. 2018. Cross-linguistic data formats, advancing data sharing and re-use in comparative linguistics. *Scientific Data* 5(1). 1–10. <https://doi.org/10.1038/sdata.2018.205>.
- Futrell, Richard, Kyle Mahowald & Edward Gibson. 2015. Quantifying word order freedom in dependency corpora. *Proceedings of the Third International Conference on Dependency Linguistics (Depling 2015)* 91–100.
- Gelman, Andrew, John Carlin, Hal Stern, David Dunson, Aki Vehtari & Donald Rubin. 2004. *Bayesian data analysis*. Boca Raton, FL: CRC press.
- Gelman, Andrew & Jennifer Hill. 2006. *Data analysis using regression and multilevel/hierarchical models*. Cambridge, UK: Cambridge University Press.
- Greenberg, Joseph H. 1963. Some universals of grammar with particular reference to the order of meaningful elements. In Joseph H. Greenberg (ed.), *Universals of language*, 73–113. Cambridge, MA: MIT Press.
- Heath, Jeffrey. 1999. *A grammar of Koyraboro (Koroboro) Senni: The Songhay of Gao, Mali* (vol. 19 Westafrikanische Studien). Cologne: Rüdiger Köppe.
- Iggesen, Oliver A. 2013. Number of cases. In Matthew S. Dryer & Martin Haspelmath (eds.), *The world atlas of language structures online*, Leipzig: Max Planck Institute for Evolutionary Anthropology. <https://wals.info/chapter/49>.
- Kiparsky, Paul. 1997. The rise of positional licensing. In Ans van Kemenade & Nigel Vincent (eds.), *Parameters of morphosyntactic change*, 460–494. Cambridge, UK: Cambridge University Press.
- Kutsch Lojenga, Constance. 1994. *Ngiti: A Central-Sudanic language of Zaire*. Cologne: Rüdiger Köppe.
- Linn, Mary Sarah. 2001. *A grammar of Euchee (Yuchi)*. Lawrence, KA: University of Kansas dissertation.

- Mace, John. 2003. *Persian grammar: For reference and revision*. London: Routledge.
- McFadden, Thomas. 2003. On morphological case and word-order freedom. *Annual meeting of the Berkeley Linguistics Society (BLS)* 29(1). 295–306. <https://doi.org/10.3765/bls.v29i1.997>.
- Norris, Mark. 2012. Towards an analysis of concord (in Icelandic). In Jaehoon Choi, E. Alan Hogue, Jeffrey Punske, Deniz Tat, Jessamyn Schertz & Alex Trueman (eds.), *Proceedings of the 29th West Coast Conference on Formal Linguistics*, 205–213. Somerville, MA: Cascadilla Proceedings Project.
- Norris, Mark. 2014. *A theory of nominal concord*. Santa Cruz, CA: University of California dissertation.
- Norris, Mark. 2017. Description and analyses of nominal concord (parts I–II). *Language and Linguistics Compass* 11(11). <https://doi.org/10.1111/lnc3.12267>.
- Norris, Mark. 2018. The morphosyntax of number in Estonian numeral-noun constructions. *Proceedings of the Linguistic Society of America* 3(45). 1–13. <https://doi.org/10.3765/plsa.v3i1.4337>.
- Norris, Mark. 2019a. A typological perspective on nominal concord. *Proceedings of the Linguistic Society of America* 4(12). 1–15. <https://doi.org/10.3765/plsa.v4i1.4515>.
- Norris, Mark. 2019b. A typological perspective on nominal concord (archived data). <https://shareok.org/handle/11244/320354>.
- Norris, Mark. 2020. Typology of nominal concord (archived data). <https://doi.org/10.17605/OSF.IO/TM49Q>.
- Pfaff, Alexander. 2017. Adjectival inflection as diagnostic for structural position: inside and outside the Icelandic definiteness domain. *The Journal of Comparative Germanic Linguistics* 20(3). 283–322. <http://www.jstor.org/stable/44651878>.
- Schad, Daniel J., Michael Betancourt & Shravan Vasishth. 2020. Toward a principled Bayesian workflow in cognitive science. *Psychological Methods* 26(1). 103–126. <https://doi.org/10.1037/met0000275>.
- Vasishth, Shravan, Bruno Nicenboim, Mary E. Beckman, Fangfang Li & Eun Jong Kong. 2018. Bayesian data analysis in the phonetic sciences: A tutorial introduction. *Journal of Phonetics* 71. 147–161. <https://doi.org/10.1016/j.wocn.2018.07.008>.