

Exploring horizontal homophony in pronominal paradigms: A case study where cross-linguistic regularities defy individual learning biases

Carmen Saldana (carmencatalina.saldana@upf.edu)

Department of Translation and Language Sciences, Universitat Pompeu Fabra, Spain

Department of Comparative Language Science, University of Zurich, Switzerland

Center for the Interdisciplinary Study of Language Evolution, University of Zurich, Switzerland

Mora Maldonado (mora.maldonado@univ-nantes.fr)

UMR 6310 Laboratoire de Linguistique de Nantes, CNRS/Nantes Université, France

Abstract

Homophony (i.e., multiple meanings expressed by the same form) is ubiquitous across the world's languages. Despite its pervasiveness, not all instances of homophony are equally likely, which suggests that homophony is unlikely to be accidental. There is a growing body of literature which aims to thoroughly examine cross-linguistic regularities in patterns of homophony and explain these from constraints in language learning and use, both at the lexical and morphosyntactic levels. Here, we examine a specific case of homophony in pronominal paradigms, that is, the lack of a number distinction (singular vs plural) for a given person value (first, second and third), a phenomenon coined as horizontal homophony. Cysouw (2003) suggested that a lack of number distinction is more likely to be found in third person (i.e., 3SG=3PL) than in second (i.e., 2SG=2PL), and it is least frequently found in first person (i.e., 1SG=1PL). We refer to this generalisation as the Horizontal Homophony Hierarchy: $3 > 2 > 1$ (where $>$ represents frequency inequality). This generalisation was nevertheless only made via qualitative description and by raw counts, and merely described without motivated explanation. In this study we take a step back and present additional evidence supporting the Horizontal Homophony Hierarchy. First, we ascertain the robustness of this typological tendency through a statistical analysis using the largest cross-linguistic database of pronominal paradigms to date (926 languages from 229 different families). Next, we explore whether the Horizontal Homophony Hierarchy has a corresponding learning correlate, which would indicate that this asymmetry is at least partly rooted in a cognitive bias. Specifically, we examine asymmetries in how easily adult humans learn different types of horizontal homophony in an artificial language learning experiment. The results from our typological analysis corroborate a hierarchy of horizontal homophony $3 > 2 > 1$ in the world's languages. However, our experimental results provide evidence against a learning bias underlying the hierarchy, thus suggesting that motivated explanations of the typology (if any) are more likely to be found in alternative pressures such as communicative need and efficiency.

Keywords: horizontal homophony; person; number; personal pronouns; morphology; semantics; quantitative typology; artificial language learning; cognitive biases; language evolution

Introduction

Homophony (i.e., multiple meanings expressed by the same form) is ubiquitous across the world's languages. We find it across lexemes (e.g., 'draw' in English can refer to the action of pulling or to that of depicting with lines onto a surface) as well as across different sets of morphosyntactic values (e.g., English 'you' can refer to a second person singular or plural). Despite its pervasiveness, not all instances of homophony are equally likely, which suggests that homophony

is unlikely to be accidental. There is a growing body of literature which aims to thoroughly examine cross-linguistic regularities in patterns of homophony and explain these from constraints in language learning and use, both at the lexical (F. Xu & Tenenbaum, 2007; Brochhagen et al., 2023; Y. Xu et al., 2020; F. Xu & Tenenbaum, 2007; Dautriche et al., 2016; Dautriche & Chemla, 2016; Carr et al., 2020; Landau & Shipley, 2001; Pothos et al., 2004; Maldonado & Culbertson, 2021) and morphosyntactic levels (Johnson et al., 2021; Pertsova, 2014; Nevins, 2015; Nevins et al., 2015; Saldana et al., 2022; Maldonado & Culbertson, 2021).

Drawing inspiration from a longstanding tradition in linguistics (e.g., Jakobson, 1936; Noyer, 1992; Baerman et al., 2005; Gardenfors, 2004; Quine, 1960), a big portion of this work proposes a bias favouring formal identity (i.e., homophony) in expressing similar meanings. In this view, forms extending over contexts with a common meaning and function are more often phonologically identical than those extending over contexts with no meaning in common, which ultimately is taken to show a preference for *similarity-based* structure in the lexicon and morphology (Saldana et al., 2022; Herce et al., 2023; Dautriche et al., 2017; Brochhagen et al., 2023; Maldonado & Culbertson, 2021; Mansfield et al., 2022). In morphology, semantic similarity of this sort is typically cashed out by positing shared morphosyntactic values (Jakobson, 1936; Bierwisch, 1967; Noyer, 1992; Müller, 2004; Baerman et al., 2005; Corbett, 2012). For example, the pronouns 'we' and 'they' are taken to share a PLURAL value in virtue of expressing plurality. In contrast, the pronouns 'we' and 'he' have little meaning in common and thus are not thought to share any morphological value.

Interestingly, the cross-linguistic tendency towards similarity-based homophony is often thought to be rooted in individuals' constraints in language learning. In a number of laboratory studies, it has been shown that adults find it easier to learn patterns of homophony grounded on semantic similarity (Maldonado & Culbertson, 2021; Johnson et al., 2021; Pertsova, 2014, 2011; Nevins, 2015; Nevins et al., 2015; Saldana et al., 2022).¹ These studies thus support the idea (sometimes referred to as the *Typological Prevalence Hypothesis*; Gentner & Bowerman, 2009) that

¹The cross-linguistic recurrence and learnability of some morphological patterns are actually proportional to the degree of semantic similarity, as shown in Saldana et al. (2022)

	Koasati (Muskogean)		Ngkolmpu (Yam)		Balinese (Austronesian)	
	SG	PL	SG	PL	SG	PL
1	anók	kosnók	ŋko	ni	hicaŋ	
2	isnók	hasnók	mpu		cahi/ñahi	
3	ibisnók		pi		hiya	

Table 1: Different types of patterns of horizontal homophony in pronominal paradigms. We show the pronominal paradigms for S arguments (i.e., sole argument in intransitive clauses) with the phonological forms. All paradigms are extracted from the Parabank Pronoun Database (Greenhill, 2022) and references therein—i.e., from Kimball (1991), Carroll (2016) and Barber (1977) for Koasati, Ngkolmpu and Balinese respectively.

cross-linguistically more frequent patterns tend to be those which are more natural and thus easier to learn for humans. The underlying assumption here is that learning biases such as the aforementioned similarity-based bias are replicated in linguistic transmission, ultimately shaping cross-linguistic distributions over time (Bickel, 2015; Smith, 2018; Reali & Griffiths, 2009; Blythe & Croft, 2021; Culbertson & Smolensky, 2012).

Crucially, not all typological regularities regarding homophony patterns can be as easily explained by a similarity-based bias. We focus on one of such cases: homophony across singular and plural exponents of personal pronouns. Personal pronouns serve to refer to individuals as defined by the communicative context: first person pronouns (1) indicate that the reference includes the speaker, second person pronouns (2) indicate that the reference includes the addressee (but excludes the speaker), and third person pronouns (3) refer to individuals that are neither speaker nor addressee (Heim, 2008). Here, we examine the lack of number distinction (singular or plural) for a given person value (first, second and third), a phenomenon coined as *horizontal homophony* by Cysouw (2003)² (see Table 1 for examples of horizontal homophony in pronominal systems). Over two decades ago, Cysouw (2003) highlighted an interesting cross-linguistic generalisation regarding horizontal homophony: A lack of number distinction is more likely to be found in 3 (i.e., 3SG=3PL) than in 2 (i.e., 2SG=2PL), and it is least frequently found in 1 (i.e., 1SG=1PL). We refer to this generalisation as the *Horizontal Homophony Hierarchy*: $3 > 2 > 1$ (where $>$ represents frequency inequality).

The Horizontal Homophony Hierarchy cannot be easily captured by a similarity-based bias. All things being equal, there is no reason to expect differences in number distinctions within a given person category, as number values (i.e., SG and PL) and their assumed meanings are consistent across persons: SG always indicates an atomic referent, while PL indicates a non-atomic referent (Heim, 2008). One then wonders if there is some additional property that makes 3SG and

²Because in person-number paradigms person is usually represented in rows and number in columns, and thus any homophony across singular and plural would be depicted horizontally across columns

3PL more likely to be collapsed as a category than 2SG and 2PL or 1SG and 1PL and could explain such typological asymmetry. There are different avenues for exploration, such as the observation that plurality is more likely to be interpreted as associative (i.e., as referring to a group including the focal referent plus associates) in 1PL and 2PL than in 3PL (which is more likely to be interpreted additively, i.e., referring to a homogenous group of atoms of the same non-participant role; Maldonado & Saldana, 2022). Another aspect worth considering is that expressions for 1PL, 2PL, and 3PL might not be equally distributed or necessary in real communication. However, neither of these avenues has been fully developed to clearly account for the Horizontal Homophony Hierarchy.

In this paper, our goal is to take a step back and present additional evidence supporting the Horizontal Homophony Hierarchy. We offer two sources of evidence. First, we ascertain the robustness of this typological tendency through a statistical analysis using the largest cross-linguistic database of pronominal paradigms to date (926 languages from 229 different families). Next, we explore whether the Horizontal Homophony Hierarchy has a corresponding learning correlate, which would indicate that this asymmetry is at least partly rooted in a cognitive bias. Specifically, we examine whether adult learners treat different patterns of horizontal homophony differently in an artificial language learning experiment. This approach controls for communicative need and distributional information, enabling us to directly evaluate the existence of learning biases.

The typology of horizontal homophony

We survey the instances of horizontal homophony for each person category in pronominal paradigms across the three largest openly-available databases: the Parabank Pronoun Database (with a focus on Austronesian, Greenhill, 2022), and ATLAS (Areal Typology of the Languages of the Americas, Chousou-Polydouri et al., 2023), and PRONOM (Pertsova, 2022). For consistency across databases, we only analyse the paradigms of personal pronouns gathered for S arguments (i.e., the sole argument of intransitive clauses, which will align with subjects of transitive clauses in nominative-accusative systems or with objects of transitive clauses in ergative-absolutive systems).³ Note that we only take into account PERSON:1,2,3 and NUMBER:SG,PL values, and we disregard whether or not systems contain a clusivity distinction (i.e., we always take the default 1PL.EXCL as 1PL across systems).

By pooling these databases together we gather pronominal paradigms of 926 languages from 229 different families, of which 67 are isolates. Only 10.5% of the sampled languages ($N = 96$) contain horizontal homophony. Fig. 1 shows the geographical distribution of these languages, coloured by the type of horizontal homophony they manifest (i.e., the per-

³Within the PRONOM database (Pertsova, 2022), we only survey those paradigms classified as personal pronouns without variation across alignment or those in nominative case.

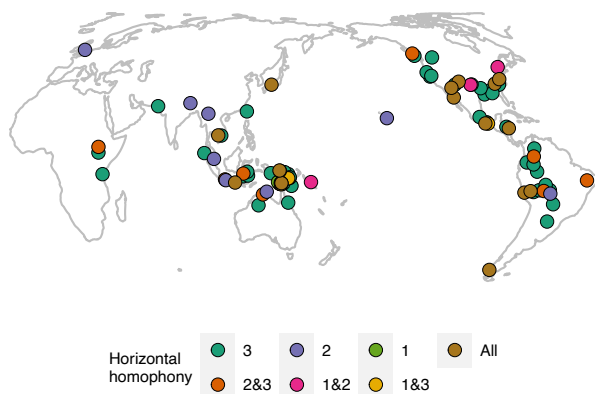


Figure 1: Geographical distribution of the languages in our sample containing horizontal homophony. Languages are colour coded according to the person categories with horizontal homophony found in the pronominal paradigm: only in third person (3), only in second person (2), only in first person (1), in both second and third person (2&3), in both first and second (1&2), in both first and third (1&3), or across all person categories.

TYPE							
COUNT	45 (47%)	21 (22%)	15 (16%)	8 (8%)	3 (3%)	2 (2%)	2 (2%)

Table 2: Frequency of each type of horizontal homophony pattern in S pronominal paradigms. Pattern types are ordered by frequency, in descending order from left to right.

son values with homophony across SG and PL values). Table 2 summarises the frequency of occurrence of the different types of horizontal homophony patterns in descending order.⁴ We observe that paradigms where only third person is homophonous constitute almost half of the instances of paradigms with horizontal homophony. The two patterns following up in frequency of occurrence are those with horizontal homophony across all person categories or only in second and third. These three pattern types (i.e., , , and) account for 83% of the paradigms with horizontal homophony. Therefore, horizontal homophony in third person is most frequent than in second as it occurs across all three types, and horizontal homophony in 1st is the least frequent as it only occurs in one of the three, and not the most frequent.

In order to test the prevalence of the Horizontal Homophony Hierarchy across the world’s languages we need to move beyond simple counts and control for linguistic relatedness. We run a Bayesian binomial regression model predicting the presence or absence of horizontal homophony in pronominal paradigms by person category. Our dependent variable is the presence or absence of horizontal homophony for a given person category in a pronominal paradigm. As fixed effects, we include person category (categorical variable with three levels: 2 as baseline, and 3 and 1). As random effects, we include intercepts for language and fam-

⁴Note that here and in our model, we exclude Ngiti in Parabank when counting systems with horizontal homophony because it has several S paradigms with different recorded instances of horizontal homophony.

ily, and a by-family random slope for the effect of person. Further details on all analyses reported here can be found in the analysis script available in the OSF repository. The model’s results suggest that among paradigms with horizontal homophony, homophony in the third person is more likely than in second person ($\hat{\beta} = 2.875$, $90\%CI = [1.665, 4.447]$, $P(\hat{\beta} > 0) = 1$), and it is less likely in first person than in second ($\hat{\beta} = -1.213$, $90\%CI = [-1.929, -0.532]$, $P(\hat{\beta} < 0) = 0.999$). These results are in line with a cross-linguistic Horizontal Homophony Hierarchy $3 > 2 > 1$, which until now had only been described qualitatively and by raw frequency counts (Cysouw, 2003).

Learnability of horizontal homophony patterns

In order to establish a connection between the Horizontal Homophony Hierarchy, attested cross-linguistically, and human cognition, we assess whether typological frequency correlates with learnability. We test whether the ease of learning aligns with cross-linguistic frequency asymmetries – that is, whether learners find it easier to learn homophony patterns that are more common than those that are less common.

We use an ease-of-learning artificial language learning paradigm where we train and test participants on one of a set of novel pronominal systems with varying patterns of horizontal homophony. We then compare how fast and accurately they learn it over time. Assuming a correlate between cross-linguistic frequency and learnability, we predict that amongst systems of a single instance of horizontal homophony, homophony only in the third person should be more learnable than horizontal homophony in second or first person only: > > . Within systems with two instances of horizontal homophony, we predict that systems where only first person contains a number distinction are most learnable, that is, more than systems where number distinctions are only found in third or only in second: > / . We focus only on comparing the learnability across systems with the same number of instances of horizontal homophony because any other difference could be driven by the varying amount of lexical items to be learned (i.e., less in paradigms with two instances of horizontal homophony than in those with one).

Materials and methods

Experimental Conditions We run six experimental conditions, each consisting of a 3×2 person-number pronominal system involving a specific horizontal homophony pattern. Fig. 2 illustrates the patterns of homophony for each experimental condition. All pronominal forms are non-compositional, that is, they do not have separative but cumulative person-number morphology. From the six conditions, three contain horizontal homophony for a single person category (upper-panel Fig. 2), and three, for two person categories (lower-panel Fig. 2).

Artificial Lexicon The lexicon is composed of four verbal forms, and of four to five pronominal forms depending on the condition (see Fig. 2). Verbal forms are taken from the fol-

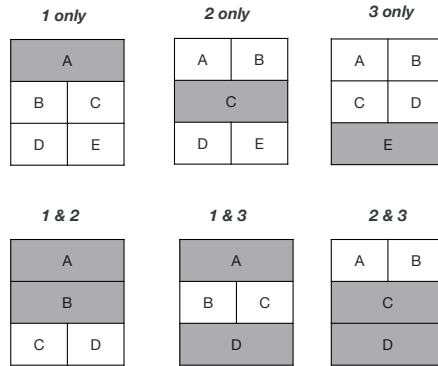


Figure 2: Experimental conditions, each containing a different pattern type of horizontal homophony. Each capital letter represents a different pronominal form. Grey cells indicate homophony.

lowing array {*solate, daduca, toloholu, fanepare*}, and are randomly assigned to actions (i.e., cooking, driving, running and reading) for each participant. The pronominal forms are taken from the array {*kurai, kusui, kupei*} for singular forms, and from {*kunaki, kufuli, kutegi*} for plural forms. When horizontal homophony is in place for a given person category, the singular form is used across singular and plural referents. In order to ease learning, when a number distinction is present in a given person category, the singular and plural forms share vowels, that is, the singular *kurai* is paired with *kunaki*, *kusui* is paired with *kufuli*, and *kupei* is paired with *kutegi*. Singular forms are randomly assigned to person categories at the beginning of the experiment.

Procedure There are two phases in the experiment: a first verbal vocabulary training phase, and a second critical pronominal testing phase. In the verbal vocabulary phase, participants see images of four different actions and the corresponding verbal (infinitival) form in the artificial languages. After 12 exposure trials (three trials per action), they are tested on these verbal forms. Over 12 trials (three trials per action), they see the action depicted and they have to choose among two options the verbal form that correctly describes it.

In the critical pronominal phase, participants are taught the personal pronominal system of the artificial language through feedback learning. At each trial, participants are shown the communicative context and a sentence (composed by a subject pronoun and a verbal form from the training). They are asked to select the referents of the pronoun contained in the phrase among two pictures, a target and a foil, differing on whether they depict the referent conveyed by the relevant person value. Participants receive full feedback after their selection, which allows them to learn the form-meaning mappings in the novel language as they go. They go over 2 blocks of 24 trials each, and 4 sub-blocks of 6 trials each (containing the 6 different pronominal referents).

In order to convey the speech act roles, participants are told they would see themselves at a party (see Fig. 3), depicted as a cartoon character (i.e., red avatar in Fig. 3), chatting to some other character (i.e., white avatar in Fig. 3). The commu-




Figure 3: Example test trial (translated from Spanish into English). The participant is told that they utter a certain sentence in the novel language within a given communicative context. They are then required to select the image they think shows the referents of the predicate out of an array of two.

nicative context is presented as a GIF, where there are thus three focal characters: the participant's character which acts as a *focal speaker*, the character they are talking to, which acts as a *focal addressee* and a character further away from these two, apparently not participating in the interaction (*focal other*, i.e., yellow avatar in Fig. 3). These three stationary characters are joined by moving partners who stand next to them for a second and can then be interpreted as a non-focal additional addressees or others.

In each trial, participants are further provided with the sentence that they uttered to their addressee, which consists of a pronoun plus a verbal form in the artificial language. Sentences involve one of six possible personal pronouns, varying in their person-number values. Singular pronouns refer to a single individual, always one of the focal characters. Plural pronouns always refer to a group of three individuals. 1PL trials, the target referent is always an *heterogeneous* group composed by the unique speaker and two additional individuals (necessarily be non-speakers). In 2PL trials, the target referent always includes the focal addressee but can additionally involve characters present in the context and close to the speaker or addressee (which makes them potential addressees) or characters that are not present in the context but are associated to the addressee. In 3PL trials, the target referent exclusively involves *others* (characters which are neither speakers nor addressees), and these can be either present in the context and next to the focal other, or absent and associated to the focal other. Figure 4 shows the different images used for these plural trials within the block.

Plural referents with and without characters absent from the communicative context are provided (across targets and foils) with different probabilities depending on the person category: for first and second person referents, it is 50/50,






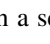



Figure 4: Illustration of images used for trials involving plural pronouns. These images correspond to the example communicative context in Fig. 3, where the focal speaker, addressee and other singular referents are the red, white and yellow avatars respectively.

and for third person referents, 75% of the time all characters are present in the context (cf. Figure 4). These probabilities are approximated from a previous norming experiment run with the same stimuli where Spanish-speaking participants selected their preferred depiction of plural pronouns in Spanish (Maldonado & Saldana, 2022). In that study, results show that the interpretation of plurality in Spanish pronouns as homogeneous-like groups follows the hierarchy $3 > 2 > 1$. These proportions are implemented across each block of 24 trials (4 trials per pronoun).




After they have completed these trials, they go over another couple of blocks of 24 each, with the same characteristics but with switched participant roles (and they are told so explicitly): the participant in the experiment is now the addressee. In this new context, pronouns for first and second person do not correspond to the same colour avatars and that allows us to ensure that participants correctly learn the indexicality of pronominal forms. The third person is also assigned another focal avatar with a different colour (i.e., instead of the yellow avatar as in the previous block, we assign the turquoise avatar as the the focal other).




Participants Participants (N = 240 in total, 40 per condition, 6 conditions) are adult Spanish speakers (located in Spain) recruited via Prolific for an experimental session lasting a median of 15 minutes. We exclude the data from participants who fail to provide at least 75% of correct responses in the vocabulary testing during the last block in the verb training phase. Participants are compensated with a base rate of £2.5 for their participation in the experiment, and can get up to £4.12 according to performance (£0.015 per correct trial).

Data analysis We run two different Bayesian binomial mixed-effects models. In the first model we analyse the three conditions with horizontal homophony across a single person category (i.e., ,  and ). In a second model, we analyse the data from the three conditions with horizontal homophony across two person categories (i.e., , , and ). Our dependent variable across the two models is whether or

not participants select the correct referent of the pronominal forms in each trial of the testing phase (coded as 1 if correct, 0 if incorrect). Fixed effects include condition and block (of 6 trials, including all pronominal forms) with an interaction term. Condition is a categorical variable with three levels (treatment coded, with second person as reference). Block is a centered continuous predictor, so the intercept is between blocks 7 and 8, and thus at the end of the trials with the first communicative context before the switch of participant roles. As random effects, we include intercepts for participant and by-participant slopes for the effect of block. Further details on all analyses reported here can be found in the analysis script available on the OSF repository.

Results

Following the preregistered hypotheses for systems with a single instance of horizontal homophony (see preregistration here), we first test whether learnability matches the observed typology and thus follows the hierarchy $3 > 2 > 1$. Fig. 5.1 shows participants' accuracy scores in 1 only () 2 only () and 3 only () conditions, along with the regression model's predicted mean accuracy scores. A visual inspection of the results suggests that, contrary to our predictions, paradigms with horizontal homophony in third person are the least learnable, and horizontal homophony in first and second is comparably learnable. Note that learnability here is defined by the overall accuracy and/or by the increase over time. The posterior distributions obtained from the Bayesian binomial regression model confirm this. We find that the accuracy at the intercept for paradigms with homophony in 3 is lower than of those with homophony in 2 ($\hat{\beta} = -0.426$, $90\%CI = [-0.853, -0.014]$, $P(\hat{\beta} > 0) = 0.954$), and the accuracy for paradigms with homophony in 1 is comparable to those with homophony in 2 ($\hat{\beta} = -0.166$, $90\%CI = [-0.604, 0.258]$, $P(\hat{\beta} < 0) = 0.744$). We also observe that accuracy increases by block as expected ($\hat{\beta} = 0.151$, $90\%CI = [0.116, 0.186]$, $P(\hat{\beta} > 0) = 1$) equally across conditions, although the learning rate seems to be slowest for paradigms with homophony in 1 (1 vs 2: $\hat{\beta} = -0.041$, $90\%CI = [-0.090, 0.008]$, $P(\hat{\beta} > 0) = 0.917$; 3 vs 2: $\hat{\beta} = -0.022$, $90\%CI = [-0.069, 0.026]$, $P(\hat{\beta} > 0) = 0.774$).

In conditions with two instances of horizontal homophony, we test whether systems with horizontal homophony in 2&3 are most learnable as suggested by the typology (see Table 2). A visual inspection of the results (see Fig. 5.2) suggests that this is not the case as paradigms with horizontal homophony in 1&2 () are the most learnable, and results from 1&3 () and 2&3 () are comparable. The posterior distributions obtained from the Bayesian binomial regression model confirm this. We find no difference in 1&3 vs 2&3 ($\hat{\beta} = -0.253$, $90\%CI = [-0.739, 0.245]$, $P(\hat{\beta} < 0) = 0.801$), but accuracy is higher for 1&2 ($\hat{\beta} = 0.561$, $90\%CI = [0.092, 1.065]$, $P(\hat{\beta} > 0) = 0.974$). We also observe that accuracy increases by block as expected ($\hat{\beta} = 0.145$, $90\%CI = [0.107, 0.184]$, $P(\hat{\beta} > 0) = 1$). This increase is




Figure 5: Experimental results of conditions with one instance of horizontal homophony (1) and with two instances (2). (A) Accuracy by testing block for each of the three conditions. Dots represent participants’ individual scores, and larger dots represent more individuals with the same accuracy score; thick lines represent the model’s predicted accuracy means conditioned on experimental condition and block. The shaded area shows the 90% credible intervals. (B) Overall accuracy by condition. Shaded dots represent participants’ individual scores; black dots represent the model’s predicted mean accuracy scores and the error bars represent the model’s predicted 90% credible intervals.

comparable across 2&3 and 1&3 conditions ($\hat{\beta} = -0.013$, $90\%CI = [-0.066, 0.039]$, $P(\hat{\beta} < 0) = 0.660$), but it is also higher for 1&2 ($\hat{\beta} = 0.058$, $90\%CI = [0.005, 0.111]$, $P(\hat{\beta} > 0) = 0.962$).

Discussion

In this study we first set out to bring new evidence supporting the cross-linguistic prevalence of the Horizontal Homophony Hierarchy proposed in Cysouw (2003). Prior to this study, the hierarchy had only been described qualitatively or supported by raw frequency counts, which were based on limited data and did not consider linguistic relatedness. Here, we extend the available typological evidence and attempt to find additional cognitive evidence that could explain this tendency. We first analysed the hitherto largest typological dataset of pronominal paradigms (926 language, of which 10.5% contain horizontal homophony), taking into account the variance across different person values within language families—and thus considering linguistic relatedness and ensuring that the results are not driven by individual families. Results corroborate a hierarchy of horizontal homophony: $3 > 2 > 1$ (where $>$ represents frequency inequality) in the world’s languages documented today. Whether or not this hierarchy can also be found diachronically (as we would expect with a robust typological bias) remains to be explored in future work using phylogenetic modelling techniques (Cathcart, 2018; Greenhill et al., 2020; Bickel, 2015; Jäger & Wahle, 2021).

In our experimental study, we assess whether this cross-linguistic pattern has a cognitive correlate during language learning which in turn could explain it, as observed for other typological patterns (see above; Gentner & Bowerman, 2009). We found that learners’ preferences do not mirror cross-linguistic regularities. Systems with horizontal homophony in third person are harder to acquire than those with homophony in first or second, and systems with homophony in first and second are more easily acquired than those with

second and third, or with first and third. If we assume that learnability reflects semantic similarity, these results suggest that it is actually the third person where singular and plural should be less similar to each other. This is not surprising if one considers the semantics of person values assumed here (see e.g., Heim, 2008). Singular and plural first person pronouns are both defined by the inclusion of the speaker, and likewise, and second person involves the addressee as main and focal referent in both singular and plural alternatives. In contrast, as long as the third person is defined by opposition (i.e., it means *neither* speaker nor addressee), it is not entirely clear what is common between singular and plural referents with the same person value.

These experimental results thus provide evidence against a learning bias underlying the Horizontal Homophony Hierarchy in typology, and thus suggest that motivated explanations (if any) are more likely to be found in alternative pressures such as communicative need and efficiency. For example, one could imagine that it might be more important in a conversation to distinguish among referents that involve the speaker or the addressee because these are speech-act participants, while it might be less problematic to not be able to distinguish between one or more other not present in the context (i.e., third person). Likewise, these person asymmetries could be accounted for by just positing a bias based on efficient information processing. Unlike the first and the second persons—which are true indexicals—the referent for the third person can typically be recovered from the linguistic context and stated as an unambiguous referential expression (e.g., a proper name or a definite description). As a result, number information about the third person may be easier to retrieve from the discourse than number information for the first or second persons, so it might be less inefficient to lose the overt expression of number in this case than for first or second pronouns. Evidence supporting these alternative hypotheses remains to be provided in future work.

Transparency and Openness All experimental materials, data and analysis reported here are available on OSF at osf.io/4fcu3/, and the preregistered design and analysis plan is accessible at osf.io/6xk2m. Experiments were coded using the *jsPsych* JavaScript library (De Leeuw, 2015). For data analysis, we used the *brms* package developed in R (Bürkner, 2017, 2018) as an interface to Stan (Carpenter et al., 2017).

Ethics The study was approved by the Ethics Committee of the School of Philosophy at the University of Zurich (Authorisation Nr. 22.6.11). Research practices follow the ethical guidelines for psychologists of the Swiss Society for Psychology (SGP), and the Ethical Principles of Psychologists and Code of Conduct of the American Psychological Association (APA).

Acknowledgements CS is currently supported by the Beatriu de Pinós programme funded by the Generalitat de Catalunya (agreement No. 2022 BP 00101) and the European Union’s Horizon 2020 research and innovation programme under the MSCA grant agreement No 801370. The work has also been partially funded by the NCCR Evolving Language, Swiss National Science Foundation (grant agreement Nr. 51NF40_180888).

References

- Baerman, M., Brown, D., & Corbett, G. G. (2005). *The syntax-morphology interface: A study of syncretism* (Vol. 109). Cambridge: Cambridge University Press.
- Barber, C. C. (1977). *A grammar of the Balinese language*. Aberdeen University Library.
- Bickel, B. (2015). Distributional typology: Statistical inquiries into the dynamics of linguistic diversity. In B. Heine & H. Narrog (Eds.), *The Oxford Handbook of Linguistic analysis, 2nd edition* (pp. 901–923). Oxford University Press.
- Bierwisch, M. (1967). Syntactic features in morphology: General problems of so-called pronominal inflection in German. In *To Honour Roman Jakobson: Essays on the occasion of his seventieth birthday* (pp. 239–270). De Gruyter Mouton.
- Blythe, R. A., & Croft, W. (2021). How individuals change language. *PLoS ONE*, *16*(6), e0252582. doi: 10.1371/journal.pone.0252582
- Brochhagen, T., Boleda, G., Gualdoni, E., & Xu, Y. (2023). From language development to language evolution: A unified view of human lexical creativity. *Science*, *381*(6656), 431–436. doi: 10.1126/science.ade7981
- Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, *80*(1), 1–28. doi: 10.18637/jss.v080.i01
- Bürkner, P.-C. (2018). Advanced Bayesian multilevel modeling with the R package brms. *The R Journal*, *10*(1), 395–411. doi: 10.32614/RJ-2018-017
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, *76*(1).
- Carr, J. W., Smith, K., Culbertson, J., & Kirby, S. (2020). Simplicity and informativeness in semantic category systems. *Cognition*, *202*, 104289. doi: 10.1016/j.cognition.2020.104289
- Carroll, M. J. (2016). *The Ngkolmpu Language with special reference to distributed exponence*. Unpublished doctoral dissertation, Australian National University.
- Cathcart, C. A. (2018). Modeling linguistic evolution: A look under the hood. *Linguistics Vanguard*, *4*(1). doi: 10.1515/lingvan-2017-0043
- Chousou-Polydouri, N., Inman, D., Huber, T. C., & Bickel, B. (2023). Multi-variate coding for possession: methodology and preliminary results. *Linguistics*, *61*(6), 1365–1402. doi: 10.1515/ling-2022-0021
- Corbett, G. G. (2012). *Features*. Cambridge University Press.
- Culbertson, J., & Smolensky, P. (2012). A Bayesian model of biases in artificial language learning: The case of a word-order universal. *Cognitive Science*, *36*(8), 1468–1498. doi: 10.1111/j.1551-6709.2012.01264.x
- Cysouw, M. (2003). *The paradigmatic structure of person marking*. Oxford: Oxford University Press.
- Dautriche, I., & Chemla, E. (2016). What homophones say about words. *PLoS one*, *11*(9), e0162176. doi: 10.1371/journal.pone.0162176
- Dautriche, I., Chemla, E., & Christophe, A. (2016). Word learning: Homophony and the distribution of learning exemplars. *Language Learning and Development*, *12*(3), 231–251. doi: 10.1080/15475441.2015.1127163
- Dautriche, I., Mahowald, K., Gibson, E., & Piantadosi, S. T. (2017). Wordform similarity increases with semantic similarity: An analysis of 100 languages. *Cognitive science*, *41*(8), 2149–2169. doi: 10.1111/cogs.12453
- De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior research methods*, *47*(1), 1–12. doi: 10.3758/s13428-014-0458-y
- Gardenfors, P. (2004). Conceptual spaces as a framework for knowledge representation. *Mind and Matter*, *2*(2), 9–27.
- Gentner, D., & Bowerman, M. (2009). Why some spatial semantic categories are harder to learn than others: The typological prevalence hypothesis. In E. L. J. Guo, N. Budwig, S. Ervin-Tripp, K. Nakamura, & S. Ozcaliskan (Eds.), (Vol. 4652009, pp. 465–480). Psychology Press.
- Greenhill, S. J. (2022). *Pronoun paradigms database*. Retrieved from <http://pdb.simon.net.nz/>
- Greenhill, S. J., Heggarty, P., & Gray, R. D. (2020). Bayesian phylolinguistics. In R. D. Janda, B. D. Joseph, & B. S. Vance (Eds.), *The Handbook of Historical Linguistics* (Vol. II, pp. 226–253). Wiley Online Library.

- Heim, I. (2008). Features on bound pronouns. In D. Adger, S. Bejar, & D. Harbour (Eds.), *Phi theory: Phi-features across modules and interfaces* (Vol. 35, p. 56). Oxford University Press.
- Herce, B., Saldana, C., Mansfield, J., & Bickel, B. (2023). Naturalness is gradient in morphological paradigms: Evidence from positional splits. *Glossa*, 8(1), online. doi: 10.16995/glossa.9280
- Jäger, G., & Wahle, J. (2021). Phylogenetic typology. *Frontiers in Psychology*, 2852. doi: 10.3389/fpsyg.2021.682132
- Jakobson, R. (1936). Beitrag zur allgemeinen Kasuslehre: Gesamtbedeutungen der russischen Kasus. *Travaux du Cercle Linguistique de Prague* 6, 240–288.
- Johnson, T., Gao, K., Smith, K., Rabagliati, H., & Culbertson, J. (2021). Investigating the effects of i-complexity and e-complexity on the learnability of morphological systems. *Journal of Language Modelling*, 9(1), 97–150. doi: 10.15398/jlm.v9i1.259
- Kimball, G. D. (1991). *Koasati grammar*. Univ. of Nebraska Press.
- Landau, B., & Shipley, E. (2001). Labelling patterns and object naming. *Developmental science*, 4(1), 109–118. doi: 10.1111/1467-7687.00155
- Maldonado, M., & Culbertson, J. (2021). Person of interest: Experimental investigations into the learnability of person systems. *Linguistic Inquiry*. doi: 10.1162/ling_a.00406
- Maldonado, M., & Saldana, C. (2022). Interpretations of plurality in personal pronouns differ across person categories: An experimental study. In *Proceedings of the 23rd amsterdam colloquium*.
- Mansfield, J., Saldana, C., Hurst, P., Nordlinger, R., Stoll, S., Bickel, B., & Perfors, A. (2022). Category clustering and morphological learning. *Cognitive Science*, 46(2), e13107. doi: 10.1111/cogs.13107
- Müller, G. (2004). A distributed morphology approach to syncretism in russian noun inflection. In *Proceedings of fast* (Vol. 12, pp. 353–373).
- Nevins, A. (2015). Productivity and portuguese morphology: How experiments enable hypothesis-testing. In E. Aboh, J. Schaeffer, & P. Sleeman (Eds.), *Romance languages and linguistic theory* (pp. 175–201).
- Nevins, A., Rodrigues, C., & Tang, K. (2015). The rise and fall of the L-shaped morpheme: Diachronic and experimental studies. *Probus*, 27(1), 101–155. doi: 10.1515/probus-2015-0002
- Noyer, R. R. (1992). *Features, positions and affixes in autonomous morphological structure*. Unpublished doctoral dissertation, Massachusetts Institute of Technology.
- Pertsova, K. (2011). Grounding systematic syncretism in learning. *Linguistic Inquiry*, 42(2), 225–266. doi: 10.1162/LING_a.00041
- Pertsova, K. (2014). Logical complexity in morphological learning: effects of structure and null/overt affixation on learning paradigms. In *Annual Meeting of the Berkeley Linguistics Society* (Vol. 38, pp. 401–413).
- Pertsova, K. (2022). *PRONOM*. doi: 10.17615/h78a-9w72
- Pothos, E. M., Chater, N., & Stewart, A. J. (2004). Information about the logical structure of a category affects generalization. *British Journal of Psychology*, 95(3), 371–386. doi: 10.1348/0007126041528158
- Quine, W. V. O. (1960). *Word and object*. MIT press.
- Realí, F., & Griffiths, T. L. (2009). The evolution of frequency distributions: Relating regularization to inductive biases through iterated learning. *Cognition*, 111(3), 317–328. doi: https://doi.org/10.1016/j.cognition.2009.02.012
- Saldana, C., Herce, B., & Bickel, B. (2022). More or less unnatural: Semantic similarity shapes the learnability and cross-linguistic distribution of unnatural syncretism in morphological paradigms. *Open Mind*. doi: 10.1162/opmi_a.00062
- Smith, K. (2018). The cognitive prerequisites for language: insights from iterated learning. *Current opinion in behavioral sciences*, 21, 154–160. doi: https://doi.org/10.1016/j.cobeha.2018.05.003
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2), 245. doi: 10.1037/0033-295X.114.2.245
- Xu, Y., Duong, K., Malt, B. C., Jiang, S., & Srinivasan, M. (2020). Conceptual relations predict colexification across languages. *Cognition*, 201, 104280. doi: 10.1016/j.cognition.2020.104280