Cognition 171 (2018) 194–201

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

Social biases modulate the loss of redundant forms in the cultural evolution of language

Gareth Roberts[a,*,1], Maryia Fedzechkina[b,c,d,1]

[a] Department of Linguistics, University of Pennsylvania, United States
[b] Department of Linguistics, University of Arizona, United States
[c] Graduate Interdisciplinary Program in Cognitive Science, University of Arizona, United States
[d] Graduate Interdisciplinary Program in Second Language Acquisition and Teaching, University of Arizona, United States

ARTICLE INFO

Keywords:
Language change
Learning biases
Social biases
Iterated language learning
Cultural evolution
Artificial language learning

ABSTRACT

According to the competitive exclusion principle (Gause, 1934), competition for the same niche must eventually lead one competitor to extinction or the occupation of a new niche. This principle applies in both biology and the cultural evolution of language, where different words and structures compete for the same function or meaning (Aronoff, 2016). Across languages, for example, word order trades off with case marking as a means of indicating who did what to whom in a sentence. Previous experimental work has shed light on how such trade-offs come about as languages adapt to human biases through learning and production, with biases becoming amplified through iterated learning over generations. At the same time, a large body of work has documented the impact of social biases on language change. However, little work has investigated how social biases interact with learning and production biases. In particular, the social dimension of language may provide alternative niches for otherwise redundant forms, preventing or slowing their extinction. We tested this hypothesis in an iterated-learning experiment in which participants were exposed to a language with two dialects, both of which had fixed word order, but differed in whether they employed case markers. In one condition, we biased participants socially towards speakers of the dialect that employed case; in other conditions we provided no bias, or biased participants for or against the dialect without case. As expected under our hypothesis, the use of case markers declined over time in all conditions, but the social bias in favor of case-dialect speakers slowed the decline.

1. Introduction

In language there is often more than one way to say the same thing (Labov, 1972). Indeed, the arbitrariness of linguistic form-meaning mappings means that words can vary infinitely in form: The words *dinner*, *vecheria*, *swper*, *Abendessen*, *cena*, *wanfan*, for instance, represent a tiny subset of the many words that exist for the evening meal. But this variation is constrained. In the case of individual words, it is vanishingly rare to find perfect synonyms. For example, *dinner* and *supper* might refer to the same meal in different dialects, but almost never carry exactly the same meaning for the same speaker – suggesting that it is very hard for two different words to occupy the same semantic niche (Taylor, 2002, p. 471). This phenomenon can be understood in evolutionary terms with reference to the competitive exclusion principle, according to which two competing entities cannot coexist indefinitely in the same niche (Gause, 1934; Hardin, 1960). The result of competition is either that one of the forms drives the other to extinction (as with English *yes* and *no* versus *yea* and *nay*) or that they come to occupy different niches, as with *regal*, *royal* and *kingly* (see Aronoff, 2016, for a discussion of the competitive exclusion principle in language).

As with the lexicon, the grammars of natural languages exhibit both variation between languages and competition between different structures within the same language. For example, languages differ in their preferred order of the subject and object in simple transitive sentences and in the amount of flexibility they allow in ordering these constituents (Dryer & Haspelmath, 2013). Latin and Russian are examples of languages with relatively flexible word order, while English and French are examples of languages with relatively fixed word order. Within a language, word order competes with other cues to grammatical role assignment (i.e., who is doing what to whom in a sentence). The most obvious competitor is morphological case marking (i.e., changes to the form of certain categories of words to indicate their grammatical role), though the same information can also be conveyed

* Corresponding author.
E-mail address: gareth.roberts@ling.upenn.edu (G. Roberts).
[1] Both authors contributed equally to this work.

https://doi.org/10.1016/j.cognition.2017.11.005
Received 27 July 2017; Received in revised form 10 November 2017; Accepted 10 November 2017

by such means as agreement, prosody, or pragmatics (van Everbroeck, 2003). Importantly, no language uses all of these means at the same time. In fact, cues to grammatical role assignment have long been observed to trade off across languages, with case marking and word order providing a good example (Blake, 2001, p. 15; Sapir, 1921, p. 66). In the Modern English sentence *Brutus killed Caesar*, word order unambiguously conveys who did the killing. In Latin, however, *Brutus Caesarem interfecit* and *Caesarem Brutus interfecit* both mean "Brutus killed Caesar", and it is the case markers on *Caesar* and *Brutus* that indicate their grammatical roles. Further support for the existence of this trade-off comes from language change. In later Latin, for instance, case marking became less consistent and word order became more fixed, so that the modern Romance languages typically exhibit relatively fixed SVO word order and little case (Kabatek & Pusch, 2011). A similar process occurred in the history of English (Marchand, 1951).

The trade-off between cues to grammatical role assignment has been argued to stem from a trade-off between robust message transmission – the speaker's goal to be understood – and production effort (Fedzechkina, Jaeger, & Newport, 2012, 2016; Kurumada & Jaeger, 2015). In languages with relatively fixed word order, grammatical role assignment can usually be inferred based on word order alone, rendering case marking redundant. In flexible-word-order languages, however, case provides important information about sentence meaning, as word order is less informative of grammatical role assignment. As case marking requires effort to produce, the implication is that languages evolve culturally such that case marking is maintained predominantly in those languages when its utility is high (i.e., those with flexible word order). Experimental support for this claim comes from work using the miniature artificial language learning paradigm. For example, Fedzechkina et al. (2016) showed that participants introduced cross-linguistic patterns of case and word order trade-offs into novel miniature languages if the input grammars were not consistent with such trade-offs. In particular, learners of a miniature language with flexible word order were more likely to maintain case marking in their own productions, while learners of a language with fixed word order tended to drop it.

Work on cumulative cultural evolution has linked the cognitive biases influencing individual learners to long-term patterns of language change. This work suggests that even small learning and production biases that are too weak to be detected in one generation of learners can have sizeable effects on the linguistic system over multiple generations (Kirby, Griffiths, & Smith, 2014). Supporting evidence for this claim comes in particular from studies using the iterated-learning paradigm, in which the linguistic output of one learner is used as the input for another learner, who is either a simulated agent (e.g., Kirby, 1999; Reali & Griffiths, 2009) or a human participant (e.g., Kirby, Tamariz, Cornish, & Smith, 2015; Smith & Wonnacott, 2010).

The linguistic system is also subject to a range of social biases alongside the learning and production biases discussed above (Labov, 2001). One's choice of name for the evening meal, for example, may communicate more than which meal is being referred to: In certain parts of Great Britain and Ireland, the use of "tea" in this sense is associated with working-class speakers and thus implies different things about the speaker's origins and social identity than the more middle-class option, "dinner". Over the last century a large body of work has documented the important role of social factors in language change (Bailey, Cameron, & Lucas, 2013; Labov, 2001), including the role of local identity (e.g., Pope, Meyerhoff, & Ladd, 2007), ethnicity (Lanehart, 2015), gender (Holmes & Meyerhoff, 2003), and class (Rampton, 2010).

While the impact of social factors on language change has also been acknowledged by researchers taking an explicitly cultural-evolutionary approach (e.g., Croft, 2000), experimental work on the cultural evolution of language has primarily focused on the role of learning and production biases in language change (e.g., Fedzechkina et al., 2016; Kirby et al., 2015) and has paid relatively little attention to the role of social biases. This is unfortunate, as social and non-social biases are likely to interact, jointly shaping the process of language change. For example, we know from previous experimental work (Fedzechkina et al., 2016) that, as a result of biases for efficient communication, learners of miniature languages are more likely to drop redundant case markers if word order is fixed than if it is flexible. However, such change tends to lead to variation between speakers (because some adopt the change sooner than others), and this variation provides raw material for social meaning. That is, if there is more than one way of saying the same thing, one of those ways may acquire social significance. This, in turn, might influence the trajectory of language change so that it differs from what we would predict based on learning and production biases alone. A natural-language example of this seems to be provided by English *whom*, the object form of *who*. In modern English, *whom* has become redundant and competes with *who* in the object position. In evolutionary terms, its niche has been invaded by a clearly fitter competitor: Outside some narrow contexts in formal written English, *who* is acceptable everywhere that *whom* is, and a speaker who does not acquire *whom* is at no serious disadvantage. The reverse is not true. So why has *whom* not disappeared yet? One possibility suggested by the competitive exclusion principle (Aronoff, 2016; Gause, 1934) is that it has found a new niche. A word like *whom* tends to be associated with more educated speakers (cf. Milroy & Milroy, 2012), and may thus serve as a means of signaling group identity, which may slow its disappearance from the linguistic system.

The purpose of this paper is to experimentally test the hypothesis that social biases can interact with learning and production biases during the cultural evolution of language, modulating the loss of a redundant form (such as case marking in a language with fixed word order) that would otherwise be expected to disappear. In evolutionary terms, the question is whether a form under threat from a fitter competitor for the same communicative niche (such as word order in later Latin, which – by becoming more fixed – had become a more informative cue) will survive longer if it has a new social niche – e.g., a desirable social meaning – to take refuge in. We tested this hypothesis using a miniature artificial language learning paradigm, which has been shown to be well-suited to studying both the influences of individual-level biases on language structure (Culbertson, Smolensky, & Legendre, 2012; Fedzechkina et al., 2012; Hupp, Sloutsky, & Culicover, 2009; see Fedzechkina, Newport, & Jaeger, 2016, for a review) and the role of social biases in language change (Samara, Smith, Brown, & Wonnacott, 2017; Sneller & Roberts, 2018; see Roberts, 2017, for a review).

In our study, participants learned a miniature "alien language" with fixed word order. The language had two dialects (indexed by the color of the alien speakers), one of which consistently employed redundant case marking, while the other had no case marking at all. We manipulated social biases acting on participants by varying the information provided about the different alien groups, encouraging participants to feel positively or negatively oriented towards one color of alien compared with the other (cf. Labov, 1963). This yielded four conditions in total: first, a condition biasing participants towards the aliens who used case markers; second, a condition biasing participants *against* the aliens who used case markers; third, a condition biasing participants towards the aliens who did *not* use case markers; and, fourth, a control condition with no bias towards either alien group. We simulated the generational transmission of language by using iterated learning, in which the output of learners is used to generate the input to other learners, creating "chains" of participants (Kirby et al., 2014). Given that case marking was in constant competition with word order, and taking into account evidence from previous work for a bias against maintaining excessive redundancy (Fedzechkina et al., 2016), we expected that the case markers would disappear from the language over time. We predicted, however, that this process would be modulated by social biases, and that the redundant case marking would persist over more generations when there was a social bias to feel positively oriented towards the group of aliens who used case in their dialect.

2. Method

2.1. Participants

A total of 359 participants were recruited online via Amazon Mechanical Turk using a custom-designed flash applet developed by Hal Tily (Tily, Frank, & Jaeger, 2011). All participants had US IP addresses and were self-reported speakers of North-American English. They received $4 for their participation in the experiment, which lasted approximately 35 min. Some participants' data were discarded and replaced by new data; this occurred for two separate reasons. First, 108 participants were replaced due to low learning accuracy (see Section 2.3.2 for details). Second, owing to an oversight, several participants took part in the experiment more than once. The data from their first participations were kept, but the data from all subsequent participations, along with the data from later generations of the same chains, were discarded and new participants recruited. This led to the replacement of data from 51 participants in total. The final sample submitted for analysis included 200 participants.

2.2. Miniature input language

Participants were instructed that they would be learning a novel miniature "alien" language with the help of alien informants (two for each dialect), by watching short videos accompanied by their descriptions in the novel language (presented both auditorily and in writing). The language contained 11 novel words (Table 1): six nouns referring to humanoid characters (CHEF, REFEREE, BANDIT, CONDUCTOR, HUNTER, MOUNTIE), four transitive verbs (TAP, HUG, KICK, ROCK), and a case marker 'di' that followed the object in the sentence. The lexicon was adopted from Fedzechkina, Chu, and Jaeger (in press). Individual words were synthesized separately using the Apple speech synthesizer (voice 'Alex') and were concatenated into sentences in real time during the experiment. To prevent unwanted associations that could interfere with learning, the assignment of novel word labels to lexical items was rotated across two randomly generated lists, such that, for instance, *peza* meant "chef" for some participants and "conductor" for others (see Supplementary materials). All nouns occurred equally often in the subject and object position with each verb.

The language had two "dialects". In the first generation, both dialects had fixed subject-object-verb (SOV) constituent order. This order was chosen because default SOV order commonly occurs with case marking typologically (Dryer & Haspelmath, 2013). The dialects differed in one respect only: One dialect (the case dialect) had a case marker 'di' that followed the object 100% of the time (there was no case marker on the subject), while the other dialect (the no-case dialect) had no case marking at all. The two dialects occurred equally frequently in the input. Case marking was thus present in 50% of the input sentences in the first generation. Since constituent order always perfectly disambiguated grammatical role assignment, case marking provided little information about sentence meaning, but required effort to produce.

Participants were told explicitly that the language had two dialects, and that one of the dialects was used by blue aliens, while the other was used by orange aliens. Assignment of color to dialect was counterbalanced across participants.

Table 1
Words of the miniature language, organized by part of speech.

Nouns	Verbs	Case marker
barsa	skroop	di
forpah	velmik	
doakla	kyse	
rizba	tegud	
koofta		
peza		

2.3. Procedure

2.3.1. Language learning task

At the start of the experiment participants read an introductory scenario informing them that they were on a mission to an alien planet and would be learning the language of the alien species living on the planet. Participants were further informed that there were two different groups of aliens on the planet, who could be identified by their color, and that each group spoke a different dialect of the alien language. Additional information was then provided about the different alien groups depending on the experimental condition (see Section 2.4 and Fig. 2). This introduction was followed by a miniature artificial language learning task adapted from Fedzechkina et al. (2016), which consisted of several blocks, as detailed below and in Fig. 1. In all blocks of the experiment, the stimuli in the novel language were presented both auditorily and in writing. The experiment lasted approximately 35 min in total.

Noun exposure. At the start of the experiment, participants were presented with pictures of the six humanoid characters one at a time, accompanied by their labels in the novel language. Each character was presented twice, with the order of presentation randomized for each participant.

Noun comprehension. Participants' recognition of humanoid character names was tested. Participants were presented with pictures of two characters accompanied by an alien-language word referring to one of them and were asked to click on the correct picture. Each character was presented twice within the block. Participants did not receive feedback on each trial, but they were provided with a summary of their performance at the end of the block. Participants had to score at least 90% correct on this block to move to the next block; otherwise they had to repeat the noun exposure block.

Noun identification. Participants' ability to identify the names of the humanoid characters was tested. Each character was presented in isolation (twice within the block in random order) along with the entire lexicon of the language at the top of the screen. Participants were instructed to click on the alien word corresponding to the character's name. As in the noun comprehension block, participants received a summary of their performance at the end of the block. They were also required to score at least 90% to move to the next block; otherwise they had to repeat the noun exposure block.

Sentence exposure. Participants were shown short animations depicting simple transitive events performed by characters on one another, accompanied by sentences describing them in the alien language. Every animation included a blue or an orange alien at the bottom of the video as a cue to the dialect being used (Fig. 1). The sentence and animation played concurrently. Participants could replay each animation (along with the sentence) as many times as they wanted by clicking a "Replay" button. The block included 24 different animations.

Sentence comprehension. Participants were shown two previously seen animations depicting the same characters and action, but with the actor in one of the videos taking the role of patient in the other. The animations were accompanied by a single sentence in the novel language, and participants were instructed to click on the animation that corresponded to it. As in sentence exposure, each animation included an alien as a cue to the dialect used. No feedback was provided at the end of either the trial or the block.

Participants were presented with two sets of two sentence exposure blocks and one sentence comprehension block (24 trials each). The association between the animation and alien dialect (cued by alien color) was counterbalanced across the blocks (e.g., an animation for which participants heard the dialect of the blue aliens in the first and third exposure blocks would be shown in the dialect of the orange aliens in the second and fourth blocks).

Sentence production. Participants were shown the entire lexicon of the language at the top of the screen (all nouns were shown as non-case-marked and the case marker was shown as a separate word) and

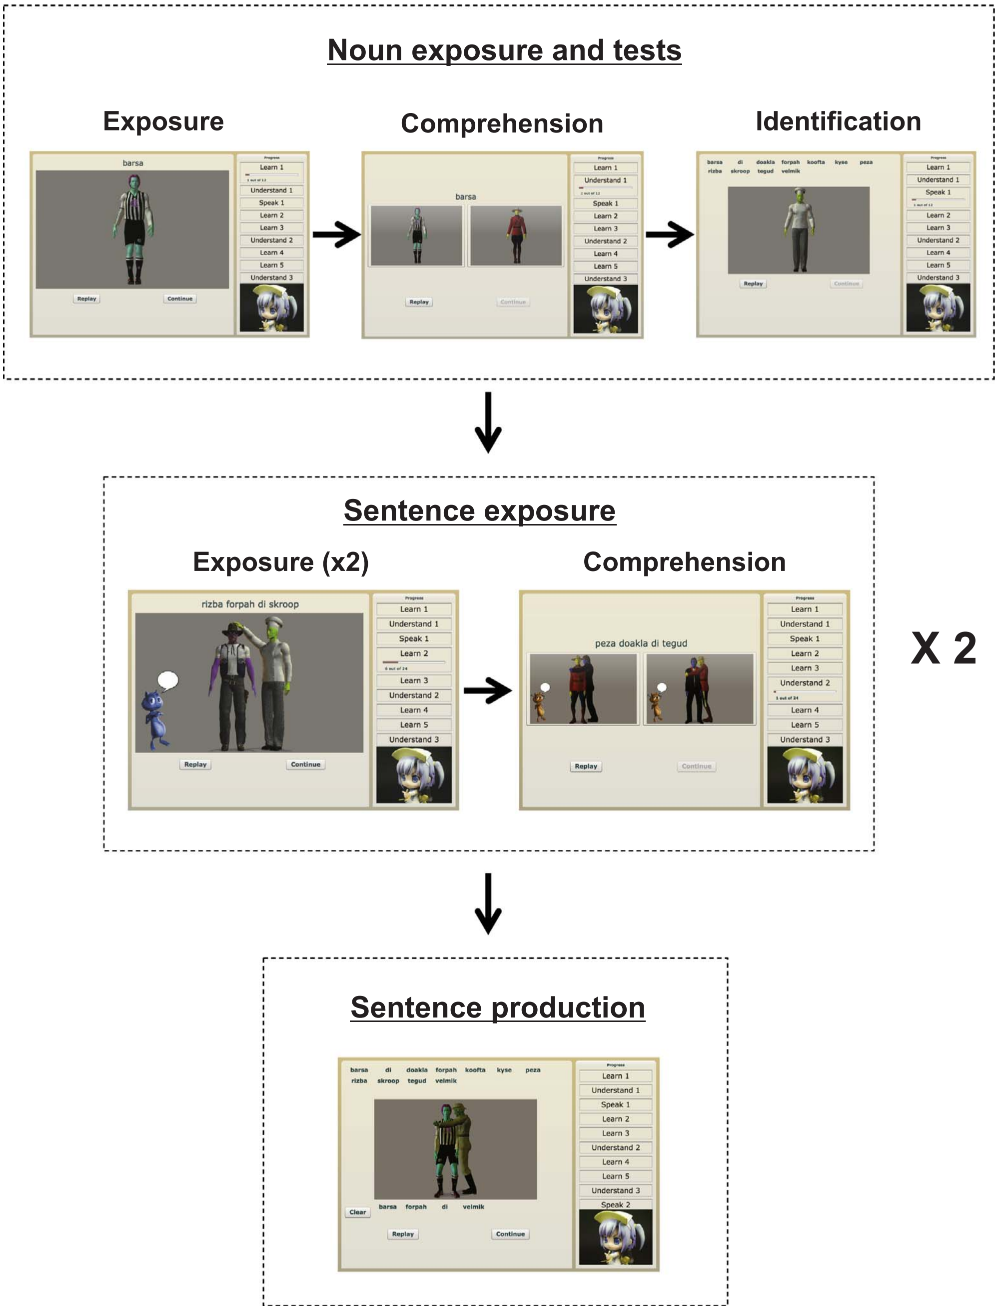

Fig. 1. Experimental procedure. Images represent screenshots of sample trials in each of the experimental blocks. Arrows indicate the succession of the blocks – the experiment started with noun exposure and ended with sentence production.

asked to describe a series of previously unseen animations (24 in total) by clicking on the corresponding lexical items. To simulate additional effort associated with case production (in line with natural language, where case use requires additional linguistic material), the case marker was treated as a separate word (with its own button). Thus, case use required one additional click compared to a non-case-marked noun. There was no image of an alien as dialect cue in this block (Fig. 1). No feedback or feedback summary was provided at the end of either the trial or the block.

2.3.2. *Iterated learning*

In each condition, there were ten chains with a maximum of five generations each, with two participants in each generation per chain (yielding a total of 20 participants per generation). As described in Section 2.2, participants in the first generation were taught a language with 100% SOV word order, and two dialects, one in which the object was always marked by the case marker *di* and one in which there was no case marking. For later generations, the syntax of the case dialect was generated based on the previous participants' responses in the sentence production block, with each participant's responses being used to generate the language for one alien informant (as noted above, there were two alien informants for each dialect).

Because retention of the case marker was the main variable of interest, it was important to be confident that the loss of case in a participant's productions was not simply due to a lack of attention during training. Therefore, if a participant's overall mean success rate in the noun comprehension or noun identification block was below 50% – implying poor attention to the learning task – their data were discarded and a new participant recruited. The same was done if more than 25% of a participant's sentences during the sentence production block were problematic, that is, did not consist of exactly three words (not counting the case marker *di*), or contained a lexical error. In total, 108 participants were replaced based on these criteria.

For the participants who were not replaced, the proportions of different sentence patterns (e.g., SOV, SOcV, where c represents the case marker *di*) were calculated based on non-problematic sentences in the

***No bias* condition:**

We are **keen to trade with the aliens**. They seem to be on our side, and they have important resources. We should **try to impress them**.

***Bias for case & Bias for no case* conditions:**

We are especially **keen to trade with the blue** aliens. They seem to be on our side, and **they have important resources**. We should **try to impress these blue aliens** in particular. The orange aliens are not so friendly, and they don't have much to offer us.

***Bias against case* condition:**

We are especially **reluctant to trade with the blue aliens**. We **don't think we can trust them**, and **they have no important resources**. We should be **very wary of these blue aliens** in particular. The orange aliens are not as bad, but we don't know if they have much to offer us.

Fig. 2. Example instructions for experimental conditions. Relevant parts of the instructions are highlighted in bold for convenience (participants saw no highlighting). Note that the color of alien that corresponded to a particular dialect (and was thus subject to a particular bias) was counterbalanced.

sentence production block. For this purpose, sentences with a single lexical error, in which word order could still be determined, were counted towards the proportion. The proportions calculated from each participant were used to generate the syntax of the case dialect for one of the two alien informants in the next generation of the same chain. Sentences in the no-case dialect were not based on the output of participants, and simply had SOV word order with no case marker in every generation. This was done to ensure that the case-dialect would always be in competition with a dialect that consistently had no case marking.

Based on prior work (Smith & Wonnacott, 2010), participants were recruited until there were five generations in each chain, or until case marking had entirely disappeared. It is important to understand that participants produced sentences by clicking on words that they had been exposed to during training (with the case marker treated as a separate word). This meant that, if neither member of a generation used the case marker at all, it would not appear in training for the next generation, and so would not be available during the sentence production block for that generation. Thus, once case marking disappeared in a chain, it could not reappear. In such cases no more participants were recruited for that chain, and any later generations in the chain were simply treated as having no case marker in the case use analysis. There were 200 such "dummy" participants (50% of the data) in our experiment.

2.4. Conditions

Participants were randomly assigned to one of four conditions. In the *No bias* condition, participants were made aware of the two alien groups speaking different dialects but were not encouraged to feel more or less positive about either alien group relative to the other. In two other conditions a social bias was introduced in the experiment instructions to encourage participants to feel positively oriented towards one group of aliens (see Fig. 2 for the instructions used in the experiment). In the *Bias for no case* condition, participants were biased towards the aliens who spoke the dialect with no redundant case marking. In the *Bias for case* condition, participants were biased towards the speakers of the dialect with redundant case marking. As Fig. 2 shows, each of the biasing conditions not only introduced a social bias towards one of the two alien groups, but also explicitly mentioned that group, possibly making it more salient to the participants. Thus, learners of the *Bias for case* condition might have been more likely to use case in their own productions simply because our instructions prompted them to pay more attention to the speech of the aliens speaking this dialect. To control for this potential confound, we ran a fourth condition, the *Bias against case* condition, in which we explicitly biased participants against the dialect with case marking (while at the same time explicitly mentioning this group in the instructions). In all cases, the bias was expressed with respect to the color of the aliens in question, rather than with respect to the dialect.

2.5. Predictions

Since case marking in the input language was a redundant cue to grammatical role assignment, but required additional effort to produce (involving an extra button press), we expected, in line with previous work (Fedzechkina et al., 2016), that case markers would become less frequent in all conditions, potentially leading to a system with consistent zero-marking. We hypothesized, however, that this preference to use case efficiently would be modulated by social pressures. Specifically, we predicted that case would be maintained in the language longer if there was a social bias in its favor. In other words, we expected that case marking would be maintained over more generations of learners in the *Bias for case* condition compared to all other conditions as this condition encouraged participants to feel positively oriented towards the aliens speaking a dialect with consistent case marking.

3. Results

3.1. Word order in production

Before turning to our central hypothesis – the retention of redundant case marking as a result of a positive social bias – we analyzed participants' word order preferences in production. Although the case marker was redundant in the initial language, it was possible for participants to make it informative by introducing flexible word order into the languages they produce. This might lead learners to maintain case marking in the language regardless of whether there was a social bias in its favor. However, this did not happen[2]: Although participants occasionally produced other word orders, they showed an overall strong preference for SOV (96% SOV across all generations in the *No bias* condition; 96% SOV across all generations in the *Bias for no case* condition; 91% SOV across all generations in the *Bias for case* condition; 99% across all generations in the *Bias against case* condition). This preference for matching the input word order distribution replicates previous findings from a single generation of learners (Fedzechkina et al., 2016).

To test whether participants' word order preferences differed significantly across conditions (which might influence case retention), we conducted a mixed effects logistic regression analysis to predict SOV use from condition and generation as fixed effects. Generation was

[2] "Dummy" participants created for the purposes of case-use analysis in chains where case had disappeared (see Section 2.3.2) were excluded from word order analysis as their word order use was impossible to determine.

treated as a continuous variable and was centered. Condition was Helmert coded (a coding scheme in which each level of the variable is compared with the mean of the preceding levels; Fox, 2002) as this coding scheme best reflected the comparisons predicted by our hypothesis (i.e., the prediction that learners of the *Bias for case* condition would have different learning outcomes compared with all other conditions). The first contrast compared *Bias for no case* and *No bias* conditions; the second contrast compared *Bias against case* condition to the joint mean of *No bias* condition and *Bias for no case* conditions; the third contrast compared *Bias for case* to the joint mean of all other conditions. The model included the maximal random effects structure justified by the data based on backwards model comparison (random intercepts for participant and object-noun). The results did not change when the maximal converging random effects structure was used. Participants' word order use did not differ depending on condition (*No bias* vs. *Bias for no case*: $\hat{\beta} = -0.24, z = -0.24, p = .811$; *Bias against case* vs. the mean of *No bias* and *Bias for no case*: $\hat{\beta} = 0.28, z = 0.39, p = .694$; *Bias for case* vs. the mean of all other conditions: $\hat{\beta} = -0.17, z = -0.51, p = .614$). There was no main effect of generation, suggesting that participants' word order preferences remained stable over time ($\hat{\beta} = -0.22, z = -0.44, p = .661$).

These data confirm that, as expected, participants did not respond to the competition between fixed word order and case marking by introducing more variable word order. The following section presents participants' case marking preferences over time, our main variable of interest.

3.2. Case marking in production

Our central hypothesis was that learners would be more likely to maintain case marking if it carried social meaning. To test this, we used Cox mixed effects regression, a type of survival analysis, to predict case omission over time (i.e., over generations of learners) from social bias condition (Helmert coded as *No bias* vs. *Bias for no case* condition; *Bias against case* vs. the joint mean of *No bias* and *Bias for no case* conditions; *Bias for case* condition vs. the joint mean of all other conditions). The model included the maximal random effects structure justified by the data based on backwards model comparison (random intercepts for participant and object-noun). The results did not change when the maximal converging random effects structure was used.

As expected under our hypothesis, while case marking tended to be lost over time in all conditions, this loss was significantly slower in the *Bias for case* condition compared with all other conditions ($\hat{\beta} = -0.42$, $z = -11.02$, $p < .001$, see Fig. 3), suggesting that a positive social bias in favor of the case-dialect speakers caused case marking to be maintained in the language for longer. There was no significant difference in case loss over time between the *No bias* and the *Bias for no case* conditions ($\hat{\beta} = -0.04, z = -0.42, p = .67$) or between the *Bias against case* and the joint mean of the *No bias* and the *Bias for no case* conditions ($\hat{\beta} = 0.002$, $z = 0.04$, $p = .97$).

These findings support our hypothesis: Case marking in all dialects was redundant and, consistent with previous results from one generation of learners (Fedzechkina et al., 2016), tended to be lost in all conditions over time. However, this trend, stemming from biases for efficient communication, was modulated by a positive social bias in favor of case-dialect users. Consistent with the competitive exclusion principle (Aronoff, 2016; Gause, 1934), case tended to persist in the language over more generations of learners when it had a social niche to occupy.

4. Discussion

We tested the hypothesis, derived from the competitive exclusion principle (Aronoff, 2016; Gause, 1934; Hardin, 1960), that a linguistic form in competition with a fitter form for the same communicative niche would either eventually become extinct or could survive by occupying a new niche. More specifically, we hypothesized that, as a result of biases for efficient communication, a redundant case marker would be lost through the process of repeated transmission, but that it would be retained longer if there were a social bias in favor of its users. This is what we found, and our results are consistent with the competitive exclusion principle. Given the competition between two means of conveying the same information (case marking and word order in our study), one (case marking) was either forced to extinction (as in three of our conditions) or came to occupy a new social niche (as in the *Bias for case* condition, where such a niche was available). The trajectory of the case marker use over time in the *Bias for case* condition looks strikingly different from its trajectory in the other three conditions, which do not differ significantly from each other. In the default situation, given no particular social bias, the case marker disappeared rapidly and was entirely lost by the fourth generation in all chains. It also declined at the same rate if there was an explicit bias against the case dialect or an explicit bias in favor of the no-case dialect. The only bias that influenced case use compared with no bias at all was the social bias in favor of the case dialect, and the effect of this was dramatic.

It is important to note that the frequency of case marker use decreased over time in all four conditions. This includes the *Bias for case* condition, where it was used the majority of the time in the first generation of all chains (in striking contrast to the results for the other chains). The mean rate of case use in this condition still eventually fell below 25%, and it disappeared entirely in five of the ten chains. The social bias, therefore, only modulated the decline of case marking over time. One should be careful not to interpret this finding in terms of learning and production biases having a stronger effect on language change than social biases. Our question was whether a social bias could help case survive against the odds, so we stacked the deck against case as a matter of design: The no-case dialect taught to participants was always the same (i.e., unlike the case dialect, it was not based on previous participants' production), and thus sentences in that dialect never contained case marking in any generation.

Similarly, our findings should not be taken to speak to the relative strength of case marking and word order as cues for grammatical role assignment. It is not surprising that participants introduced changes in case as opposed to word order distributions. Previous work in miniature-artificial-language learning has shown that adult participants tend not to introduce changes to the input distributions of cues that contain no variability (e.g., Fedzechkina et al., 2016; St Clair, Monaghan, & Ramscar, 2009; Tily et al., 2011). The fact that word order was consistent and case marking varied in the initial input could have biased participants to introduce changes into the case system as opposed to word order. Participants were also exposed to only simple transitive sentences presented out of context. Given a different kind of task, in which, for instance, participants were encouraged to emphasize different elements in the sentence, we might see greater flexibility emerging in word order, as with focus- or topic-marking systems in many natural languages (Givón, 1988; Roberts & Stevens, 2017; Yokayama, 1986).

There are two potential limitations to our study. One is the relative simplicity of the social bias in our experiment. Participants were simply encouraged to feel positively or negatively oriented towards one of the groups of aliens. This lacks the nuance of some sociolinguistic work. Social meaning, for example, had been argued to exist in a complex "indexical field" – a constellation of related social meanings in which the same form can have multiple simultaneous meanings, whose relative importance varies according to context (Eckert, 2008). However, while it would certainly be valuable to explore more complex biases in future work, we consider the simplicity of the bias in the current study to be valuable in itself. In particular, our simple manipulation produced very clear results that correspond intuitively to patterns seen in natural language, providing a useful baseline for comparison with data resulting from more complex biases. Our approach is also consistent with

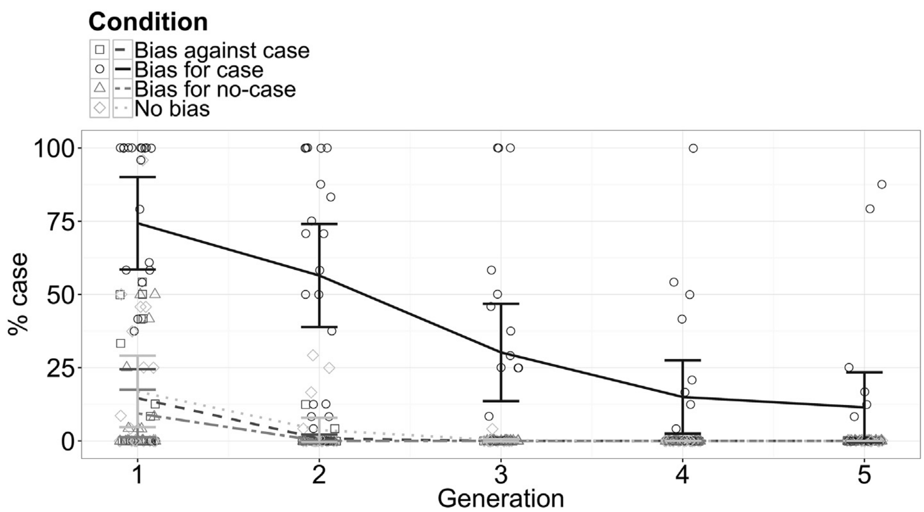

Fig. 3. Case marker use by condition. Lines represent mean case use at each generation; shapes represent individual participants. The error bars represent bootstrapped 95% confidence intervals.

established paradigms in social psychology, and in particular with the minimal group paradigm, in which participants take part in tasks with partners who are stated to be either similar or different to them in some task-irrelevant way (Tajfel, 1982). A typical (and robust) finding within this paradigm is that human beings are surprisingly susceptible to making simplistic in-group/out-group distinctions (Hogg, 2016). While our study does not fit straightforwardly within this paradigm, our results are conceptually consistent with this observation.

A second potential limitation is that the experiment tested social and communicative biases but involved only learning, with no real communication or social interaction. Including direct communication in a genuine social context into an experimental approach of this kind (see, e.g., Fay, Garrod, Roberts, & Swoboda, 2010; Galantucci, 2005; Kirby et al., 2015; Roberts, 2010; Sneller & Roberts, 2018) would increase the ecological validity of the study. It would allow us to investigate the role of such processes as accommodation or interactive alignment, whereby interacting individuals converge along various behavioral dimensions, including language (Giles, Coupland, & Coupland, 1991; Pickering & Garrod, 2004). Based on prior work in this area (e.g., Pardo, 2006; Walker & Campbell-Kibler, 2015), we predict that, while the general effect observed in our study would remain the same, case marker use would be subject to interlocutor influences (cf. Sneller & Roberts, 2018), such that the overall pattern of case loss or retention would be more or less pronounced depending on patterns of interaction. Like the inclusion of greater complexity of social biases, including communication would thus be a valuable step in future studies.

As it is, however, our study provides – via a simple manipulation – a compelling model of how Gause's competitive exclusion principle might operate in language, an explanation for the survival of certain redundant linguistic forms, and a demonstration of how the role of competing biases in the cultural evolution of language can be tested experimentally.

Acknowledgements

We thank Andrew Watts and Elisha Cooper for their help with setting up the study.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.cognition.2017.11.005. This includes a link to the dataset for this study, archived at the Open Science Framework (Roberts & Fedzechkina, 2017).

References

Aronoff, M. (2016). Competition and the lexicon. In A. Elia, C. Iacobino, & M. Voghera (Eds.). *Livelli di analisi e fenomeni di interfaccia. Atti del XLVII congresso internazionale della società di linguistica Italiana* (pp. 39–52). Rome: Bulzoni Editore.
Bailey, R., Cameron, R., & Lucas, C. (Eds.). (2013). *Oxford handbook of sociolinguistics*. Oxford: Oxford University Press.
Blake, B. J. (2001). *Case.* Cambridge: Cambridge University Press.
Croft, W. (2000). *Explaining language change.* Harlow: Pearson Education Ltd.
Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition, 122*, 306–329.
Dryer, M. S., & Haspelmath, M. (Eds.). (2013). *The World Atlas of Language Structures Online*Leipzig: Max Planck Institute for Evolutionary Anthropology.
Eckert, P. (2008). Variation and the indexical field. *Journal of Sociolinguistics, 12*(4), 453–476.
Fay, N., Garrod, S., Roberts, L., & Swoboda, N. (2010). The interactive evolution of human communication s ystems. *Cognitive Science,* 1–36.
Fedzechkina, M., Chu, B., & Jaeger, T. F. (2017). Human information processing shapes language change. *Psychological Science* (in press).
Fedzechkina, M., Jaeger, T. F., & Newport, E. (2012). Language learners restructure their input to facilitate efficient communication. *Proceedings of the National Academy of Sciences*.
Fedzechkina, M., Jaeger, T. F., & Newport, E. (2016). Balancing effort and information transmission during language acquisition: Evidence from word order and case-marking. *Cognitive Science*.
Fedzechkina, M., Newport, E., & Jaeger, T. F. (2016). The miniature artificial language learning paradigm as a complement to typological data. In L. Ortega, A. Tyler, H. I. Park, & M. Uno (Eds.). *The Usage-based Study of Language Learning and Multilingualism*. Georgetown: Georgetown University Press.
Fox, J. (2002). *An R and S-plus companion to applied regression.* Thousand Oaks, CA: Sage Publications Inc.
Galantucci, B. (2005). An experimental study of the emergence of human communication systems. *Cognitive Science, 29*(5), 737–767.
Gause, G. F. (1934). *The struggle for existence.* Baltimore: Williams and Wilkins.
Giles, H., Coupland, J., & Coupland, N. (1991). Accommodation theory: Communication, context and consequence. In H. Giles, J. Coupland, & N. Coupland (Eds.). *Contexts of accommodation: Developments in applied sociolinguistics* (pp. 1–68). Cambridge: Cambridge University Press.
Givón, T. (1988). The pragmatics of word order: Predictability, importance and attention. In M. Hammond, E. A. Moravcsik, & J. Wirth (Eds.). *Studies in syntactic typology* (pp. 243–284). Amsterdam: John Benjamins.
Hardin, G. (1960). The competitive exclusion principle. *Science, 131*, 1292–1297.
Hogg, M. A. (2016). Social identity theory. In S. McKeown, R. Haji, & N. Ferguson (Eds.). *Understanding peace and conflict through social identity theory* (pp. 3–17). Cham: Springer International Publishing.
Holmes, J., & Meyerhoff, M. (Eds.). (2003). *The handbook of language and gender*. Malden, MA: Blackwell.
Hupp, J. M., Sloutsky, V. M., & Culicover, P. W. (2009). Evidence for a domain-general mechanism underlying the suffixation preference in language. *Language and Cognitive Processes, 24*(6), 876–909.
Kabatek, J., & Pusch, C. D. (2011). The Romance languages: Typology. In van der Auwera, & B. Kortmann (Eds.). *The languages and linguistics of Europe: A comprehensive guide* (pp. 69–96). Berlin/New York: Mouton de Gruyter.
Kirby, S. (1999). *Function, selection and innateness: The emergence of language universals.* Oxford: Oxford University Press.
Kirby, S., Griffiths, T., & Smith, K. (2014). Iterated learning and the evolution of language. *Current Opinion in Neurobiology, 28*, 108–114.
Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition, 141*, 87–102.

Kurumada, C., & Jaeger, T. F. (2015). Communicative efficiency in language production: Optional case-marking in Japanese. *Journal of Memory and Language, 83*, 152–178.
Labov, W. (1963). The social motivation of a sound change. *Word, 19*, 273–309.
Labov, W. (1972). *Sociolinguistic patterns.* Philadelphia: University of Pennsylvania Press.
Labov, W. (2001). *Principles of linguistic changeSocial factorsVol. 2*. Malden, MA: Blackwell.
Lanehart, S. (Ed.). (2015). *The Oxford handbook of African American language*. Oxford: Oxford University Press.
Marchand, H. (1951). The syntactical change from inflectional to word order system and some effects of this change on the relation "verb/object" in English. A diachronic-synchronic interpretation. *Anglia-Zeitschrift Für Englische Philologie, 70*, 70–89.
Milroy, J., & Milroy, L. (2012). *Authority in language: Investigating language prescription and standardisation* (3rd ed.). London: Routledge.
Pardo, J. S. (2006). On phonetic convergence during conversational interaction. *Journal of the Acoustic Society of America, 119*(4), 2382–2393.
Pickering, M., & Garrod, S. (2004). Towards a mechanistic psychology of dialogue. *Behavioral and Brain Sciences, 27*, 169–226.
Pope, J., Meyerhoff, M., & Ladd, D. R. (2007). Forty years of language change on Martha's Vineyard. *Language, 83*(3), 615–627.
Rampton, B. (2010). Social class and sociolinguistics. *Applied Linguistics Review, 1*(1), 1–22.
Reali, F., & Griffiths, T. (2009). The evolution of frequency distributions: Relating regularization to inductive biases through iterated learning. *Cognition, 111*(3), 317–328.
Roberts, G. (2010). An experimental study of the role of social selection and frequency of interaction in linguistic diversity. *Interaction Studies, 11*(10), 138–159.
Roberts, G. (2017). The linguist's Drosophila: Experiments in language change. *Linguistics Vanguard, 3*(1), 20160086.
Roberts, G., & Fedzechkina, M. (2017). Social biases modulate the loss of redundant forms in the cultural evolution of language. *Open Science Framework*. http://dx.doi.org/10.17605/OSF.IO/XHB5G.
Roberts, G., & Stevens, J. (2017). Information theoretic factors in marking linguistic focus: A laboratory-language approach. *Proceedings of the 39th annual conference of the cognitive science society. Austin, TX*.
Samara, A., Smith, K., Brown, H., & Wonnacott, E. (2017). Acquiring variation in an artificial language: Children and adults are sensitive to socially conditioned linguistic variation. *Cognitive Psychology, 94*, 85–114.
Sapir, E. (1921). *Language: An introduction to the study of speech.* New York: Harcourt, Brace Retrieved from < http://www.bartleby.com/186/ > .
Smith, K., & Wonnacott, E. (2010). Eliminating unpredictable variation through iterated learning. *Cognition, 116*(3), 444–449.
Sneller, B., & Roberts, G. (2018). Why some behaviors spread while others don't: A laboratory simulation of dialect contact. *Cognition, 170C*, 298–311.
St Clair, M. C., Monaghan, P., & Ramscar, M. (2009). Relationships between language structure and language learning: The suffixing preference and grammatical categorization. *Cognitive Science, 33*(7), 1317–1329.
Tajfel, H. (1982). Social psychology of intergroup relations. *Annual Review of Psychology, 33*, 1–39.
Taylor, J. R. (2002). *Cognitive grammar.* Oxford: Oxford University Press.
Tily, H., Frank, M., & Jaeger, T. F. (2011). The learnability of constructed languages reflects typological patterns. In L. Carlson, C. Hoelscher, & T. Shipley (Eds.). *Proceedings of the 33rd annual conference of the cognitive science society* (pp. 1364–1369). Austin, TX: Cognitive Science Society.
van Everbroeck, E. (2003). Language type frequency and learnability from a connectionist perspective. *Linguistic Typology, 7*(1), 1–50.
Walker, A., & Campbell-Kibler, K. (2015). Repeat what after whom? Exploring variable selectivity in a cross-dialectal shadowing task. *Frontiers in Psychology, 6*, 546.
Yokayama, O. T. (1986). *Discourse and word order.* Amsterdam: John Benjamins.