UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

How communicative efficiency and social biases shape language in autistic and allistic learners

Permalink

Journal

Authors

Fletcher, Lauren E F
Culbertson, Jennifer
Rabagliati, Hugh

Publication Date

2023

Peer reviewed

How communicative efficiency and social biases shape language in autistic and allistic learners

L. E. F. Fletcher (l.e.f.fletcher@sms.ed.ac.uk)
CDT Natural Language Processing and Centre for Language Evolution
University of Edinburgh, 3 Charles Street, Edinburgh, EH8 9AD, UK

Jennifer Culbertson (jennifer.culbertson@ed.ac.uk)
Centre for Language Evolution
University of Edinburgh, 3 Charles Street, Edinburgh, EH8 9AD, UK

Hugh Rabagliati (hugh.rabagliati@ed.ac.uk)
Department of Psychology
University of Edinburgh, 3 Charles Street, Edinburgh, EH8 9AD, UK

Abstract

In natural languages and in experimental studies of artificial language learning, case marking of grammatical arguments is more likely to be used in languages with flexible word order due to an efficiency trade-off between production effort and communicative accuracy. However, experimental evidence suggests that language learners are less efficient when there is a social bias in favour of a group whose productions are inefficient. Here, we examine the impact of autistic traits on efficient communication. We find that autistic people's use of case in the absence of a social bias is comparative to their neurotypical peers. However, we also find evidence that autistic people adhere more to social biases; they increase production effort in order to behave more like the group they are biased towards. We argue that some autistic people may be more likely to adhere to a social bias as a result of learnt social behaviours. More generally, these results underscore the importance of studying more diverse populations in language evolution research.

Keywords: language evolution; neurodivergence; autism; artificial language learning; social biases; cultural evolution

Introduction

The trade-off between minimising production effort and maximising robust communication is argued to be a key driving force in shaping the features of language (Fedzechkina, Jaeger, & Newport, 2012; Kurumada & Jaeger, 2015; Fedzechkina & Jaeger, 2020). For example, this trade-off has been used to explain the inverse correlation between whether a language has a fixed word order, and whether it uses case to mark grammatical role assignment (Sinnemäki, 2008). Both artificial language learning studies and large-scale typological analyses have suggested that this inverse correlation is related to the trade-off between minimising production effort and maximising robust communication (Fedzechkina, Newport, & Jaeger, 2017; Koplenig, Meyer, Wolfer, & Müller-Spitzer, 2017; Levshina, 2021; Fedzechkina et al., 2012). Specifically, languages typically do not have case when they also have a fixed word order, due to the fact that the fixed word order alone is enough of a cue to grammatical role, and producing case requires extra effort. Likewise, languages typically have more case when they also have a more flexible word order, as case is required to understand the grammatical roles of the sentence, despite the extra production effort.

Recently, a set of experiments have used artificial language learning to examine the impact of wider social factors on this trade-off (Roberts & Fedzechkina, 2018; Fedzechkina & Roberts, 2020). Their studies show that, in the absence of any bias, participants trained on a language with fixed word order and (redundant) object case use case significantly less often than it is present in the input. At the same time, participants trained on a language with flexible word order and (informative) case use case significantly *more* than the input. In other words, learners restructure the language to make it more efficient (Fedzechkina et al., 2012, 2017). When social biases are introduced, they influence this efficiency. For example, they find that a social bias in favour of a group that uses redundant case leads learners to produce significantly more redundant case. Similarly, a social bias in favour of a group that does not use informative case leads learners to produce significantly less informative case. In other words, individuals are willing to sacrifice production efficiency or robust communication as a result of a social bias.

As with most work in experimental psychology and cognitive science, the vast majority of experiments exploring how biases like communicative efficiency shape language are focused on relatively homogeneous groups of neurotypical adults. However, there are large numbers of neurodivergent people, many of whom are fully immersed in a language community. Neurodivergent individuals therefore have the potential both to adapt to and to change the language used around them by neurotypical individuals. Thus, whether neurodivergent individuals display the same trade-off between reducing effort and maximising communication could have consequences for how languages look both within neurodivergent communities and the wider language-using community. In this paper, we specifically focus on autistic individuals.

Autism Spectrum Disorder (ASD) is formally characterised as a highly heterogeneous neurodevelopmental social-communicative disorder (American Psychiatric Association, 2013). The NHS and WHO estimate that approximately 1/100 people are autistic, though more recent studies on school children in England suggest that 1/57 children are diagnosed autistic (Roman-Urrestarazu et al., 2021), with a further unknown percentage of children likely to be undiag-

2825

nosed. Autistic individuals experience the world and social situations in different ways than allistic individuals[1]. Autistic people typically do not prefer person-first language, such as 'person with autism' (Botha, Hanlon, & Williams, 2021; Kenny et al., 2016; Bury, Jellett, Spoor, & Hedley, 2020). As such, 'autistic people' and 'autistic individuals' will be the preferred terms in this paper.

Language difficulties of a variety of types are common in individuals with ASD (Levy et al., 2010). Language differences in autistic individuals have been associated with pragmatic and communication differences, which are characteristic issues of ASD (American Psychiatric Association, 2013). Differences in pragmatics may persist throughout the individual's lifespan (Eigsti, Marchena, Schuh, & Kelley, 2011), but are highly heterogeneous.

Importantly, because autistic people have difficulties with social communication, their behavior in experiments investigating how efficiency and social bias shape language may differ from the behavior of allistic people. Here we replicate and expand the findings of Roberts and Fedzechkina (2018); Fedzechkina and Roberts (2020) with both neurotypical and autistic populations. These studies target the impact of social biases on the trade-off between production effort and robust communication, and it is unknown if this trade-off impacts autistic people's language productions in the same way it does neurotypical people's. Indeed, it is unclear whether autistic people behave in the same way in artificial language learning experiments at all (although see Obeid, Brooks, Powers, & Lum, 2016; Brown, Aczel, Jiménez, Kaufman, & Grant, 2010). Therefore, we first aim to establish whether autistic people show a tendency to restructure non-efficient input languages—i.e., to reduce redundant case, and maintain or increase informative case—by partially replicating Roberts and Fedzechkina (2018); Fedzechkina and Roberts (2020).

Secondly, we aim to replicate Roberts and Fedzechkina (2018); Fedzechkina and Roberts (2020)'s social bias conditions in the autistic population. We predict that autistic people will behave differently to their neurotypical peers in the presence of social biases. We entertain two hypotheses as to what this might look like. First, they may pay less attention to the social bias, on the basis of the general characterisation of autism as a deficit in socio-communication (e.g. American Psychiatric Association (2013)). However, autistic people may in fact pay *more* attention to the social bias, as a consequence of their experience paying increased attention to social cues in order to compensate and function socially. Various studies have examined the role of 'masking' or 'camouflaging' in autistic people (see, e.g., (Pearson & Rose, 2021)). Masking entails the minimisation of the appearance of autistic traits. While non-autistic people perform some similar behaviours, such as altering their communication to be more like their partner's, due to the desire to 'fit in' socially, masking for autistic people can be considered a

matter of survival, as differences in social behaviour can lead to them being marginalised in society (Ai, Cunningham, & Lai, 2022; Cook, Hull, Crane, & Mandy, 2021; Miller, Rees, & Pearson, 2021). Autistic people employ a range of masking behaviours in communicative situations, such as imitating speech and body language (Hull et al., 2017; Miller et al., 2021). If autistic participants employ these strategies in our study, they may imitate the speech of the target population even more closely, and therefore produce less efficient output than allistic participants.

Experiment 1: Efficient restructuring of language in autistic people

Design and Materials

Participants   A total of 112 participants were recruited via Prolific[2]. Participants were required to indicate their consent by a button press. Participants were aged 18 or over, and were self-reported native speakers of English.

50 participants reported to Prolific that they met at least one of the following criteria: diagnosed with ASD as a child, diagnosed with ASD as an adult, currently undergoing an ASD diagnosis, or self-identify as being autistic[3]. We allowed for self-identification due to the fact that many autistic people are unlikely to be formally diagnosed, such as older adults (Stagg & Belcher, 2019), women (Green, Travers, Howe, & McDougle, 2019), and people of colour (Steinbrenner et al., 2022; Imm, White, & Durkin, 2019). Further, some autistic people do not seek diagnosis due to the possibility of discrimination as a result of a diagnosis (for discussion of discriminations faced by autistic people, see Aylward, Gal-Szabo, and Taraman (2021); Shkedy, Shkedy, and Sandoval-Norton (2021); Griffiths et al. (2019); Nicolaidis et al. (2015); Romualdez, Heasman, Walker, Davies, and Remington (2021); Jury, Perrin, Desombre, and Rohmer (2021)).

62 participants reported to Prolific that they did not have a diagnosis of ASD, did not identify as being autistic, or didn't know. The participants were compensated £8/ hour for their participation. The data of 14 participants (12 of whom were self-reported not to be autistic) was discarded due to a technical error, leaving a total of 98 participants whose data was used for analysis.

Materials   Stimuli consisted of pictures of individual referents (chef, burglar, clown, police) and scenes of events (kick, hug) involving two referents, written descriptions in the artificial language, and synthesised speech (generated with Amazon Polly using the voice 'British English Amy'). The vocabulary consisted of 7 words: 4 nouns, 2 verbs, and 1 object case marker. Each noun referent and verb action was paired

---

[1]The term 'allistic' is a neologism coined by autistic individuals to describe non-autistic people.

[3]The number of participants who reported during our experimental questionnaire that they were autistic, or who scored 6 or above on the AQ-10, was below the number of participants recruited who reported to Prolific that they were autistic on Prolific's screening measures

randomly with a non-word from the set: {*koofta, rizba, peza, barsa, velmik, tegud*}. The object marker was always *di*. The referents, actions, and non-words are a subset of those used in Roberts and Fedzechkina (2018); Fedzechkina and Roberts (2020)[4]. All possible three-content word sentences, both with and without case, were generated, with the restriction that the object and subject could not be the same referent. For the resulting 48 sentences, S(ubject)O)bject)V(erb) and OSV versions were created. Each set was pseudo-randomly split into exposure and tests sets, with the aim of keeping vocabulary spread even across the two sets.

There were two conditions in this study. First, the redundant case condition, in which the language had a fixed SOV word order, making the object case marking unnecessary to understand the language. Second, the informative case condition, in which the language had flexible SOV (50%) and OSV (50%) word order, making object case marking necessary to understand the language. These two conditions replicate those in Roberts and Fedzechkina (2018).

The language was presented as being used by two groups of different coloured aliens (blue and orange), who each appeared for 50% of trials. One colour of aliens used case, whilst the other did not, in both conditions. In both conditions, therefore, object case marking in the input was used in 50% of the sentences.

Procedure The experiment was presented in participants' browser using jsPsych (de Leeuw, 2015).

Participants were first introduced to the task, and to the aliens; participants were told that two groups of aliens lived on the same planet, and that we were keen to trade with both groups of aliens as they had important resources and were friendly. They were then taught the noun lexicon (see Figure 1 for example trials). In the noun exposure phase, on each trial, a picture of a referent was displayed along with its noun label (presented orthographically and auditorily). Participants were able to view the word-image pair for as long as they wished, and trials were progressed by a button press. Each referent was displayed 2 times for a total of 8 trials.

In the noun comprehension phrase, on each trial, four pictures (all referents) appeared along with one of the nouns in the artificial language. Participants were asked to click on the picture that corresponded to the noun. Participants were given feedback on their answers. Each noun was displayed 2 times, for a total of 8 trials.

In the noun production phase, on each trial, one referent was presented, along with all four nouns in the artificial language. Participants were asked to click on the noun that corresponded to the picture. Participants were given feedback on their answers. Each noun was presented twice in a random order. In order to proceed to the next phase, participants were required to achieve at least 80% accuracy in the noun production stage. Participants were informed of their accu-

racy at the end of the block of trials, and either told that they would move to the next part of the study, or repeat the noun phases again. Each referent was displayed 2 times, for a total of 8 trials.
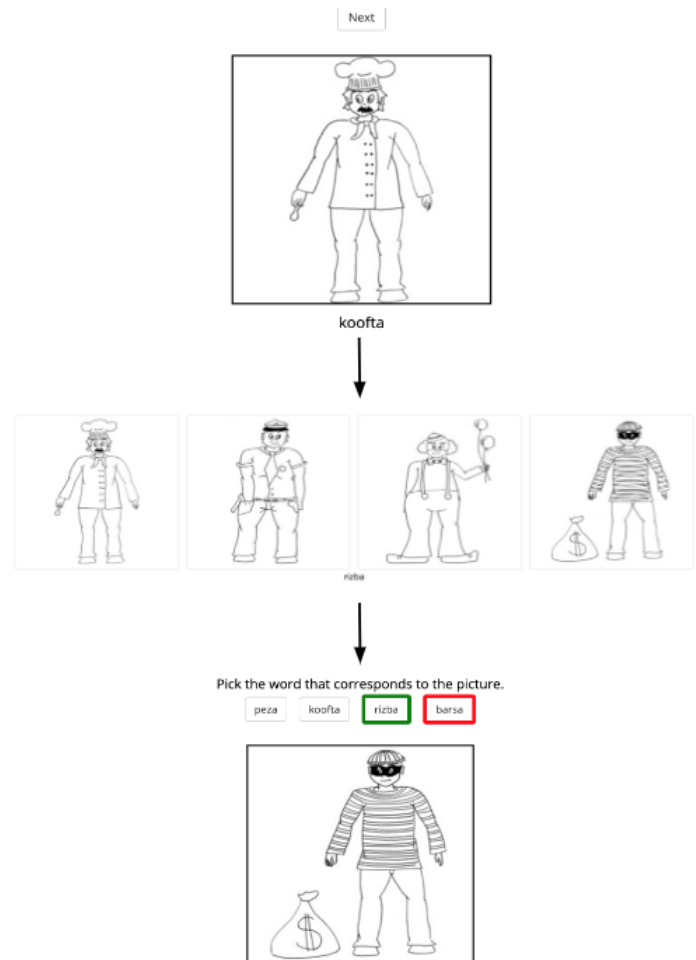

Figure 1: Example noun learning trials.

Participants were then trained and tested on sentences in the language (see Figure 2 for example trials). In the sentence exposure phase, on each trial an image of an event involving two of the referents was displayed along with a sentence (presented orthographically and auditorily). Alongside each image, an alien of the group that would use this sentence (dependent on the presence of case) was shown on screen. Each sentence was displayed 2 times, for a total of 48 trials.

In the sentence comprehension phase, on each trial two event images appeared along with a sentence (presented orthographically and auditorily), and the corresponding alien who would use it. Participants were asked to select the image that corresponded to the sentence. Participants were given feedback in this phase. Each pair was displayed twice, for a total of 48 trials.

In the sentence production phase, on each trial a single event image appeared along with buttons for each lexical item

in the language. Participants were asked to construct sentences by selecting the words in order. Participants were required to use at least three words (the minimum required to form a correct sentence without case marking). They were able to reset their choices and start again as many times as they wished for each sentence. Images were always of a previously unseen referents-event pair. No alien appeared in these trials, following (Roberts & Fedzechkina, 2018; Fedzechkina & Roberts, 2020). There were a total of 24, previously unseen, trials.


Figure 2: Example sentence learning trials. The alien shown represents one of the two species participants are introduced to; one who uses case, and one who does not (the alien species in this example does use case).

After the study, participants were asked to complete the Autism Spectrum Quotient 10 (AQ-10), a screening tool used to make quick screening decisions for autism diagnosis referrals on the NHS (Allison, Auyeung, & Baron-Cohen, 2012; NICE, 2021). The AQ-10 consists of the 10 most discriminating (in terms of distinguishing between an autistic and an allistic individual) questions from the longer Autism Spectrum

Quotient (AQ) (Baron-Cohen, Wheelwright, Skinner, Martin, & Clubley, 2001), retaining predictive validity, sensitivity, and specificity (Allison et al., 2012). The AQ-10 consists of 10 short statements with four choices of answer for each question: 'definitely agree', 'slightly agree', 'slightly disagree', and 'definitely disagree'[5].

We used the AQ-10 to measure autistic traits in a continuous manner to reflect the fact that autism is highly heterogeneous. It is important to note that we are not suggesting that all individuals within this study (or indeed, the general population) are autistic, just with various degrees of traits. Rather, we acknowledge that within the group of autistic individuals, there is a huge array of different manifestations of autistic traits, and some individuals will require more support in some areas than other. As such, it is possible that using a binary measure simply of autistic versus allistic would obscure the effects that individual autistic traits have on performance in our studies. The AQ-10's short length allows us to gain an insight into the impact of autistic traits within a limited experimental setting, rather than enforcing a single, binary definition of what an autistic person looks like.

Finally, participants completed a short questionnaire reporting their native language, any other language experience, age, gender, level of motivation during the experiment, and whether they themselves would use their data as a researcher. Finally, they were asked whether they had ever received a diagnosis of ASD, whether they identified as being autistic, and whether they had a close family history of ASD.

Results

All analyses reported here were conducted in R (R Core Team, 2022), using a combination of methods including binomial generalised linear mixed effects models with the lme4 package (Bates, Machler, Bolker, & Walker, 2015). All models were fitted with a random by-participant intercept. Contrasts were sum-coded and AQ-10 scores were standardised to scale between -2 and 2.

First, the overall use of redundant and informative case was examined to determine whether our results replicate those of previous studies (Roberts & Fedzechkina, 2018; Fedzechkina & Roberts, 2020).

We fit two intercept-only models predicting the use of the case marker in the redundant case and informative case conditions respectively. These models indicated that participants in the redundant case condition used case at a rate significantly below chance ($\beta = -1.75, p < 0.05$). By contrast, participants in the informative case condition used case numerically more than 50% of the time, but this was not significantly above chance ($\beta = 0.42, p > 0.05$). A third model predicting case by condition revealed a significant difference between the two conditions such that participants in the redundant case condition were significantly less likely to use case ($\beta = -0.95, p < 0.05$).

---

[5]It should be noted that these choices are treated as a binary choice of agree and disagree, with no scoring difference between selecting 'slightly' or 'definitely' for each option.

Figure 3 shows the impact of autistic traits, as measured by the AQ-10, on the use of redundant and informative case. In line with our main study goals, we fit a model predicting case use by AQ-10 score. This model revealed no significant impact of AQ-10 score on the use of case marking ($\beta = -0.41, p > 0.1$). A model including both AQ-10 score and condition as fixed-effects revealed no significant interaction between AQ-10 score and condition ($\beta = -0.40, p > 0.05$).

To summarise, the results of this study largely replicate previous findings. In the absence of any social bias favouring one alien group over the other, participants use case differently depending on whether word order is fixed or flexible, with less case used when it is redundant (i.e., word order is fixed), and more case used when it is informative (i.e., word order is flexible). Further, we found no indication that case-marking behaviour was impacted by AQ-10 score.
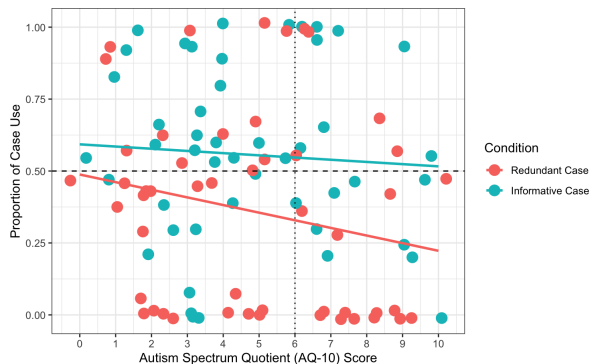

Figure 3: Proportion of case use in both conditions by AQ-10 score. The dashed horizontal line cutting the y axis indicates the proportion of case use in the input. The dotted vertical line cutting the x axis represents the cut-off point on the AQ-10 for consideration for referral for diagnosis on the NHS.

Experiment 2: Examining the role of social biases in autistic people

Design and Materials

Participants A total of 125 participants were recruited via Prolific[6]. Participants were required to indicate their consent by a button press. Participants were aged 18 or over, and were self-reported native speakers of English. 75 participants reported to Prolific that they met at least one of the following criteria: diagnosed with ASD as a child, diagnosed with ASD as an adult, currently undergoing an ASD diagnosis, or self-identify as being autistic[7]. 50 participants reported to Prolific that they did not have a diagnosis of ASD, did not identify

---

[6]This study was granted ethical approval by the University of Edinburgh PPLS Ethics Committee (reference number 209-2122/2)

[7]The number of participants who reported during our experimental questionnaire that they were autistic, or who scored 6 or above on the AQ-10, was below the number of participants recruited who reported to Prolific that they were autistic on Prolific's screening measures, despite over-sampling into this population

as being autistic, or didn't know. The participants were compensated £8/hour for their participation. The data of 5 participants was discarded due to a technical error, leaving a total of 120 participants whose data was used for analysis.

Materials The materials used were the same as in Experiment 1. The key difference in Experiment 2 is that participants were biased towards one of the two colours of aliens, based on two of Roberts and Fedzechkina (2018); Fedzechkina and Roberts (2020)'s bias conditions. In the first condition, participants were biased towards the aliens who *do use redundant case* (Roberts & Fedzechkina, 2018) (this condition is henceforth referred to as 'bias for redundant case'). In the second condition, participants were biased towards the aliens who *do not use informative case* (this condition is henceforth referred to as 'bias for no informative case'). Details of the biasing procedure can be found below.

In the bias for redundant case condition, word order was fixed (100% SOV). In the bias for no informative case condition, word order was flexible (50% SOV, 50% OSV). In both conditions, one colour of aliens used case, while the other did not, leading to 50% case input overall.

Procedure The experiment phases were the same as described in Experiment 1, with two exceptions. First, before the beginning of the noun exposure phase, participants were introduced to the intended social bias. This page indicated that one of the two alien species (identified by colour) had more to offer in trade and were friendlier, whilst the other species had less to offer and was not friendly. The alien species that the participants were biased towards depended on condition. The text used was taken directly from Roberts and Fedzechkina (2018); Fedzechkina and Roberts (2020).

Second, before the beginning of the sentence exposure phase, participants were asked to indicate which one of the aliens was considered to have more to offer and to be more friendly, by clicking on the appropriate alien. This functioned both as an attention check and as a means to reinforce the bias. If participants answered correctly, their answer was highlighted in green and accompanying text told them that they were correct. If participants answered incorrectly, their answer was highlighted in red, the correct answer was highlighted in green, and they were told by accompanying text that they were wrong and that the other alien was the one that we wish to trade with. This reminder was not used in previous studies, but was added to ensure that participants had not forgotten about the social bias in the course of training. Only one participant failed this check.

Results

First, we examined the overall use of case in the two bias conditions, to check whether our results replicate those of Roberts and Fedzechkina (2018); Fedzechkina and Roberts (2020). Recall that in their study, they found that redundant case was more likely to be retained when participants were biased in favour of the aliens who used redundant case, and

informative case was more likely to be dropped when participants were biased in favour of the aliens who did not use informative case. We fit two intercept-only models predicting the use of the case marker in the redundant and informative conditions respectively. The first model indicated that participants used case in the bias for redundant case condition significantly more than predicted by chance ($\beta = 4.45, p < 0.05$). The second model indicated that participants used case in the bias for no informative case condition significantly less than predicted by chance ($\beta = -1.10, p < 0.05$). These results replicate the findings of Roberts and Fedzechkina (2018); Fedzechkina and Roberts (2020), showing that the manipulation of social biases significantly impacts participants' case marking behaviour.

Figure 4 shows the impact of autistic traits in the bias conditions, as measured by the AQ-10. To analyse these results, we fit a model with AQ-10 score as a fixed effect on case marker usage. This model showed that there was no significant overall impact of AQ-10 score on case marking ($\beta = 0.29, p > 0.05$). However, a model including both AQ-10 score and condition as fixed-effects revealed a significant interaction between AQ-10 score driven by behaviour in the bias for redundant case condition ($\beta = 0.73, p < 0.05$).
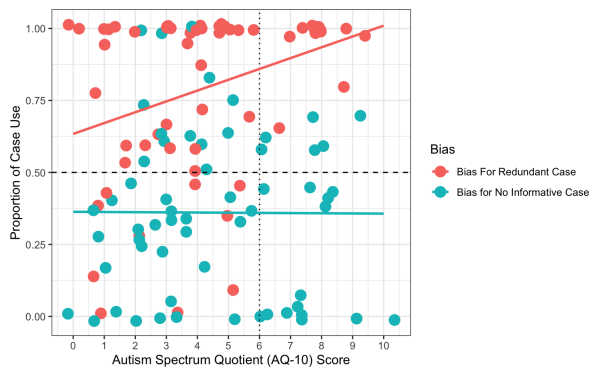

Figure 4: Proportion of case use in both conditions by AQ-10 score. The dashed horizontal line cutting the y axis indicates the proportion of case use in the input. The dotted vertical line cutting the x axis represents the cut-off point on the AQ-10 for consideration for referral for diagnosis on the NHS.

Discussion

Here we have conducted the first experiments specifically target the autistic population in an artificial language learning experiment investigating the pressures claimed to shape language evolution. We investigated whether the trade-off between minimising production effort and robust communication differs in this population. We targeted the impact of social biases on this trade-off, as social differences are well-documented in autistic individuals. We partially replicated the experiments of Roberts and Fedzechkina (2018); Fedzechkina and Roberts (2020) to achieve this goal, investigating the use of case marking in different bias conditions.

We used a continuous measure of autistic traits, the AQ-10 to explore the impact on efficiency along, and in the presence of social biases.

Our results suggest that when there is no social bias, people with higher levels of autistic traits behave similarly to their allistic peers. That is, they retain informative case, which is necessary to understand the language, and reduce their use of redundant case, which is unnecessary to understand the language. This result concurs with previous findings that statistical learning is intact in autistic people (Obeid et al., 2016; Brown et al., 2010).

However, when social biases are introduced, our results suggest that autistic people act differently than their neurotypical peers. Autistic traits were associated with an increased adherence to the social bias in the bias for redundant case condition. People with higher levels of autistic traits retained significantly more redundant case when biased towards a population that used it, despite it being unnecessary for communication and requiring more production effort. Autistic traits do not, however, lead to a significant reduction in the retention of informative case in the bias for no informative case condition. Participants with higher autistic traits, as with their lower autistic traits peers, do use less informative case in this condition, and participants are across the board unlikely to reach the extreme of no usage. This may reflect the fact that case is necessary to understand the language; without it, it was not possible to reliably understand which referent was the object and which was the subject. On the other hand, redundant case requires more effort to produce, but including it does not obscure understanding of the target language. These results indicate that autistic people may weigh communicative accuracy over minimising production effort when these two trade off.

However, it is worth noting again that autism is highly variable, and our autistic participants represent only a subsection of the autistic population as a whole. Given the nature of our recruitment method, it is likely that participants with higher levels of autistic traits are able to function in day-to-day social activities, through masking (Pearson & Rose, 2021) and other techniques. This population is thus likely to be accustomed to attempting to 'fit in' socially, and may therefore be particularly likely to pay more attention to social biases compared to participants with lower autistic traits. It is possible that autistic individuals with higher social support needs would pay less attention to social biases than their non-autistic counterparts. Regardless, our results show that autistic individuals may contribute to the evolution of language in a distinct way. This points to the importance of understanding and accounting for the role of neurodiverse populations in shaping language.

Acknowledgments

References

Ai, W., Cunningham, W. A., & Lai, M.-C. (2022). Reconsidering autistic 'camouflaging' as transactional impression management. *Trends in Cognitive Sciences*, *26*, 631-645. doi: 10.1016/J.TICS.2022.05.002

Allison, C., Auyeung, B., & Baron-Cohen, S. (2012). Toward brief "red flags" for autism screening: the short autism spectrum quotient and the short quantitative checklist for autism in toddlers in 1,000 cases and 3,000 controls [corrected]. *Journal of the American Academy of Child and Adolescent Psychiatry*, *51*, 202-212. doi: 10.1016/j.jaac.2011.11.003

American Psychiatric Association, D. S. M. T. F. (2013). *Diagnostic and statistical manual of mental disorders : Dsm-5*. Arlington, VA: American Psychiatric Association.

Aylward, B. S., Gal-Szabo, D. E., & Taraman, S. (2021). Racial, ethnic, and sociodemographic disparities in diagnosis of children with autism spectrum disorder. , *42*(8), 682–689. doi: 10.1097/DBP.0000000000000996

Baron-Cohen, S., Wheelwright, S., Skinner, R., Martin, J., & Clubley, E. (2001). The autism-spectrum quotient (aq): Evidence from asperger syndrome/high-functioning autism, males and females, scientists and mathematicians. *Journal of Autism and Developmental Disorders*, *31*, 5-17. doi: 10.1023/A:1005653411471/METRICS

Bates, D., Machler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*.

Botha, M., Hanlon, J., & Williams, G. L. (2021). Does language matter? identity-first versus person-first language use in autism research: A response to vivanti [Journal Article]. *Journal of autism and developmental disorders*, 1-9. doi: 10.1007/s10803-020-04858-w

Brown, J., Aczel, B., Jiménez, L., Kaufman, S. B., & Grant, K. P. (2010). Intact implicit learning in autism spectrum conditions. *Quarterly Journal of Experimental Psychology*, *63*, 1789-1812. doi: 10.1080/17470210903536910

Bury, S. M., Jellett, R., Spoor, J. R., & Hedley, D. (2020). "it defines who i am" or "it's something i have": What language do [autistic] australian adults [on the autism spectrum] prefer? [Journal Article]. *J Autism Dev Disord*. doi: 10.1007/s10803-020-04425-3

Cook, J., Hull, L., Crane, L., & Mandy, W. (2021). Camouflaging in autism: A systematic review. *Clinical Psychology Review*, *89*, 102080. doi: 10.1016/J.CPR.2021.102080

de Leeuw, J. R. (2015). jsPsych: a JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, *47*, 1–12. doi: 10.3758/s13428-014-0458-y

Eigsti, I.-M., Marchena, A. B., Schuh, J. M., & Kelley, E. (2011). Language acquisition in autism spectrum disorders: a developmental review. *Research in Autism Spectrum Disorders*, *5*, 681-691. doi: 10.1016/j.rasd.2010.09.001

Fedzechkina, M., & Jaeger, T. F. (2020). Production efficiency can cause grammatical change: Learners deviate from the input to better balance efficiency against robust message transmission. *Cognition*, *196*, 104115. doi: 10.1016/J.COGNITION.2019.104115

Fedzechkina, M., Jaeger, T. F., & Newport, E. L. (2012). Language learners restructure their input to facilitate efficient communication. *Proceedings of the National Academy of Sciences of the United States of America*, *109*, 17897-17902. doi: 10.1073/PNAS.1215776109

Fedzechkina, M., Newport, E. L., & Jaeger, T. F. (2017). Balancing effort and information transmission during language acquisition: Evidence from word order and case marking. *Cognitive Science*, *41*, 416-446. doi: 10.1111/COGS.12346

Fedzechkina, M., & Roberts, G. (2020). Learners sacrifice robust communication as a result of a social bias. In *Proceedings of the 42nd annual conference of the cognitive science society.* Austin, TX: Cognitive Science Society.

Green, R. M., Travers, A. M., Howe, Y., & McDougle, C. J. (2019). Women and autism spectrum disorder: Diagnosis and implications for treatment of adolescents and adults. *Current Psychiatry Reports*, *21*(4), 22. doi: 10.1007/s11920-019-1006-3

Griffiths, S., Allison, C., Kenny, R., Holt, R., Smith, P., & Baron-Cohen, S. (2019). The vulnerability experiences quotient (VEQ): A study of vulnerability, mental health and life satisfaction in autistic adults. , *12*(10), 1516–1528. doi: 10.1002/aur.2162

Hull, L., Petrides, K. V., Allison, C., Smith, P., Baron-Cohen, S., Meng, , . . . Mandy, W. (2017). "putting on my best normal": Social camouflaging in adults with autism spectrum conditions. , *47*, 2519-2534. doi: 10.1007/s10803-017-3166-5

Imm, P., White, T., & Durkin, M. S. (2019). Assessment of racial and ethnic bias in autism spectrum disorder prevalence estimates from a US surveillance system. , *23*(8), 1927–1935. doi: 10.1177/1362361319827510

Jury, M., Perrin, A.-L., Desombre, C., & Rohmer, O. (2021). Teachers' attitudes toward the inclusion of students with autism spectrum disorder: Impact of students' difficulties. , *83*, 101746. doi: 10.1016/j.rasd.2021.101746

Kenny, L., Hattersley, C., Molins, B., Buckley, C., Povey, C., & Pellicano, E. (2016). Which terms should be used to describe autism? perspectives from the uk autism community [Journal Article]. *Autism*, *20*(4), 442-62. doi: 10.1177/1362361315588200

Koplenig, A., Meyer, P., Wolfer, S., & Müller-Spitzer, C. (2017). The statistical trade-off between word order

and word structure – large-scale evidence for the principle of least effort. *PLOS ONE*, *12*, e0173614. doi: 10.1371/JOURNAL.PONE.0173614

Kurumada, C., & Jaeger, T. F. (2015). Communicative efficiency in language production: Optional case-marking in japanese. *Journal of Memory and Language*, *83*, 152-178. doi: 10.1016/J.JML.2015.03.003

Levshina, N. (2021). Communicative efficiency and differential case marking: a reverse-engineering approach. *Linguistics vanguard : multimodal online journal*, *7*. doi: 10.1515/LINGVAN-2019-0087

Levy, S. E., Giarelli, E., Lee, L. C., Schieve, L. A., Kirby, R. S., Cunniff, C., ... Rice, C. E. (2010). Autism spectrum disorder and co-occurring developmental, psychiatric, and medical conditions among children in multiple populations of the united states. *Journal of Developmental and Behavioral Pediatrics*, *31*, 267-275. doi: 10.1097/DBP.0B013E3181D5D03B

Miller, D., Rees, J., & Pearson, A. (2021). "masking is life": Experiences of masking in autistic and nonautistic adults. *https://home-liebertpub-com.ezproxy.is.ed.ac.uk/aut*, *3*, 330-338. doi: 10.1089/AUT.2020.0083

NICE. (2021, June). *Autism spectrum disorder in adults: diagnosis and management.* https://www.nice.org.uk/guidance/CG142.

Nicolaidis, C., Raymaker, D. M., Ashkenazy, E., McDonald, K. E., Dern, S., Baggs, A. E., ... Boisclair, W. C. (2015). "respect the way i need to communicate with you": Healthcare experiences of adults on the autism spectrum. , *19*(7), 824–831. doi: 10.1177/1362361315576221

Obeid, R., Brooks, P. J., Powers, G.-L. K., K. L., & Lum, J. A. (2016). Statistical learning in specific language impairment and autism spectrum disorder: A meta-analysis. *Frontiers in psychology*, *7*, 1245.

Pearson, A., & Rose, K. (2021). A conceptual analysis of autistic masking: Understanding the narrative of stigma and the illusion of choice. *Autism in Adulthood: Challenges and Management*, *3*, 52. doi: 10.1089/AUT.2020.0043

R Core Team. (2022). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from https://www.R-project.org/

Roberts, G., & Fedzechkina, M. (2018). Social biases modulate the loss of redundant forms in the cultural evolution of language. *Cognition*, *171*.

Roman-Urrestarazu, A., van Kessel, R., Allison, C., Matthews, F. E., Brayne, C., & Baron-Cohen, S. (2021). Association of race/ethnicity and social disadvantage with autism prevalence in 7 million school children in england [Journal Article]. *JAMA Pediatrics*, *175*(6), e210054. doi: 10.1001/jamapediatrics.2021.0054

Romualdez, A. M., Heasman, B., Walker, Z., Davies, J., & Remington, A. (2021). "people might understand me better": Diagnostic disclosure experiences of autistic individuals in the workplace. , *3*(2), 157–167. doi:

10.1089/aut.2020.0063

Shkedy, G., Shkedy, D., & Sandoval-Norton, A. H. (2021). Long-term ABA therapy is abusive: A response to gorycki, ruppel, and zane. , *5*(2), 126–134. Retrieved from https://doi.org/10.1007/s41252-021-00201-1 doi: 10.1007/s41252-021-00201-1

Sinnemäki, K. (2008). Complexity trade-offs in core argument marking. , 67-88. doi: 10.1075/SLCS.94.06SIN

Stagg, S. D., & Belcher, H. (2019). Living with autism without knowing: receiving a diagnosis in later life. *Health Psychology and Behavioral Medicine*, *7*(1), 348-361. doi: 10.1080/21642850.2019.1684920

Steinbrenner, J. R., McIntyre, N., Rentschler, L. F., Pearson, J. N., Luelmo, P., Jaramillo, M. E., ... Hume, K. A. (2022). Patterns in reporting and participant inclusion related to race and ethnicity in autism intervention literature: Data from a large-scale systematic review of evidence-based practices. *Autism*, 1-15. doi: 10.1177/13623613211072593