


Daniel Karczewski

 University of Białystok

 dakar@uwb.edu.pl

 <https://orcid.org/0000-0001-8255-6018>

Generic overgeneralization effect or quantifier domain restriction? Evidence from paraphrasing in Polish

Abstract

This study examines how native speakers of Polish paraphrase two types of quantified statements: (i) minority statements (e.g., *Wszystkie kangury mają torby* ‘All kangaroos have pouches’), that is, false universal generalizations about properties possessed by only a subset of a category, and (ii) statements involving pragmatically implied quantifier domain restriction (e.g., *Wszyscy przynieśli prezenty dla Hani* ‘Everyone brought presents for Hania’), where the quantifier is understood to apply to a limited set. The aim is to contribute to the ongoing discussion on overgeneralization bias, that is, the tendency of adults to endorse false universal statements. Two competing accounts have been proposed for this phenomenon: the generic overgeneralization effect hypothesis (Leslie et al., 2011) and the quantifier domain restriction hypothesis (Lazaridou-Chatzigoga et al., 2017, 2019). The analysis of paraphrases showed that participants rarely referred to a restricted subset of the category. Instead, they either retained a strong universal quantifier or omitted the quantifier in their paraphrases. These findings lend support to the generic overgeneralization effect hypothesis over the domain restriction account.

Keywords

false universals, generic overgeneralization effect, minority statements, paraphrase, quantifier domain restriction.

Abstrakt

Artykuł jest poświęcony analizie, w jaki sposób rodzimi użytkownicy języka polskiego parafrazują dwa typy zdań z kwantyfikatorem: (i) fałszywe zdania ogólne, w których cechę reprezentatywną przypisuje się jedynie części kategorii (np. *Wszystkie kangury*

mają torby) oraz (ii) zdania z kwantyfikacją restryktywną (np. *Wszyscy przynieśli prezenty dla Hani*), w których kwantyfikator jest interpretowany jako odnoszący się tylko do ograniczonego zbioru. Badanie wpisuje się w dyskusję na temat błędu nadmiernego uogólnienia, czyli tendencji dorosłych użytkowników języka do uznawania fałszywych twierdzeń uniwersalnych za prawdziwe. W literaturze zaproponowano dwie konkurencyjne hipotezy wyjaśniające to zjawisko: hipotezę o nadmiernym uogólnieniu (ang. *the generic overgeneralization effect*) (Leslie et al., 2011) oraz hipotezę dotyczącą ograniczenia zakresu nazwy (ang. *quantifier domain restriction*) (Lazaridou-Chatzigoga et al., 2017, 2019). Analiza parafraz wykazała, że uczestnicy rzadko odnosili się do ograniczonego podzbioru kategorii w swoich parafrazach. Zamiast tego posiłkowali się innym kwantyfikatorem lub całkowicie go pomijali w parafrazach. Wyniki te przemawiają na korzyść hipotezy o nadmiernym uogólnieniu w większym stopniu niż hipotezy o ograniczeniu zakresu nazwy.

Słowa kluczowe

błąd nadmiernego uogólnienia, fałszywe zdania ogólne, ograniczenie zakresu nazwy, parafraza, zdania z cechą reprezentatywną przypisywaną części kategorii.

1. Introduction

Overgeneralization bias refers to the tendency of individuals to accept false universals, such as *Wszystkie kangury mają torby* ('All kangaroos have pouches') or *Wszystkie lwy mają grzywy* ('All lions have manes') as true, even when they are aware that such statements are false. Two competing explanations have been proposed for this phenomenon: the generic overgeneralization effect hypothesis (Leslie et al., 2011) and the quantifier domain restriction hypothesis (Lazaridou-Chatzigoga et al., 2017, 2019). The former hypothesis holds that individuals endorse false universal statements because they rely on the easily accessible corresponding generic statements, such as *Kangury mają torby* ('Kangaroos have pouches') or *Lwy mają grzywy* ('Lions have manes'). The generic overgeneralization effect hypothesis provides support for the generics-as-default hypothesis (Gelman, 2010; Leslie, 2007, 2008), which posits that generic statements reflect a cognitively basic and easily accessible mode of generalization. In this view, comprehending and producing generic generalizations relies on a low-effort cognitive mechanism, while quantified statements require more deliberate and analytical processing. Consequently, people may conflate a true generic such as *Kangury mają torby* ('Kangaroos have pouches') with its false universal counterpart *Wszystkie kangury mają torby* ('All kangaroos have pouches') in situations where analytical reasoning is not engaged. The generics-as-default hypothesis therefore predicts the overgeneralization error observed among adult language users. Lazaridou-Chatzigoga and her colleagues (2017, 2019), by contrast, argue that quantifier domain restriction is the main factor

contributing to overgeneralization bias. According to this account, individuals accept statements such as *Wszystkie kangury mają torby* ('All kangaroos have pouches') because they mentally restrict the domain of kangaroos to female kangaroos only. They also caution against overlooking other possible explanations that may contribute to overgeneralization bias, such as ignorance of facts, sub-kind interpretation and atypical behaviour of the quantifier 'all'.

So far, this tendency has been investigated using different methodological designs across three languages: Greek, English and Polish (Karczewski et al., 2020; Lazaridou-Chatzigoga et al., 2019; Leslie et al., 2011). The results of these studies have consistently shown similar patterns, yet the proponents of the two competing views offer different interpretations of the same phenomenon. This article aims to contribute to the ongoing discussion of overgeneralization bias by analysing how native speakers of Polish paraphrase two types of quantified statements: minority statements¹ and statements involving pragmatically implied quantifier domain restriction. Whereas the latter type is widely acknowledged as involving contextual narrowing, Lazaridou-Chatzigoga et al. (2019) also treat minority statements as cases involving quantifier domain restriction.

The present study tests this claim by comparing paraphrasing strategies across both sentence types. The analysis is based on data from Polish, whose configuration of nominals involved in genericity and definiteness differs from that of English and Greek. Following Lazaridou-Chatzigoga et al. (2019), who emphasize the importance of studying genericity and quantification cross-linguistically, Polish provides a testing context for investigating overgeneralization bias in a language that does not mark definiteness or genericity overtly (Grochowski, 2018; Grzegorzczkowska, 2010; Karczewski, 2016; Świączkowska, 2004; Topolińska, 1976, 1977, 1981).

The following section will discuss some major theoretical issues concerning overgeneralization bias and will review some of the findings stemming from the previous studies. The remainder of the paper describes the experimental design, presents the results and discusses their implications for the competing accounts.

¹ See, e.g., Leslie (2007, 2008) and Leslie et al. (2011) for different types of generics. A summary of this taxonomy is offered in Karczewski (2016).

2. The generic overgeneralization effect and competing explanations for overgeneralization bias

According to Leslie et al. (2011, p. 17), the generic overgeneralization effect (GOG) “involves overgeneralizing from the truth of a generic to the truth of the corresponding universal statement”. This means that individuals endorse false universals, such as *Wszystkie kangury mają torby* (‘All kangaroos have pouches’), because they conflate these with generic statements such as *Kangury mają torby* (‘Kangaroos have pouches’). The GOG effect provides support for the generics-as-default hypothesis (Gelman, 2010; Leslie, 2007, 2008), which posits that generics give voice to more basic, intuitive generalizations than quantified statements. According to this view, generics function as a cognitive default and are linked to System 1, which involves a fast, automatic and intuitive mode of thinking. In contrast, quantified statements are associated with System 2, which involves slower, more effortful reasoning (see Kahneman & Frederick, 2002, for a discussion of the dual system theory of cognition). Initially, support for this hypothesis came from studies on children, which showed a tendency to interpret universal and existential quantified statements as generics (Hollander et al., 2002; Mannheim et al., 2010; Tardif et al., 2012). However, recent research on large language models (Allaway et al., 2024; Ralethe & Buys, 2022) illustrates that overgeneralization bias extends to systems trained on human-produced linguistic data. Because these models reproduce patterns present in natural language, their output supports the view that the tendency to conflate generics with universals reflects a structural property of human language use rather than a task-specific effect. As shown by Allaway et al. (2024), large language models frequently generate universal quantifiers even when exceptions are implied, which parallels the human GOG effect and reinforces the generics-as-default hypothesis. This tendency is consistent with Kirkpatrick and Sterken’s (2025) observation that most formal and computational models misinterpret generics as universal statements rather than defeasible generalizations, reflecting a bias that is also evident in human language production.

In their 2011 paper, Leslie and her colleagues experimentally ruled out four alternative explanations for overgeneralization bias: (i) quantifier domain restriction (QDR), which refers to how context affects the interpretation of quantifiers in sentences such as *Wszystkie kangury mają torby* (‘All kangaroos have pouches’) as applying only to a relevant subset of kangaroos, namely female kangaroos; (ii) ignorance of facts, which refers to participants being unaware of cases that contradict the statement (for instance, not knowing that male kangaroos do not have pouches); (iii) sub-kind interpretation, which suggests that individuals interpret statements about kangaroos

as referring to all kinds of kangaroos, e.g., Red Kangaroos, Eastern Grey Kangaroos or Western Grey Kangaroos; and (iv) the atypical behaviour of the quantifier ‘all’, which may also contribute to overgeneralization bias.

Among other concerns, Lazaridou-Chatzigoga et al. (2019) questioned the validity of the method used by Leslie et al. (2011) to rule out QDR. Specifically, they challenged the practice of providing participants with population numbers (e.g., “Suppose the following is true: there are 431 million ducks in the world. Do you agree with the following: all ducks lay eggs?”) as a means of eliminating domain restriction. According to them, this approach may not be sufficient to override participants’ tendency to interpret ‘all’ as referring only to a contextually restricted subset.

By contrast, Lazaridou-Chatzigoga and her colleagues (2017, 2019) argue that QDR is the main factor underlying overgeneralization bias. QDR refers to the influence of context on the interpretation of quantifiers, as illustrated by examples such as *Wszystkim podobała się moja fryzura* (‘Everybody loved my hairdo’). Here, *wszystkim* ‘everybody’ is normally interpreted as referring to a specific group of people rather than to the entire category of people, thereby illustrating how quantifier domain restriction operates in natural language.

Building on this line of reasoning, Lazaridou-Chatzigoga et al. (2019) claim that individuals accept statements such as *Wszystkie kangury mają torby* (‘All kangaroos have pouches’) because they mentally restrict the domain of kangaroos to female kangaroos only. However, the role of context becomes particularly problematic in the case of universal generalizations with ‘all’ or ‘every’ since sentences like *Wszystkie kangury mają torby* (‘All kangaroos have pouches’) or *Każdy kangur ma torbę* (‘Every kangaroo has a pouch’) are typically interpreted to mean that every individual member of the category has a pouch. This reading stands in stark contrast to context-dependent sentences such as *Wszystkim podobała się moja fryzura* (‘Everybody loved my hairdo’).

Although Lazaridou-Chatzigoga et al. (2019) attribute overgeneralization bias to quantifier domain restriction, domain restriction may be a more appropriate label. In particular, domain restriction can occur independently of overt quantifiers.² In Polish, for example, even unquantified statements such as *Kaczki mają obcięte skrzydła* (‘The ducks have clipped wings’) undergo pragmatic restriction to contextually salient sets. By contrast,

² Evidence that generics themselves can undergo pragmatic domain restriction has been presented by Sterken (2015). She argues that generics such as *Birds lay eggs* or *Dobermans have floppy ears* are best understood as generalizations over contextually restricted subsets (for example, adult female birds), rather than as genuinely exception-tolerant statements. This observation complements the present distinction between domain restriction and quantifier domain restriction, suggesting that context sensitivity is not limited to quantified expressions.

generalizations such as *Wszystkie kaczki znoszą jaja* ('All ducks lay eggs') seem to involve a restriction based on conceptual knowledge about the kind (i.e., that only female ducks lay eggs) rather than narrowing of the quantifier domain. Irrespective of whether the phenomenon is referred to as quantifier domain restriction or domain restriction, it remains an open empirical question whether speakers actually restrict the domain when reformulating quantified statements. In the present study, this distinction is reflected in the analysis: paraphrases of minority statements are examined in terms of domain restriction, while paraphrases of pragmatically restricted statements are analysed as cases of quantifier domain restriction.

In current linguistic experiments it is not feasible to determine the cognitive processes engaged in assessing the truth value of universal generalizations. Therefore, although the explicit demonstration of the restriction in the domain impacts test scores, it does not follow that individuals engage in the process of restricting the domain beyond experimental conditions. It is important to acknowledge the significance of tacit knowledge, as emphasized by Lazaridou-Chatzigoga et al. (2017, 2019); however, if the quantifier domain restriction is claimed to be a primary cause for overgeneralization bias, it should be demonstrated by other experimental methods. To this end, the current study employs a paraphrase task whose aim is to assess whether participants restrict the domain when reformulating a generalization. While the paraphrase task parallels Experiment 2b of Leslie et al. (2011), the present study extends that design in two important ways. First, it distinguishes between two kinds of quantified statements: minority statements and statements involving pragmatically implied quantifier domain restriction, whereas Leslie et al.'s study examined only the former. Second, its aim is not to replicate the GOG effect itself but to investigate the interpretive mechanisms that give rise to it, by comparing how speakers reformulate both types of statements. In Leslie et al.'s (2011) Experiment 2b, a small sample of English-speaking participants evaluated and paraphrased universal statements such as *All ducks lay eggs*. Their responses were analysed for signs of domain restriction, such as references to subtypes (e.g., female ducks or kinds of ducks). However, such paraphrases were rare. Most participants either preserved the universal form or reformulated the statement as a generic, thus providing support for the generic overgeneralization effect hypothesis.

The present study examines how native speakers of Polish interpret and reformulate two types of quantified statements that have been central to the debate on overgeneralization bias: minority statements and statements involving pragmatically implied quantifier domain restriction. Lazaridou-Chatzigoga et al. (2019) treat both types as instances of quantifier domain restriction, extending this mechanism to minority statements. The present study tests this claim by comparing how speakers paraphrase both types and

by examining whether the two kinds of statements show similar or different paraphrasing patterns.

3. Current study

3.1 Aims and research questions

Building on previous findings, the present study uses a paraphrase task to explore the mechanisms underlying overgeneralization bias. The aim was not to re-demonstrate the effect itself, which has been repeatedly shown in truth-value judgement studies (e.g., Karczewski et al., 2020; Lazaridou-Chatzigoga et al., 2019; Leslie et al., 2011), but to examine whether the two types of quantified statements elicit distinct interpretive strategies. Specifically, the study compares domain restriction in minority statements with quantifier domain restriction in pragmatically restricted statements. The following research questions guided the analysis:

RQ1: What paraphrasing strategies do speakers apply to minority statements?

RQ2: To what extent do these strategies involve domain restriction (e.g., referencing subcategories such as female kangaroos)?

RQ3: What paraphrasing strategies do speakers apply to pragmatically restricted statements?

RQ4: To what extent do these strategies involve quantifier domain restriction (e.g., narrowing the quantifier to a contextually relevant set)?

3.2 Participants

A total of 65 Polish-speaking English-philology undergraduates ($M = 20.78$ years, $SD = 1.49$, range = 18–25) from the Faculty of Philology at the University of Białystok completed the paraphrasing task. The sample was almost evenly divided between women ($n = 31$, 47.8%) and men ($n = 28$, 43%); three participants did not specify their gender ($n = 3$, 4.6%) and three identified as another gender ($n = 3$, 4.6%).

3.3 Materials and stimuli

The experiment employed 18 critical sentences in Polish, distributed across three categories: minority statements (MS), such as *Wszystkie kangury mają torby* ‘All kangaroos have pouches’; statements in the past involving

pragmatically implied quantifier domain restriction (PAST; e.g., *Wszyscy przynieśli prezenty dla Hani* ‘Everyone brought presents for Hania’); and statements in the present involving pragmatically implied quantifier domain restriction (PRES; e.g., *Wszyscy kelnerzy trzymają tace* ‘All the waiters are holding trays’). Present-tense statements were included as a comparison condition because they do not suggest any specific situation or event that could help narrow down the reference of the quantifier. Their role was to test whether participants would still attempt to restrict the domain when no such clues were available. All items were matched for length, and the full list of critical items appears in the Appendix. The 18 items were counter-balanced across three Latin-square lists so that each participant saw six critical sentences (two from each category) and six identical filler items (quantified comparative statements). Filler responses were excluded from the statistical analyses.

3.4 Procedure

Participants first read the instruction that asked them to rewrite each sentence so that its meaning was preserved without using the quantifiers *każdy/każda/każde* ‘every’. Each stimulus was printed with a blank line beneath it, and participants hand-wrote their paraphrases directly below the original sentence. The three versions of the test were distributed at random. It took approximately 15 minutes to complete the entire task. All participants provided informed oral consent, and no personal data were recorded. Participants were not instructed to treat the sentences as true or false, as the task was designed to elicit their spontaneous reformulations rather than their truth-value judgements.

3.5 Data-analysis procedure and coding scheme

The author of this paper, together with four MA students, initially coded a small practice set during a seminar. Discrepancies were discussed, and a written coding protocol was established. One of the students then applied the protocol to code the full dataset. The coded file was subsequently reviewed independently by two linguistically trained coders, including the author. All discrepancies were resolved through discussion and consensus; therefore, no statistical reliability coefficient was calculated. The paraphrases assigned to codes 1–4 were considered valid reformulations and included in the analysis. Items assigned code 5 were excluded. When a paraphrase combined two strategies, for example dropping the quantifier and restricting the domain (*Female kangaroos have pouches*), it was classified under domain restriction, since this represents the more specific

interpretive strategy. The present coding scheme followed the classification principles described by Leslie et al. (2011, Experiment 2b), who distinguished paraphrases based on whether they retained universality, reflected domain restriction, or omitted the quantifier. Table 1 presents the coding scheme, including operational definitions and examples drawn from participants' responses.

Table 1. Coding scheme for quantitative analyses.

Code	Label	Operational definition and example
1	Universal	The paraphrase retains a strong universal quantifier or a near-synonym and does not mention any narrower subcategory or condition (e.g., <i>bez wyjątku</i> 'without exception'). PL: <i>Faktem jest, że węże znoszą jaja, bez wyjątku.</i> EN: 'It is a fact that snakes lay eggs, without exception'
2	Domain restriction	The paraphrase limits the scope of the statement by referencing specific subcategories or conditions (e.g., <i>samice kangurów</i> 'female kangaroos'). PL: <i>Wszystkie samice kangurów mają torby.</i> EN: 'All female kangaroos have pouches'
3	Quantifier domain restriction	The paraphrase limits the scope of the statement by referencing specific subcategories or conditions (e.g., <i>wszyscy zaproszeni goście</i> 'all invited guests'). PL: <i>Wszyscy zaproszeni goście przynieśli prezenty dla Hani.</i> EN: 'All invited guests brought presents for Hania'
4	Dropping a quantifier	The paraphrase drops any explicit quantifier. PL: <i>Lwy mają grzywy.</i> EN: 'Lions have manes'
5	Other/unrelated	The paraphrase either changes the original meaning, includes irrelevant information, or is effectively a near-exact repetition without genuine rewording. PL: <i>Torba jest charakterystycznym elementem ciała kangura.</i> EN: 'The pouch is a characteristic part of the kangaroo's body'

Source: Author's own elaboration.

3.6 Results

In total, 284 paraphrases were collected, of which 15 (5.3%) were excluded as incomplete or ambiguous, resulting in 269 valid responses analysed in an experimental task across three statement categories: minority statements, past-tense statements involving pragmatically implied quantifier domain restriction and present-tense statements with such restriction, as detailed in Section 3.3. The paraphrase types evaluated were universal, domain restriction, quantifier domain restriction and dropping a quantifier, as defined in Table 1. The distribution of paraphrase frequencies within each category is summarized in Table 2 and illustrated in Figure 1.

Across the entire experiment, the frequency of chosen paraphrases offers information about the paraphrasing preferences of Polish speakers. The universal paraphrase was the most frequently selected, with 155 instances accounting for 57.6% of all responses across categories. This high frequency emphasizes a strong tendency among participants towards unrestricted, generalized interpretations. In contrast, the domain restriction paraphrase was the least frequent, with only eight selections, all occurring within the minority statement category. The dropping a quantifier paraphrase was selected 79 times, accounting for 29.4% of responses. The quantifier domain restriction paraphrase was chosen 27 times, with responses concentrated in past-tense contexts.

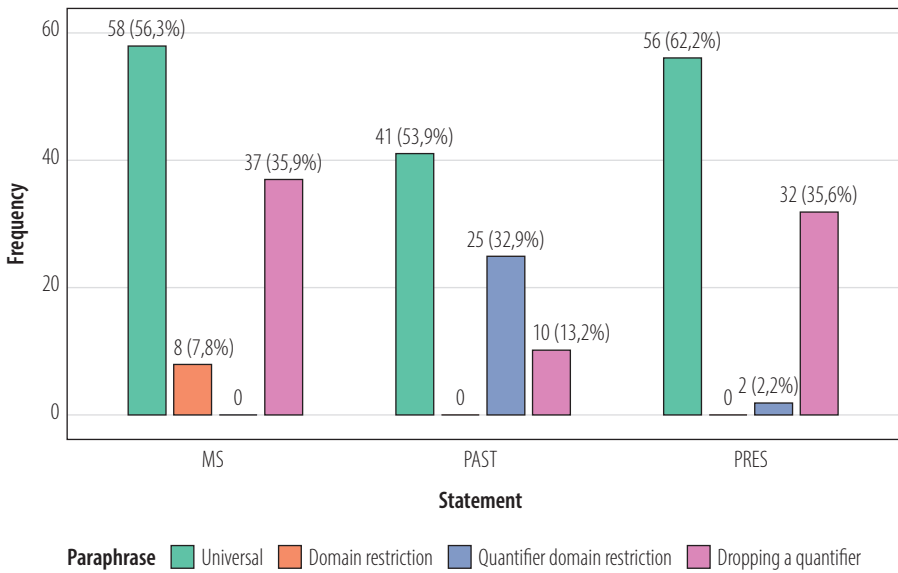


Figure 1. Distribution of paraphrase preferences across statement types.

Source: Author’s own data and elaboration.

Table 2. Analysis of paraphrase preferences among Polish speakers by statement type.

Statement type	Paraphrases				P
	Universal (N = 155)	Domain restriction (N = 8)	Dropping a quantifier (N = 79)	Quantifier domain restriction (N = 27)	
Minority (MS), N = 103	58 (56.3%)	8 (7.8%)	37 (35.9%)	—	< 0.001
Pragmatically Restricted (PAST), N = 76	41 (53.9%)	—	10 (13.2%)	25 (32.9%)	< 0.001
Pragmatically Restricted Statements (PRES), N = 90	56 (62.2%)	—	32 (35.6%)	2 (2.2%)	< 0.001

Note: Percentages are calculated based on row totals (N). p-values are derived from Fisher's exact test. Cells marked with "—" indicate that no instances were observed in this category.

Source: Author's own data and analysis.

Paraphrase strategies within minority statements

In the case of minority statements ($n = 103$), the universal paraphrase was the most common choice, selected 58 times (56.3%). The domain restriction paraphrase was chosen only eight times (7.8%), while the dropping a quantifier option was selected 37 times (35.9%). Post-hoc pairwise Fisher's exact tests showed significant differences between all paraphrase types (all $p_{adj.Fisher} < 0.001$), confirming a clear ranking: universal and dropping a quantifier paraphrases were preferred over the domain restriction paraphrases.

Paraphrase strategies within statements involving pragmatically implied quantifier domain restriction (PAST)

For pragmatically restricted statements in the past tense ($n = 76$), the universal paraphrase was the most common choice, with 41 selections (53.9%). The quantifier domain restriction paraphrase followed with 25 selections (32.9%), indicating that it plays a particularly important role in past-tense contexts.³ The dropping a quantifier paraphrase was selected 10 times (13.2%). Post-hoc pairwise Fisher's exact tests showed significant differences between all paraphrase types (all $p_{adj.Fisher} < 0.001$), confirming a clear pattern of preferences. Although the universal paraphrase remained dominant, the quantifier

³ One of the reviewers observed that the higher rate of quantifier domain restriction in past-tense statements may result from participants' use of world knowledge rather than from tense itself. This possibility does not affect the interpretation offered here, which focuses on the fact that different types of statements appear to rely on distinct interpretive mechanisms.

domain restriction paraphrase appears to play a distinct interpretive role in these contexts.

Paraphrase strategies within statements involving pragmatically implied quantifier domain restriction (PRES)

Within pragmatically restricted statements in the present tense ($n = 90$), the universal paraphrase was the most frequent choice, selected 56 times (62.2%). It was followed by the dropping a quantifier paraphrase, which was chosen 32 times (35.6%). The quantifier domain restriction paraphrase was used only twice (2.2%). Post-hoc pairwise Fisher's exact tests showed significant differences (all $p_{adj.Fisher} < 0.001$) between all groups, confirming a strong preference for universal and dropping a quantifier strategies compared to the other options.

Quantifier domain restriction across PAST and PRES conditions

Based on the aggregated data in Table 2, the quantifier domain restriction paraphrase accounted for 32.9% of paraphrases in the PAST condition (25 out of 76), compared to 2.2% in the PRES condition (2 out of 90). To assess statistical significance, Fisher's exact tests were conducted using contingency tables comparing the frequency of quantifier domain restriction for PAST versus PRES. The odds ratio was 21.57, with $p < 0.001$, indicating a significantly higher frequency of quantifier domain restriction in the PAST condition. This result provides evidence that grammatical tense influences the use of quantifier domain restriction, with the PAST condition eliciting substantially more instances of quantifier domain restriction than the PRES condition.

Frequency of domain restriction and quantifier domain restriction overall

The combined number of domain restriction and quantifier domain restriction paraphrase selections across all statement types was calculated. Out of 269 total paraphrases, eight involved domain restriction and 27 involved quantifier domain restriction, resulting in a combined total of 35 selections (13.0%). To assess whether this proportion differs meaningfully from what might be expected by chance, a binomial test was conducted. The two-sided result was not statistically significant ($p = 0.104$), and a one-sided test for a greater-than-expected proportion also did not reach significance ($p = 0.065$). The 95% confidence interval for the observed proportion ranged from 9.0 to 17.0%. These results indicate that domain restriction/quantifier domain restriction is relatively infrequent across the dataset.

Stability of paraphrasing strategies overall

The data show that many participants who preferred the universal paraphrase in minority statements continued to use it in the present-tense condition, but not as consistently in the past-tense condition. Among the 30 participants who mainly used the universal strategy in the minority statement condition, 27 (90%) also used it most frequently in the present-tense condition. The remaining three switched strategies. A statistical test confirmed that this consistency was unlikely to be due to chance ($p < 0.001$), indicating strong stability in the present-tense condition. In contrast, only 14 participants (46.7%) maintained their universal preference in the past-tense condition, while the others switched to different strategies. In this case, the statistical test did not show a significant result ($p = 0.708$), suggesting that the pattern could have occurred by chance. These findings indicate that participants' paraphrasing preferences were more stable in the present-tense context than in the past.

4. Discussion and conclusion

The aim of the study was to determine whether paraphrasing strategies applied to minority statements and to statements involving pragmatically implied quantifier domain restriction align more closely with the predictions of the generic overgeneralization effect hypothesis (Leslie et al., 2011) or with the quantifier domain restriction account (Lazaridou-Chatzigoga et al., 2017, 2019). Across the sentence types examined, participants showed a strong preference for the universal paraphrase strategy and the quantifier-dropping strategy. In contrast, domain restriction and quantifier domain restriction were used infrequently. The prevalence of the quantifier-dropping strategy, resulting in bare plural generics, provides evidence that speakers reinterpret false universals as generic statements, consistent with the generic overgeneralization hypothesis proposed by Leslie et al. (2011).

In the case of minority statements, participants rarely reformulated the statements by narrowing the scope of the noun phrase (e.g., *Wszystkie samice kangurów mają torby* 'All female kangaroos have pouches'). Domain restriction accounted for only 7.8% of paraphrases. Instead, many participants opted for the dropping a quantifier strategy, producing generic statements such as *Kangury mają torby* ('Kangaroos have pouches'). This pattern supports the claim that generic statements function as a default mode of generalization (Gelman, 2010; Leslie, 2007, 2008) and challenges the view that speakers engage in domain restriction to resolve false universal statements.

For statements involving pragmatically implied QDR, quantifier domain restriction occurred more frequently, but still did not dominate. In the past-tense condition, QDR was observed in 32.9% of paraphrases (e.g., *Wszyscy zaproszeni goście przynieśli prezenty dla Hani* ‘All invited guests brought presents for Hania’), suggesting that contextually salient cues may facilitate this interpretative strategy. In contrast, in the present-tense condition, QDR was virtually absent (2.2%), indicating that present-tense contexts are less conducive to the pragmatic narrowing of the quantifier domain. The higher frequency of quantifier domain restriction in past-tense contexts may reflect the dynamic, presuppositional nature of domain restriction (Roberts, 1995), whereby tense and event framing facilitate access to contextually restricted domains. Past-tense statements typically evoke bounded, eventive situations that foreground specific participants, while present-tense statements often receive habitual or stative readings that invite broader, more generic interpretations. This distinction aligns with Langacker’s (2008) account of aspectual construal, according to which past-tense utterances tend to profile delimited, perfective events, whereas present-tense utterances favour imperfective or habitual construals. Nevertheless, even in the past-tense condition, the universal paraphrase remained the most frequent paraphrasing strategy, followed by quantifier-dropping, with QDR representing only a minority of responses.

Taken together, these results provide little support for the claim that domain restriction is the primary strategy used to preserve the truth of minority statements, as proposed by Lazaridou-Chatzigoga et al. (2019). Instead, they align with the generic overgeneralization effect hypothesis, which holds that speakers tend to interpret false universal statements as generics (Gelman, 2010; Leslie, 2007, 2008).

While the present study contributes to the understanding of how Polish speakers paraphrase minority statements and statements involving pragmatically implied quantifier domain restriction, several limitations should be acknowledged. First, although the paraphrase task provides insight into participants’ reformulation strategies, it does not reveal the underlying reasoning processes. Future studies could complement this method with think-aloud protocols or follow-up interviews to better understand how participants interpret and evaluate the statements. Second, the design did not systematically explore item-level variation. Some statements may have elicited particular paraphrasing strategies due to differences in content familiarity or other factors. A more detailed item-level analysis would help determine whether such factors systematically influence paraphrasing choices. Third, the study included both past-tense and present-tense pragmatically restricted statements but did not analyse tense as an independent variable.

Since tense may modulate the salience of contextual restriction or influence participants' truth-value judgements, future work should examine more closely how tense interacts with quantifier interpretation. Finally, the current study focused on Polish. Given cross-linguistic differences in how quantification and definiteness are encoded, further research should investigate whether similar patterns hold in other languages, especially those with different morphosyntactic resources for marking generics and quantifiers.

References

- Allaway, E., Bhagavatula, C., Hwang, J. D., McKeown, K., & Leslie, S.-J. (2024). Exceptions, instantiations, and overgeneralization: Insights into how language models process generics. *Computational Linguistics*, 50(4), 1211–75. https://doi.org/10.1162/coli_a_00530
- Gelman, S. (2010). Generics as a window onto young children's concepts. In F. J. Pelletier (Ed.), *Kinds, things, and stuff: Mass terms and generics* (1st ed., pp. 1–22). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195382891.001.0001>
- Grochowski, M. (2018). Szyk operatorów kwantyfikacji ogólnej w języku polskim. *Slavica Wratislaviensia*, 165, 125–35. <https://doi.org/10.19195/0137-1150.165.11>
- Grzegorzczkova, R. (2010). *Wprowadzenie do semantyki językoznawczej* (Wyd. 4. rozsz.). Wydawnictwo Naukowe PWN.
- Hollander, M. A., Gelman, S. A., & Star, J. (2002). Children's interpretation of generic noun phrases. *Developmental Psychology*, 38(6), 883–94. <https://doi.org/10.1037/0012-1649.38.6.883>
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (1st ed., pp. 49–81). Cambridge University Press. <https://doi.org/10.1017/CBO9780511808098.004>
- Karczewski, D. (2016). *Generyczność w języku i w myśleniu: Studium kognitywne*. Universitas.
- Karczewski, D., Wajda, E., & Poniat, R. (2020). Do all storks fly to Africa? Universal statements and the generic overgeneralization effect. *Lingua*, 246, 102855. <https://doi.org/10.1016/j.lingua.2020.102855>
- Kirkpatrick, J. R., & Sterken, R. K. (2025). Generics and default reasoning in large language models [Preprint]. arXiv. <https://arxiv.org/abs/2508.13718>
- Langacker, R. W. (2008). *Cognitive grammar: A basic introduction*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195331967.001.0001>
- Lazaridou-Chatzigoga, D., Stockall, L., & Katsos, N. (2017). A new look at the 'Generic Overgeneralisation' effect. *Inquiry*, 66(9), 1655–81. <https://doi.org/10.1080/0020174X.2017.1285993>
- Lazaridou-Chatzigoga, D., Stockall, L., & Katsos, N. (2019). Contextualising generic and universal generalisations: Quantifier domain restriction and the generic

- overgeneralisation effect. *Journal of Semantics*, 36(4), 617–64. <https://doi.org/10.1093/jos/ffz009>
- Leslie, S.-J. (2007). Generics and the structure of the mind. *Philosophical Perspectives*, 21(1), 375–403. <https://doi.org/10.1111/j.1520-8583.2007.00138.x>
- Leslie, S.-J. (2008). Generics: Cognition and acquisition. *The Philosophical Review*, 117(1), 1–47. <https://doi.org/10.1215/00318108-2007-023>
- Leslie, S.-J., Khemlani, S., & Glucksberg, S. (2011). Do all ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Language*, 65(1), 15–31. <https://doi.org/10.1016/j.jml.2010.12.005>
- Mannheim, B., Gelman, S. A., Escalante, C., Huayhua, M., & Puma, R. (2010). A developmental analysis of generic nouns in Southern Peruvian Quechua. *Language Learning and Development*, 7(1), 1–23. <https://doi.org/10.1080/15475441003635620>
- Ralethe, S., & Buys, J. (2022). Generic overgeneralization in pre-trained language models. In *Proceedings of the 29th International Conference on Computational Linguistics* (pp. 3187–96). International Committee on Computational Linguistics.
- Roberts, C. (1995). Domain restriction in dynamic semantics. In E. Bach, E. Jelinek, A. Kratzer, & B. H. Partee (Eds.), *Quantification in natural languages* (Vol. 54, pp. 661–700). Springer Netherlands. https://doi.org/10.1007/978-94-017-2817-1_20
- Sterken, R. K. (2015). Generics, content and cognitive bias. *Analytic Philosophy*, 56(1), 75–93. <https://doi.org/10.1111/phib.12056>
- Świączkowska, H. (2004). *Semiotyka określoności*. Wydawnictwo Uniwersytetu w Białymstoku.
- Tardif, T., Gelman, S. A., Fu, X., & Zhu, L. (2012). Acquisition of generic noun phrases in Chinese: Learning about lions without an ‘-s.’ *Journal of Child Language*, 39(1), 130–61. <https://doi.org/10.1017/S0305000910000735>
- Topolińska, Z. (1976). Wyznaczoność (tj. charakterystyka referencyjna) grupy imiennej w tekście polskim, I: Uwagi ogólne. *Polonica*, 2, 33–72.
- Topolińska, Z. (1977). Wyznaczoność (tj. charakterystyka referencyjna) grupy imiennej w tekście polskim. *Polonica*, 3, 59–78.
- Topolińska, Z. (1981). *Remarks on the Slavic noun phrase*. Zakład Narodowy im. Ossolińskich.

Appendix

Minority statements

1. Wszystkie kangury mają torby.
‘All kangaroos have pouches.’
2. Wszystkie węże znoszą jaja.
‘All snakes lay eggs.’
3. Wszystkie koty karmią mlekiem młode.
‘All cats feed their young with milk.’

4. Wszystkie lwy mają grzywy.
'All lions have mane.'
5. Wszystkie kozy dają mleko.
'All goats give milk.'
6. Wszystkie jelenie mają poroże.
'All deer have antlers.'

Past-tense statements involving pragmatically implied quantifier domain restriction

1. Wszyscy przynieśli prezenty dla Hani.
'Everyone brought presents for Hania.'
2. Wszyscy polubili fryzurę Basi.
'Everyone liked Basia's hairstyle.'
3. Wszyscy chcieli poznać się z Tomkiem.
'Everyone wanted to meet Tomek.'
4. Wszyscy pogłaskali kota przed wyjściem.
'Everyone petted the cat before leaving.'
5. Wszyscy obserwowali ptaka na gałęzi.
'Everyone was watching the bird on the branch.'
6. Wszyscy nakarmili kaczki nad stawem.
'Everyone fed the ducks by the pond.'

Present-tense statements involving pragmatically implied quantifier domain restriction

1. Wszyscy uczniowie mają zielone koszulki.
'All students are wearing green T-shirts.'
2. Wszyscy goście mają puste talerze.
'All the guests have empty plates.'
3. Wszyscy kelnerzy trzymają tace.
'All the waiters are holding trays.'
4. Wszystkie wilki mają obroże z nadajnikami.
'All wolves are wearing tracking collars.'
5. Wszystkie bociany mają podcięte skrzydła.
'All storks have clipped wings.'
6. Wszystkie owce są ogolone.
'All sheep are shorn.'