

Meaningful Blurs: the Sources of Repetition-based Plurals in ASL*

Philippe Schlenker**
(Institut Jean-Nicod, CNRS; New York University)

Jonathan Lamberton

Revised, March 28, 2020

Abstract. In several sign languages, plurals can be realized with unpunctuated or punctuated repetitions of a noun, with different semantic implications; similar repetition-based plurals have been described in some homesigns and silent gestures. Unpunctuated repetitions often get approximate 'at least' readings while punctuated repetitions typically correspond to 'exactly' readings. The prevalence of these mechanisms could be thought to be a case in which Universal Grammar does not just specify the abstract properties of grammatical elements, but also their phonological realization, at least in the visual modality. We explore an alternative in which punctuated and unpunctuated repetitions arise from general properties of iconic representations. On an empirical level, we argue that in American Sign Language, punctuated and unpunctuated repetitions are unlikely to be an exclusively grammatical mechanism, as they can be found with purely iconic (word-free) representations. On a theoretical level, we argue for a modular account with three components. First, repetition-based plurals can create a simplified pictorial representation. Second, unpunctuated repetitions give rise to pictorial vagueness, resolved by way of quantification over precisifications. Third, a pragmatic process (which we model within the 'Rational Speech Act' framework) maps these vague representations onto a set of candidate linguistic meanings, including some 'at least' plural readings that are best expressed by unpunctuated repetitions.

* Authors' acknowledgments: we are very grateful to two reviewers for *Linguistics & Philosophy* for extraordinarily constructive and helpful comments, which led to several extensions and improvements.

P.S.'s acknowledgments: For discussion and references, I am indebted to Carlo Geraci, Gabe Greenberg and Jeremy Kuhn. Special thanks to Benjamin Spector for illuminating discussions of the pragmatic part. For judgments on onomatopoeias in English, I am very grateful to Brian Buccola, Jonah Katz, Jeremy Kuhn, Nathan Klinedinst, Matthew Mandelkern, Lyn Tieu. I also received much input on comics from Jonathan Bonhomme and Claire Schlenker. Many thanks to Lucie Ravaux for checking all averages and for putting the raw scores in the Excel file made available in the Supplementary Materials, as well as for working on the bibliography.

Authors' contributions: Philippe Schlenker initiated this research, constructed all examples in consultation with Jonathan Lamberton, and developed the analysis. Jonathan Lamberton was the ASL consultant for the initial phase of the work. When it was written, he checked or provided transcriptions and translations, and discussed important aspects of the analysis. Any theoretical discussion among co-authors occurred only after the data were collected and the first version of the article was written, and primarily by email because the authors were not on the same continent. The separation between the two phases of the work was intended to minimize the risk the 'theoretical contamination' of sign language judgments.

Grant acknowledgments: This research received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 788077, Orisem, PI: Schlenker). Research was conducted at Institut d'Etudes Cognitives, Ecole Normale Supérieure - PSL Research University. Institut d'Etudes Cognitives is supported by grants ANR-10-IDEX-0001-02 and FrontCog ANR-17-EURE-0017.

** Institut Jean-Nicod (ENS - EHESS - CNRS), Département d'Etudes Cognitives, Ecole Normale Supérieure, Paris, France; PSL Research University ; New York University, New York.

Keywords: sign language semantics, iconicity, plurals, unpunctuated repetitions, punctuated repetitions, pictorial semantics, iconic semantics, super semantics

| | | |
|-----------|---|-----------|
| 1 | Introduction | 3 |
| 1.1 | Goals | 3 |
| 1.2 | Repetition-based plurals in ASL | 4 |
| 1.3 | Structure | 5 |
| 1.4 | Transcription conventions and elicitation methods | 6 |
| 2 | Grammatical vs. iconic interpretation of an ambiguous constructions | 7 |
| 3 | Comparing repetitions of nouns and of purely iconic signs | 10 |
| 3.1 | Simple repetitions | 10 |
| 3.2 | Repetitions with modifications | 12 |
| 4 | Downward-monotonic environments | 15 |
| 4.1 | 'Exactly' readings in downward-monotonic environments | 15 |
| 4.2 | Intermediate conclusion | 17 |
| 4.3 | Outlook: logical and pictorial conditions | 18 |
| 5 | Pictorial semantics | 19 |
| 5.1 | Basic pictorial semantics | 19 |
| 5.2 | Pictorial truth-of | 22 |
| 6 | Blurry pictures | 23 |
| 6.1 | Semantics through precisifications | 23 |
| 6.2 | Using bivalent tolerant truth for blurry pictures | 24 |
| 6.3 | Bivalent tolerant truth via entailments | 26 |
| 6.4 | Deriving 'at least' readings | 26 |
| 6.5 | The Edge Effect revisited | 27 |
| 6.6 | Roles of vagueness | 29 |
| 7 | Revisiting the truth conditions of ASL repetition-based plurals | 30 |
| 8 | Pragmatic exploitation | 31 |
| 8.1 | Motivations | 31 |
| 8.2 | Varying sets of precisifications | 33 |
| 8.3 | Literal addressee A_0 | 33 |
| 8.4 | Strategic speaker S_1 | 35 |
| 8.5 | Strategic addressee A_1 | 36 |
| 8.6 | Intermediate conclusion | 37 |
| 9 | Modifying the marking rule? | 37 |
| 9.1 | A problem | 37 |
| 9.2 | An iconic solution | 38 |
| 10 | Possible extensions | 41 |
| 10.1 | Connection with experimental data on classifier predicates | 41 |
| 10.2 | Sign language pluractionals | 42 |
| 10.3 | Repetition-based plurals and pluractionals in gestures | 42 |
| 10.4 | Repetitions in the vocal modality | 43 |
| 10.4.1 | Onomatopoeias | 43 |
| 10.4.2 | Repetition-based plurals and pluractionals in speech | 44 |
| 10.5 | Comics and cartoons | 44 |
| 11 | Conclusion | 45 |
| 11.1 | Main results | 45 |
| 11.2 | Open questions | 45 |
| | Appendix I. Blurry pictures: truth via entailments vs. bivalent tolerant truth | 47 |
| | Appendix II. Adding cost | 48 |
| | Appendix III. Mini-survey on repeated onomatopoeias | 53 |
| | Supplementary Materials: Raw Data | 55 |
| | References | 56 |

1 Introduction

1.1 Goals

Plurals have been argued to be realized by two kinds of repetitions in several sign languages, as well as in some homesigns and in some pro-speech (= word-replacing) gestures (Pfau and Steinbach 2006, Coppola et al. 2013, Abner et al. 2015, Schlenker and Lamberton 2019, Schlenker, to appear a). Punctuated repetitions are made of discrete, clearly separable iterations of the same nominal sign in different parts of signing space, and they are usually associated with precise quantitative conditions, corresponding to the number of iterations. Unpunctuated repetitions involve iterations with shorter and less distinct breaks between them, which makes them less distinct and typically harder to count; they are associated with imprecise quantitative conditions and are often compared to spoken language plurals. The existence of both types of repetitions in a homesigner investigated with experimental means by Coppola et al. 2013 was striking because his hearing mother didn't use such repetitions: this was a grammatical-like device that he had invented. Similarly, non-signers have been argued to understand repetition-based plurals in gestures, with a distinction between punctuated and unpunctuated ones (Schlenker and Lamberton 2019, Schlenker, to appear a). The ability of some homesigners to produce and of some non-signers to understand different kinds of repetition-based plurals with little or no input raises a question: what is the source of repetition-based plurals in signs and in gestures?

One possibility is that, in this particular area, Universal Grammar does not just specify the abstract properties of some grammatical features (such as person and number), but also their phonological realization. On this view, Universal Grammar specifies a mapping between the grammatical form of plurals and certain patterns of repetition. A purely grammatical (non-iconic) version of this view is hard to maintain, at least for ASL (American Sign Language). Schlenker and Lamberton 2019 show in detail that repetition-based plurals can be associated with detailed iconic conditions: in their data, the construction doesn't just specify the number of objects involved, but also their spatial arrangement; for instance, three punctuated iterations of *TROPHY* arranged as a triangle refer to a group of three trophies arranged in this fashion. They make similar (if preliminary) claims about repetitions in gestures. For ASL, they argue that the distinction between unpunctuated and punctuated repetitions should in the end be viewed in an iconic fashion, but they do not go into the details of how the iconic process works. Still, a weakened version of the grammatical view remains a live option: it could be that the distinction between unpunctuated and punctuated repetitions is grammatical and specified by Universal Grammar, even though the iconic component isn't so specified (in part because it is gradient and modifiable at will).

We develop an alternative to these grammatical theories and argue that the realization of punctuated and unpunctuated repetitions originates in iconic cognition. Specifically, we propose that repetition-based plurals in signs (and probably gestures) are an area in which iconic semantics and pragmatic reasoning conspire to yield a grammatical effect. Grammar plays a role at an initial stage in providing a Logical Form for the surrounding expressions (e.g. by way of existential quantification over pluralities), and at the last stage in providing some candidate meanings for iconic expressions. But the mapping between these candidate meanings and various forms is determined by iconic and pragmatic principles.

Our main empirical argument is that internal to ASL, unpunctuated and punctuated repetitions can have the same types of implications for normal signs and for entirely iconic, word-free representations, including in atypical cases in which each iteration differs from the preceding one (e.g. to represent objects of increasing size).¹ On a theoretical level, we propose to start from

¹ By 'word-free representations', we mean improvised gestural representations identified by the native ASL consultant (and co-author) as corresponding to no conventional lexical form. We explore cases of repetition with modification to make it even less likely that a grammatical operation is at work (since to our knowledge standard pluralization does not involve modifications of a lexical sign).


an iconic semantics based on geometric projections, as in Greenberg's (2013, 2019) analysis of pictorial semantics. This immediately yields precise quantitative conditions for punctuated repetitions. The case of unpunctuated repetitions is more complex because it involves pictorial vagueness. We analyze it within a precisification-based 'tolerant' logic, one that makes a representation true just in case one of its precisifications is (our semantics is a 'bare bones' bivalent version of systems discussed in Cobreros et al. 2012). This semantics for blurry pictures then feeds into a pragmatic process mapping pictures to candidate meanings through a process of pragmatic enrichment which we model using the Rational Speech Act (RSA) model (e.g. Bergen et al. 2016).

The significance of these results for semantic theory is twofold. Substantively, they suggest that as logical a category as plurality may be realized through iconic means: this is an area in which logic and iconicity are intertwined at the heart of language. Methodologically, this is to our knowledge the clearest case so far in which the details of Greenberg's pictorial semantics are crucial to develop a semantic analysis of an uncontroversially semantic phenomenon. This justifies an expansion of the semanticist's toolbox – in fact, several extensions, since our data will require that pictorial semantics be combined with vagueness and a mechanism of pragmatic enrichment.

1.2 Repetition-based plurals in ASL

In a typological study, Pfau and Steinbach 2006 note that "reduplication and zero marking appear to be two basic pluralization strategies attested in all sign languages that have been investigated so far". Schlenker and Lamberton 2019 develop an analysis of repetition-based plurals in ASL, with a special focus on three types of repetitions, which they illustrate with (1) (with acceptability judgments on a 7-point scale at the beginning of each example).

- (1) HERE HAVE
'Here there



- a. ⁶ [TOILET_ TOILET TOILET]_{horizontal}.
are three (?) toilets.'
b. ⁷ TOILET-rep_{3horizontal}.
are (at least 3) toilets.'
c. ^{6,7} TOILET-cont_{horizontal}.
are (at least 3-4) toilets.'

Video: <https://youtu.be/ov2mww9Z9AU>
(ASL 33, 0522a,b,c)

As they write,

(1)a is a punctuated repetition, with three horizontal iterations of the word *TOILET*, realized as a trembled *T* of the manual alphabet, whose first iteration is illustrated (...). Characteristic is the fact that the hand goes down between the three iterations, which makes them easy to individuate: the *T* handshape disappears and goes back to a neutral position between iterations. In (1)b, there are three iterations as well, but realized a bit faster, and in an unpunctuated fashion, characterized by the fact that the hand only partly goes down between the iterations: the *T* handshape only partially disappears between them. The realization can be much faster in other examples, making it hard to count the iterations, and highlighting the contrast with punctuated repetitions (which is rather subtle in (1)). Finally, in (1)c a continuous repetition is illustrated; the hand does not go down at all between the iterations, and these are faster, more numerous, and very hard to count.

In one of their final analyses, summarized in (2), all repetition-based plurals have the same semantic core. First, bare nouns are assumed to be true of atoms and non-atoms like, as is made plausible by the fact a bare noun can have a plural reading, especially when combined with a numeral. This justifies Condition (i), which states that a repeated noun N should hold true of the denoted group.² In addition, however, repetition-based plurals crucially come with an iconic component, stated in Condition (ii), and they also have the ability to create new singular discourse referents, as seen in Condition (iii); the latter is motivated by the fact that in constrained cases individual iterations are made available for further anaphoric uptake.

- (2) **Schlenker and Lamberton 2019, "Improved Analysis 1"**: semantics of punctuated, unpunctuated and continuous repetitions without pluralization, but with the assumption that count terms are true of atoms and non-atoms

For $N\text{-iter}_X = N\text{-rep}_X$ or $[N N N]_X$ or $N\text{-cont}_X$, if P is a propositional expression $[[N\text{-iter}_X P]]^{f,s} = \text{true}$ iff for some group d ,

- (i) $[[N]]^{f,s[X \rightarrow d]}(d) = \text{true}$, and
- (ii) $N\text{-iter}$ iconically represents d given c , and
- (iii) $[[P]]^{f,s[X \rightarrow d, x_1 \rightarrow d_1, \dots, x_n \rightarrow d_n]} = \text{true}$, where x_1, \dots, x_n are discourse referents made available by the iconic semantics of $N\text{-iter}$, and x_1, \dots, x_n respectively denote d_1, \dots, d_n (which are parts of d).

This analysis is underspecified as to the difference between punctuated and unpunctuated repetitions.³ We take up this problem in the present piece, varying two parameters that are somewhat conflated in this earlier analysis: (i) whether the iterations are realized slowly or quickly, and (ii) whether the hand goes down between iterations. In clear earlier cases, unpunctuated repetitions were realized quickly and without the hand going down, while punctuated repetitions were realized slowly and with the hand going down. Since we are interested in the precise source of these readings, it will be interesting to distinguish these two properties. For this reason, we will be talking of *separated* vs. *unseparated* iterations depending on whether the hand goes down, and *slow* or *fast* iterations depending on their speed. A clear punctuated repetition would thus be slow and separated, a clear unpunctuated repetition fast and unseparated. As we will see, the traditional terminology of *punctuated* vs. *unpunctuated* makes excellent sense in the end, because what seems to matter for most (though not all) purposes is whether the iterations are presented as easy or hard to count. (We leave out continuous repetitions from the present discussion. They were primarily used by Schlenker and Lamberton to investigate iconic properties of mass nouns, but some of their continuous repetitions applied to count nouns are arguably similar to the fast unseparated iterations we study in this piece).

1.3 Structure

The rest of this article is organized as follows. After laying out our transcription conventions and elicitation methods (Section 1.4), we introduce our initial finding: one and the same construction may be ambiguous between a repetition-based plural and a purely iconic use of a repetition, but in such cases the grammatical and the iconic constructions give rise to broadly similar inferences (Section 2). We then turn to simpler and unambiguous paradigms, systematically comparing the pluralization of words and the repetition of purely iconic and improvised manual representations, and arguing that punctuated vs. unpunctuated repetitions yield the same semantic differences

² Without the assumption that nouns hold true of atoms and non-atoms alike, one would need to posit an operation of pluralization of the noun before it can hold true of non-atoms.

³ Note that some of their examples suggest that punctuated repetitions do not invariably come with quantitative conditions that correspond to the number of iterations. Specifically, the number *ten* may be followed by three punctuated iterations of *TROPHY* to mean *ten trophies spread out*. The punctuated nature of the repetition is indicative of the fact that the trophies are spread out, not that there are three of them. Schlenker and Lamberton 2019 thus take the quantitative condition to just be a default that can be overridden. We do not seek to account for these examples in the present piece.

across the two types (Sections 3-4). We develop our analysis in stages, starting with a pictorial semantics augmented with a precisification-based account of vague representations (Section 5-6), which can be integrated with a standard semantics for nominals (Section 7). We then consider an improvement that relies on a mechanism of pragmatic enrichment: while on purely iconic grounds unpunctuated repetitions are very unlikely to have 'at least' readings, this possibility is exploited and thus made likely through strategic communicative reasoning (Section 8). We refine our analysis in Section 9, and extend it to different data types (including gestures and onomatopoeias) in Section 10, before summarizing our conclusions and open questions (Section 11).

Since the argument and the analysis have several moving parts, we provide a table to help the reader gain a bird eye's view of the proposal.

(3) Structure of the proposal

| | | | |
|-------------------------------|--|--|-------------|
| Empirical analysis | Iconic vs. grammatical uses of repetitions | One and the same construction is ambiguously linguistic or purely iconic, and punctuated/unpunctuated repetitions yield the same reading types across the linguistic and iconic cases. | Section 2 |
| | | Punctuated/unpunctuated repetitions as well as repetitions with modifications have similar semantic effects across the linguistic and purely iconic cases. | Section 3 |
| | 'Exactly' vs. 'at least' readings | Punctuated repetitions can have 'exactly' readings', unpunctuated repetitions can have 'at least readings': argument from downward-monotonic environments. | Section 4 |
| Core Theory | Basic pictorial semantics | A basic pictorial semantics yields an appropriate account of punctuated repetitions. | Section 5 |
| | Pictorial semantics with vagueness | With a vague component, a pictorial semantics can yield an account of unpunctuated repetitions. | Section 6 |
| | Integrating logical and pictorial semantics | The logical and the iconic components can be integrated by way of quantification over pluralities that must resemble the repeated element. | Section 7 |
| | Pragmatic enrichment | Even if the set of precisifications needed to derive 'at least' readings is extremely unlikely, this possibility is exploited for communicative purposes. | Section 8 |
| Theoretical extensions | Refining the pictorial semantics: improving the marking rule | Slow punctuated repetitions can have 'at least' readings; this can be derived using a more discriminating marking rule. | Section 9 |
| | Adding cost to the pragmatic analysis | When 'at least' readings are preferably realized with 2-3 (rather than n) repetitions, this can be accounted for by taking into account the cost of repetitions. | Appendix II |

1.4 *Transcription conventions and elicitation methods*⁴

In the following, sign language sentences are glossed in capital letters, as is standard. Translations were chosen to reflect, to the extent possible, the inferential judgments that were given. Inferential questions appear with the examples. Raw data can be consulted in the Supplementary Materials.

A suffixed locus, as in *WORD-i*, indicates that the word points towards locus *i* (a position of signing space associated with a discourse referent). *IX* (for 'index') is a pointing sign towards a locus, while *POSS* is a possessive; they are glossed as *IX-i* and *POSS-i* if they point towards (or 'index') locus *i*; the numbers *1* and *2* correspond to the position of the signer and addressee respectively. *IX-i* is a standard way of realizing a pronoun corresponding to locus *i*, but we also use *IX* without a suffixed locus when it is repeated. Agreement verbs include loci in their realization – for instance the verb *1-GIVE-2* starts out from the first person locus *1* and targets the second person locus *2*; it means that the signer gives something to the addressee.

In the literature and in (1), the suffix *-rep* is used for unpunctuated repetitions, and in

⁴ This part is adopted from Schlenker and Lamberton 2019.

such cases *-rep3*, *-rep4*, *-rep5*, ... indicate that there are 3, 4, 5, ... iterations. When relevant, we add a subscript indicating the shape of the repetition, e.g. *-rep3_{horizontal}* for a horizontal repetition (whether in a straight line or as a horizontal arc), *-rep3_{triangle}* for a triangular-shaped repetition. The suffix *-cont* is used for continuous repetitions, and subscripts may be used as well to indicate the shape of the movement, such as *-cont_{horizontal}* or *-cont_{triangle}*. Punctuated repetitions of an expression *WORD* are encoded as [*WORD WORD WORD*] if they involve three iterations of that expression; [*WORD WORD WORD*]_{horizontal} and [*WORD WORD WORD*]_{triangle} provide information about the shape of the repetition.

In this piece, we refine these descriptions. We use *-unsep3* for 3 unseparated repetitions, *-sep3* for 3 separated repetitions, and the subscripts *fast* and *slow* for fast and slow repetitions respectively. We also encode as a subscript the shape formed by the repetitions. For instance, we use *o-unsep4_{fast-circular}* to encode 4 unseparated fast iterations of an o-like shapes, and these 4 shapes form a circle. Notational conventions will be introduced or refined as we go. We do not encode non-manual expressions, which remained as constant as possible within sentences of a given paradigm (i.e. set of examples sharing the same number).

The consultant (and co-author) is a Deaf, native signer of ASL (of Deaf, signing parents).⁵ Elicitation was conducted using the 'playback method', described for instance in Schlenker et al. 2013, Schlenker 2014, Schlenker and Lamberton 2019. It involved repeated quantitative acceptability judgments (1-7, with 7 = best), as well as well as inferential judgments, obtained in two ways: by way of open questions, and by way of quantitative judgments of inferential strength (1-7, with 1 = no inference and 7 = strongest inference). Judgments of inferential strength have become standard in experimental pragmatics, and have proven useful in ASL fieldwork (Schlenker, to appear b)⁶. For readability, only average judgments are given (complete quantitative judgments are given when there is more than a 2-point difference in the judgments obtained for a given sentence). In summary tables, we use gray shades for sentences that are degraded (average acceptability at or below 4), as inferential judgments in these cases might not be telling. Specialists are invited to consult the raw data when relevant.⁷

2 Grammatical vs. iconic interpretation of an ambiguous constructions

We now argue that punctuated and unpunctuated repetitions alike have the same kinds of semantic implication in purely iconic (word-free) representations as they do in linguistic environments, which suggests that they are based on a cognitive mechanism that is not purely linguistic.

We start with a paradigm in which, depending on the context, the same sequence can be interpreted as a grammatical construction involving a repetition-based plural, or as a purely iconic (word-free) representation. The grammatical construction is slightly different from the ones we've discussed up to this point: instead of involving the repetition of a noun, it involves the repetition

⁵ We use the term 'consultant' to refer to information obtained in elicitation sessions (and thus prior to writing the article).

⁶ As Cremers and Chemla 2017 write, that graded inferential judgments "may help detect otherwise hidden effects". This was borne out for our ASL consultant (who has years of experience with quantitative acceptability judgments): he noted in the past that the inferential part of his task was made easier, not harder, by using quantitative judgments of inferential strength. Without these, he had to reflect at length about how to categorize judgments of intermediate strength; the quantitative method allowed for less arbitrary decisions in such cases.

⁷ As mentioned in Schlenker, to appear b, one might worry that our sentences are assessed by the very signer that produced them. It is clear that our results can only bear on the idiolect of a particular individual. But the work we conduct on this idiolect is arguably more rigorous than is common in non-experimental fieldwork, since (i) we distinguish the production phase from the judgment phase, and (ii) we provide multiple quantitative data points that make it possible to assess the stability (or lack thereof) of the judgments through time. This method should of course be assessed rigorously in the future, but this holds of other elicitation methods as well – and ours has the advantage of transparency (since the raw written data are made available in the Supplementary Materials).

of a pointing sign associated with a noun. Specifically, in (4), the construction of interest is *HOUSE* followed by three pointing signs. A grammatical use of the latter construction implies the presence of several houses, with different quantitative conditions depending on the realization of the repetition. For instance, with three slow separated iterations, the implication is that there are four houses, i.e. one for the initial occurrence of *HOUSE*, and an additional one for each pointing sign. But the very same sequence can also be understood in a completely different way. Specifically, if the context specifies that I have a contemporary drawing with nothing but houses *followed by dots*, the three 'pointing signs' can be taken to be iconic representations of dots (importantly, this does not involve the word *DOT*, which would be realized by way of finger-spelling or a classifier, for instance⁸). In our case, the word *HOUSE* is taken to represent a singular house, and three slow separated iterations of the iconic representation of a dot indicate that there are three dots next to the house in the drawing. We call these three iterations 'word-free' because there is no lexical form at all used here: this is, to the best of our knowledge, an improvised gestural and iconic representation of a dot (using one's finger), similar to one that a non-signer could produce.

Using an ambiguous construction allows for a very minimal comparison between a grammatical and an iconic use of a repetition: the consultant thus assessed one and the same video twice in each judgment task, just changing the context in which it was evaluated, thereby switching the interpretation from 'several houses' to 'one house followed by three dots'.

A word should be said about the inferential judgments. They were of two kinds: open judgments about the number of relevant objects found in the drawing; and judgments of inferential strength pertaining to exactly *n*, at least *n* and approximately *n* readings (with 7 = strongest endorsement).⁹ We only report the latter, quantitative judgments in the text, but it can be checked in the Supplementary Materials that these did not miss anything essential found in the answers to the open questions. In the grammatical case, *HOUSE* followed by three slow separated iterations of *IX* yielded an inference that there were exactly 4 houses (not exactly 3), presumably because each occurrence of *IX* referred to an additional house. In our summary table, we encoded how many objects the repeated term referred to: it would be 3 in this case (since this is the contribution made by the repeated pointing sign).

- (4) *Context 1*: I have a contemporary drawing with nothing but houses.
Context 2: I have a contemporary drawing with nothing but houses followed by dots.

POSS-1 DRAWING

'My drawing has

a. *HOUSE*.
 has a house.'

b. *HOUSE IX-unsep*_{fast-horizontal}
 at least three houses.' (Context 1)
 a house next to at least four dots.' (Context 2)

c. *HOUSE IX-unsep*_{slow-horizontal}
 exactly four houses.' (Context 1)
 a house next to exactly three dots.' (Context 2)

d. *HOUSE IX-sep*_{fast-horizontal}

⁸ For some examples, see the discussion by William Vicars at <http://www.lifeprint.com/asl101/pages-signs/d/dots.htm> (retrieved November 6, 2019).

⁹ Note that the inferential possibilities were clearly understood with their pragmatic implications. For instance, in (4)e there was a high endorsement of the 'exactly three' reading and a lower endorsement of the 'approximately 3' and 'at least three' readings, no doubt because these were interpreted to exclude the more informative 'exactly 3' situations.

(deviant in both contexts)

- e. HOUSE IX-sep3_{slow-horizontal}.
 exactly four houses.' (Context 1)
 a house next to exactly three dots.' (Context 2)
 (ASL, [35, 1440](#); anonymized video: <https://youtu.be/njmMqpCshY4>)

(5) Inferential questions (sample)

Meaning 1: What do you infer about the number of (Context 1) **houses found** (Context 2) **dots found** in the contemporary drawing? (provide an estimate, e.g. n, more than n, between m and n, etc)

Meaning 3: About the objects mentioned in Meaning 1:

Context 1: How strongly do you infer that there were (i) exactly 4? (ii) approximately 4? (iii) at least 4? (iv) at least 5?

Context 2: How strongly do you infer that there were (i) exactly 3? (ii) approximately 3? (iii) at least 3? (iv) at least 4?

(6) Acceptability and inferential judgments (*Meaning 3*) obtained for (4)

| Examples (ASL, 35, 1440) | Context | Acceptability | Number of objects that the repetitions stand for | | | |
|--|---------------|---------------|--|-----|-----|--------------|
| | | | = 3 | ±3 | ≥3 | ≥4 |
| a. HOUSE | 1 grammatical | 7 | 1 | 1 | 1 | 1 |
| | 2 iconic | 5.7 (3, 7, 7) | 1 | 1 | 1.3 | 1 |
| b. HOUSE IX-unsep3 _{fast} | 1 grammatical | 6.7 | 1.3 | 2.7 | 5.3 | 4.7 |
| | 2 iconic | 5 | 1 | 1.3 | 2 | 6.3 |
| c. HOUSE IX-unsep3 _{slow} | 1 grammatical | 7 | 6 | 3.7 | 2 | 1 |
| | 2 iconic | 6.7 | 6 | 3 | 1.7 | 1 |
| d. HOUSE IX-sep3 _{fast} | 1 grammatical | 2.7 | 4 | 5 | 3.7 | 2.3 |
| | 2 iconic | 3.3 | 3.7 | 5 | 3.7 | 1.7 |
| e. HOUSE IX-sep3 _{slow} | 1 grammatical | 7 | 5.7 | 3 | 3.3 | 1 |
| | 2 iconic | 7 | 5.7 | 2.3 | 2.3 | 1 (1;1;5) |

The results are striking: whether it is interpreted as a pluralized form of the pointing sign, or as a purely iconic (gestural) representation of multiple dots, the repeated pointing gives rise to rather similar quantitative inferences (*modulo* the difference, mentioned at the outset, that in the grammatical interpretation all quantities must be incremented by 1 due to the presence of the noun *HOUSE*). The control in (4)a is just interpreted as involving one house and no dots (an un-repeated noun is often interpreted by default as involving singularity; when co-occurring with numerals, the situation radically changes, of course). The other examples all involve 3 iterations of pointing. In (4)b, they are realized as unseparated and fast, hence a typical unpunctuated repetition. When interpreted grammatically, the iterations yield an inference that there are at least 3 or 4 houses in addition to the one represented by the noun. When interpreted iconically, they yield an inference that there are at least 4 dots (thus this is not exactly the same inference as in the grammatical case, but close enough). When realized as unseparated but slow, the 3 iterations yield an inference to the existence of exactly 3 (additional) houses, or of exactly 3 dots. Here and throughout, signing the repetitions separated but roughly at the same speed as unseparated fast repetitions is deviant (this just seems inconvenient to realize); since all such examples are severely degraded, we will not discuss the inferences they give rise to. When realized separated and slow, the repetitions primarily give rise to an inference about the existence of exactly 3 (additional) houses, or exactly three dots.

Thus a paradigm constructed to be ambiguous between a grammatical and an iconic reading gives rise to the same inferences for punctuated vs. unpunctuated repetitions across the two readings. This suggests that the inferences triggered by different types of repetitions might arise outside of purely grammatical cases.

3 Comparing repetitions of nouns and of purely iconic signs


3.1 Simple repetitions

Our initial example had the advantage of allowing for the assessment of the very same of signs on two interpretations, one grammatical and one purely iconic, depending on the context. But this forced us to use a slightly more involved construction than a standard plural. In a normal plural construction, *HOUSE* would be repeated to yield a plural meaning. In the grammatical interpretation of the case we studied, a plural meaning was obtained with *HOUSE* followed by the iteration of the pointing sign – a crucial device in order to get the ambiguity we wanted. In this section, we explore simpler and unambiguous paradigms without a pointing sign, of two sorts: some involve the repetition of a normal sign, others that of a purely iconic representation.

The repetitions we consider are arranged in a circle. This choice was made to help bring out semantic differences among the examples. Our initial worry was this: plurals can easily get existential readings. Take for instance 4 punctuated iterations of a circle on a line, in the context of a drawing (assume for concreteness that these are purely iconic, word-free representations, although the problem is similar for words). This could be taken to mean that it is possible to find four circles in the drawing; this reading wouldn't preclude the presence of further circles on the line. So one would have to ask about inferences about the number of circles inferred to be *on the relevant part of the line*, which yields complications (e.g. there could be a continuation of the line that the signer decided not to represent). We sought to limit the problem by representing the repetitions in a circle, where it might be less likely that there were further objects that the signer decided to omit, as it might seem more natural to omit part of a line on a drawing than to omit part of a circle (as we will see in Section 4, this ploy was not entirely successful, in a way that might be informative).

We start with the repetition of a word, here the *G* of the manual alphabet. It does not in any way resemble the printed letter *G*, as can be seen in (7), and we will often refer to it as \mathcal{G} ¹⁰. In other words, it is a word that happens to refer – non-iconically – to a letter.

(7) Sign for *G* (signed high, in (8)a, as a control for circular repetitions that follow)



To facilitate comparison with purely iconic cases, we used a context pertaining to a contemporary abstract drawing, which may contain letters in some cases, and geometric shapes in others.

As in all other cases, there were open and closed, quantitative inferential questions. We report judgments of inferential strength for the second question in (9), and we henceforth write acceptability judgments as superscripts at the beginning of the relevant sentences. (It should be noted that in the unseparated fast condition, it is often somewhat difficult to count the manual *G*'s.)

Notation: $G\text{-unsep}4_{\text{fast-circular}}$ refers to 4 fast unseparated iterations of the manual *G* (i.e. \mathcal{G}) arranged as a circle.

(8) *Context:* I have a contemporary abstract drawing
POSS-1 DRAWING
'My drawing contains

a. \mathcal{G} .

¹⁰ There are different possible realizations of ASL *G*. This one is close to what is represented in (7) (the drawing is from <https://www.realsasl.com/learn-south-african-sign-language/144-south-african-sign-language-alphabet-abc>, retrieved on October 17, 2019).

- a. G.¹
 b. ⁷G-unsep_{fast-circular}⁴
 at least 4 or 5 G's (arranged in a circular fashion).¹
 c. ⁷G-unsep_{slow-circular}⁴
 exactly 4 G's (arranged in a circular fashion).¹
 d. ⁴G-sep_{fast-circular}⁴
 (deviant)
 e. ⁷G-sep_{slow-circular}⁴
 4 / at least 4 G's (arranged in a circular fashion).¹
 (ASL, [35, 1492](#); anonymized video: <https://youtu.be/GX0awdkEa-s>)

(9) Inferential questions (sample)

Meaning 1: What do you infer about the number of **shapes** in the contemporary drawing? (provide an estimate, e.g. n, more than n, between m and n, etc)

Meaning 3: About the objects mentioned in Meaning 1: How strongly do you infer that there were **(i) exactly 4? (ii) approximately 4? (iii) at least 4? (iv) at least 5?**


(10) Inferential judgments (= Meaning 3) obtained for (8)

| Examples (ASL, 35, 1492) | Number of objects that the repetitions stand for | | | |
|--|--|-----|-------------|-------------|
| | = 4 | ±4 | ≥4 | ≥5 |
| a. G | 1 | 1 | 1 | 1 |
| b. G-unsep _{fast-circular} ⁴ | 1.7 | 3.3 | 5.3 | 5.7 |
| c. G-unsep _{slow-circular} ⁴ | 5.7 | 3.3 | 1.7 | 1 |
| d. G-sep _{fast-circular} ⁴ | 2 | 3.7 | 4.3 (6;3;4) | 4.3 (5;2;6) |
| e. G-sep _{slow-circular} ⁴ | 4.7 | 3.7 | 5 | 2 |

With one exception, the generalizations are the same as in Section 2: 4 unseparated iterations of G signed quickly give rise to the inference that there are at least 4 or at least 5 G's in the circle (it is worth noting that the fast unseparated iterations are genuinely hard to count, which makes the contrast between *at least 4* and *at least 5* moot). By contrast, 4 slow unseparated iterations give rise to the inference that there were exactly (or possibly approximately) 4. Fast separated iterations are degraded, as before. One difference relative to the paradigm in (4) is that slow separated iterations might be ambiguous, meaning *exactly/approximately 4* or *at least 4* (in (4) the separated slow repetitions only gave rise to an 'exactly' reading). We come back to an explanation of this ambiguity in Section 9.

We turn to a minimally different case in which an iconic manual shape represents a small circle, as in (11). This is close in shape to a classifier used for small objects, but here it is just interpreted to represent a circle, in an iconic fashion.¹¹

(11) Iconic representation of a small circle (signed high, in (12)a, as a control for circular repetitions that follow)



¹¹ As Lamberton notes, the circle represented on the fingers in (11) is slightly bigger than the shape he would use for the small object classifier (for him, the diameter of the circle represented inside the thumb and index finger might be 25-50% smaller in the classifier use). In addition, the small object classifier would be accompanied by a facial expression.

(12) *Context*: I have a contemporary abstract drawing.

POSS-1 DRAWING

'My drawing contains

a. ⁷ o .

a little circle.'

b. ⁷ o-unsep₄^{fast-circular}.

at least 5 little circles (arranged in a circular fashion).'

c. ⁷ o-unsep₄^{slow-circular}.

exactly 4 little circles (arranged in a circular fashion).'

d. ^{3.7} o-sep₄^{fast-circular}.

(deviant)

e. ⁷ o-sep₄^{slow-circular}.

exactly 4 little circles (arranged in a circular fashion).'

(ASL, [35, 1466](#); anonymized videos: <https://youtu.be/d-7Cw1eaWUM>)

(13) Inferential judgments obtained for (12)

| Examples (ASL, 35, 1466) | Number of objects that the repetitions stand for | | | |
|--|--|-----|-------------|-----|
| | = 4 | ±4 | ≥4 | ≥5 |
| a. o | 1 | 1 | 1 | 1 |
| b. o-unsep ₄ ^{fast-circular} | 1 | 3.3 | 3 | 5.7 |
| c. o-unsep ₄ ^{slow-circular} | 5.7 | 3.3 | 2 | 1 |
| d. o-sep ₄ ^{fast-circular} | 4.3 | 5 | 2.7 | 1.3 |
| e. o-sep ₄ ^{slow-circular} | 5.3 | 2.7 | 3.3 (1;4;5) | 1.7 |

Rather similarly to the repetition of *G* in (8), unseparated fast iterations give rise to an 'at least 5' reading (with *G*, the 'at least 4' inference was strong as well). The difference between 'at least 4' and 'at least 5' in unseparated fast repetitions should not be overinterpreted: it is just very hard to count the number of occurrences in this case, and thus the distinction between 4 and 5 unseparated fast iterations would probably be hard to detect. With slow unseparated iterations, 'exactly 4' is the dominant reading. Fast separated iterations are, as always, deviant. Slow unseparated iterations yield an 'exactly 4' reading (the 'approximately 4' reading is a bit weaker than with the example involving *G*, as are judgments for the 'at least 4' reading, which are also unstable in this case).

In sum, there is a strong similarity between the effects of various types of repetitions with *G* on the one hand and with an iconic manual representation on the other: unseparated fast iterations give rise to 'at least' readings, unpunctuated slow iterations to 'exactly' readings', unseparated slow iterations are less clear (with an 'exactly' reading, but further interpretations as well in some cases, including an 'at least reading'); unseparated fast iterations are always deviant.

3.2 Repetitions with modifications

Our argument has a flaw, however: we can't exclude the possibility that our iconic representations are words created on the spot, which are then treated according to the grammatical rules of the language. In other words, it could be that all the interesting action is in the transformation of an improvised iconic sign into a word, not in the repetition process that pluralizes it. We attempt to address this issue by modifying the grammatical process itself: instead of a standard repetition, we investigate a repetition with a change of shape, with each iteration larger than the preceding one. We do not know of descriptions of grammatical repetitions that involve such changes of shape from one iteration to the next.

We start with the iconic case, involving again the manual representation of a small circle, as in (11). But here each new iteration will have a larger opening than the preceding one, with the inference that the circles increase in size. (In the rest of this section, we do not provide the

inferential questions, because they are similar to those used in earlier examples; they can be found in the Supplementary Materials.)

Notation: We encode the fact that the iterations are of increasing size by way of the symbol $<$. Thus $o-<unsep4_{fast-circular}$ refers to 4 fast unseparated iterations of little circles of increasing size, arranged in a circular fashion.

(14) *Context:* I have a contemporary abstract drawing.

POSS-1 DRAWING

'My drawing contains

a. ⁷ o.

a little circle.'

b. ^{5.7} o-<unsep4_{fast-circular}.

at least 5 little circles of increasing size (arranged in a circular fashion).'

c. ⁶ o-<unsep4_{slow-circular}.

exactly 4 little circles of increasing size (arranged in a circular fashion).'

d. ³ o-<sep4_{fast-circular}.

(deviant)

e. ^{6.7} o-<sep4_{slow-circular}.



exactly 4 little circles of increasing size (arranged in a circular fashion).'



(ASL, [35, 1464](#); anonymized video: <https://youtu.be/xc7pUsLAKY8>)

(15) Inferential judgments obtained for (12)


| Examples (ASL, 35, 1464) | Number of objects that the repetitions stand for | | | |
|--|--|-------------|-------------|-----|
| | = 4 | ±4 | ≥4 | ≥5 |
| a. o | 1 | 1 | 1 | 1 |
| b. o-<unsep4 _{fast-circular} | 1.3 | 3.3 | 2.7 | 5.7 |
| c. o-<unsep4 _{slow-circular} | 6 | 2.3 | 1.7 | 1 |
| d. o-<sep4 _{fast-circular} | 4.7 | 5 | 2.3 | 1.3 |
| e. o-<sep4 _{slow-circular} | 5.3 | 3.3 (2;3;5) | 3.3 (1;4;5) | 1.3 |

As before, the representation with fast separated iterations is deviant. While the slow separated iterations are acceptable, the unseparated iterations are a bit less so, possibly because changing the hand shape in this semi-continuous fashion is less than natural. In any event, the same general contrasts are found as in our earlier paradigms. First, unseparated fast repetitions give rise to an 'at least' reading – in this case: at least 5 rather than at least 4, which need not be surprising because (as mentioned before) the number of iterations is just very hard to count. Unseparated slow repetitions yield an 'exactly four' reading because of the four iterations. Separated slow repetitions also primarily give rise to an 'exactly four' reading, with greater uncertainty about the 'approximately four' and 'at least four' readings – a point we revisit in Section 9 (we will predict an ambiguity in this case).

It is interesting to compare this purely iconic example to that of a *G* repeated with a change of shape: the initial shape is a *G* of the manual alphabet, i.e. , and the rectangle formed by the thumb and index finger is gradually opened from one iteration to the next, leading to an *L* shape at the end, i.e. ¹², as illustrated in (16).

¹² Our drawing (from the same source as in fn. 10) is not consistent between  and : the former is given from the signer's perspective, the latter from the addressee's perspective (see the pictures in (16) for the 'real' addressee perspective).

(16) First and fourth manual shape appearing in (17)e (proportions are not preserved)



General contrasts are the same as in the iconic case, but the interpretation of the shape itself is of some interest: due to the change of shape from one iteration to the next, a purely iconic interpretation becomes salient, whereby the drawing contains geometric shapes (rather than letters) akin to the manual shapes displayed. Still, as explicitly noted by the consultant, this purely iconic interpretation need not fully preclude a partly iconic interpretation whereby the shapes represent letters: a *G* becoming larger and then becoming an *L* (but see the Supplementary Materials for details¹³).

(17) *Context*: I have a contemporary abstract drawing



POSS-1 DRAWING

'My drawing contains



a. ⁷G.

a G.'

b. ⁶G-<unsep4_{fast-circular}.

approximately 4 shapes of increasing size ranging from from  to  (arranged in a circular fashion).'

c. ^{6.3}G-<unsep4_{slow-circular}.

exactly 4 shapes of increasing size ranging from from  to  (arranged in a circular fashion).'

d. ²G-<sep4_{fast-circular}.

(deviant)

e. ^{6.7}G-<sep4_{slow-circular}.

exactly 4 shapes of increasing size ranging from from  to  (arranged in a circular fashion).'

(ASL, [35, 1494](#); anonymized video: <https://youtu.be/xYwDeywuw-w>)

(18) Inferential judgments obtained for (17)

| Examples (ASL, 35, 1494) | Number of objects that the repetitions stand for | | | |
|--|--|-----|-----------|-----------|
| | = 4 | ±4 | ≥4 | ≥5 |
| a. o | 1 | 1 | 1 | 1 |
| b. o-<unsep4 _{fast-circular} | 3 | 5 | 4 (3;3;6) | 3 (2;2;5) |
| c. o-<unsep4 _{slow-circular} | 6 | 2 | 1.7 | 1 |
| d. o-<sep4 _{fast-circular} | 4.3 | 5 | 4.3 | 3 |
| e. o-<sep4 _{slow-circular} | 5.3 | 3.3 | 4 | 1.7 |

In sum, we have established that purely iconic representations can be iterated by way of punctuated and unpunctuated repetitions in the same way as normal words, and with comparable semantic effects. On the one hand, purely iconic (and improvised) representations can enter in standard patterns of repetitions, including with some 'at least' readings that are characteristic of plurals. On the other hand, the repetition process itself can be modulated in iconic ways, with

¹³ One helpful remark made by the consultant (= [JL 19.10.10]) was phrased as follows: "The signs used might represent abstract shapes rather than actual letters, particularly since the intermediate shapes are clearly not letters. However, the first and last shapes strongly invoke the letters *G* and *L* and that rules for me."

each repetition larger than the preceding one – a process that deviates from grammatical repetitions. When applied to a noun (here the letter name ㄅ), the repetition with modification brings out an iconic reading pertaining to the very shape displayed, but the number and realization of the repetitions yield the same quantitative effects as in other cases.¹⁴

4 Downward-monotonic environments

One important question is whether the 'exactly' readings noted above are an intrinsic component of the literal meaning of the examples, or arise due to an implicature, by competition with other forms. Unembedded sentences cannot decide the issue. To give an example from English, if the numeral *three* means *at least three*, as on many standard treatments, it is expected to give rise to an 'exactly three' reading in unembedded cases due to competition with *four*. To decide whether its literal meaning has an *exactly three* component, we must embed the numeral in downward-monotonic environments, where implicatures are unlikely to arise. This is so because, even for theories that allow implicatures to be computed locally (e.g. Chierchia et al. 2012), the contribution of an expression in such environments is usually its literal meaning.¹⁵

4.1 'Exactly' readings in downward-monotonic environments

We consider sentences of the form *if p, I will give you \$20*, with plural-based repetitions within *p*, and inferential questions pertaining to numbers that would force the signer to give the addressee \$20. *IF*-clauses normally prevent implicatures from being computed locally,¹⁶ so we expect to be able to assess in this way the literal meaning of the constructions.¹⁷ Since we are interested in the iconic semantics of repetitions, we focus on a context involving *look like (a shape)*, which highlights this iconic component. Our first paradigm involves repetitions of the manual *G*, i.e. ㄅ , and some of the results replicate what we saw before, but now with an argument that we are assessing the literal meaning of the iconic representation rather than the result of an implicature. The inferential judgments were intended to test the quantitative condition that would force the signer to give the addressee \$20, hence the formulations in (19).

¹⁴ Since we do not know of other studies of such repetitions with modifications, our discussion leaves several questions open, including (i) how they semantically compare with standard plurals in various environments, and (ii) how they interact with other categories, such as adjectives.

¹⁵ The case of numerals in English turns out to be complicated: as discussed in Spector 2013, one of the main current theories is that numerals have an *at least* reading (in particular under downward-monotonic environments), but that they associate with exhaustivity operators more freely than other constructions, hence the appearance of *exactly* readings even in downward-monotonic environments. We revisit this question in relation to ASL repetition-based plurals at the end of Section 4.2.

¹⁶ The reason is that *IF*-clauses behave like a downward-monotonic environment, which both licenses Negative Polarity Items such as *ever* and suspends implicatures, as illustrated in (i):

- (i) a. If you ever go to China, you should visit Shanghai.
 b. If you invite Robin or Sam, I'll give you \$20.
 => if you invite Robin and Sam, I'll give you \$20.

¹⁷ Some care is needed, however, as any semantic differences might be erased if the context somehow facilitates an existential reading. Concretely: suppose one predicate, call it $=4$, means 'exactly 4 circles', and the other, call it ≥ 4 , means 'at least 4 circles'. If the target sentence tests in (i), with an existential quantifier within the *IF*-clause and a distributive predicate, the semantic difference between the two expressions will be erased. Some of the examples we investigated in our fieldwork are arguably like this, and we focus in the text on those that are not.

- (i) a. If $\exists x =4(x)$ & x can be seen, I'll give you \$20.
 b. If $\exists x \geq 4(x)$ & x can be seen, I'll give you \$20.

(19) Inferential questions (sample)

Meaning 1: What is number of **shapes** in the contemporary drawing that will force the signer to give the addressee \$20? (provide an estimate, e.g. n, more than n, between m and n, etc)

Meaning 3: About the objects mentioned in Meaning 1 and **assuming the drawing's properties other than number of shapes fit the parameters of the signed sentence:** How strongly do you infer that the signer is committed to giving the addressee \$20 in case there are

(i) exactly 4 (ii) approximately 4 (iii) at least 4 (iv) at least 5 of these objects?

(20) *Context:* The addressee has a contemporary abstract drawing

IF POSS-2 DRAWING LOOK-LIKE ____, IX-1 1-GIVE-2 20 \$.

'If your drawing looks like this, with ..., I'll give you \$20.'

a. ⁷G

a G towards the top

b. ⁷G-unsep₄^{fast-circular}

at least 4 G's arranged in a circular fashion

c. ⁷G-unsep₄^{slow-circular}

exactly 4 G's arranged in a circular fashion

d. ^{3.3}G-sep₄^{fast-circular}

(deviant)

e. ⁷G-sep₄^{slow-circular}

exactly (or at least?) 4 G's arranged in a circular fashion

(ASL, [35.1560](#))

| Examples (ASL, 35.1560) | Number of objects that the repetitions stand for | | | |
|--|--|-------------|-------------|-------------|
| | = 4 | ±4 | ≥4 | ≥5 |
| a. G | 3.7 (1;4;6) | 3.7 (1;4;6) | 3.7 (1;4;6) | 3.7 (1;4;6) |
| b. G-unsep ₄ ^{fast-circular} | 6.7 | 6.7 | 6.7 | 7 |
| c. G-unsep ₄ ^{slow-circular} | 7 | 2.7 | 1.3 | 1 |
| d. G-sep ₄ ^{fast-circular} | 6.3 | 6 | 4.3 (2;4;7) | 4.7 (1;6;7) |
| e. G-sep ₄ ^{slow-circular} | 7 | 5.3 (7;5;4) | 4.7 (7;5;2) | 4 (7;4;1) |

Unseparated fast repetitions yield inferences consistent with an *at least 4* reading, with the result that no matter whether the drawing has 4 or more than 4 G's, the signer is committed to giving the addressee \$20. This is similar to one would obtain with the following English sentence: *If your drawing displays at least 4 G's arranged as a circle, I'll give you \$20.* Any number of appropriate G's at or above 4 would commit the signer to giving the \$20. One potential difference in the ASL sentence is that there might be a vague component about the threshold, so it might be (clumsily) paraphrased as 'at least approximately 4 G's'; this might explain the high endorsement in the ±4 column in (20)a.

Importantly, unseparated slow repetitions have an 'exactly' reading: it is now only in case the drawing has exactly 4 G's that the signer must give the addressee \$20. This suggests that this 'exactly' reading is not the result of a standard implicature, which would be expected not to arise within the scope of the *IF*-clause. But this leaves open two possibilities, discussed in greater detail below: one (our preferred option) is that this 'exactly' reading stems from the literal meaning of the construction due to its iconic semantics; the other is that this 'exactly' reading stems from a local implicature akin to that of numerals, which can be computed in the scope of *if*-clauses.

In addition, a difference arises in this context between slow unseparated and slow separated repetitions: the latter seem to be ambiguous, with an 'exactly 4' and an 'at least 4' reading, and unstable inferential judgments (unlike the other cases, where judgments are mostly stable). In fact, the example in (20)a with a lone G towards the top might raise a similar problem: there seems to be uncertainty as to whether the drawing should contain a G towards the top, and nothing else, or at least a G towards the top.

Related conclusions can be drawn from the same paradigm with iconic shapes replacing letters, as in (21).

(21) *Context*: The addressee has a contemporary abstract drawing

IF POSS-2 DRAWING LOOK-LIKE ____, IX-1 1-GIVE-2 20 \$.

'If your drawing looks like this, with ..., I'll give you \$20.'

a. ⁷o

a little circle towards the top

b. ⁷o-unsep₄^{fast-circular}

at least 4 little circles arranged in a circular fashion

c. ⁷o-unsep₄^{slow-circular}

exactly 4 little circles arranged in a circular fashion

d. ^{3.7}o-sep₄^{fast-circular}

(deviant)

e. ⁷o-sep₄^{slow-circular}

exactly (or at least?) 4 little circles arranged in a circular fashion

(ASL, [35, 1562](#))

| Examples (ASL, 35, 1562) | Number of objects that the repetitions stand for | | | |
|--|--|-------------|-------------|-------------|
| | = 4 | ±4 | ≥4 | ≥5 |
| a. o | 4 (1;4;7) | 4 (1;4;7) | 4 (1;4;7) | 4 (1;4;7) |
| b. o-unsep ₄ ^{fast-circular} | 6.7 | 6.7 | 6.7 | 7 |
| c. o-unsep ₄ ^{slow-circular} | 7 | 2.3 | 1.3 | 1 |
| d. o-sep ₄ ^{fast-circular} | 7 | 5.7 | 4.7 (2;5;7) | 5 (1;7;7) |
| e. o-sep ₄ ^{slow-circular} | 7 | 5.3 (7;4;5) | 5 (7;5;3) | 4.3 (7;5;1) |

Here too, fast unseparated repetitions are compatible with an 'at least reading'. A clear 'exactly' reading arises with unseparated slow repetitions. But once again there is a distinction between these and separated slow repetitions, which seem to allow both for an 'exactly' and an 'at least' reading. Some of our unembedded examples, for example in (17), displayed contrasts that went in the same direction (clear 'exactly' reading for unseparated slow iterations, additional possibilities for separated slow iterations); others did not display this contrast. This might be because an 'at least' reading was obfuscated by the presence of an implicature ('at least 4, but not at least 5, hence exactly 4').

4.2 *Intermediate conclusion*

It is worth pausing to state our main generalizations.

(22) **Readings obtained with repetition-based plurals**

- In all environments, 4 unseparated fast iterations give rise to an 'at least' reading ('at least 4' or 'at least 5' – the latter might particularly arise when the iterations are particularly hard to count).
- In all environments, 4 unseparated slow iterations give rise to an 'exactly 4' reading.
- In unembedded environments, 4 separated slow iterations give rise to an 'exactly 4' reading, and they seem to be ambiguous between an 'exactly' and an 'at least' reading under *IF*-clauses.
- In all cases, 4 separated fast iterations are degraded.

The contrast between separated and unseparated slow repetitions is a theoretical challenge. There are two directions one could explore.

Analysis I. Local implicatures in downward-monotonic environments: First, one could posit that both have an 'at least' literal meaning, but that in the case of slow unseparated repetitions it gets enriched by default through a local implicature, computed within the scope of the *IF*-clause, whereas this is optional in the case of slow separated repetitions. This line of analysis raises two

issues: (1) Why are locally computed implicatures possible or even near-obligatory in this case, despite the fact that the environment is downward-monotonic? (2) And why is there a difference in this respect between these two types of repetitions?

The first question might be less thorny than it looks, as 'exactly' readings also arise with garden-variety numerals in English inside *if*-clauses (e.g. *If Sam is 14 years old, he must be difficult to deal with*). The theory of numerals proposed by Spector 2013 is precisely that numerals differ from other scalar expressions in that they more easily associate with embedded exhaustivity operators, including in downward-monotonic environments. But the second question is much trickier, as it is entirely unclear why one construction should give rise to embedded implicatures in downward-monotonic environments while the other doesn't.¹⁸


Analysis II. A semantic difference between unseparated and separated repetitions: As an alternative, one could posit that the literal meaning of unseparated and separated slow repetitions is different: the former should have an 'exactly' reading, while the latter should be ambiguous between an 'exactly' and an 'at least' reading. We will develop a proposal along these lines in Section 9.

We turn to an explanation of these patterns, based on a combination of pictorial semantics (outlined in Section 5) combined with a theory of pictorial vagueness (Section 6). This will yield an initial derivation of the desired truth conditions (Section 7), which will be refined through a pragmatic procedure (Section 8) and a revision of one aspect of pictorial semantics (Section (10)). (The latter point will provide an argument for Analysis II.)

For most of the discussion, we will assume that the crucial factor is whether iterations are easy or hard to count, which justifies going back to the traditional terminology of 'punctuated' (and thus easy to count) vs. 'unpunctuated' repetitions (harder to count). But we will refine the analysis Section 9, where we try to explain a potential difference between unseparated and separated slow iterations.

4.3 Outlook: logical and pictorial conditions

In the rest of this piece, we will successively develop a simple pictorial semantics (to deal with punctuated repetitions) and a pictorial semantics with vagueness (to deal with unpunctuated repetitions). This pictorial component will then be integrated with a logical semantics involving quantification over pluralities that must resemble the repetition (in the precise sense of pictorial semantics). Various important refinements, including in the pragmatics, will be developed after the core semantics has been developed.

To illustrate, take the punctuated repetition in (8)c, repeated as (23)a. It involves four punctuated (here: unseparated slow) iterations of the manual *G* (i.e. ) , arranged in a circular fashion. We will posit the Logical Form in (23)a', where the boxed part is an existential quantifier over pluralities of *G*'s resembling the repetition, and introducing a variable *X* (ranging over pluralities).


- (23) a. POSS-1 DRAWING G -unsep $4_{\text{slow-circular}}$.
 'My drawing contains exactly 4 *G*'s (arranged in a circular fashion).'
 a'. $\exists X$ [my-drawing has X]
 b. $\llbracket (a') \rrbracket^{\text{s}} = \text{true}$ iff ... for some group d ,
 (i) $\llbracket G \rrbracket^{\text{f}, \text{s}[X \rightarrow d]}(d) = \text{true}$, i.e. d is a group of *G*'s, and
 (ii) G -unsep $4_{\text{slow-circular}}$ is (pictorially) true of d , and
 (iii) $\llbracket \text{my-drawing has } X \rrbracket^{\text{f}, \text{s}[X \rightarrow d]}(d) = \text{true}$, i.e. my drawing has d .

¹⁸ This problem pertains to the contrast between two types of 'easy to count' repetitions, slow unseparated and slow separated ones. A further issue for this theory is why unpunctuated repetitions, i.e. fast unseparated repetitions, should fail to give rise to an embedded implicature.

The schematic truth conditions we will derive are given in (23)b. (23)b(i), (iii) are standard logical conditions: the existential statement is made true only if one can find a group d which is made of G 's, and satisfies $[my-drawing\ has\ X]$ when the variable X denotes d . On the other hand, the condition (23)b(ii) will arise from a pictorial semantics: we will consider $G-unsep4_{slow-circular}$ as a picture, or rather a picture part, which can be true of certain objects, and we will require that d should be one such (plural) object. Our pictorial semantics will explicate in detail what it means for a picture part to be true of an object, and thus the schematic truth conditions in (23)b will in the end aggregate logical and pictorial conditions.

The details of the pictorial semantics will greatly matter at several junctures, but one point should already be kept in mind. The difference between a word-based and a purely iconic repetition is that in the former case (but not in the latter), the linguistic symbol need not resemble what it denotes, although the *arrangement* of the symbols provides information about the arrangement about the denoted group. This will be handled in due course by modulating the marking rule, i.e. the rule by which pixels of a picture are marked in view of what projects onto them. To anticipate: we will take the marking rule for the case of word repetitions to be semi-iconic and semi-symbolic. Assuming for simplicity that the picture/signing space on the left is made of 9 pixels, a ☞ will appear in a pixel just in case a G from the drawing projects onto it, as illustrated in (24). (The case of purely iconic repetitions will be simpler, with a pictorial component and no symbolic component at all.)

- (24) A semi-iconic and semi-symbolic marking rule: the ☞ 's do not resemble their denotations (G 's on a drawing), but their arrangement does provide iconic information about the arrangement of the denoted group.



5 Pictorial semantics

We turn to the first component of our analysis, which is based on a semantics for pictorial representations. But to see where we are going, let us first outline how a pictorial semantics


5.1 Basic pictorial semantics

Greenberg 2013, 2019 develops an explicit semantics for pictures, crucially based on the notion of a geometric projection: for him, a picture is true in a world w relative to a viewpoint v if w projects onto the picture relative to v . In Greenberg's words (2019),


a simple type of PERSPECTIVE PROJECTION is illustrated in [(25)a]. Here we begin with a concrete 3-dimensional region of spacetime (possible or actual), which I'll think of as a possible world. In the example below, the world contains only a cube. Next, a PROJECTION SOURCE is located within the space of the world. A projection source is thought of simply as a geometric point in space and time. This in turn defines a system of PROJECTION LINES, which link each point in the world to the source. Finally, a PICTURE PLANE is introduced into this spray of lines, and they are used to map spatially distributed features of the scene back to surface features of the picture plane itself— in this case, the lines of the line drawing. The result of such a projection is displayed at right below.

(25) **Examples of a projection-based semantics for pictures (Greenberg 2019a)**

a. Perspective projection



b. Parallel projection



As summarized in Schlenker 2019, the viewpoint includes both the projection source and the position of the picture. The relevant projection lines are given by the system of projection assumed: a system of "linear projection" is represented in (25)a, and is "characterized by the fact that the projection source is a point, to which all projection lines converge" (Greenberg 2019a); other systems can be used, for instance "parallel projection", illustrated in (25)b, whereby "the projection lines, rather than converging on a single point, all run perpendicular to the projection source, hence parallel to one another" (Greenberg 2019a).

Simplifying somewhat, this allows Greenberg to define the truth of a picture in a world relative to a viewpoint, as in (26).


(26) **Pictorial truth (Greenberg 2019a)**

A picture P is true in world w relative to viewpoint v along the system of projection S iff¹⁹ w projects to P from viewpoint v along S , or in other words:
 $\text{proj}_S(w, v) = P$

As Greenberg 2019a notes, a projection system has two components: a "projection condition", which "establishes the systematic spatial relationship between points in the scene and points on the picture plane, as embodied by projection lines" (parallel projection in the case illustrated in (25)); and a "marking" rule, which "associates properties of the scene with types of mark on the picture plane, following the spatial distribution specified by the first stage".

We will consider two simplified marking conditions, pertaining to black-and-white pictures, or equivalently pictures manually sketched with handshapes appearing in various parts of a 2-dimensional space. One marking condition is purely iconic, and corresponds to a case discussed by Greenberg 2019a, in which "the features associated with each point in the array correspond to those features of the environment which would project to the chromatic properties of the corresponding point in the picture. For example, in a system of line drawing, where edge-points project to line-points, every line-point on the picture plane will be associated with the feature edge in the corresponding point in the array". An illustration from Greenberg 2014 is given in (27): only visible edges are represented (as lines) on the picture. Further possible marking conventions are discussed in Greenberg 2014, but in the purely iconic case we will stick to edge marking because it suffices to provide a simple initial model of manual marking in signing space.

¹⁹ Greenberg 2019a is committed to the 'only if' part but not necessarily to the 'if' part, because further principles might enrich the content of a picture. In case some situations are plausible while others are extremely implausible, the latter might be excluded from consideration when computing the truth conditions of pictures. (Thanks to G. Greenberg for discussion of this point.)

(27) **Illustration of a simple marking convention: only visible edges are marked** (Greenberg 2014)

To be concrete, we will take the 2-dimensional space to be made of pixels, which may be filled or empty, as in (28) (one may think of 'filled' as 'black' and 'empty' as 'white', except that we be interested in signing space, where what matters is where parts of the hand appear).

(28) **Purely iconic marking rule**

Assuming a method of projection, mark a pixel as 'black' if and only if its projection line meets the edge of an object.

An alternative marking method is only semi-iconic, because the marking rule involves putting a linguistic symbol on a kind of map, following the statement in (29).

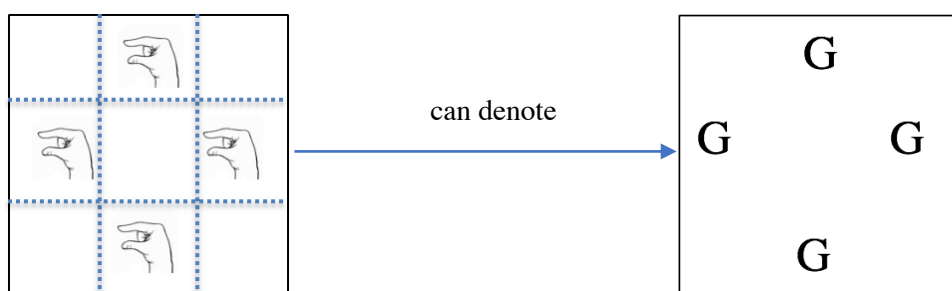
(29) **Mixed marking rule**

Assuming a method a projection, mark a pixel (which may be taken to be a large part of the 2-dimensional space) with a given word w just if and only if its projection line meets an object in the denotation of w .

To see the difference between these two methods, consider first the mixed marking rule in (29). We assume for greatest simplicity that the picture, here a vertical array in signing space, has only 9 pixels, and that projection is parallel. We restrict attention to the case discussed in the text: a G of the manual alphabet, i.e. the shape 𐄂 , appears on a (large) pixel just in case a G from the 2-dimensional contemporary drawing described can project onto it. We hasten to say that the representation on the left is a highly idealized representation of signing space. To give but an example, in unseparated repetitions the 𐄂 shape appears continuously through space (this is what makes the iterations unseparated!), but a little forward movement of the hand indicates which parts of space are *marked* with 𐄂 (see for instance the video associated with (8)).

(30) **Example of the mixed marking rule for our ASL example in (8)**

A highly simplified representation of signing space appears on the left, and the drawing which is discussed is represented on the right.




Importantly, the arrangement of the manual letters iconically represents the arrangement of their denotations, but the shape of the manual letters is entirely different from that of their denotations. This is because the marking rule is only semi-iconic: whether something gets represented in a pixel depends on the geometric projection, but what gets put in the pixel is a linguistic symbol.

Things are different in the purely iconic marking rule in (28), which is best illustrated with many more pixels, akin to a black-and-white photograph. Here the representation on the left can denote a drawing on the right just in case the latter genuinely resembles the former down to the details of the shapes involved. In the case discussed in the text, where a contemporary

drawing is represented in signing space, the manual representation will roughly resemble what appears in the drawing.

(31) **Example of the iconic marking rule for our ASL example in (12)**

A highly simplified representation of signing space appears on the left, and the drawing which is discussed is represented on the right.




5.2 Pictorial truth-of

The interpretive rule for repetition-based plurals in (2) makes use of the notion of a noun 'resembling' an object. This would be relatively easy to make precise if we had a definition of what it means for part of a picture to represent an object. But our definition of truth for pictures doesn't quite yield this: it just defines what it means for a picture to be true of a viewpoint and of a world, not what it means for a picture *part* to be true of *an object*, a viewpoint and a world. Fortunately, this slightly different notion was briefly explored in Greenberg 2014. We will call it *pictorial truth-of* to highlight the similarity with pictorial truth, but it is just the notion of a part of a picture *representing* an object. (The terminology we choose will be helpful because we will need to extend our analysis of truth for blurry pictures to truth-of.)

We follow Greenberg 2014 in taking as primitive a notion of grapheme, a "semantically significant" part of a picture, and in extending projection conditions to this notion to this notion of truth-of. We do not seek to define more precisely what it means for an object d to project onto a grapheme μ , but an illustration is given in (32).

(32) The picture on the left is true of the situation on the right, and the left-most shape denotes the top-most square (from Greenberg 2014)



The definition of pictorial truth-of can then be given in (33).²⁰

(33) **Pictorial truth-of (inspired by Greenberg 2014)**

A grapheme which is a part μ of a picture is true-of object d in world w relative to viewpoint v along the system of projection S iff d projects to μ from viewpoint v along S , or in other words:
 μ is true-of d in world w from viewpoint v along system of projection S iff $\text{projs}(w, v, d) = \mu$.

Truth-of will be useful in our final analysis to make precise part (ii) of the semantics in (2), i.e. the requirement that the repeated noun should 'resemble' its denotation: resemblance will be

²⁰ Greenberg 2014 proposes (in our terms) to define pictorial truth in terms of pictorial truth-of, but we will not be concerned with this reduction in the present piece.

cached out in terms of 'truth-of'. But before we get there, we will develop most of the analysis in terms of truth (rather than truth-of) of pictures.²¹

6 Blurry pictures

A key distinction in ASL repetition-based plurals is whether the iterations are easy or hard to count. The former case can be handled by the simple picture semantics summarized in the preceding section. The latter case will be handled by a theory of blurry pictures, to which we now turn.

6.1 Semantics through precisifications

How should blurry pictures be interpreted? We develop an account within a precisification-based framework. Precisifications are commonly used in the analysis of vague terms in trivalent semantics, and they usually involve supervaluations (e.g. van Fraassen 1966); we will see that refinements are needed in the case of pictures.

Take the schematic picture in (34)a, which has a blurry center and clear edges consisting of one cup each. We take the blurriness to consist in the fact that the center could correspond to one cup, yielding a total of three cups, as in (34)b, or to two cups, bringing the total to four, as in (34)c. From a precisification-based perspective, the basic intuition is that the blurry picture in (34)a can be made precise through (34)b as well as (34)c.

Notations: P is the blurry picture in (34)a, P₃ the precisification in (34)b and P₄ the precisification in (34)c. , P₃, P₄

(34) Two precisifications (in b. and c.) of a blurry picture (in a.)



Formally, then, we take a blurry picture to be associated with a set of precisifications, as stated in (35). In standard treatments of vagueness, precisifications are *interpretations*, but for simplicity we will take precisifications to be themselves non-blurry pictures. (It is clear that from a non-blurry picture one can infer an interpretation. We leave for future research a statement in terms of precise interpretations rather than precise pictures.)

²¹ This paper was already written when Greenberg 2019b appeared. It provides a semantics for tagging, a hybrid of symbols and pictures, used for instance in maps and in labelled pictures. We leave for future research the investigation of potential connections between the present piece and Greenberg's analysis of tagging.

(35) Precisifications of a blurry picture

A blurry picture P is associated with a set of precisifications $\Pi(P) = \{P_1, P_2, \dots\}$, each of which is a picture which (i) has the same frame as P , and (ii) is non-blurry.

The next step is to define the truth conditions of blurry pictures. There are two salient approaches to truth via precisifications. The most common view is based on strict truth, and it is associated with supervaluations. As stated in (36)a, a vague expression is strictly true just in case it is true on *each* of its precisifications. But we will have to make use of a more liberal notion, which is sometimes called tolerant truth (and is associated with 'subvaluations'). As stated in (32)b, a vague expression is tolerantly true just in case it is true on *at least one* of its precisifications (see for instance Varzi 1995, Hyde 1997, Cobreros et al. 2012).

(36) Truth of vague expressions

A vague expression is

- a. strictly true iff it is true on each of its precisifications; strictly false iff it is false on each of its precisifications.
- b. it is tolerantly true on at least one of its precisifications; tolerantly false iff it is false on at least one of its precisifications.

In view of the definition of falsity in (36)a, it is clear that an expression may be neither strictly true nor strictly false (in case it is true on some precisifications and false of others), which requires the use of truth gaps. And in view of (36)b, an expression may be both tolerantly true and tolerantly false (also in case it is true on some precisifications and false of others), which requires truth 'gluts'. Since the logics involved in the treatment of truth gluts are complex, we will use a simplified, bivalent notion of tolerant truth, defined in (37): on this analysis, an expression which is not tolerantly true is just false (see Cobreros et al. 2012 for a discussion of this approach and glut-based improvements on it).

(37) Bivalent tolerant truth

A vague expression is tolerantly true on at least one of its precisifications; tolerantly false otherwise.

On this simplified notion, an expression is always true or false, and it behaves in this respect much like a modal statement (with 'tolerantly true' interpreted as 'possibly true', 'tolerantly false' as 'not possibly false').

6.2 Using bivalent tolerant truth for blurry pictures

Restricting attention to truth (rather than falsity) conditions, we can apply both notions of the truth to blurry pictures, as in (38) (in a bivalent system, falsity conditions will immediately follow).

(38) Truth of a blurry picture

Relative to a world w , a viewpoint v and a system of projection S , a picture P associated with precisifications $\pi(P) = \{P_1, P_2, \dots\}$ is

- a. strictly true iff for each P_i in $\pi(P)$, $\text{proj}_S(w, v) = P_i$;
- b. tolerantly true iff for at least one P_i in $\pi(P)$, $\text{proj}_S(w, v) = P_i$.

In standard linguistic applications, one often makes use of strict truth. But in the present case, this won't do: on the strict notion, a blurry picture will almost never²² be true of anything. The reason is that the precisifications provide incompatible information. Concretely, consider again the blurry picture in (34)a, assuming for concreteness that it is taken by a camera. The precisification in (34)b requires that there should be exactly three trophies in front the camera,

²² We write 'almost never' rather than 'never' because there could be special cases to consider. For instance, with marking rule we will adopt in Section 9.2, a picture may be 'silent' on what happens in part of the situation of evaluation. If so, a different picture may be true of the same situation relative to the very same viewpoint (e.g. if the second picture is identical to the first one except that it provides information on a component of the situation that the first one is silent on).

whereas the precisification in (34)c requires that there should be exactly four trophies in front of the camera. These are incompatible conditions, so the picture will never be strictly true.

This pictorial case is different from standard cases of vagueness investigated in language. For *bald*, different precisifications establish different upper limits on the number of hairs that can be tolerated for the predicate to be true of the argument. If we call these precisifications *bald*_{≤n} for different values of the upper limit n, different precisifications are compatible with each other; for instance a person (call him Peter) who has 0 hair will count as 'bald' on all the precisifications, and thus *Peter is bald* will in that case be strictly true.

There are two conclusions one could draw. One is that the correct notion of truth in the case at hand is that of tolerant truth, not strict truth (since the latter is essentially useless). An alternative would be to take the primitive notion to be the information *entailed* by a vague picture, rather than the worlds in which it is true. We will see that in a certain sense, both solutions boil down to the same thing. But let us first define the notion of tolerant truth for blurry pictures we will be using throughout. For simplicity, we keep it bivalent, leaving a multivalent account for future research.

(39) Bivalent tolerant truth for a blurry picture

Relative to a world *w*, a viewpoint *v* and a system of projection *S*, a picture *P* associated with precisifications $\pi(P) = \{P_1, P_2, \dots\}$ is tolerantly true iff for at least one P_i in $\pi(P)$, $\text{proj}_S(w, v) = P_i$; tolerantly false otherwise.

Since our initial semantic rule for repetition-based plurals in (2) involves "discourse referents made available by the iconic semantics" of the repeated noun, it will be useful to have a corresponding notion of truth-of, because the discourse referents in question will be analyzed as picture parts that denote certain objects. An initial definition is given in (40). We assume that, just like pictures, parts of pictures are associated with precisifications. Specifically, if a picture *P* has precisifications P_1, P_2, \dots , then a part $\mu(P)$ of *P* has as possible precisifications the corresponding parts $\mu(P_1), \mu(P_2), \dots$ (the notion of a picture part could presumably be analyzed in terms of sets of pixels of an array, but we will take the notion as primitive here).

(40) Bivalent tolerant truth-of for a blurry picture – simple case

Relative to a world *w*, a viewpoint *v* and a system of projection *S*, a grapheme which is a part μ of a picture *P* associated with precisifications $\pi(P) = \{P_1, P_2, \dots\}$ is tolerantly true-of an object *d* iff for at least one P_i in $\pi(P)$, for the part μ_i of P_i that correspondsto μ , $\text{proj}_S(w, v, d) = \mu_i$. Otherwise, μ is tolerantly false of *d*.

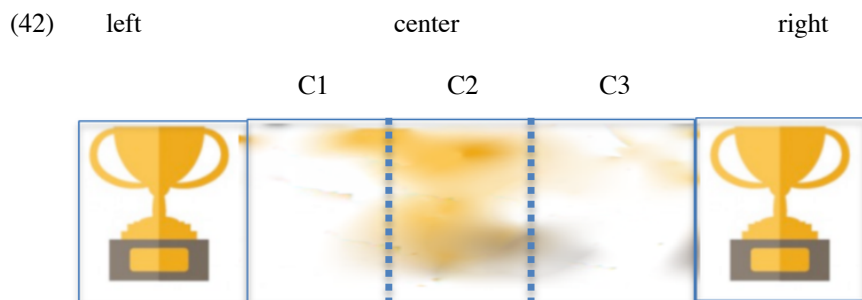
The semantics in (2) involves the simultaneous creation of several discourse referents, which will be analyzed as several picture parts that each have a denotation. This requires a refinement to the definition in (40) to the case of tuples of discourse referents, as defined in (41).

(41) Bivalent tolerant truth-of for a blurry picture – general case

Relative to a world *w*, a viewpoint *v* and a system of projection *S*, a sequence of graphemes $\langle \mu^1, \dots, \mu^n \rangle$ which are part of a spicture *P* associated with precisifications $\pi(P) = \{P_1, P_2, \dots\}$ is tolerantly true-of a sequence of objects $\langle d_1, \dots, d_n \rangle$ in world *w* iff for at least one P_i in $\pi(P)$, for the parts μ_i^1, \dots, μ_i^n that correspond to μ^1, \dots, μ^n respectively, $\text{proj}_S(w, v, d_1) = \mu_i^1, \dots, \text{proj}_S(w, v, d_n) = \mu_i^n$. Otherwise, $\langle \mu^1, \dots, \mu^n \rangle$ is tolerantly false of $\langle d_1, \dots, d_n \rangle$.

To see why this simultaneous evaluation of picture parts is needed, consider again the blurry picture in (34), with the rough division into parts displayed in (42). It makes sense to say that C1 and C3 can each denote a trophy (on the precisification in (34)c), and that C2 can denote a trophy as well (on the precisification in (34)b), but it won't do to say that C1, C2 and C3 can, together, denote 3 trophies, as this would bring the total to 5.²³ It is thus important to evaluate denotations of tuples simultaneously, as we do in (41).

²³ Technically, the problem is that we would get that the mereological sum $C1+C2+C3$ of the three relevant parts of the array can only denote 1 or 2 cups, whereas the sum of the denotations of C1, C2 and C3 includes 3 cups.



6.3 Bivalent tolerant truth via entailments

As we saw, strict truth as defined in (38)a is almost useless when it comes to blurry pictures. An alternative to using tolerant truth would be to take the primitive notion to be the information *entailed* by a vague picture, rather than the worlds in which it is true. Concretely, consider again the blurry picture in (34)a. Intuitively, it does not entail that there are exactly three trophies in front of the camera, nor that there are exactly four trophies, but it *does* entail that there are at least three, and in fact it entails that there exactly three or exactly four.

This notion could be cashed out in terms that are congenial to the strict notion of truth, as in (43). (43)a states the general case: a vague expression E entails an expression E' just in case for each precisification π of the interpretation, $\pi(E)$ entails $\pi(E')$. But we are interested in the more special case in (43)b, in which E is a picture and E' is a standard, non-vague propositional expression.

(43) Bivalent entailment for vague expressions

a. General case

A vague expression E entails E' iff for each precisification π , $\pi(E)$ entails $\pi(E')$. (Otherwise, E doesn't entail E' .)

b. Case in which E is a blurry picture and E' is not vague

Consider a viewpoint v and a system of projection S , let P be a picture with precisifications $\pi(P)$ and let E' be a non-pictorial, non-vague expression (which may depend both on a world and a viewpoint of evaluation).

P entails E' just in case for every world w , for every precisification P_i in $\pi(P)$, if $\text{projs}(w, v) = P_i$, E is true in w from viewpoint v . (Otherwise, P doesn't entail E' .)

One could thus approach the truth of blurry pictures indirectly: we start by finding all the entailments E' of a blurry picture P , using the recipe in (43)b. Then we decree that P is true in a world w just in case each of its entailments is.

Concretely, consider again the blurry picture P in (34)a, and let us restrict attention to four worlds w_1, w_2, w_3 and w_4 , with respectively 1, 2, 3 and 4 trophies arranged 'in the right way' in front of the camera. It is clear that P entails the proposition $\{w_3, w_4\}$ (or anything weaker), but nothing stronger. The indirect notion of truth specifies that P is true in a world w just in case each of its entailments is, or equivalently, just in case its strongest entailment $\{w_3, w_4\}$ is. In the end, P is true in w_3 and w_4 . But this is also what the tolerant notion of truth would give us.

As we show in Appendix I, the two notions of truth – bivalent tolerant truth, and truth via entailments – yield the same result. This gives us a modicum of confidence that this notion of truth is rather natural in the case at hand. It remains an overly simple notion, since it obliterates trivalence and genuine vagueness, and thus refinements should be explored in the future. But we henceforth use this notion of bivalent tolerant truth in our analysis.

6.4 Deriving 'at least' readings

With a theory of bivalent tolerant truth for blurry pictures in hand, we can derive the existence of 'at least' readings for unpunctuated repetitions, but we need one additional hypothesis: instead of

considering just two precisifications, as in our initial example in (34), we must consider infinitely many. Specifically, for n unpunctuated iterations, we will consider precisifications with i iterations for each $i \geq n$, as is stated in (44).

(44) **Precisifications to derive 'at least' readings**

If P involves n unpunctuated (i.e. hard to count) iterations of an expression E (which may be iconic or conventional), then $\pi(P) = \{P_i: i \geq n\}$, where P_i is a precisification with exactly i iterations of E .

To illustrate how (39) and (44) can be combined, consider (8)b. It involves 4 unseparated fast iterations of the manual letter G , which are thus particularly hard to count – a clear example of an unpunctuated repetition, which we write as $G\text{-rep}4_{\text{circular}}$. If we write as G^i the clear precisifications of $G\text{-rep}4_{\text{circular}}$ that involve exactly i iterations of G , arranged in the same (circular) shape as $G\text{-rep}4_{\text{circular}}$, we can define the precisifications specified by (44):

(45) **Precisifications of $G\text{-rep}4_{\text{circular}}$ (as in (8)b)**

$\pi(G\text{-rep}4_{\text{circular}}) = \{G^i: i \geq 4\}$, where G_i is a precisification with exactly i iterations of $G\text{-rep}4_{\text{circular}}$.

We can now say under what conditions $G\text{-rep}4_{\text{circular}}$ is (bivalently tolerant) true. Assuming the mixed marking rule in (29), each G^i will come with a requirement that there should be exactly i G 's arranged in a circular fashion. The bivalent tolerant truth conditions are given in (46).

(46) **Tolerant truth of $G\text{-rep}4_{\text{circular}}$ with the precisifications in (45)**

Relative to a world w , a viewpoint v and a system of projection S , $G\text{-rep}4_{\text{circular}}$ is tolerantly true iff for at least one G^i in $\pi(G\text{-rep}4_{\text{circular}})$, $\text{projs}(w, v) = G^i$,
iff (roughly) for some $i \geq 4$, there are exactly i G 's arranged in a circular fashion.
Otherwise, $G\text{-rep}4_{\text{circular}}$ is tolerantly false.

6.5 The Edge Effect revisited

Our approach based on precisifications makes it possible to revisit a constraint on anaphora noted by Schlenker and Lamberton 2019. As they argue, punctuated repetitions introduce by default one discourse referent per iteration, but arguably unpunctuated repetitions only introduce discourse referents at the edges (they call this the 'Edge Effect'). The main effect is illustrated in (47): a possessive pronoun indexing the right-most iteration of the unpunctuated repetition of *TROPHY* is acceptable, and refers to the right-most trophy; a possessive pronoun indexing the intermediate iteration is degraded. When the unpunctuated repetition is replaced with a punctuated repetition, any pattern of indexing is available.

(47) HERE HAVE TROPHY- $\text{rep}3_{\text{horizontal}}$.

'Here there are several trophies arranged in a row.'

a. ⁷ [no additional sentence]

b. ⁷ POSS-right SHAPE STRANGE.

'The right-most trophy has a weird shape.'

=> the right-most trophy has a weird shape

c. ⁵ POSS-middle SHAPE STRANGE.

'The intermediate trophy has a weird shape.'

=> the intermediate trophy has a weird shape


(ASL, 33, 0592; 3 judgments; Schlenker and Lamberton 2019)

In triangular representations, each vertex of the triangle introduces a separate discourse reference (hence the 'Edge' effect really pertains to vertices, but the latter terminology is a bit too technical to be optimally useful).

Schlenker and Lamberton 2019 suggest that the source of the contrast is that in unpunctuated repetitions the intermediate iteration (out of 3) is somehow blurry, as illustrated in

(48), but they do not give a formal account (see their discussion for connections with theories of iconically generated loci in arc-based plural pronouns).

(48) Successful pointing (on the left) and unsuccessful pointing (in the middle)



We propose that a constraint on pointing is that the outer frontiers of the area pointed to should be clear rather than blurry. This can be cashed out in terms of precisifications. In a nutshell, we will assume that pointing is always towards a part μ of a picture which (i) has a natural shape, and also (ii) has clear edges, i.e. parts that are edges in all precisifications of the blurry picture. This is made more precise in (49). Importantly, we leave open several notions, such as what a 'natural shape' is, and what it means for a point of a representation to be 'between' points of an edge. This is because these cognitive notions are likely complex, and here we are just interested in understanding how they interact with pictorial vagueness.

(49) **Availability for pointing of picture parts**

- a. A part of a picture is available for pointing only if (i) it has a natural shape, (ii) it is delimited by clear edges.
- b. A part μ of a (possibly blurry) picture P is delimited by clear edges just in case there is a set ε of pixels of μ such that for each precisification P_i in $\pi(P)$, in P_i :
 - each point of ε is marked;
 - each point of μ which is marked is either within ε , or included between two points of ε .

The intuition we develop is that leftward pointing in (48) is acceptable because the outer edges of the left-most trophy are sharp. Pointing towards the center is less acceptable because the outer edges of the intermediate trophy are not sharp. Still, we wouldn't want to prohibit all pointing towards a representation that has a blurry component, as one can certainly point towards the entire set of trophies in (48) (the same observation extends to sign language: it is clear that an unpunctuated repetition as a whole can introduce a discourse referent, even though its blurry central component doesn't²⁴). We could try to develop the analysis in semantic terms, by talking

²⁴ Still, the analogy with the behavior of a whole unpunctuated repetition in ASL shouldn't be exaggerated. It might be through grammatical means that a plural discourse referent is introduced by the entire set of repetitions. By contrast, iconic means might be needed to introduce singular discourse referents at the edges. This possible claim was tested by Schlenker and Lamberton 2019 with a paradigm that included (i). The idea was that plural reference to the entire set of sentences described by (i) could be effected grammatically, while singular reference to an individual sentence should require an iconic interpretation and an edge. As they wrote, "the first sentence evaluated on its own only provides weak iconic information. Adding the second clause with a right-edge possessive pronoun strengthens the iconic inference in (...) [(i)a)], but the effects are not entirely clear".

(i) YESTERDAY YOU WRITE _____ MEANING VAGUE.

'Yesterday you wrote sentences. The meaning of the right-most sentence/of all these sentences was vague.'

- a. ^{6,3} ___ =SENTENCE-rep_{3horizontal} = POSS-right
3/3 judgments: the sentences are more likely to be in a row when the 2nd clause is added
- b. ⁷ ___ =SENTENCE-rep_{3horizontal} = POSS-arc
1/3 judgment: the sentences are more likely to be in a row when the 2nd clause is added
(ASL, 34, 2256c,d; 3 judgments)

about the edges of the denoted groups, but at this point we find it simpler to talk about the edges of the representation rather than of its denotations (we leave a comparison for future research).

To apply the theory, we consider once again the blurry picture in (34)a with the two precisifications in (34)b,c, and consider in turn different patterns of pointing involving the three regions depicted in (50).

(50) Three parts of the blurry picture in (34)a



It is clear that left pointing is acceptable because the left-most part of the picture is entirely sharp: all precisifications agree on the left-most part, and as a result the outer edges of the left-most trophy are sharp. The same conclusion applies to the outer edges of the right-most trophy. Pointing towards the center won't work in the same way, however. The reason is that no set of pixels counts (on any reasonable definition) as a sharp edge of the intermediate trophy or trophies. This is so because the outer edges of the intermediate component in (51)a (continuous line) are just different from the outer edges of the intermediate component in (51)b (also represented with a continuous line), so there is no way these edges could be sharp. Furthermore, we could use as sharp edges the innermost edges of the left-most and right-most trophies (hatched lines), but the resulting area certainly doesn't correspond to a natural shape (it would be made of the intermediate trophy/trophies *with* the innermost edges of the extreme trophies). Finally, nothing in this account prohibits pointing towards the entire set of trophies: although their center is blurry, their edges are sharp, and thus the condition for pointing in (49) is satisfied. (In the triangular case, we would need to say that the three sharp vertices are enough to satisfy this condition, which makes intuitive sense.)

(51) Precisifications of the picture in (50)



Both the empirical and the theoretical side of the Edge Effect will have to be refined in future research. There are certainly alternative ways to develop the analysis. But we hope to have shown that an approach based on precisifications can help make potential theories sharper and possibly more explanatory.

6.6 Roles of vagueness

At this point, we have seen that blurry pictures can be analyzed by way of a precisification-based semantics. A notion of strict truth is of little interest because the various precisifications of a blurry picture usually make incompatible demands: one may specify that there are exactly three trophies in front of the camera, the other that there should be exactly four, and no world can satisfy both. We must thus appeal to a tolerant notion of truth, one on which a blurry picture is (tolerantly) true just in case one of its precisifications is true. The bivalent notion we used obliterates trivalence, since a blurry picture is either tolerantly true or tolerantly false, never indeterminate. We saw, however, that the same notion can be motivated in a different way, by starting from the propositional entailments of a picture. An additional benefit of the precisification-based approach

is to offer tools to state possible hypotheses about the source of the Edge Effect described by Schlenker and Lamberton 2019. While we have not attempted to give the final word on this issue, we used a notion of 'sharp edges' that might provide the beginning of a solution.

7 Revisiting the truth conditions of ASL repetition-based plurals

While we have provided an account of the truth of iconic representations involving repetitions, we have not explained how our initial definition of truth for sentences with repetition-based plurals should be amended in light of this analysis. Two components of the original definition in (2), repeated in (52), should be taken care of: the part pertaining to iconic representation in (52)(ii), and that pertaining to the discourse referents made available by the iconic semantics of the repeated expression in (52)(iii).

- (52) For $N\text{-iter}_X = N\text{-rep}_X$ or $[N N N]_X$ or $N\text{-cont}_X$, if P is a propositional expression $[[N\text{-iter}_X P]]^{f,s} = \text{true}$ iff for some group d ,
- (i) $[[N]]^{f,s[X \rightarrow d]}(d) = \text{true}$, and
 - (ii) $N\text{-iter}$ iconically represents d given c , and
 - (iii) $[[P]]^{f,s[X \rightarrow d, x_1 \rightarrow d_1, \dots, x_n \rightarrow d_n]} = \text{true}$, where x_1, \dots, x_n are discourse referents made available by the iconic semantics of $N\text{-iter}$, and x_1, \dots, x_n respectively denote d_1, \dots, d_n (which are parts of d).

We propose to make this interpretive rule more precise as follows (further refinements will have to be explored in the future).

- (53) For $N\text{-iter}_X = N\text{-rep}_X$ or $[N N N]_X$ or $N\text{-cont}_X$, if P is a propositional expression, $[[N\text{-iter}_X P]]^{f,s} = \text{true}$ iff for some viewpoint v and for system of projection S made available by c , for some group d ,
- (i) $[[N]]^{f,s[X \rightarrow d]}(d) = \text{true}$, and
 - (ii) $N\text{-iter}$ is tolerantly true of d , and
 - (iii) $[[P]]^{f,s[X \rightarrow d, x_1 \rightarrow d_1, \dots, x_n \rightarrow d_n]} = \text{true}$, where x_1, \dots, x_n is a maximal set²⁵ of parts of $N\text{-iter}$ which are available for pointing in $N\text{-iter}$ and $\langle x_1, \dots, x_n \rangle$ is tolerantly true of $\langle d_1, \dots, d_n \rangle$ and d_1, \dots, d_n are parts of d .

In the case, discussed by Schlenker and Lamberton 2019, in which N is a word (rather than an iconic representation), the system of projection S will have to involve the kind of mixed marking rule exemplified in (29).²⁶

To illustrate the new statement in (53)(ii), consider first a derivation of the truth conditions of (8)a, repeated as (55)a, and analyzed with the simplified Logical Form in (55)b. Because the repetitions are arranged in a circular fashion, there are no edges and we assume that the rule in (53)(iii) has no effect: no individual discourse referents are introduced. In line with our earlier discussions, we take the precisifications of $G\text{-unsep}_{fast-circular}^4$ to be all the clear (parts of) pictures with i iterations of G arranged in a circular fashion, for $i \geq 4$; this is defined in (54).

- (54) Precisifications of $G\text{-unsep}_{fast-circular}^4$
Writing G^i for i iterations of G arranged in the same circular shape as in $G\text{-unsep}_{fast-circular}^4$, the

²⁵ We add to (52)(iii) a maximality condition to make sure that all parts that are available for pointing are taken into account.

²⁶ Note that the mixed marking rule yields a partial redundancy between conditions (53)(i) and (53)(ii). Take for instance a circular punctuated repetition of manual G 's, i.e. \mathfrak{G} , as in (40). By (53)(i), the denoted group d will have to satisfy the condition that $[[\mathfrak{G}]]^{f,s[X \rightarrow d]}(d) = \text{true}$, or in other words: d is a (singular or plural) group of G 's. By (53)(ii), the \mathfrak{G} 's in signing space have to pictorially represent G 's, hence the same result. Future research could seek to eliminate the redundancy by doing without condition (53)(i), but one point should be kept in mind: (53)(i) makes reference to a context and an assignment function, whereas (53)(ii) doesn't. Thus if a repeated noun has a context- or assignment-sensitive component, getting rid of (53)(i) might create problems.

precisifications of $G\text{-unsep}4_{\text{fast-circular}}$ are given by:

$$\pi(G\text{-unsep}4_{\text{fast-circular}}) = \{G^i: i \geq 4\}$$

- (55) a. POSS-1 DRAWING $G\text{-unsep}4_{\text{fast-circular}}$
 'My drawing contains at least 4 or 5 G's (arranged in a circular fashion).'
 a'. $G\text{-unsep}4_{\text{fast-circular}} X$ [my-drawing has X]
 b. $\llbracket (a') \rrbracket^{c,s} = \text{true}$ iff for some viewpoint v and for system of projection S made available by c , for some group d ,
 (i) $\llbracket G \rrbracket^{c,s[X \rightarrow d]}(d) = \text{true}$, i.e. d is a groupe of G's, and
 (ii) $G\text{-unsep}4_{\text{fast}}$ is tolerantly true of d , and
 (iii) $\llbracket \text{my-drawing has } X \rrbracket^{c,s[X \rightarrow d]}(d) = \text{true}$, i.e. my drawing has d .
 In view of (54), (ii) is equivalent to: for some $i \geq 4$, $\text{proj}_S(w, v, d) = G^i$, i.e. (roughly) d is made of at least 4 G's arranged as a circle.

Let us now consider the case of 3 unpunctuated iterations, as in (47)a, repeated as (56)a, with the schematic Logical Form in (56)a', for a complex predicate P (which will contain the variable X). The precisifications of $TROPHY\text{-rep}3_{\text{horizontal}}$ are given in (56). The derivation of the truth conditions now involves the creation of new singular discourse referents, which could have an effect if P contained singular pronouns.

- (56) Precisifications of $TROPHY\text{-rep}3_{\text{horizontal}}$
 Writing $TROPHY^i_{\text{horizontal}}$ for i iterations of G arranged in the same horizontal shape as in $TROPHY\text{-rep}3_{\text{horizontal}}$, the precisifications of $TROPHY\text{-rep}3_{\text{horizontal}}$ are given by:
 $\pi(TROPHY\text{-rep}3_{\text{horizontal}}) = \{TROPHY^i_{\text{horizontal}}: i \geq 3\}$

- (57) a. HERE HAVE $TROPHY\text{-rep}3_{\text{horizontal}}$.
 'Here there are several trophies arranged in a row.'
 a'. $TROPHY\text{-rep}3_{\text{horizontal}} X$ F
 b. $\llbracket (a') \rrbracket^{c,s} = \text{true}$ iff
 (i) $\llbracket TROPHY \rrbracket^{c,s[X \rightarrow d]}(d) = \text{true}$, i.e. d is a group of trophies, and
 (ii) $TROPHY\text{-rep}3_{\text{horizontal}}$ is tolerantly true of d , i.e. for some $i \geq 3$, $\text{proj}_S(w, v, d) = G^i_o$, i.e. (roughly) d is made of at least 3 trophies arranged horizontally,
 (iii) $\llbracket P \rrbracket^{c,s[X \rightarrow d, x_1 \rightarrow d_1, \dots, x_n \rightarrow d_n]} = \text{true}$, where x_1, \dots, x_n is a maximal set of parts that are available for pointing in $TROPHY\text{-rep}3_{\text{horizontal}}$ and $\langle x_1, \dots, x_n \rangle$ is tolerantly true of $\langle d_1, \dots, d_n \rangle$ relative to v and S and d_1, \dots, d_n are parts of d .

By the same reasoning as in (50), only the two extreme iterations have clear edges and are thus available for pointing, which means that (iii) reduces to (iii'):

$$(iii') \llbracket P \rrbracket^{c,s[X \rightarrow d, x_1 \rightarrow d_1, x_2 \rightarrow d_2]} = \text{true}, \text{ where } x_1 \text{ and } x_2 \text{ are the extreme iterations of } TROPHY\text{-rep}3_{\text{horizontal}} \text{ and } \langle x_1, x_2 \rangle \text{ is tolerantly true of } \langle d_1, d_2 \rangle \text{ and } d_1 \text{ and } d_2 \text{ are parts of } d.$$

Here we have followed Schlenker and Lamberton 2019 in giving a static analysis of discourse referents, but to handle anaphoric examples like (47), we would need to add to the analysis a dynamic account of discourse referent transmission across sentences, as already argued for ASL in Schlenker 2011. Since the dynamification of the present system is a separate (and rather standard, if complex) issue, we do not discuss it further in this piece.

8 Pragmatic exploitation

8.1 Motivations

Our analysis based on vagueness and tolerant truth derives the main data but suffers from a drawback: when evaluating the number of iterations produced in cases of unpunctuated repetitions, it just isn't very plausible that there could be an arbitrary number of them. Take three blurry iterations, as in (34)a. Vagueness yields uncertainty about their number: there could be 3, or 4. In cases of unpunctuated repetitions that are realized as fast unseparated iterations, as in

(8)b, there might be greater uncertainty, e.g. 3, 4, 5 or 6 iterations. This wouldn't yield an 'at least' reading, but rather an 'approximately' reading. In other words, there is a discrepancy between the number iterations that one infers on iconic grounds (uncertainty between 3, 4, 5 and 6, for instance), and what the theory requires in order to derive an 'at least' reading (uncertainty among all numbers greater than or equal to 3). A more plausible view would be that there are various possible sets of precisifications, and that the one that yields the 'at least' reading is possible but very unlikely, while there are several others. But if we stopped here, we would generate far too many readings, so this measure would have to be constrained.

Another potential challenge for the analysis developed up to this point is that it just doesn't yield any 'approximately' readings. But some are arguably instantiated. Take (17)b, the unseparated fast repetition of a manual 𐄂 , which gradually opens from one iteration to the next (hence a preferentially iconic reading, pertaining to the very shape displayed). One of the dominant inferences is that the relevant drawing has approximately four of these shapes (rather than exactly four). Now this fact could be analyzed by way of a standard implicature: one would start from a reading akin to 'at least four or so' (where *or so* represents a vague component that we don't quite derive with bivalent tolerant truth). This would then compete with five unseparated fast iterations of the same kind, corresponding to a quantity of 'at least five or so', and by negating this more informative alternative, we might conceivably get something akin to 'approximately four'.

But this reasoning won't work as well in (8)c, the unseparated slow repetition of the manual *G*. The problem is that this expression doesn't seem to have an 'at least' reading. The dominant reading is an 'exactly' reading, but there also seems to be a weaker 'approximately' reading. To derive it, we need to use precisifications, but we shouldn't use the infinite set of precisifications that gives rise to the 'at least' reading. In this case, the small number of precisifications we considered in our initial formal example in (34)b would seem to do the job rather well.

The upshot is that on grounds of plausibility and in view of some of our data (admittedly ones that require further investigation), it makes sense to consider further possible sets of precisifications. But this raises an immediately challenge: if we consider that the two-membered set of precisifications in (34) is one possibility, while the infinite set of precisifications used in our 'official' theory is another, there should be numerous other readings as well, and we must find a way to constrain them.

We propose to address these problems by positing that there are indeed multiple possible sets of precisifications, but that different sets come with different probabilities. The set that gives rise to the 'at least reading' will, on iconic grounds, be very unlikely. But pragmatic reasoning will raise the availability of this reading. We develop the analysis within (a simplified version of) the Rational Speech Act (RSA) of pragmatic enrichment (e.g. Bergen et al. 2016), for the following reason: we wish to explain why a reading that is extremely unlikely on iconic grounds (the 'at least' reading of unpunctuated repetitions) becomes very likely when on reasons on what the speaker²⁷ might have wanted to say. This requires a strategic component: it is because the speaker has no better alternative to express the 'at least' reading that this reading is raised to prominence. One also needs a probabilistic component: an unpunctuated repetition can also have an 'approximately' reading, and the relative likelihood of the 'at least' and the 'approximately' readings will come from one's prior beliefs about what the speaker wishes to express.

More precisely, the structure of our pragmatic analysis is as follows. We start from the notion of tolerant truth applied to a blurry picture, but we add three components. One is that different readings are possible with different sets of precisifications. The second is that only some salient readings are considered. Notably, we will consider meanings like *exactly three G's*,

²⁷ In our theoretical discussions of the RSA model applied to sign language, we use the term 'speaker' because it is more common than 'signer' in this context, and also because as a matter of linguistic practice, Deaf signers do occasionally say things like *I'll speak with Sam*, thus extending the root *speak* in the same way as we do here.

approximately three G's, at least three G's, but not between 4 and 9 G's, for instance. The third and crucial component is that interpretation is enriched by pragmatic reasoning: we start from a literal addressee A_0 , then derive what a strategic speaker S_1 would do to convey the intended information to A_0 , and then how a strategic addressee A_1 would react to S_1 's utterance. To simplify the discussion, we assume that the speaker and addressee are both perfectly rational (within the RSA model, this corresponds to an infinite value for the 'temperature parameter' often notated as λ ; the RSA model often does not posit perfect rationality).

8.2 Varying sets of precisifications

We assume that a literal addressee who sees a pictorial signal P seeks to map it to a meaning, i.e. to truth conditions. There is uncertainty about the precisifications that are in fact used, and thus several sets of precisifications are considered, with different probabilities assigned to them. For the literal addressee, the probabilities of the various readings are just the probabilities of these sets of precisifications.

Consider again the blurry picture P , with three sets of precisifications:

$$\pi_1(P) = \{P_3\}$$

$$\pi_2(P) = \{P_4\}$$

$$\pi_3(P) = \{P_3, P_4\}$$

Each of these sets is assigned a probability, for instance the following (which sum to 1):

$$\Pr(\pi_1) = .1$$

$$\Pr(\pi_2) = .1$$

$$\Pr(\pi_3) = .8$$

This means that there are 3 meanings to consider:

with 10% chance, the meaning is that of P_3 ;

with 10% chance, the meaning is that of P_4 ;

with 80% chance, the meaning is that of the disjunction of $\{P_3, P_4\}$.

8.3 Literal addressee A_0

We will postulate that only precisifications corresponding to salient meanings can be considered, constrained by quantitative conditions corresponding to: *exactly n* , *approximately n* (simplified as: n or $n+1$), and *at least n* . Thus we take the literal addressee to infer possible meanings on the basis of (i) the distribution of probabilities on precisifications, and (ii) which meanings are considered possible. This is given by the equation in (58): the probability of a possible meaning M is the probabilistic weight of the corresponding precisification P relative to all the precisifications P' that yield usable meanings.

$$(58) \text{L(meaning is } M \mid \text{picture is } P) = \Pr(\text{precisification}(P) = M) / \sum_{M': M' \text{ a possible meaning}} \Pr(\text{precisification}(P) = M')$$

To make things concrete, we start from the probabilities in (59), where a given set of precisifications of 3 unpunctuated iterations has a greater probability to be interpreted with a meaning type pertaining to 3 (i.e. 'approximately 3', 'at least three', 'exactly three') than with a meaning type pertaining to 4 (i.e. 'approximately 4'). The probabilities are somewhat arbitrary but were picked so as to highlight the fact that an 'at least' interpretation is fairly unlikely (... and also to make the computations as simple as possible).

Notation: we write $=n$ for 'exactly n ', $\geq n$ for 'at least n ', and $\neq n$ for 'approximately n ', analyzed for simplicity as n or $n+1$.

(59) Probabilities of various sets of precisifications

| Precisifications | Corresponding meaning | Probability π of the precisifications |
|----------------------|-----------------------|---|
| $\pi_1(P) = \{P_3\}$ | $=3$ | .2 |
| $\pi_2(P) = \{P_4\}$ | $=4$ | .1 |

| | | |
|--------------------------------|------------------------|-----|
| $\pi_3(P) = \{P_3, P_4\}$ | ± 3 | .2 |
| $\pi_4(P) = \{P_4, P_5\}$ | ± 4 | .1 |
| $\pi_5(P) = \{P_3, P_4, P_5\}$ | not a possible meaning | .1 |
| $\pi_6(P) = \{P_4, P_5, P_6\}$ | not a possible meaning | .1 |
| $\pi_7(P) = \{P_k: k \geq 3\}$ | ≥ 3 | .1 |
| $\pi_8(P) = \{P_k: k \geq 4\}$ | ≥ 4 | .05 |
| $\pi_9(P) = \{P_k: k \geq 5\}$ | ≥ 5 | .05 |

We then compute the probabilities of different meanings for the literal addressee. Here it matters that π_5 and π_6 do not yield a possible meaning, as the quantitative conditions they give rise to are akin to *3, 4 or 5* and *4, 5 or 6* respectively (they appear in gray in the table). As a result, these meanings are not considered, and probability weights are redistributed among the other meanings by equation (58), yielding the results in (60).²⁸

(60) **Probability of the meanings obtained for 3 unpunctuated iterations**

For $i = 5, 6$ $\Pr(\text{picture has meaning } \pi_i) = 0$

For $i \neq 5, 6$,

$$\Pr(\text{picture has meaning } \pi_i) = \Pr(\pi_i) / (\Pr(\pi_1) + \Pr(\pi_2) + \Pr(\pi_3) + \Pr(\pi_4) + \Pr(\pi_7) + \Pr(\pi_8) + \Pr(\pi_9)) \\ = \Pr(\pi_i) / .8$$

| Precisifications | Corresponding meaning | Probability π of the precisifications | Probability π^* of the corresponding meaning, with $\pi^* = \pi / .8$ |
|--------------------------------|-----------------------|---|---|
| $\pi_1(P) = \{P_3\}$ | $=3$ | .2 | 1/4 |
| $\pi_2(P) = \{P_4\}$ | $=4$ | .1 | 1/8 |
| $\pi_3(P) = \{P_3, P_4\}$ | ± 3 | .2 | 1/4 |
| $\pi_4(P) = \{P_4, P_5\}$ | ± 4 | .1 | 1/8 |
| $\pi_7(P) = \{P_k: k \geq 3\}$ | ≥ 3 | .1 | 1/8 |
| $\pi_8(P) = \{P_k: k \geq 4\}$ | ≥ 4 | .05 | 1/16 |
| $\pi_9(P) = \{P_k: k \geq 5\}$ | ≥ 5 | .05 | 1/16 |

We must also provide an interpretive rule for a punctuated repetition. We take it to be unambiguously interpreted by an 'exactly' reading, as in (61):

(61) Probability of the meanings for a punctuated repetition

$$\Pr(\text{meaning is } =3 \mid \text{picture is } G G G) = 1$$

We assume that the interpretation of $G\text{-rep}_4$ and $G G G G$ is the same as $G\text{-rep}_3$ and $G G G$, *modulo* an increment of 1 for all the picture indices and meanings, e.g. by replacing the shaded row of (60) with:

²⁸ As noted by B. Spector (p.c.), it makes sense to posit that for the naive addressee A_0 the probability of a meaning reduces to the probability of the corresponding set of precisifications. This is because each set of precisifications corresponds to exactly one meaning. To illustrate, A_0 's probability of being faced with an 'exactly 3' meaning is the probability that the speaker is using the set of precisifications $\{P_3\}$. This is made clear in (i), where we use rep_3 for: the speaker uses 3 unpunctuated iterations; *meaning is =3* for: the intended meaning is 'exactly 3'; *precisification is $\{P_3\}$* for: the set of precisifications if $\{P_3\}$.

(i) $\Pr(\text{meaning is } =3 \mid \text{rep}_3)$

$$= \Pr(\text{meaning is } =3 \text{ and precisification is } \{P_3\} \mid \text{rep}_3) + \Pr(\text{meaning is } =3 \text{ and precisification is not } \{P_3\} \mid \text{rep}_3)$$

$$= \underline{\Pr(\text{meaning is } =3 \mid \text{precisification is } \{P_3\} \text{ and } \text{rep}_3)} * \Pr(\text{precisification is } \{P_3\} \mid \text{rep}_3) + \boxed{\Pr(\text{meaning is } =3 \mid \text{precisification is not } \{P_3\} \text{ and } \text{rep}_3)} * \Pr(\text{precisification is not } \{P_3\} \mid \text{rep}_3)$$

It is immediate that the underlined part is equal to 1, and that the boxed part is equal to 0 (since in (60) no set of precisifications besides $\{P_3\}$ yields an 'exactly 3' meaning). In the end, we have that $\Pr(\text{meaning is } =3 \mid \text{rep}_3) = \Pr(\text{precisification is } \{P_3\} \mid \text{rep}_3)$: the probability of being faced with the 'exactly 3' meaning is just the probability that the set of precisifications of $\{P_3\}$.

| | | | |
|----------------------|----|----|-----|
| $\pi_1(P) = \{P_4\}$ | =4 | .2 | 1/4 |
|----------------------|----|----|-----|

8.4 Strategic speaker S_1

Next, we consider the actions of a strategic speaker S_1 , who takes into account the behavior of L_0 to optimally convey the meaning that S_1 has in mind. We will just consider the 5 meanings that already appear in (60). We depart from the standard RSA model in positing a completely rational behavior on the part of the speaker (corresponding to $\lambda = \infty$ in the RSA model). We consider just 4 signals, involving 3 or 4 punctuated or unpunctuated repetitions, as in (62).

(62) Signals considered by S_1

G G G
 G-rep3
 G G G G
 G-rep4

The strategic speaker S_1 seeks to maximize the chance that the literal listener L_0 obtains the intended meaning m . This means that the selected signal s^* should be an s that maximizes the probability $L_0(\text{meaning is } m \mid \text{signal is } s)$; this is stated in (63).

(63) Behavior of the strategic speaker S_1

Given a meaning m , find a signal s^* such that

$s^* \in \text{Argmax}_s L_0(\text{meaning is } m \mid \text{signal is } s)$ (i.e. s^* maximizes the value of $L_0(\text{meaning is } m \mid \text{signal is } s)$), or in other words:

$L_0(\text{meaning is } m \mid \text{signal is } s^*) \in \text{Max}\{L_0(\text{meaning is } m \mid \text{signal is } s) : s \text{ is a possible signal}\}$

When there is only one s^* such that $L_0(\text{meaning is } m \mid \text{signal is } s^*) = \text{Max}_s(\text{meaning is } m \mid \text{signal is } s)$, the signal is fully determined, otherwise any probability distribution over the maximizing signals can be picked (e.g. one may randomize over these).

It is clear that the meanings =3 and =4 are best conveyed by $G G G$ and $G G G G$ respectively. 'Approximately n ' and 'at least n ' readings are best conveyed by $G\text{-rep}n$.

(64) Optimal choices to convey different meanings

| Target meaning | Optimal signal |
|----------------|----------------|
| =3 | G G G |
| =4 | G G G G |
| ±3 | G-rep3 |
| ±4 | G-rep4 |
| ≥3 | G-rep3 |
| ≥4 | G-rep4 |

At this point, it is worth noting that punctuated repetitions give rise to unambiguous 'exactly' meanings,²⁹ whereas unpunctuated repetitions are ambiguous: they can have 'approximately' or 'at least' readings. Importantly, the rest of the discussion will entirely derive from this result, and thus the details of the probabilities selected in (60) need not matter. In particular, the initial probability of an 'at least n ' interpretation of $G\text{-rep}n$ could have been taken to be even lower, say .01 instead of .1, as long as $G\text{-rep}n$ remained the best way of conveying 'at least n ' (this would involve tweaking the other probability assignments to ensure that, for instance, 'at least 4' is best expressed by $G\text{-rep}4$ rather than $G\text{-rep}3$, for instance).

²⁹ We will refine the analysis in Section 9 to explain why separated slow repetitions give rise to an ambiguity (despite the fact that they are easy to count and should be classified as punctuated).

8.5 Strategic addressee A_1

Finally, we consider the behavior of an addressee that takes into account the strategic speaker's behavior. At this point, this strategic addressee A_1 needs to have some prior beliefs about the meanings that are likely to be used by the speaker. A_1 will then compute the probability of the meaning in view of the signal by using Bayes's rule, stated in (65):

$$(65) \Pr(\text{meaning is } m \mid \text{signal is } s) = \frac{\Pr(\text{signal is } s \mid \text{meaning is } m) * \Pr(\text{meaning is } m)}{\sum_{m'} \Pr(\text{signal is } s \mid \text{meaning is } m') * \Pr(\text{meaning is } m')}$$

We make the following assumptions about the addressee's prior probabilities that different meanings are used. They incorporate the assumption that 'at least' and 'exactly' readings are more likely than 'approximately' readings.

(66) Assumptions about priors

| Target meaning | Optimal signal | Prior probability |
|----------------|----------------|-------------------|
| =3 | G G G | .2 |
| =4 | G G G G | .2 |
| ±3 | G-rep3 | .1 |
| ±4 | G-rep4 | .1 |
| ≥3 | G-rep3 | .2 |
| ≥4 | G-rep4 | .2 |

Using (65) and (66), we can compute the posterior probability of different interpretations by L_1 of our four target expressions, as in (67), with the result in

$$(67) \begin{aligned} \Pr(=3 \mid G G G) &= 1 \\ \Pr(=4 \mid G G G G) &= 1 \\ \Pr(\pm 3 \mid G\text{-rep}3) &= \frac{\Pr(\text{rep}3 \mid \pm 3) * \Pr(\pm 3)}{\Pr(\text{rep}3 \mid \pm 3) * \Pr(\pm 3) + \Pr(\text{rep}3 \mid \geq 3) * \Pr(\geq 3)} \\ &= \frac{1 * .1}{1 * .1 + 1 * .2} \\ &= 1/3 \\ \Pr(\pm 4 \mid G\text{-rep}4) &= 1/3 \quad (\text{by a similar computation as on the preceding line}) \\ \Pr(\geq 3 \mid G\text{-rep}3) &= \frac{\Pr(\text{rep}3 \mid \geq 3) * \Pr(\geq 3)}{\Pr(\text{rep}3 \mid \pm 3) * \Pr(\pm 3) + \Pr(\text{rep}3 \mid \geq 3) * \Pr(\geq 3)} \\ &= 2/3 \quad (\text{also obtained as: } 1 - \Pr(\pm 3 \mid G\text{-rep}3)) \\ \Pr(\geq 4 \mid G\text{-rep}4) &= 2/3 \quad (\text{by a similar computation as on the preceding line}) \end{aligned}$$

(68) A_1 's interpretation strategies

| Signal | Meaning | Posterior probability $\Pr(\text{meaning} \mid \text{signal})$ |
|---------|---------|--|
| G G G | =3 | 1 |
| G G G G | =4 | 1 |
| G-rep3 | ±3 | 1/3 |
| G-rep4 | ±4 | 1/3 |
| G-rep3 | ≥3 | 2/3 |
| G-rep4 | ≥4 | 2/3 |

It must be noted that a strategic speaker S_2 who determined the optimal signal in view of the interpretive behavior of A_1 in (68) would do the very same thing that S_1 did in (64), using punctuated repetitions in an unambiguous fashion (with an 'exactly' reading) and unpunctuated repetitions in an ambiguous fashion (with an 'approximately' or an 'at least' reading). Thus further levels of recursion in the strategic reasoning will not be modify the results obtained at this point.

8.6 *Intermediate conclusion*

Several points must be highlighted. First, despite the initial implausibility of an 'at least' reading for the naive addressee A_0 , this reading becomes highly available for unpunctuated repetitions. The reason is that this is the best way to convey the 'at least' reading. Second, an 'approximately' reading exists alongside the 'at least' reading, which is consistent with some of our data. Third, however, when an unpunctuated repetition is used, the 'at least' reading is more likely to be derived than the 'approximately reading'. But this is entirely due to the addressee's prior beliefs about which meanings are more likely to be employed. This incorporates a kind of linguistic assumption on the salience of certain readings: as it stands, our analysis of repetition-based plurals is not purely based on iconicity and cognition.

In sum, we have improved our initial model of the iconic semantics of repetition-based plurals, along two dimensions: we start from a purely iconic interpretation on which 'at least' readings are unlikely, but this possibility is exploited for communicative purposes; second, our analysis now allows for 'approximately' readings that our earlier analysis didn't derive.

Let us add a more general point. Because our pragmatic theory pertains to a mapping between meanings and constituents, it can function in embedded contexts just as well as in matrix contexts. It is in this sense different from theories that enrich the truth conditions of a sentence by reasoning on the situation the speaker is in, as is standard in neo-Gricean theories as well as in the RSA model: these only allow for the strengthening of the truth conditions of matrix clauses. By reasoning on the meaning the speaker wishes to associate with a signal, we allow the pragmatic process to affect in the same way embedded and unembedded constituents. This is a positive feature of the present analysis since all the relevant readings arise in the scope of *IF*-clauses, as we saw in Section 4.

Finally, it should be mentioned that this highly simplified analysis has one flaw: it disregards the fact that producing 15 iterations when one aims for an 'exactly 15' or an 'at least 15' reading might be... taxing. For punctuated repetitions, this need not radically affect the analysis: excessive cost may lead one to avoid large numbers of punctuated repetitions, but *when* they are used, it is plausible that they get 'exactly' readings. For unpunctuated repetitions, the problem is more severe: in cases (such as unembedded clauses) in which one can use an 'at least 3' reading (= 3 unpunctuated iterations) as well as an 'at least 15' reading (= 15 unpunctuated iterations), it is likely that one will go for the first option for reasons of cost. This problem, raised by a reviewer, is particularly sharp in production data from another sign language, Nicaraguan Sign Language: different numbers of punctuated iterations are used for different 'exactly' readings, but for (what seem to be) 'at least' readings, only 2-3 unpunctuated iterations are employed. To avoid cluttering an already complex discussion, we investigate the problem and its solution in Appendix II. Suffice it to say that a very simple cost component can account for these contrasts. A more involved analysis in which cost is integrated to our RSA model could also lead to an empirically interesting result: due to the trade-off between cost and benefit, 3 unpunctuated iterations might come to have an 'at least 4' reading, 4 unpunctuated iterations an 'at least 7' reading, etc, with small increments in the number of iterations leading in some cases to big increments in the thresholds.

9 **Modifying the marking rule?**

9.1 *A problem*

We return to the inferential data pertaining to embedding under *IF*-clauses in Section 4. These pose a puzzle. On the one hand, unseparated repetitions give rise to a clear pattern: 4 unseparated fast iterations yield an *at least 4* reading, as in (20)b, while 4 unseparated slow iterations give rise to an *exactly 4* reading, as in (20)c. Setting aside separated fast iterations, which are always deviant (e.g. in (20)d), it is striking that separated slow iterations, as in (20)e, give rise to entirely unstable inferential judgments: they seem to be compatible with an *at least* or with an *exactly*

reading. This is all the more surprising since, if anything, these iterations are easier to count than the unseparated slow repetitions which give rise to a clear *exactly* reading.

One could posit that these ambiguities are due to the optional presence of an existential quantifier over pluralities in the scope of the *IF*-clause. In a nutshell, starting from a reading akin to *exactly 4 G's forming a circle* as in (20)b, we could get an analysis such as (69), which just requires that it is possible to find a group of exactly 4 *G*'s arranged in a circle. With this existential quantifier, the antecedent will remain true even if there are supernumerary *G*'s in the drawing.

(69) If [$\exists X$: X is made of exactly 4 *G*'s arranged in a circle] your drawing displays X, I will give you \$20.

There are two reasons to doubt this account, however. The main worry is this: if the ambiguity of the separated slow iterations is due to the optional presence of an existential quantifier over pluralities within the *IF*-clause, why isn't the same possibility open with unseparated slow iterations as well? In our data, however, these do not give rise to instability in the judgments.

There is also a more minor worry. The predicate we used was *LOOK-LIKE*, which was designed to make the interpretation of the iterations as iconic as possible. Specifically, the analysis with an existential quantifier should really give rise to the truth conditions in (70).

(70) If [$\exists X$: X is made of exactly 4 *G*'s arranged in a circle] your drawing looks like X, I will give you \$20.

But it is not entirely clear that these truth conditions are satisfied in a situation in which there are 5 *G*'s in the picture. The reason is that the picture with the 4 separated slow iterations does *not* look like the iconic representation realized in signs. Rather, a subpart of the picture looks like this iconic representation, which should be something different.

9.2 An iconic solution

We propose that the problem we encountered is due to an ambiguity in the iconic conventions used by the signer. The question is what the absence of a marking in signing space corresponds to. Here we conjecture that there are two possibilities: if the hand does not appear in some area of signing space, this may be because no object projects onto that area, or because the representation is partial, i.e. says nothing at all about what may or may not project onto that particular area.

More specifically, we take the marking rule to have the following properties:

- (i) if the signer's hand traces a circle or a \curvearrowright in an area, there should be an object with appropriate properties (corresponding to a circle, or to a *G*) that projects onto it;
- (ii) if the signer's hand does not go down but does not trace anything, there should be no object that projects onto that area;
- (iii) if the signer's hand does not appear at all in a part of signing space, there could be an ambiguity: either (1) there should be no object that projects onto that part of signing space, or (2) anything goes, in the sense that the representation is silent on that part of the manual picture.


In the case of unseparated repetitions, the signer's hand does not go down between iterations of the circle or of the \curvearrowright . While the representation may be underspecified with respect to parts of space in which the hand does not appear, for those parts in which the hand is up but does not trace anything, there should be no object that project onto them.

In the case of separated repetitions, the hand goes down between the iterations. This gives rise to the following ambiguity pertaining to the space found between the iterations: according to (iii)(1), there should be no object with relevant properties that projects onto that space, hence an 'exactly' readings (= the objects arranged in a circle are exactly those that are depicted in signing space); according to (iii)(2), the representation is silent about the space that is found between the iterations, hence an 'at least' reading (= the objects arranged in a circle include at least those that are depicted in signing space).


If we write as a continuous line the parts of signing space in which the hand is up, we have the result depicted in (71)a: apart from the 4 circles traced in the representation, no object should project onto the continuous line, and this blocks a situation in which 5 rather than 4 circles are arranged in a circular shape. When the hand goes down between iterations, as in (71), a further reading arises, one on which the representation is silent concerning points of signing space in which the hand goes down (if we apply the rule in (iii)(2) above).

(71) Possible denotations with a liberal marking rule on which the representation is silent when the hand goes down

a. The hand doesn't go down between the repetitions of a circle



b. The hand goes down between the repetitions of a circle



Thus we propose to account for the ambiguity of separated but not unseparated repetitions by way of the ambiguity of the marking rule as specified in (72).

(72) **Marking rule for iconic sign representations**

A **manual** projective feature map is accurate at a scene *s* if and only if, for each **atomic part** *p* of the array,

a. **Strict version**

p is manually marked with feature *F* if and only if some object with feature *F* in *s* projects onto *p*

b. **Liberal version**

if the hand goes through an area surrounding *p*, *p* is manually marked with feature *F* if and only if some object with feature *F* in *s* projects onto *p*.

On the strict version, if the hand fails to mark the array in signing space with a feature, i.e. it fails to trace anything, no object should project onto it. This holds whether the hand is up or down in the relevant part of space. On this strict marking rule, the separated repetitions on the left-hand side of (71)b cannot denote the situation (itself a drawing) that appears on the right-hand side. On the liberal marking rule, on the other hand, the 4 separated iterations can denote the picture with 5 circles because the picture drawn in signing space is silent on those parts of the array in which the hand goes down. On both the strict and the liberal rule, the 4 unseparated iterations in (71)a fail to denote the drawing with the 5 circles on the right.

Importantly, the liberal version doesn't require that the hand go through an atomic part *p* for the 'if and only if' condition to be in force, but only *through an area surrounding p* – an expression which is left intentionally vague. The reason we use *through an area surrounding p* rather than *through p* is that otherwise our rule would be too weak. Consider the bunny at the top of the drawing on the right-hand side of (73). It includes an O as a proper part. Thus if the marking

rule made it possible to select arbitrarily what is represented and what is omitted, the representation on the left should be deemed accurate (it would just happen to be silent about the eyes, nose and ears of the bunny). Intuitively, this is the wrong result.

(73) **A possible denotation relation on an excessively liberal marking rule**



The requirement that the 'if and only if' condition should hold as soon as the hand goes through an area *surrounding* p is intended to block this case: since the hand depicts the top-most O, it should also depict any edges surrounding it, hence also the eyes, nose and ears of the bunny. Of course this condition would have to be made more precise in future work.

In sum, we have a possible solution to our puzzle. The problem was that separated repetitions are more ambiguous (between an 'exactly' and an 'at least' reading) than unseparated repetitions, despite the fact that the former are easier to count than the latter. We proposed a solution that crucially involves the details of a pictorial semantics. The main point is that when the hand goes down in a part of the array, this could be for two separate reasons depending on the marking rule: because no object projects onto that part of the array, or because the representation is silent on that part of the array. When the hand is up but fails to trace anything, no ambiguity arises, as both marking rules require that no object project onto that part of the array.

One issue we leave for future research is the development of a proper account of examples such as those in (74) (discussed in Schlenker and Lamberton 2019), where the numerals 4 or 10 co-occur with three punctuated iterations of *TROPHY*.

(74) MUSEUM HAVE

a. ⁷ 4 [TROPHY TROPHY TROPHY]_{horizontal_arc}.

'The museum has 4 trophies (spread out).'

b. ⁷ 10 [TROPHY TROPHY TROPHY]_{horizontal_arc}.

'The museum has 10 trophies (spread out).'

(Schlenker and Lamberton 2019, *ASL* 34, 2216, f, h; 3 judgments. Anonymized video of the entire paradigm: <https://youtu.be/029qrrytaPQ>. Note that a. and b. are the 6th and 8th examples; see Schlenker and Lamberton 2019 for a transcription of the full paradigm.)

Importantly, in this case the iterations are realized as unseparated and slow, and yet the number of trophies depicted is probably not interpreted literally, or else it would clash with the numerals.³⁰ The pictorial component seems to contribute the information that the trophies are spread out. One possibility that the repetitions are in this case a pictorial representation of a *subpart* of the denoted

³⁰ This depends on one's specific view on numeral semantics. The problem we refer to will arise with the rather natural semantics in (i), in which the effect of the numeral 4 is to introduce a maximal group of objects that satisfy the noun and the predicate (this is akin to the 'reference set' discussed in connection with anaphora in Nouwen 2003, Chapter 3). If *N-iter* is made of 3 slow unseparated iterations of *TROPHY*, it could not be a picture (part) true of a group of exactly 4 trophies, unless some trophies are somehow hidden.

(i) $[[[4 \text{ N-iter}]_x P]]^{c,s} = \text{true}$ iff for some group d ,

(a) d is the maximal group d' such that $[[N]]^{c,s[X \rightarrow d']}(d') = \text{true}$ and $[[P]]^{c,s[X \rightarrow d]} = \text{true}$,

(b) d contains exactly four objects, and

(c) *N-iter* is tolerantly true of d .

group (maybe because part of the objects are hidden from view); but why this is and how this phenomenon should be constrained remains to be seen.³¹

10 Possible extensions

Our analysis of repetition-based plurals in sign language relies on three modules: pictorial semantics, vagueness, and pragmatic enrichment (which interacts with a linguistic bias for some salient readings). All three modules could exist in a variety of communicative situations, and thus repetition-based plurals should be found in a variety of iconic forms. We consider potential extensions to sign language pluractionals, to repetition-based plurals and pluractionals in gestures, to vocal repetitions, and to pictures.

10.1 Connection with experimental data on classifier predicates³²

As noted by a reviewer, our conclusions on 'exactly' readings of punctuated repetitions (especially unseparated slow repetitions) dovetail with experimental results from Davidson and Mayberry 2015. In a paper devoted to the scalar implicatures in ASL, they used as a control two or three iterations of an animal classifier (transcribed as *CL:5(claw)*) in different parts of signing space, to represent either two bears, as in (75), or three bears (with one additional iteration of the classifiers).

(75) BEAR CL:5(*claw*)_Y, CL:5(*claw*)_Z.
'There is a (toy) bear here (at location X) and here (at Y).'



(Davidson and Mayberry 2015, Figure 2b)

ASL subjects included two populations: early Deaf learners, who strongly computed the 'not all' implicatures triggered by *SOME* (which forms a scale with ASL *ALL*), and late (non-native) Deaf learners, who computed this implicature less strongly, thus endorsing some under-informative sentences. Strikingly, both populations were at ceiling with respect to the 'exactly two' reading of the classifier in (75): if there were three objects on the table and a sentence described the situation with just two iterations of the classifier, it was rejected 100% of the time by early learners and 98% of the time by late learners. One possible interpretation is that the 'exactly' reading just isn't due to a scalar implicature but is the literal meaning of the construction: this is an 'exactly' reading because the meaning is based on a pictorial semantics, and that objects that fail to be represented are likely absent.

Relevant as well are experimental results from Davidson 2014, which compared underinformative sentences in English and in ASL. In both languages, subjects were presented with a situation with three objects, only two of which were mentioned. Thus in the presence of a candle, a globe and a wallet, they had to assess an English sentence such as *There is a candle*

³¹ This analytical direction need not conflict with the marking rule in (72)b if the 3 iterations of *TROPHY* depict a (possibly typical) *subpart* of the denoted group. In (74)a,b, the hand does not go down between the iterations, and this is presumably responsible for the inference that the trophies are spread out (by (72)b, the areas through which the hand goes without tracing *TROPHY* are ones in which there is nothing, hence the 'spread out' inference). Things are different in (71)b, where the hand goes down between the iterations, with the result that the marking rule in (72)b is silent about what there is between the depicted objects.

³² Many thanks to a reviewer for *Linguistics & Philosophy* for highlighting the relevance of the references mentioned in this section.

and a globe on the table. ASL native signers were asked to assess an ASL near-analogue, but one that crucially involving classifiers for the nouns.³³ In this case, there was "quantitatively a higher percentage" of rejections of the uninformative sentences in ASL (88% rejection) than in English (54% rejection; the difference between English and ASL was "marginally significant"). One possibility entertained by Davidson is that there is in this case a difference between the two languages in terms of the derivation of ad hoc implicatures. An alternative that could be investigated is that in ASL the 'exactly' reading is due to pictorial semantics, a different mechanism.³⁴

More generally, predicate classifiers are usually treated as lexical forms whose position or movement in signing space is interpreted highly iconically (e.g. Emmorey and Herzig 2003, Zucchi 2011, Davidson 2015). It would be natural in the future to extend to predicate classifiers (not just repeated ones, as in Davidson and Mayberry's work) the pictorial semantics developed in this piece.³⁵

10.2 Sign language pluractionals

In sign languages such as ASL and LSF, repetitions can be used with a pluractional meaning (referring to a plurality of events rather than entities) Kuhn 2015 particularly focuses on the interaction between logical and iconic conditions in this domain. He thus discusses examples such as (76), where *GIVE* is repeated at an accelerating pace to refer to an action that accelerates; Kuhn and Aristodemo 2017 make similar remarks about LSF.

(76) IF JOHN PAPERS GIVE-alt-speeding-up, SECRETARY WILL HAPPY.

'If John gives papers at an accelerating rate, the secretary will be happy.' (ASL, Kuhn 2015)

Focusing on ASL, Fischer 1973 notes the existence of two varieties of repetition-based pluractionals, which she calls 'slow' and 'fast' reduplication. These might correspond to what is termed punctuated and unpunctuated repetitions, but unfortunately we know of no detailed semantic investigation of the interpretive difference between them, and thus this is a topic we must leave for future research (but see relevant discussions of verbal reduplication in Klima and Bellugi 1979, Cormier 2002, Liddell 2003, among others).

10.3 Repetition-based plurals and pluractionals in gestures

Schlenker and Lamberton 2019 and Schlenker, to appear argue that the distinction between punctuated and unpunctuated repetitions can be replicated with gestures that have a plural or a pluractional reading (see Feldstein 2015 for an earlier attempt to study plurals in gestures). Based


³³ An ASL sentence cited in part in Davidson 2014 (Figure 3a) appears in (i):

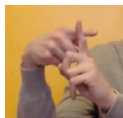
(i) THERE CANDLE CL-C_Y GLOBE CL:5(claw)_Z
'There is a candle and a globe.'

³⁴ Davidson 2014 sketches an account based on the interaction between Questions under Discussion and classifiers (which bias the question towards the configuration of the objects). This congenial to but distinct from the pictorial analysis envisaged here. As in our discussion of (74), however, one issue to keep in mind is that a picture need not depict the entirety of the situation if some objects are hidden from view; this might allow some ASL native signers to accept the under-informative statement.

³⁵ Tangentially, a reviewer asks about further uses of repetitions in ASL. Lamberton can think of these: (i) First, a case that is almost certainly unrelated to the present investigation: a single repetition (i.e. a total of 2 iterations) is a standard and well-known nominalization strategy in ASL, and serves for instance to distinguish *CHAIR* (nominal, 2 iterations) from *SIT* (verbal, 1 iteration). (ii) As in English, repetition can serve to intensify, as in *I never never never did that*. (iii) Far closer to our topic, there are temporal uses of repetitions: to say that someone won many trophies in their career, one might sign *TROPHY* repeatedly from one's shoulder down to a neutral space – hence: many trophies were won *over time*. We leave these diverse uses for future research.

on a detailed survey with three consultants, Schlenker, to appear a discusses in particular the contrast in (77), pertaining to plurals, and that in (78), pertaining to pluractionals.



Notation: *CROSS* represents the gesture , repeated in several parts of gestural space, in a punctuated fashion for *CROSS CROSS CROSS*, in an unpunctuated fashion for *CROSS-rep3*.

SLAP is a slapping gesture produced in neutral position (in front of the speaker). *SLAP, SLAP, SLAP* encodes three repetitions with clear pauses, *SLAP-repn* encodes n repetitions without clear pauses; *accelerating* and *decelerating* indicate that the repetitions are produced at an accelerating and decelerating speed respectively. (Judgments are averaged over 3 consultants on a 7-point scale, with 7 = best.)

(77) *Context:* The addressee is taking part in a treasure hunt in churches. The speaker provides an indication about the location of the treasure.

a. ^{6,7} If you enter a room and you see [*CROSS CROSS CROSS*]_{horizontal}, you have reached the prize.

=> when he reaches the prize, the speaker will see exactly three crosses (precise) arranged horizontally

b. ⁵ If you enter a room and you see *CROSS-rep3*_{horizontal}, you have reached the prize.

=> when he reaches the prize, the speaker will see ? three crosses or more (imprecise) arranged horizontally
(Video 4001, examples a. and c. <https://youtu.be/DXMsNtZ0-U>)

(78) My opponent, I am going to

a. ^{5,7} *SLAP, SLAP, SLAP*.

=> the speaker will slap his opponent exactly three times slowly [3/3 consultants] / in a very deliberate fashion [1/3 consultants]

b. ^{5,3} *SLAP-rep3*.

=> the speaker will slap his opponent two/three times [2/3 consultants] / exactly three times [1/3 consultants] / three times or more [1/3 consultants]

fast

(Video 4013, examples a. and b. <https://youtu.be/FQar0ELhrBA>)

If confirmed, these and other data (pertaining to the arrangement of the iterations in iconic ways) display a striking similarity to the sign language data in Schlenker and Lamberton 2019. This similarity is expected from the present perspective, but requires a more detailed empirical investigation.

10.4 Repetitions in the vocal modality

10.4.1 Onomatopoeias³⁶

The distinction between punctuated and unpunctuated repetitions can arguably be replicated in the vocal medium. In the context of a shooter video game, one may use an onomatopoeia akin to *prh* to evoke a gunshot. But our impression is that different quantitative implications are obtained

³⁶ Thanks to Jeremy Kuhn (p.c.) for helpful remarks.

depending not just on the number of repetitions, but also on their realization, as illustrated in (79): if 4 iterations are separated by pauses and easy to count, one can get an 'exactly four' reading. If the 4 iterations are produced without pauses and are harder to count, one might get instead an 'approximately four' or an 'at least four' reading.

- (79) *Context*: the addressee is playing a shooter video game.
- a. In the next 10 seconds, you should [prh - prh - prh - prh]_{slow, pauses}
 - b. In the next 10 seconds, you should prh-rep4_{fast, no pauses}
 - c. If in the next 10 seconds you [prh - prh - prh - prh]_{slow, pauses}, I'll give you \$20.
 - d. If in the next 10 seconds you prh-rep4_{fast, no pauses}, I'll give you \$20

We describe in Appendix II a mini-survey (with 6 linguists) that was designed to test these impressions, with inferential questions that were very close to those we used in ASL. The findings pertaining to punctuated repetitions indeed argue for the existence of 'exactly' readings, including in the scope of *if*. The findings pertaining to unpunctuated repetitions are harder to interpret, in part because several consultants interpreted them as involving a machine gun in which pulling the trigger once gives rise to a *burst* of shots. When this is taken into account, it seems that 'at least 4' readings were sometimes obtained, although additional readings appeared as well (possibly including an 'at least 1' reading). We refer the reader to Appendix II for details.³⁷

10.4.2 Repetition-based plurals and pluractionals in speech

As mentioned in Pfau and Steinbach 2006 and in Kuhn and Aristodemo 2017, reduplication is a common means of plural and pluractional formation in spoken language, as illustrated in (80).

- (80) a. Plural formation by reduplication
Warlpiri: *kurdu* 'child' => *kurdu-kurdu* 'children' (Pfau and Steinbach 2006, citing Nash 1986)
 b. Pluractional formation by reduplication
Hausa: *kiraa* 'keep on calling' => *kirkiraa* 'keep on calling'/'call many people' (Kuhn and Aristodemo 2017, citing Newman 2012)

It would be interesting to explore in the future what relation, if any, these bear to the repetition-based plurals discussed in this piece.

10.5 Comics and cartoons

We conjecture that hard-to-count iterations might be used in comics to evoke pluralities of indefinite numerosity. A possible example is reproduced in (81), where the number of coins appearing at the surface of the 'sea of money' is indeterminate, but large, and greater than the number of individual coins that can be distinguished (see Cohen 2013, chapter 2, for other uses of repetition/reduplication in comics).

³⁷ Several directions could be explored in future research. First, as suggested by J. Kuhn (p.c.), it would be useful to create stimuli by splicing one and the same realization of an onomatopoeia in several linguistic contexts so as to ensure that the repeated sounds in (79)a,b (or (79)c,d) are genuinely identical. Second, it would be useful to investigate nominal (rather than pluractional) uses of repeated onomatopoeias so as to have a possibly closer point of comparison for ASL plurals. In a different context, Guerrini and Migotti 2019 used repeated onomatopoeias pluralities of champagne bottle (the sound was evocative of a cork popping). We are not sure how to realize these sounds in an unpunctuated fashion, but this direction would be worth exploring.

(81) Logo for Disney's *Ducktales*³⁸



11 Conclusion

11.1 Main results

We started with a puzzle: the distinction between punctuated and unpunctuated repetitions seems to be invented and reinvented without effort across signs and gestures, in some cases without obvious input. Within ASL, the distinction arises in purely iconic representations just as it does in words, including in cases in which the repetition involves a gradual modification of the repeated shape (unlike in standard grammatical cases). A related distinction can arguably be found in other mediums, such as vocal onomatopoeias. While these observations could lead to the conclusion that, in this domain, Universal Grammar specifies a specific mapping between form and meaning, we developed an alternative account based on the interaction between different modules. In our final analysis, an iconic semantics was combined with a theory of pictorial vagueness to yield multiple possible readings (depending on the set of precisifications considered), including an 'exactly' reading for punctuated repetitions, and an extremely unlikely 'at least' reading akin to plurals for unpunctuated repetitions. By pragmatic reasoning, the latter possibility was exploited for communicative purposes, making it dominant in some cases (while still allowing for an 'approximately' reading). More speculatively, we argued that an unexpected ambiguity found with slow separated iterations (but not with unseparated iterations) can be traced to the details of the 'marking rule' used in pictorial semantics.

Two broader conclusions can be drawn. First, a full understanding of linguistic meaning requires a detailed theory of iconic representations (in fact, of blurry iconic representations), including in an area thought to be as grammatical as plurals. Second, there might be a counterpart of plurals in non-linguistic systems because the combination of means used in unpunctuated repetitions (iconicity, vagueness, pragmatic exploitation) are available more broadly.

11.2 Open questions

Our analysis leaves open several questions. On an empirical level, our data ought to be tested with further consultants in ASL and extended to other sign languages. It would also be interesting to see if similar distinctions can be found in repetition-based pluractionals – something we sketched for gestural pluractionals but not for sign pluractionals.³⁹ Various extensions beyond signs and gestures should be considered as well, as we sketched for vocal onomatopoeias and comics; music could be considered as well. Two components of our analysis particularly require further

³⁸ Retrieved on November 23, 2019 from [https://en.wikipedia.org/wiki/DuckTales#/media/File:DuckTales_\(Main_title\).jpg](https://en.wikipedia.org/wiki/DuckTales#/media/File:DuckTales_(Main_title).jpg).

³⁹ We also said nothing about the potential role of speed of repetition *per se*. In Schlenker and Lamberton 2019, it was mentioned that, all other things being equal, greater speed might be associated with larger quantities. In the present theory, speed only plays an indirect role, by making the iterations harder to count. The theory might have to be refined to account for a more direct relation between speed and quantities.

investigation: the existence of 'approximately' readings in downward-monotonic environments; and the subtle semantic distinction between unseparated vs. separated slow repetitions.

On a theoretical level, several important questions require further research. First, for reasons of simplicity we adopted a bivalent tolerant interpretation of blurry pictures, but this eliminates vagueness *per se* (which would require that unclear cases be deemed 'neither true nor false' or 'both true and false'). As a result, both 'approximately' and 'at least' readings came with clear numerical conditions, which seems inaccurate (there should be clear and unclear cases, as in genuine theories of vagueness). More sophisticated accounts ought to be explored. The pragmatic component of our analysis was also reduced to its bare bones: besides assuming perfect rationality, we made rather strong assumptions about the meanings considered (only 'exactly', 'approximately' and 'at least' readings); in effect, this implemented a linguistic bias in pictorial semantics/pragmatics, since in principle many further potential readings could have been considered. It would be interesting to explore what happens when these strong assumptions are relaxed (see Appendix II for a discussion of ways to add a cost component to the analysis).

Appendix I. Blurry pictures: truth via entailments vs. bivalent tolerant truth

We reiterate in (82) our definition of bivalent tolerant truth for blurry pictures, and in (83) the definition of truth via entailments for blurry pictures. For readability we boldface propositions, and thus if P is a picture (evaluated relative to a viewpoint v and a system of projection S), we will write as \mathbf{P} the set of worlds that make P true.

(82) **Bivalent tolerant truth for a blurry picture**

Relative to a world w , a viewpoint v and a system of projection S , a picture P associated with precisifications $\pi(P) = \{P_1, P_2, \dots\}$ is tolerantly true iff for at least one P_i in $\pi(P)$, $\text{projs}(w, v) = P_i$; tolerantly false otherwise.

(83) **Truth via entailments - case of a blurry picture**

Relative to a world w , a viewpoint v and a system of projection S , a picture P associated with precisifications $\pi(P) = \{P_1, P_2, \dots\}$ is true* iff for each proposition \mathbf{E} such that P entails \mathbf{E} , \mathbf{E} is true in w ; false* otherwise.

In this case, bivalent tolerant truth yields the same result truth via entailments, as stated in (84).

(84) Let P be a picture evaluated relative to a world w , a viewpoint v and a system of projection S , and let $\pi(P) = \{P_1, P_2, \dots\}$ be the set of its precisifications.
 P is true* in w if and only if P is bivalently tolerantly true in w .

Proof

Relative to a viewpoint v and a system of projection S :

P is true* in w iff [1] for each proposition \mathbf{E} such that P entails \mathbf{E} , $w \in \mathbf{E}$
iff [2] for each proposition \mathbf{E} such that for each precisification P' of $\pi(P)$, P' entails \mathbf{E} , $w \in \mathbf{E}$
entails \mathbf{E} , $w \in \mathbf{E}$ iff [3] for each proposition \mathbf{E} such that for each precisification P' of $\pi(P)$, $\mathbf{P}' \subseteq \mathbf{E}$, $w \in \mathbf{E}$
iff [4] $w \in \bigcup_{P' \in \pi(P)} \mathbf{P}'$ ⁴⁰
iff [5] for some $P' \in \pi(P)$, $w \in \mathbf{P}'$
iff [6] P is tolerantly true in w .

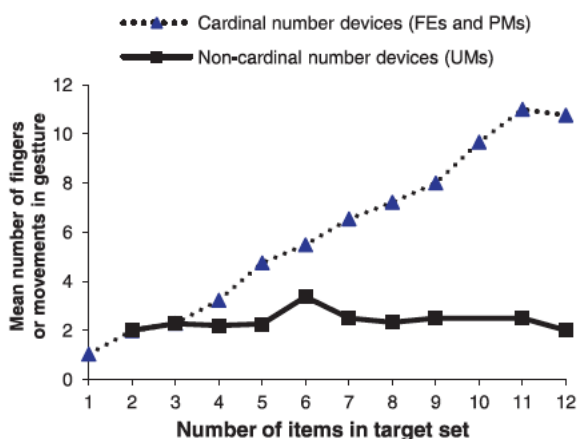
⁴⁰ In greater detail: assume w satisfies [3]. For $\mathbf{E} = \bigcup_{P' \in \pi(P)} \mathbf{P}'$, for each precisification P' of $\pi(P)$, P' entails \mathbf{E} , hence $w \in \mathbf{E}$, i.e. [4]. Now assume [4]. Thus for some $P' \in \pi(P)$, $w \in \mathbf{P}'$. It follows that for each proposition \mathbf{E} such that for each precisification P' of $\pi(P)$, $\mathbf{P}' \subseteq \mathbf{E}$, $\mathbf{P}' \subseteq \mathbf{E}$ and therefore $w \in \mathbf{P}'$.

Appendix II. Adding cost⁴¹

□ An apparent problem

Both the iconic and the pragmatic component of our analysis could be expected to apply to other systems, in particular other sign languages, and also homesigns. As a reviewer notes, the production data of Nicaraguan homesigners reported in Coppola et al. 2013 might seem to raise a worry for our analysis. In a production experiment, homesigners "were videotaped describing ten 10- to 20-s vignettes with varying numbers of objects and events". The authors coded three types of number marking: "Finger Extensions (FEs) were handshape configurations in which each extended finger indicated a member of the target set"; "Punctuated Movements (PMs) were a series of discrete movements, each referring to an entity or action in the vignette"; "Unpunctuated Movements (UMs) were movements produced in rapid succession with no clear break between them". The latter two categories seem to correspond rather closely to the punctuated and unpunctuated repetitions we described for ASL. In view of the model of Section 8, we would expect that, when they are used, punctuated repetitions should only have 'exactly' reading. This is roughly consistent with the data in (85): when punctuated movements were used, there were roughly as many iterations as there were objects to describe. What might be concerning, on the other hand, is that the number of unpunctuated iterations does not seem to depend on the number of objects to be described.

(85) Use of finger extensions (FE), punctuated movements (PM) and unpunctuated movements (UM) in a production experiment on Nicaraguan homesigners (Coppola et al. 2013)



There are two issues to be addressed. One is the apparent absence of 'approximately' readings of unpunctuated repetitions. Since this is not the dominant reading we derive in (68) (in essence, because we postulated that such a reading would be used less often than the 'at least' reading), we might posit that it is either unavailable or not deemed useful in this production situation and/or in this language. With 2-3 unpunctuated iterations, the 'at least' readings we derive in (68) are correctly predicted to be true in the situations in which they are used (for present purposes, we treat two unpunctuated iterations along the same model as the three unpunctuated iterations in (68), with an 'at least 2' reading). But this leaves another question open: why do subjects not produce 'at least 3' readings (with 3 unpunctuated iterations) when there are 3 objects or more, 'at least 4' readings (with 4 unpunctuated iterations) when there are 4 objects or more, etc?

One possibility is that homesigners' unpunctuated repetitions are just a realization of English-style plurals, grammaticalized as 2-3 unpunctuated iterations. This is unlikely, since the Coppola et al. 2013 write that "homesigners' Unpunctuated Movements often captured the global

⁴¹ Many thanks to a reviewer for *Linguistics & Philosophy* for raising questions addressed in this appendix.

spatial properties of an array of objects or a series of events (as well as their multiplicity), as do analogous plural predicates in ASL". A second possibility is that these are plurals with an iconic component, as in some of the analyses of Schlenker et al. 2019; this need not predict that larger number of unpunctuated iterations can be produced. A third possibility, however, is that the analysis we developed for ASL applies to homesigners as well, but that the pragmatic component of the theory must be refined. In essence, we will now argue that long unpunctuated forms are avoided in this production situation because they are costly; shorter unpunctuated forms still convey useful information, but at a lesser cost, and they are thus preferred.

This general issue highlights the absence of a cost component in our main analysis: intuitively, the longer the expression, the greater its cost. For unpunctuated repetitions, when there are 4 objects, one may elect to produce 2 unpunctuated iterations, or 3, or 4, as the resulting readings will all be true, but one might prefer to produce just 2 iteration in order to minimize cost. For punctuated repetitions, one has no choice: because only 'exactly' readings are produced (see (68)), one can only produce as many iterations as there are objects. Cost may lead one not to use these forms at all, but *if* one uses them, there will be as many iterations as objects.

□ *External cost module: unembedded clauses*

The simplest way to take cost into account is as a separate module from the rest of our analysis. For simplicity, we illustrate it for the case in which there are 10 objects, and we just compare the utility of different numbers of unpunctuated iterations (we disregard punctuated iterations because we already know that *if* they are used, their number must correspond to the number of objects present). Since the data in (85) only pertain to unembedded clauses, we restrict attention to the choice of a quantitative expression in an unembedded clause. The absence of embedding matters because it entails that when there are 10 objects, one can truly use an 'at least 2', an 'at least 3', ..., and an 'at least 10' reading.

We write as $U_{10}(\text{rep}_x)$ the utility of producing x unpunctuated iterations in this situation involving 10 objects, as $B_{10}(\text{rep}_x)$ the informational benefit of doing so, and as $C(\text{rep}_x)$ the corresponding cost (which only depends on the number of iterations used). We assume for simplicity that the x unpunctuated iterations only yield an 'at least x ' reading (thus disregarding the 'approximately x ' readings we derived in Section 8). As is standard, the utility is the benefit, i.e. the informational value, minus the cost (see (86)a). As stated in (86)b, if there are more than 10 iterations, the 'at least x ' reading is false and thus the benefit is $-\infty$ (which means that one never says something one knows to be false). For simplicity, we assume in (86)c that the cost is proportional to the number of iterations (with unit cost .5).


- (86) Utility $U_{10}(\text{rep}_x)$, benefit $B_{10}(\text{rep}_x)$ and cost of x unpunctuated repetitions to describe a situation with 10 objects
- a. $U_{10}(\text{rep}_x) = B_{10}(\text{rep}_x) - C(\text{rep}_x)$
 - b. If $x > 10$, $B_{10}(\text{rep}_x) = -\infty$ because the 'at least x ' reading is false.
 - c. $C(\text{rep}_x) = .5x$

To determine how many unpunctuated iterations are produced, we need some assumptions about how the benefit of using x iterations changes as x varies between 2 and 10 (beyond 10, the benefit falls to 0 because the reading is false). It is natural to posit that this is an increasing function: when there are 10 objects, it is better to say that there are at least 9 than to say that there are at least 2 (as both statements are true but the former is more informative). Three cases are worth discussing, illustrated in (87). In (87)a,b, one gets a fixed benefit (of 5) for saying something true, but in addition one gets additional 'points' as the number x of iterations nears 10. The difference is that in (87)a the benefit grows more quickly than the cost, with the result that the optimal strategy is to produce 10 iterations; whereas in (87)b, the benefit grows less quickly than the cost, and thus the optimal strategy is to produce the minimum number of iterations, here 2. The latter case would appropriately capture the production data from Nicaraguan Sign Language.

(87) $B_{10}(\text{rep}_x)$ (top line/curve) vs. $C(\text{rep}_x)$ (bottom line) for different values of n , $2 \leq x \leq 10$

In all cases, $C(\text{rep}_x) = .5x$. We consider several values of the benefit function.

In c., we include an additional curve corresponding to the utility, i.e. $U_{10}(\text{rep}_x) = B_{10}(\text{rep}_x) - C(\text{rep}_x)$;



In (87)c, something more subtle happens: initially (for low values of x), the benefit grows more quickly than the cost, but then it grows less quickly than it. As a result, the optimal number of iterations is of 4 or 5 (because the maximal utility is obtained for x near 4.5). This situation arises with a benefit function in $B_{10}(\text{rep}_x) = 10 - 1/(x/10)$, which can be obtained as follows: the maximum benefit is obtained among these value for $x = 10$, with a benefit of $10 - 1/1 = 9$. When x is smaller, the proportion $x/10$ is smaller as well, and one loses 'points' in proportion to the inverse of that value. As a result, due to the trade-off between cost and benefit, one might prefer to use just 4-5 iterations to describe a situation with 10 objects.

The upshot is that a simple analysis that takes cost into account can easily capture the Nicaraguan Sign Language production data, by positing a benefit function that grows less quickly than the cost function, as in (87)b. But further interesting cases arise as well, as illustrated in (87)a,c. Of course further research would be needed to establish the shape of the benefit and cost functions on independent grounds: they are stipulated in the present analysis.

□ *Internalizing cost*

The solution we just discussed leaves the pragmatic module of Section 8 unchanged, and adds an external cost module to it. But it would be more in keeping with the spirit of the RSA model to make cost part of the strategic reasoning discussed earlier. We won't attempt a full analysis, but it is worth discussing a highly simplified case, in which the literal iconic meaning of an unpunctuated repetition rep_x is an 'at least x ' meaning. As before, we assume that the speaker seeks to find the best form to express the meaning of a (possibly embedded) constituent. The addressee seeks to recover the meaning (of the relevant constituent) that the speaker had in mind.⁴³

The literal addressee A_0 just associates with rep_x an 'at least x ' meaning. We assume that the strategic speaker S_1 seeks to express an 'at least n ' meaning, but must take cost into account, with the following constraints, which are generalized versions of those illustrated in (86):⁴⁴

⁴² As discussed in the next section, the maximum is at $(2*n)^{1/2}$ for $n = 10$, i.e. $20^{1/2} \approx 4.5$.

⁴³ As noted by B. Spector (p.c.), it would be in keeping with standard RSA models to take rep_x to represent the speaker's epistemic state, which would correspond to the case of unembedded sentences. But since some of our key findings pertain to embedded cases, we continue to take rep_x to express the meaning of a constituent, not speaker's epistemic state.

⁴⁴ One important issue should be mentioned. In the model with external cost (previous section), which just pertained to unembedded sentences, it was natural to posit a huge benefit difference between the case in which one produces too few iterations ($x < n$) and too many ($x > n$) when there were n objects: in the former case, rep_x was true (if under-informative), whereas in the latter case it was false. In the present model, both cases ($x < n$ and $n > x$) fail to adequately represent the target meaning of the constituent, but truth isn't a consideration (since we are considering potentially embedded constituents). One could try to justify the asymmetry by positing that additional processes of pragmatic enrichment can take an 'at least x ' meaning closer to an 'at least n ' meaning when $x < n$ but not when $x > n$ (because pragmatic enrichment can strengthen but not weaken meaning). We leave this as an open issue.

- (88) For $n \geq 2$:
- $U(\text{rep}_x | \geq n) = B_n(\text{rep}_x) - C(\text{rep}_x)$
 - If $x > n$, $B_n(\text{rep}_x) = -\infty$ because the 'at least x ' reading is false.
 - $C(\text{rep}_x) = .5x$

The strategic speaker S_1 seeks to maximize her utility, and if we take the cost as given by (88)c, everything hinges on the benefit function. We assume once again that the benefit increases as the number x of iterations gets closer to n (from below) when one seeks to express an 'at least n ' meaning. If the benefit grows linearly, two cases arise, as in (87)a,b. If the benefit grows more quickly than the cost, the speaker will produce n iterations to expression an 'at least m ' meaning. If the benefit grows less quickly than the cost, the speaker will produce the minimum number of iterations that yield a plural, say 2.

- (89) a. Linear benefit
For $n \geq 2$:
If $2 \leq x \leq n$, $B_n(\text{rep}_x) = a + bx$, with $a \geq 0$, $b > 0$, and thus $U(\text{rep}_x | \geq n) = B_n(\text{rep}_x) - C(\text{rep}_x) = a + (b-.5)x$
If $x > n$, $B_n(\text{rep}_x) = -\infty$ because the 'at least x ' reading is false.
- b. Case 1: $b > .5 \Rightarrow U(\text{rep}_x | \geq n)$ is maximum for $x = n$, i.e. the speaker produces n iterations to represent the meaning 'at least n '
- c. Case 2: $b < .5 \Rightarrow U(\text{rep}_x | \geq n)$ is maximum for $x = 2$, i.e. the speaker produces 2 iterations to represent the meaning 'at least n ', no matter what n is.

But there is another case which is more interesting. In (90), we consider a generalization of the case illustrated in (87)c. The benefit of using x iterations to express an 'at least n ' reading (with $x \leq n$) has the form $10-1/(x/n)$: when $x = n$, the benefit is 9; when $x/n = 1/3$, the benefit is $10-3 = 7$, for instance.


- (90) a. Non-linear benefit: an example
For $n \geq 2$:
If $2 \leq x \leq n$, $B_n(\text{rep}_x) = 10-1/(x/n)$, thus $U(\text{rep}_x | \geq n) = B_n(\text{rep}_x) - C(\text{rep}_x) = 10-1/(x/n) - x/2$
(= $10-n/x-x/2$)
If $x > n$, $B_n(\text{rep}_x) = -\infty$ because the 'at least x ' reading is false.
- b. Maximum of the speaker's utility
 $dU(\text{rep}_x | \geq n)/dx = n/x^2 - 1/2$, hence $dU(\text{rep}_x | \geq n)/dx = 0$ iff $n/x^2 = 1/2$, i.e. $x^2 = 2n$, i.e. $x = (2n)^{1/2}$. It can be checked that this is a maximum (because the second derivative is negative⁴⁵), and always $\leq n$ for $n \geq 2$,⁴⁶ which means that $B_n(\text{rep}_x)$ indeed has the form $10-1/(x/n)$ (rather than $-\infty$, as when $x > n$).

The shape of the utility function $10-n/x-x/2$ as depending on the number x of iterations is illustrated for $n = 3, 5, 10$ in (91). The maxima are given (in view of (90)b) by $(2*3)^{1/2}$, $(2*5)^{1/2}$, $(2*10)^{1/2}$, i.e. 2.4, 3.2 and 4.5. In other words, for larger n , the optimal number of iterations grows, but less quickly than n .

⁴⁵ $d^2U(\text{rep}_x | \geq n)/dx^2 = -2n/x^3$, which is negative for all $n, x > 0$.


⁴⁶ $(2n)^{1/2} \leq n$ iff $2n \leq n^2$ iff $n(n-2) \geq 0$, iff $n \geq 2$ (for $n > 0$).

(91) Value of $10-n/x-x/2$ as a function of x for $n = 3$ (top curve), $n = 5$ (intermediate curve) and $n = 10$



The values of x that maximize the function $U(\text{rep}_x | \geq n)$ can be computed as in (90)b, by computing the derivative $dU(\text{rep}_x | \geq n)/dx$ and finding points at which it is equal to 0. This is the function $f(n) = (2n)^{1/2}$, represented in (92): it can be seen that the optimal number of iterations used to represent an 'at least n ' reading increases far less quickly than n .

(92) Values of x that maximize the function $10-n/x-x/2$ for different values of n , represented by the function $f(n) = (2n)^{1/2}$ (= bottom curve) It can be checked for that $n \geq 2$, $(2n)^{1/2} \leq n$ ($g(n) = n$ is represented by the straight line)



Roughly, then, the strategic speaker S_1 will have the following behavior:
 for target meanings $\geq 2, \geq 3$, S_1 will use rep_2 , i.e. 2 unpunctuated iterations;
 for target meanings $\geq 4, \geq 5, \geq 6$, use rep_3 , i.e. 3 unpunctuated iterations;
 for target meanings $\geq 7, \dots, \geq 10$, use rep_4 , i.e. 4 unpunctuated iterations; etc.

The strategic addressee A_1 who takes into account S_1 's strategy will then associate 2 iterations with an 'at least 2' or 'at least 3' meaning, 3 unpunctuated iterations with 'at least 4', 'at least 5' or 'at least 6' meanings, etc. This will give rise to a pragmatic enrichment whereby the more repetitions are used, the greater the quantity, but with fewer than n iterations for an 'at least n ' meaning. This might prove to be an interesting result if it can be shown that in some contexts 5 unpunctuated repetitions do not have an 'at least 5' meaning but an 'at least n ' meaning for $n > 5$.

We leave it for future research to explore an integration of the cost function to the simplified RSA model we sketched in Section 8. It must be added that we kept things as simple as possible by taking the speaker and addressee to be entirely rational; relaxing this assumption would also be interesting in the future (and in keeping with the RSA model).

Appendix III. Mini-survey on repeated onomatopoeias

We repeat in (93) the target paradigm involving repeated onomatopoeias, from (79).

- (93) *Context*: the addressee is playing a shooter video game.
- In the next 10 seconds, you should [prh - prh - prh - prh]_{slow, pauses}
 - In the next 10 seconds, you should prh-rep4_{fast, no pauses}
 - If in the next 10 seconds you [prh - prh - prh - prh]_{slow, pauses}, I'll give you \$20.
 - If in the next 10 seconds you prh-rep4_{fast, no pauses}, I'll give you \$20

We tested this paradigm by way of audio files and an online survey with 6 consultants (all native speakers of American English, all linguists, including one working on gestures and one working on sign language).⁴⁷ Acceptability and (open or quantitative) inferential questions were modeled after those we used in ASL, with quantitative questions displayed in (94). Average acceptability scores (with 7 = best) and inferential scores (with 7 = strongest) are reported in (95). They should be interpreted with caution because there was a lot of variation among consultants (see the Supplementary Materials).

(94) **Questions**

a. **Acceptability**: How acceptable is the sentence you heard on a 7-point scale (with 1 = worst, 7 = best)?

b. **Quantitative inferential questions for (79)a,b**

How strongly do you draw the following inferences? (with 1 = no inference; 7 = strongest inference)

In the next 10 seconds, the addressee should shoot...

- ...exactly 4 times.
- ...approximately 4 times.
- ...at least 4 times.
- ...at least 5 times.

c. **Quantitative inferential questions for (79)c,d**

Assuming that the shots are spaced appropriately (if relevant), how strongly do you infer that the speaker is committed to giving the addressee \$20 in case the addressee shoots...

- ...exactly 4 times.
- ...approximately 4 times.
- ...at least 4 times.
- ...at least 5 times.

(95) **Inferential results** (averages, mini-survey of 6 linguists, see the Supplementary Materials)

Dominant inferences (i.e. those with the highest endorsement for a given sentence) are boldfaced.

| Sentences | Acceptability | Number of objects that the repetitions stand for | | | |
|--|---------------|--|------------|------------|-----|
| | | = 4 | ±4 | ≥4 | ≥5 |
| a. [prh - prh - prh - prh] _{slow, pauses} | 5.5 | 5.7 | 3.7 | 4.8 | 1 |
| b. prh-rep4 _{fast, no pauses} | 6.5 | 2 | 3.3 | 3.8 | 2 |
| c. [prh - prh - prh - prh] _{slow, pauses} | 5.2 | 5.7 | 2.3 | 4.5 | 3.2 |
| d. prh-rep4 _{fast, no pauses} | 6 | 4.5 | 4.5 | 3.8 | 3.5 |

In unembedded contexts and under *if*, the strongest reading of a punctuated repetition (= (95)a,c) is an 'exactly' reading. Inferences about unpunctuated repetitions are surprisingly weak, possibly because some consultants obtained a reading on which shooting once with a machine gun gave rise to a burst of shots. Here it is particularly informative to consider the 6 responses to 2 open questions, reproduced in (96)-(97). Most answers mention multiple shots and/or a machine

⁴⁷ Before recording the paradigm, we discussed its acceptability with one of the consultants.

gun (unlike in the case of punctuated repetitions; see the Supplementary Materials for full answers), although not necessarily with an 'at least 4' threshold.

- (96) About (93)b: What should the addressee do in the next 10 seconds? (provide a description, not necessarily limited to the number of times the addressee should shoot)
1. Shoot rapidly, at least four times
 2. fire a machine gun
 3. The addressee should fire a burst of shots. I'm imagining an automatic-style gun, where when you pull the trigger (even without holding it for very long), a burst of several shots (around 3-4) get fired very fast.
 4. The addressee should shoot.
 5. fire a short burst of shots
 6. Shoot rapidly as necessary.
- (97) About (93)d: Under what conditions must the speaker give the addressee \$20? (provide a description, not necessarily limited to the number of times the addressee should shoot)6 responses
1. Provided he fire a number of shots
 2. fire two pistols alternatively, one in each hand
 3. Minimally, the addressee has to fire a burst of around 4 shots. Beyond that, it depends on more context.
 4. if the addressee fires their gun / shoots (not necessarily at a target?)
 5. fire a short burst of about four shots
 6. If the player fires at least one burst of shots, or possibly a single shot.

It is worth noting, however, that Consultant 4 apparently obtained an 'at least 1' reading on which the unpunctuated repetition can be satisfied by a single shot. (Consultants also commented on possible effects of intonation, which we did not seek to control for, and differences in realization between one occurrence of the punctuated repetition and the other.)

Supplementary Materials: Raw Data

An Excel files with the raw quantitative scores and averages can be downloaded at the following URL (averages have been underlined when there was more than a 2-point difference among the averaged numbers):⁴⁸

https://drive.google.com/file/d/1hs3sBtrkYglQA_0Zka27anIjnWwAwkfy/view?usp=sharing

Raw ASL data can be downloaded in .docx format at the following URL:

<https://drive.google.com/file/d/1ekfXy98xINkdOuKdfzdpIy4nMIKmxIa/view?usp=sharing>

The mini-survey on onomatopoeias can be found at the following URL (with links to the audio files):

<https://drive.google.com/file/d/1NGRDbyRLYttrlxAX01V55kCYD9sN0EC/view?usp=sharing>

The judgments of 6 linguists (all native speakers of American English, including 1 working on sign language and 1 working on gestures) can be found here:

Anonymized summary: https://drive.google.com/file/d/10WRhYSbISdpbFIQUYMHgR6MIQ8PHxgM_/view?usp=sharing

Anonymized Excel file: <https://drive.google.com/file/d/1F1hvw2Cg5Peeg2vK4vT0xj0SGf9cszY-/view?usp=sharing>

⁴⁸ Many thanks to Lucie Ravaux for putting together this Excel files from the scores that appear in the raw data obtained in elicitation sessions.

References

- Abner, N., Namboodiripad, S., Spaepen, E. & Goldin-Meadow, S. (2015, January). Morphology in Child Homesign: Evidence from Number Marking. (Slides of a talk given at the 2015 Annual Meeting of the Linguistics Society of America, Portland, Oregon)
- Bergen, Leon; Levy, Roger; Goodman, Noah D.: 2016, Pragmatic reasoning through semantic inference. *Semantics & Pragmatics*, 9. <https://semprag.org/article/view/sp.9.20>
- Chierchia, Gennaro; Fox, Danny; Spector, Benjamin: 2012, Scalar implicature as a grammatical phenomenon. In C. Maienborn, P. Portner, and K. von Stechow (Eds.), *Semantics: An International Handbook of Natural Language Meaning*, Volume 3, pp. 2297-2331. de Gruyter.
- Cobrerros, Pablo; Egré, Paul; Ripley, David; van Rooij, Robert: 2012, Tolerant, Classical, Strict. *Journal of Philosophical Logic* 41 (2):347-385.
- Cohn, Neil: 2013, *The Visual Language of Comics: Introduction to the Structure and Cognition of Sequential Images*. London: Bloomsbury.
- Coppola, M., Spaepen, E., & Goldin-Meadow, S. (2013). Communicating about quantity without a language model: Number devices in homesign grammar. *Cognitive Psychology*, 67, 1–25
- Cormier, K. (2002). Grammaticization of indexic signs: How American Sign Language expresses numerosity. Doctoral dissertation, University of Texas, Austin.
- Cremers, Alexandre and Chemla, Emmanuel: 2017, Experiments on the acceptability and possible readings of questions embedded under emotive-factives. *Natural Language Semantics*, 25(3): 223--261.
- Davidson, Kathryn: 2014, Scalar implicatures in a signed language. *Sign Language & Linguistics* 17(1). 1-19
- Davidson, Kathryn: 2015, Quotation, Demonstration, and Iconicity. *Linguistics & Philosophy* 38(6), 477-520.
- Davidson, Kathryn and Mayberry, Rachel I.: 2015, Do adults show an effect of delayed first language acquisition when calculating scalar implicatures? *Language acquisition* 22(4): 329-354.
- Emmorey, Karen & Herzig, Melissa : 2003, Categorical versus gradient properties of classifier constructions in ASL. In K. Emmorey (Ed.), *Perspectives on Classifier Constructions in Signed Languages*, Lawrence Erlbaum Associates, Mahwah NJ, pp. 222–246
- Feldstein, E. (2015). The development of grammatical number and space: Reconsidering evidence from child language and homesign through adult gesture. Short manuscript, Harvard University.
- Fischer, S. (1973). Two Processes of Reduplication in the American Sign Language. *Foundations of Language*, 9: 469-480.
- Greenberg, Gabriel: 2013, Beyond Resemblance. *Philosophical Review*, 122, 2
- Greenberg, Gabriel: 2014, Reference and Predication in Pictorial Representation. Handout of a talk given at the London Aesthetics Forum (February 19, 2014).
- Greenberg, Gabriel: 2019a, Semantics of Pictorial Space. Manuscript, UCLA.
- Greenberg, Gabriel: 2019b, Tagging: Semantics at the Iconic/Symbolic Interface. *Proceedings of the Amsterdam Colloquium 2020*.
- Guerrini, Janek and Migotti, Léo: 2019, Musical gestures in the typology of linguistic inferences. Talk given at the workshop on Linguistic Investigations Beyond Language, Berlin, March 11, 2019.
- Hyde, D.: 1997, From heaps and gaps to heaps of gluts. *Mind*, 106, 641–660. 14.
- Klima, Edward S. and Ursula Bellugi. 1979. *The Signs of Language*, with Robbin Battison, Penny Boyes Braem, Susan Fischer, Nancy Frishberg, H. Lane, Ella M. Lentz, Don Newkirk, E. L. Newport, C. C. Pedersen, and Patricia Siple. Cambridge, MA: Harvard University Press.
- Kuhn, J. (2015). *Cross-categorical singular and plural reference in sign language*. Doctoral dissertation, New York University

- Kuhn, Jeremy and Aristodemo, Valentina: 2017, Pluractionality, iconicity, and scope in French Sign Language. *Semantics & Pragmatics* 10(6).
- Liddell, Scott K.: 2003, *Grammar, gesture and meaning in American Sign Language*. Cambridge: Cambridge University Press.
- Nash, David G.: 1986, *Topics in Warlpiri Grammar*. New York: Garland.
- Newman, Paul: 2012, Pluractional verbs: an overview. In Cabredo Hofherr, P. and Laca, B., editors, *Verbal plurality and distributivity*. De Gruyter, Berlin, Boston.
- Nouwen Rick: 2003, *Plural pronominal anaphora in context*. PhD thesis, Neth. Grad. Sch. Linguist., Univ. Utrecht
- Pfau, R. & Steinbach, M. (2006). Pluralization in sign and in speech: A cross-modal typological study. *Linguistic Typology*, 10, 49–135
- Schlenker, P. (2011). Donkey Anaphora: the View from Sign Language (ASL and LSF). *Linguistics and Philosophy*, 34(4), 341-395
- Schlenker, P. (2014). Iconic Features. *Natural Language Semantics* 22, 4, 299-356
- Schlenker, Philippe: 2019, What is Super Semantics? *Philosophical Perspectives* 32, 1: 365-453 <https://doi.org/10.1111/phpe.12122>
- Schlenker, Philippe: to appear a, Gestural Grammar. *Natural Language & Linguistic Theory*.
- Schlenker, Philippe: to appear b, Iconic Presuppositions. Accepted with minor revisions, *Natural Language & Linguistic Theory*.
- Schlenker, Philippe and Lamberton, Jonathan: 2019, Iconic Plurality. *Linguistics & Philosophy* 42: 45-108 <https://doi.org/10.1007/s10988-018-9236-0>
- Schlenker, P., Lamberton, J. & Santoro, M. (2013). Iconic Variables. *Linguistics & Philosophy* 36(2), 91-149
- Spector, B.: 2013, Bare Numerals and Scalar Implicatures. *Language & Linguistics Compass*, 7(5), 273-294. doi:10.1111/lnc3.12018
- van Fraassen, Bas C.: 1966, Singular Terms, Truth Value Gaps and Free Logic. *Journal of Philosophy*, 63: 481–95
- Varzi, A.: 1995, Supervaluationism and paraconsistency. In D.Batens (ed.), *Frontiers in Paraconsistent Logic*, pp. 279-297. Baldock: Research Study Press, 2003.
- Zucchi, Sandro: 2011. Event Descriptions and Classifier Predicates in Sign Languages. Presentation at FEAST (Formal and Experimental Advances in Sign language Theory) in Venice, June 21, 2011.