



Characterizing the Dynamics of Learning in Repeated Reference Games

Robert D. Hawkins,^a Michael C. Frank,^a Noah D. Goodman^{a,b}

^a*Department of Psychology, Stanford University*

^b*Department of Computer Science, Stanford University*

Received 16 December 2019; received in revised form 4 March 2020; accepted 6 April 2020

Abstract

The language we use over the course of conversation changes as we establish common ground and learn what our partner finds meaningful. Here we draw upon recent advances in natural language processing to provide a finer-grained characterization of the dynamics of this learning process. We release an open corpus (>15,000 utterances) of extended dyadic interactions in a classic repeated reference game task where pairs of participants had to coordinate on how to refer to initially difficult-to-describe tangram stimuli. We find that different pairs discover a wide variety of idiosyncratic but efficient and stable solutions to the problem of reference. Furthermore, these conventions are shaped by the communicative context: words that are more discriminative in the initial context (i.e., that are used for one target more than others) are more likely to persist through the final repetition. Finally, we find systematic structure in how a speaker's referring expressions become more efficient over time: Syntactic units drop out in clusters following positive feedback from the listener, eventually leaving short labels containing open-class parts of speech. These findings provide a higher resolution look at the quantitative dynamics of ad hoc convention formation and support further development of computational models of learning in communication.

Keywords: Language use; Interaction; Conventions; Meaning; Syntax; Natural language processing; Semantic embeddings

1. Introduction

Human language use is remarkably flexible. We are able to coax new meanings out of existing words—or even coin new ones—to handle the diverse challenges encountered in everyday communication (Clark, 1983; Davidson, 1986). This flexibility is partially explained by de novo pragmatic reasoning, which allows listeners to use context to infer

an intended meaning even in cases of ambiguous or non-literal usage (Glucksberg, 2001; Goodman & Frank, 2016; Lascarides & Copestake, 1998). However, a rich theoretical thread has suggested that learning mechanisms may also play an important role, allowing speakers and listeners to dynamically adapt their representations of meaning over the course of an interaction (Brennan & Clark, 1996; Delaney-Busch, Morgan, Lau, & Kuperberg, 2019; Pickering & Garrod, 2004).

Two functional considerations motivate the need for continued learning in communication, even among adults. First, just as there is substantial phonetic variability across speakers with different accents (Kleinschmidt, 2019), words may vary in meaning from speaker to speaker. This variability is clear for cases like slang, technical lingo, nicknames, or colloquialisms (e.g. Clark, 1998), but it may extend even to more ordinary nouns and adjectives. It may be difficult to know at the outset of an interaction exactly which meanings will be shared and which will not, requiring ongoing adaptation. Second, because we live in a changing environment, we often experience novel entities, events, thoughts, and feelings that we want to talk about but do not already share (literal) words to express. Both of these obstacles can be overcome using feedback from one's partner to dynamically re-calibrate expectations about meaning.

The *repeated reference game* task has provided a natural and productive paradigm for eliciting behavior under such conditions. In this task, pairs of participants are presented with arrays of novel images. On each trial, one player (the director) is privately shown a *target* object and must produce a referring expression allowing their partner (the matcher) to correctly select that object from the array. The director is then given feedback at the end of each trial about which object the matcher selected, and the matcher is given feedback about the true target object. Critically, each object appears as the target multiple times in the trial sequence, allowing the experimenter to examine how referring expressions change as the director and matcher accumulate shared experience. To the extent that the director and matcher converge on an accurate system of stable referring expressions, and these referring expressions differ from the ones that were initially produced, it may be claimed that *ad hoc conventions* or *pacts* have formed within the dyad (Hawkins, Goodman, & Goldstone, 2019).

One of the earliest and most intriguing phenomena observed in this task is that descriptions are dramatically shortened across repetitions: An initial description like “the one that looks like an upside-down martini glass in a wire stand” may gradually converge to “martini” by the end (Krauss & Weinheimer, 1964). That is, speakers are able to communicate the same referential content much more efficiently over time. Subsequent work has established a number of signature properties of this process through careful experimental manipulation. First, the extent to which descriptions are shortened is contingent on evidence of understanding from the matcher (Hupet & Chantraine, 1992; Krauss, Garlock, Bricker, & McMahon, 1977; Krauss & Weinheimer, 1966), and it is therefore not easily explained as a mere practice or repetition effect. Second, the resulting labels are *partner-specific* in the sense that they do not transfer if a novel matcher is introduced (Brennan & Hanna, 2009; Metzger & Brennan, 2003; Wilkes-Gibbs & Clark, 1992). Third, they are *sticky* in the sense that they persist through precedent with the same partner even after

the referential context changes (Brennan & Clark, 1996) and are readily extended to similar objects (Markman & Makin, 1998). These qualitative effects provide an empirical backbone for theories of communication to explain. However, as theories are increasingly formalized as computational models making more precise quantitative predictions, setting criteria to distinguish between them will depend critically upon resolving more detailed theoretical questions about the dynamics of adaptation in natural language communication.

In this paper, rather than arguing in favor of a particular theory, we release a new, open corpus of repeated reference games and conduct a variety of analyses to address current gaps in measurement and establish a firmer theoretical foundation facilitating future modeling work. In particular, we address two methodological challenges that have limited the ability of previous studies to provide a sufficiently fine-grained characterization of behavior. First, we need more data. Recent technical developments have allowed interactive multi-player experiments to be run on the web (Hawkins, 2015), boosting sample sizes by an order of magnitude. For comparison, seminal work by Clark and Wilkes-Gibbs (1986) used a sample of eight pairs of participants, while our confirmatory sample alone contains 83 pairs. Second, the computational techniques needed to work with rich natural language data were limited at the time of prior work, but they have become newly tractable, given developments in natural language processing (NLP).

Our analyses roughly divide into two broad categories, corresponding to the dynamics of syntactic *structure* and semantic *content*. Our investigations of syntactic structure in Section 3 focus on the process by which referring expressions are shortened to communicate the same idea more efficiently. One particularly simple model, for example, might predict that shortening is purely driven by a random corruption process: at each repetition, each word from the previous repetition's utterance has some probability of being dropped. Raw word counts alone are not sufficient for disambiguating this simple model from more cognitively complex proposals. To move beyond word counts, we extracted part-of-speech tags and syntax trees from the text to understand which parts of utterances were being dropped, and in which sequence. In contrast to the predictions of the random corruption model, we find that clauses and modifiers tend to be dropped in clusters, preferentially leaving open-class parts of speech (e.g., an adjective and noun) by the final repetition, and that the choice to shorten an utterance or not depends on sources of listener feedback.

In Section 4, we examine the semantic content of utterances over the course of this shortening process. Our analyses revolve around the theoretical constructs of *arbitrariness* and *stability*, which have been central to accounts of convention since Lewis (1969). Arbitrariness refers to the claim that multiple equally successful solutions exist in the space of possible conventions: There is no single optimal solution that all speakers should objectively use. Stability refers to the claim that, once a solution has been found, speakers should not deviate from it. Our contribution is to operationalize these claims in the high-dimensional space of vector embeddings for referring expressions (i.e., GloVe embeddings). By measuring the similarity between referring expressions in this space, we find that signatures of arbitrariness and stability gradually increase over the course of the

interaction. We also clarify the (non-arbitrary) processes shaping which words eventually become conventions. In particular, we test the prediction that pragmatic pressures to be informative in context lead more discriminative words to conventionalize (Gibson et al., 2017; Kirby, Tamariz, Cornish, & Smith, 2015). Taken together, our findings characterize core processes operating within the microcosm of dyadic, natural-language interactions. These processes may ultimately contribute to the adaptive properties of conventions shared across language communities.

2. Methods: Repeated reference experiment

We implemented two variants of the repeated reference task used in classic work by Clark and Wilkes-Gibbs (1986): a relatively unconstrained *free-matching* version that more closely replicates classic in-lab designs, and a more tightly controlled *cued* version. Most importantly, the *cued* version allows us to identify which object each utterance refers to, supporting higher-resolution analyses at the object-by-object level. We considered the free-matching version to be an exploratory pilot sample and subsequently pre-registered planned analyses for the *cued* version at <https://osf.io/vzvmf>. While we are releasing the corpora from both versions, we restrict our analyses to the *cued* version throughout the paper as a cleaner confirmatory sample. Any exploratory analyses and changes to our pre-registered plan are described as relevant in each section below.

2.1. Participants

A total of 480 participants (218 in the *free-matching* version and 262 in the *cued* version) were recruited from Amazon's Mechanical Turk and paired into dyads to play a real-time communication game.

2.2. Stimuli and procedure

On every trial, participants were shown a 6×2 grid containing 12 tangram shapes (see Fig. 1), reproduced from Clark and Wilkes-Gibbs (1986). After passing a short quiz about task instructions, participants were randomly assigned the role of either "director" or "matcher" and automatically paired into virtual rooms containing a chat box and the grid of stimuli. The chat box was a standard modern messaging interface. Participants saw a binary indicator that their partner was currently typing, but only saw the message once their partner hit the *Enter* key or clicked the *Send* button. Use of the chat box was completely unrestricted in both versions of the task: Both participants could freely use the chat box to communicate at any time, and there was no limit on the number or length of messages that could be sent.

In the *free-matching* version, our procedure closely followed Clark and Wilkes-Gibbs (1986). The director and matcher began each trial with scrambled boards. The director's tangrams were fixed in place, but the matcher's could be clicked and dragged into new positions. The players were instructed to communicate through the chat box such that the

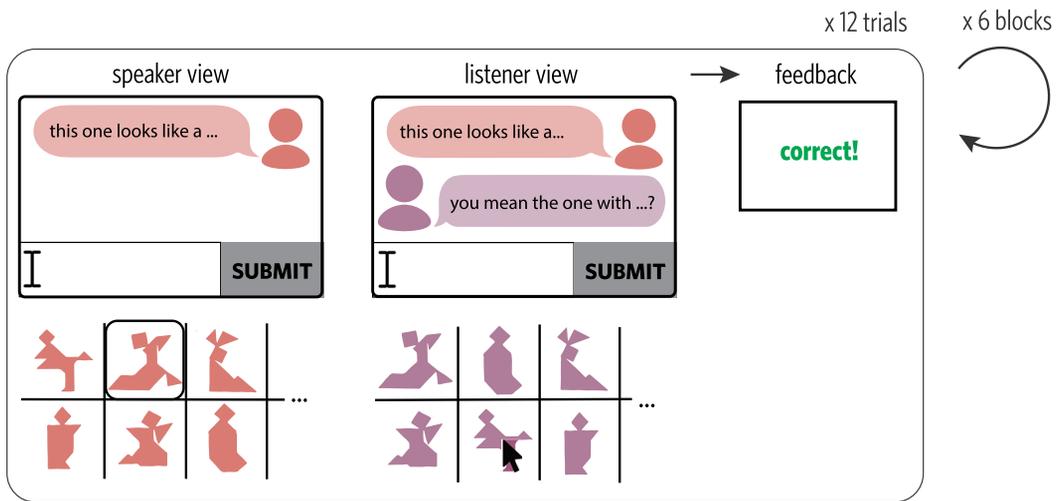


Fig. 1. Display and procedure for the repeated reference game task.

matcher could rearrange their shapes to match the order of the director's board. When the players were satisfied that their boards matched, the matcher clicked a "submit" button that gave players batched feedback on their score (out of 12) and scrambled the tangrams for the next round. After six rounds, players were redirected to a short exit survey. Cells were labeled with fixed numbers from 1 to 12 in order to help participants easily refer to locations in the grid.

While this replicated design allowed highly naturalistic interaction, it posed several problems for text-based analyses. First, utterances contained not only descriptions of the tangrams but also information about the intended location (e.g., "*number 10* is the ..."). Additionally, because there were no constraints on the sequence, participants could revisit tangrams out of order or mention multiple tangrams in a single message, making it difficult to isolate exactly which utterances referred to which tangrams without extensive hand-annotation. Finally, the design of the "submit" button made it easy for players to occasionally advance to the next round without referring to all 12 tangrams.

To address these problems, we designed a more straightforwardly sequential *cued* variant of the task design where directors were privately cued to refer to targets one-by-one and feedback was given on each trial (Fig. 1). This additional structure allowed us to conduct analyses at the object-by-object level. On each trial, one of the 12 tangrams was privately highlighted for the director as the *target*. Instead of clicking and dragging into place, matchers simply clicked the one they believed was the target. They were not allowed to click until after a message was sent by the director, but were not restricted in their use of the chat box. We constructed a sequence of six blocks of 12 trials (for a total of 72 trials), where each tangram appeared once per block. Because targets were cued one at a time, numbers labeling each square in the grid were irrelevant and we removed them. The grid of tangrams was scrambled on every trial, and participants were given

full, immediate feedback: The director saw which tangram their partner clicked, and the matcher saw the intended tangram.

2.3. Exclusion criteria

After applying our pre-registered exclusion criteria for games that terminated before the completion of the experiment due to server error or network disconnection (40 in *free matching* and 33 in *cued*), as well as games where participants reported a native language different from English (2 in *free matching* and 3 in *cued*), we implemented an additional exclusion criterion based on accuracy. We used a 66/66 rule, excluding pairs that got fewer than 66% of trials correct (≤ 8 of 12) on more than 66% of blocks (≥ 4 of 6). While most pairs were near ceiling accuracy by the final repetition, this rule excluded 11 in *free matching* and 8 in *cued* who appeared to be guessing or rushing to completion.

2.4. Data pre-processing

We used a three-step pre-processing pipeline to prepare this corpus for subsequent analyses. Unless otherwise noted, we used the open-source Python package spaCy (version 2.2) to implement all NLP analyses.

1. **Spell-checking and regularization:** We conservatively extracted all tokens that did not exist in the vocabulary of the smallest available ($\sim 50,000$ word) spaCy model and passed them through the SymSpell spell-checker.¹ These suggested corrections were then sequentially presented to the first author and either accepted or overridden at their judgment. This process constructed a spell-correction dictionary containing 677 corrections.
2. **Cleaning unrelated discourse:** Because we allowed our participants to interact in real-time through the chat box, many pairs produced text unrelated to the task of referring to the current target (e.g., greeting one another, asking personal questions, commenting on the length of the task or the results of previous trials). We wanted to ensure that our results were not confounded by patterns in this kind of discourse across the task, and that the semantic content we observe on a particular trial is in fact being used to refer to the current target rather than task-irrelevant topics or, as we found in some cases, referring to other tangrams while debriefing previous errors. We therefore conducted a manual review removing any text not directly referring to the current target. For example, utterances like “the dancing woman” and “this is the one we got wrong last time” were kept in because they were referring to properties of the current tangram, but words like “yeah” or “ok” and messages like “good job” and “they’ll go quicker if you remember what I say!” were not. This review affected 1,448 messages, and we also saved these corrections in a dictionary.
3. **Collapsing multiple messages within a trial:** Finally, some directors used our chat box like a texting interface, hitting the enter key between every micro-phrase of text. Because many of our analysis techniques depend upon having a single referring

expression for each trial, we collapsed all messages sent by each participant within a trial into a single message by inserting commas between messages. We chose to use commas because it tends to maintain grammaticality and does not inflate word counts. Importantly, we used this strategy even on trials where there were multiple rounds of interaction between the director and matcher, joining together a director's original message with their follow-up to the matcher's response. This strategy had the advantage of preserving all of the referential content produced by the director on a trial, compared to a "purer" strategy of removing messages sent in response to the matcher, but the disadvantage of losing the precise interactional structure of the dialogue.

After implementing exclusions and cleaning, we were left with a *free matching* corpus containing a total of 8,639 ($\approx 50,000$ words) messages over 56 complete games and a *cued* corpus containing 7,867 messages ($\approx 46,000$ words) over 83 games.

3. Results: Characterizing the dynamics of structure

Our first set of analyses examines how the structure of participants' utterances changes over the course of our experiment. We begin with the observation that the mean number of words used by directors for each tangram decreases strongly over time (see Fig. 2A).² This result holds whether we count the total number of words produced over the whole trial (i.e., including words produced in follow-up responses to matchers) or only the words in the initial description, prior to any feedback from the matcher, so it is not purely an effect of dialogue. This reduction replicates a highly reliable effect found throughout the literature on repeated reference games (e.g., Brennan & Clark, 1996; Krauss & Weinheimer, 1964), though participants in our task used fewer words overall than reported by Clark and Wilkes-Gibbs (1986). This difference is likely due to the text-

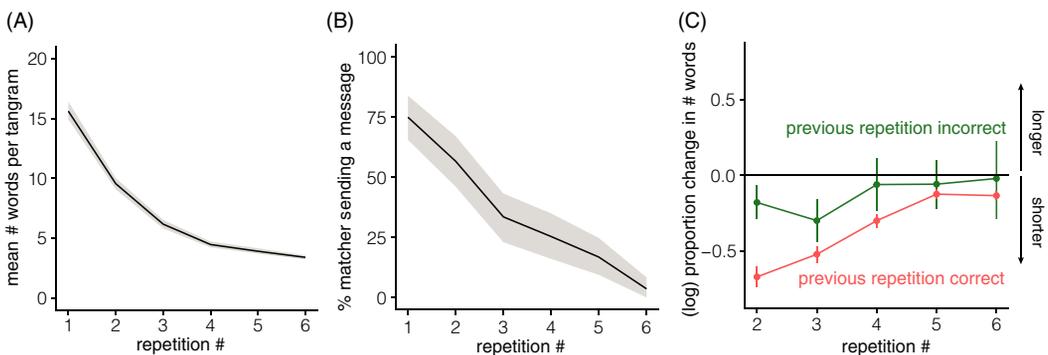


Fig. 2. (A) Directors use fewer words per tangram over time, (B) matchers are less likely to send messages over time, and (C) directors are sensitive to feedback from the matcher's selection, modulating the reduction in message length on the subsequent repetition of a tangram after an error is made.

based (vs. spoken) interface. The following analyses break down this general gain in efficiency into a finer-grained set of phenomena concerning the *structure* of referring expressions over time. What sequence of transformations do descriptions undergo over the course of repeated reference?

3.1. *The effect of listener feedback on reduction*

Conventions are formed *collaboratively*, not in isolation (Clark & Wilkes-Gibbs, 1986; Schober & Clark, 1989), and thus depend on some form of social feedback. We consider two channels of feedback that were available to speakers. First, matchers could voluntarily initiate a bi-directional feedback process at any point within a trial by asking follow-up questions, suggesting corrections, and acknowledging or verbally confirming their own understanding through a backchannel. Second, we automatically supplied ground-truth feedback about the matcher's selection and the true target at the end of each trial. We expect that speakers integrate both sources of evidence about the listener's understanding to determine how to shorten their utterances. If either feedback channel is restricted, or if they provide conflicting evidence, descriptions may not reduce as strongly (Garrod, Fay, Lee, Oberlander, & MacLeod, 2007; Hupet & Chantraine, 1992; Krauss & Weinheimer, 1966).

We predicted that the matcher's use of backchannel feedback should be highest on the first repetition and drop off once meanings are agreed upon, consistent with the patterns observed by Clark and Wilkes-Gibbs (1986). To test this prediction, we coded whether the matcher sent a message or not on each trial and fit a mixed-effects logistic regression model with a fixed effect of repetition, random intercepts, and slopes for each pair of participants, and a random intercept for each target. We found that the probability of the matcher sending a message decreased significantly over the game ($b = -0.84$, $t = -9.1$, $p < .001$). While usage of the backchannel in our online text-based task was less frequent overall than reported in previous verbal lab experiments, we nonetheless strongly replicated the overall trend. In aggregate, 75% of matchers responded with at least one message over the 12 trials in the first repetition block, but only 4% sent any message at all over the course of the last block (see Fig. 2B). These messages were frequently questions: as a lower bound, we observed that 49% of matcher messages explicitly contained question marks (e.g., "is it standing?"). Other messages simply echoed the director's label or suggested alternative labels.³

Next, as an exploratory analysis, we examined the extent to which directors were sensitive to the ground-truth feedback that was provided at the end of each trial about which tangram the matcher actually *selected*. If the matcher failed to select the correct target, the director may take this as evidence that their description was insufficient and attempt to provide more detail the next time they must refer to the same tangram. If the matcher is correct, on the other hand, the director may take this as evidence of understanding and reduce their level of detail when the tangram next appears. Note that ground-truth feedback provided the speaker distinct information from backchannel feedback within the trial: Backchannel feedback did not guarantee a correct response, and matchers often

made the correct response without sending any messages in response. For example, errors on the first repetition block were only slightly less likely when matchers engaged in dialogue through the chatbox (20%) than when they stayed silent (23%; difference not significant, $\chi^2(1) = 0.56$, $p = .45$), although matchers may also have been more likely to initiate backchannel responses on more difficult trials.

We tested the speaker's sensitivity to ground-truth feedback by comparing the proportional change in utterance length (i.e., $\log(n_t/n_{t-1})$) on the block after an error against the change after a correct response. This measure could be positive, indicating a net increase in utterance length, or negative, indicating a reduction. We fit a mixed-effects regression model predicting this measure with a categorical fixed effect of the matcher's response for the same at the previous repetition block (correct vs. incorrect) and a (centered) continuous effect of repetition block number, including maximum random effects at the speaker level. We found a significant main effect of feedback, controlling for block number: Utterance length decreased more after correct responses than after incorrect responses, $b = -0.28$, $t = -7.2$, $p < .001$ (see Fig. 2C).

Although appearances of the same tangram were spaced out by block, it is still possible that this effect is not item-specific but the result of lower level attentional or affective mechanisms triggered in the aftermath of an error signal. To evaluate this possibility, we also measured the proportional change in utterance length on the following *trial*, when feedback about the matcher's response would be freshest but the target tangram would be different. We then constructed a second regression model including categorical fixed effects of matcher response (correct vs. incorrect) and item-specificity (change measured relative to previous trial vs. previous repetition block), as well as their interaction, with no random effects. We found a significant cross-over interaction, $b = -0.32$, $t = -6.2$, $p < .001$.⁴ The sensitivity to feedback we observed on the subsequent repetition block is not present on the subsequent trial: Speakers are equally likely to use more or fewer words immediately after a correct response, and actually use slightly fewer words on the trial immediately after an *incorrect* response due to a regression to the mean: Statistically, more words than average are used for harder tangrams. This pattern of results is consistent with sensitivity to *tangram-specific* evidence of the matcher's understanding when deciding to modify referring expressions.

3.2. Breaking down the structure of reduction

3.2.1. Reduction in parts of speech

Having established the matcher-dependent conditions under which directors are willing to shorten their utterances, we now examine the way they are shortened in more detail. First, we explore which *kinds* of words are most likely to be dropped. We used the SpaCy part-of-speech tagger to count the number of words belonging to different parts of speech in each message sent by the director.⁵ In Fig. 3A, we show the shifting proportions of different parts of speech at each repetition. We find that nouns account for

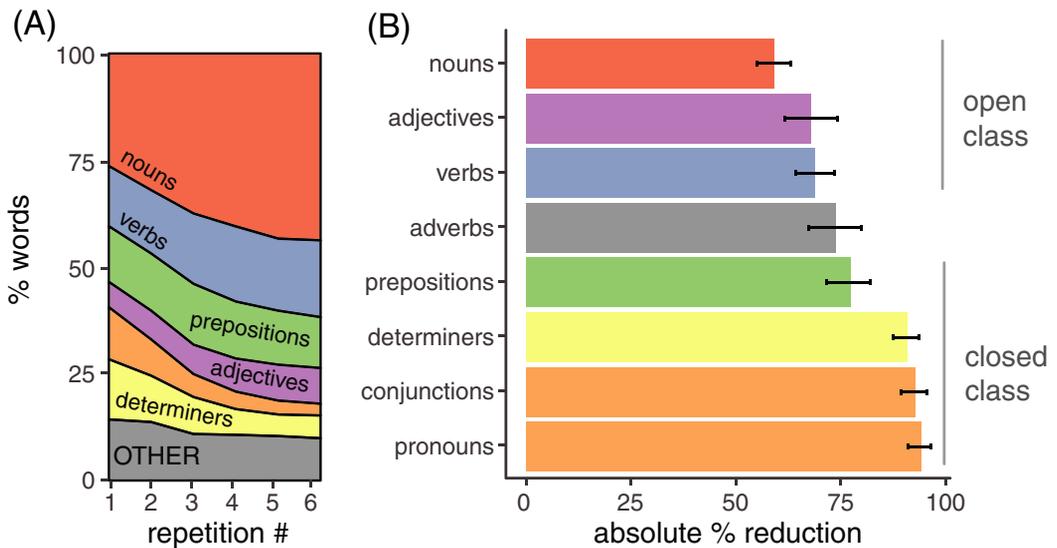


Fig. 3. (A) Proportion of words from different part of speech at each repetition block. For legibility, pronouns and conjunctions are combined in the orange strip while adverbs were grouped into “OTHER”. (B) Closed-class parts of speech are more likely to be dropped than open-class parts of speech. Note that the classification of adverbs is controversial, as many common adverbs are considered closed-class items (e.g., “only,” “now,” “there”) while others are open. Error bars are bootstrapped 95% confidence intervals.

proportionally more of the words being used over time, while determiners and prepositions account for fewer.

To test which kinds of words are more likely to be dropped, we measured the percent reduction in the number of words in each part of speech from the first repetition to the sixth repetition (i.e. $(n_1 - n_6)/n_1$). We find that pronouns (“it”, “he”), conjunctions (“and”, “that”), and determiners (“the”, “a”, “an”) are the most likely classes of words to be dropped (94%, 93% and 91%, respectively) and nouns (“dancer”, “rabbit”) are the least likely to be dropped (59%). More generally, closed-class parts of speech, including function words, appear strictly more likely to be dropped than open-class parts of speech (Fig. 3B). To statistically test this claim, we constructed a mixed-effects model predicting the reduction of each part of speech from the first round to the final round. We included the binary fixed effect of open class versus closed class (excluding adverbs), as well as random intercepts for each pair of participants, and random intercepts for each individual part of speech within each class. We found a significant effect of class, $b = 0.24$, $t = 4.5$, $p = .006$, supporting the observation that closed class parts of speech are more likely to be dropped.

One possible interpretation of these findings is that reduction may be driven mostly by the loss of function words as directors shift to a less-grammatical shorthand over the course of the task. However, when examining the n-grams most likely to be dropped (see Table 1), we noticed that many of the most dropped closed-class words are used to form

prepositional phases (“of”, “with”) or combine different clauses (“and”). Others are modifiers (“the right ...”). These examples suggest an alternative explanation: The higher reduction of closed-class function words may be a consequence of entire meaningful grammatical units (e.g. clauses, prepositional phrases) being dropped at once.

3.2.2. Reduction in syntactic units

If initial descriptions tend to be syntactically complex because they combine multiple pieces of identifying information about the target, then the director may gradually omit meaningful “chunks” as they become informationally redundant. We explicitly tested this hypothesis by examining whether pairs of words dropped from one reference to the next tend to come from the same syntactic units, relative to what would be expected by random deletions. We quantified the extent to which dropped words “cluster” by examining closeness between the dropped words in a dependency parse tree (see Fig. 4). Dependency grammars represent syntax in terms of directed dependency relations between pairs of words (Corbett, Fraser, & McGlashan, 1993; Hudson, 1984), and dependency parsing, the task of extracting a tree of dependency relations from raw input text, has become a canonical task in computational linguistics (Jurafsky & Martin, 2014; Kübler, McDonald, & Nivre, 2009).⁶

We used the state-of-the-art statistical dependency parser from SpaCy, using the ClearNLP dependency labels (Choi & Palmer, 2012), to extract a parse tree for every utterance in our corpus (see Honnibal & Johnson, 2015; Kiperwasser & Goldberg, 2016, for details on neural network architectures for parsing). We then compared each referring expression to the one produced on the subsequent repetition block to determine which words were dropped and which reappeared. For each pair of words that were dropped, we used a

Table 1

Top 15 unigrams, bigrams, and trigrams with the highest numeric reduction from first repetition to last repetition. Text was lemmatized before computing n-grams, which also mapped all pronouns to the “-PRON-” token

	Unigrams	Bigrams	Trigrams
#1	a	look like	look like a
#2	the	like a	look like -PRON-
#3	-PRON-	to the	to the right
#4	like	this one	like a person
#5	look	the right	to the left
#6	be	the left	one look like
#7	on	like -PRON-	this one look
#8	one	on the	like -PRON- be
#9	with	with a	this one be
#10	to	a person	-PRON- look like
#11	and	-PRON- be	look like someone
#12	right	on top	diamond on top
#13	this	a diamond	in the air
#14	of	in the	on top of
#15	head	one look	a diamond on

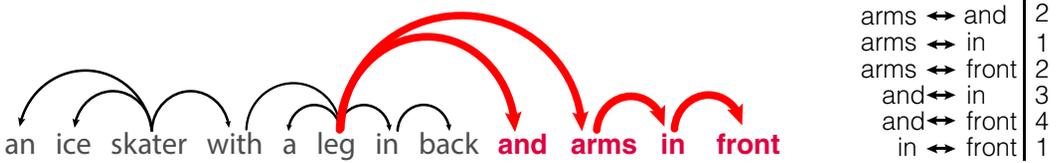


Fig. 4. Example dependency parse for a referring expression in our task. If the words “and arms in front” were dropped, we would find a mean path length of 2.17 among the dropped words.

standard shortest-path algorithm to find the minimal distance between them in the dependency parse tree we extracted (see Fig. 4). Finally, we computed the mean path lengths between all such pairs of dropped words on each given trial, and then took the mean across all trials (excluding blocks where no words were dropped). This method weights each utterance evenly, preventing trials with more words from dominating the global average.

We compared this empirical “syntactic clustering” statistic to two baselines. For the *random* baseline, instead of examining dependency lengths between the words that were actually dropped, we randomly sampled the same number of words from the referring expression and computed the dependency length between them. We repeated this procedure 100 times to obtain a null distribution of the mean dependency length that would be expected if words were being dropped *randomly* from anywhere in the message. For the *function words* baseline, we were specifically interested in the null distribution that should be expected if function words were preferentially dropped independent of the syntactic subunits they belong to. We first sampled from the set of function words in the utterance, and if this set was smaller than the total number of words dropped, we filled the remainder with random non-function words.

We found a mean empirical dependency length of 2.77, which lay outside both the random null distribution (range: [2.90, 2.99]) and the function word null distribution (range: [3.03, 3.08]), indicating a reliable effect of syntactic clustering among the words that were dropped on each round. That is, these words tended to be closer to one another in the dependency parse than expected by total chance or by preferentially dropping function words independently of their corresponding syntactic units. Furthermore, while overall dependency lengths get smaller as utterances become shorter, this result holds within every repetition block (see Fig. S2), and other statistics gave similar results, including the minimum dependency length and the raw distance in the sequence of words. This result accords with earlier observations by Carroll (1980), who reanalyzed transcripts from Krauss and Weinheimer (1964). In those data, the short names that participants converged upon were prominent in some syntactic construction at the beginning of the session, often as a head noun that was initially modified or qualified by other information.

4. Results: Characterizing the dynamics of semantic content

So far we have examined the increasing efficiency of referring expressions in terms of their (syntactic) structure. We next explore how the *semantic content* of referring

expressions changes over repeated reference. Which words from a speaker's initial description are most likely to become conventionalized in their final labels? Why do all dyads not end up with the same conventions? And, once efficient conventions are formed, are they stable? In exploring these questions, we find support for a view of adaptation as a path-dependent process of gradually paring down redundant information and coalescing around the most diagnostic features for the given context.

4.1. *Initially distinctive words are more likely to conventionalize*

Two general computational principles guide our exploration of which content is dropped and which is preserved. First, Gricean principles suggest that a good referring expression is one that applies more strongly to the target than to the distractors; in contrast, those expressions that apply to multiple objects will be less informative. Second, principles of cross-situational learning suggest that these informativity considerations will be strengthened over time. The exclusive usage of a word with one tangram and no others should reinforce the specificity of that meaning in the local discourse context, even if the matcher may be *a priori* willing to extend it to other targets. Conversely, if a particular word has been successfully used with several different referents, its specificity may be weakened in the local context. Putting these principles together, we hypothesized that the labels that conventionalize should not be a random draw from the initial description. Instead, more *initially distinctive* words should be more likely to conventionalize.

For each pair of participants, we quantified the distinctiveness of a word w as n_w : the number of tangrams that it was used to describe on the first repetition. A word that is only used in the description of a single tangram (e.g., a descriptive noun like "rabbit") would be very distinctive, while a word used with all 12 tangrams (e.g., an article like "the") would be not distinctive at all. While this formulation is easy to state in words, it is closely related (up to a simple deterministic transformation) to two popular and theoretically motivated measures of distinctiveness used in natural language processing (Salton & Buckley, 1988): *tf-idf* and *PPMI*.⁷ Given this simple but principled measure of word distinctiveness at the speaker-by-speaker level, we were interested in the extent to which it accounts for conventionalization: the probability that a word in the director's initial description is preserved until the end of the game. More than half of the words used to refer to a tangram on the final repetition (57%) appeared in the initial utterance.⁸ We thus restricted our attention to this subset of words, coding them with a 1 if they later appeared at the final repetition and 0 if they did not. We then ran a mixed-effects logistic regression including a fixed effect of initial distinctiveness and maximal random effect structure with intercepts and slopes for each tangram and pair of participants. We found a significant positive effect of distinctiveness: Words that were used with a larger number of tangrams on the first repetition were less likely to conventionalize, $b = -0.23$, $z = -6.1$ (see Fig. 5). Similar results are found explicitly using the *tf-idf* measure.

To further evaluate how distinctiveness is related to eventual conventionalization, we conducted a non-parametric permutation test. For each speaker and tangram, we extracted the word with the *highest distinctiveness value* and computed the mean probability of this

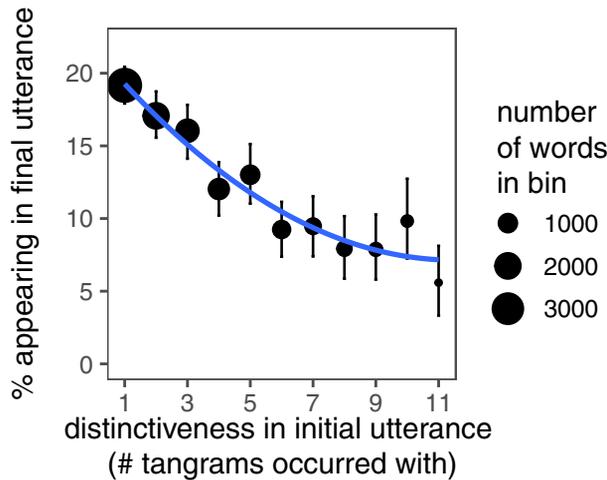


Fig. 5. More distinctive words are more likely to conventionalize. Points represent estimates of the mean probability of conventionalizing across all words with a given distinctiveness value. Size of points represents the number of words at that value. Curve shows regression fit; error bars are bootstrapped 95% CIs.

word also being used on the final repetition. In the case of a tie, we sampled from the set of words sharing the highest distinctiveness value. After repeating this procedure 1,000 times, we found a distribution ranging from 24% to 31%. As a baseline null model, we randomly sampled from the list of all words contained in the initial utterance instead of the most distinctive one. Repeating this procedure yielded a null distribution ranging from 2.5% to 6.6%, which was significantly lower than the one derived from the most distinctive words. These results are also consistent with our earlier finding that open-class parts of speech are more likely to be preserved on the final repetition than closed-class parts of speech (see Section 3.2), since open-class parts of speech are statistically more likely to supply distinctive words.

4.2. Semantic meaning diverges across pairs and stabilizes within pairs

Conventions are characterized by their arbitrariness and stability (Lewis, 1969). Our remaining predictions concern the dynamics of these properties. First, due to sources of variability in the population of speakers, we predict that the referring expressions used by different pairs will increasingly diverge to different, idiosyncratic labels. In other words, different pairs will find different but equally successful equilibria in the space of possible linguistic conventions. Second, as directors learn and gradually strengthen their expectations about how their partner will interpret their referring expressions, the labels used within each pair for each tangram will stabilize. In other words, once there is evidence that a particular label is successfully understood, there should be little reason to deviate from it.

To operationalize these constructs, we used a measure of similarity based on distances computed between continuous vector space embeddings of referring expressions.⁹

Although the idea of using such representations of words to measure similarity is an old one (Bengio, Ducharme, Vincent, & Jauvin, 2003; Landauer & Dumais, 1997; Osgood, 1952), recent progress in machine learning has yielded substantial improvements in the quality of these representations. To quantify the dynamics of semantic content in referring expressions across and within games, we therefore first extracted the 300-dimensional GloVe vector (Pennington, Socher, & Manning, 2014) for each word in the messages produced by the director on each trial. We limited our analysis to messages produced by the director, rather than collapsing together director and matcher messages, in order to preserve consistency in the source of semantic content. We then averaged these word vectors to obtain a single sentence vector for each trial. Variations on such simple averaging methods are surprisingly strong baselines for sentence representations (Arora, Liang, & Ma, 2017), providing better downstream task performance than whole-sentence encoders (Kiros et al., 2015).

To avoid artifacts from function words, we only included open-class content words (nouns, adjectives, verbs, and adverbs) in this average. Because different forms of a word may have slightly different representation, we also applied a lemmatizer to further standardize the input. Lemmatization maps multiple morphological variants (e.g., “played,” “playing,” “plays”) to the same stem (“play”). We did not want an observed difference between two pairs to be driven simply by different forms of the same word. We then defined a similarity metric between any pair of sentence vectors $\langle u_i, u_j \rangle$. Our results are robust to several choices of metric, but for simplicity we will use cosine similarity throughout the presentation below:

$$\langle u_i, u_j \rangle = \cos \theta_{ij} = \frac{u_i \cdot u_j}{\|u_i\| \|u_j\|}$$

We begin by visualizing the trajectories taken by each pair of participants when referring to a particular example tangram (see Fig. S3 for similar plots for the other items). To create this visualization, we took the first 50 components recovered by running principal components analysis (PCA) on the 300-dimensional embeddings for all utterances used to refer to this tangram, including all speakers and all repetition blocks. We then used t-SNE (Maaten & Hinton, 2008) to stochastically embed the lower-dimensional PCA representation of these utterances in a common 2D vector space.¹⁰ Finally, we connected the first and last utterance a particular pair used to refer to this tangram with an arrow (Fig. 6), and annotated utterances in several regions of the space.

Most strikingly, we observed that the initial utterances of each game tend to cluster tightly near the center of the space and the final utterances are *dispersed* more widely around the edges. This pattern is consistent with the hypothesis that early descriptions may overlap before each speaker hones in on more distinctive different equilibria later in the game. Indeed, pairs often initially mentioned multiple properties (e.g., “person raising their arms up like a choir singer”) before breaking the symmetry and collapsing to one of these properties (“choir singer”). Our example also shows the variety of different solutions discovered by different speakers. A handful of semantically distinct labels served as

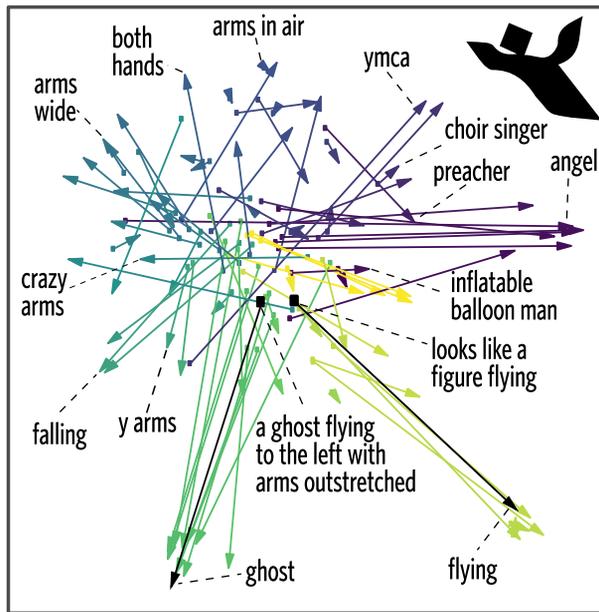


Fig. 6. 2D projection of semantic embeddings for example tangram using t-SNE. Each arrow represents the trajectory between the first repetition to last repetition for a distinct pair of participants. Color represents the rotational angle of the final location to more easily see where each pair began. Annotations are provided for select utterances, representing different equilibria found by different participants. Arrows in black highlight a pair of trajectories where the initial utterances were similar (e.g., both containing "flying"), but the final equilibria were differentiated. Because t-SNE is a stochastic algorithm, even identical words (e.g., the many instances of "ghost") will map to slightly different locations.

equilibria for a number of pairs ("ghost," "flying," "angel") while many more idiosyncratic labels spread out more widely in space. In the remainder of this section, we quantitatively substantiate these observations.

4.2.1. Utterances are more similar overall within interactions than across interactions

Before examining the dynamics of how these vectors change over time, we test the basic prediction that referring expressions used by a single speaker *within* a game are more similar overall than those used by different speakers *across* games. For each tangram, we computed the pairwise similarities between all utterances used by a speaker to refer to that tangram at different times *within* a game and also between all utterances used by different speakers *across* games. The distributions of these values are shown in Fig. 7. We estimated the distance between these distributions using the standard normalized sensitivity $d' = \frac{\mu_A - \mu_W}{\sqrt{1/2(\sigma_A^2 + \sigma_W^2)}} = 2.65$. To compare this estimated difference against

the null hypotheses that within- and across-game similarities are drawn from the same distribution, we conducted a permutation test by scrambling "within" and "across" labels for each similarity value and re-computing d' 1000 times. We found that our observed statistic was extremely unlikely under this null distribution, 95% CI: [-0.09, 0.09],

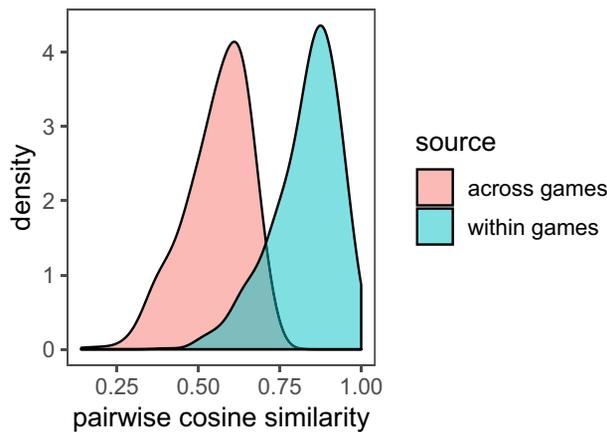


Fig. 7. Distribution of similarities between different utterances within and across different games.

$p < .001$. In other words, the referring expressions a single pair uses for a particular tangram tend to cluster together in semantic space while those used by different pairs are spread out in different parts of the space. This observation leaves open the question of whether pairs start out semantically similar and become different through the conventionalization process (as predicted by the theory of conventions), or simply come into the experiment with idiosyncratic differences. To explore this question, we conducted analyses on how the semantic vectors changed over time.

4.2.2. Utterances become increasingly consistent within interactions

As directors modified their utterances across successive repetitions, we hypothesized that they would converge on increasingly consistent, stable ways of referring to each tangram. To test this prediction, we computed the cosine similarity between successive utterances produced by each speaker (see Fig. 8A). A mixed-effects model with (orthogonalized) linear and quadratic fixed effects of repetition number and maximal random effects for both tangram and pair of participants showed that similarity between successive utterances increased substantially throughout an interaction ($b = 2.8$, $t = 11.0$, $p < .001$). The quadratic term was also weakly significant ($b = -0.44$, $t = -2.2$, $p < 0.035$).

4.2.3. Utterances become increasingly different across different interactions

Finally, we predicted that although the referring expressions used by different pairs may begin with substantial overlap, they would become increasingly dissimilar from each other across time, gradually diverging into different equilibria. We tested this prediction by computing the mean similarity between referring expressions used by different speakers. The large sample of similarities ($N = 245,016 = 12 \text{ tangrams} \times 6 \text{ repetitions} \times \frac{83-82}{2} \text{ distinct pairs}$) presented both advantages and disadvantages for this analysis. On one hand, we could obtain highly reliable estimates of mean similarity. On the other hand,

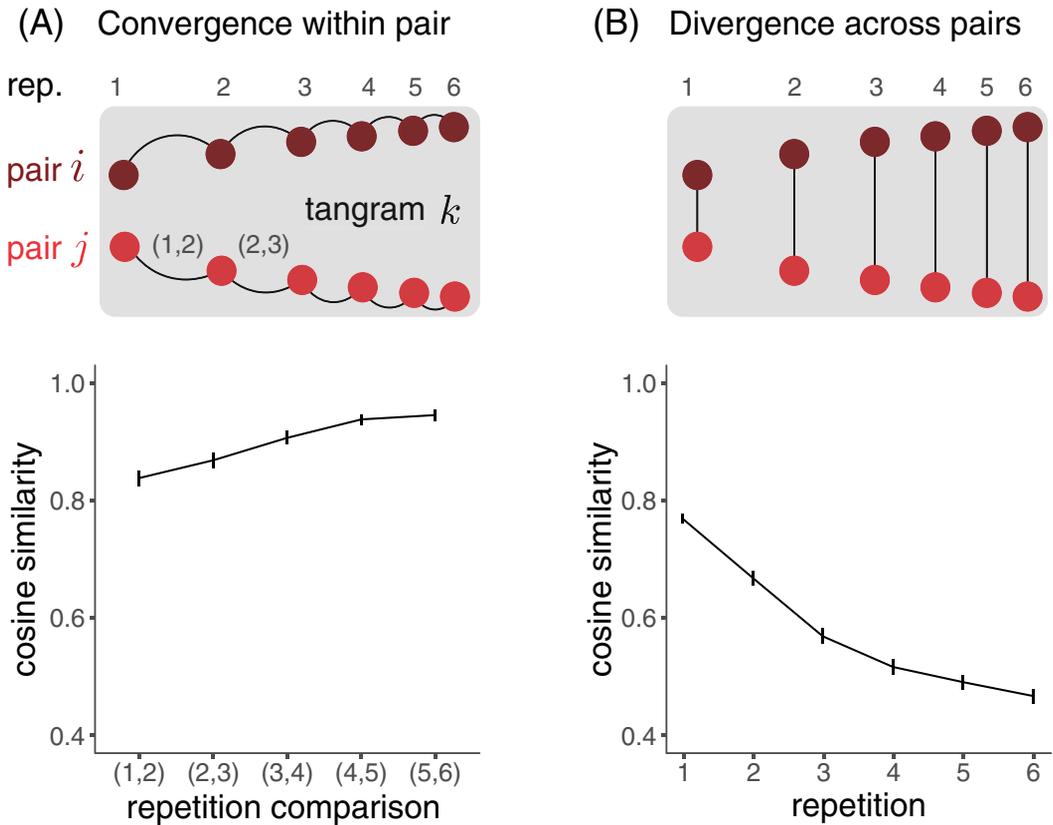


Fig. 8. (A) Utterances within a pair become more similar to successive utterances on later repetitions, converging on a stable convention, but (B) utterances across pairs become steadily more dissimilar, diverging to different solutions. These patterns are depicted schematically by dots within a pair changing less over time while dots in different pairs move further apart. Error bars are bootstrapped 95% CIs over participant-level means.

larger random-effects structures led to convergence problems. We therefore ran a mixed-effects regression model including linear and quadratic fixed effects of repetition number and including random effects only at the tangram-level. We found a strong negative linear fixed effect of repetition on between-game semantic similarity ($b = -48.6$, $t = -16.7$, $p < .001$) as well as a significant quadratic effect ($b = 15.2$, $t = 10.9$), indicating that divergence slows over time (see Fig. 8B).

5. General discussion

Our language changes as we get to know a social partner through repeated interactions. We gradually learn what is meaningful to them and establish common ground. In this paper, we characterized the quantitative dynamics of this process by examining behavior

in a new corpus of natural-language repeated reference games. This corpus is sufficiently large to provide new traction toward resolving theoretical questions about the nature of adaptation in communication. Our study illustrates the general point that larger datasets enable more precise measurements, which in turn drive theory development (Frank, 2018).

In our corpus, we replicate the classic finding that directors reduce the length of their descriptions over the course of the task. But we also show that they do so in a way that is sensitive to evidence of matcher understanding and structured to omit redundant syntactic chunks of information, leaving eventually only the most distinguishing words. The resulting labels display quantitative signatures of increasing *arbitrariness* in the sense that different pairs increasingly diverge to distinct solutions, and *stability* in the sense that speakers do not deviate from a solution once it is discovered. Taken together, these findings clarify the desiderata for theories of *ad hoc convention formation*. For a model of communication to explain how general-purpose meanings are systematically tailored to the needs of the current interaction in the way we observed, it must provide a mechanism to select and combine syntactic phrases that are initially distinctive and to prune them over time, modifying them if they are unsuccessful.

Our findings also raise new and subtle questions about the cognitive mechanisms giving rise to these properties. One such question concerns the mechanisms supporting arbitrariness over such short interactions: What breaks the “symmetry” among different possible descriptions and leads different pairs to diverge from one another? One possibility is that each individual speaker may initially have strong but idiosyncratic initial preferences for short labels, and arbitrariness emerges from variability in these preferences throughout the population. Under this possibility, speakers begin with long, elaborated descriptions due to uncertainty about whether their preferred label will be understood, but in the absence of misunderstandings will proceed with their pre-meditated label. A second possibility is that speakers *themselves* may be unclear about a mutually understandable way to refer to these unfamiliar objects. If uncertain speakers initially sample an utterance from a broad distribution of acceptable labels, and update their distribution on subsequent repetitions conditioned on feedback, different pairs may end up in different equilibria due to randomness in sampling from a more or less shared initial distribution. This latter mechanism has been proposed in recent probabilistic models of convention formation (Brochhagen, 2017; Hawkins, Frank, & Goodman, 2017; Smith, Goodman, & Frank, 2013), which present simulations reproducing several of the properties we observed in our data. These two possibilities—strong but idiosyncratic initial preferences or initial shared uncertainty—are not mutually exclusive. Our results rule out the possibility of universally shared strong preferences, but it is possible that some speakers have different strong preferences about labels while others are initially more uncertain. One way for future work to disentangle these possibilities is to elicit better measurements of speakers’ initial beliefs about appropriate labels. For instance, an approach proposed by Fussell and Krauss (1989) asked directors to either produce descriptions for others or for themselves in the future, and Bayesian truth serum approaches (Prelec, 2004) compare an

individual's own subjective preferences with their expectations about whether these would be shared by others.

Although our analyses go beyond previous work by using recent vector-space semantic models, they still face several limitations. We address two such limitations with supplemental analyses included in the Appendix. First, measures of similarity relying on vector space representations like GloVe are fundamentally limited by the quality of the semantic space that has been learned. To address this concern, we provide converging evidence for the properties of arbitrariness and stability using the *discrete distributions* of word tokens appearing in each utterance instead of continuous utterance embeddings (see Appendix A). Second, a related concern is that the gradual divergence we observed between different interactions could be an artifact of the way we constructed utterance embeddings by averaging word embeddings. If averaging together more words creates a distinctive type of “washed-out” utterance embedding, and early descriptions contain more words, then high initial similarity across interactions may reflect utterance length rather than actual semantic overlap. We address this concern by providing an additional permutation test baseline that scrambles words across different tangrams prior to averaging word embeddings (see Appendix B). This baseline also presents an opportunity to compare the divergence we observed across different interactions against the divergence within a *single speaker's* descriptions of their 12 different tangrams. Just as different speakers initially include many of the same attributes in their descriptions for a tangram but eventually (unknowingly) diverge to distinct labels, a single speaker also begins by re-using certain attributes for several tangrams but (knowingly) prunes them down to be as distinctive as possible due to informativity pressures, consistent with our findings in Section 4.1.

Our use of classic tangram stimuli also raises an important question about how our findings would be expected to apply to other spaces of novel objects. In particular, it is likely that participants converge to distinctive “names” because the target of reference were distinctive objects. If the targets of communication instead varied along clear latent dimensions (e.g., Nölle, Staib, Fusaroli, & Tylén, 2018), contained multiple objects in relation to one another, or depicted events or activities unfolding over time, participants might instead have converged on more compositional systems making use of adjectives, verbs, and prepositions. Similarly, the generality of our results is limited by the population we sampled. Our use of online data collection allowed us to create a relatively large number of arbitrary dyads within a convenience population, but also limits opportunities for studying these dyads over longer time periods. It will be important to determine how the *ad hoc* meanings formed in one novel context generalize to other contexts with the same partner. Furthermore, though our dyads are likely diverse in many ways relative to the US national population (Levay, Freese, & Druckman, 2016), they are not representative of either the US population or any broader population. Thus, further cross-cultural work examining the validity of our conclusions across populations, and in different languages, would be a valuable contribution for future work.

Finally, there are important differences between face-to-face verbal communication and online text-mediated communication. Text message interfaces do not transmit the

message character by character as it is being written; messages are seen in their entirety upon being sent. Meanwhile, speech is processed incrementally, allowing for a real-time “backchannel” where the director may be interrupted as soon as the matcher is confident they understand. The timing of verbal interruptions may provide stronger evidence of understanding than a reply to a fully composed message. Still, we believe that extending classic in-person work to a text-based chat modality is a strength of our study, as instant text messaging is an increasingly common site of communication in the real world. All of the key results from classic work in face-to-face verbal conversations were strongly replicated in our text-based data, suggesting that these processes are modality-general features of communication. Moreover, while a simple heuristic to keep talking until being interrupted could in principle suffice to explain speaker reduction in face-to-face conversation, our evidence shows that speakers *plan* shorter utterances even when they cannot be interrupted; social feedback is being used in a more sophisticated way as latent evidence of understanding. While our analyses of referential content were limited to those produced by the director to preserve the consistency of their source, it is important for future work to more systematically examine the matcher’s contributions to establishing reference. For instance, in some cases the matcher may supply the label that goes on to be conventionalized and may systematically produce different discourse acts in different contexts.

The rapid timescale of adaptation we have investigated in dyadic reference games is not only of interest in its own right for theories of meaning and social coordination; it is a key building block toward grounding the adaptiveness and efficiency of larger-scale human language in the cognitive mechanisms of individual minds (Gibson et al., 2019; Kirby et al., 2015). If community-wide conventions emerge from agents generalizing across different dyadic interactions, then local learning mechanisms leading to efficiency and informativity *within* a dyad may explain how a community’s conventions remain well-calibrated to the demands of the current environment. The sensitivity of such calibration has been previously tested using small artificial languages in the lab (e.g., Winters, Kirby, & Smith, 2014), but our observation of similar dynamics in ordinary natural language use emphasizes that local learning may be an ongoing and pervasive influence. In sum, the resolution provided by the larger corpus we have collected, in combination with recent advances in natural language processing techniques, provides a new window into the quantitative dynamics of adaptation in dyadic communication and beyond. We hope that both the new corpus and new analytic techniques contribute to the testing and elaboration of theories of human language.

Open Research badges



This article has earned Open Data and Open Materials badges. Data and materials are available at <https://osf.io/vzvmf/>.

Notes

1. Available at <https://github.com/wolfgarbe/SymSpell>. We used the smallest model because larger models include typo forms (e.g., “the”) in their vocabulary and thus are not useful for catching errors.
2. A similar reduction curve was found in the “free matching” version of the task, though it required more words overall. Participants needed to additionally mention the spatial location where the tangram needed to be moved (i.e., “number 3 is the ...”).
3. We also pre-registered an analysis examining whether a higher rate of listener messages on early rounds would lead to greater overall reduction in speaker descriptions, but later realized that this analysis was confounded. When listeners send more messages on early rounds, the speaker produced more words in response, which led by definition to greater reduction from this higher initial verbosity.
4. We report the results of a traditional linear regression model because even the most minimal random effect structure encountered singularity issues during optimization. Because matcher errors were relatively infrequent, these singularities were likely caused by an asymmetry in cell size between the correct and incorrect levels of the matcher response variable. However, when we fit a Bayesian regression with maximal random effects, using the default priors implemented by the *brms* package to prevent variances from collapsing to boundary values, we found a nearly identical estimate of the interaction coefficient, $b = -0.31$, 95% credible interval: $[-0.42, -0.19]$.
5. The SpaCy tagger is statistical, meaning it uses a neural network to return the (uniquely) most probable assignment of tags under its model. It obtains comparable accuracy (~97%) to other modern taggers (Manning, 2011). However, it is important to note that the language used in our task likely differs from the tagger’s training sample, containing higher rates of sentence fragments, bare NPs, and “ungrammatical” language that human annotators might also find difficult to classify into standard parts of speech. The SpaCy dependency parser (see below) is similarly statistical and unlikely to be perfect. For example, neural parsers experience particular difficulty determining exactly where to attach conjunction relations (Ficler and Goldberg, 2017). As the state-of-the-art techniques for these tasks continue to improve, we hope future work using our corpus will be able to achieve more precise estimates.
6. Dependency grammars are an alternative to phrase structure grammars, which represent *constituency* relations between abstract units like noun phrases and verb phrases. We pre-registered a different version of this analysis using the pure prevalence of different higher-order clausal features extracted from a constituency parse at each repetition (e.g., adjectival clauses or subordinate clauses). While we found evidence for our pre-registered predictions, we later realized that this measure failed to capture the relevant notion of clustering: Clausal features can disappear

due to rephrasing without the words in the clause necessarily being dropped together, and conversely, entire meaningful units can be dropped at once (e.g., noun phrases) without any clausal features being present at all. Path lengths in the dependency tree are better suited for assessing these questions at the word-by-word level.

7. The first is *term frequency-inverse document frequency* (tf-idf, Sparck Jones, 1972), which multiplies the term frequency $tf(w,d)$ of a word w in a document d by a “global” term $\log(N/n_w)$, where N is the total number of documents and n_w is the number of documents containing w . In our case, the “documents” are just the referring expressions used for a distinct tangram on the first repetition, so $N = 12$ and we can take $tf(w,d)$ to be a boolean for simplicity: 1 if the word occurs, 0 if it does not. We can thus retrieve our simpler measure by exponentiating, dividing by 12, and taking the inverse. The second is *positive point-wise mutual information* (PPMI). Point-wise mutual information compares the joint probability of a word occurring with a particular tangram to the probability of the two occurring independently: $PMI_{\text{word,tangram}} = \log \frac{P(\text{word,tangram})}{P(\text{word})P(\text{tangram})}$. Positive point-wise mutual information is given by $\min(0, PMI)$, restricting the lower bound to 0. It can be shown for our case that *tf-idf* is the maximum likelihood estimator for PPMI: The numerator reduces to a boolean when we only have one observation per tangram (Robertson, 2004). Our pre-registered analysis plan used PPMI, and our result holds using either measure; however, we decided to report the pure tangram count measure as more interpretable.
8. The 43% of final repetition words that did not exactly match were sometimes synonyms or otherwise semantically related to words used on the first repetition, for example, “foot” on the first repetition vs. “leg” on the last. In other cases, the labels used at the end were introduced after the first repetition, for example, one pair only started using the conventionalized label “portrait” on repetition 3.
9. We pre-registered discrete analogs of these analyses, which we have included in Appendix A, but we later decided that the continuous vector space measures present a more interpretable and direct test of our hypothesis.
10. t-SNE is a stochastic, non-linear dimensionality reduction technique which focuses on keeping neighboring points in the high-dimensional space close together in the lower-dimensional space. An initial linear reduction to an intermediate dimensionality is commonly used to speed up computation and reduce noise in the high-dimensional space, compared to applying t-SNE directly to the 300-dimensional vectors. Conversely, the advantage of using t-SNE over projecting directly to two dimensions with a linear technique like PCA is its ability to preserve non-linear structure in the high-dimensional space.

References

- Arora, S., Liang, Y., & Ma, T. (2017). A simple but tough-to-beat baseline for sentence embeddings. In International Conference on Learning Representations (ICLR). Toulon, France, April 24-26.
- Bengio, Y., Ducharme, R., Vincent, P., & Jauvin, C. (2003). A neural probabilistic language model. *Journal of Machine Learning Research*, 3, 1137–1155.
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1482.
- Brennan, S. E., & Hanna, J. E. (2009). Partner-specific adaptation in dialog. *Topics in Cognitive Science*, 1 (2), 274–291.
- Brochhagen, T. (2017). Signalling under uncertainty: Interpretative alignment without a common prior. *The British Journal for the Philosophy of Science*, 71(2), 471–496. <https://doi.org/10.1093/bjps/axx058>
- Carroll, J. M. (1980). Naming and describing in social communication. *Language and Speech*, 23, 309–322.
- Choi, J. D., & Palmer, M. (2012). Guidelines for the clear style constituent to dependency conversion. *Technical Report*. Boulder, CO: University of Colorado.
- Clark, H. H. (1983). Making sense of nonce sense. *The Process of Language Understanding*, 297–331.
- Clark, H. H. (1998). Communal lexicons. In K. Malmkjaer & J. Williams (Eds.), *Context in language learning and language understanding* (Vol. 4, pp. 63–87). Cambridge: Cambridge University Press.
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22, 1–39.
- Corbett, G. G., Fraser, N. M., McGlashan, S. (Eds.), (1993). *Heads in grammatical theory*. Cambridge: Cambridge University Press.
- Davidson, D. (1986). A nice derangement of epitaphs. *Philosophical Grounds of Rationality: Intentions, Categories, Ends*, 4, 157–174.
- Delaney-Busch, N., Morgan, E., Lau, E., & Kuperberg, G. R. (2019). Neural evidence for bayesian trial-by-trial adaptation on the n400 during semantic priming. *Cognition*, 187, 10–20.
- Ficler, J., & Goldberg, Y. (2017). Improving a strong neural parser with conjunction-specific features. In M. Lapata, P. Blunsom, & A. Koller (Eds.), *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers* (pp. 343–348). Valencia, Spain: Association for Computational Linguistics.
- Frank, M. C. (2018). With great data comes great (theoretical) opportunity. *Trends in Cognitive Sciences*, 22, 669–671.
- Fussell, S. R., & Krauss, R. M. (1989). The effects of intended audience on message production and comprehension: Reference in a common ground framework. *Journal of Experimental Social Psychology*, 25, 203–219.
- Garrod, S., Fay, N., Lee, J., Oberlander, J., & MacLeod, T. (2007). Foundations of representation: Where might graphical symbol systems come from? *Cognitive Science*, 31, 961–987.
- Gibson, E., Futrell, R., Jara-Ettinger, J., Mahowald, K., Bergen, L., Ratnasingam, S., Gibson, M., Piantadosi, S. T., & Conway, B. R. (2017). Color naming across languages reflects color use. *Proceedings of the National Academy of Sciences*, 114, 10785–10790.
- Gibson, E., Futrell, R., Piantadosi, S. T., Dautriche, I., Mahowald, K., Bergen, L., & Levy, R. (2019). How efficiency shapes human language. *Trends in Cognitive Sciences*, 23(5), 389–407.
- Glucksberg, S. (2001). *Understanding figurative language: From metaphor to idioms*. Oxford: Oxford University Press.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20, 818–829.
- Hawkins, R. X. D. (2015). Conducting real-time multiplayer experiments on the web. *Behavior Research Methods*, 47, 966–976.

- Hawkins, R. X. D., Frank, M. C., & Goodman, N. D. (2017). Convention-formation in iterated reference games. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th annual meeting of the Cognitive Science Society*. (pp. 482–487).
- Hawkins, R. X. D., Goodman, N. D., & Goldstone, R. L. (2019). The emergence of social norms and conventions. *Trends in Cognitive Sciences*, 23, 158–169.
- Honnibal, M., & Johnson, M. (2015). An improved non-monotonic transition system for dependency parsing. In L. Marquez, C. Callison-Burch, J. Su, D. Pighin, & Y. Marton (Eds.), *Proceedings of the 2015 conference on empirical methods in natural language processing* (pp. 1373–1378). Association for Computational Linguistics.
- Hudson, R. A. (1984). *Word grammar*. Oxford: Blackwell.
- Hupet, M., & Chantraine, Y. (1992). Changes in repeated references: Collaboration or repetition effects? *Journal of Psycholinguistic Research*, 21, 485–496.
- Jurafsky, D., & Martin, J. H. (2014). *Speech and language processing* (Vol. 3, p. 3). London: Pearson.
- Kiperwasser, E., & Goldberg, Y. (2016). Simple and accurate dependency parsing using bidirectional lstm feature representations. *Transactions of the Association for Computational Linguistics*, 4, 313–327.
- Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141, 87–102.
- Kiros, R., Zhu, Y., Salakhutdinov, R. R., Zemel, R., Urtasun, R., Torralba, A., & Fidler, S. (2015). Skip-thought vectors. *Advances in Neural Information Processing Systems*, 3294–3302.
- Kleinschmidt, D. F. (2019). Structure in talker variability: How much is there and how much can it help? *Language, Cognition and Neuroscience*, 34, 43–68.
- Krauss, R. M., Garlock, C. M., Bricker, P. D., & McMahon, L. E. (1977). The role of audible and visible back-channel responses in interpersonal communication. *Journal of Personality and Social Psychology*, 35, 523.
- Krauss, R. M., & Weinheimer, S. (1964). Changes in reference phrases as a function of frequency of usage in social interaction: A preliminary study. *Psychonomic Science*, 1, 113–114.
- Krauss, R. M., & Weinheimer, S. (1966). Concurrent feedback, confirmation, and the encoding of referents in verbal communication. *Journal of Personality and Social Psychology*, 4, 343.
- Kübler, S., McDonald, R., & Nivre, J. (2009). Dependency parsing. *Synthesis Lectures on Human Language Technologies*, 1, 1–127.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211.
- Lascarides, A., & Copestake, A. (1998). Pragmatics and word meaning. *Journal of Linguistics*, 34, 387–414.
- Levy, K. E., Freese, J., & Druckman, J. N. (2016). The demographic and political composition of mechanical turk samples. *Sage Open*, 6, 2158244016636433.
- Lewis, D. (1969). *Convention: A philosophical study*. Cambridge, MA: Harvard University Press.
- Maaten, L. V. D., & Hinton, G. (2008). Visualizing data using t-sne. *Journal of Machine Learning Research*, 9, 2579–2605.
- Manning, C. D. (2011). Part-of-speech tagging from 97% to 100%: Is it time for some linguistics? In A. F. Gelbukh (Ed.), *International conference on intelligent text processing and computational linguistics* (Vol. 6608, pp. 171–189). Springer.
- Markman, A. B., & Makin, V. S. (1998). Referential communication and category acquisition. *Journal of Experimental Psychology: General*, 127, 331.
- Metzing, C., & Brennan, S. E. (2003). When conceptual pacts are broken: Partner-specific effects on the comprehension of referring expressions. *Journal of Memory and Language*, 49, 201–213.
- Nölle, J., Staib, M., Fusaroli, R., & Tylén, K. (2018). The emergence of systematicity: how environmental and communicative factors shape a novel communication system. *Cognition*, 181, 93–104.
- Osgood, C. E. (1952). The nature and measurement of meaning. *Psychological Bulletin*, 49, 197.

- Pennington, J., Socher, R., & Manning, C. (2014). Glove: Global vectors for word representation. In A. Moschitti, B. Pang, & W. Daelemans (Eds.), *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1532–1543). ACL.
- Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27, 169–190.
- Prelec, D. (2004). A bayesian truth serum for subjective data. *Science*, 306, 462–466.
- Robertson, S. (2004). Understanding inverse document frequency: On theoretical arguments for IDF. *Journal of Documentation*, 60, 503–520.
- Salton, G., & Buckley, C. (1988). Term-weighting approaches in automatic text retrieval. *Information Processing & Management*, 24, 513–523.
- Schober, M. F., & Clark, H. H. (1989). Understanding by addressees and overhearers. *Cognitive Psychology*, 21, 211–232.
- Smith, N. J., Goodman, N. D., & Frank, M. (2013). Learning and using language via recursive pragmatic reasoning about other agents. In C. J. C. Burges, L. Bottou, Z. Ghahramani, & K. Q. Weinberger (Eds.), *27th Annual Conference on Neural Information Processing Systems*, (pp. 3039–3047).
- Sparck Jones, K. (1972). A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation*, 28, 11–21.
- Wilkes-Gibbs, D., & Clark, H. H. (1992). Coordinating beliefs in conversation. *Journal of Memory and Language*, 31, 183–194.
- Winters, J., Kirby, S., & Smith, K. (2014). Languages adapt to their contextual niche. *Language and Cognition*, 1–35.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article:

Fig. S1. Permuting utterances across pairs increases entropy of word distribution, consistent with internal stability and multiple equilibria. Mean empirical entropy (red) and mean permuted entropy (blue) are shown for each tangram. Error bars are 95% CIs for bootstrapped empirical entropy and the permuted distribution, respectively.

Fig. S2. The empirical dependency lengths between dropped words are lower than expected under two baselines for every repetition block. Samples from the baselines are shown as densities.

Fig. S3. t-SNE visualizations of utterance trajectories for all 12 tangrams; panel C is annotated in Fig 6.

Appendix A: Discrete word distributions

Here we examine an alternative approach to evaluating claims of arbitrariness and stability using discrete word distributions instead of the continuous vector space measure used in the main text. We begin by examining the discrete *distribution of words* that each pair uses to refer to each tangram, excluding stop words. This distribution is a *unigram* distribution over the vector of words that appear throughout the utterances produced by a given speaker to refer to a particular object (the modal support size of this distribution is 7 words.) If a pair of participants converges on stable labels for a tangram, this stability should manifest in a highly structured distribution over words throughout the game for that pair. If different speakers discover diverging conventions, this idiosyncrasy should manifest in differing word distributions. We formalize these intuitions by examining entropy, an information-theoretic measure:

$$H(W) = \sum_w P(w) \log P(w)$$

The entropy of the word distribution for a pair is maximized when all words are used equally often and declines as the distribution becomes more structured, that is, when the probability mass is more concentrated on a subset of words. It also increases as a function of the support size; because in principle we consider this an important signature of a game, we focus on this unnormalized measure; however, the results hold if we control for the support size (i.e., divide the entropy by $\log(N)$ so that a uniform distribution will always have the maximum value of one.)

To compare word distributions across games, we use a permutation test methodology. By scrambling referring expressions for each tangram across games and recomputing the entropy of the scrambled word distribution, we effectively disrupt any structure within each pair. There are two important inferences we can draw from this test. First, in a null scenario where different pairs did *not* diverge as predicted and instead every pair coordinated on roughly the same (optimal) convention for each tangram, this permutation operation would have no effect since it would be mixing together copies of the same distribution. Second, in another null scenario where pairs did not converge and instead varied wildly in the words they used from repetition to repetition, then permuting across games would also have no effect since it would simply mix together word distributions that already have high entropy. Hence, scrambling should *increase* the average game's entropy only in the case where both predictions hold: each game's idiosyncratic but concentrated distribution of words would be mixed together to form more heterogeneous and therefore high-entropy distributions.

Following this logic, we computed the average within-game entropy for 1000 different permutations of director utterances. We permuted utterances within repetition blocks rather than across the entire data set to control for the fact that earlier trials may generically differ from later ones (e.g. in utterance length). Because we are permuting and

measuring entropy at the tangram-level, this yields 12 permuted distributions (see Fig. S1). We found that the mean empirical entropy lay well outside the null distribution for all 12 tangrams, $p < .001$, consistent with our predictions of internal stability within pairs and multiple equilibria across pairs.

Finally, it is worth noting some advantages and disadvantages of this discrete measure compared to the continuous vector space measure used in the main text. A key advantage is that the entropy is not dependent on any particular choice of pre-trained vector embedding. Due to biases in the vocabulary of their training corpora, vector representations also may not capture some of the more idiosyncratic conventions that participants converge on (e.g. “zig zag” or “Frank” —short for “Frankenstein”). Thus, to the extent we find converging results, the discrete measure may address concerns about the quality of the continuous representation. A key disadvantage, on the other hand, is that our permutation test methodology is more indirect and does not have a natural scale. We can strongly reject the complete absence of arbitrariness and stability—a lower bound—but there is no clear derivation for a corresponding upper bound showing exactly how strong these effects are. Directly measuring divergence between word distributions is technically possible using divergence measures, but would not be informative at the fine granularity required for these analyses (i.e., at the level of single utterances). Most utterances use entirely disjoint sets of words, and on later repetitions, the distribution may only contain one or two contentful words. A final disadvantage is that discrete analyses treat even close synonyms as entirely distinct tokens in the word distribution because they are based entirely on the frequency of tokens rather than semantic content. In summary, these two approaches provide complementary and converging evidence.

Appendix B: Additional baselines for evaluating divergence

Could the divergence effect reported in Section 3.2 be explained away as an artifact of our procedure for computing utterance embeddings? If averaging together greater numbers of word vectors generically causes the resulting utterance vectors to be washed out and more similar one another, then the decrease in semantic similarity could be explained by a decrease in utterance *length* (see Section 4.2) rather than divergence in content. We tested this null hypothesis using a further permutation test. We reasoned that if the effect is in fact driven by length, then the similarity measured across interactions should be invariant to re-sampling utterance *content*—the individual words that will be averaged together—within interactions. We thus scrambled the words used by a participant across all 12 tangrams at each repetition, destroying any tangram-specific semantic content, but preserving utterance length. By repeating this procedure 100 times, we found that the true mean similarity across pairs was higher than predicted under the null distribution at all six repetitions, $p < .01$, suggesting that the divergence effect is not solely driven by utterance length.

At the same time, we observed that this permutation test disrupted the mean similarity less than expected. On the first repetition, for instance, the range of the null samples was [0.754, 0.764], only slightly lower (in absolute terms) than the empirical value of 0.774.

Why would this be the case, and how should we interpret the absolute degree of divergence? One possibility is that there is already substantial semantic overlap on the first round in how a single speaker refers to different tangrams, so that scrambling does not dramatically disrupt the semantic content. This possibility suggests examining the divergence between utterances used to refer to *different tangrams within an interaction* as a useful baseline. Based on our results in Section 3.1, we predicted that pragmatic pressures would lead labels for different tangrams *within* an interaction to diverge more strongly than those for the same tangram *across* interactions, despite starting with roughly similar overlap. Indeed, we found that the average semantic similarity within an interaction was indistinguishable from the similarity across interactions on the first repetition (*paired difference*: 0.003) but the gap appeared to widen over subsequent rounds, indicating that the pressure to distinguish tangrams leads to greater divergence for a single speaker than the neutral divergence across different speakers would predict.

To test the statistical significance of this observation, we conducted a model comparison between mixed-effects models. The dependent variable in both models is the difference score between mean within-speaker and across-speaker similarities (aggregated at the level of the speaker). In the null model, we include only an intercept, which allows for a non-zero difference but does not allow this difference to increase or decrease with time. In the full model, we additionally include a linear term for repetition number. Because we have a mean difference score for each speaker, we also include random intercepts at the speaker level for both models. A likelihood ratio test between these models shows that the full model fits the data significantly better, controlling for the additional degree of freedom, $\chi^2(1) = 10.8$, $p < .001$. We have focused on this comparison to hold random effects constant, but including an additional term for a quadratic effect of repetition and an additional random effect of repetition are also supported by likelihood ratio tests. In this full model, we find a marginally significant linear effect of repetition, $b = 0.15$, $t = 1.9$, $p = .059$.

To summarize, we suggested the use of a baseline to better interpret our core result showing divergence in labels across different speakers as pairs discover different conventions. This baseline—the divergence in a speaker's *own* utterances for different tangrams—begins at a similar level, indicating that the initial utterances used by different speakers overlap approximately as much as the different initial utterances used by a single speaker. While both subsequent trajectories indicate divergence, the different labels used by a single speaker rapidly spread out in vector space and become more distinct from one another than the labels used by different speakers.