Anaphoric reference in descriptions of surprising events[1]

Asya ACHIMOVA — *University of Tübingen*
Christian STEGEMANN-PHILIPPS — *University of Tübingen*
Susanne WINKLER — *University of Tübingen*
Martin V. BUTZ — *University of Tübingen*

Abstract. The choice of anaphoric reference is a complex process regulated by a combination of linguistic and cognitive constraints. This paper experimentally addresses the impact of world knowledge on the types of references speakers produce, focusing on the predictability of event progressions. In order to avoid confounding event predictability and the predictability of words in their descriptions, we created an artificial virtual world and trained participants to recognize typical event progressions within. Speakers then described novel scenes, which either conformed to their expectations or violated them, in a free production experiment. The data reveal that surprising event progressions lead to a more frequent production of definite noun phrases, in contrast to reduced linguistic expressions (pronouns and zero anaphors). We further introduce a Bayesian inference model, which offers an explanation of why a definite noun phrase and not a pronoun is used for the description of surprising events. We conclude that speakers choose more informative descriptions to talk about surprising events because this strategy allows them to avoids misunderstanding with higher likelihood.

Keywords: Referring expressions; anaphors; speech production experiment; event-predictive cognition; world knowledge.

1. Introduction

Anaphoric reference—reference to a previously mentioned entity—can take at least three distinct forms: a definite noun phrase (1a), a pronoun (1b), and a null subject (1c).

(1) A friend of mine sent me a letter three weeks ago.
   a. The letter finally arrived today.
   b. It finally arrived today.

[1]We would like to thank Tizian Thieringer, Lena Holzwarth, Jannis Strecker, and Loredana Columbo for their help in setting up the experiment and annotating the data. We are also grateful to the audience of Sinn und Bedeutung 26 for their insightful feedback. This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project number 198647426. Martin V. Butz is also a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645.

c. Finally arrived today[2].

When choosing referring expressions, speakers face competing pressures. On the one hand, they select those forms of reference that maximize the chance of the listener to correctly identify the intended object or person. In that sense, providing longer, more detailed descriptions might be beneficial. On the other hand, speakers minimize their production efforts: they choose the least expensive means to express the desired content. Both, the theoretical pragmatic literature starting with the work of Grice (1989) and game-theoretic modeling (Franke, 2009; Frank and Goodman, 2012), emphasize the role of comprehensibility of the utterance for the listener as a driving force determining speaker choices.

In this paper, we experimentally manipulate the predictability of described events and investigate how their predictability status affects their linguistic encoding. Our goal here is to determine how world knowledge, operationalized as knowledge of common event patterns, affects the production of anaphoric references.

2. The choice of anaphoric references

A number of syntactic, semantic, and pragmatic constraints regulate the choice of referring expressions in situations where there is more than one potential antecedent for the reference. On the one hand, in line with the Gricean maxim of manner, to be clear but brief Grice (1989), speakers are expected to refrain from using a full noun phrase when referring to a previously mentioned entity. In fact, repeating a name causes a disruption in discourse coherence and slows down reading times. This effect is known as repeated-name penalty (Gordon et al., 1993) and has been shown to be robust in a number of languages (Almor, 1999; Gordon and Hendrick, 1998). In a separate line of work, Williams (1997) argued that choosing a definite NP as a form of anaphoric reference is legitimate only if the less phonologically prominent forms are blocked.

On the other hand, using a pronoun or a zero anaphor as a form of reference comes with their own challenges. A pronoun may be potentially compatible with multiple antecedents that match it in number and gender. Pronoun interpretation is regulated both by linguistic constraints and contextual factors, including semantic knowledge, or world knowledge. Thus, in (2), world knowledge guides the only permitted co-indexing in (2a) and (2b).

[2]While the form in (1c) may not be available in standard English, it is common in telegraphic speech (Barton, 1998; Haegeman, 2013).

(2) a. The ball$_i$ fits in the suitcase$_j$ because it$_{i/*j}$ is small.
b. The ball$_i$ fits in the suitcase$_j$ because it$_{*i/j}$ is large.

(Butz, 2017: 6)

Participants in speech production experiments seem to be sensitive to potential ambiguity of pronouns. For instance, they refrain from producing a pronoun if the potential referents for it share relevant characteristics (Fukumura et al., 2011).

Unlike pronouns, zero anaphors do not introduce ambiguity in a sentence, because they cannot be freely co-indexed both with the subject and the object of the previous phrase. Thus, (3) can only mean that the virus ran away.

(3) The green virus$_i$ attacked the red cell$_j$ and $t_{i/*j}$ ran away.

At the same time, zero anaphors introduce a challenge of a different sort: a combination of a noisy auditory channel (Shannon, 1948) and a surprising event structure may prompt a phonological repair on the side of the listener. Evidence from eye-tracking experiments suggest that readers maintain uncertainty over the identity of previously read words and reinterpret these words in accordance with the upcoming linguistic material in case of a conflict (Levy et al., 2009). Computational modeling and experimental evidence suggest that this integration process can be viewed as a rational statistical inference assuming a noisy channel (Gibson et al., 2013). For reduced anaphoric references, such as zero anaphors in particular, these models of input integration suggest that zero anaphors can be reinterpreted as some other linguistic material if the message cannot be successfully processed without the repair—a situation quite possible when the utterance refers to an unpredictable event.

3. Effect of Predictability

Predictability has been shown to affect multiple levels of linguistic encoding. Thus, on the phonological level, words that mark suprising information may carry higher acoustic prominence (Lam and Watson, 2010). On the syntactic level, predictability affects the use of overt complementizers, such as *that*: they are more likely to appear in a sentence if speakers judge its content to be less predictable (Jaeger, 2010). At the level of morphological structure, case markings have been shown to be placed strategically when sentence properties, such as the animacy of the object or plausibility in general, suggest a different structure (Kurumada and Jaeger, 2015). Corpus studies further reveal that referential expressions in writ-

ten text alternate between full descriptive reference, proper name, and pronoun depending on predictability (Tily and Piantadosi, 2009).

Thus, multiple experimental studies suggest that predictability of the upcoming linguistic material affects different aspects of linguistic encoding. In this paper, we ask whether predictability can be viewed as a cognitive category that acts at the level of world knowledge. We focus on a specific aspect of world knowledge: the structure of events. Following Zacks and Tversky (2001); Zacks et al. (2007), we define an event as a spatio-temporal unit that has a clear beginning and an end. Event representations provide a structure for much of our general ability to predict what happens in the world (Bunger et al., 2013; Zacks and Tversky, 2001; Zacks et al., 2007; Radvansky and Zacks, 2014; Franklin et al., 2020). Such knowledge allows the speaker to form expectations about events that are likely to occur, and induce the effect of surprise when we are confronted with unexpected or new events (Baldwin and Kosie, 2021; Butz et al., 2021; Kuperberg, 2021). Predictability is relevant to the concept of events on multiple levels, including their internal dynamics defined by a set of participants and possible interactions, and the predictability of event progressions.

Our experiment is aimed at testing whether the types of references change depending on whether participants describe a predictable or a surprising event. We further develop a Bayesian model of reference choice that factors in predictability of event progressions and includes uncertainty over the speech stream as one of its critical components.

4. Production experiment

Studying the effect of event predictability on linguistic choices is often confounded by the predictability of words in the descriptions of such events. In this section, we present a more immediate way to explore the influence of world knowledge by training participants to recognize the patterns of novel events in an artificial virtual world (Stegemann-Philipps et al., 2021)[3]. The participants are then asked to describe the events verbally in a free-production task[4]. This type of narration differs from scripted story continuation adopted in other psycholinguistic studies (e.g. Rosa and Arnold (2017)).

[3]Stegemann-Philipps et al. (2021) presented preliminary analysis of the data that we report in this paper.

[4]The experiment is available at `https://stegemann.connotations.de/mturk-prod3/mturkexp-prod3.html`

4.1. Design

Our main experimental manipulation concerned the predictability of the event progressions. In half of the trials, the speakers had to describe familiar events, while in the other half an unexpected event progression was shown. We predict that this manipulation will affect the types of references speakers produce: to be more precise, we expect a higher proportion of definite NPs as a reference to the agent of surprising events.

The scenes created in the *BrainControl* virtual world (Schrodt et al., 017b), featured three types of moving entities: long blue cells, spiky red cells, and round green viruses[5]. A collision between two entities forces one of them to flee: a blue cell fled from the green virus, and the green virus fled from the red cell. The cells themselves did not interact. We include all possible character interactions in (4) and show a possible scen in Figure 1.

(4) a. Green virus attacks blue cell; blue cell runs away.
b. Green virus attacks red cell; green virus runs away.
c. Red cell attacks green virus; green virus runs away.
d. Blue cell attacks green virus; blue cell runs away.

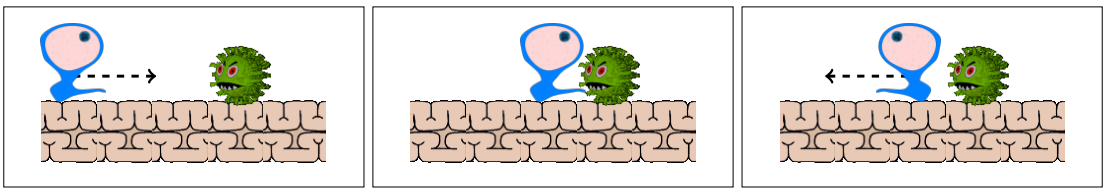

Figure 1: Exemplar trained interaction event: the blue cell attacks the virus and flees.

4.2. Testing procedure

The experiment started with a learning phase, were participants were then shown the four interactions given in (4a-d). We then tested the participants' knowledge of the basic event patterns in (4).

The production phase included four possible interactions in (4), as well as their surprising counterparts, where the other character fled, i.e. a red cell fled from a green virus and a green virus fled from a blue cell. Figure 2 provides an example

[5]We collected the data for this project before the Covid-19 pandemic unfolded. It is an unfortunate coincidence that our stories feature viruses as characters.

of such a surprising version, essentially directly reversing the pattern shown in Figure 1.

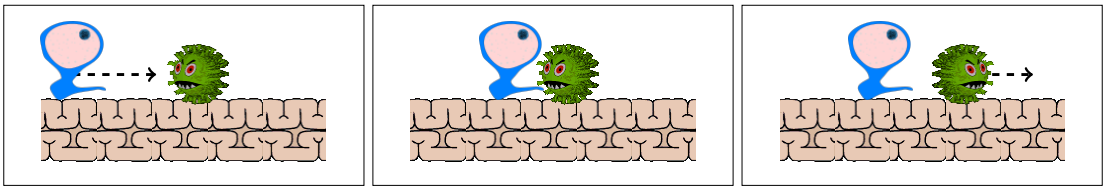

Figure 2: Exemplar surprising interaction event: the blue cell attacks the virus; the virus flees.

The participants were asked to follow the prompt: *Simply describe what is going on!* The participants' linguistic descriptions involved unconstrained spontaneous speech production. We explicitly avoided priming the use of NPs or reduced utterances as well as any specific grammatical constructions. We anticipated that to describe the events shown in Figure 1, the participants might produce one of the utterances in (5):

(5) a. The blue cell attacks the virus and the blue cell flees.
b. The blue cell attacks the virus and it flees.
c. The blue cell attacks the virus and flees.

Table 1: The 2 x 2 design of the trial stimuli

	agent of 1st sub-event fleeing	patient of 1st sub-event fleeing
predictable	B attacks A, B flees *Condition 1*	A attacks B, B flees *Condition 3*
surprising	A attacks B, A flees *Condition 2*	B attacks A, A flees *Condition 4*

The experiment featured a 2 x 2 design (Table 1), where we manipulated the thematic role of the fleeing character in the first sub-event (agent or patient) and the predictability pattern. In Figure 2, it is the patient of the first sub-event that flees in the second sub-event. In Figure 1, however, the same character acts as agent in both sub-events. The predictability pattern could take one of two values: the second sub-event could either be predictable (Figure 1) or surprising (Figure 2). Participants viewed and described two scenes for each of the four conditions.

We conducted the experiment online and recruited 300 participants with US IP addresses through the online platform *Amazon Mechanical Turk*. Participants received compensation for their participation with a rate of $10/hour. Out of 300 participants, we excluded data from 31 people because of bad audio quality and 25 people because they did not pass the learning test block. We further excluded

11 participants because their descriptions did not conform to our schema at all, mostly because they produced predictions instead of descriptions. Data from 231 participants (1864 descriptions) was included in the analysis.

4.3. Annotation and Analysis

Technical assistants blind to the purpose of the experiment transcribed the audio recordings that we obtained from the study participants. Despite the fact that speech production was fully unconstrained, participants produced a range of utterances that conformed to our expectations (e.g. 6a - 6d).

(6) Sample descriptions:
   a. The red cell just attacked the virus and it ran away.
   b. The virus attacks the blue cell, the blue cell runs away to the left.
   c. The green virus attacks and hits the blue cell, but then runs away.
   d. Virus eats the red blood cell and then the virus runs.

We annotated the utterances as to whether the subject of the second verb in the event description was a full noun phrase (NP) or a pronoun. As an alternative, sometimes the speakers produced a conjoined verb phrase (5c). We counted such responses as zero anaphors.

Responses that did not fit any of the categories were tallied separately. Such descriptions either did not conform to subject verb structure, which allows the use of a definite NP or a pronoun (7), or the speakers provided a summary of a whole scene rather than describing the sub-events (8).

(7) Virus and cell attack each other.

(8) The cell defeats the virus.

If the predictability of an event affects the form of reference, we expect more definite NPs in descriptions of surprizing events. This effect would be similar to an m-implicature (manner implicature): speakers choose a marked form of a linguistic expression to describe an unusual event (Levinson, 2000).

4.4. Results

Participants produced both definite NPs and reduced forms when both sub-events contained the same character as the agent (conditions 1 and 2), using reduced

expressions in 34.4% and definite NPs in 65.6% of the cases[6].

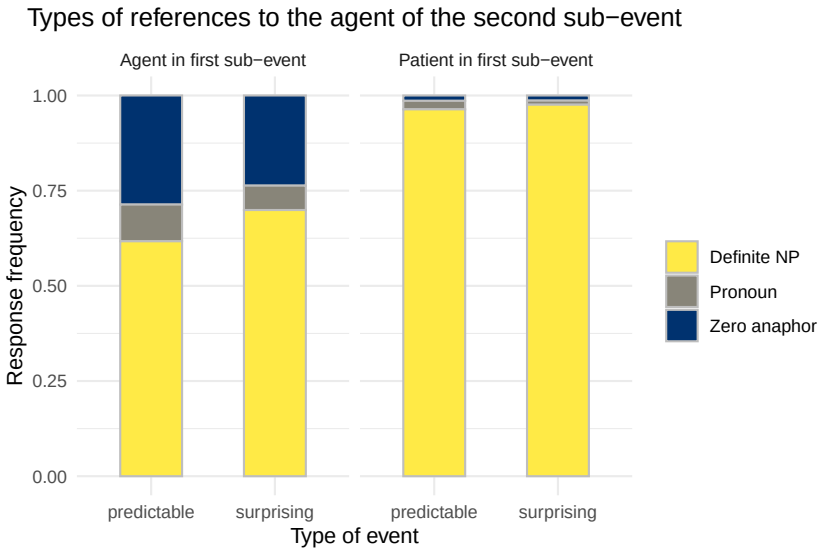

Figure 3: Production of referring expressions for the agent of the second sub-event

In the conditions where the patient of the first sub-event became the agent of the second sub-event (conditions 3 and 4), speakers overwhelmingly produced definite NPs to refer to the agent of the second sub-event (97% of the time), independent of whether the event structure violated the patterns they learned or not. Conditions that feature the agent vs. patient of the previous sub-event differ not only in the thematic structure of the events themselves but also in discourse relations (Elman and McRae, 2019; Kehler et al., 2008) between the two sub-events. Moreover, the topic structure between these conditions is also non-identical. The events, where the patient of the first sub-event becomes the agent in the second sub-event are associated with a possible topic shift, priming a definite NP to refer to the agent in the second sub-event. All these factors may have contributed to the overwhelming preference for definite NPs in conditions 3 and 4. Due to this ceiling effect in these conditions, we do not analyze this data further and now concentrate on conditions 1 and 2, where the same character acted as agent in both sub-events.

Figure 3 shows a distribution of different referring expressions depending on the condition. We fitted a cumulative link mixed-effects model using the `ordinal` package in R (Christensen, 2019) to analyze the data. The type of referring expression served as the dependent variable with three levels: definite NPs, pro-

[6]Data and analysis files are available at
`https://osf.io/u43jn/?view_only=1aec01b2249c419eb59af0d448d68dbb`

nouns, and zero anaphors. Predictability of the event and trial order were treated as independent variables. The random effect structure included subjects and items as random intercepts, as well as random slopes for surprise per subject. The data reveal that the choice of the referring expression depends on the predictability of an event ($\beta = 0.639$, $SE = 0.257$, $z = 2.586$, $p = 0.013$). Trial order also affected the types of responses ($\beta = 0.146$, $SE = 0.045$, $z = 3.273$, $p = 0.001$). Subsequent inspection revealed that speakers were overall more likely to produce a pronoun or a zero anaphor for the first trial than for consecutive trials. We then excluded the first trials from the analysis and repeated model fitting, the results were qualitatively similar to the ones we report for the full sample.

An alternative analysis, where we grouped pronouns and zero anaphors together to form a 'reduced' category to contrast with definite NPs also confirmed the effects we observed with a multinomial regression analysis. We fitted a binomial mixed effects model with event predictability and order as independent variables and random intercepts for participants, as well as random slopes for the effect of predictability per participant. Models that also included a random intercept for items resulted in singular fits. Speakers were more likely to use a definite NP for surprising events ($\beta = 0.566$, $SE = 0.204$, $z = 2.78$, $p = 0.005$), confirming the results of ordinal mixed effect modeling.

Overall, the results corroborate further evidence that speakers tend to produce more precise descriptions (definite NPs) when event participants engage in surprising interactions. In the next section, we propose a computational model that formalizes the benefit of this behavioral tendency.

5. Computational Modeling

Our formalization grows out of two lines of modeling: the Rational Speech Act (RSA) framework (Frank and Goodman, 2012; Goodman and Frank, 2016) and the event-predictive cognition approach (Butz, 2016; Butz and Kutter, 2017; Butz et al., 2021). The RSA architecture provides the core components of the model—the functions of speaker and listener, while the event-predictive coding approach defines the goal of utterance interpretation as reaching predictive consistency. The RSA architecture also offers tools for incorporating the effect of world knowledge on language comprehension and production since as a Bayesian model, RSA allows for building in prior expectations into the calculations (Degen and Tonhauser, 2021). Bayesian models have been successfully used to model pronoun interpretation Kehler et al. (2008). We further relate our model to computational

models of utterance comprehension under the noisy channel (Gibson et al., 2013) and pursue the idea of integrating prior knowledge (in our case the knowledge of possible event progressions) with probabilistic linguistic evidence.

Within the RSA framework utterance choice is modeled by considering how the listener would potentially interpret the utterance. The goal of the speaker is to convey a particular message m in such a way that the listener's interpretation i of the speaker's utterance u corresponds to m. Meanwhile, though, the speaker wants to stay brief, which is typically modeled by a general penalty term $c(u)$ for a particular utterance choice u. Thus, the utility of an utterance u can be quantified as the exponential of the probability to that an utterance is understood correctly minus the cost for producing the utterance:

$$U_{speaker}(u;m) \propto \quad e^{\alpha P(I=m|u)} - c(u), \tag{1}$$

where factor α emphasizes the importance of being understood. We denote the interpretation of the listener given an utterance u by $P(I|u)$ as a probability density over all possible interpretations I (assumed to be discrete for simplicity reasons). Being understood corresponds to the case when the utterance is interpreted as intended, that is, $P(I = m|u)$. Utterance choice then is governed by maximizing utterance utility. Thus, the policy π of choosing a particular utterance u when intending to convey message m can be denoted by:

$$u \leftarrow \pi(m) \leftarrow \quad \arg\max_u U_{speaker}(u;m), \tag{2}$$

where we assume that the speaker essentially uses her own model as a proxy for inferring the listener's interpretation probability $P(m|u)$.

In our particular case, we need to formalize how the listener may end up at a particular event interpretation $i \in I$ given the utterance u of the speaker. The speaker's intention is that the actual event interpretation $i^t \in I$ at the current point in time t will correspond to the intended message m^t. We assume that the listener interprets the received utterance u using her event-predictive knowledge.

Syntactic and semantic constraints on co-reference determine whether one or both antecedents may serve as the agent of the second sub-event. We denote these conditional probabilities by:

$$P(I_u^t|u^t), \tag{3}$$

conditioning on the utterance. We assume that definite NPs select the referent

unambiguously[7]. Pronouns, however, may potentially refer to any of the two antecedents but exhibit a co-reference structural bias to the entity encoded as the grammatical subject (Crawley et al., 1990). We denote this pronoun bias by ρ, where $\rho = 0.5$ corresponds to an equal referent probability, while $\rho = 1$ would correspond to an unambiguous reference to the agent. In the case of a zero anaphor, grammatically the reference is only possible to the agent of the previous sub-event.

Meanwhile, for the events under consideration, we define the event knowledge as prior probabilities for a particular next event given the chosen previous event interpretation i^{t-1}:

$$P(I_e^t | i^{t-1}), \tag{4}$$

which would in our case indicate that after an attack it is more likely that the blue cell flees from the green virus, while the green virus flees from the red cell. We determine this temporal prior of events co-occurring by a model parameter ϕ, where $\phi = 0.5$ corresponds to no expectation and $\phi = 1$ corresponds to a full prior belief that the expected progression will unfold.

We can now compute the joint interpretation probability by fusing the temporal prior event probability with the event interpretation encoded in the linguistic form of the utterance. Assuming information independence, the model essentially fuses the temporal prior $P(I_e^t | i^{t-1})$ with the utterance-based evidence $P(I_u^t | u^t)$:

$$P(I^t | u^t, i^{t-1}) = w(u) P(I_u^t | u^t) + (1 - w(u)) P(I_e^t | i^{t-1}) \tag{5}$$

We assume that $w(u)$ differs only between the full noun phrases and the reduced forms. For the full noun phrases, we assume maximal evidence from the utterance, setting $w(u = NP) = 1$. For the reduced forms, we are interested in parameter variations. Thus, we introduce a fourth model parameter: $w(u \neq NP) = \omega$.

With these specifications, the model enables us to compute exact probabilities of being understood with $P(i^t = m | u^t, i^{t-1})$, that is, the probability of understanding the utterance u^t as the intended message m^t. Figure 4 illustrates the process embedded in the complete speaker-listener interaction. Thus, we model the influence of temporal event regularities as well as the independent influence of the linguistic encodings.

[7]In our set up, there is no definite NP that could simultaneously refer to both interacting entities and thus be ambiguous.

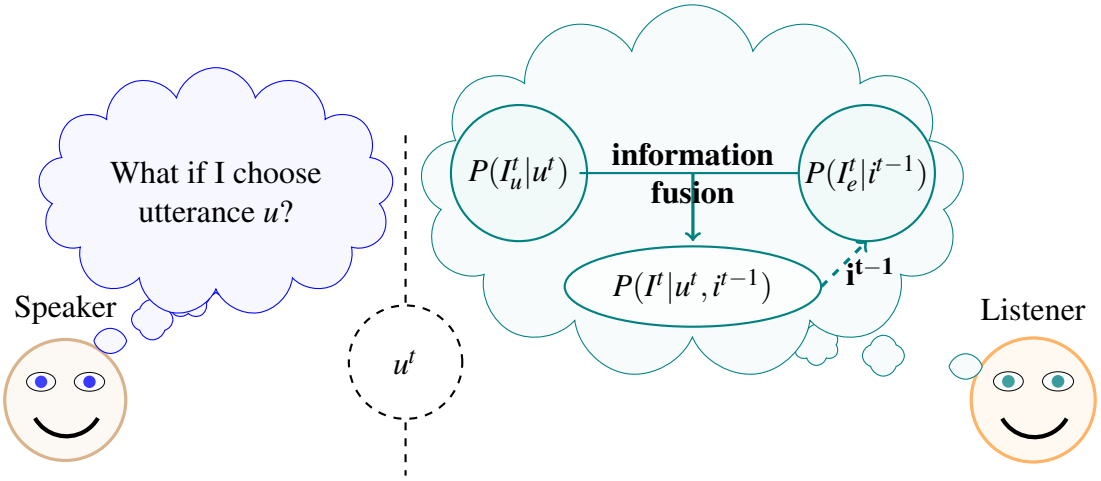

Figure 4: Considered speaker and listener interaction and assumed inference model in the listener, which the speaker may simulate to infer an utterance choice.

Our model thus contains four parameters: the exponential factor α scales the importance of being understood, ρ specifies the strength of subject bias (Crawley et al., 1990) in the pronoun interpretation[8], ϕ denotes the bias of the temporal event prior towards favoring the learned event progression, while ω denotes the signal certainty. The latter defines the fusion weight of trusting what the person heard over knowing which events are possible. Considering the eight data points from the averaged experimental results, a model parameterization with $\alpha = 6$, $\rho = .75$, $\phi = .6$, and $\omega = .7$ yields the results shown in Figure 5, which closely match to the the data from the experiment presented above in Figure 3[9].

To illustrate the role of different parameters in the model, we plot the effect of changing their values on model predictions in Figure 6. A very small exponential factor of $\alpha = 1$ decreases the utility of choosing a definite NP, while an even much larger factor of $\alpha = 20$ yields a close to 100% NP preference in all four event progression types. If we remove the subject bias ($\rho = .5$), the pronoun becomes fully ambiguous and it becomes less likely to be chosen. At the same time, if we increase the subject bias to $\rho = 1$, then likelihood of choosing a pronoun

[8]We opted for subject bias as a simple heuristic that affects pronoun interpretation in this model. This parameter could be replaced with a component that takes into account coherence relations (Elman et al., 2006; Kehler and Rohde, 2013, 2019) and therefore modulates the bias depending on the type of coherence relations between the two sub-events.

[9]We assume zero cost in Equation (1), although a prior preference for shorter descriptions expecting to follow Grice's maxim of brevity (Grice, 1989) may also be part of the model. At the moment, though, without further data, variations in the cost for the noun phrase relative to the reduced anaphors would make the model parameters non-identifiable. We thus do not allow varying the cost c, setting it to zero for all utterances.

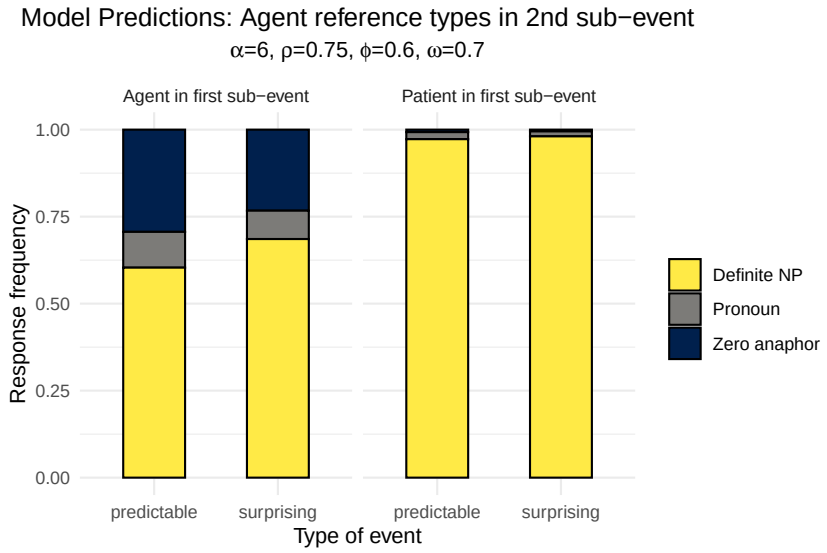

Figure 5: Model results with approximately optimal parameter values

equals the zero anaphor case. No particular temporal expectation of the event progression ($\phi = .5$) results in no distinction between predicted and surprising cases, while a strong expectation ($\phi = 1$) yields a stronger contrast between the two cases. Finally, setting the signal certainty to 0 ($\omega = 0$), makes the model ignore the linguistic evidence of reduced anaphoric references, fully relying on the prior (event knowledge). The linguistic evidence produced by definite NPs, on the other hand, is fixed to 1, that is, full evidence, essentially ignoring the event knowledge. When full evidence is granted also to the reduced expressions ($\omega = 1$), zero-anaphors act like definite NPs[10].

In sum, the proposed utterance choice model relies directly on the probability of temporal event progressions and suggests a fusion process integrating independent temporal and utterance-given information sources. As a result, the model predicts that in the case of a predictable event progression, speakers tend to choose reduced reference more frequently than in the case of unexpected event progressions, because the chances of being misunderstood in the former case are smaller than in the latter case.

6. Conclusion

We have presented empirical evidence demonstrating that speakers are more likely to use definite NPs to refer to agents of surprising events. By creating an artificial world set up, we were able to isolate the effect of knowledge about event pro-

[10]To explore the effect of different parameter values of the model on its predictions, please visit `https://cognitivemodeling.shinyapps.io/shinyanaphor/`

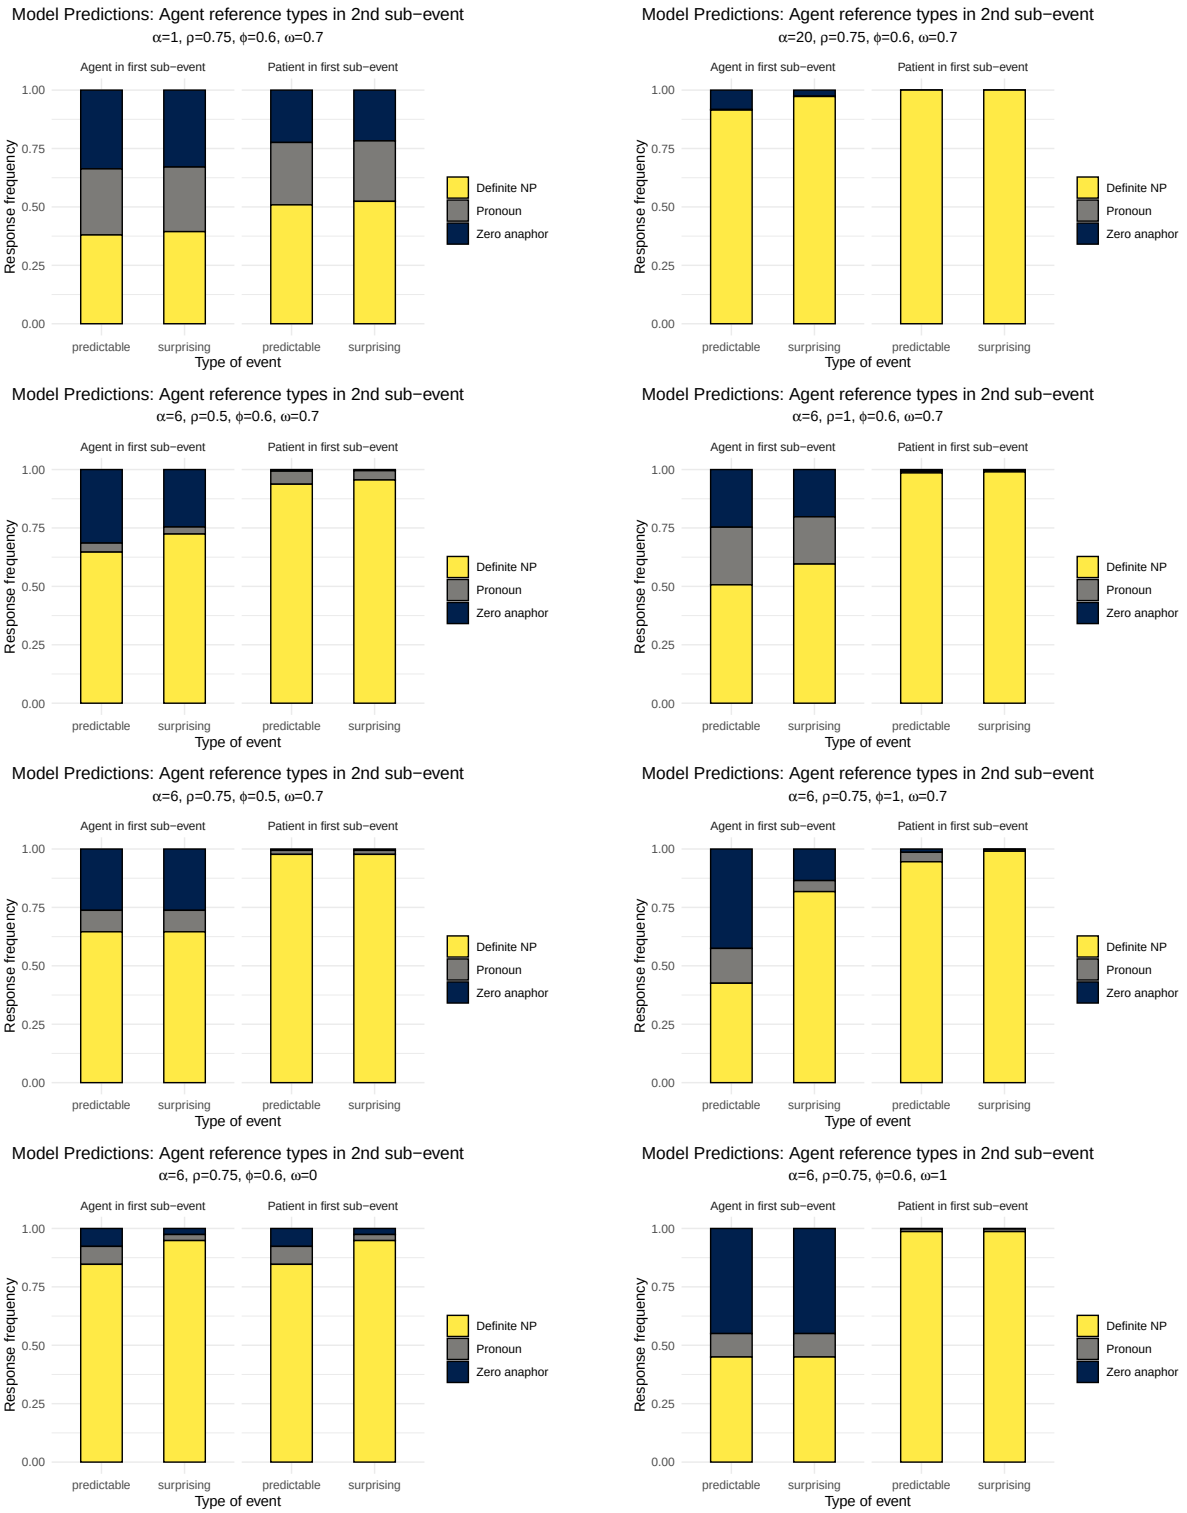

Figure 6: Model results when systematically varying the parameters from the optimum: rows consider variations in parameter α, ρ, ϕ, and ω, respectively. Columns show effect of decreasing and increasing the considered parameter value.

gressions from the distributional regularities of the language stream. Therefore, our work highlights effects of world knowledge on the production of referential expressions, and language production more broadly.

We have furthermore proposed an anaphoric reference choice model, which determines the probability of choosing distinct referring expressions to describe particular types of events. The probabilistic model augments the RSA framework with a temporal event prediction model. It essentially determines the utility of different utterances by assessing the probability of the listener arriving at the intended interpretation, given the utterance and prior world knowledge about event progressions. Thus, we have modelled influences of both linguistic information and world knowledge on interpretation and production processes.

Our probabilistic anaphoric reference choice model allows the choice of suboptimal forms of reference. Furthermore, speakers may unintentionally pick less than optimal forms of reference if the speaker and listener do not share the same event knowledge. Finally, a higher production cost of more precise utterances may favor shorter but less reliable descriptions. In all of these situations, the listener may face a low overall predictive consistency of the message that she believes was transmitted. If the message clashes with the listener's event knowledge, she may solve this inconsistency problem in one of the two ways: either revise her prior beliefs or reinterpret the linguistic evidence. Our model assumes that depending on the phonological salience, the types of reference differ in the level of signal certainty. For the interpretation process this means that less reliable signals are more likely to be reinterpreted. This proposed possibility of a revision by the listener views utterance interpretation an iterative process, seeking a consistent interpretation attractor (Butz, 2017). In the process of interpretation, the listener strives to reach overall consistency given the perceived message, possibly taking multiple revision steps to reach this goal.

References

Almor, A. (1999). Noun-phrase anaphora and focus: The informational load hypothesis. *Psychological review 106*(4), 748.

Baldwin, D. A. and J. E. Kosie (2021). How does the mind render streaming experience as events? *Topics in Cognitive Science 13*(1), 79–105.

Barton, E. L. (1998). The grammar of telegraphic structures: sentential and nonsentential derivation. *Journal of English Linguistics 26*(1), 37–67.

Bunger, A., A. Papafragou, and J. C. Trueswell (2013). Event structure influ-

ences language production: Evidence from structural priming in motion event description. *Journal of Memory and Language 69*(3), 299–323.

Butz, M. V. (2016). Toward a unified sub-symbolic computational theory of cognition. *Frontiers in psychology 7*, 925.

Butz, M. V. (2017). Which structures are out there. In T. K. Metzinger and W. Wiese (Eds.), *Philosophy and Predictive Processing*. MIND Group.

Butz, M. V., A. Achimova, D. Bilkey, and A. Knott (2021). Event-predictive cognition: A root for conceptual human thought. *Topics in Cognitive Science 13*, 10–24.

Butz, M. V. and E. F. Kutter (2017). *How the mind comes into being: introducing cognitive science from a functional and computational perspective* (First edition ed.). Oxford University Press.

Christensen, R. H. B. (2019). ordinal—regression models for ordinal data. R package version 2019.12-10. https://CRAN.R-project.org/package=ordinal.

Crawley, R. A., R. J. Stevenson, and D. Kleinman (1990). The use of heuristic strategies in the interpretation of pronouns. *Journal of Psycholinguistic Research 19*(4), 245–264.

Degen, J. and J. Tonhauser (2021, 09). Prior Beliefs Modulate Projection. *Open Mind 5*, 59–70.

Elman, J. L., A. Kehler, and H. Rohde (2006). Event structure and discourse coherence biases in pronoun interpretation. In *Proceedings of the annual meeting of the Cognitive Science Society*, Volume 28.

Elman, J. L. and K. McRae (2019). A model of event knowledge. *Psychological Review 126*(2), 252–291.

Frank, M. C. and N. D. Goodman (2012). Predicting Pragmatic Reasoning in Language Games. *Science 336*, 998–998.

Franke, M. (2009). *Signal to act: Game theory in pragmatics*, Volume DS-2009-11 of *ILLC dissertation series*. Amsterdam: Institute for Logic, Language and Computation and Universiteit van Amsterdam.

Franklin, N. T., K. A. Norman, C. Ranganath, J. M. Zacks, and S. J. Gershman (2020). Structured event memory: A neuro-symbolic model of event cognition. *Psychological Review 127*(3), 327–361.

Fukumura, K., R. P. G. van Gompel, T. Harley, and M. J. Pickering (2011). How does similarity-based interference affect the choice of referring expression? *Journal of Memory and Language 65*(3), 331–344.

Gibson, E., L. Bergen, and S. T. Piantadosi (2013). Rational integration of noisy evidence and prior semantic expectations in sentence interpretation. *Proceed-*

ings of the National Academy of Sciences 110(20), 8051–8056.

Goodman, N. D. and M. C. Frank (2016). Pragmatic Language Interpretation as Probabilistic Inference. *Trends in Cognitive Sciences 20*(11), 818–829.

Gordon, P. C., B. J. Grosz, and L. A. Gilliom (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science 17*(3), 311–347.

Gordon, P. C. and R. Hendrick (1998). The representation and processing of coreference in discourse. *Cognitive science 22*(4), 389–424.

Grice, P. (1989). *Studying in the Way of Words*. Cambridge, MA: Harvard University Press.

Haegeman, L. (2013). The syntax of registers: Diary subject omission and the privilege of the root. *Lingua 130*, 88–110.

Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology 61*(1), 23–62.

Kehler, A., L. Kertz, H. Rohde, and J. L. Elman (2008). Coherence and coreference revisited. *Journal of Semantics 25*(1), 1–44.

Kehler, A. and H. Rohde (2013). A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. *Theoretical Linguistics 39*(1-2), 1–37.

Kehler, A. and H. Rohde (2019). Prominence and coherence in a bayesian theory of pronoun interpretation. *Journal of Pragmatics 154*, 63–78.

Kuperberg, G. R. (2021). Tea with milk? a hierarchical generative framework of sequential event comprehension. *Topics in Cognitive Science 13*(1), 256–298.

Kurumada, C. and T. F. Jaeger (2015). Communicative efficiency in language production: Optional case-marking in Japanese. *Journal of Memory and Language 83*, 152–178.

Lam, T. Q. and D. G. Watson (2010). Repetition is easy: Why repeated referents have reduced prominence. *Memory & cognition 38*(8), 1137–1146.

Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. MIT press.

Levy, R., K. Bicknell, T. Slattery, and K. Rayner (2009). Eye movement evidence that readers maintain and act on uncertainty about past linguistic input. *Proceedings of the National Academy of Sciences 106*(50), 21086–21090.

Radvansky, G. A. and J. M. Zacks (2014). *Event cognition*. Oxford University Press.

Rosa, E. C. and J. E. Arnold (2017). Predictability affects production: Thematic roles can affect reference form selection. *Journal of Memory and Language 94*, 43–60.

Schrodt, F., J. Kneissler, S. Ehrenfeld, and M. V. Butz (2017b). Mario becomes cognitive. *Topics in Cognitive Science 9*(2), 343–373.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal 27*(3), 379–423.

Stegemann-Philipps, C., M. V. Butz, S. Winkler, and A. Achimova (2021). Speakers use more informative referring expressions to describe surprising events. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Volume 43, pp. 2897–2903.

Tily, H. and S. Piantadosi (2009). Refer efficiently: Use less informative expressions for more predictable meanings. In *Proceedings of the Workshop on the Production of Referring Expressions: Bridging the Gap between Computational and Empirical Approaches to Reference*.

Williams, E. (1997). Blocking and anaphora. *Linguistic Inquiry 28*(4), 577–628.

Zacks, J. M., N. K. Speer, K. M. Swallow, T. S. Braver, and J. R. Reynolds (2007). Event perception: A mind/brain perspective. *Psychological bulletin 133*(2), 273–293.

Zacks, J. M. and B. Tversky (2001). Event structure in perception and conception. *Psychological Bulletin 127*(1), 3–21.