Causal Priming: How a Language Production Mechanism Guides Representation

Caitlin M. Fausey (cmfausey@psych.stanford.edu)[a]
Neal Snider (snider@stanford.edu)[b]
Lera Boroditsky (lera@psych.stanford.edu)[a]

[a]Psychology Department, Stanford University, Stanford, CA 94305 USA
[b]Linguistics Department, Stanford University, Stanford, CA 94305 USA

Abstract

How do people talk about causal events? One constraint on causal event descriptions may be the local linguistic environment in which the description occurs. In this paper, we report results from one corpus study and one experimental study that examined the role of priming in the production of agentive language such as "*She broke the vase*" and non-agentive language such as "*The vase broke*". In both natural language use as well as in a more constrained experiment, English speakers were more likely to describe a causal event using the same kind of expression as had been used in a previous causal event description than to change the form of their description. Given evidence that people remember and reason about causal agents differently depending on whether they comprehend agentive or non-agentive language (Fausey & Boroditsky, 2007; 2008), understanding the language production mechanisms that create agentive and non-agentive linguistic environments is an important piece of a broader puzzle about how people integrate linguistic and non-linguistic sources of information in the representation and processing of causal events.

Keywords: Priming; causal language; causal explanation

Introduction

Imagine a woman tidying up rooms in her house, listening on the phone to a friend who is recounting the morning events with her toddler: "*First the cheerios stuck to his face, then the bib slipped off, and then the spoon dropped on the floor*". Upon entering her own daughter's bedroom, the woman sees a broken vase on the floor. Does she exclaim "*The vase broke!*" or "*She broke the vase!*"? Would her reaction be different if her friend had just told her, "*First, he stuck the cheerios to his face, then he slipped off the bib, and then he dropped the spoon on the floor*"?

Given the ubiquity of causal events in people's lives, and variation in their possible description, examining causal language use is likely to provide interesting insights about language production and comprehension, and also the relations between event construal and language. In this paper, we test whether local linguistic context influences how English speakers describe causal events.

Why examine causal language production? This question about language production is motivated by findings that the form of a causal event description impacts other inferences about the event. For example, event descriptions influence whether people more readily imagine causal antecedents or consequents (e.g., Majid, Sanford & Pickering, 2007), infer one or two events in a causal chain (e.g., Wolff, 2003), or attribute outcomes to the described agent or patient (e.g., Semin & Marsman, 1994). Recent findings also suggest that agentive event descriptions like "*She broke the vase*" lead people to punish agents more harshly, and to better remember agents, than do non-agentive descriptions like "*The vase broke*" (Fausey & Boroditsky, in preparation; 2008). Why do different event descriptions lead to different reasoning and memory for causal events?

The effects of causal language on causal event construal are presumably based in people's experience learning and using language across various non-linguistic contexts. Beyond this broad statement, however, we understand very little about the mechanisms that lead to the interesting language comprehension effects. For a more detailed account, we may need to better understand how causal language is produced in a range of non-linguistic contexts. Examining production may help us to understand the contexts in which agentive and non-agentive expressions are likely to be used and when they are especially likely to influence causal event construal. One relevant production mechanism may be linguistic priming.

Priming. Previous research has suggested that the form of an utterance is influenced by the form a previous utterance. For example, people are more likely to describe an event using the passive voice if they have recently heard a passive description than if they have recently heard an active description (e.g., Bock, 1986; Sankoff & Laberge, 1978). We purposely use the phrase "linguistic" priming in this paper because we remain agnostic about the nature of the priming effects that we report (see Pickering & Ferreira, 2008, for one review of the many sources of priming effects). Previous research has suggested that the production of active versus passive transitives (e.g., Bock, 1986; Sankoff & Laberge, 1978), prepositional versus object datives (e.g., Bock, 1986) and relative clause attachment (e.g., Scheepers, 2003) may be primed. Could the agentive and non-agentive language use that impacts attribution and memory for causal agents also be primed?

Priming causal event descriptions in natural language and experiments. One may examine linguistic priming from at least two perspectives. First, it is interesting to consider a particular kind of language use – such as causal verbs that may appear in both agentive and non-agentive expressions – and analyze patterns in language use *per se*. Second, it is interesting to consider the perspective that people might take on causal events – such as inferring causal agents or not – and use language production as one behavioral indicator of a

particular perspective. In this paper, we examine priming in both ways, using a natural language corpus to address the first question and a more constrained experiment to address the second question.

Both corpus analyses and experiments have advantages and disadvantages as paradigms for examining language production. Corpus analyses are more ecologically valid and also may include more tokens per speaker than constrained experiments. With respect to priming, it may be easier to find evidence of priming in natural language use than in experimental language use because of many levels of context that could mutually reinforce each other to lead to priming (e.g., Pickering & Garrod, 2004). Though experimental contexts are often limited to few items and utterances per participant and are rarely fully conversational, it is possible to control the non-linguistic (and linguistic) context of experiments and therefore they are amenable to testing more precise mechanistic hypotheses. Converging evidence from both natural and experimental language use, of course, would most strongly provide evidence for the behavior of interest.

Therefore, we did both a corpus study and an experiment. Both are analyses of fairly unconstrained language use. The corpus study used natural dialogue that was constrained only by the general topic of conversation. The experiment merely required participants to describe a set of pictures however they saw fit.

By examining whether agentive and non-agentive expressions may be primed, we broaden our understanding of what kinds of expressions are susceptible to local linguistic context effects and also reveal a potential mechanism for generating descriptions that may influence one's own and others' subsequent reasoning. Language production is both a reflection of the current event representation as well as influential in continuing to construct that representation and influence further processing. Can linguistic context influence the production of agentive and non-agentive language?

Study 1: Causal language in the wild

This study investigates whether causal language is primed in the course of natural language use. The data come from a corpus of conversations. Evidence of priming in natural language data would suggest that representations involved in processing causal events might persist over time.

Data

The corpus used in this study was the manually parsed Penn Treebank (release 3, approximately, 800,000 words, see Marcus et al., 1999) portion of the Switchboard corpus (Godfrey et al., 1992), a syntactically-annotated corpus of conversational speech. All intransitive and monotransitive forms of 24 alternating agentive/non-agentive verbs were automatically extracted using *tgrep2*[1]. These 24 verbs were chosen because they were found to alternate in a cross-linguistic study of agentive/non-agentive verbs by Haspelmath (1993). Haspelmath's study actually examined 31 verbs, but only 24 of theses are found in the Treebank Switchboard. The 24 verbs in the study are: *begin, boil, break, burn, change, close, connect, develop, dry, fill, finish, freeze, gather, improve, melt, open, rock, roll, sink, split, spread, stop, turn,* and *wake*, with *change* occurring the most often in the dataset (266 times) and *rock* occurring the least often (2 times), with a mean of 37.2 occurrences per verb. There are a total of 889 tokens in the dataset. 72 of these tokens were passives and were excluded from analyses.

Each token was automatically coded as agentive if it occurred monotransitively in active voice (where the agent of the action was explicitly mentioned), and automatically coded as non-agentive if it occurred intransitively (with no agent explicitly mentioned). Next, the form of the prime construction (agentive or non-agentive) was added to the database. The prime is defined as the nearest previous usage of one of the 24 alternating verbs. Primes that occurred in the same utterance as the target were excluded (n = 30), as these are more likely to be speech errors or stylistic repetitions. Prime-target distances ranged from 1 turn to 332 turns, with a mean of 60.78.

Method

To test the effect of priming in choice of agentive versus non-agentive descriptions, the data set was analyzed with mixed logit models (Bates & Sarkar, 2006; Breslow & Clayton, 1993). Mixed logit models can be thought of as an extension of logistic regression that include modeling of random effects. Inclusion of random effects, such as speaker or participant, is necessary to generalize beyond the participants in the current data set (Clark, 1973).

The dependent variable was structure choice, agentive or non-agentive, with agentive coded as the positive choice. The primary independent variable of interest was the prime construction, which had three levels: agentive, non-agentive, and no-prime. The no-prime cases are when an agentive or non-agentive verb occurred at the beginning of the conversation, so the prime is unknown. The no-prime tokens act as a baseline, allowing one to determine whether there are independent effects of both agentive and non-agentive primes relative to a baseline level of agentive versus non-agentive use.

In addition to the fixed effect of the prime construction, several other factors were added to the model. These factors are commonly examined in priming research: conversation topic, agentive bias of individual verbs, prime-target distance, whether or not the prime and target were produced by the same speaker, and whether or not the prime and target expression shared the same verb.

Each conversation in the Switchboard corpus had a pre-defined topic, and we observed in our exploration that some topics were more likely to contain agentive descriptions. Thus, we divided the topics into three roughly equal-sized categories (agentive-bias, no-bias, and non-agentive-bias) based on the proportion of agentive descriptions used in each topic. We also observed that some verbs were used

[1] tedlab.mit.edu/~dr/Tgrep2/

agentively more often than others, so we added a factor to account for the agentive bias of the target verb.

We also added several interaction terms. We added an interaction between priming and the topic bias, to investigate the possibility that priming may be weaker when the conversation has a general bias towards or against agentive descriptions. We analyzed whether priming varied depending on the distance in utterances between the prime and target to investigate possible decay of priming. We also added a factor indicating whether the prime and target were produced by the same speaker, because of some suggestions that priming may be stronger in production than in comprehension (e.g., Bock & Griffin, 2000). We also examined the strength of priming when the prime and target verbs were the same and when they were different, given previous results that suggest stronger priming when the prime and target constructions share the same verb (e.g., Pickering & Branigan, 1998).

A random effect of speaker was added to ensure our result generalizes beyond the speakers in this sample, allowing each speaker a different base rate of agentive language use.

Results

Figure 1 shows that in the absence of a prime, speakers were about as likely to use an agentive expression (.52 of these tokens) as a non-agentive expression (.48). Following agentive primes, speakers were more likely to describe events using agentive language (.64) and following non-agentive primes people were more likely to use non-agentive language (.62). Regression analyses permit a more detailed examination of these data.

In the regression analysis, we report the coefficient for the independent variable and its levels of significance. Coefficients in logistic regression models are given in log-odds (the space in which logistic models are fitted to the data). For categorical factors, significant positive coefficients mean that the positive level (agentive in this case) is more likely in the tested level of the variable than in the other level. For example, if the coefficient of prime=agentive is positive, then having a prime that is agentive makes an agentive form more likely in the target description. Negative coefficients mean the opposite. Below, we also report the difference in odds between conditions (as the name suggests, odds are simply $e^{log-odds}$). Odds range from 0 (for proportions of 0) to positive infinity (for proportions of 1), with proportions of 0.5 corresponding to odds of 1. Odds are a multiplicative scale, so we talk about an x-fold increase or decrease in odds between conditions.

The coefficients in log-odds and standard errors associated with the factors are given in the second and third column of Table 1. The corresponding odds coefficients are given the fourth columns. The fifth and sixth columns summarize the Wald's Z statistic, which tests whether the coefficients are significantly different from zero. Finally, the last two columns give the χ^2 over the change in data likelihood ($\Delta_x(\Lambda)$) associated with the removal of the predictor (x) from the final model. The latter test is more robust against collinearity in the model (Agresti, 2002). The χ^2 value, which literally corresponds to the difference in the model's data likelihood without the predictor, can be seen as a measure of the predictor's importance in the model.

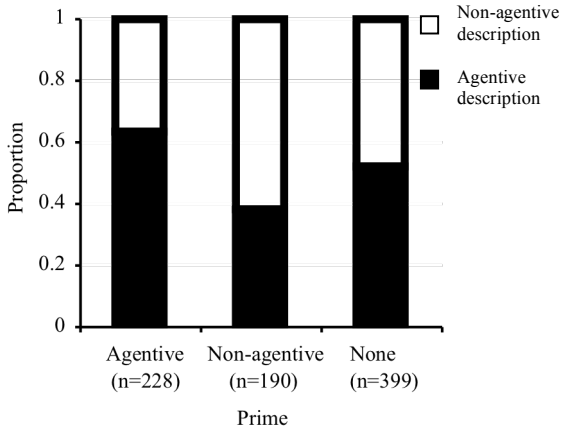

Figure 1: Causal event descriptions in natural language.

The Wald-test is included because it implicitly tests the directionality of the effect (unlike the χ^2 over the change in data likelihood). In the case of factors with multiple levels, Table 1 lists the coefficients for the separate levels of significant factors only.

Predictor (independent variable)	***Parameter estimates*** Log-odds	S.E.	Odds	***Wald's test*** *Z*	*p*	$\Delta_x(\Lambda)$***-test*** χ^2	*p*
Intercept	-0.12	0.15	0.88	.87	<.4		
Topic=Agentive-biased	0.69	0.19	1.99	3.6	<.001	55.2	<.001
Topic=Nonagentive-biased	-0.71	0.19	0.49	-3.7	<.001		
Verb bias	1.60	0.25	4.95	6.5	<.001	42.8	<.001
Prime construction						0.82	<.7
Same verb x Agentive prime	0.88	0.28	2.4	3.2	<.002	17.7	<.001
Same verb x Nonagentive prime	-1.33	0.32	0.26	-4.1	<.001		
Prime x distance						0.67	<.5
Prime x same speaker						4.47	<.15
Prime x Topic						0.5	<1

Table 1. Mixed logit model results of corpus analysis.

The first thing to note about the results from the regression analysis is that priming varies depending on whether or not the same verb is used in the prime and target description. When the same verb is used, there is an effect of the prime description. People are 2.4 times as likely to describe an event using agentive language when they have recently heard a similarly agentive event description, compared to when there was no prime. Further, they are 3.8 (=1/0.26) times less likely to describe an event agentively if they have just heard a non-agentive description. When the prime and target verbs do not match, there is no significant priming effect, as revealed by the analysis in which the prime construction factor does not significantly improve the model when the interaction between prime and verb identity has been controlled. There were no other interactions between priming and topic, distance, or speaker.

Given previous corpus and experimental results that priming does not require identical verbs in the prime and target description, the lack of a main effect of priming may seem surprising. This finding may be due to lack of power from the relatively small data set compared to other corpus studies of priming (e.g., Szmrecsanyi (2005) and Jaeger and Snider (2007) had at least 1000 tokens in all studies). In order to make sure that the null effects of the other priming interactions were not simply due to the generally weaker priming on non-verb-matching tokens, a separate model was run using only those tokens where the prime and target verbs matched. In these 168 tokens, the main effect of the prime construction is still highly significant ($\chi^2 = 26.9$, $p <$. 001), but the interactions with verb bias, topic, distance, and speaker identity are not (p's > .1).

Discussion

This study provides evidence that people are more likely to use agentive language in natural conversation if they have previously heard an agentive description, and to use non-agentive language after a non-agentive description, at least when the prime and target verb are the same. We find these effects even when controlling for effects of individual speakers, verbs, and the conversation topic. These results lend support to the notion that priming is a psychologically real mechanism in actual language use and could contribute to reinforcing particular description patterns both locally as well as over a lifetime of language use.

Knowing the conditions under which repetition of expressions occurs in natural language use is valuable information, but because structures may be repeated for a variety of reasons it is challenging to interpret this repetition. To further examine when and why priming of agentive and non-agentive descriptions might occur, perhaps even when the prime and target verbs are different, we also conducted a more constrained experiment.

Study 2: Causal language in the lab

The way that people describe causal events during natural language use appears to depend on the local linguistic environment: The results of our corpus analysis suggest that a particular causal event description is more likely to match the form of the previous causal event description than to change forms, at least when the prime and target descriptions share a verb. In natural language corpora, however, each target description can only be primed by either an agentive prime description or a non-agentive description. That is, it is impossible to manipulate the prime status for any given target description. Because different causal events were described following agentive and non-agentive language, patterns revealed by our corpus analysis provide only weak evidence that linguistic context may generally influence how people frame subsequent events. Understanding whether this is possible will be a key component of a broader account about how linguistic context may be able to bias descriptions and/or construals of events. Could descriptions of the same causal event be changed by linguistic primes? An experimental paradigm constrains the environment of the language user so that answering this question is possible: participants in our second study described the exact same event following minimally-paired linguistic primes.

In this study, participants viewed a pair of pictures depicting the beginning and end states of a causal event and then described the event. Evidence about only the beginning and end states of an event leaves ample room for observers to infer additional details about the event. For example, they might infer that an agent was involved or they might not. Many real world scenarios are similarly ambiguous about the causal status of potential agents (consider our initial story, in which the mother knew only the beginning and end state of the vase). Some participants in our experiments described events in the absence of a linguistic prime and other participants described events after reading either an agentive or a non-agentive linguistic prime. We hypothesized that people's descriptions of the causal event would match the form of an unrelated prime sentence that they had read prior to viewing the event depiction. Assuming that people's language production is related to how they construe the beginning and end state pictures of the event, this paradigm may reveal one way in which local linguistic context shapes causal event construal.

Our corpus analyses suggest that an implicit priming mechanism may operate during natural language use of agentive and non-agentive expressions. Can the production of these expressions be linguistically primed even when participants describe the very same causal event?

Participants

338 students at the University of California, Merced participated for course credit.

Materials

Linguistic primes. Participants in the prime conditions read one prime sentence, drawn from a set of eight sentences. The full set consisted of an agentive and non-agentive description of four events (see Table 2).

Visually-depicted events. Participants described a pair of pictures depicting either a paint-splattering event or a vase-breaking event. Each visual depiction consisted of a beginning-state frame and an end-state frame (see Figure 2).

Design

After reading an agentive prime sentence, a non-agentive prime sentence, or no prime sentence, participants viewed a pair of pictures and described the visually-depicted event. Primes and pictured events were fully crossed. The dependent measure was whether people described the pictured event using agentive or non-agentive language.

Table 2. Linguistic prime stimuli.

Agentive primes	**Non-agentive primes**
He popped the balloon.	The balloon popped.
He opened the umbrella.	The umbrella opened.
He unfastened the necklace.	The necklace unfastened.
He blew out the match.	The match blew out.

Figure 2: Pictured events.

Procedure

Participants in the prime conditions completed a two-sided survey that was presented among several other unrelated surveys. On the front side of the page, participants read one sentence and were asked to "please continue the story for another sentence or two" on blank lines that appeared below the prime sentence. This encouraged them to actually process the prime sentence. On the back side of the page, participants saw the beginning and end of a causal event and were asked to "please describe this event".

Participants in the no-prime condition completed just the target event description side of the survey.

Results

Coding. Each event description was coded as agentive or non-agentive. Transitive sentences, both active (71%) and passive (29%), were coded as agentive[2]. Intransitive sentences were coded as non-agentive (see Table 3).

Overall analyses. People described target events differently depending on the linguistic prime (see Figure 3). Without any prime, people were about as likely to use agentive language (.48 of these tokens) as non-agentive language. Following agentive primes, people were more likely to describe an event using agentive language (.65) than non-agentive language. Following non-agentive primes, people were more likely to describe an event using non-agentive language (.59) than agentive language. Chi-square analyses confirmed the different distributions of target event descriptions following each kind of prime, $\chi^2(1) = 9.79$, $p = .002$. Unprimed descriptions differed from agentive-primed descriptions, $\chi^2(1) = 6.38$, $p = .012$, and not from non-agentive primed descriptions, $\chi^2(1) = .27$, $p = .60$.

Item-specific analyses. English speakers had different overall biases with respect to agentivity in describing the paint-splattering and the vase-breaking events.

Table 3. Example causal event descriptions.

Agentive responses
Somebody broke the vase.
Someone took the paint and splattered it on the wall.
Non-agentive responses
The pretty antique vase broke and shattered into pieces.
The paint was in the buckets then it spilled onto the wall.

As shown in panels (b) and (c) of Figure 3, English speakers preferred to describe the paint-splattering event agentively and the vase-breaking event non-agentively. The priming effect, however, was in the same direction for both events, marginally for paint-splattering ($\chi^2(1) = 2.90$, $p = .089$) and significantly for vase-breaking ($\chi^2(1) = 6.71$, $p = .01$). The four prime items influenced causal event descriptions in similar ways.[3]

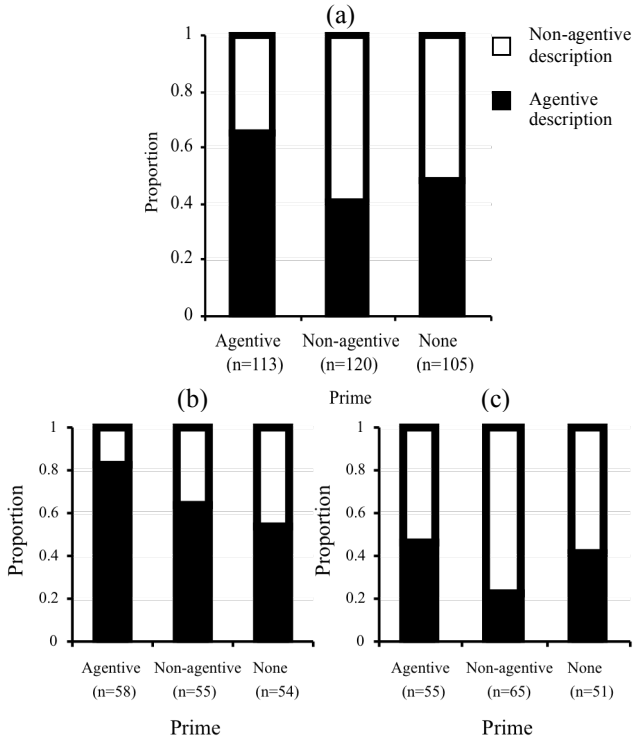

Figure 3. Primed descriptions of causal events.
(a) overall, (b) paint-splattering,(c) vase-breaking.

Discussion

In a controlled experimental setting, people were more likely to talk about the agent of a causal event in the presence of an unrelated event description that was agentive rather than non-agentive. Descriptions of the very same visual stimulus depended on the local linguistic context.

General Discussion

Our studies provide evidence that agentive and non-agentive causal event descriptions can be primed. In both experimental settings and in natural dialogue, people are more likely to mention an agent when describing a causal event when they have recently encountered an agentive event description than when they have recently encountered a non-agentive event description. This result is especially compelling because it comes from converging evidence in both conversational and controlled experimental settings.

One implication of these findings is the need to consider language production in theories about the relationship

[2] We considered several schemes for classifying passive descriptions and the choice of scheme did not change the results.

[3] The only exception was the umbrella prime-paint target combination in which 3 of 28 people used non-agentive language.

between language use and causal reasoning. Priming clearly influences linguistic choices in everyday conversation using agentive and non-agentive event descriptions and other research suggests that these expressions impact further reasoning about causal events (e.g., Fausey & Boroditsky, 2007; 2008; in prep). Thus, priming may be a mechanism by which linguistic experience influences reasoning more generally. Future experiments will combine the to-date separate approaches of production and comprehension in order to more directly understand the relationship between linguistic priming and causal reasoning.

This work also informs priming research in general, as we have shown that intransitive expressions prime under certain circumstances. The nature of the non-linguistic context in our experiment may have facilitated some non-agentive language use: People saw an event depiction without an agent and therefore could felicitously be described by an intransitive expression. Interestingly, participants could infer a causal agent of the event, and were more likely to do so in the presence of an unrelated agentive expression. In future experiments, we will examine a wider range of stimuli and in future corpus studies, we will explore how intransitives of non-causal verbs may affect causal language use.

What influences the description of causal events? Our results suggest that whether someone describes a causal event using agentive language like "*She broke the vase"* or non-agentive language like "*The vase broke"* depends in part on the local linguistic context.

Acknowledgments

We thank SPLAT! attendees for helpful comments, and T. Matlock and J. Matthews for help with data collection.

References

Agresti, A. (2002). *Categorical Data Analysis*. New Jersey: Wiley & Sons.

Bates, D., & Sarkar, D. (2006). lme4: Linear mixed-effects models using S4 classes. R package version 0.995-2.

Bock, J.K. (1986). Syntactic persistence in language production. *Cognitive Psychology, 18*, 355-387.

Bock, J.K., & Griffin, Z.M. (2000). The persistence of structural priming: Transient activation or implicit learning? *Journal of Experimental Psychology: General, 129*, 177-192.

Breslow, N.E., & Clayton, D.G. (1993). Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association 88*, 9–25.

Clark, H.H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior, 12*, 335–359.

Fausey, C.M., & Boroditsky, L. (in prep). Speaking of accidents: "He did it" invites more punishment than "It happened".

Fausey, C.M., & Boroditsky, L. (2008). English and Spanish speakers remember causal agents differently. Proceedings of the *30th annual meeting of the Cognitive Science Society*, Washington, D.C.

Fausey, C.M., & Boroditsky, L. (2007). Language changes causal attributions about agents and objects. Poster presented at the *29th annual meeting of the Cognitive Science Society*, Nashville, TN.

Godfrey, J., Holliman, E., & McDaniel, J. (1992). Switchboard: Telephone speech corpus for research and development. In *IEEE ICASSP*. 517-520.

Gries, S.T. (2005). Syntactic priming: A Corpus-based Approach. *Journal of Psycholinguistic Research, 34*, 365–399.

Jaeger, F., & Snider, N. (2007). Implicit learning and syntactic persistence: Surprisal and cumulativity. *Rochester Working Papers in the Language Sciences*, *3(1)*. Rochester, NY.

Haspelmath, M. (1993). More on the typology of inchoative/causative verb alternations. In *Causatives and transitivity* (Ed. B. Comrie & M. Polinsky), John Benjamins: Amsterdam.

Majid, A., Sanford, A. J. & Pickering, M. J. (2007). The linguistic description of minimal social scenarios affects the extent of causal inference-making. *Journal of Experimental Social Psychology, 43*, 918-932.

Marcus, M.P., Santorini, B., & Marcinkiewicz, M.A. (1994). Building a large annotated corpus of English: The Penn treebank. *Computational Linguistics, 19*, 313–330.

Pickering, M.J., & Branigan, H.P. (1998). The representation of verbs: Evidence from syntactic priming in language production. *Journal of Memory and Language, 39*, 633-651.

Pickering, M.J., & Ferreira, V.S. (2008). Structural priming: A critical review. *Psychological Bulletin, 134,* 427-459.

Pickering, M., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences, 27*, 169-226.

Sankoff, D., & Laberge, S. (1978). Statistical dependence among successive occurrences of a variable in discourse. In D. Sankoff and S. Laberge (Eds.): *Linguistic Variation: Methods and Models*. Academic Press, NY.

Scheepers, C. (2003). Syntactic priming of relative clause attachments: Persistence of structural configuration in sentence production. *Cognition*, *89*, 179—205.

Semin, G.R., & Marsman, G. (1994). "Multiple inference-inviting properties" of interpersonal verbs: Event instigation, dispositional inference, and implicit causality. *Journal of Personality and Social Psychology, 67,* 836-849.

Szmrecsanyi, B. (2005). Language users as creatures of habit: a corpus-based analysis of persistence in spoken English. *Corpus Linguistics and Linguistic Theory, 1*, 113–149.

Wolff, P. (2003). Direct causation in the linguistic coding and individuation of causal events. *Cognition, 88*, 1-48.