Running: Structural Priming as Implicit Learning

Structural Priming as Implicit Learning: Cumulative Priming Effects and Individual Differences

Michael P. Kaschak, Timothy J. Kutta, and John L. Jones

Florida State University

[in press, *Psychonomic Bulletin and Review*]

Word Counts:

Abstract:	141
Text, Notes, and References:	4000

Correspondence to:

Michael Kaschak

Department of Psychology

Florida State University

Tallahassee, FL 32306

Phone: 850-644-9363

Email: kaschak@psy.fsu.edu

Abstract

We explored the claim that structural priming is a case of implicit learning within the language production system. The experiment began with a baseline phase in which we assessed participants' rate of production for the double object and prepositional object constructions. Then, participants were biased toward the production of either the double object or prepositional object construction. Finally, we again assessed participants' rates of production for the target constructions. Consistent with claims that structural priming is a case of implicit learning, we found that biasing participants toward the prepositional object construction produced stronger cumulative priming effects than biasing participants toward the double object construction. We also found that individual differences in implicit learning were marginally correlated with overall rates of production for the double object construction. Participants who scored better on the learning task tended to produce fewer double object constructions.

Structural priming refers to the tendency for individuals to repeat syntactic structures across utterances (Bock, 1986). For example, a person who produced a double object construction (DO; *Meghan gave Michael a toy*) is more likely to produce another DO when describing a transfer situation (*John sent Tim the files*) than to produce a prepositional object (PO) construction describing the same event (*John sent the files to Tim*). Structural priming has been widely studied, as the presence (or absence) of priming between prime and target sentences is revealing about linguistic representations. As one illustration, researchers have used structural priming to explore the relationship between verbs and syntactic constructions (e.g., Kaschak & Borreggine, 2008; Coyle & Kaschak, 2008; Bernolet & Hartsuiker, 2010).

Structural priming is used to explore linguistic representations, but can be studied as a phenomenon unto itself to understand the mechanisms through which the language system adapts its behavior (e.g., Bock & Griffin, 2000). Structural priming has been explained as transient spreading activation (e.g., Pickering & Branigan, 1998), implicit learning (e.g., Chang, Dell & Bock, 2006), or a combination of both mechanisms (e.g., Reitter, Keller & Moore, 2011). There is evidence for both mechanisms. Transient increases in priming arising when lexical items are repeated across sentences suggests an activation-type mechanism (e.g., Hartsuiker et al., 2008), and long-lasting priming (e.g., Kaschak, Kutta & Schatschneider, 2011) indicates a role for implicit learning.

Our study explores two aspects of the claim that long-range structural priming effects are examples of implicit learning. First, we examine the *inverse frequency effect*, or the finding that lower-frequency constructions produce stronger priming than higher-frequency constructions (e.g., Scheepers, 2003; Jaeger & Snider, 2008; Hartsuiker & Westenberg, 2000; Ferreira, 2003). The inverse frequency effect arises from long-range accumulations of production experience, and is predicted by implicit learning accounts of structural priming (e.g., Chang et al., 2006). Despite being predicted by implicit learning accounts, few demonstrations of inverse frequency effects have been reported. Reitter et al. (2011) examined the inverse frequency effect in the context of cumulative structural priming (i.e., structural priming accumulating over many tokens of one construction; e.g., Kaschak et al., 2006). Using an ACT-R-based model of language production, they simulated cumulative priming for the DO and PO constructions over a range of 1 – 25 prime sentences. They assessed how much structural priming increases as the number of prime constructions produced prior to the target sentences increases. Their simulation generated an intriguing result: there was robust cumulative priming for the lower-frequency PO construction (priming increased as the number of primes produced increased), and weaker, non-significant priming for the higher-frequency DO construction.

Reitter et al.’s (2011) simulation suggests an interaction between the long-range frequency of constructions and the extent to which priming accumulates over shorter ranges of experience. If this prediction is confirmed empirically, it puts constraints on the mechanisms that can explain structural priming. Although inverse frequency effects have been reported in the literature (e.g., Scheepers, 2003), including cases where priming accumulates over several sentences (e.g., Hartsuiker & Westenberg, 2000; Kaschak, 2007), the previous studies were not designed so that structure-based differences in cumulative priming could be observed. The relevant studies either had multiple constructions produced throughout the experiment (thus not producing the sort of cumulative priming simulated by Reitter et al., 2011; e.g., Hartsuiker & Westenberg, 2000), or require cross-experiment comparisons to examine the strength of cumulative priming. Thus, additional work is needed to assess Reitter et al.’s (2011) prediction.

Our study employed a modification of Kaschak's (2007) cumulative priming paradigm to test whether Reitter et al.'s (2011) simulation result can be observed using written language production. Participants completed a series of *target stems* that allowed the production of either the DO or PO construction (e.g., *Meghan gave…*). This provided baseline information about use of the DO and PO constructions. Bock and Griffin (2000) suggest that American English speakers have a 2:1 bias toward the DO construction, and something close to these relative frequencies should be observed here. Next, cumulative priming was instantiated by presenting *prime stems* that disposed participants toward producing either the DO (*Meghan gave Michael...*) or PO (*Meghan gave the toy...*) construction. Participants saw either 100% DO primes, or 100% PO primes. Finally, we measured cumulative priming by presenting another set of target stems. Reitter et al.'s (2011) model predicts that DO-biased participants will show less change in production behavior between the first and last phases of the study than PO-biased participants.

Our second goal was to test the structural-priming-as-implicit-learning view by assessing how individual differences in implicit learning relate to cumulative structural priming. It had long been thought that there were no important individual differences in implicit learning (e.g., Reber, 1993), but subsequent work has shown that such differences do exist (e.g., Kaufman et al., 2010). Moreover, individual differences in implicit learning appear to be related to language comprehension skill (e.g., Conway et al, 2010; Misyak & Christiansen, in press). If cumulative structural priming is a case of implicit learning, the magnitude of the observed cumulative priming effects should be positively correlated with participants' implicit learning ability. Based on Reitter et al. (2011), this relationship may be particularly strong for participants biased toward the PO construction (as the DO construction may produce weak cumulative priming).

Method

Participants. 113 undergraduate psychology students participated for partial course credit. As in previous work (e.g., Kaschak, 2007), we ensured the integrity of the cumulative priming manipulation by eliminating data from participants who did not complete at least 80% of the prime stems as intended (e.g., participants biased toward the DO construction must complete > 80% of the prime stems as a DO). All participants met this requirement, producing the intended construction 97% of the time. Furthermore, because we were interested in changes in rates of DO and PO production across time, we excluded participants who did not produce DO or PO constructions on 50% (i.e., 3 out of 6) of target trials in both the Pre- and Post-bias phases of the experiment. Eleven participants were excluded on this criterion.

Materials. Twelve target stems were constructed (e.g., *Meghan gave…*). These could be completed as either DO or PO constructions. Target stems were arranged into a fixed order, and then split into two sets of 6 (A and B). Half of the participants saw set A in the Pre-bias phase, and set B in the Post-bias phase; the other half saw the sets in the opposite order. Fourteen pairs of prime stems were constructed. One pair member elicited the DO construction (*Meghan gave Michael…*) and the other elicited the PO construction (*Meghan gave the toy…*). We constructed 104 filler stems that could not easily be completed as DO or PO constructions. Fillers were presented throughout the experiment, with 4 fillers separating each prime or target stem.

Our measure of individual differences in implicit learning was the “Simon” task described by Conway et al. (2010). Following the Milton Bradley game, the computer screen had four colored squares (red, green, blue, yellow). On each trial, the squares lit up in sequence (e.g., red-green-red-yellow). Participants used labeled keys on the keyboard to reproduce the sequence. Sequences were scored as correct if participants reproduced the entire sequence correctly. The initial game trials trained participants on the probabilistic rules used to generate the sequences. Afterwards, participants saw sequences that followed the rules, and sequences that did not. Sequence length ranged from 4 to 8 responses. Our task employed the same rules, training and test sequences as those employed by Conway et al. (2010). For sequences produced correctly in the test phase, participants received one point for each response in the sequence (e.g., an 8-response sequence was worth 8 points). The measure of implicit learning (IL score) was generated by subtracting the number of points earned on rule-violating items from the number of points earned on rule-following trials. IL scores ranged from -8 (participant performed better on non-rule following items than on rule-following items) to 74 (mean=31.5, SD=15.1).

Procedure. Participants were randomly assigned to the DO or PO Bias condition, with half of the participants in each condition. Participants were told that they would see a series of sentence stems on the computer screen, and that they should complete each as a grammatical sentence. The experiment began with the Pre-bias phase, where participants completed 6 target stems. Next, came the Bias phase. Participants saw 14 prime stems that elicited either the DO or PO construction; participants were biased 100% toward one construction. Then, participants completed 6 more target stems (Post-bias phase). Finally, participants performed the Simon task.

Scoring. Prime and target stem completions were scored DO, PO, or “other” as described in Kaschak (2007).

Design and Analysis. Trials scored as “other” (PO bias=13%; DO bias=16%) were excluded from the analysis, creating a binary dependent variable (DO responses coded “1,” PO responses coded “0”). Mixed logit analysis was conducted to predict the log odds of producing a DO target completion. Analysis was done with the lme4 package (Bates, Maechler, & Bolker, 2011) of R (R development core team, 2011). All predictors in all analyses were centered. For every model reported here, collinearity was assessed using the collin.fnc() procedure in R. All obtained Kappa values were < 2.64. Thus, collinearity was not a concern. Models included participants and items as crossed random factors. Intercepts could vary across participants and items. We employed a model comparison approach to determine the best fitting model for our data. We first ran a model including Bias condition (DO bias=1, PO bias=0), Time (Pre-bias=0, Post-bias=1), and the Bias x Time interaction without random slopes, and the same model with the full complement of random slopes. Because the random slopes marginally improved model fit (p=.07), we tested whether each individual random slope improved model fit. Subsequent to this, we assessed whether the addition of IL score and its associated interactions improved model fit. The best fitting model included Bias condition, Time, and the Bias x Time interaction, with random slopes for participants on the Time variable.

Results

The best-fitting mixed logit model predicting the log odds of producing a DO target completion is presented in Table 1. The Time x Bias condition interaction was significant (p<.001). In a mixed logit model analyzing only DO-biased participants, the effect of Time was not significant (p=.30). Participants produced similar rates of DO responses in the Pre- and Post-bias phases (see Table 1). The effect of Time was significant (p<.001) in a model assessing PO-biased participants. Participants showed cumulative priming, producing fewer DOs in the Post-bias phase than in the Pre-bias phase. Adding IL score to the model in Table 1 showed that IL score was a marginally significant predictor (B=-.02, p=.06), and that it marginally improved model fit ($X^2(1)=3.52$, p=.07). As IL scores increased, the likelihood of producing a DO completion decreased. Interactions involving IL score did not improve model fit.

We observed an inverse frequency effect: the lower-frequency PO construction produced stronger cumulative priming than the higher-frequency DO construction. However, our results do not directly test Reitter et al.'s (2011) claim about changes in cumulative priming based on the current run of DO and PO constructions (more consecutive POs = stronger priming; less change in priming based on the number of consecutive DOs). To assess this hypothesis, we analyzed responses from the Post-bias phase, coding for Bias condition and IL score, plus the current run of consecutive DO or PO constructions (ignoring trials with "other" responses) produced prior to each target stem. This latter measure approximates Reitter et al.'s cumulative priming manipulation. Participants could produce DO or PO constructions at any time (even in the Bias phase), and thus the run of consecutive DOs or POs varied across trials. Current run was coded as a positive number when the preceding trial(s) used the construction toward which the participant was biased (called the *bias* construction). These runs could vary from 1 (n-1 trial used the bias construction, but n-2 trial did not) to 25 (on the final Post-bias phase trial, all 25 prior prime and target stem completions employed the bias construction). The current run was coded as a negative number when the preceding trial(s) used the construction opposite the one toward which the participant was biased. These runs could vary from -1 (trial n-1 used the non-bias construction, but the trial n-2 used the bias construction) to -25 (on the final Post-bias phase trial, all preceding trials were completed as the non-bias construction), though the presence of prime stems in the Bias phase made it unlikely that negative runs would exceed 5 non-bias constructions. Current run measures were similar across Bias conditions (DO bias: range=-5–25, mean=8.49, SD=8.9; PO bias: range =-5–22, mean=8.06, SD=8.9).

The best fitting model for this analysis is presented in Table 2. As before, the inclusion of IL score did not significantly improve model fit (*p*=.06). The critical result is a significant Bias condition x Current run interaction (*p*<.001). We explored the interaction by performing separate mixed logit analyses on the DO and PO Bias conditions. Beyond the Current run variable, we coded for the cumulative probability of producing a DO construction across the experiment (i.e., up to a given trial, what proportion of prime and target stems were completed as a DO?; see Jaeger & Snider, 2008). Cumulative probability is a proportion, where higher values are expected to predict higher rates of DO production. [To yield more interpretable coefficients, we converted this proportion to a percentage for the analysis. Both predictors were centered within each Bias condition prior to analysis]. We included Cumulative probability in this model to determine whether effects of Current run would be observed after accounting for participants' overall likelihood of producing the DO.

In the DO condition, Cumulative probability predicts the likelihood of producing a DO, but Current run does not (see Table 2; Figure 1, top). This finding supports our main analysis – experiencing longer runs of the DO construction does not significantly affect the likelihood of producing a DO completion. The PO condition shows a different pattern: both Cumulative probability and Current run predict the likelihood of a DO completion (Table 2; Figure 1, bottom). Higher cumulative probabilities lead to higher odds of producing a DO response, and longer runs of the PO construction lead to lower odds of producing a DO response.

One concern about our conclusions regarding the Current run variable is that these results may be constrained by the patterns in our main analysis: Current runs cannot show an effect when there is no cumulative priming on an experiment-wide level (DO condition), but can show an effect when there is cumulative priming (PO condition). Our data argue against this concern. The range, average length, and variance in runs are similar across conditions, as is the proportion of bias constructions produced (63% DOs in the DO condition; 55% POs in the PO condition). Based on the similarities seen across conditions, if the global patterns of performance were driving the effect of Current runs, we would expect a similar runs effect across conditions. This was not the case.

Discussion

We explored the view that cumulative structural priming reflects implicit learning in the language production system. Our data confirm one element of this claim: as predicted by Reitter et al. (2011), the main analysis showed cumulative priming for the (lower-frequency) PO construction, and weaker priming for the (higher-frequency) DO construction (i.e., an inverse frequency effect). Subsequent analysis supported this result, as PO-biased participants showed an effect of Current run, but DO-biased participants did not.

We have two additional comments on our data. First, the rate of DO completions in Pre- and Post-bias phases of the DO Bias condition (.61 and .63) are similar to the Post-bias phase rates of DO production in our other studies using written production (Kaschak, 2007: Exp. 1=.64, Exp. 2=.66; Kaschak et al., 2011: .66). Rates of DO production across these Post-bias phases, viewed in light of our Pre-bias data, suggest a re-interpretation of our earlier results, namely that the previous cumulative priming effects were largely driven by the PO bias condition. DO bias effects were probably small. Second, rates of DO production in these studies (and natural language use) suggest that DO production is not at ceiling. The lack of a Current run effect for this construction is therefore striking, as there is ample room for DO production to increase. Although striking, the weakness of the Current runs effect for the DO is predicted by at least one model (Reitter et al., 2011). The result is also consistent with the spirit of other implicit learning accounts of structural priming. Whether Current runs ultimately has no effect on DO production, or a weak one (note that data in the DO Bias condition hint at a weak effect of runs), it is clear that inverse frequency effects occur for alternations that do not have stark frequency differences between members. Models such as Reitter et al.'s (2011) appear to provide mechanisms for understanding such effects. However, in their current form, the mechanisms posited, and the modeling assumptions that are made, may not be optimally sensitive to some types of local regularities. For example, it is noteworthy that the strength of the cumulative priming generated by Reitter et al.'s (2011) model is weaker than that observed here and in other studies (a point raised in their paper).

Our results are consistent with studies showing that language comprehenders (e.g., Kaschak & Glenberg, 2004; Kaschak, 2006; Farmer et al., 2011; Wells et al., 2009) and producers (e.g., Kaschak, 2007; Jaeger & Snider, 2008) rapidly adapt to the probabilistic use of syntactic structures within their linguistic environment The present data provide an important qualification to demonstrations of flexibility in language use. Specifically, the difference in the effect of Current run in the DO and PO bias conditions suggests that not all experiences with language on a given time-scale count equally in modifying language processing (c.f., Jaeger & Snider, 2008; Kraljic et al., 2008). It will be important to determine whether our findings hold for other structural alternations. If so, interactions between long-range frequency and shorter-range cumulative priming will place constraints on the mechanisms posited to explain structural priming and other adaptations in language processing.

A second aim of our study was to determine whether the magnitude of cumulative priming is predicted by individual differences in implicit learning. We expected that a) stronger implicit learning performance would predict stronger cumulative priming, and b) this relationship would be stronger in the PO bias condition. These predictions were not confirmed. Although one interpretation of these findings is that implicit learning does not drive cumulative structural priming, this interpretation may be premature. Correlations between implicit learning measures are often low (e.g., Gebauer & Mackintosh, 2007; Misyak & Christiansen, in press), raising the possibility that different implicit learning measures might correlate with cumulative priming. Furthermore, there are distinctions between implicit-learning-based models of language processing that instantiate learning through either procedural (e.g., Chang et al., 2006) or declarative (e.g., Reitter et al., 2011) memory. If different learning mechanisms underlie Simon task performance and language production, it may explain our null result.

A final possibility that may be worth exploring is that individual differences in implicit learning relate to a more nuanced sensitivity to statistical structure in language than we assessed. We examined how global frequencies of use for the DO and PO constructions interact with local frequencies of use within the experiment in shaping language production. We did not examine finer-grained levels of statistical structure, such as that found on the word level (e.g., verb biases; Bernolet & Hartsuiker, 2010), and individual variation in implicit learning may have a stronger effect on that level. Indeed, Conway et al.'s (2010) finding that Simon task performance relates to language processing employed a task where word-level expectancies were central to performance. The relationship between implicit learning tasks and language processing tasks may not be straightforward, and may depend on the degree to which the learning demands of the two tasks overlap.

References

Bates, D.M., Maechler, M., & Bolker, B. (2011). lme4: Linear mixed-effects models using S4 classes, R package version 0.999375-39.

Bernolet, S., & Hartsuiker, R. J. (2010). Does verb bias modulate syntactic priming? *Cognition*, *114*, 455-461.

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.

Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review*, *113*, 234-272.

Conway, C. M., Bauernschmidt, A., Huang, S. S., & Pisoni, D. B. (2010). Implicit statistical learning in language processing: Word predictability is the key. *Cognition*, *114*, 356-371.

Coyle, J. M., & Kaschak, M. P. (2008). Patterns of experience with verbs affect long-term cumulative structural priming. *Psychonomic Bulletin and Review, 15*, 967-970.

Farmer, T., Fine, A.B. & Jaeger, T.F. (2011). Implicit Context-Specific Learning Leads to Rapid Shifts in Syntactic Expectations. *33rd Annual Meeting of the Cognitive Science Society*. Boston, MA.

Ferreira, V. S. (2003). The persistence of optional complementizer mention: Why saying a “that” is not saying “that” at all. *Journal of Memory and Language, 48,* 379-398.

Gebauer, G. F., & Mackintosh, N. J. (2007). Psychometric intelligence dissociates implicit and explicit learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 34-54.

Hartsuiker, R. J., Bernolet, S., Schoonbaert, S., Speybroeck, S., & Vanderelst, D. (2008). Syntactic priming persists while the lexical boost decays: Evidence from written and spoken dialogue. *Journal of Memory and Language*, *58*, 214-238.

Hartsuiker, R. J., & Westenberg, C. (2000). Word order priming in written and spoken sentence production. *Cognition*, *75*, B27-B39.

Jaeger, T. F., & Snider, N. (2008). Implicit learning and syntactic persistence: Surprisal and cumulativity. In B. C. Love, K. McRae & V. M. Sloutsky (Eds). *Proceedings of the 30th Annual Conference of the Cognitive Science Society*. Washington, DC: Cognitive Science Society.

Kaschak, M. P. (2006). What this construction needs is generalized. *Memory and Cognition*, *34*, 368-379.

Kaschak, M. P. (2007). Long-term structural priming affects subsequent patterns of language production. *Memory and Cognition*, *35*, 925-937.

Kaschak, M. P., & Borreggine, K. L. (2008). Is long-term structural priming affected by patterns of experience with individual verbs? *Journal of Memory and Language*, *58*, 862-878.

Kaschak, M. P., & Glenberg, A. M. (2004). This construction needs learned. *Journal of Experimental Psychology: General*, *133*, 450-467.

Kaschak, M. P., Kutta, T. J,. & Schatschneider, C. (2011). Long-term cumulative structural priming persists for (at least) one week. *Memory and Cognition*. *39*, 381-388.

Kaschak, M. P., Loney, R. A., & Borreggine, K. L. (2006). Recent experience affects the strength of structural priming. *Cognition*, *99*, B73-B82.

Kaufman, S. B., DeYoung, C. G., Gray, J. R., Jiminez, L., Brown, J., & Mackintosh, N. (2010). Implicit learning as an ability. *Cognition*, *116*, 321-340.

Kraljic, T., Brennan, S. E., & Samuel, A. G. (2008). First impressions and last resorts: How listeners adjust to speaker variability. *Psychological Science*, *19*, 332-338.

Misyak, J. B., & Christiansen, M. H. (in press). Statistical learning and language: An individual differences study. *Language Learning*.

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.

R development core team (2007). R: A language and environment for statistical computing. Vienna: R Foundation for Statistical Computing, http://www.R-project.org.

Reber, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious*. Oxford University Press.

Reitter, D., Keller, F., & Moore, J. D. (2011). A computational cognitive model of syntactic priming. *Cognitive Science*.

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

Wells, J. B., Christiansen, M. H., Race, D. S., Acheson, D. J., & MacDonald, M. C. (2009). Experience and sentence comprehension: Statistical learning and relative clause comprehension. *Cognitive Psychology, 58*, 250-271.

Acknowledgments

Michael Kaschak, Timothy Kutta, and John Jones, Department of Psychology, Florida State University, Tallahassee, FL 32306. The research was supported by NSF grant 0842620. Correspondence can be addressed to Michael Kaschak (kaschak@psy.fsu.edu).

Table 1

Mixed Logit Analysis with Raw and Estimated Means

Mixed Logit Model

Predictor	Coefficient	SE	*Z*	*p*-value
Intercept	.42	.38	1.09	.28
Time[a]	-.35	.19	-1.86	.06
Bias condition	.58	.37	1.56	.12
Time x Bias	1.51	.38	4.01	< .001

DO-only Analysis

Intercept	.69	.38	1.83	.07
Time[a]	.24	.23	1.04	.30

PO-only Analysis

Intercept	.08	.47	<1	.86
Time[a]	-1.16	.30	-3.88	<.001

Raw and Estimated Mean Proportion of DO Target Completions

	Raw Means		Estimated Means	
Bias Condition:	DO	PO	DO	PO
Pre-Bias	.61	.63	.62	.66
Post-Bias	.63	.45	.71	.40

Note: Coefficients express log odds. [a] = random slopes for participants

Table 2

Follow-up Mixed Logit Analysis

Mixed Logit Model

Predictor	Coefficient	SE	*Z*	*p*-value
Intercept	.17	.31	<1	.59
Bias condition	1.13	.30	3.81	<.001
Current run	-.017	.016	-1.09	.27
Bias x Current run	.19	.03	6.33	< .001

DO-only analysis

Intercept	.75	.30	2.50	.01
Cumulative Priming	.05	.018	2.80	< .01
Current run	.04	.03	1.46	.14

PO-only analysis

Intercept	-.42	.37	-1.13	.26
Cumulative Priming	.083	.025	3.26	.001
Current run	-.09	.028	-3.22	.001

Raw and Estimated Mean Proportion of DO Target Completions from Main Analysis

	Raw Means			Estimated Means		
Run length:	-5- -1	1-15	16-25	-5- -1	1-15	16-25
DO Bias	.36	.63	.87	.49	.68	.82

PO Bias	.65	.45	.19		.67	.40	.18

Note: Coefficients express log odds. For the purposes of generating raw and estimated means, we split the Current run variable into three bins, with one bin representing runs of the opposite construction (i.e., negative numbers), and the remaining two bins splitting the positive runs. The number of observations per bin is as follows. DO Bias: -5 - -1 = 70; 1 – 15 = 106; 16 – 25 = 76; PO Bias: -5 - -1 = 97; 1 – 15 = 97; 16 – 25 = 78. The number of observations varies across bins because a) we needed to keep all observations with the same run length in the same bin, and b) we wanted to have the same bin parameters for both conditions. To generate the estimated means, we used run values corresponding to the mean run length within each bin. The mean run length within each bin was slightly different for the DO and PO Bias conditions, so we rounded to the nearest whole number between the two condition means, and then converted this value to its corresponding “centered” value (as the variables in the regression were centered). Thus, we employed run values of -10, 0, and 10 to generate the estimated means.

Figure Caption

Figure 1. Estimated proportion of DO target completions as a function of Current run and Cumulative probability for the DO (top) and PO (bottom) Bias conditions. Predicted values are derived from the regression models at the bottom of Table 2. The variables of interest were centered in the analyses, but we have converted the scales back to raw run length and raw cumulative probability for the figure. The probability and run lengths selected represent the mean of each variable within each Bias condition, plus the value 1 SD above and 1 SD below the mean (rounded to whole numbers). “Probability” refers to the probability of producing a DO. In the DO condition, negative run lengths refer to runs of the PO, and positive run lengths refer to runs of the DO. In the PO condition, negative run lengths refer to runs of the DO, and positive run lengths refer to runs of the PO.

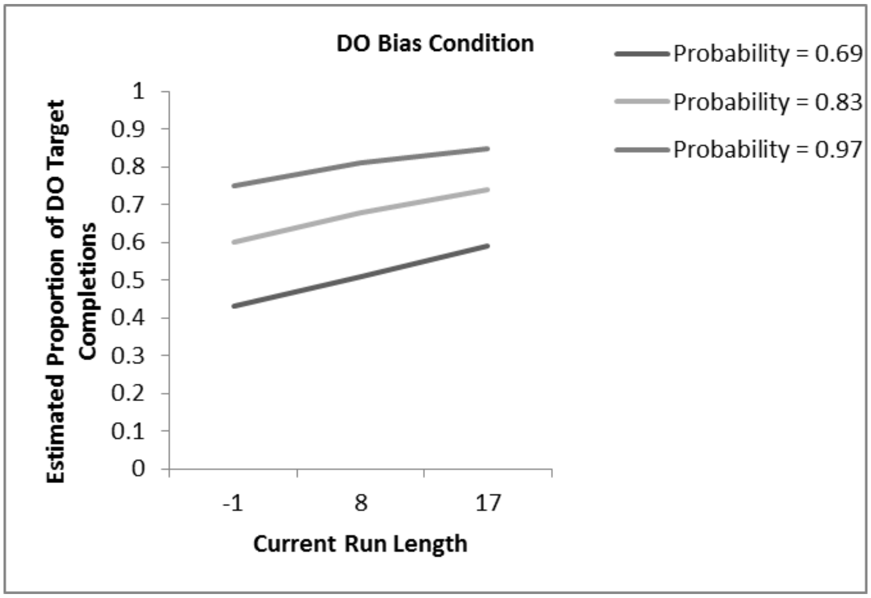
DO Bias Condition
Probability = 0.69
Probability = 0.83
Probability = 0.97
Estimated Proportion of DO Target Completions
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-1
8
17
Current Run Length

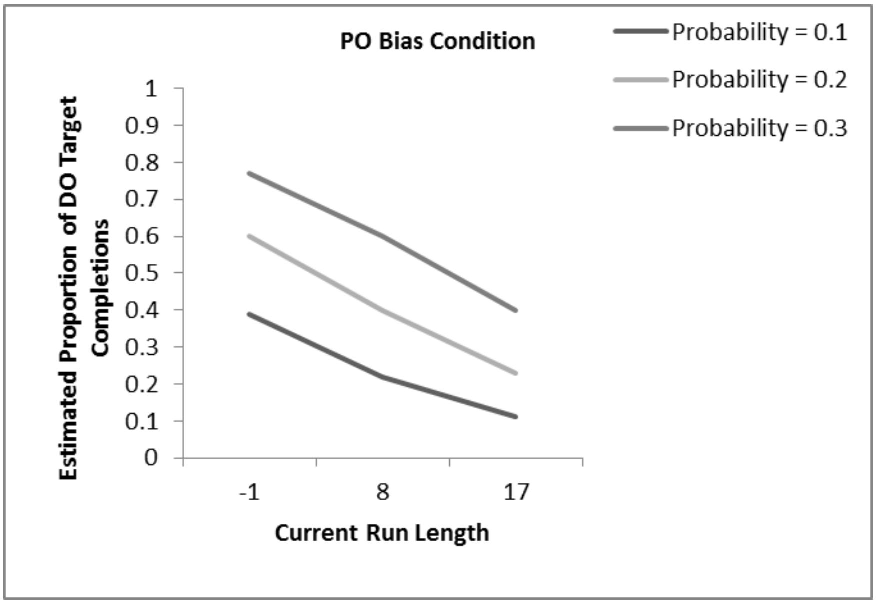
PO Bias Condition
Probability = 0.1
Probability = 0.2
Probability = 0.3
Estimated Proportion of DO Target Completions
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-1
8
17
Current Run Length