Reply to Farrell and Lewandowsky: Recency–contiguity interactions predicted by the temporal context model

Marc W. Howard
Syracuse University, Syracuse, New York

Per B. Sederberg
Princeton University, Princeton, New Jersey

and

Michael J. Kahana
University of Pennsylvania, Philadelphia, Pennsylvania

Farrell and Lewandowsky (2008) argued that the temporal context model (TCM; Howard & Kahana, 2002) cannot explain nonmonotonicities in the contiguity effect seen at extreme lags. However, TCM actually predicts these nonmonotonicities to the extent that end-of-list context persists as a retrieval cue during recall and to the extent that end-of-list context generates a recency effect. We show that the observed nonmonotonicity in the contiguity effect interacts with the recency effect, as predicted by TCM. In conditions, such as immediate and continual-distractor free recall, that exhibit strong recency, one observes more prominent nonmonotonicities in the contiguity effect than in conditions, such as delayed free recall, that attenuate recency. The nonmonotonicities in the contiguity effect at extreme lags, and the interactions between recency and contiguity, result from the role of end-of-list context as a retrieval cue in TCM. Results of an additional simulation based on the Howard and Kahana (2002) version of TCM may be downloaded from http://pbr.psychonomic-journals.org/content/supplemental.

In free recall studies, participants are presented with a list of words and are then instructed to recall them in the order in which they come to mind. Because the order of recall is unconstrained by the experimenter, regularities in the transition probabilities reflect properties of the organization of memory. Perhaps the most important of these regularities in constraining models of episodic memory retrieval is the conditional response probability (CRP) as a function of lag (i.e., lag-CRP; see Kahana, 1996). Given that a participant has just recalled an item from serial position i, the lag-CRP estimates the probability that the next item recalled will be from serial position $i +$ lag, attempting to control for the availability of potential recalls. The lag-CRP has a canonical shape, exhibiting a strong contiguity effect favoring adjacent transitions over remote transitions and an asymmetry favoring forward transitions over backward transitions (see Kahana, Howard, & Polyn, 2008, for a review).

According to the temporal context model (TCM; Howard, Fotedar, Datey, & Hasselmo, 2005; Howard & Kahana, 2002; Howard, Kahana, & Wingfield, 2006; Polyn, Norman, & Kahana, 2009; Sederberg, Howard, & Kahana, 2008), the contiguity effect arises because items from neighboring serial positions are associated with similar contexts. In the TCM, recall of an item retrieves its associated contexts, which, in turn, cue items studied in nearby list positions. By assuming that context evolves as a function of the sequence of presented or recalled items, TCM predicts a forward bias in the contiguity effect (see Howard & Kahana, 2002, for details). The experience-dependent nature of contextual drift also implies that the context at the start of an immediate recall test will be more similar to the contexts associated with recently studied items than to those studied farther in the past. In this way, TCM accounts for the recency effect as well as the contiguity effect.

Since Howard and Kahana's (2002) original formulation of TCM, the model has been extended to address the effects of hippocampal lesion on transitive associations (Howard et al., 2005), the effect of aging on the shape of lag-CRP curves (Howard et al., 2006), the dynamics of immediate, delayed, and continual-distractor free recall (Sederberg et al., 2008), and the interactions between temporal, semantic, and source clustering in recall (Polyn et al., 2009). All of these applications of TCM have preserved the model's basic assumptions concerning how contextual dynamics give rise to recency and contiguity.

Reanalyzing data from previously published free recall studies, Farrell and Lewandowsky (2008) showed that, in addition to the contiguity effect, participants also tend to make transitions to the ends of the list more frequently than transitions of intermediate lag (see Figure 2 in their article). This tendency, which appears as a nonmonotonicity in the shape of the lag-CRP functions, is particularly striking for forward transitions. If we follow the forward lag-CRP outward from zero, eventually a persistent recency effect overcomes the contiguity effect, resulting in an increase in the lag-CRP for large positive lags. Similarly, a persistent primacy effect (Howard & Kahana, 1999) would result in an increase in transitions to early list items and hence in an increase in the lag-CRP for large negative lags.[1]

The present reply seeks to clarify the empirical and theoretical analysis of nonmonotonicities in the lag-CRP. Empirically, we start by describing the source of the nonmonotonicity predicted by TCM and then turn to an analysis of the variables that affect the degree to which these nonmonotonicities are present in the data. On the theoretical side, we present modeling results indicating that the pattern of results observed in the data is consistent with the predictions of TCM.

M. J. Kahana, kahana@psych.upenn.edu

Nonmonotonicity in the Lag-CRP

According to TCM, temporal context changes gradually as a function of the sequence of studied and recalled items. At the time of recall, the current state of context serves as a retrieval cue for items studied in similar contexts. Specifically, the current state of context activates items in memory to the extent that their associated contexts are similar to the current context. When an item is recalled, it retrieves its associated contexts, which combine with the current context and together serve as a cue for the next recall. This enables the model to account for contiguity effects: When a studied item is recalled, the contextual state it retrieves will resemble that of the encoding context of neighboring list items, resulting in an increased tendency to recall neighbors of the recalled item. These basic ideas are common to all previous applications of TCM and to its extensions (see Howard et al., 2005; Howard & Kahana, 2002; Howard et al., 2006; Polyn et al., 2009; Rao & Howard, 2008; Sederberg et al., 2008).

TCM predicts a change in the shape of the lag-CRP to the extent that end-of-list context persists as part of the retrieval cue and to the extent that this end-of-list context supports a recency effect.[2] In TCM, the degree of contextual drift at any given time step is a function of the amount of information that is provided as input to temporal context. This leads to the interesting property that, when no input is provided, there is no change in the state of temporal context. This accounts for the finding that recency is preserved when the delay between study and test is not filled with distracting activity (Baddeley & Hitch, 1977; Murdock, 1963). Given that the amount of information provided as input to temporal context may differ between encoding and retrieval of an item, the rate of contextual drift during study may differ from the rate of drift during retrieval. In their analysis of TCM, Farrell and Lewandowsky (2008) considered two special cases, which they treated as dichotomous. They referred to the case in which temporal context drifts at the same rate during study and during recall as TCM_{evo}; the case in which temporal context at recall is uncorrelated with the contexts associated with study items they referred to as TCM_{pub}.

The following simulations use the Howard et al. (2006) implementation of TCM (see the Appendix). Other variants of TCM, including the Sederberg et al. (2008) and the Howard and Kahana (2002) models, exhibit similar behavior. Farrell and Lewandowsky (2008) used a version of TCM based closely on the version presented in Howard and Kahana (2002). Simulation results for the Howard and Kahana (2002) model, along with a derivation that sheds light on why Farrell and Lewandowsky did not observe the qualitative patterns we report here, are provided as supplemental materials that can be downloaded from http://pbr .psychonomic-journals.org/content/supplemental.

Figure 1A illustrates the predicted lag-CRP functions for immediate free recall as a function of the rate of contextual drift during the test phase of an experiment. The model parameter ρ_{test} controls the rate of contextual drift at test by specifying the degree to which the preceding state of context persists in the subsequent state of context. Large values

of ρ_{test} imply a slow rate of drift at test; high values of ρ_{test} imply a faster rate of drift. When ρ_{test} is near zero, context will have been effectively reset at the time of test and there will be no recency effect. When the rate of contextual drift is identical at study and at test (i.e., $\rho_{test} = \rho_{study}$), context at the time of test will be similar to that for end-of-list items, increasing their probability of being recalled and producing the nonmonotonicity and skew seen in the lag-CRP functions. For backward recalls, both retrieved context and time-of-test context favor retrieval of items contiguous to the item just recalled. In contrast, for retrievals in the forward direction, these cues are in conflict, resulting in both contiguity and recency in the lag-CRP.

As revealed by the preceding analysis, TCM predicts forward nonmonotonicities to the extent that time-of-test context facilitates retrieval of end-of-list items. If recall of an item produces a large change in the current state of context, then the next transition will not exhibit a strong bias toward end-of-list items. Rather, the recall transitions will primarily cluster around the item just recalled. The same logic applies to the case of delayed free recall, where participants are given a demanding distractor task after the presentation of the study list and before the recall phase. Figure 1B shows the lag-CRP predicted by TCM for immediate free recall (black curve) and for delayed free recall with increasingly long retention intervals (successively lighter curves). In this simulation, $\rho_{test} = \rho_{study}$ and the input to the context vector during the distractor task is assumed to be orthogonal to the representations of the list items that compete for recall. It is well known that increasing the retention interval between study of the last item and test results in a decrease of the recency effect in free recall (e.g., Postman & Phillips, 1965). As can be seen from Figure 1B, skew and nonmonotonicity are present when the test is immediate and gradually decrease with the increase in the retention interval.

Note that the results shown in Figure 1B falsify one of the claims of Farrell and Lewandowsky (2008) regarding the predictions of TCM:

> An initial examination of the model revealed a striking non-monotonicity of the forward lag-CRP functions in TCM_{evo}. Irrespective of whether recall was immediate or delayed or involved a continuous distractor task, lags greater than 5 attracted nearly as many—or indeed more—transitions than did lags +1. (p. 1243)

The light gray lines in Figure 1B decrease monotonically. At extreme lags, they do not in any way approach—let alone exceed—the much higher values observed at lags near zero. Farrell and Lewandowsky's (2008) conclusion is not a property of TCM in general, but rather is a property of the parameter values they chose (see the supplemental materials for a direct counterexample to their claim).

Skew and Nonmonotonicity of the Lag-CRP Across Delay Conditions

Here we examine the effects of different experimental manipulations on the nonmonotonicity and skew in the

A

B
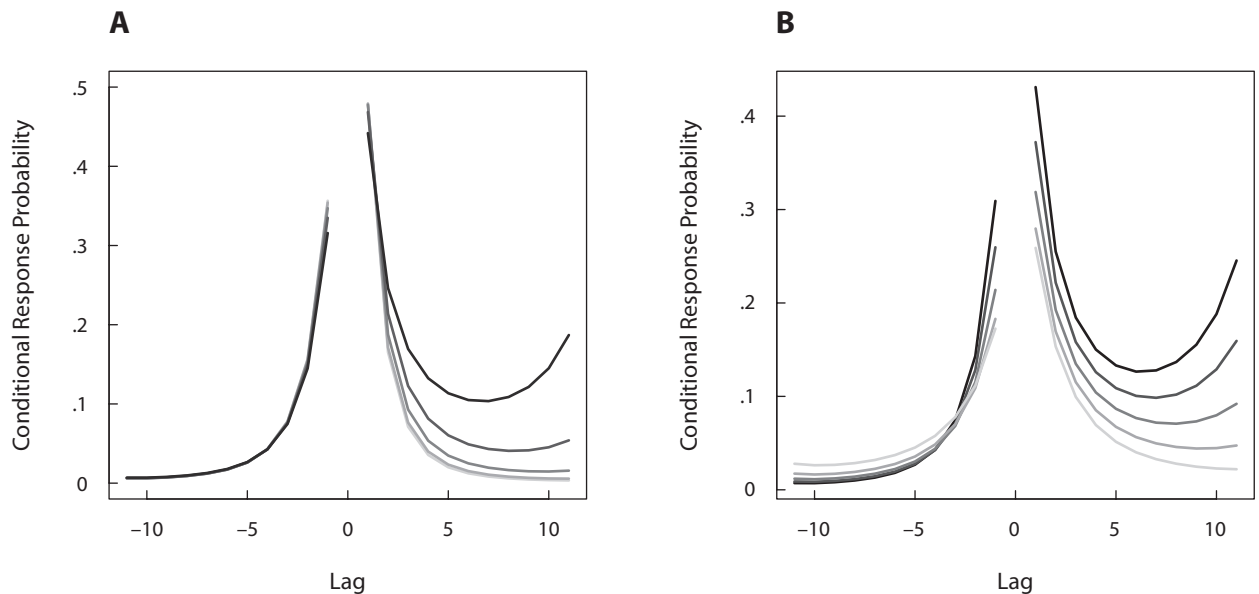
Figure 1. Recall transitions favor both contiguous and recent items. (A) Predicted effect of ρ_{test} on the lag-CRP for the first recall transition. The value of ρ_{test} is varied from ρ_{study} (black line) down to zero (indicated by the lightest shaded line). (B) Predicted effect of delay interval on the lag-CRP for $\rho_{test} = \rho_{study}$. Retention interval was varied between zero (immediate free recall, black line) to an infinite delay (lightest shaded line). For both panels, $\rho = .85$, $\gamma = .8$, and $\tau = .3$; ρ controls the rate of contextual drift, and high values of ρ imply a slow rate of drift; γ controls the degree to which recalled items recover their temporal context and thus the degree of asymmetry observed in the lag-CRP; and τ controls the sensitivity of the Luce choice retrieval rule that maps activations onto probability of recall (see the Appendix for details).

lag-CRP. In order to control as many variables as possible, we restrict our attention to published experiments from our labs in which relatively short lists of words were presented visually under conditions designed to minimize rehearsal and for which verbal free recall was collected. These analyses use a subset of the experiments examined by Farrell and Lewandowsky (2008) as well as final free recall (FFR) data (reported in Howard, Youker, & Venkatadass, 2008) from the Howard, Venkatadass, Norman, and Kahana (2007) immediate free recall study. In an attempt to maximize the amount of available data, we examine lag-CRPs collapsed across all recall transitions rather than only considering the first recall transition. Our analyses suggest that skew and nonmonotonicity in the lag-CRP are observed to the extent that there is a recency effect, regardless of whether recall is immediate or delayed. In order to establish this, we first summarize the recency effect observed in these studies.

Figure 2 shows the recency effect observed in the experimental data we consider here. Figure 2A shows the probability of first recall (PFR) curves from Experiments 1 and 2 of Howard and Kahana (1999), exhibiting immediate, delayed, and continual-distractor free recall. In *immediate free recall*, the list is presented and test immediately follows presentation of the last item, and a strong recency effect is observed in the PFR. In *delayed free recall*, a delay intervenes between study of the last item and test, and the recency effect is attenuated. In *continual-distractor free recall*, a delay intervenes between each list item and also occurs at the end of the list. The recency effect in the PFR is larger in continual-distractor free recall than in delayed free

recall. The long-term recency effect in continual-distractor free recall is also observed when FFR is examined across lists (Glenberg et al., 1980; Tzeng, 1973).

Figure 2B illustrates recency on the much longer time scale of entire lists. This graph shows the PFR obtained when participants were asked to freely recall items presented on 48 distinct study–test lists at the end of the experimental session (Howard et al., 2007). A recency effect was observed in immediate free recall of each list, but as shown in Figure 2B, the across-list recency effect extends back by at least five lists.

Figure 3 illustrates lag-CRPs from immediate, delayed, and continual-distractor free recall, as well as FFR across lists. Whereas prior work has focused on documenting the existence of a contiguity effect by focusing on lags around zero (e.g., Howard & Kahana, 1999; Howard et al., 2008), here we examine all possible lags, as suggested by Farrell and Lewandowsky (2008). Figure 3A compares immediate with delayed free recall from Experiment 1 of Howard and Kahana (1999). There is a boost in the contiguity effect, as well as what appears to be a larger nonmonotonicity in immediate free recall than with delayed free recall.

Figure 3B compares the 16-sec continual-distractor condition of Experiment 2 of Howard and Kahana (1999) with the delayed free recall condition from the same experiment (the latter condition has no interitem distraction). Continual-distractor free recall shows a larger recency effect than does delayed free recall (Figure 2A). Although the contiguity effect is similar in magnitude across the conditions, the nonmonotonicity exhibited at extreme

A

B
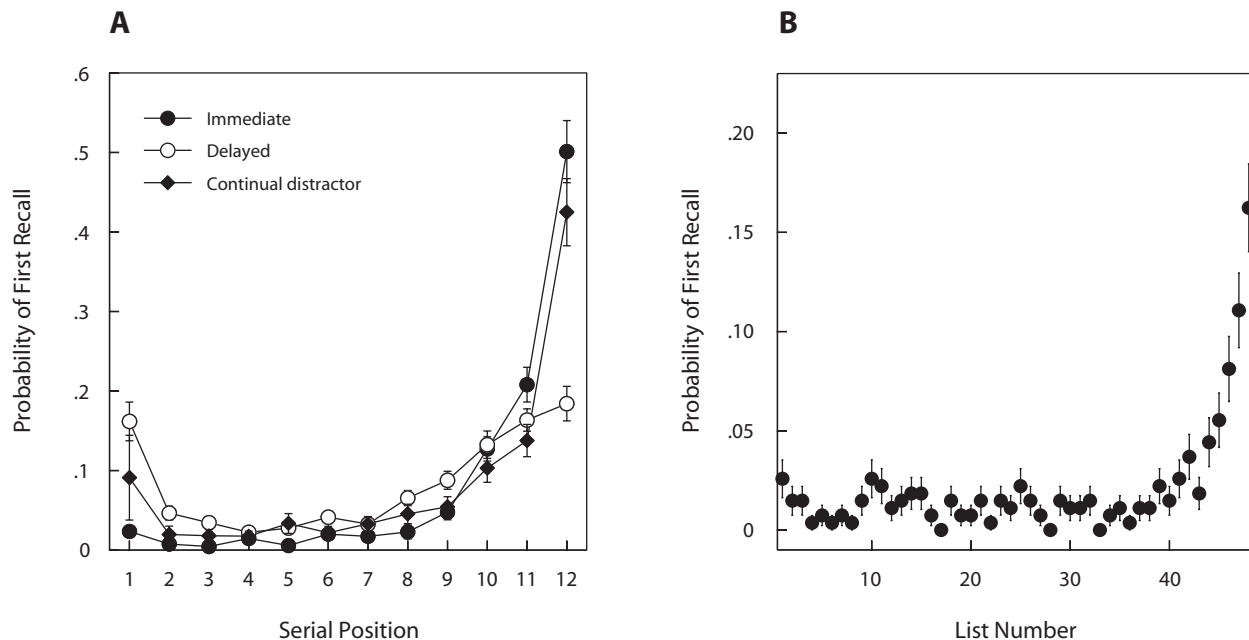
Figure 2. Recency across time scales. (A) The recency effect, as illustrated by the probability of first recall (PFR), from immediate, delayed, and continual-distractor free recall from Experiments 1 and 2 from Howard and Kahana (1999). Whereas the recency effect is attenuated in delayed free recall, it is amplified in continual-distractor free recall. (B) The recency effect in the PFR across lists from Howard, Youker, and Venkatadass (2008). Error bars represent ±1 standard error of the mean.

positive lags appears larger in continual-distractor free recall than in delayed free recall.

Figure 3C compares the lag-CRP from immediate free recall in the Howard et al. (2007) data with the within-list lag-CRP observed in FFR of the same items. The FFR data is comparable with data for delayed free recall with an extremely long delay and thus does not demonstrate a within-list recency effect. The results appear to be consistent with Figure 3A. Again, the contiguity effect is larger in magnitude in immediate recall. In addition, a nonmonotonicity in the forward direction is observed in immediate free recall. With the delay, and leaving aside the primacy effect observed at extreme backward lags, the lag-CRP has a consistent degree of asymmetry in contrast to the skew observed in the immediate free recall data. Figure 3D illustrates the across-list lag-CRP observed by Howard et al. (2008). There appears to be a steep non-monotonicity observed at extreme positive lags, presumably corresponding to the strong across-list recency effect observed in the same data (see Figure 2B).

TCM Predictions

We next examined whether the effects of delay on the lag-CRP curves are consistent with the predictions of TCM. Our strategy was to use a common set of parameters that illustrate the qualitative behavior of the model across conditions, which we compare with the pattern of results observed across experiments. The goal of this approach is to provide insight as to whether the source of the skew and nonmonotonicity in the lag-CRP is consistent with the origin predicted by TCM. Note that our goal is not to find the parameters that maximize the likelihood of a single

data set under the assumption that the model is correct, but to provide insight into the basis of a phenomenon by describing a variety of experimental manipulations across multiple studies with a single set of parameters.

Farrell and Lewandowsky (2008) evaluated maximum likelihood fits of a two-parameter TCM model and found that the predictions of the model deviated significantly from the data. This is not surprising, given the many sources of variability that are not included in their two-parameter formulation. For instance, it is known that free recall is strongly affected by the degree of proactive interference the items are subject to (Goodwin, 1976), the duration of the delay interval (Postman & Phillips, 1965), and the semantic organization of the list (Glanzer, Koppenaal, & Nelson, 1972; Romney, Brewer, & Batchelder, 1993). As shown by Polyn et al. (2009), an elaborated version of TCM that handles semantic memory effects and task manipulation effects, among others, provides a nice qualitative fit to the full range of lag-CRP phenomena (see Figure 3 in Polyn et al., 2009).

In the simulations that follow, we use a reduced version of the Howard et al. (2006) instantiation of TCM with $\rho_{\text{test}} = \rho_{\text{study}}$ (see the Appendix for details). Simulations for the Sederberg et al. (2008) version of TCM are reported in the body of the article; results for the Howard and Kahana (2002) version, used by Farrell and Lewandowsky (2008), are reported as supplemental materials. In reevaluating TCM, we informally searched for a set of parameters that would exhibit recency and contiguity effects similar to those shown in Figures 2A and 3A–3C. A relatively broad range of parameters exhibit the same basic properties. These parameters, with $\rho = .85$, $\gamma = .8$,
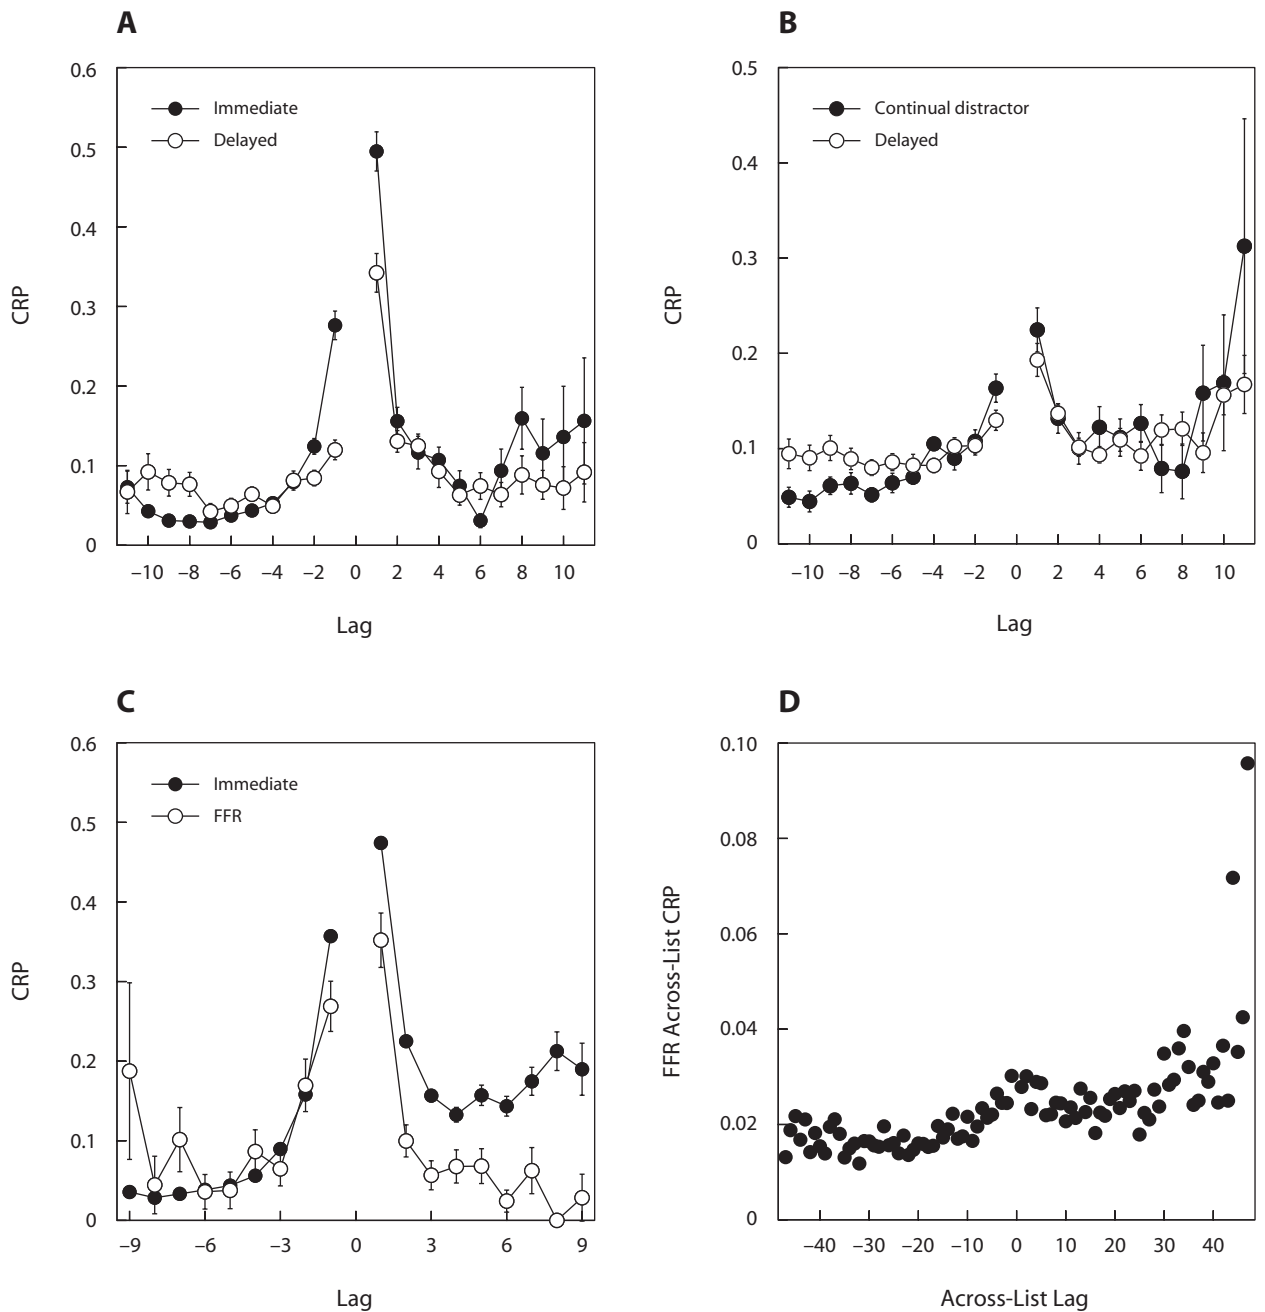
Figure 3. Nonmonotonicities in the lag-CRP across manipulations of the recency effect. (A) Immediate and delayed conditions of Howard and Kahana (1999, Experiment 1). (B) Continual-distractor and delayed conditions of Howard and Kahana (1999, Experiment 2). (C) Immediate and final free recall (FFR) data from Howard et al. (2007, control lists). (D) Across-list lag-CRP (Howard et al., 2008). In all panels, lag-CRPs are calculated across all recall transitions. Error bars in panels A–C represent ±1 standard error of the mean.

$\tau = .3$, and $\rho_D = .4$, were used in the predictions generated in Figures 1A and 1B, Figure 4A, and Figures 5A–5C (see the Appendix for further details): γ controls the degree to which recalled items recover their temporal context and thus the degree of asymmetry observed in the lag-CRP, τ controls the sensitivity of the Luce choice retrieval rule that maps activations onto probability of recall, and ρ_D controls the effective length of the distractor interval

where appropriate (see the Appendix). Because we treated the lists as single items in the across-list FFR simulations, a separate set of parameters was chosen for the across-list free recall data ($\rho = .8$, $\gamma = .8$, and $\tau = .6$). In all of these predictions, ρ_{study} was equal to ρ_{test}.

As is shown in Figure 4A, TCM correctly predicts a strong recency effect in immediate free recall, an attenuation of the recency effect in delayed free recall, and a

strong recency effect in continual-distractor free recall (cf. the empirical results shown in Figure 2A; see also Howard & Kahana, 2002). In modeling the recency effect in FFR across lists, TCM can also predict a recency effect that extends across multiple lists (see Figure 4B).

In evaluating TCM's ability to account for the skew and nonmonotonicities in the lag-CRP, we focused on the first transition in recall. This simplification is necessary, because the early versions of TCM considered by Farrell and Lewandowsky (2008) did not specify mechanisms for dealing with resampling and termination of recall. These limitations are overcome in more recent implementations of TCM that replace the Luce choice decision rule with a set of leaky accumulators and that specify a set of resampling rules for dealing with multiple retrievals (Polyn et al., 2009; Sederberg et al., 2008).

Figure 5 shows the lag-CRPs for the first recall transition based on the same parameters that generated the corresponding PFR curves in Figure 4. As can be seen from Figure 5A, TCM correctly predicts that the lag-CRP from immediate free recall should have a larger contiguity effect, somewhat greater skew, and a larger nonmonotonicity in the forward direction than does the lag-CRP from delayed free recall. In continual-distractor free recall (Figure 5B), the model correctly predicts that, although the contiguity effect should be of similar magnitude across conditions, the difference between the lag-CRP from continual-distractor free recall and delayed free recall should be manifested as a large and sharp nonmonotonicity at extreme positive lags in continual-distractor free recall relative to delayed free recall.[3]

Some caution should be exercised in comparing the predictions of TCM with the empirical data shown in Figure 3, which collapse across recall transitions. Whereas

the overall shape of the lag-CRP is similar for early and late recall transitions, the magnitude of the contiguity effect in immediate free recall is much greater for the first several responses (Howard & Kahana, 1999; Kahana, Howard, Zaromb, & Wingfield, 2002). TCM underpredicts the magnitude of the contiguity effect at early recall transitions in immediate free recall (Howard & Kahana, 2002). The use of lag-CRPs from the first recall transition in the modeling and the use of lag-CRPs collapsed across recall transitions in the data tends to obscure this difference in immediate free recall (but see Polyn et al., 2009; Sederberg et al., 2008).

Figure 5C shows predictions comparing the lag-CRP from the immediate free recall within-list lag-CRP with the lag-CRP from FFR of the same lists. To simulate the very long delay between study of a typical list from the experiment and the FFR session, we simply set the effective length of the retention interval for the within-list lag-CRP to be infinite (the gray lag-CRP curve in Figure 5C), rather than modeling delayed free recall as reflecting a small residual recency effect (Figures 2A and 4A). Indeed, there is no evidence for a within-list recency effect or a primacy effect in these FFR data (Howard et al., 2008). The correspondence between the predictions (Figure 5C) and the empirical observations (Figure 3C) in this case is particularly strong. The model not only correctly predicts a nonmonotonicity in the immediate lag-CRP that is larger than in FFR, but it also describes a continually increasing difference between immediate free recall and delayed FFR for increasing values of lag. Moreover, the model also correctly predicts a benefit for backward transitions in delayed FFR.

Comparing Figure 5D, which shows predicted values of across-list lag-CRPs, with Figure 3D, which shows empiri-
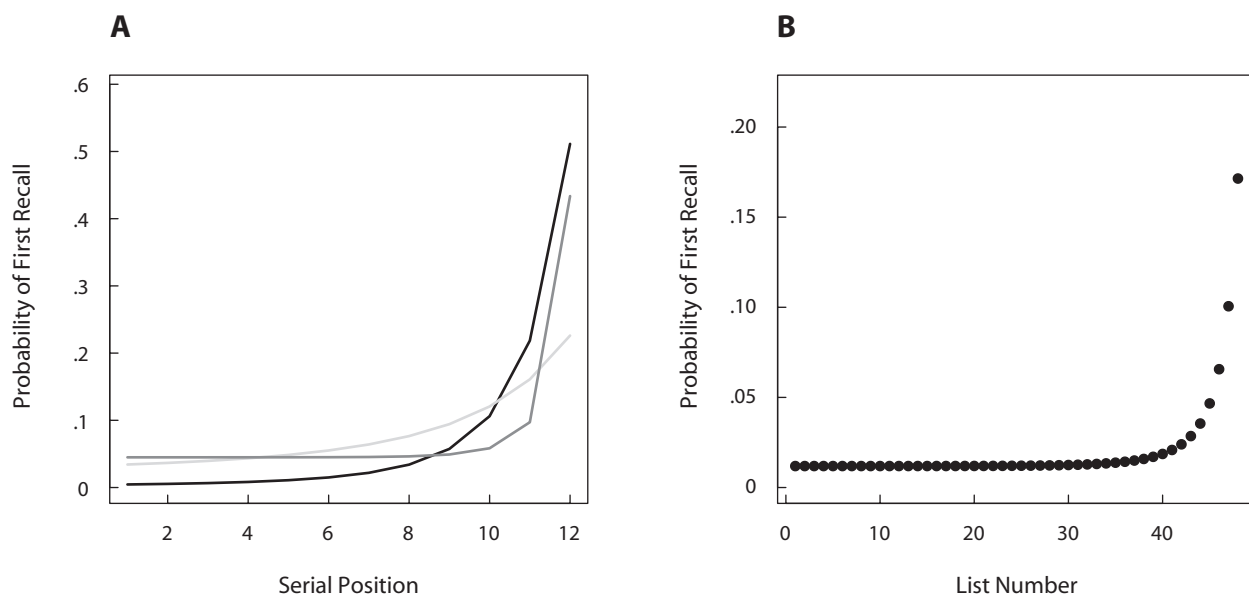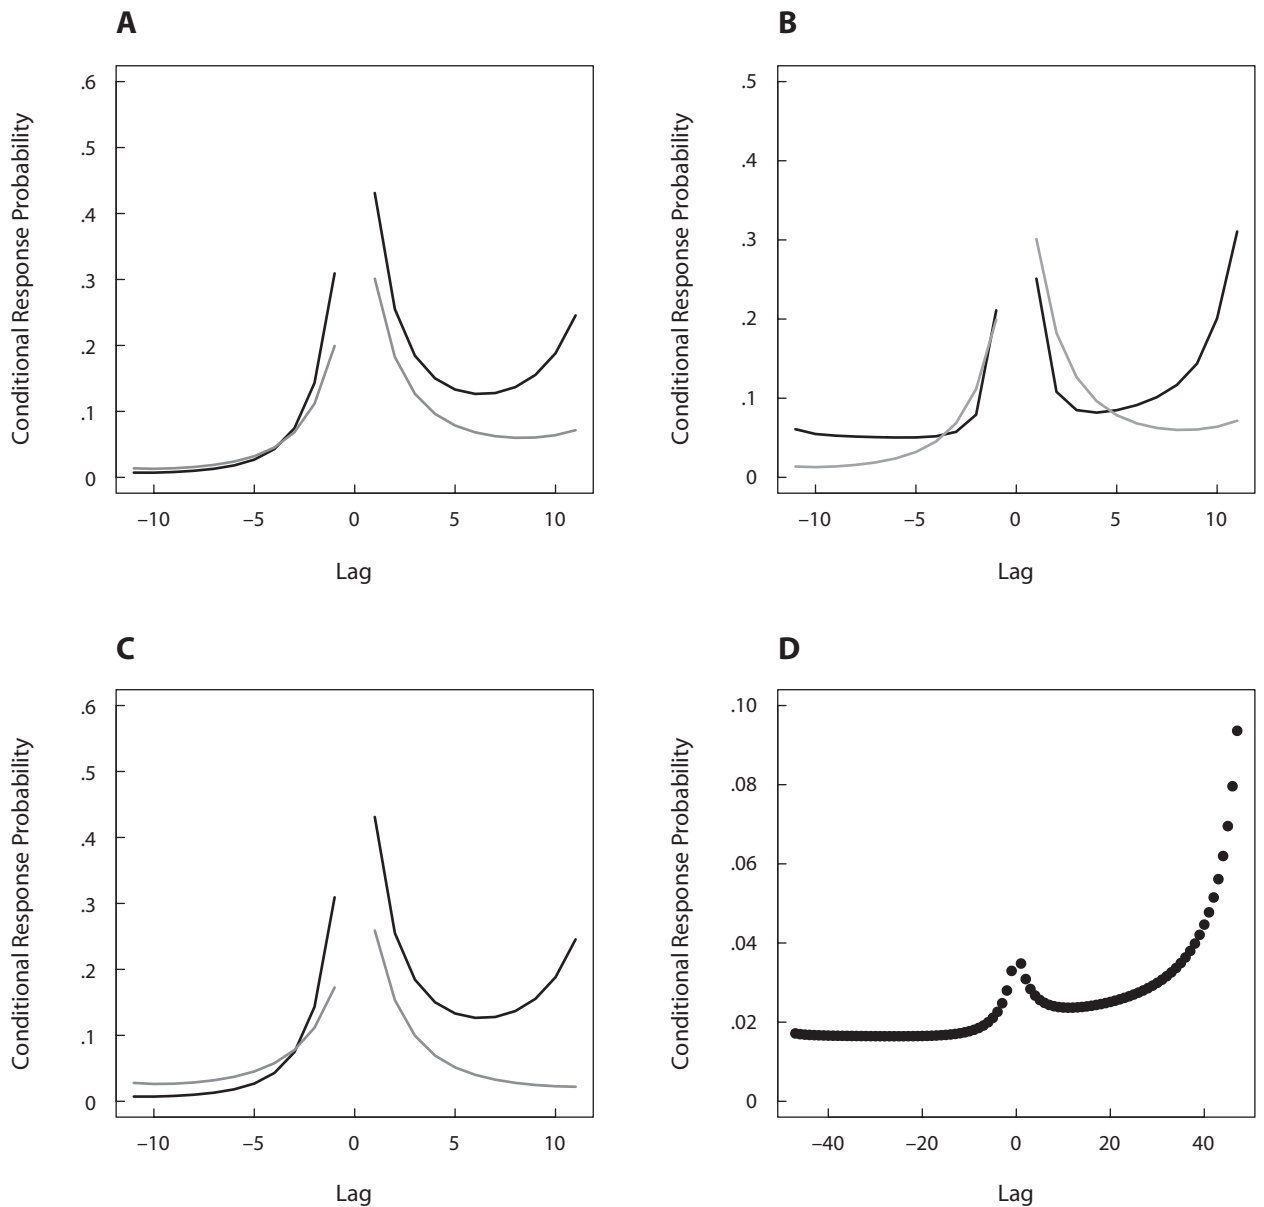

Figure 4. Recency across time scales in the temporal context model (TCM). Compare to Figure 2. (A) Probability of first recall (PFR) functions for immediate (black), delayed (light gray), and continual-distractor (dark gray) free recall. The same parameters were used in Figures 5A–5C. (B) Across-list PFR. The same parameters were used in Figure 5D.

Figure 5. Simulations of lag-CRPs from the first recall transition using TCM with the Luce choice retrieval rule. Compare to Figure 3. (A) Immediate free recall (black) and delayed free recall. (B) Continual-distractor free recall (black) and delayed free recall (gray). (C) Immediate free recall compared with delayed free recall with an infinite delay. (D) Across-list CRP.

cally observed across-list lag-CRPs, we see that the model has successfully captured several aspects of the data. First, the model predicts a contiguity effect across several lists that is also exhibited in the data.[4] Second, there is a large nonmonotonicity in the forward direction, such that extremely large across-list lags are recalled better than adjacent across-list lags are. Third, there is a persistent skew to the entire curve that appears to be reflected in the data.

From these analyses, the model appears to correctly predict the range of shapes of lag-CRP curves that are observed, and the variation in the skew and nonmonotonicity observed across experiments. In particular, these seem to co-occur with the recency effect, as predicted by TCM.

Rather than ruling out TCM as a description of recency and contiguity across scales, the data from examining the entire range of lag-CRPs across conditions are generally consistent with the predictions of TCM. Even though we achieved a good description of the data under the assumption that $\rho_{test} = \rho_{study}$ (referred to as TCM_{evo} by Farrell & Lewandowsky, 2008), it is quite possible that other values of ρ_{test} would provide even better fits.

Contiguity in TCM Is Not an Averaging Artifact

Whereas the experimental data show a strong contiguity effect following recall of items from nearly all serial positions, Farrell and Lewandowsky (2008) argued that
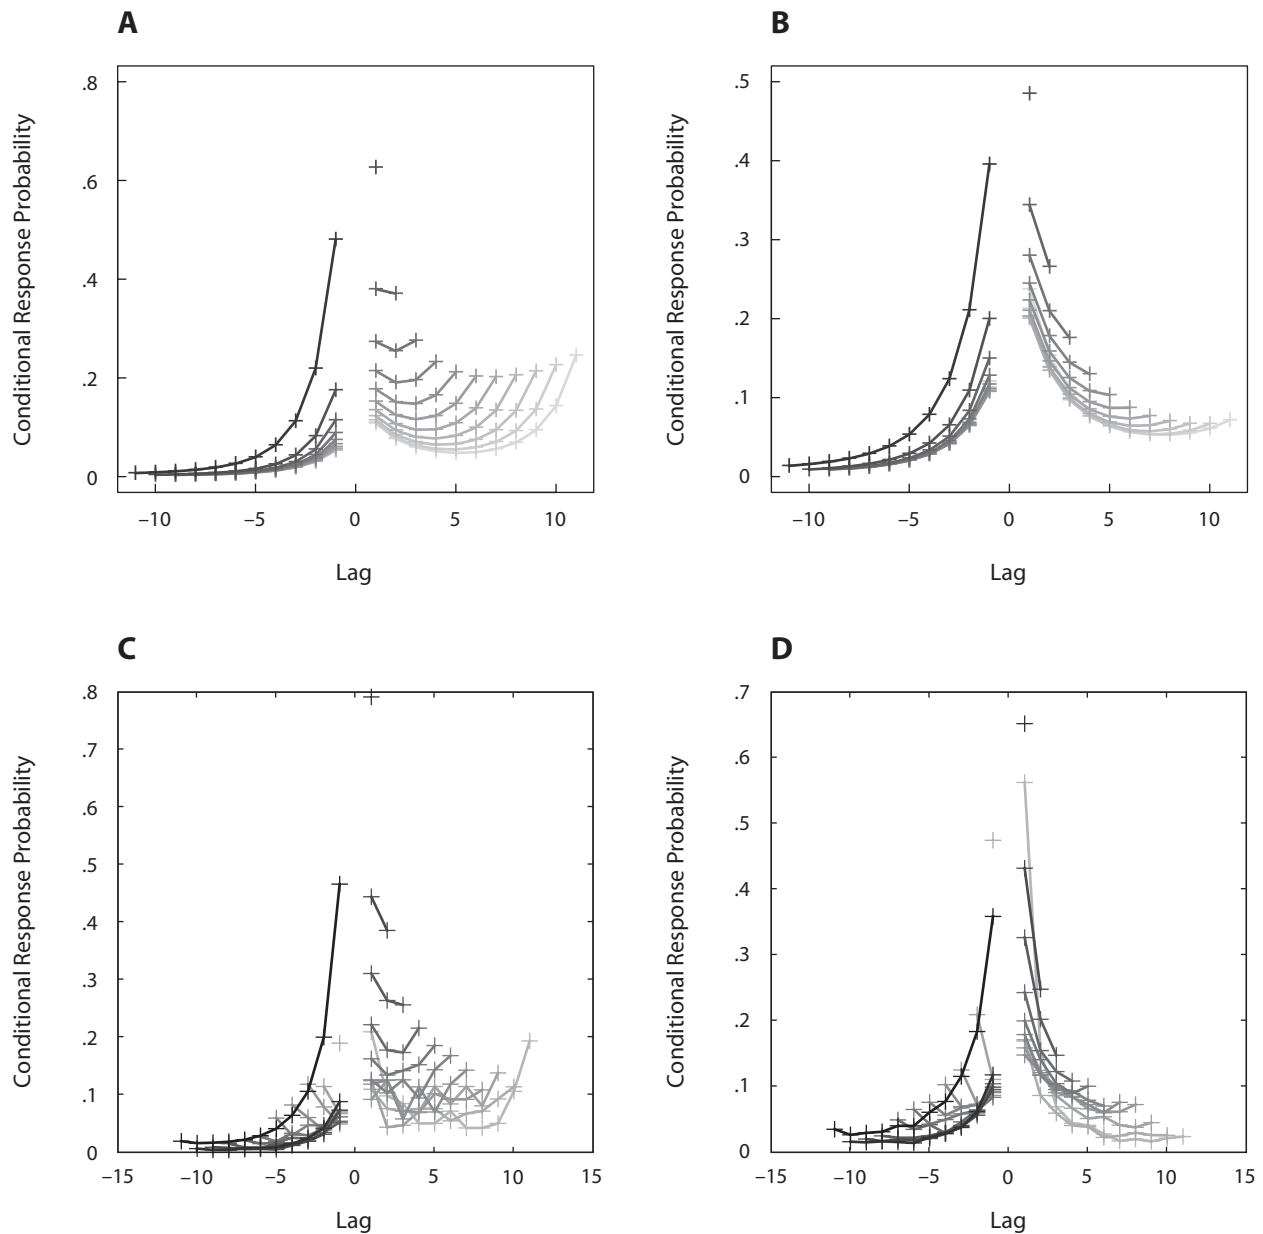
Figure 6. The CRP predicted by TCM in immediate free recall need not be an artifact of the recency effect. Compare with Figure 7 of Farrell and Lewandowsky (2008). Panels on the left (A and C) correspond to immediate free recall. Panels on the right (B and D) correspond to delayed free recall (note change of scale). Panels A and B: Simulations of the Howard, Kahana, and Wingfield (2006) version of TCM with $\rho_{test} = \rho_{study}$. The same parameters were used as in Figures 4A and Figure 5A–5C. Panels C and D: Simulations from the TCM-A model of Sederberg, Howard, and Kahana (2008). CRP curves broken down by serial position of the item just recalled are shown. These curves average over all output positions. Parameter values are as reported by Sederberg et al.

TCM does not reproduce this critical feature of the data. Rather, they claimed that TCM exhibits a pure recency effect in the lag-CRP computed separately for items recalled for different serial positions. In other words, they argued that the contiguity effect predicted by TCM in immediate free recall arises as an artifact of averaging lag-CRP functions calculated for different serial positions. Here we show that Farrell and Lewandowsky's finding is an artifact of the particular choice of parameters used in generating the predictions and does not reflect a general property

of the model (see the supplemental materials for a direct evaluation of the version of TCM they used).

Figures 6A and 6B show TCM predictions using the same parameters as Figures 4 and 5. The model clearly shows a contiguity effect at each serial position in addition to a recency effect that appears as a nonmonotonicity at extreme lags. To determine whether these results were due to some detail of the implementation of TCM used here, we examined a number of alternative implementations and found that all exhibited the same qualitative behavior. Fig-

ure S3 (supplemental materials) shows that the Howard and Kahana (2002) version of TCM produces nonartifactual lag-CRPs for appropriate parameter values, but that it can also produce an artifactual lag-CRP for the parameter values chosen by Farrell and Lewandowsky (2008).

Figures 6C and 6D show the lag-CRP segregated by serial position of the just-recalled word generated by TCM-A (Sederberg et al., 2008)—a variant of TCM in which the Luce choice rule is replaced with a set of leaky competing accumulators. These simulations employed the same set of parameters reported by Sederberg et al. As is shown in Figure 6C, TCM-A predicts a contiguity effect, as well as a nonmonotonicity in the forward direction, even when the lag-CRP is conditionalized on the serial position of the just-recalled item. Unlike the basic TCM predictions shown previously, the predictions generated by TCM-A collapse across all recall transitions, as was done in our analysis of the empirical data. Although it is possible to find parameters for TCM that generate an artifactual contiguity effect, this is not a weakness of the model per se so much as it is a weakness of the specific choice of parameters used by Farrell and Lewandowsky (2008). Although one may note that their parameters were chosen by means of maximum likelihood estimation, they manipulated only a subset of the model's parameters in fitting these data. As such, their fits are not representative of the range of behavior the model can produce.

Discussion

By analyzing free recall data aggregated across a large number of diverse studies, Farrell and Lewandowsky (2008) showed that participants' recall transitions reveal not only the standard forward-asymmetric contiguity effect (e.g., Kahana, 1996) but also an increased attraction to end-of-list items (e.g., Figure 2 of their article). They argued that this nonmonotonicity in the lag-CRP function is not predicted by the TCM. Further, they argued that TCM fails to account for critical features of the contiguity effect itself. Here we re-examine both their empirical and theoretical claims.

Regarding their empirical discovery, the analyses reported in the present article (Figure 3) replicate their basic findings and suggest that the nonmonotonicity in the contiguity effect for extreme lags is driven by persistent serial position effects. Specifically, experimental conditions that produce strong recency also exhibit strong nonmonotonicities and forward skew in the lag-CRP function. Studies that show minimal recency do not appear to exhibit significant nonmonotonicities in the lag-CRP.

Contrary to Farrell and Lewandowsky's (2008) theoretical claims, our analyses show that TCM predicts the nonmonotonicity in the lag-CRP revealed by their analysis and the tendency for this effect to increase under conditions that favor retrieval of recent list items (Figure 5). We also show that the contiguity effect in TCM is not an artifact of averaging for many appropriate parameter values. The possibility of such artifacts arising for certain parameter values, however, highlights the importance of considering serial position as a factor in any analysis of the contiguity effect.

Farrell and Lewandowsky (2008) have provided a service by pointing out that the entirety of the lag-CRP curve can yield important constraints on models of free recall. When end-of-list context is not allowed to persist ($\rho_{test} = 0$) and recall is immediate, TCM does not accurately reproduce the pattern of recall transitions to items of varying lags. However, the effect of the value of ρ_{test} is apparent only to the extent that there is a recency effect in the serial position curve, so that the error induced by setting $\rho_{test} = 0$ (or analogously by assuming an infinite delay prior to presentation of the cue) is minimal, if one is examining delayed free recall (Howard et al., 2006).

Whereas Howard and Kahana (2002) and Howard et al. (2006) developed highly simplified models of free recall based on the core ideas of context evolution and retrieval embodied in TCM, two recent TCM-based models of free recall have been developed to explain a broad array of empirical phenomena. The model of Sederberg et al. (2008) provides a description of dissociations between immediate recency and long-term recency. The model of Polyn et al. (2009) accounts for temporal, semantic, and source-based clustering effects in free recall, along with their interactions. Both of these models use a much richer retrieval process that is able to produce more realistic serial position effects and interresponse time phenomena than was possible in the reduced-form variants of free recall presented in earlier work (e.g., Howard & Kahana, 2002). As our empirical understanding of the recall process continues to advance, we will uncover further limitations of these models that will pave the way for future advances.

Note—In the original submission of their target article "Empirical and Theoretical Limits on Lag-Recency in Free Recall," S. Farrell and S. Lewandowsky (2008) had included a figure (and associated analysis) with a format similar to the one provided by the present authors in their Figure 1B. Unlike the present Figure 1B, the model-based predictions from the Farrell and Lewandowsky figure deviated qualitatively and quantitatively from the target data. Because of space limitations, the Editor required Farrell and Lewandowsky to drop the figure and the associated analysis from their original submission.

—*Editor*

REFERENCES

BADDELEY, A. D., & HITCH, G. J. (1977). Recency reexamined. In S. Dornic (Ed.), *Attention and performance VI* (pp. 647-667). Hillsdale, NJ: Erlbaum.

FARRELL, S., & LEWANDOWSKY, S. (2008). Empirical and theoretical limits on lag-recency in free recall. *Psychonomic Bulletin & Review*, **15**, 1236-1250.

GLANZER, M., KOPPENAAL, L., & NELSON, R. (1972). Effects of relations between words on short-term storage and long-term storage. *Journal of Verbal Learning & Verbal Behavior*, **11**, 403-416.

GLENBERG, A. M., BRADLEY, M. M., STEVENSON, J. A., KRAUS, T. A., TKACHUK, M. J., & GRETZ, A. L. (1980). A two-process account of

long-term serial position effects. *Journal of Experimental Psychology: Human Learning & Memory*, **6**, 355-369.

GOODWIN, J. (1976). Changes in primacy and recency with practice in single-trial free recall. *Journal of Verbal Learning & Verbal Behavior*, **15**, 119-132.

HOWARD, M. W. (1999). *Temporal context in human memory*. Unpublished doctoral dissertation, Brandeis University.

HOWARD, M. W., FOTEDAR, M. S., DATEY, A. V., & HASSELMO, M. E. (2005). The temporal context model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. *Psychological Review*, **112**, 75-116.

HOWARD, M. W., JING, B., RAO, V. A., PROVYN, J. P., & DATEY, A. V. (2009). Bridging the gap: Transitive associations between items presented in similar temporal contexts. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **35**, 391-407.

HOWARD, M. W., & KAHANA, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **25**, 923-941.

HOWARD, M. W., & KAHANA, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, **46**, 269-299.

HOWARD, M. W., KAHANA, M. J., & WINGFIELD, A. (2006). Aging and contextual binding: Modeling recency and lag recency effects with the temporal context model. *Psychonomic Bulletin & Review*, **13**, 439-445.

HOWARD, M. W., VENKATADASS, V. S., NORMAN, K. A., & KAHANA, M. J. (2007). Associative processes in immediate recency. *Memory & Cognition*, **35**, 1700-1711.

HOWARD, M. W., YOUKER, T. E., & VENKATADASS, V. S. (2008). The persistence of memory: Contiguity effects across hundreds of seconds. *Psychonomic Bulletin & Review*, **15**, 58-63.

KAHANA, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, **24**, 103-109.

KAHANA, M. J., HOWARD, M. [W.], & POLYN, S. (2008). Associative processes in episodic memory. In J. H. Byrne (Series Ed.) & H. L. Roediger III (Vol. Ed.), *Learning and memory: A comprehensive reference. Vol. 2: Cognitive psychology of memory* (pp. 476-490). Amsterdam: Elsevier.

KAHANA, M. J., HOWARD, M. W., ZAROMB, F., & WINGFIELD, A. (2002). Age dissociates recency and lag recency effects in free recall. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **28**, 530-540.

MURDOCK, B. B., JR. (1963). Short-term memory and paired-associate learning. *Journal of Verbal Learning & Verbal Behavior*, **2**, 320-328.

POLYN, S. M., NORMAN, K. A., & KAHANA, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, **116**, 129-156.

POSTMAN, L., & PHILLIPS, L. W. (1965). Short-term temporal changes in free recall. *Quarterly Journal of Experimental Psychology*, **17**, 132-138.

RAO, V. A., & HOWARD, M. W. (2008). Retrieved context and the discovery of semantic structure. In Y. S. J. C. Platt, D. Koller, & S. Roweis (Eds.), *Advances in neural information processing systems* (Vol. 20, pp. 1193-1200). Cambridge, MA: MIT Press.

ROMNEY, A. K., BREWER, D. D., & BATCHELDER, W. H. (1993). Predicting clustering from semantic structure. *Psychological Science*, **4**, 28-34.

SEDERBERG, P. B., HOWARD, M. W., & KAHANA, M. J. (2008). A context-based theory of recency and contiguity in free recall. *Psychological Review*, **115**, 893-912.

SHANKAR, K. H., JAGADISAN, U. K. K., & HOWARD, M. W. (in press). Sequential learning using temporal context. *Journal of Mathematical Psychology*.

TZENG, O. J. L. (1973). Positive recency in delayed free recall. *Journal of Verbal Learning & Verbal Behavior*, **12**, 436-439.

NOTES

1. In their analysis of individual data sets, Farrell and Lewandowsky (2008) evaluated nonmonotonicities in the lag-CRP by comparing the fits of monotonic and nonmonotonic functions with the lag-CRP. Whereas their nonmonotonic functions included parameters that produced a nonzero asymptote, their monotonic functions did not. Their analysis thus leaves open the question of whether their nonmonotonic functions provided superior fit because of their ability to capture the nonzero asymptote that is clearly evident in the data, or whether they actually detected nonmonotonicities in individual data sets. This statistical issue does not challenge their finding of an overall nonmonotonic lag-CRP for data collapsed across many free recall studies.

2. The primacy effect is also a prominent feature of free recall that we have modeled in recent versions of TCM (e.g., Polyn et al., 2009; Sederberg et al., 2008). The present analysis, however, focuses on recency, because the nonmonotonicity of the lag-CRP indicates a stronger bias in transitions to recency than to primacy items and because the original version of TCM analyzed by Farrell and Lewandowsky (2008) did not include a mechanism for explaining the primacy effect.

3. The TCM-A model of Sederberg et al. (2008) shows the same pattern of results at extreme lags across delay conditions using the published parameter settings.

4. Howard et al. (2008) conducted analyses on a surrogate data set to confirm that the boost in the lag-CRP was not an artifact of a persistent recency effect or, in fact, any other variation in encoding across the experimental session.

SUPPLEMENTAL MATERIALS

Results of an additional simulation based on the Howard and Kahana (2002) version of the TCM may be downloaded from http://pbr.psychonomic-journals.org/content/supplemental.

APPENDIX
Simulation Methods

In simulating free recall using TCM, we used a reduced-form version of the implementation described by Howard et al. (2006). To minimize the number of free parameters, we turned off the primacy gradient and we assumed that there was no retrieval of noise. Because they only fit lag-CRP functions in delayed recall, Howard et al. (2006) assumed that the retention interval was infinity (i.e., $\rho_{test} = 0$). Because we fit lag-CRPs in both immediate and delayed free recall, we adopted the assumption that $\rho_{test} = \rho_{study}$ (i.e., TCM_{evo}).

In TCM, the current state of context $\mathbf{t}_i$ is generated from the previous state of context $\mathbf{t}_{i-1}$ and the current input $\mathbf{t}_i^{IN}$ according to

$$\mathbf{t}_i = \rho_i \mathbf{t}_{i-1} + \beta \mathbf{t}_i^{IN}, \tag{A1}$$

where β is a free parameter and ρ_i is chosen such that the length of $\mathbf{t}_i$ is unity. Note that this implies that the rate of contextual drift depends on the amount and nature of the input vector $\mathbf{t}_i^{IN}$. We treated the asymptotic rate of drift

$$\rho \equiv \sqrt{1 - \beta^2}$$

as the parameter.

Here the Euclidean norm is used for $\mathbf{t}$, rather than the $\mathcal{L}^1$ norm used in some recent articles (Rao & Howard, 2008; Shankar, Jagadisan, & Howard, in press).

In TCM, items are encoded in their temporal contexts by means of a Hebbian outer product matrix $\mathbf{M}$, which is updated using

$$\Delta \mathbf{M} = \mathbf{f}_i \mathbf{t}'_{i-1}, \tag{A2}$$

where $\mathbf{f}_i$ is the item representation of item i and the prime denotes the transpose. That is, we associated item i to the state of context that preceded it, $\mathbf{t}_{i-1}$. The use of $\mathbf{t}_{i-1}$ in this equation, rather than $\mathbf{t}_i$, is consistent with recent treatments of TCM (e.g., Howard, Jing, Rao, Provyn, & Datey, 2009; Howard et al., 2006; Rao & Howard, 2008; Sederberg et al., 2008) but different from earlier work (Howard, 1999; Howard et al., 2005; Howard & Kahana, 2002). The use of $\mathbf{t}_{i-1}$ is an important development of the model. Because $\mathbf{t}_i$ includes $\mathbf{t}_i^{IN}$ (see Equation A1), learning $\mathbf{t}_i$ in the matrix $\mathbf{M}$ places constraints on the long-term behavior of the model.

In TCM, context cues for retrieval of items. Each item is activated to some extent by the state of context used as a cue and then the items compete to be retrieved. Given a contextual cue $\mathbf{t}$, item i is activated by

$$a_i \equiv \mathbf{f}_i' \mathbf{M} \mathbf{t}. \tag{A3}$$

The distractors for the retention interval in the case of delayed free recall and continual-distractor free recall were implemented by adding a vector orthogonal to all preceding states of context weighted by β and multiplying the preceding context vector by ρ_D. This is conceptually identical to all previous implementations of TCM, although the parameterization is a bit different.

In TCM, repetition of a study item during retrieval causes a change of the contextual state, which in turn leads to a contiguity effect. This contextual state recovered by the remembered item contains two components. As in all previous treatments of TCM, these two components are differentially responsible for the asymmetry observed in the lag-CRP. Here the parameter γ was used to weight the degree of contextual retrieval upon repetition of an item. Although Farrell and Lewandowsky (2008) did not do so, allowing the proportion of the two components of retrieved context to vary has long been a feature of TCM (Howard et al., 2005; Howard et al., 2006; Sederberg et al., 2008). The parameterization of the weighting here is somewhat different from previous work. Here, the consistent part of the retrieved context vector is an orthogonal vector $\mathbf{c}$ chosen separately for each item in the list. The vector $\mathbf{c}$ for each item remains fixed throughout. Recovery of $\mathbf{c}$ leads to a strong forward association and is closely analogous to the preexperimental component of previous implementations of TCM (Howard et al., 2005; Howard & Kahana, 2002; Howard et al., 2006). In addition to the fixed component, each item is also associated with an $\mathbf{h}$ vector. The $\mathbf{h}$ vector recovers a contextual state during list presentation and thus serves the same role as the "newly learned" context in previous treatments of TCM (Howard et al., 2005; Howard & Kahana, 2002; Howard et al., 2006). The $\mathbf{h}$ vector provides a symmetric retrieval cue for the neighbors of the recalled item. Following recent treatments of TCM (Howard et al., 2009; Rao & Howard, 2008), if item A is presented (or recalled) at time step i, then the input pattern at time step i is given by

$$\mathbf{t}_i^{IN} \propto (1 - \gamma) \mathbf{c}_A + \gamma \mathbf{h}_A, \tag{A4}$$

where the proportionality symbol indicates that $\mathbf{t}_i^{IN}$ is normalized to be of unit length before entering into the evolution equation above. Each item's $\mathbf{h}$ vector is initialized to zero and then updated according to

$$\Delta \mathbf{h}_A = \mathbf{t}_{i-1}. \tag{A5}$$

In TCM, contiguity effects are a consequence of the similarity that the $\mathbf{c}$ and $\mathbf{h}$ vectors caused by repeated (recalled) items have with the encoding contexts of items presented at similar times. Note that because $\mathbf{c}_A$ is orthonormal to $\mathbf{c}_B$ for other items, the function of $\mathbf{c}_A$ is identical to the preexperimental context retrieved on

---

APPENDIX (Continued)

---

the initial presentation of an item in previous implementations of TCM. Similarly, $\mathbf{h}_A$ is identical to the newly learned context described in previous versions of TCM, post-2006. That is, one could rewrite Equation A4 as

$$\mathbf{t}_r^{IN} \propto (1 - \gamma)\mathbf{t}_i^{IN} + \gamma \mathbf{t}_{i-1} \tag{A6}$$

(compare with Howard et al., 2006), and the predictions of the model would be identical as long as items are repeated at most once. The use of $\mathbf{c}$ and $\mathbf{h}$ is differentiable from the use of $\mathbf{t}^{IN}$ and $\mathbf{t}$ when items are repeated multiple times (Howard et al., 2009; Rao & Howard, 2008), where it is able to avoid serious problems with the old formulation (see the General Discussion in Howard et al., 2005).

In some earlier treatments of TCM, we have parameterized the proportion of the two components according to a ratio of the two coefficients (e.g., Howard et al., 2005; Howard et al., 2006; see also Farrell & Lewandowsky, 2008). Changing the definition of γ in this way does not affect the output of the model, although it does affect the value of the number. Rather than taking on values from zero to infinity, γ is restricted to the range from zero to one in the current treatment.

In TCM, after a word is recalled, it provides input to the context vector, and the resulting context vector is used as a cue for recall of the other items. Each item is activated to some extent. After the activations have been calculated, it is necessary to translate this number into a probability of recall in order to compare with behavioral data. As in early treatments of TCM (Howard & Kahana, 2002; Howard et al., 2006), but not as in more recent studies that use accumulators to model the retrieval process (Polyn et al., 2009; Sederberg et al., 2008) and as in Farrell and Lewandowsky (2008), we used the Luce choice rule to map the activation of an item a_i onto the probability of recalling that item:

$$P_R(i) = \frac{\exp a_i / \tau}{\sum_j \exp a_j / \tau}. \tag{A7}$$

The parameter τ controls the sensitivity of this mapping.

Predicted lag-CRP values were generated by first calculating the PFR for each serial position. The probability of second recall for each other serial position conditionalized on the first recall was then calculated for each item. These values were aggregated by lag, weighted by the PFR of the first item.