Running head: LIMITS ON LAG-RECENCY SUPPLEMENT

Supplementary material for

"Empirical and theoretical limits on lag recency in free recall"

by Farrell and Lewandowsky

Simon Farrell

University of Bristol

Stephan Lewandowsky

University of Western Australia

March, 2008

Dr. Simon Farrell

Department of Experimental Psychology

University of Bristol

12a Priory Road

Clifton, Bristol BS8 1TU, UK

e-mail: Simon.Farrell@bristol.ac.uk

fax: +44 (0)117 928 8588

Abstract

This supplementary document reports detailed results of the simulations conducted by Farrell and Lewandowsky (2008) in their exploration of the behavior of the Temporal Context Model (TCM; Howard & Kahana, 2002). This document is not a stand-alone paper and hence is not intended to be self-explanatory; please refer to Farrell and Lewandowsky (2008) for further explanation of the principles of the analysis and modelling.

Supplementary material for

"Empirical and theoretical limits on lag recency in free recall"

by Farrell and Lewandowsky

Farrell and Lewandowsky (2008) have presented a re-analysis of published data that have previously been used to explore the characteristics of recency and "lag-recency" in free recall data. Farrell and Lewandowsky (2008) noted that conditional transitions in free recall (i.e., lag-CRP functions) have usually only been analyzed for relatively small lags, and showed that lag-CRP functions (described in detail in their article) often display non-monotonicities when a larger range of lags is examined. This indicates that conventional primacy and recency can bleed into estimates of lag-recency in free recall. This result is potentially problematic for the Temporal Context Model (TCM: Howard & Kahana, 2002), for which recency and lag-recency are usually taken as key pieces of evidence.

To determine whether TCM can qualitatively and quantitatively account for the data when examined over longer lags, and to provide a greater understanding of the behavior of the model, Farrell and Lewandowsky (2008) reported summaries of fits of two versions of the model to data from a number of free recall experiments. In the following, we report details of the data analysis and statistical characeterization of the data, and details of simulations of TCM, that could not be included in the main paper.

Procedures for data analysis and model fitting

Construction of FRP's and lag-CRP's

The FRP and lag-CRP functions were calculated following descriptions in published reports (e.g., Howard & Kahana, 1999; Kahana, 1996). For the FRP functions, for each participant a record was kept of (a) the number of legitimate first responses, and (b) the

frequency with which the item from each serial position was produced as the first recall. The FRP function for a participant is simply (b) divided by (a), with group averages being calculated as the average FRP for each serial position.

Following Howard and Kahana (1999) and others, the lag-CRP function was generally calculated without respect to the serial position of the first-recalled item (the exception is the simulation reported in the section "Recency in TCM$_{evo}$" in the main paper). For each participant p, two vectors were updated across trials. Both vectors contained $2L-1$ elements, where L is the list length, with each element representing a lag of a particular displacement, viz. from $-(L-1)$ to $+(L-1)$. The first vector, $\mathbf{r}_p$, kept track of the frequency of displacements at the second output position. For each correct response at that output position, the lag was calculated, and the cell corresponding to that lag incremented by one. The second vector, $\mathbf{d}_p$, was a denominator vector and kept track of all lags that could have been produced at the second output position. The set of possible lags was defined by (a) the possible lags given the item recalled first (e.g., if the second item from a 12-item list was recalled first, then the possible lags for the second output position ranged from -1 to $+10$); and (b) following Howard and Kahana (1999), it was assumed that already recalled items could not be recalled again; thus, the denominator for lag $= 0$ was never incremented. All possible lags in $\mathbf{d}_p$ were also incremented by one whenever a cell in $\mathbf{r}_p$ was incremented. The lag-CRP for a participant was then obtained by dividing $\mathbf{r}_p$ by $\mathbf{d}_p$ in an element-wise fashion, and the average of individual lag-CRP's across participants yielded the mean lag-CRP functions shown in the figures. In cases in which a participant could not have produced responses at a lag of a particular distance, their data was not included in the average for that lag. In consequence, the lag-CRP functions can occasionally display slight discontinuities when different subsets of participants contribute to different lags.

Fitting of descriptive models

The descriptive models were fit to positive and negative lags separately by minimizing the negative log-likelihood discrepancy between the data and the model for each participant. Each model predicted a probability $f(l)$ of making a transition of a particular lag l. For computational convenience, l was a subscript that ran from 1 (or -1) to its maximum (minimum) value even when the smallest lag was $\pm$ 2. Those predictions were then corrected for the number of opportunities with which a particular lag l could be produced. That is, because the empirical lag-CRP's were obtained by dividing $\mathbf{r}_p$ by $\mathbf{d}_p$, the predicted probabilities for the descriptive model, $f(l)$, were multiplied (element-wise) by $\mathbf{d}_p$ and then normalized to sum to unity, to give $p(l)$. The values of $p(l)$ were thus commensurate with the raw number of observed transitions, permitting computation of the likelihood by:

$$lnL = \sum_{l=\,\pm\,1}^{\pm(L-1)} r_p(l)\log(p(l)), \tag{1}$$

where $r_p(l)$ was the observed number of transitions for a given lag l and $\log(p(l))$ the corresponding prediction of the descriptive function.

For graphical presentation of the predictions, values of $p(l)$ were multiplied by the total number of responses (in the forward or backward direction, depending on the analysis) to obtain expected frequencies, $\mathbf{m}_p$, equivalent to $\mathbf{r}_p$ for the data. Exactly as with the data, the final predicted lag-CRP's were then calculated as the element-wise division of $\mathbf{m}_p$ and $\mathbf{d}_p$, before being averaged across participants to yield the mean predictions shown in the figures.

Fitting of TCM_{pub} and TCM_{evo}

TCM was fit to the data using maximum-likelihood estimation. Following Howard and Kahana (2002), and paralleling our re-analysis, we fit only the first two output

positions. The first output position determines the FRP's, and the transition between the two output positions determines the lag-CRP. Data were fit on a response-by-response basis. On each trial, a likelihood was calculated for the first output position directly from the probability distribution across items given by Equation 5 using the activations from Equation 12. A likelihood for the second output position was calculated by feeding the activations from Equation 13 into Equation 5 based on the previous recall on that trial by that participant.

To calculate the FRP's and lag-CRP's, vectors that tracked first and second predicted recalls were incremented across trials in exactly the same manner as for the observed recalls. For the first output position, all elements in the numerator were incremented by an amount given by Equation 5 using the activations from Equation 12, with the numerator being identical to that used to calculate empirical FRP's. For the lag-CRP function, the predicted response probabilities were calculated using the activations from Equation 13 (TCM_{pub}) or Equation 14 (TCM_{evo}), where the item i used in these equations was the item actually recalled first by the participant at OP 1, and where response repetitions were disallowed. The predicted response probabilities were then re-expressed as a function of lag (given i) and added to vector $\mathbf{m}_p$. The model lag-CRP's were then calculated as the element-wise division of $\mathbf{m}_p$ and $\mathbf{d}_p$, and then averaged across participants.

One benefit of using the log-likelihood method is that each individual response is given an equal weighting in contributing to the goodness-of-fit. This is important for assessing the contribution and degree of non-monotonicity, given the smaller number of observations at extreme lags. It may be argued that if the observed non-monotonicity is due to the contribution of only a few data points, then it should possibly be disregarded for the purposes of model evaluation. Because the log-likelihood estimation treats each responses equally, the data at more extreme lags automatically have less influence on the

summed log-likelihood, there being fewer of them, thus countering an argument based on the weighting of extreme lags.

Formal description of the temporal context model

TCM assumes that, at any point in time i, an item will be represented by a vector $\mathbf{f}_i$ across vector space F, and that the temporal context present at time i is represented by $\mathbf{t}_i$ across space T. These spaces, which we here treat as layers in a connectionist network (Howard & Kahana, 2002), are connected by two matrices of unidirectional weights: one matrix $\mathbf{M}^{TF}$ maps from T to F, while another matrix $\mathbf{M}^{FT}$ stores associations in the other direction, such that items in F can be used as cues to temporal contexts in T.

List learning is primarily accomplished by adjusting the weights between T and F such that states across T can later be used as a cue to retrieve items in F. TCM assumes Hebbian learning, such that the change in the weights in $\mathbf{M}^{TF}$ on presentation of an item $\mathbf{f}_i$ at time i is given by:

$$\Delta\mathbf{M}^{TF} = \mathbf{f}_i\mathbf{t}_i', \tag{2}$$

where $\mathbf{t}_i$ is the temporal context present at time i. An item can then be retrieved at time k by using the temporal context prevailing at time k, $\mathbf{t}_k$, as a cue to $\mathbf{M}^{TF}$:

$$\mathbf{f}^{IN} = \mathbf{M}^{TF}\mathbf{t}_k. \tag{3}$$

Under the assumption that all items on the space F are orthonormal, Howard and Kahana (2002) show that the "activation" a_i of item $\mathbf{f}_i$ when the temporal context at j is used as a cue is given by:

$$a_i = \sum_{\mathbf{f}_k=\mathbf{f}_i} \mathbf{t}_j \bullet \mathbf{t}_k, \tag{4}$$

where $\bullet$ is the inner (or dot) product, the activation a_i is defined by $\mathbf{f}^{IN} \bullet \mathbf{f}_i$, and k indexes the previous presentations of i.

Following published applications of TCM (Howard & Kahana, 2002; Howard, 2004; Howard, Fotedar, Datey, & Hasselmo, 2005), the activations a_i for all i are converted into recall probabilities using the Luce-Shepard choice function (Luce, 1963; Shepard, 1957); under the assumption of orthonormality, this reduces to:

$$P(\mathbf{f}_i|\mathbf{f}^{IN}) = \frac{\exp\left(\frac{2a_i}{\tau}\right)}{\sum_j \exp\left(\frac{2a_j}{\tau}\right)}, \tag{5}$$

where τ is a free parameter (Howard & Kahana, 2002).

Central to TCM is the specification of how the context space, T, is updated. The updated state of T at time i, $\mathbf{t}_i$, is the weighted sum of the temporal context at time step $i-1$ (called the *carry-over* context here and in the main article) and the context retrieved by the item just presented or recalled, $\mathbf{t}_i^{IN}$ (called the *retrieved* context):

$$\mathbf{t}_i = \rho \mathbf{t}_{i-1} + \beta \mathbf{t}_i^{IN}, \tag{6}$$

where ρ_i is given by:

$$\rho_i = \sqrt{1 + \beta^2[(\mathbf{t}_{i-1} \bullet \mathbf{t}_i^{IN})^2 - 1]} - \beta(\mathbf{t}_{i-1} \bullet \mathbf{t}_i^{IN}). \tag{7}$$

The retrieved component $\mathbf{t}_i^{IN}$ in Equation 6 is obtained by cueing the weight matrix mapping items to contexts, $\mathbf{M}^{FT}$:

$$\mathbf{t}_i^{IN} = \mathbf{M}_{i-1}^{FT} \mathbf{f}_i. \tag{8}$$

Two versions of Equation 8 have been presented by Howard and colleagues; these differ in whether $\mathbf{M}_{i-1}^{FT}$ or $\mathbf{M}_i^{FT}$ is probed by item $\mathbf{f}_i$ to obtain a retrieved context (Howard & Kahana, 2002; Howard et al., 2005). We assume these equations are identical in

intention and differ according to how the numerous events occurring at presentation of an item are broken into time steps.

The matrix $\mathbf{M}^{FT}$ is assumed to be updated on presentation of item $\mathbf{f}_i$, such that the context retrieved on a subsequent encounter of that item is a linear combination of the pre-experimental context associated with the item ($\mathbf{t}_i^{IN}$), and the context present on the prior encounter of the item ($\mathbf{t}_i$):

$$\mathbf{t}_r^{IN} = A_i \mathbf{t}_i^{IN} + B_i \mathbf{t}_i. \tag{9}$$

The weights A_i and B_i are chosen such that the vector norm of $\mathbf{t}_r^{IN}$ is 1, and are constrained such that $A_i = \gamma B_i$. This gives the solution (Howard & Kahana, 2002):

$$B_i = \frac{1}{\gamma^2 + 2\gamma(\mathbf{t}^{IN} \bullet \mathbf{t}_i) + 1}, \tag{10}$$

allowing a useful re-expression of Equation 9 as:

$$\mathbf{t}_r^{IN} = \frac{\gamma \mathbf{t}_i^{IN} + \mathbf{t}_i}{\gamma^2 + 2\gamma(\mathbf{t}^{IN} \bullet \mathbf{t}_i) + 1}. \tag{11}$$

Learning in the weight matrix $\mathbf{M}^{FT}$ is accomplished by a rule more complicated than simple Hebbian learning, and is detailed in Howard and Kahana (2002). For the present simulations, we assume such learning has taken place and directly apply Equations 9 to 11.

Standard implementation: TCM$_{pub}$

In the standard implementation that underlies all published work to date (called TCM$_{pub}$ here and in the main article), it is assumed that the temporal context immediately prior to recall is carried over into the recall to provide a cue for the first retrieval. With the assumption that lists do not contain repetitions, Equation 7 reduces to $\rho = \sqrt{1 - \beta^2}$ (Howard & Kahana, 2002), and activations of list items are directly given by:

$$a_i = \rho^{((L-i)(d_i+1)+d_r)}, \tag{12}$$

where L is the list length, d_i is the effective duration of the distractor activity, if any, following each list item (assumed to be constant across the list), and d_r is the effective duration of distractor activity intervening between presentation of the last list item and the cue to recall (cf. Equation 12 in Howard, 2004).

Lag-recency has been examined in TCM$_{pub}$ by studying transitions between pairs of items under the assumption that an infinite retention interval intervenes prior to recall of the first item in the pair (even though modeling has consistently focused on the first two output positions, in which case this assumption is not necessarily appropriate).

Assuming an infinite retention interval means that only the second term of Equation 6 will have any role distinguishing between items, as $\mathbf{t}_i$ will be orthogonal to list contexts. Accordingly, the activation of all items j following recall of item i ($j \neq i$) is given by:

$$a_j = \begin{cases} A\beta^2\rho^{(1+d_i)(j-i)} + B\beta\rho^{(1+d_i)|j-i|}, & j > i, \\ B\beta\rho^{(1+d_i)|j-i|}, & j < i. \end{cases} \tag{13}$$

(see Equations 30 and 31 in Howard, 2004).

Continuous evolution version: TCM$_{evo}$

In the continuously evolving version (TCM$_{evo}$) that was developed by the present authors for the purposes of this article, it is assumed that the temporal context is carried over between items following Equation 6. This does not allow for a solution as in Equation 13, because we can no longer assume that there have been no repetitions: items recalled prior to recall of item j will have influenced the temporal context on their initial presentation, as well as when they were recalled.

First-recall probabilities are again given by Equation 12. All following recall probabilities can then be obtained by updating the activations with information given by the retrieved context:

$$\begin{aligned} a_j &= \mathbf{t}_{T+1} \bullet \mathbf{t}_j \\ &= (\rho_T \mathbf{t}_T + \beta \mathbf{t}_r^{IN}) \bullet \mathbf{t}_j \\ &= \rho_T \mathbf{t}_T \bullet \mathbf{t}_j + A\beta \mathbf{t}_i^{IN} \bullet \mathbf{t}_j + B\beta \mathbf{t}_i \bullet \mathbf{t}_j. \end{aligned} \tag{14}$$

The assumption that repetitions of items are not separated by an infinite time interval requires the use of Equation 7 to calculate a value of ρ for each time step at retrieval. Accordingly, we need to calculate a value for $\mathbf{t_i^{IN}} \bullet \mathbf{t_{T-1}}$. Using Equation 9, this dot product is given by:

$$\mathbf{t}_R^{IN} \bullet \mathbf{t}_{T-1} = A_i \mathbf{t}_i^{IN} \bullet \mathbf{t}_{T-1} + B_i \mathbf{t}_i \bullet \mathbf{t}_{T-1}, \tag{15}$$

where the term $\mathbf{t_i} \bullet \mathbf{t_{T-1}}$ is given by a_i from the previous retrieval step, and $\mathbf{t_i^{IN}} \bullet \mathbf{t_{T-1}}$ is given by re-expressing Equation 6:

$$\mathbf{t}_i^{IN} \bullet \mathbf{t}_{T-1} = \frac{1}{\beta} \left(\mathbf{t}_i \bullet \mathbf{t}_{T-1} - \rho_l \mathbf{t}_{i-1} \bullet \mathbf{t}_{T-1} \right), \tag{16}$$

where ρ_l is the constant value of ρ applying throughout list presentation. As for Equation 15, $\mathbf{t_i} \bullet \mathbf{t_{T-1}}$ is given by a_i from the previous retrieval time step. In the simulations, to obtain $\mathbf{t_{i-1}} \bullet \mathbf{t_{T-1}}$ we calculated a_{i-1} at the first retrieval attempt. These values are already given by the above equations for all items except the first; calculating and updating a_{i-1} explicitly throughout recall allowed generalization to cases with distractor intervals between items and following the list.

Fits to individual experiments

Standard implementation: TCM$_{pub}$

Figures 1 to 7 show the predictions of TCM$_{pub}$ (the published version of TCM) for individual data sets, averaging across the predictions for each participant under their best-fitting parameters. Although TCM$_{pub}$ approximated the empirical first-recall probabilities (FRP's) quite well in some cases (e.g., Experiments 1 and 2 of Howard and Kahana, 1999, in Figures 1 and 2; the data of Howard, Venkatadass, Norman, & Kahana, in press, in Figure 7), in other cases the model over-predicted the recency in the FRP's (Figures 4 and 5, which respectively show the fits for Murdock & Okada, 1970, and Murdock, 1962).

Concerning lag-CRPs, although the fit of the model was satisfactory in some cases (e.g., the delayed recall condition in Figure 1), in most cases the model's predictions did not accord with the empirical results. TCM$_{pub}$ sometimes under-predicted the occurrence of +1 transitions; this can be seen in the immediate condition of Figure 1, the 16 s ISI condition of Figure 3, in Figure 4, for all conditions in Figure 6, and in Figure 7. The model also failed to capture the non-monotonicity evidenced in the data (as re-analyzed by Farrell & Lewandowsky, 2008); obvious deviations are apparent for the shorter intra-list distractor durations in Figure 3, in several of the conditions in Figure 6, and in Figure 7.

Applying TCM with continuously evolving context to data

The predictions of TCM$_{evo}$, in which temporal context is assumed to continuously evolve across list presentation and recall as described earlier, are shown in Figures 8 to 14. Inspection of Figures 8 and 9 shows that TCM$_{evo}$ qualitatively captured the recency in the FRP functions of the data of Howard and Kahana (1999), although the model failed to provide a good quantitative account of some of these functions. With the exception of the delayed-recall condition of Experiment 1, Figures 8 and 10 show that the model did

not quantitatively handle the empirical lag-CRP functions. Although TCM$_{evo}$ was able to produce the non-monotonicity in the lag-CRP functions (in some cases excessively so), this often occurred at the cost of over-predicting these responses generally, and under-predicting the standard lag-recency effect (that is, transitions of small lag, particularly immediate $\pm$ 1 transitions).

This behavior was also generally apparent for the fits to the data from Murdock and Okada (1970) and Murdock (1962). Figures 11 and 13 show that the model produced the appropriate non-monotonicity in the lag-CRP functions; however, the frequency of more extreme transitions was generally over-predicted, with forward transitions of lag +1 being severely under-predicted. Additionally, TCM$_{evo}$ over-predicted the extent of recency in the FRP functions for these two studies (Figures 11 and 12).

Finally, Figure 14 shows the fits of TCM$_{evo}$ to the data from Howard et al. (in press). The model captured the qualitative effects of recency in the FRP's and non-monotonicity in the lag-CRPs, but it again under-predicted the frequency of short transitions and over-predicted longer transitions.

Recency in TCM$_{evo}$

Farrell and Lewandowsky (2008) also showed that the lag-recency produced by TCM$_{evo}$ in immediate recall is an artifactual consequence of the manner in which the data are typically analyzed (i.e., by averaging across serial positions). Farrell and Lewandowsky showed that when TCM$_{evo}$'s behavior is instead considered at the level of individual serial positions (i.e., by obtaining separate lag-CRP's for each serial position of the first-recalled item), the model shows a complete absence of lag-recency and is solely characterized by recency. That is, regardless of which item was just recalled, an item from the terminal list positions tends to be recalled next. In consequence, TCM$_{evo}$'s lag-CRP functions were found to be monotonically increasing rather than decreasing.

Although Farrell and Lewandowsky (2008) demonstrated this behavior under maximum likelihood parameter estimates, those estimates were obtained by fitting the first-recall probabilities and lag-CRP function simultaneously. In consequence, the monotonically increasing lag-CRP functions may have arisen from the requirement to satisfy both constraints simultaneously, with first-recall probabilities winning out over the lag-CRP functions. To explore this possibility, an additional simulation was run in which TCM_{evo} was fit to the data of Howard et al. (in press) only for output position two. This meant that the model was not required to account for the first-recall probabilities, and was only fit to the data entering into construction of the lag-CRP function. The resulting parameter estimates ($\beta = 0.50$, $\tau = 0.17$, with a maximum $\ln L$ of -13237) were similar to those obtained when fitting the first two output positions together. Accordingly, Figure 15 shows that even when seeking to account for the lag-CRP's alone, without regard to first-recall probabilities, the TCM_{evo} produced almost identical behavior to that demonstrated by Farrell and Lewandowsky (2008) in their fit to the first two output positions.

References

Farrell, S., & Lewandowsky, S. (2008). Empirical and theoretical limits on lag recency in free recall. *Manuscript submitted for publication*.

Howard, M. W. (2004). Scaling behavior in the temporal context model. *Journal of Mathematical Psychology*, *48*, 230-238.

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., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and 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., Venkatadass, V., Norman, K. A., & Kahana, M. J. (in press). Associative processes in immediate recency. *Memory & Cognition*.

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

Luce, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1, p. 103-189). New York: Wiley.

Murdock, B. B. (1962). Direction of recall in short-term memory. *Journal of Verbal Learning and Verbal Behavior*, *1*, 119-124.

Murdock, B. B., & Okada, R. (1970). Interresponse times in single-trial free recall. *Journal of Experimental Psychology*, *86*, 263–267.

Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, *22*, 325-345.

Author Note

Collaboration on this project was assisted by an Australian Research Council International Linkage Grant to Stephan Lewandowsky, Gordon Brown, and Simon Farrell. The second author was supported by an Australian Professorial Fellowship from the Australian Research Council. We thank Marc Howard for his clarification of some details of the operation of the model. Correspondence should be addressed to Simon Farrell, Department of Psychology, University of Bristol, 12a Priory Road, Clifton, Bristol BS8 1TU, UK; e-mail: Simon.Farrell@bristol.ac.uk.

Figure Captions

Figure 1. TCM_{pub} maximum likelihood fits (connected circles; crosses are data) for Experiment 1 of Howard and Kahana (1999), from maximum likelihood parameter estimates using the standard TCM implementation. The top row shows the predicted FRP functions; the bottom row shows predicted lag-CRP functions (2nd output position only). The left column shows predictions for immediate recall, while the right column shows predictions from delayed recall.

Figure 2. FRP functions predicted by TCM_{pub} (connected circles; crosses are data) for various durations of intra-list distractor activity (ISI). The panels show the ISI=0 (top-left), 2.5s (top right), 8s (bottom left) and 16s (bottom right) conditions in Experiment 2 of Howard and Kahana (1999).

Figure 3. Lag-CRP functions predicted by TCM_{pub} (connected circles; crosses are data) for the ISI=0 (top-left), 2.5s (top right), 8s (bottom left) and 16s (bottom right) conditions in Experiment 2 of Howard and Kahana (1999).

Figure 4. FRP (left panel) and lag-CRP (right panel) functions predicted by TCM_{pub} (connected circles; crosses are data) for the free recall data of Murdock and Okada (1970).

Figure 5. FRP functions predicted by TCM_{pub} (connected circles; crosses are data) for the six conditions of Murdock (1962): 10 item lists, 2s per item (top-left); 15 item lists, 2s per item (top-right); 20 item lists, 1s per item (middle-left); 20 item lists, 2s per item (middle-right); 30 item lists, 1s per item (bottom-left); 40 item lists, 1s per item (bottom-right).

Figure 6. Lag-CRP functions predicted by TCM_{pub} (connected circles; crosses are data) for the six conditions of Murdock (1962): 10 item lists, 2s per item (top-left); 15 item lists,

2s per item (top-right); 20 item lists, 1s per item (middle-left); 20 item lists, 2s per item (middle-right); 30 item lists, 1s per item (bottom-left); 40 item lists, 1s per item (bottom-right).

Figure 7. FRP (left panel) and lag-CRP (right panel) functions predicted by TCM_{pub} (connected circles; crosses are data) for the free recall data of (Howard et al., in press).

Figure 8. TCM_{evo} maximum likelihood fits (connected circles; crosses are data) for Experiment 1 of Howard and Kahana (1999). The top row shows the predicted FRP functions; the bottom row shows predicted CRP functions (2nd output position only). The left column shows predictions for immediate recall, while the right column shows predictions from delayed recall.

Figure 9. FRP functions predicted by TCM_{evo} (connected circles; crosses are data) for the ISI=0 (top-left), 2.5s (top right), 8s (bottom left) and 16s (bottom right) conditions in Experiment 2 of Howard and Kahana (1999).

Figure 10. Lag-CRP functions predicted by TCM_{evo} (connected circles; crosses are data) for the ISI=0 (top-left), 2.5s (top right), 8s (bottom left) and 16s (bottom right) conditions in Experiment 2 of Howard and Kahana (1999).

Figure 11. FRP (left panel) and lag-CRP (right panel) functions predicted by TCM_{evo} (connected circles; crosses are data) for the free recall data of Murdock and Okada (1970).

Figure 12. FRP functions predicted by TCM_{evo} (connected circles; crosses are data) for the six conditions of Murdock (1962): 10 item lists, 2s per item (top-left); 15 item lists, 2s per item (top-right); 20 item lists, 1s per item (middle-left); 20 item lists, 2s per item (middle-right); 30 item lists, 1s per item (bottom-left); 40 item lists, 1s per item (bottom-right).

Figure 13. Lag-CRP functions predicted by TCM_{evo} (connected circles; crosses are data) for the six conditions of Murdock (1962): 10 item lists, 2s per item (top-left); 15 item lists, 2s per item (top-right); 20 item lists, 1s per item (middle-left); 20 item lists, 2s per item (middle-right); 30 item lists, 1s per item (bottom-left); 40 item lists, 1s per item (bottom-right).

Figure 14. FRP (left panel) and lag-CRP (right panel) functions predicted by TCM_{evo} (connected circles; crosses are data) for the free recall data of Howard et al. (in press).

Figure 15. Lag-CRP function for TCM_{evo} for the fit to the data of Howard et al. (in press), broken down by the serial position of the first recalled item. Only responses at the second output position were considered when estimating parameters. The darkest line represents the case in which the last item was recalled first (and is only present for negative lags because no forward transitions are possible after recall of the final item), and successively lighter lines represent successive serial positions back in the list of the first-recalled item (the single data point represents the case in which the penultimate list item is recalled first, for which the only possible forward transition is +1).

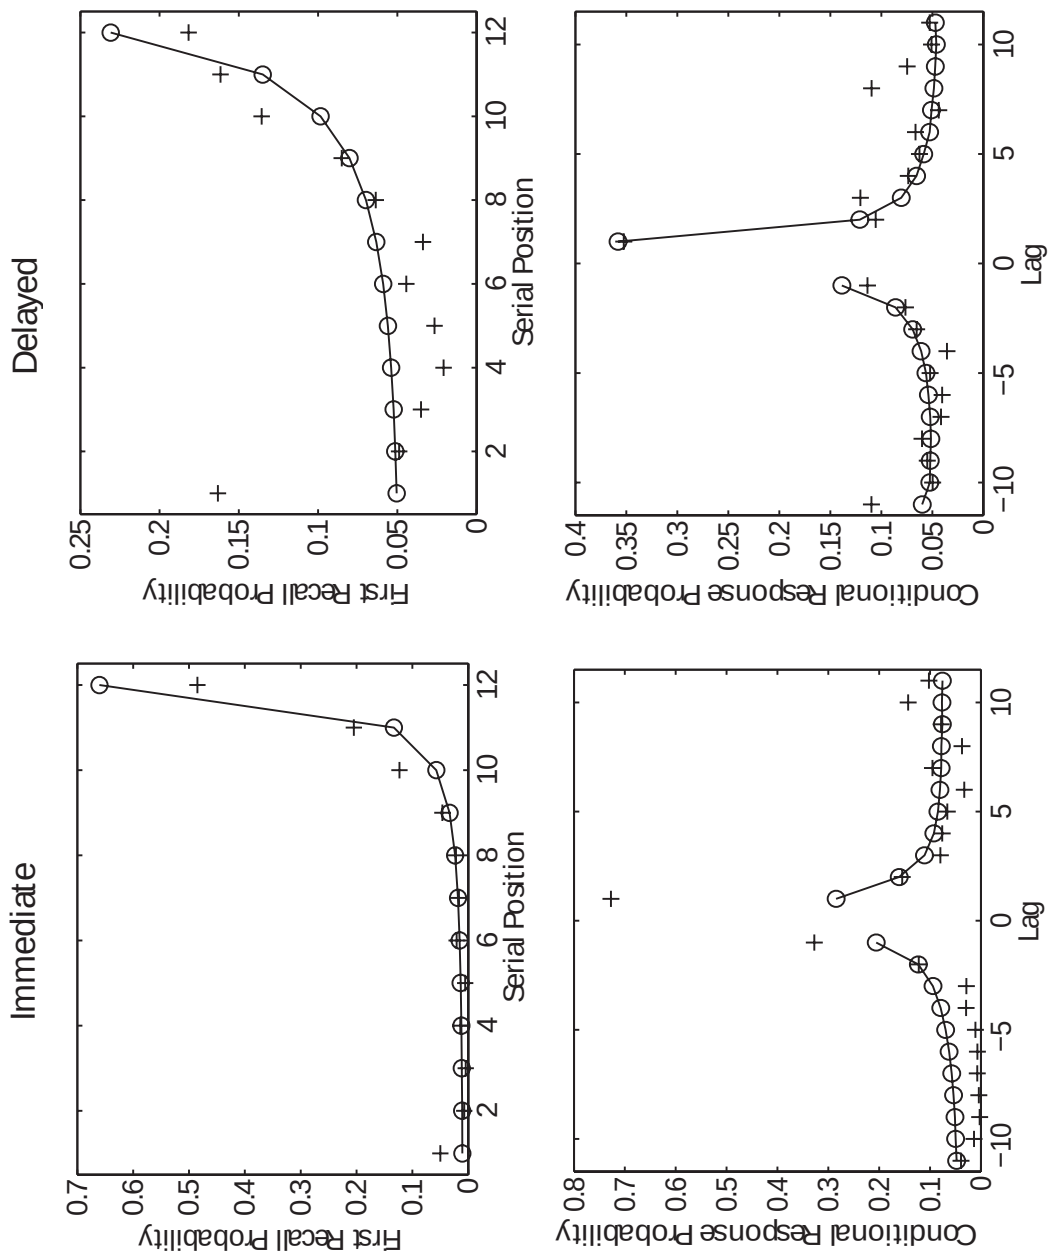
Immediate
Delayed
First Recall Probability
Serial Position
Conditional Response Probability
Lag

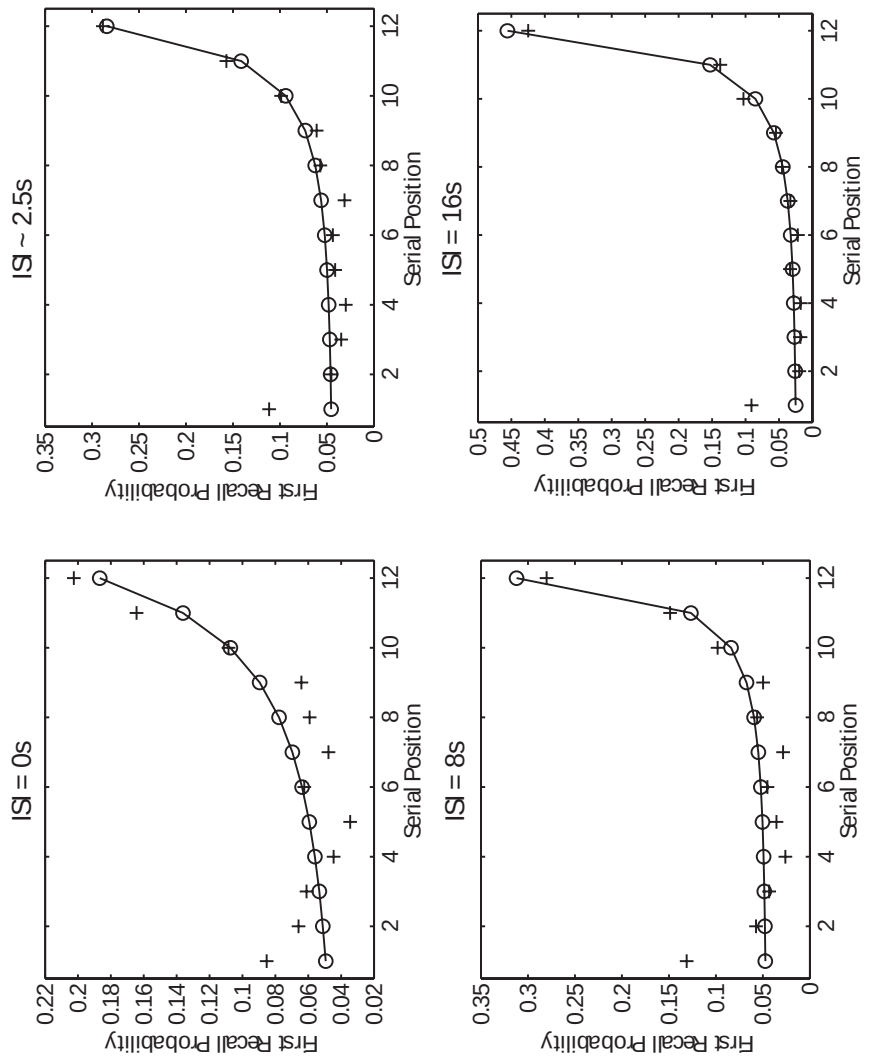
ISI = 0s
ISI ~ 2.5s
ISI = 8s
ISI = 16s
First Recall Probability
Serial Position

Limits on lag-recency supplement, Figure 3

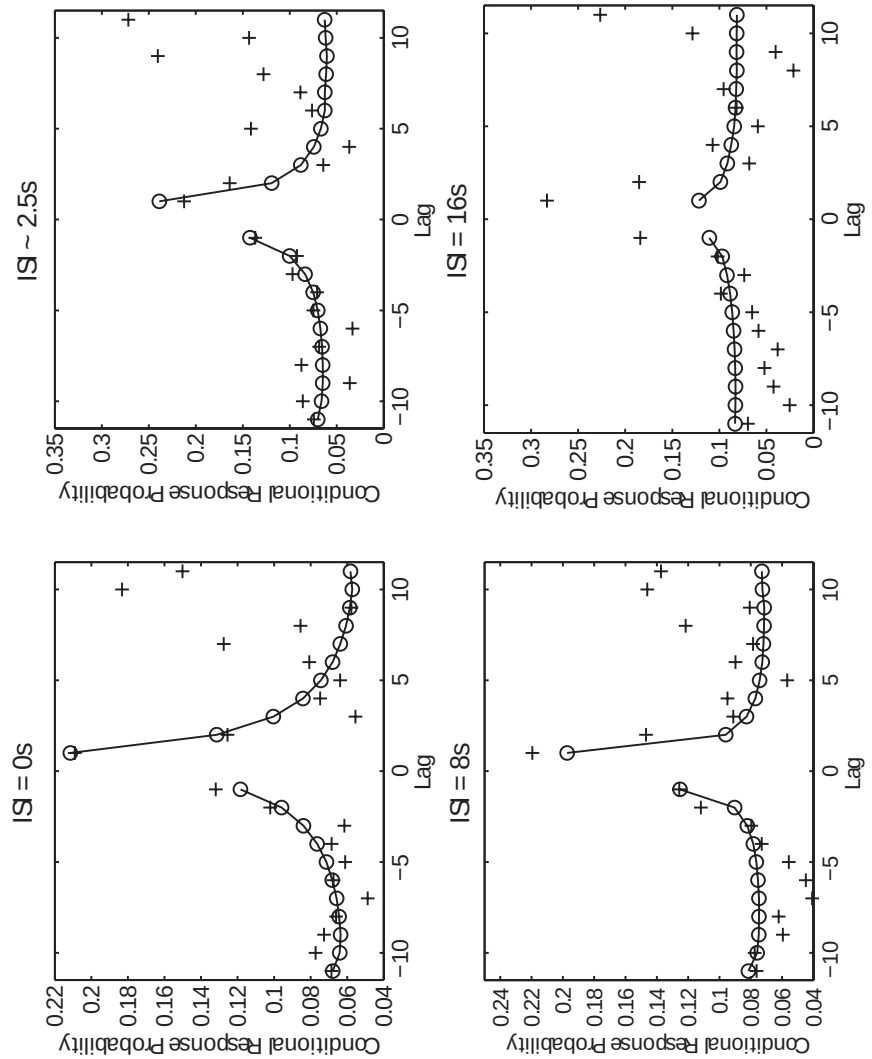

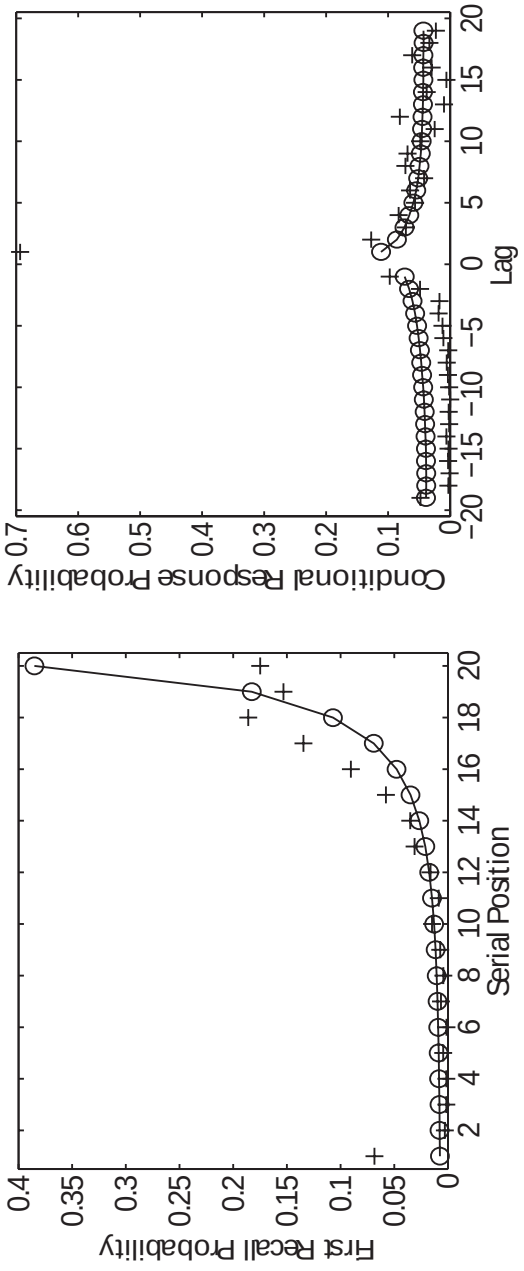
Conditional Response Probability
Lag
First Recall Probability
Serial Position

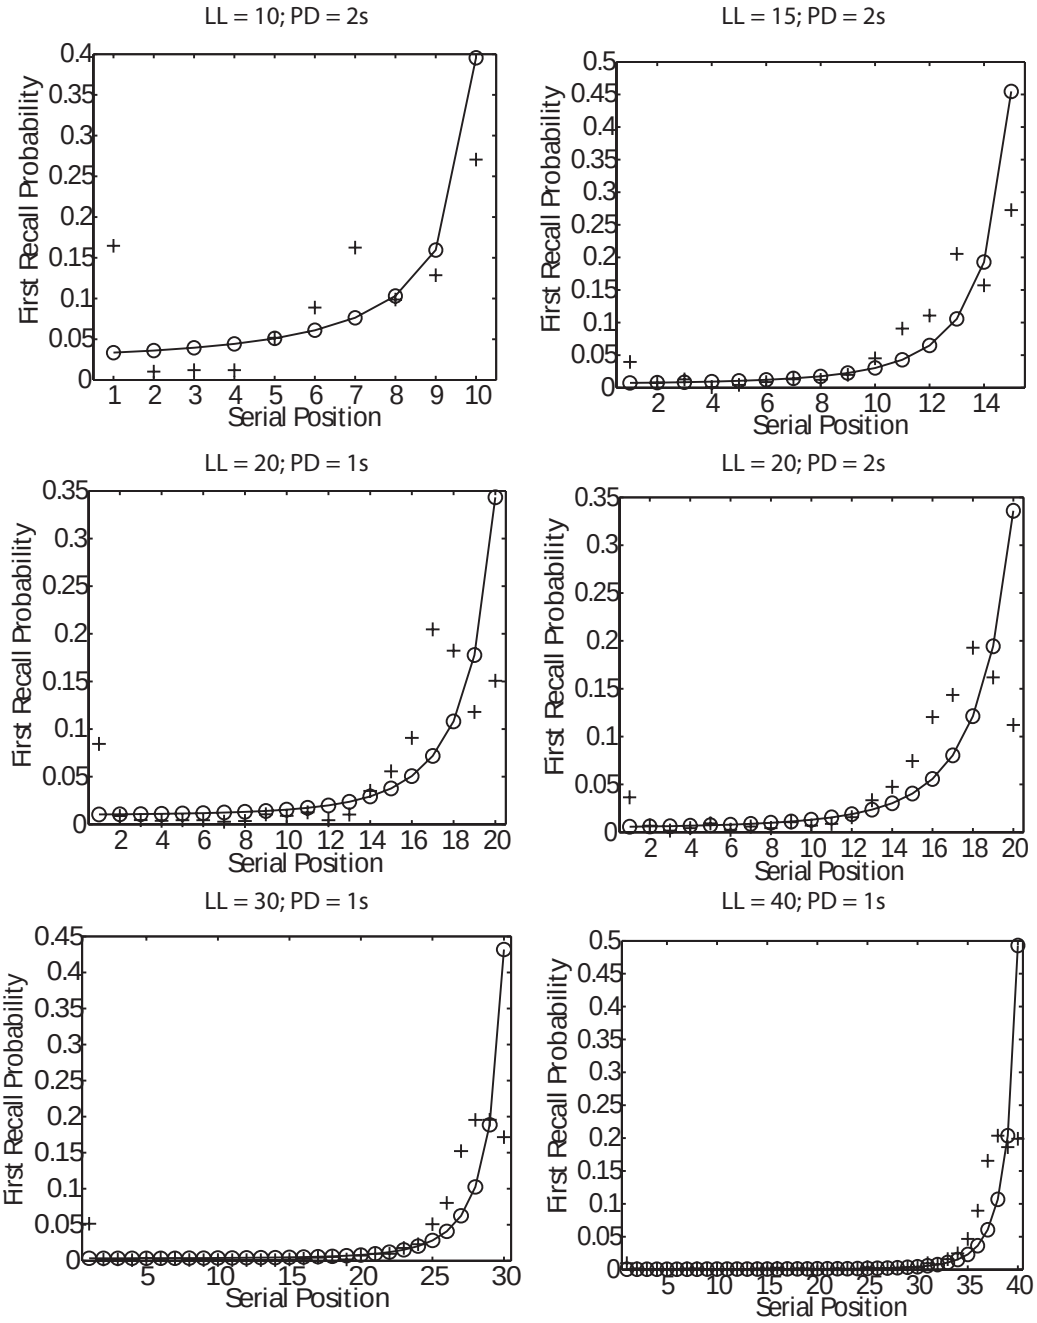
LL = 10; PD = 2s
LL = 15; PD = 2s
LL = 20; PD = 1s
LL = 20; PD = 2s
LL = 30; PD = 1s
LL = 40; PD = 1s
First Recall Probability
Serial Position

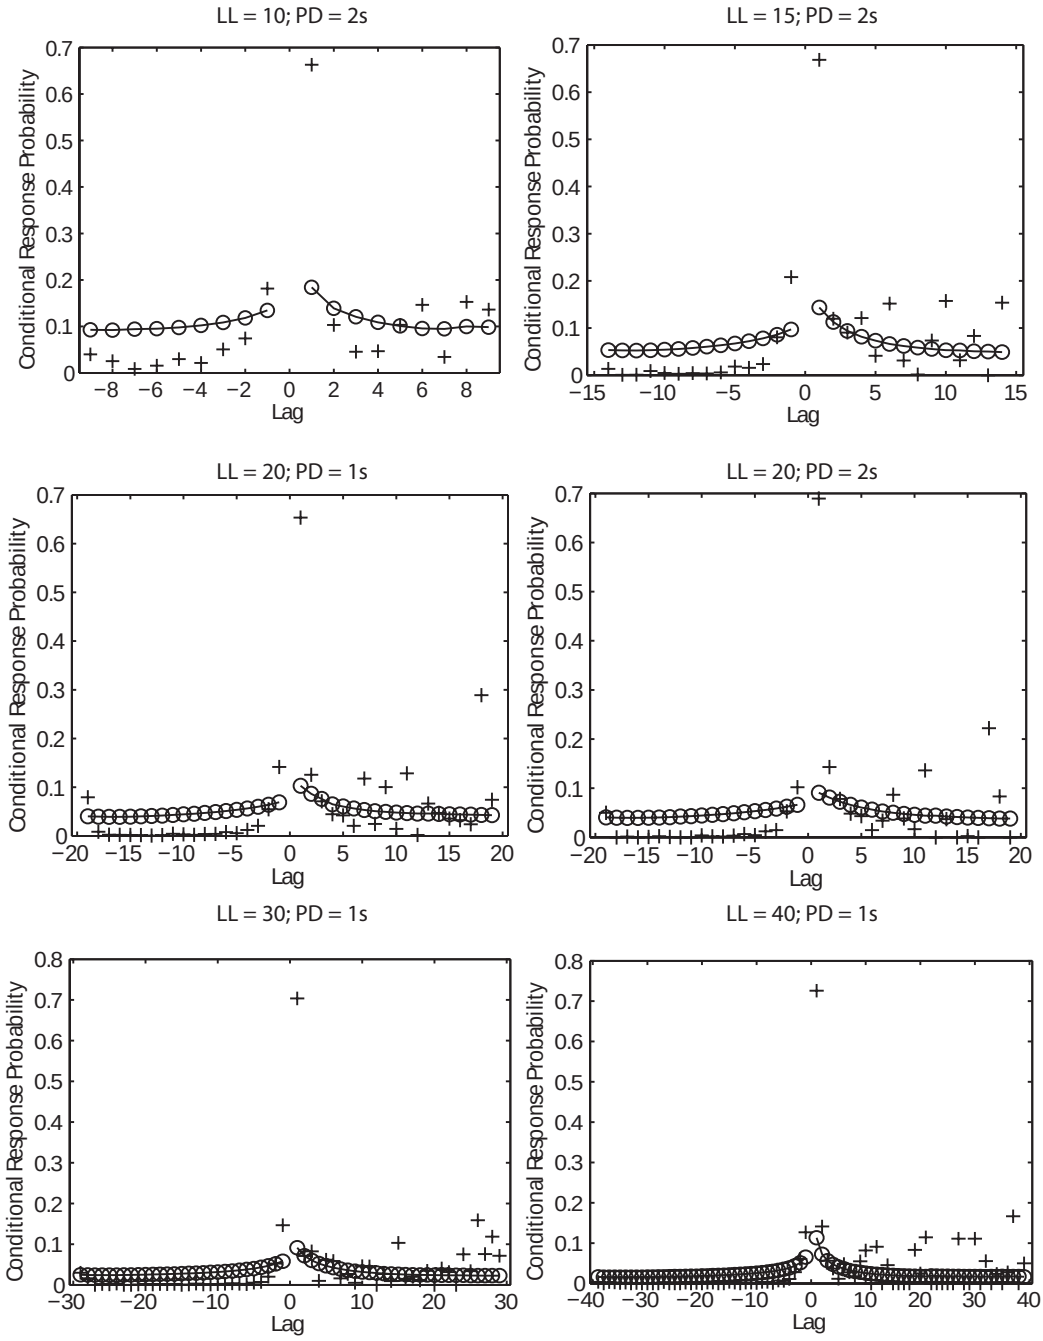
LL = 10; PD = 2s
LL = 15; PD = 2s
LL = 20; PD = 1s
LL = 20; PD = 2s
LL = 30; PD = 1s
LL = 40; PD = 1s
Conditional Response Probability
Lag

Limits on lag-recency supplement, Figure 7

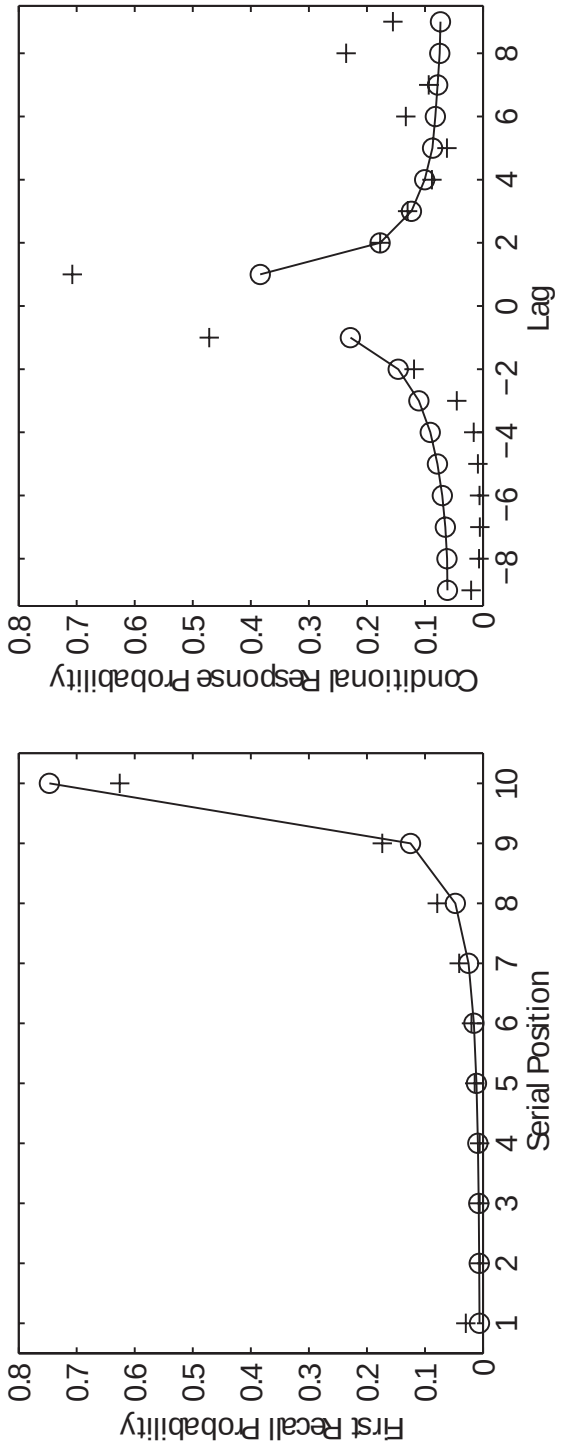

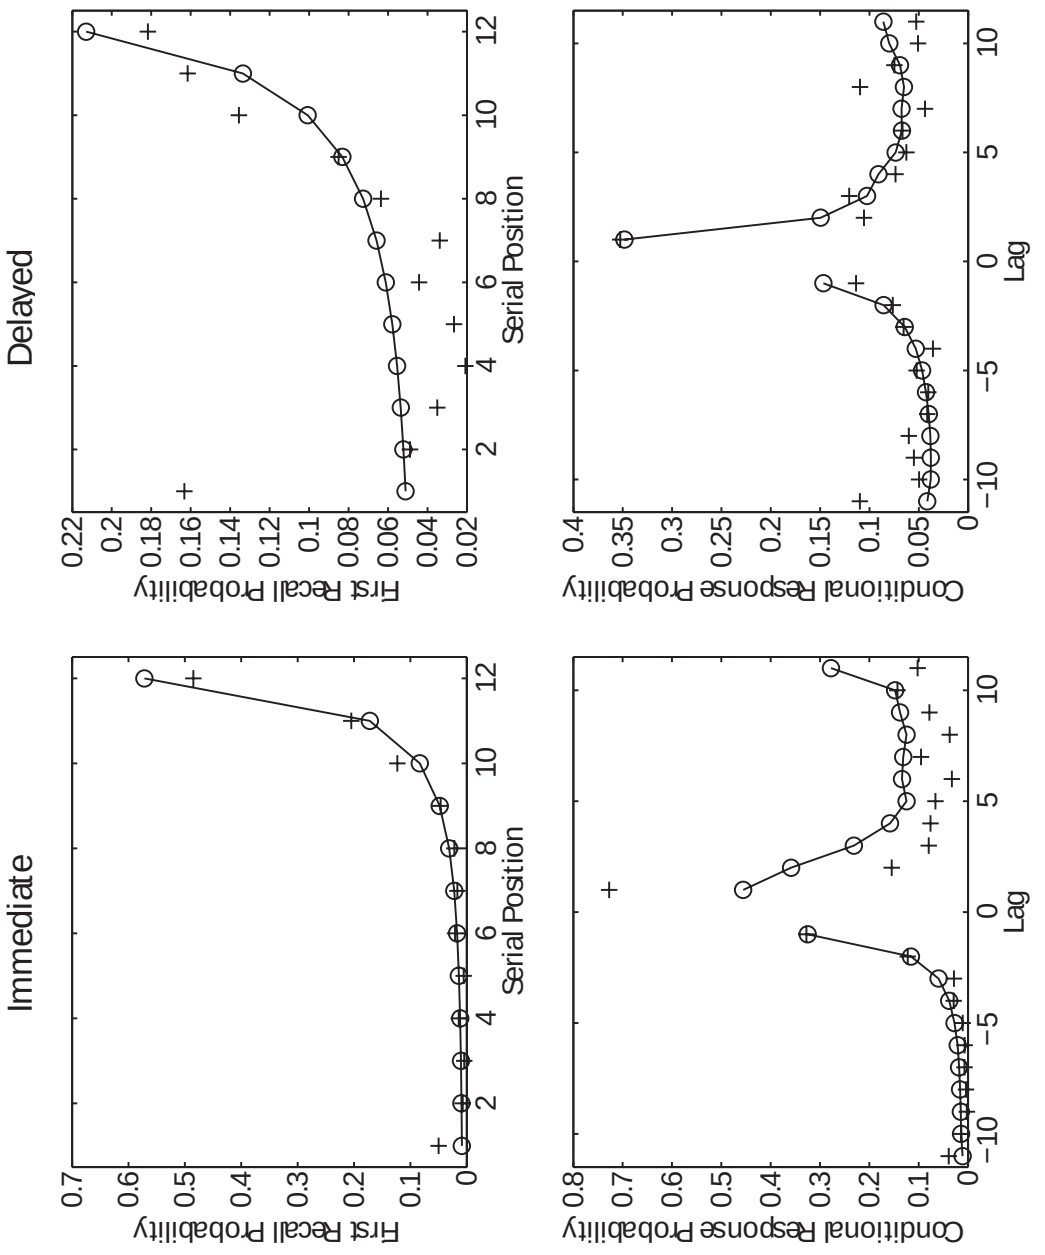
Immediate
Delayed
First Recall Probability
Serial Position
Conditional Response Probability
Lag

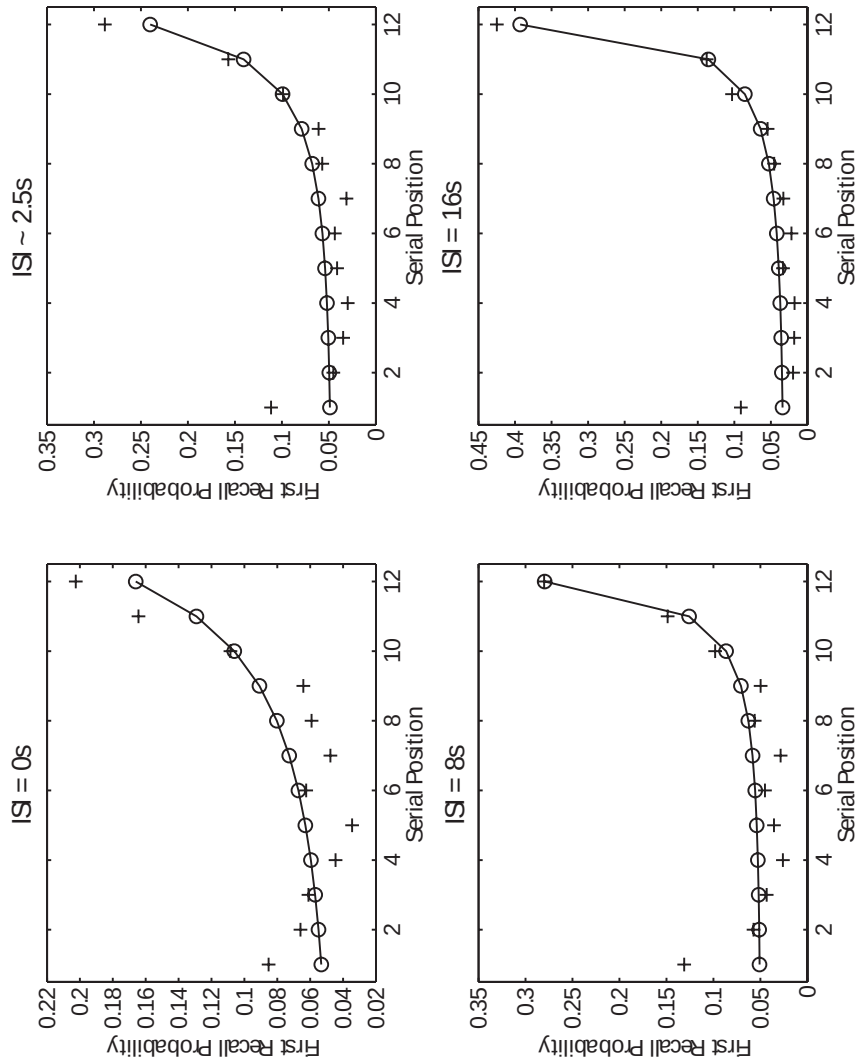
ISI = 0s
ISI ~ 2.5s
ISI = 8s
ISI = 16s
First Recall Probability
Serial Position

Limits on lag-recency supplement, Figure 10

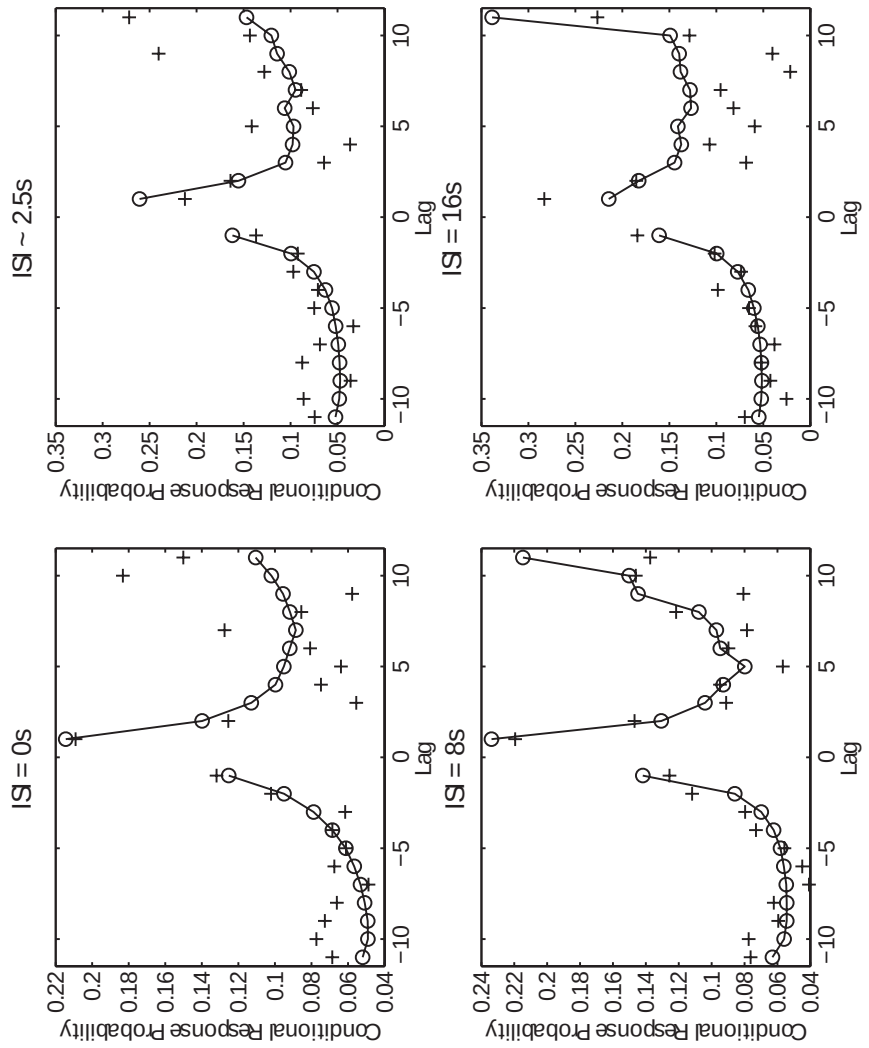

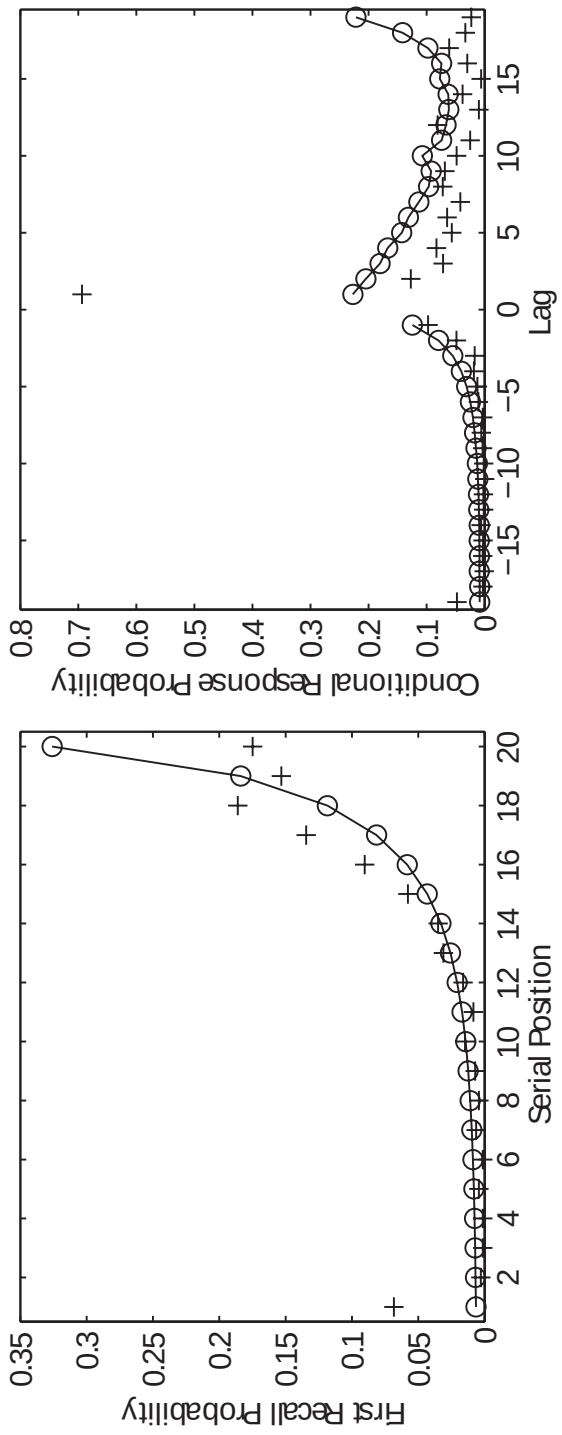
First Recall Probability
Serial Position
Conditional Response Probability
Lag

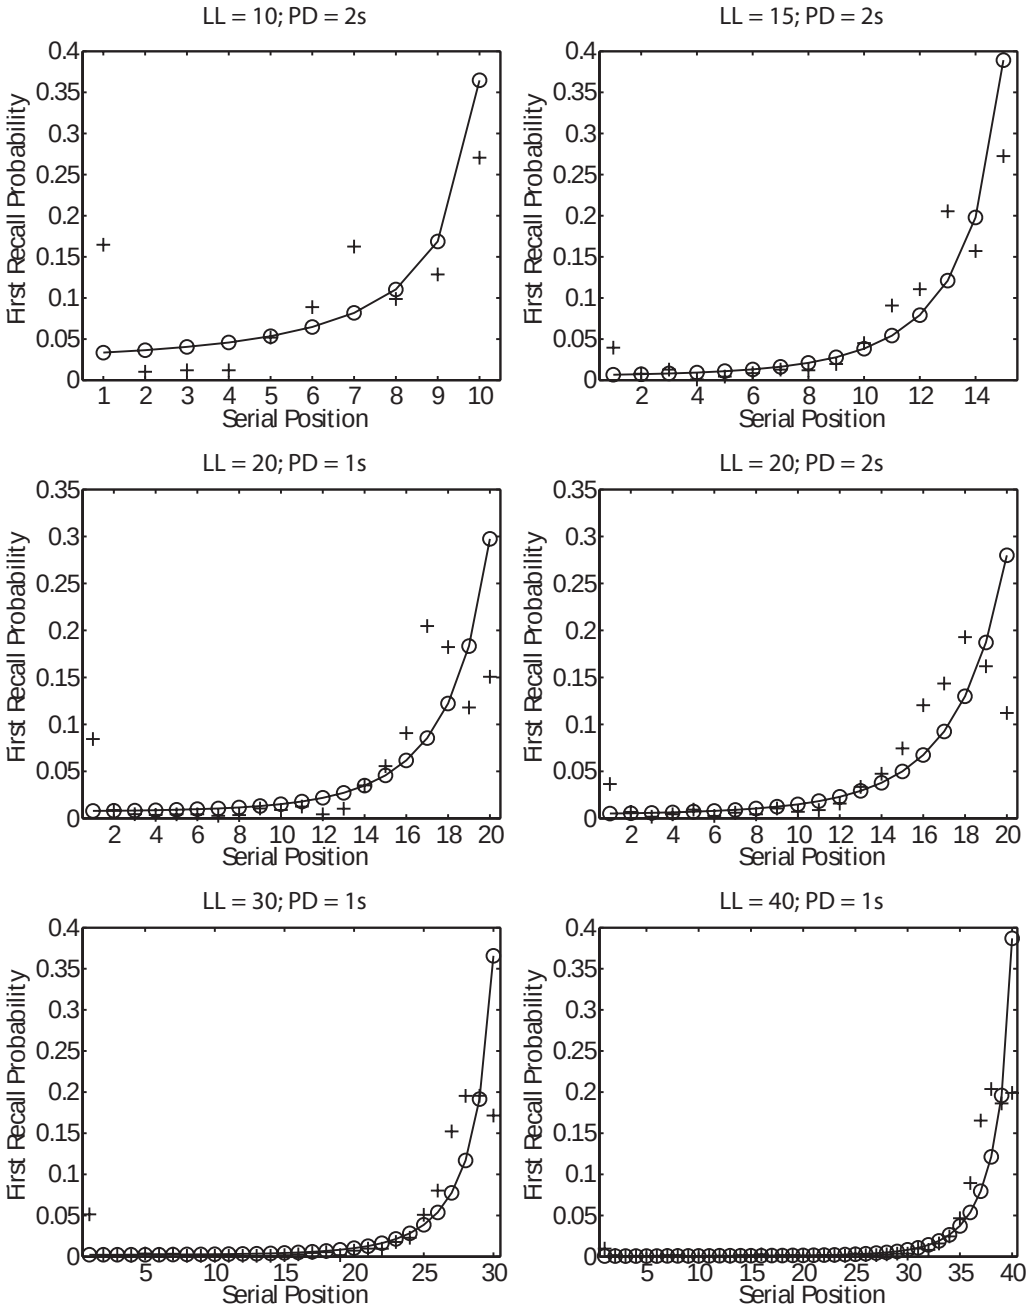
LL = 10; PD = 2s
LL = 15; PD = 2s
LL = 20; PD = 1s
LL = 20; PD = 2s
LL = 30; PD = 1s
LL = 40; PD = 1s
First Recall Probability
Serial Position

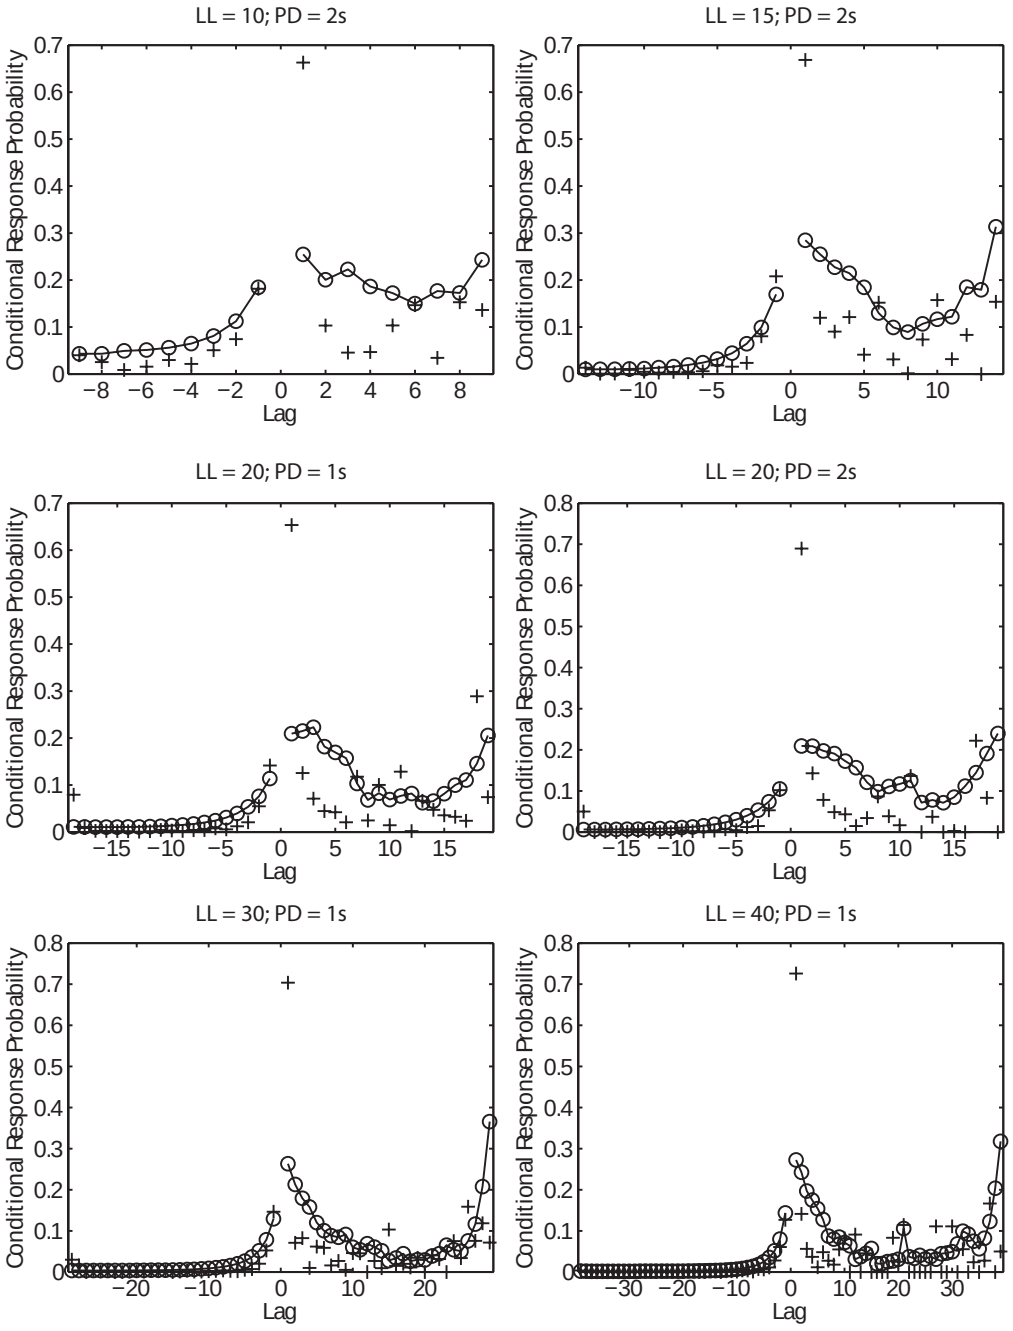
LL = 10; PD = 2s
LL = 15; PD = 2s
LL = 20; PD = 1s
LL = 20; PD = 2s
LL = 30; PD = 1s
LL = 40; PD = 1s
Conditional Response Probability
Lag

Limits on lag-recency supplement, Figure 14

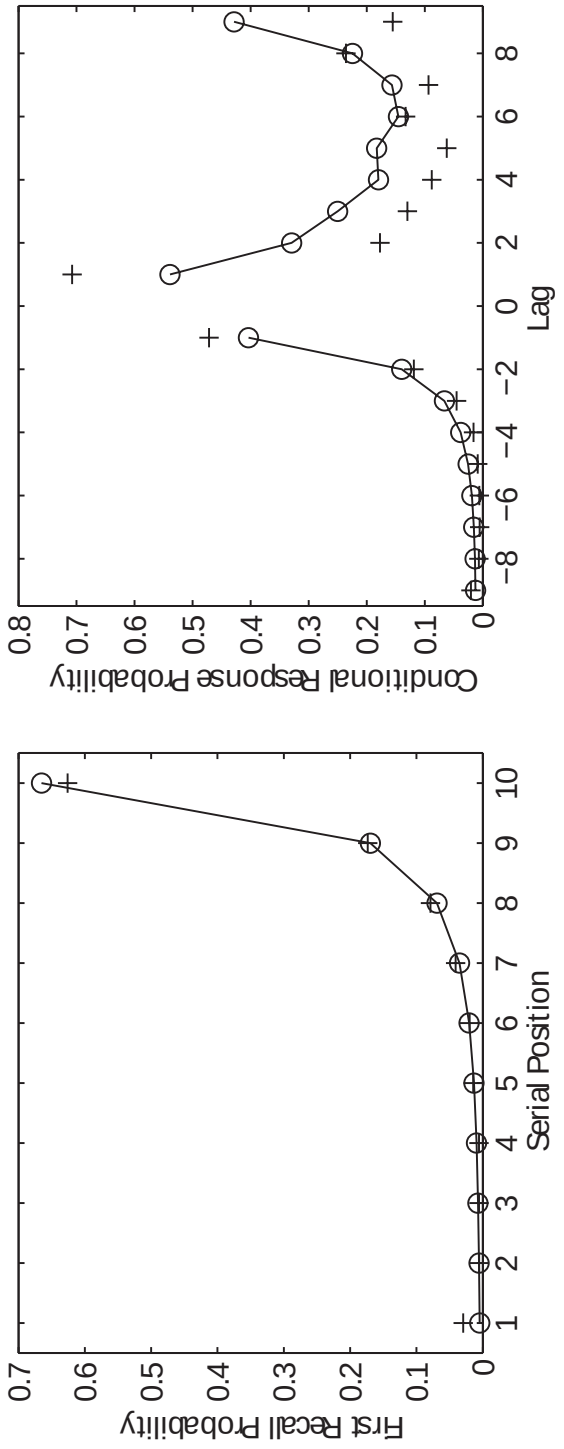

Limits on lag-recency supplement, Figure 15