

RESEARCH REPORT

Criterion Noise in Ratings-Based Recognition: Evidence From the Effects of Response Scale Length on Recognition Accuracy

Aaron S. Benjamin and Jonathan G. Tullis
University of Illinois at Urbana–Champaign

Ji Hae Lee
Washington University in St. Louis

Rating scales are a standard measurement tool in psychological research. However, research has suggested that the cognitive burden involved in maintaining the criteria used to parcel subjective evidence into ratings introduces *decision noise* and affects estimates of performance in the underlying task. There has been debate over whether such decision noise is evident in recognition, with some authors arguing that it is substantial and others arguing that it is trivial or nonexistent. Here we directly assess the presence of decision noise by evaluating whether the length of a rating scale on which recognition judgments are provided is inversely related to performance on the recognition task. That prediction was confirmed: Rating scales with more options led to lower estimates of recognition than did scales with fewer options. This result supports the claim that decision noise contributes to recognition judgments and additionally suggests that caution is warranted when using rating scales more generally.

Keywords: recognition, ratings, criteria, criterion noise, decision noise

Rating scales are among the most widely used measurement tools in psychology. They provide the basis for a majority of absolute and relative judgment tasks in perception and cognition, often provide the fundamental data for exercises in scaling, and, most importantly for present purposes, provide a means of estimating multiple points on a single detection or discrimination function. That function is often called an *isosensitivity function*, or receiver-operating characteristic, and the points along it represent equivalent discrimination but different underlying decision criteria. Isosensitivity functions play a prominent role in theoretical development, particularly in research on recognition memory, so it is important to examine closely the assumptions that underlie the translation between the shape and location of the isosensitivity function and the nature of the evidence that yields that function.

Here we consider the contrasting implications of the standard view of the decision process being static and nonvariable, as in classical Theory of Signal Detection (TSD; Green & Swets, 1966; Macmillan & Creelman, 2005), and a view with a noisy decision process (Benjamin, Diaz, & Wee, 2009; Malmberg & Xu, 2006; Mueller & Weidemann, 2008). In particular, we test the prediction

of the Noisy Decision Theory of Signal Detection (ND-TSD; Benjamin et al., 2009) that the location of the isosensitivity function should vary with the length of the rating scale used to assess it. The claim that recognition involves a noisy decision process has been controversial (Kellen, Klauer, & Singmann, 2012), so a direct evaluation of the predictions of ND-TSD is important. First, we briefly review the rating-scale methodology in recognition research and the way in which isosensitivity functions are generated from those data.

The isosensitivity function is a theoretical function relating true and false positives across the range of all potential decision criteria. It is useful because the shape and location of that function can be related to the shapes and locations of probabilistic evidence distributions that are thought to underlie the decision. Examples are provided in Figure 1, in which different isosensitivity functions are shown, along with the probabilistic distributions of noise and signal that generated them. Rectangular distributions with thresholds yield functions that are linear and intersect the axes away from one or both of the corners. Gaussian distributions yield curvilinear functions, with the degree of asymmetry indexing differences in variability between the distributions.

The isosensitivity function is estimated by collecting multiple points along the function in one of two ways. In one procedure, bias is manipulated by having participants respond under different payoff procedures; these payoffs induce more conservative or liberal responding by virtue of the cost/benefit tradeoff of different types of errors. In the second procedure, more common in recognition research, confidence ratings are taken during the response task, and these ratings are treated as criteria partitioning the evidence space. Some have argued that the inclusion of ratings so perverts the shape of the isosensitivity function that the broad

This article was published Online First February 18, 2013.

Aaron S. Benjamin and Jonathan G. Tullis, Department of Psychology, University of Illinois at Urbana–Champaign; Ji Hae Lee, Department of Psychology, Washington University in St. Louis.

Portions of this research were made possible by Grant NIA 026263 from the National Institutes of Health.

Correspondence concerning this article should be addressed to Aaron S. Benjamin, Department of Psychology, University of Illinois at Urbana–Champaign, 603 East Daniel Street, Champaign, IL 61820. E-mail: asbenjam@illinois.edu

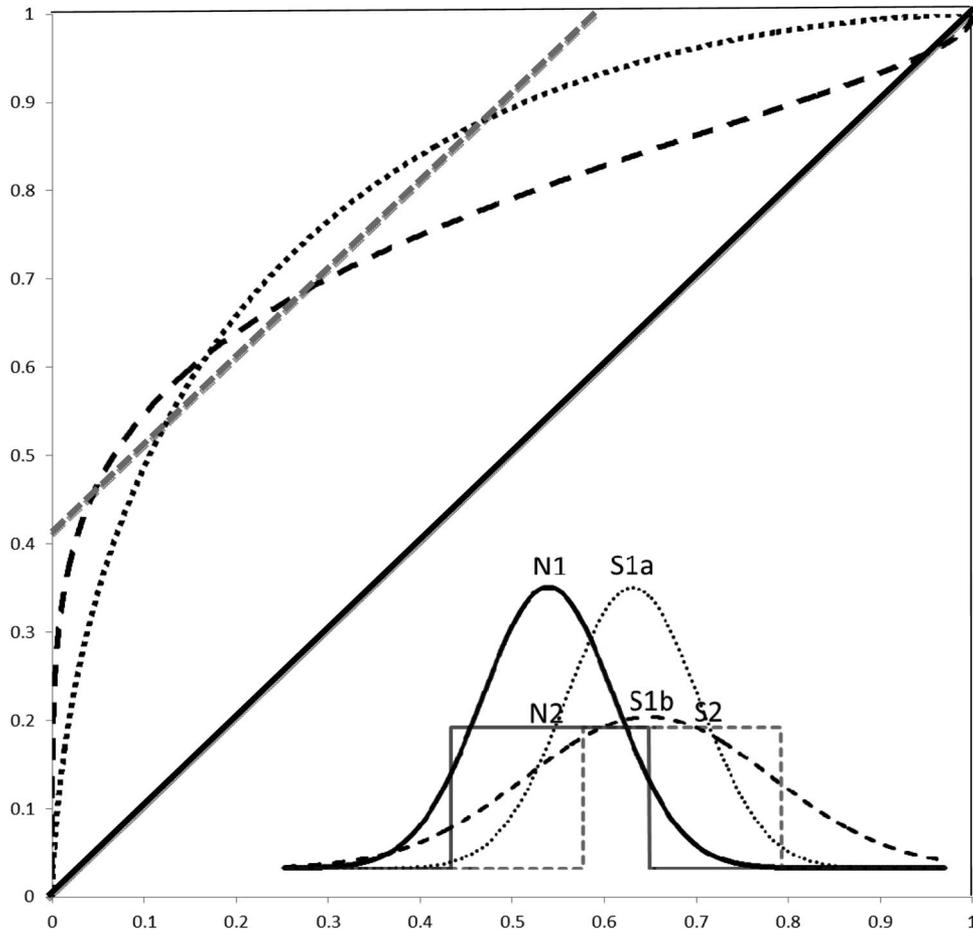


Figure 1. Theoretical isosensitivity functions and their generating distributions. Gaussian signal and noise distributions of equal variance (N1 and S1a) yield bowed and symmetric isosensitivity functions (black dotted curve). Gaussian signal and noise distributions of unequal variance (N1 and S1b) yield bowed and asymmetric functions (black dashed curve). Rectangular threshold functions (N2 and S2) yield straight functions (gray line).

consensus that evidence is graded rather than thresholded is wrong (Bröder & Schütz, 2009). However, there is compelling evidence that isosensitivity functions are in fact curvilinear even when estimated from manipulations of bias (Dube & Rotello, 2012; Koen & Yonelinas, 2011), indicating that the assumption of graded evidence is indeed correct. However, the exact shape of isosensitivity functions estimated from ratings does differ across conditions of differential bias (Van Zandt, 2000), so there is reason for concern that the underlying information available to the recognizer might not be equivalent in the two cases.

Malmberg and Xu (2006; Malmberg, 2002) noted that variations and suboptimalities in the decision process corrode the relationship between the isosensitivity function and the underlying evidence, and they suggested that the theoretical agenda of trying to discern the nature of the evidence in recognition from the isosensitivity function may be fundamentally flawed, particularly when that function is estimated from confidence ratings. Benjamin et al. (2009) echoed this sentiment and further provided estimates of decision noise within recognition that were sufficiently large to merit concern. In their study, decision noise was estimated to be of

approximately the same magnitude as stimulus noise—that is, decision noise contributed as much to the recognition decision as did the differences across stimuli within the experiment.

Kellen et al. (2012) provided new data, using a direct comparison between forced-choice and yes/no recognition, and came to the opposite conclusion: that decision noise played no meaningful role in recognition judgments. The goal of the present experiment was to examine in as direct a manner as possible the claim that the criteria that recognizers set have some variability, or noisiness, associated with them. We did so by evaluating whether rating scales with more options, and consequently more criteria to discriminate between those options, yield “poorer” recognition performance than do scales with fewer options. If this prediction were confirmed, it would suggest that the higher number of criteria engender a greater amount of decision noise that plays out in estimates of “poorer” performance. *Poorer* is placed in scare quotes here because the core process of recognition is not presumed to be impaired by the use of rating scales, only the translation of that evidence into judgments via a noisy decision process.

Wickelgren (1968) suggested that noisy criteria could lead the point corresponding to yes/no discrimination to lie above the isosensitivity function estimated from multiple points, indicating “poorer” performance in the multiple-rating than the two-rating case. The small amount of research on this topic is mixed and almost entirely from research in perception, with some results confirming this claim (Swets, Tanner, & Birdsall, 1961, Figure 14) and others not (Baranski & Petrusic, 2001; Egan, Schulman, & Greenberg, 1959). Complicating the issue further is the fact that, even when comparisons do not yield evidence for differences in discriminability, they may lead to different response times (Petrusic & Baranski, 2003). The closest result in the literature comes from a report by Koen and Yonelinas (2011), in which they directly compared isosensitivity functions estimated from confidence ratings with ones estimated from yes/no responses with a payoff manipulation between conditions. They found similar functions across those two conditions, but several aspects of their procedure are not ideal for our purposes. First, the payoff manipulation may invite a novel memory demand that depresses performance in that yes/no condition relative to a case in which payoffs are unvaried. Second, their study had relatively low power due to the between-subjects nature of the manipulation and the relatively small sample size ($n = 20$ and 22 in the two relevant conditions). Here we provide a means of evaluating the effects of rating-scale length powerfully and incisively, without additional manipulations of payoff. We measured recognition performance for previously studied words under conditions in which participants made yes/no judgments (a 2-point rating scale), 4-point confidence rating scale judgments, and 8-point confidence rating scale judgments.

Method

Participants

Sixty undergraduate students from the University of Illinois participated in this experiment as part of a course requirement.

Materials

Six hundred words were chosen from the English Lexicon Project (Balota et al., 2002). To ensure a wide range of preexperimental familiarities, half of the words were chosen to be of relatively high frequency (mean log HAL frequency 11.66), and half were chosen to be of relatively low frequency (mean log HAL frequency 8.84). The words ranged in length between four and eight letters (mean-high = 5.44, mean-low = 5.59). A total of six sets of 60 words were chosen pseudorandomly from the pool without replacement for each participant, with the condition that half of each set of words was high frequency and the other half low frequency. Three of the six subsets were designated to be study lists, and the remaining lists served as distractors for the tests.

Design and Procedure

Participants experienced three study–test cycles, each of which implemented a different rating scale condition—a two-alternative (yes/no), a four-alternative, or an eight-alternative condition. The order of the conditions was counterbalanced, and partici-

pants were informed of the nature of the rating scale immediately prior to the relevant test.

Participants were individually placed in a well-lit room, seated approximately 40 cm away from a computer monitor. Presentation of stimuli and recording of responses was controlled by Matlab with the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Before the first study phase, participants were told that they would study a list of words for a later memory test. During the study phase, each word was presented for 3,000 ms in white in the center of the black screen, with a 300-ms interstimulus interval (ISI). After each study phase, participants were given another set of instructions telling them that they would now be tested on the words they had just studied. Participants were told to choose a boxed number that corresponded to their response decision regarding their memory of a given word.

During the test, participants were presented with either a previously studied or a previously unstudied word and were asked to pick a response to the question “Do you remember studying this word?” The response options, which were presented as boxed numbers, appeared right below the word being tested and remained on the screen until a response was made. Only the cues *sure no* (on the far left) and *sure yes* (on the far right) were presented below the boxed numbers, to ensure that the order of the rating scale was clear. Test stimuli remained on the screen until a response was provided.

The presentation of words in both study and test blocks was designed so that no more than four words from the same frequency category were presented in a row. Also, the presentation of studied and nonstudied words in each test block was pseudorandomly ordered so that no more than four studied or nonstudied words were presented in a row.

Results

The data are shown in the left panel of Figure 2. The data points show the average hit rates and false-alarm rates averaged across participants for the three conditions.¹ It can be seen that the performance level for yes/no discrimination is higher (i.e., lies further in the upper left quadrant of the space) than the points from the 4-point judgments and that performance for the 4-point judgments is higher than for the 8-point judgments. However, direct comparison is difficult. Performance in the conditions with more than two options on the rating scales is traditionally summarized with a measure of performance that takes into account the (potential) differential variability between the signal and noise evidence (such as d_a), but that measure cannot be computed for performance in the two-option condition. To ensure that all conditions could be compared against one another on equal footing, we developed a novel technique for comparison.

The logic of this test is as follows. For each participant, we compared the obtained hit rate for a given condition (and for a given rating value, when appropriate) to the predicted hit rate estimated from the isosensitivity function for a condition with a higher number of ratings. Equivalently, we took a given point (say, the yes/no point), drew a vertical line to the isosensitivity function with a higher number of ratings (say the 4-rating condition), and compared the y values of those two points. If the predicted hit rates

¹ Scores were adjusted by adding 0.5 to the count of hits and false alarms and adding 1 to the total number of relevant items.

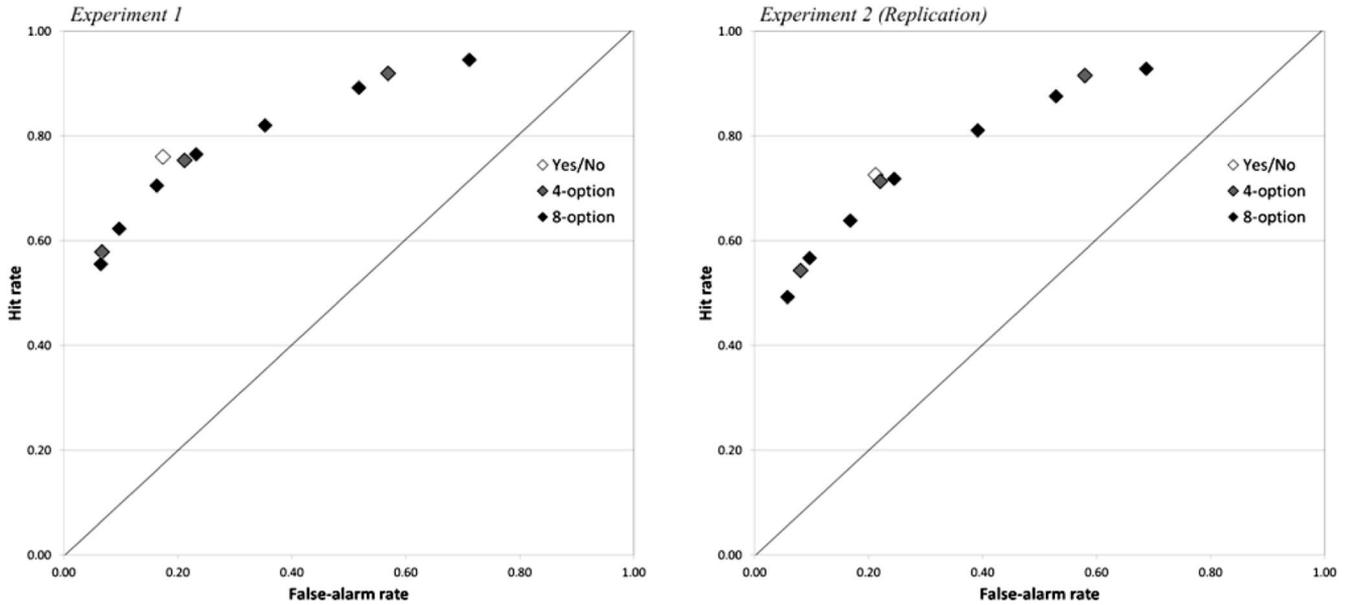


Figure 2. Mean performance as a function of rating scale length for Experiment 1 (left) and Experiment 2 (right). Endorsement rates are averaged over participants.

were reliably lower than the obtained hit rates—that is, if the isosensitivity function from the condition with the higher number of ratings lied consistently below the obtained score—then those conditions differed in discriminability in the predicted direction. Because the conditions with higher ratings were hypothesized to involve more decision noise, those conditions should have yielded lower performance and thus underestimated performance in conditions with fewer rating options.

This technique also has the quality that it conditionalizes on an exact false-alarm rate for each comparison. This is an advantage because the conditions may induce different response biases, rendering direct comparison between the empirically obtained hit/false-alarm rate pairs across conditions difficult.

The results from this analysis are shown in Figure 3. For the yes/no condition, theoretical hit rates estimated from the isosensitivity functions of the four-option, $t(59) = 2.17$, and the eight-option, $t(59) = 2.37$, conditions underestimated performance. The average effect size d for these two comparisons was 0.30, indicating a small-to-medium effect (Cohen, 1988). For the four-option condition, estimates from the eight-option condition underestimated the most liberal cumulative ratings category, $t(59) = 2.24$, but not the other two, $t(59) = 1.04$ and 0.28. The effect size d for these three tests was 0.16, indicating a small effect.

One potential concern here lies in the self-paced nature of the test: As noted in our earlier review, longer rating scales might elicit longer response times (Petrusic & Baranski, 2003). This did occur in this experiment ($M_{RT} = 1.29$ s, 1.60 s, and 1.80 s, for the two-, four- and eight-rating conditions, respectively). There are several effects that this confound could have on our results. First, longer scales might foster the use of more conservative points on the speed-accuracy tradeoff function. Second, longer scales might introduce a sufficiently large delay to induce a functionally greater retention interval. The first possibility can be ruled out because it makes a prediction opposite

to what actually occurred in the experiment: Longer ratings scales led to lower, not higher, accuracy. To evaluate whether retention interval played a role in the effect reported here, we directly compared d_a across the four- and eight-rating conditions for the first and second halves of the test. If each individual test trial introduced a longer delay in the longer than the shorter rating scale conditions, then the advantage for shorter rating scales should be greater in the second half of the test, when a greater difference owing to this differential “slack” has accumulated. However, the effect of rating scale was actually numerically greater in the first half ($\Delta d_a = 0.34$) than the second half ($\Delta d_a = 0.27$) of the test, thus allaying any concern that differences in retention interval played a role in the effect described above.

Results of Replication (Experiment 2)

In order to ensure the validity of the claim that shorter rating scales induce poorer performance, we conducted an exact replication of the experiment. Only one replication was conducted (i.e., we did not conduct multiple tests and select the one with the most promising results), and the only difference between the two experiments is that 64 participants were included in the replication.

The right panel in Figure 2 and the middle panels in Figure 3 display the results from the replication experiment. In the replication, theoretical hit rates estimated from the isosensitivity functions of the four-option, $t(63) = 0.86$, and the eight-option, $t(63) = 2.54$, conditions underestimated performance in the yes/no condition, but only the latter was significant. The average effect size d for these two comparisons was 0.22, indicating a small effect (Cohen, 1988). For the four-option condition, estimates from the eight-option condition underestimated the most liberal cumulative ratings category, $t(63) = 3.33$, and the middle category, $t(63) =$

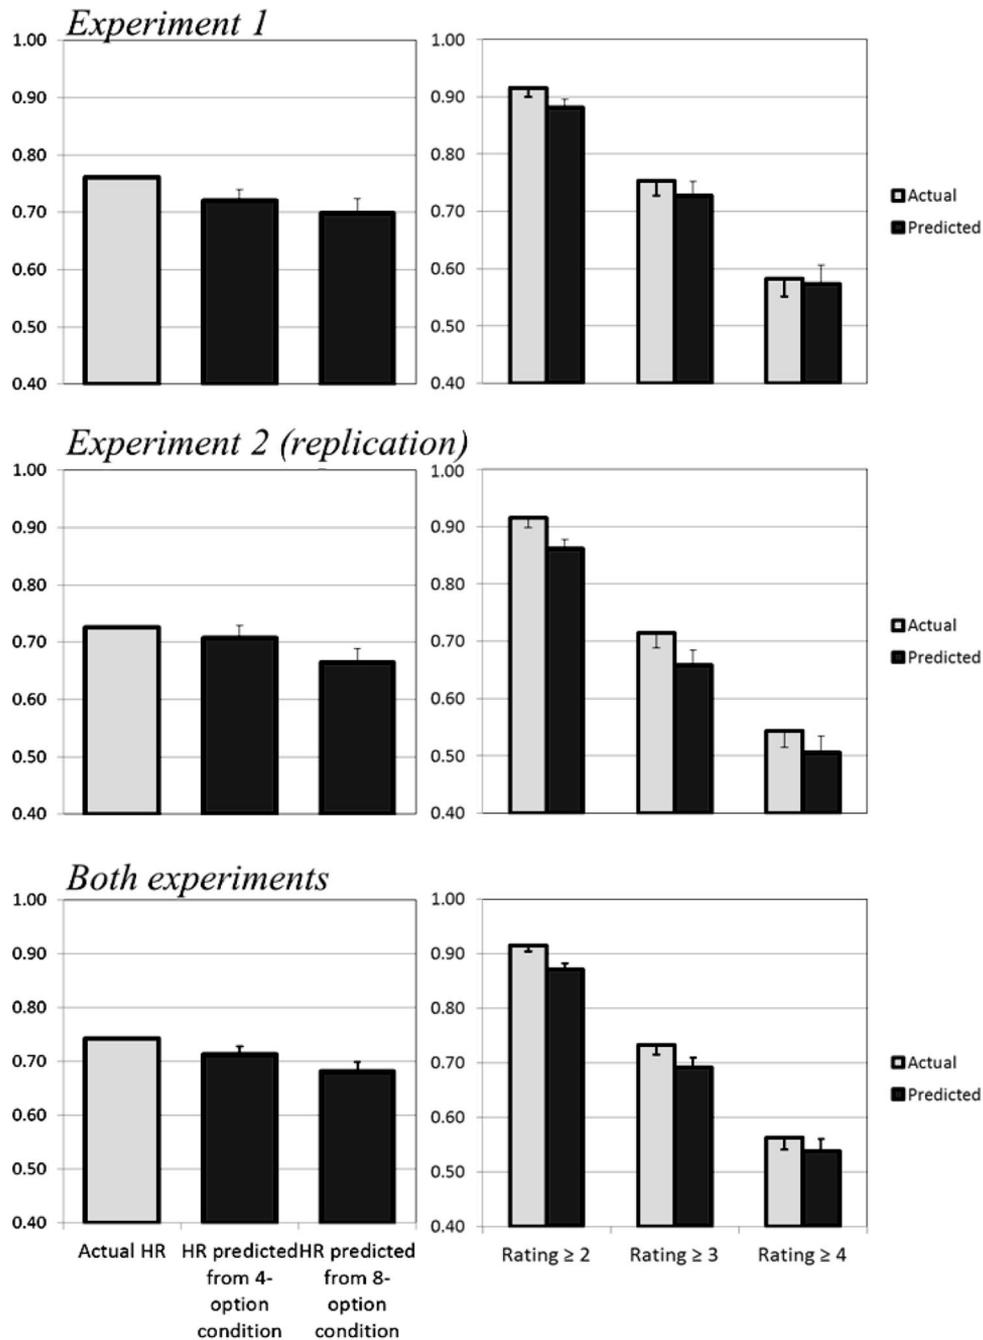


Figure 3. Actual and predicted hit rates for yes/no (left panels) and four-option (right panels) rating scales. Error bars represent one standard error of the mean of the within-subject difference score.

2.16, but not the most conservative category, $t(63) = 1.33$. The effect size d across the three tests was 0.29, indicating a small-to-medium effect.

Combined analysis

The data from both experiments were combined in order to increase power for each of the component comparisons. The re-

sults are shown in Figure 4 and in the bottom panel of Figure 3. In that analysis, obtained hit rates for the yes/no condition were higher than theoretical hit rates estimated from either the four-option, $t(123) = 2.02$, or the eight-option, $t(123) = 3.49$, condition. In addition, obtained hit rates were higher in the four-option condition than those predicted by the eight-option condition for all cumulative rating categories, $t_s(123) = 3.98$ and 2.29, respectively, except for the most conservative one, $t(123) = 1.13$. The

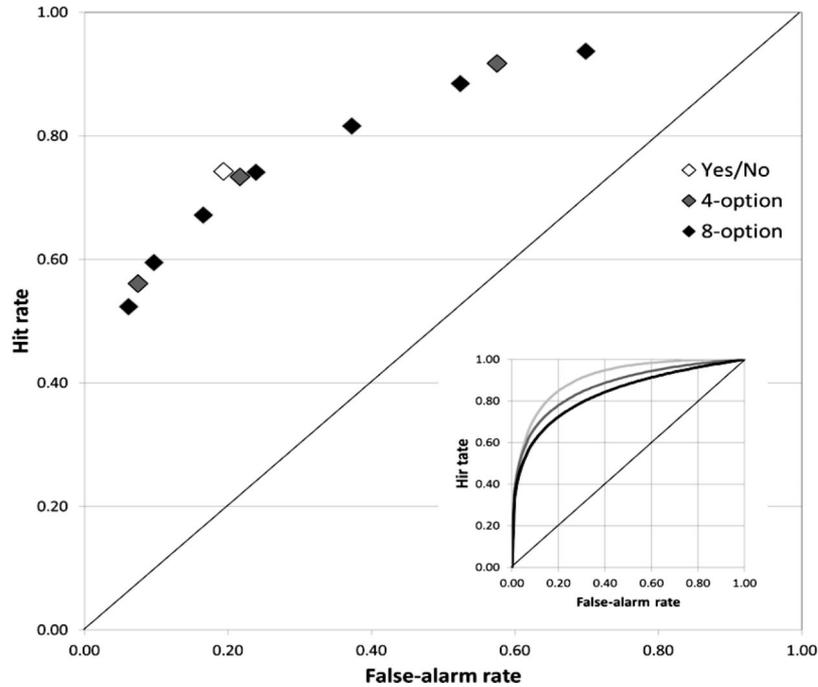


Figure 4. Mean performance as a function of rating scale length, collapsed across both experiments. The larger figure shows endorsement rates averaged over participants; the inset shows the functions from the average parameters estimated from the model fit to individual participants. Median parameter values were used because the model occasionally failed to converge on reasonable solutions for the eight-option condition.

effect sizes (d) for each of these comparisons were 0.18, 0.31, 0.29, 0.14, and 0.04, respectively. These results indicate that the comparison across more distant conditions (two scale options vs. eight) yielded larger effects and also that effects are more pronounced in the more liberal response portion of the scale (the higher side of the isosensitivity function). As a final check on the effect of interest, d_a ratings were directly compared between the four-option ($M = 1.54$) and the eight-option ($M = 1.33$) condition, and they were significantly higher in the four-rating condition, $t(123) = 3.43$.

The inset of Figure 4 shows the isosensitivity functions based on the median parameters generated from a maximum-likelihood fit of the unequal-variance signal-detection (UVSD) model to the four- and eight-rating condition and generated from a fit of the equal-variance signal-detection (EVSD) model to the yes/no condition. The yes/no curve is symmetric because the EVSD model cannot support the estimation of differential variance between signal and noise. These group isosensitivity functions are shown only for ease of visualization, not for analysis, but it can also be seen in those functions that scales with a higher number of options led to poorer performance.

Discussion

The fact that rating scales with more options lead to lower estimates of recognition performance has major implications for theoretical views of the decision process underlying recognition and for the practical value of using rating scales in psychological experiments. Here we must conclude either that the nature of the rating scale somehow affects memory for the material being

tested—an unlikely option—or, as suggested by Benjamin et al. (2009), that each point on the rating scale introduces some amount of variability to the decision process and undermines recognition performance.

The idea that maintaining criteria poses a burden to memory—and thus that maintaining more criteria poses a greater burden—is consistent with a large range of evidence in memory and psychophysics, including response autocorrelations (Treisman & Williams, 1984), inconsistencies in the relationship between forced-choice and yes/no recognition procedures (Green & Moses, 1966), variability in the slope of the function across learning conditions (Glanzer, Kim, Hilford, & Adams, 1999), differences between confidence-rating and bias-induction recognition procedures (Van Zandt, 2000), probability matching (Lee, 1963), response conservatism in response to manipulations of base rates (Healy & Kubovy, 1978), effects of aging (Kapucu, Rotello, Ready, & Seidl, 2008), and variation in the slope of the z receiver operating characteristic for “remembered” items (Wixted & Stretch, 2004). In addition, criterion variability has been revealed in perceptual tasks (Bonnell & Miller, 1994; Nosofsky, 1983) and is incorporated into sampling models of recognition (Ratcliff & Rouder, 1998). Though the traditional application of TSD to perceptual and mnemonic tasks leaves no room for such decision noise, models incorporating a role for decision noise are available (e.g., Benjamin et al., 2009; Nosofsky, 1983; Wickelgren, 1968), and these are entirely consistent with the spirit of detection theory.

If criteria do place a burden on memory, what is the nature of that burden? ND-TSD treats criteria as random samples from

normal distributions but is agnostic with respect to the nature of the source of variance. There is good evidence that criteria are not maintained as singular entities but rather tied to the range of evidence that the recognizer experiences (e.g., Benjamin, 2003, 2005; Hirshman, 1995; Stretch & Wixted, 1998; Tullis & Benjamin, 2012; Turner, Van Zandt, & Brown, 2011). Mapping evidence onto responses thus requires a plan in which quantiles are determined from the number of response options available and are updated as the range of evidence changes. This updating is one source of noise, and the shifting of criteria induced by changes in range may also introduce memory failures in which, for example, a recognizer mistakenly fails to use an updated criterion value. Such a process would also be consistent with the presence of response dependencies (Malmberg & Annis, 2012; Treisman & Williams, 1984). In our original experiment, sequential dependencies were apparent and reliable (mean response autocorrelation = 0.06; standard error of the mean = 0.015) but did not differ across the ratings conditions. It would thus be difficult to attribute our condition effect exclusively to any source of noise that would lead to different levels of sequential dependencies.

The fact that decision noise can influence the shape and location of the isosensitivity function does suggest limits on the use of those functions in theoretical development. In recognition memory research, major theoretical debates over the number and nature of the processes that contribute to recognition have played out on a battlefield of isosensitivity functions (e.g., Wixted, 2007; Yonelinas, 1999), in which relatively subtle variations in form are taken to have substantial theoretical relevance. The presence of decision noise, and the unknown individual differences it brings with it, suggests that an overreliance on such methodological tools may be dangerous.

We do not wish to suggest that the general use of rating scales in psychology is fundamentally flawed. In many cases, the addition of an unknown amount of decision noise does not meaningfully affect the types of conclusions researchers wish to draw from their data. The problems introduced by decision noise can perhaps be characterized best as a bias in estimation rather than comparison. Measures of performance that include decision noise are likely inaccurate estimators of the underlying perceptual or mnemonic skill. But comparisons between similar conditions that are affected to the same degree by decision noise should not be dramatically hampered by the presence of that noise. More accurately, such comparisons are affected by criterion noise in the same way that they are affected by the many other unavoidable forms of uncontrolled noise in such experiments.

We end by noting that the decision noise debate echoes an earlier debate over the use of rating scales in individual-difference research (e.g., Garner, 1960). There it has been known for a long time that rating scales with many options provide little benefit for measurement when compared to scales with fewer options (Symonds, 1924). Some authors have even shown a loss in reliability with a higher number of scale options (Bendig, 1953), and others have argued that any increases in reliability that might come from increasing the scale length do not benefit the validity of the instrument (Cronbach, 1950).

As in that field, it is worth remembering that the best scale is the one that optimizes the tradeoff between the coarseness of the measurement and the limited discriminating precision of the rater. Discriminating among levels of subjective confidence or evidence

in service of a recognition decision is certainly no less fraught with uncertainty over the boundary between response categories than discriminating among options in response to a personality or educational instrument. The seminal article by Miller (1956) is often remembered for its review of limitations on short-term memory, but it was in fact more substantively concerned with limits on absolute identification—a limitation that would be profitable to remember when designing response instruments. We ignore decision noise at our own peril.

References

- Balota, D. A., Cortese, M. J., Hutchison, K. A., Neely, J. H., Nelson, D., Simpson, G. B., & Treiman, R. (2002). *The English Lexicon Project: A web-based repository of descriptive and behavioral measures for 40,481 English words and nonwords*. Retrieved from Washington University website: <http://elexicon.wustl.edu/>
- Baranski, J. V., & Petrusic, W. M. (2001). Testing architectures of the decision-confidence relation. *Canadian Journal of Experimental Psychology*, *55*, 195–206. doi:10.1037/h0087366
- Bendig, A. W. (1953). The reliability of self-ratings as a function of the amount of verbal anchoring and the number of categories on the scale. *Journal of Applied Psychology*, *37*, 38–41. doi:10.1037/h0057911
- Benjamin, A. S. (2003). Predicting and postdicting the effects of word frequency on memory. *Memory & Cognition*, *31*, 297–305. doi:10.3758/BF03194388
- Benjamin, A. S. (2005). Recognition memory and introspective remember/know judgments: Evidence for the influence of distractor plausibility on “remembering” and a caution about purportedly nonparametric measures. *Memory & Cognition*, *33*, 261–269. doi:10.3758/BF03195315
- Benjamin, A. S., Diaz, M. L., & Wee, S. (2009). Signal detection with criterion noise: Applications to recognition memory. *Psychological Review*, *116*, 84–115. doi:10.1037/a0014351
- Bonnel, A.-M., & Miller, J. (1994). Attentional effects on concurrent psychophysical discriminations: Investigations of a sample-size model. *Perception & Psychophysics*, *55*, 162–179. doi:10.3758/BF03211664
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436. doi:10.1163/156856897X00357
- Bröder, A., & Schütz, J. (2009). Recognition ROCs are curvilinear—Or are they? On premature arguments against the two-high-threshold model of recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 587–606. doi:10.1037/a0015279
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Hillsdale, NJ: Erlbaum.
- Cronbach, L. J. (1950). Further evidence on response sets and test design. *Educational and Psychological Measurement*, *10*, 3–31. doi:10.1177/001316445001000101
- Dube, C., & Rotello, C. M. (2012). Binary ROCs in perception and recognition memory are curved. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*, 130–151. doi:10.1037/a0024957
- Egan, J. P., Schulman, A. I., & Greenberg, G. Z. (1959). Operating characteristics determined by binary decisions and by ratings. *Journal of the Acoustical Society of America*, *31*, 768–773.
- Garner, W. R. (1960). Rating scales, discriminability, and information transmission. *Psychological Review*, *67*, 343–352. doi:10.1037/h0043047
- Glanzer, M., Kim, K., Hilford, A., & Adams, J. K. (1999). Slope of the receiver-operating characteristic in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 500–513. doi:10.1037/0278-7393.25.2.500
- Green, D. M., & Moses, F. L. (1966). On the equivalence of two recognition measures of short-term memory. *Psychological Bulletin*, *66*, 228–234. doi:10.1037/h0023645

- Green, D. M., & Swets, J. A. (1966). *Signal Detection Theory and psychophysics*. New York, NY: Wiley.
- Healy, A. F., & Kubovy, M. (1978). The effects of payoffs and prior probabilities on indices of performance and cutoff location in recognition memory. *Memory & Cognition*, *6*, 544–553. doi:10.3758/BF03198243
- Hirshman, E. (1995). Decision processes in recognition memory: Criterion shifts and the list-strength paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 302–313.
- Kapucu, A., Rotello, C. M., Ready, R. E., & Seidl, K. N. (2008). Response bias in “remembering” emotional stimuli: A new perspective on age differences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 703–711.
- Kellen, D., Klauer, K. C., & Singmann, H. (2012). On the measurement of criterion noise in Signal Detection Theory: The case of recognition memory. *Psychological Review*, *119*, 457–479. doi:10.1037/a0027727
- Koen, J. D., & Yonelinas, A. P. (2011). From humans to rats and back again: Bridging the divide between human and animal studies of recognition memory with receiver operating characteristics. *Learning & Memory*, *18*, 519–522. doi:10.1101/lm.2214511
- Lee, W. (1963). Choosing among confusably distributed stimuli with specified likelihood ratios. *Perceptual and Motor Skills*, *16*, 445–467. doi:10.2466/pms.1963.16.2.445
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, NJ: Erlbaum.
- Malmberg, K. J. (2002). On the form of ROCs constructed from confidence ratings. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 380–387. doi:10.1037/0278-7393.28.2.380
- Malmberg, K. J., & Annis, J. (2012). On the relationship between memory and perception: Sequential dependencies in recognition memory testing. *Journal of Experimental Psychology: General*, *141*, 233–259. doi:10.1037/a0025277
- Malmberg, K. J., & Xu, J. (2006). The influence of averaging and noisy decision strategies on the recognition memory ROC. *Psychonomic Bulletin & Review*, *13*, 99–105. doi:10.3758/BF03193819
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97.
- Mueller, S. T., & Weidemann, C. T. (2008). Decision noise: An explanation for observed violations of Signal Detection Theory. *Psychonomic Bulletin & Review*, *15*, 465–494. doi:10.3758/PBR.15.3.465
- Nosofsky, R. M. (1983). Information integration and the identification of stimulus noise and criterial noise in absolute judgment. *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 299–309. doi:10.1037/0096-1523.9.2.299
- Pelli, D. G. (1997). The Video Toolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442. doi:10.1163/156856897X00366
- Petrusic, W. M., & Baranski, J. V. (2003). Judging confidence influences decision processing in comparative judgments. *Psychonomic Bulletin & Review*, *10*, 177–183. doi:10.3758/BF03196482
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*, 347–356. doi:10.1111/1467-9280.00067
- Stretch, V., & Wixted, J. T. (1998). Decision rules for recognition memory confidence judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1397–1410. doi:10.1037/0278-7393.24.6.1397
- Swets, J. A., Tanner, W. P., Jr., & Birdsall, T. G. (1961). Decision processes in perception. *Psychological Review*, *68*, 301–340. doi:10.1037/h0040547
- Symonds, P. M. (1924). On the loss of reliability due to coarseness of the scale. *Journal of Experimental Psychology*, *7*, 456–461. doi:10.1037/h0074469
- Treisman, M., & Williams, T. C. (1984). A theory of criterion setting with an application to sequential dependencies. *Psychological Review*, *91*, 68–111. doi:10.1037/0033-295X.91.1.68
- Tullis, J. G., & Benjamin, A. S. (2012). The effectiveness of updating metacognitive knowledge in the elderly: Evidence from metamnemonic judgments of word frequency. *Psychology and Aging*, *27*, 683–690.
- Turner, B. M., Van Zandt, T., & Brown, S. (2011). A dynamic, stimulus-driven model of signal detection. *Psychological Review*, *118*, 583–613. doi:10.1037/a0025191
- Van Zandt, T. (2000). ROC curves and confidence judgments in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 582–600. doi:10.1037/0278-7393.26.3.582
- Wickelgren, W. A. (1968). Unidimensional Strength Theory and component analysis of noise in absolute and comparative judgments. *Journal of Mathematical Psychology*, *5*, 102–122. doi:10.1016/0022-2496(68)90059-X
- Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*, *114*, 152–176. doi:10.1037/0033-295X.114.1.152
- Wixted, J. T., & Stretch, V. (2004). In defense of the signal detection interpretation of remember/know judgments. *Psychonomic Bulletin & Review*, *11*, 616–641. doi:10.3758/BF03196616
- Yonelinas, A. P. (1999). The contribution of recollection and familiarity to recognition and source-memory judgments: A formal dual-process model and an analysis of receiver operating characteristics. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 1415–1434. doi:10.1037/0278-7393.25.6.1415

Received December 7, 2012

Revision received January 2, 2013

Accepted January 7, 2013 ■