Running head: FROM ROC CURVES TO PSYCHOLOGICAL THEORY

From ROC Curves To Psychological Theory

Jeffrey N. Rouder, Jordan M. Province, April R. Swagman, Jonathan E. Thiele

University of Missouri

Jeff Rouder

rouderj@missouri.edu

Abstract

ROC plots are a common data representation for drawing conclusions from behavioral data about underlying mental representations and processes. According to broadly accepted conventions, the curvature, symmetry and detailed patterns of single curves are indicative of whether processing is mediated by continuous latent strengths, by discrete states, or by a dual-process mixture of the two. These conventions, however, are based on difficult-to-justify and untestable assumptions, and, in our opinion, are in need of substantial revisions. We show here that without these assumptions, there is no link between ROC patterns and processing architecture, and that latent-strength and discrete-state models can predict any ROC curve. We show that although these models may predict any single ROC curve in isolation, there is much constraint when considering a family of related curves, say those drawn by manipulating a stimulus-strength variable. The relations between different ROC curves rather than the shape or symmetry of any one are informative of underlying processing. We define here two novel ROC properties, *discrete-state representability* and *shift representability*, which are signatures of processing mediated by discrete states and continuous latent-strength models. We show how these properties provide for superior principled inference about processing without recourse to tenuous assumptions, and that existing recognition memory data from our lab are compatible with a discrete-state model.

From ROC Curves To Psychological Theory

A goal of cognitive psychology is to describe the processes and representations underlying mental phenomena. One approach to exploring these processes and representations is to ask people to respond to difficult stimuli for which they will make errors. To account for these errors, cognitive theories often stipulate that cognitive processes and representations are perturbed by internal, latent noise. One popular theoretical theme is what we term here graceful degradation, and a quintessential example of a graceful degradation is the theory of signal detection (Green & Swets, 1966). A stimulus is represented by latent strength. This strength faithfully represents physical strength variables, such as duration and intensity, except it is perturbed by noise. Another example of graceful degradation is the diffusion model of perception (Ratcliff, 1978) where the stimulus effect is gradually perturbed in real time. Further examples of graceful degradation include multidimensional representational theories where the effect of the stimulus on a given trial is perturbed from a veridical point in a multidimensional space (Ashby, 1992) and ideal-observer models where participants ideally account for perturbations (Attneave, 1959; Ma, 2012; Ma, Beck, Latham, & Pouget, 2006; Sims, Jacobs, & Knill, 2012). The key concept in graceful degradation is that the latent representation reflects the stimulus on each trial, even if this representation is degraded by some degree from internal noise. Most models of cognition describe graceful degradation of stimulus information. A contrasting theme to graceful degradation is all-or-none loss, and an example of all-or-none loss is the high-threshold model (Blackwell, 1953) where an item is processed fully or not at all. Another example of all-or-none loss is Zhang and Luck's (2009) sudden-death theory of forgetting, where, at one point in time, all stimulus information about a certain item is lost. This distinction between graceful degradation or all-or-none loss remains topical in many domains including recognition memory (e.g., Broder & Schutz, 2009; Dube & Rotello, 2012; Yonelinas & Parks, 2007; J. Wixted & Mickes, 2010), working memory (e.g., Bays & Husain, 2008; Miller, 1956; Rouder, Tuerlinckx, Speckman, Lu, & Gomez, 2008), letter and word identification (e.g., Massaro, 1998; Massaro & Oden, 1979; Rouder, 2004; Townsend, 1971; Townsend & Landon, 1982), and perception (Ma, 2012).

One classic approach to assess whether responses in a task are better accounted for by continuous or all-or-none loss is to inspect the resulting *receiver operating characteristic* (ROC) plot. ROC plots, which are described in the next section, provide a graphical representation of how response choices change with varying experimenter-imposed response demands. In the experimental psychology literature, there are well-established conventions for drawing conclusions about processing architectures from the data patterns in these plots. In fact, these conventions form a common core knowledge in the field: they have been canonized in articles and books (selected examples include Egan, 1975; Green & Swets, 1966; Macmillan & Creelman, 2005; Yonelinas & Parks, 2007), and are often part of the curriculum for graduate students to master en route to becoming successful cognitive psychologists. We refer to these conventions for interpreting ROCs as *the standard story.* In this paper, we argue that the standard story is based on a collection of tenuous assumptions; some of these are unnecessarily restrictive while others have no psychological content. Consequently, the standard story is in need of revision, and previous ROC-based conclusions about continuous or all-or-none loss should be heavily qualified if not completely ignored.

The Standard Story

It is helpful to describe the standard story with reference to a specified task. Here, we refer to an *old-new recognition memory* task, though the standard story applies analogously for other cognitive tasks. In an old-new recognition-memory task, the participant first studies a list of items. Then, after a suitable delay, these items as well as new ones are presented at test, one at a time. For each test item, the participant judges whether it was *old*, that is, previously studied, or *new*, that is, novel. The rate of old-item responses comprises the data of interest: the *hit rate* is the rate of old-item responses conditional on old-item test trials, and *the false alarm rate* is the same conditional on new-item test trials. An ROC plot is a graph of the hit rate as a function of the false alarm rate.

A key element for drawing inferences about processing in ROC analysis is the *isosensitivity curve,* which is sometimes just referred to as an ROC curve. To draw an ROC curve, the experimenter imposes varying response demands on the participant. For example, in one condition, a neutral condition, the experimenter might reward the participant 5¢ for each hit and correct rejection (where correct rejections are new-item responses conditional on new-item trials and are the complement of a false alarm). In a different condition, the experimenter might reward the participant 10¢ for each hit and only 1¢ for each correct rejection. In such a case, participants may be biased to produce old-item responses as these are better rewarded. Likewise, in a third condition, the inverse response demand may be imposed: participants could be rewarded 1¢ for each hit and 10¢ for each correct rejection. Each condition yields a pair of hit and false-alarm rates or a point in the ROC plot. An ROC curve is composed of several such points where the points correspond to conditions of different response bias. The standard story in experimental psychology is that different processing models yield different predictions for the shape and asymmetries of ROC curves. Figure 1 captures the standard story for a number of formal models as discussed below:

I. A High-Threshold Model of All-or-None Information Loss. Figure 1A shows a double high threshold model (Egan, 1975). When the item is old (left tree), the item may either be detected from memory (with probability d_s), or if detection fails (with probability $1 - d_s$), the participant guesses. If the item is detected, then the participant responds "old;" if not, then the participant responds "old" with probability g and responds "new" with probability $1 - g$. When the item is new (right tree), the item may be detected as new with probability d_n, or the participant may guess with probability $1 - d_n$. In this model, response biases from different demand conditions are captured by g, the probability of producing an old-item response when guessing. Let h_i, f_i be the hit and false-alarm rate for the ith bias condition. Then,

$$h_i = d_s + (1 - d_s)g_i \tag{1}$$

$$f_i = (1 - d_n)g_i, \tag{2}$$

where g_i is the bias from the ith condition. As the payoffs increasingly favor old-item responses, g becomes greater. To understand the ROC predictions of the model, the hit rate is expressed as a function of the false alarm rate. After substitution,

$$h_i = \left(\frac{1 - d_s}{1 - d_n}\right) f_i + d_s.$$

This equation describes a straight line, that is, hit rates are a linear function of false alarm rates (see Figure 1B).

II. The Signal-Detection Model of Graded Information Loss. Figure 1C shows the signal-detection model where the mnemonic effect of test items are represented by graded normal distributions on a latent strength axis. These strengths are centered at zero with unit standard deviation when the test item is new. In the most popular version of this model, this distribution increases in mean and variance for old items. Participants reach decisions by setting a criterial bound on this strength axis: if the strength on a test trial is greater than bound, then an item is judged as old, otherwise it is judged as new. The hit and false alarm rates are

$$h_i = \Phi((d' - c_i)/\sigma) \tag{3}$$

$$f_i = \Phi(-c_i), \tag{4}$$

where d' and σ are the mean and standard deviation of strength for old items, and c_i is the criterial bound for the ith condition. The function Φ is the *cumulative distribution function* (CDF) of the normal distribution, and the supplement to this article provides a brief review of the relationship between *density functions*, CDFs, and the inverse of CDF, called *quantile functions*. The above equations directly lead to predictions for ROC curves: substituting (4) into (3) yields

$$h_i = \Phi([d' + \Phi^{-1}(f_i)]/\sigma).$$

The equation implies a familiar curved ROC contour, which is shown in Figure 1D. The asymmetry of the curve relative to the negative diagonal serves as a measure of the variance of the old-item distribution. Signal detection is considered a simple, elegant, single-process account of phenomena and serves as a contrast for more complex, dual-process or separate multiple-systems accounts.

III. Dual Process Models: Modern memory theories often stipulate that memory is governed by several distinct processes working in parallel (Schacter & Tulving, 1994). The most common distinction is between *familiarity*, a quick, unconsciousness, automatic process, and *recollection*, a slow, conscious, and deliberative process (Mandler, 1980; Jacoby, 1991), and this distinction has had far reaching effects throughout cognitive psychology, cognitive neuroscience, and social psychology. Yonelinas (1994) provides a dual-process account of recognition memory whose predictions have become part of the standard story. Accordingly, old items may be recollected, which leads to an "old" response, or may be mediated by familiarity, which is modeled as a signal detection process. New items cannot be recollected and are mediated only by familiarity. The hit and false alarm rates for the Yonelinas dual-process model are

$$h_i = r + (1 - r)\Phi(d' - c_i) \quad (5)$$

$$f_i = \Phi(-c_i), \quad (6)$$

where d', and c_i are defined previously and r, the new parameter, is the probability of recollection. In this model, d' serves as a measure of the strength of familiarity. Figure 1E shows a graphical representation of the model. The model looks like the high-threshold model in this representation, but it is quite different. The main difference is on the lower branches of the processing trees, and the terms $\tilde{h}_i$ and $\tilde{f}_i$ are the hit and false alarm rates from the signal detection model. These differ from the guessing in the high-threshold model in that $\tilde{h}_i$ depends on stimulus strength as well as response criteria. Consequently, the dual-process model predicts curved ROC plots as shown in Figure 1F.

There have been literally hundreds of experiments that draw ROC curves. There are two nearly ubiquitous regularities for old-new recognition memory tasks: First, ROCs by and large are not straight lines, that is, they display a notable degree of curvature. Second, ROCs show an asymmetry much like that in Figure 1D and Figure 1F. The following two conclusions follow from the standard story: **1.** Because ROCs are curved, high-threshold models are not tenable. **2**. Because ROCs are curved and asymmetric, the unequal-variance signal detection model and dual process models are tenable. To tell whether processing is mediated by one or more-than-one process, some researchers study very fine details of ROC curves to see if the degree of curvature is better predicted by the unequal-variance signal detection model, which would indicate a single process, or by the dual-process model, which would indicate two processes (Yonelinas & Parks, 2007; Howard, Bessette-Symons, Zhang, & Hoyer, 2006).

The standard story provides for a focus on the shape and asymmetry of single ROC curves. By focusing on whether individual curves are straight or truly curved, or have asymmetry, researchers can draw fine-grained conclusions about process. We argue here that this focus is misplaced and that discrete-state and signal-detection models can produce ROC curves of any shapes and asymmetries. We propose as an alternative that the focus should be on the relations between curves, and relational properties provide for more principled and stringent assessments of underlying process.

Revising the Link Between Discrete-State Models and ROC plots

What's Wrong With The Standard Story

In this section, we argue that one part of the standard story, the claim that models of all-or-none information loss predict straight line ROCs, is too narrow. The high-threshold model is an example of a discrete-state model where there is one state for detection and another for guessing. In the guessing state, participants distribute their responses, and this distribution changes depending on response demands. For example, guesses favor old-item responses when hits are better rewarded than correct rejections and favor new-item responses when correct rejections are better rewarded than hits. The responses under detection, in contrast, are fixed rather than distributed. If an item is detected as either old or new, then the corresponding response is always made. We call this assumption where a state always leads to a specific response as the *certainty assumption.*

Luce (1963) was the first to challenge the certainty assumption as being too restrictive. He posited that responses are distributed in all states, not just in the guessing state. It is possible, reasoned Luce, that people entered states outside of their volition and occasionally in error. Although people may not be able to control how or when they entered states, they certainly can control what responses they make conditional on a state. They therefore may adjust the distribution of responses to meet real or even perceived response demands. For example, in a condition with payoffs that favor a hit, participants may occasionally produce an old-item response even in the detect-new state. Broadbent (1966) offered the same critique of the certainty assumption—he describes the possibility that participants switch response strategies from trial to trial conditional on a detection state. Malmberg (2002) and Erdfelder & Buchner (1998) make a similar point in the context of confidence ratings; participants may not endorse the highest degree of confidence conditional on a detect state.

The certainty assumption is seemingly incompatible with the well known and well replicated phenomena of *probability matching* (Myers, 1976; Shanks, Tunney, & McCarthy, 2002; Vulcan, 2000). Starting in the 1950s, learning theorists were stunned that people routinely chose suboptimal strategies in simple decision learning problems. Consider, for example, a task where the participant is presented a light that turns either red or green on each trial, and must make either a left-hand or right-hand response. The participant is not told which color is associated with which response, but can easily learn so from feedback. If the feedback is consistent, that is, if a specific color is always associated with a specific response, then learning is fast, and participants are very consistent in their response, which is optimal. The interesting case is when the feedback is variable. For instance, conditional on the light turning red, the correct response is "left" on 3/4 of such trials and "right" on the remaining 1/4, and conditional on on the light turning green, the correct response is "right" on 3/4 of such trials and "left" on the remaining 1/4. In this case, the optimal strategy is to always respond "left" for red lights and "right" for green lights, and this strategy leads to an overall accuracy of .75, which serves as an upper limit. This strategy corresponds to a certainty assumption—when detecting a red light, the response should always be "left," and when detecting a green light, the response should always be "right." People, however, do not obey this certainty assumption. Instead, they tend to probability match (Healy & Kubovy, 1981). That is, when presented with the red light, they choose the left response 3/4 of the time and the right response 1/4 of the time, and when presented with the green light they reverse these proportions. The certainty assumption, therefore, does not hold in a broad array of contexts even when it describes the ideal response rule to maximize accuracy.

Given that people often violate certainty assumptions, it is judicious to consider the

ROC predictions of all-or-none loss without the certainty assumption. Figure 2A provides an example of a discrete state model where certainty is not assumed. Conditional on detection of an old item, the probability of "old" and "new" responses is a_i and $1 - a_i$, respectively; conditional on detection of a new item, the same probabilities are b_i and $1 - b_i$. Hit and false alarms are given by

$$h_i = d_s a_i + (1 - d) g_i, \tag{7}$$

$$f_i = d_n b_i + (1 - d) g_i. \tag{8}$$

In the above model, detection and guessing are treated in an analogous fashion. Not only are responses distributed, the distributions depend on the payoffs. For each payoff condition, there are three free parameters g_i, a_i, and b_i that may be freely adjusted to meet response demands imposed by the payoff structure. Because detection will lead to many correct responses, it is expected that in general $a_i > .5$ and $b_i < .5$, and it is reasonable to insist that $a_i > b_i$. In this model, the guessing state leads to the same responses regardless of whether the item was old or new; that is, the responses are em stimulus independent. Detection, in contrast, leads to stimulus-dependent responses—there is a different distribution of responses dependent on whether an old or new item was detected.

Because three parameters (g_i, a_i, b_i) may be adjusted with changing response demands from payoffs, the model is no longer constrained to predict straight-line ROC curves. To see how the model may predict curvature, consider the case where g_i, a_i, and b_i are each linked to a common bias parameter z_i, which reflects the payoffs in the ith condition. To constrain g_i, a_i, and b_i to be between 0 and 1, consider the following relationship between the log-odds of each parameter and z_i:

$$\log\left(\frac{g_i}{1-g_i}\right) = g_0 + z_i, \tag{9}$$

$$\log\left(\frac{a_i}{1-a_i}\right) = a_0 + z_i, \tag{10}$$

$$\log\left(\frac{b_i}{1-b_i}\right) = b_0 + z_i. \tag{11}$$

Figure 2B shows the ROC curve as z_i is varied across a wide range. As can be seen, this model well accounts for curved ROCs.

The fact that the discrete-state model may account for curvature is not so surprising given that the responses from any or all states may be shaped to account for response demands. In fact, the above model is so richly parameterized that it can account for any observed ROC curve! To see this, suppose the experimenter manipulates payoffs through I levels to produce I different points that lie on a single curve. The model may account for these I points perfectly. The equivalence may be seen by setting $d_s = d_n = 1$, a_i to the observed hit rate for the ith payoff condition, and b_i to the observed false alarm rate in the ith payoff condition. Therefore, the discrete-state model can always account for any single ROC curve from binary-choice (yes/no, 2AFC) paradigms. The comparable point that discrete-state models may account for any single ROC curve in confidence ratings has been made previously with analogous logic by Broder & Schutz (2009), Erdfelder & Buchner (1998), and Malmberg (2002).

We anticipate that readers may be skeptical of dispensing with the certainty assumption, especially in binary-choice paradigms, for the following reasons: First, readers may ask about the wisdom of advocating for a model that does not constrain the shape of ROC curves. We show below that the discrete-state model does provide for strong and testable constraints on the relationship among several different ROC curves rather than any one in isolation. In fact, these discrete-state constraints are stronger than the straight-line property even though they provide for a test of a more general model.

Second, readers may ask whether the concept of detection is meaningful if certainty is not assumed. We think so. The key point is that discrete-state models without certainty embed a complete-information loss theory and remain in stark contrast to gradual-degradation models. Those unhappy with the term "detection" for states without certainty are welcome to rename detect and guessing states as states where responses are stimulus dependent and stimulus independent, respectively. We see no reason to exclude *a priori* discrete-state models without the certainty assumption as they are theoretically important and, as will be shown, highly constrained in appropriate paradigms.

Conditional Independence and ROC patterns

Although the discrete-state model does not place any constraint on any one individual curve, there is strong and testable constraint on the relationship among several different ROC curves. The key psychological constraint in a discrete state model may be termed *conditional independence.* Conditional independence is a statement about the select influence of stimulus factors. Stimulus factors will certainly affect the probability of entering a given state—strong stimuli, for example, will result in increased probabilities of entering detect states. The key constraint is that once a participant enters a state, the distribution of responses does not depend on the stimulus. Consider an experiment in which stimuli are repeated at study, and the number of repetitions is manipulated systematically. Repetition will certainly affect the probability of detection with higher values of d_s corresponding to increased repetition. Conditional on an item being detected, however, the distribution of responses does not depend on repetition, the distribution in fact only depends on a_i. If a participant fortuitously enters the detection state for an item repeated once at study, then the distribution of responses will be no different than for the detection of an item presented multiple times at study. Stimulus factors such as repetition only affect the probability of entering states; response factors such as payoffs only affect the probability of specific responses conditional on states. Conditional independence is, therefore, the core psychological property of discrete mental states.

This psychologically-motivated property of conditional independence leads directly to constraints on families of ROC plots. Figure 2C provides a graphical demonstration of the constraint, and there are hypothetical ROC plots for three repetition conditions and five payoff conditions. The pattern of constraint is easiest to describe when old-item and new-item detection are the same; i.e., $d_s = d_n = d$. This equality condition is most plausible in 2AFC paradigms where detection of items in one position, say to the left of fixation, is the same as the detection of items in the other position. The curve with points (a, b, c, d, e) is from when $d = 0$, and, consequently, the points lie on the diagonal. The specific locations of these points on the diagonal reflect the values of $\boldsymbol{g}$, where $\boldsymbol{g} = (g_1, \ldots, g_I)$, the collection of guessing parameters. The curve with points (A, B, C, D, E), is from the detection state $(d = 1)$, and the specific location of the points reflect the values of $\boldsymbol{a}$ and $\boldsymbol{b}$, the collection of response probabilities from the detect-old and detect-new states, respectively. Connecting the points are *isobias* lines $\vec{aA}$, $\vec{bB}$, $\vec{cC}$, $\vec{dD}$ and $\vec{eE}$. The discrete-state model constraint is that points for any observed ROC must lie on specific locations on these isobias line reflecting the value of d. Figure 2C shows three ROC curves for the three repetition conditions (dashed lines). For one repetition, the value of d is .25, and the first point is 25% up the isobias line from $\vec{aA}$; the second point is 25% up the isobias line from bB, and so on. For the two repetitions, the value of d is .5, and each points is half way up the respective isobias line, and for four repetitions, the value of d is .75, and each point is 75% up the respective isobias line. Figure 2D provides a different example of curves that obey conditional independence. The curves for $d = 0$ (points a, b, c, d, e) and $d = 1$ (points A, B, C, D, E) are different than the corresponding curves in Figure 2C, reflecting differences in parameters $\boldsymbol{g}$, $\boldsymbol{a}$, and $\boldsymbol{b}$. Nonetheless, the same constraint holds for observed ROCs. Points not only lie on $\vec{aA}$, $\vec{bB}$, $\vec{cC}$, $\vec{dD}$ and $\vec{eE}$, but are a constant proportion that reflects d. We refer to ROC families that obey the conditional independence constraint as *discrete-state representable.* The supplement provides the proof that discrete-state models lead to these ROC constraints. There are conditional-independence constraints when $d_s \neq d_n$, though these are more difficult to describe graphically though they still may be assessed by model comparison methods.

The standard story that discrete-state models predict straight-line ROCs is based on the certainty assumption. When this assumption is relaxed, the discrete-state model may predict any single ROC curve in isolation. There is strong constraint from discrete-state models without the certainty assumption, and this constraint holds across a family of several ROC curves. Previous claims about the inapplicability of discrete-state models are based on observed curvature of empirical ROC curves (Dube & Rotello, 2012; Kinchla, 1994). These claims are overstated because there is no reason to impose the certainty assumption, especially since discrete-state models without the certainty assumption make testable predictions.

Conditional Independence In Confidence-Ratings Distribution

Discrete-state representability yields an equivalent property for confidence ratings that we have found particularly easy to assess in practice. In Province and Rouder (2012), we showed participants a long study list of words, with some words repeated once, twice, or four times. On most test trials, participants were shown a studied word either to the left or right of fixation, and a lure on the other side. Participants judged their confidence by moving a slider between an anchor on the left which indicated that the participant was sure the left word was old and an anchor on the right which indicated that the participant was sure the right word was old. On a small minority of test trials, however, participants were shown two new words and asked to just which was old. This condition is useful for understanding responses from the guessing distribution should it exist. We refer to it as studying a word for zero repetitions.

Figure 3 shows the conditional independence constraint on the distributions of confidence ratings. In this case, *FROG* was the studied word and *TABLE* is the lure. The top row shows hypothetical confidence-ratings distributions for a guess state (Panel A) and a detect state (Panel B). The middle row shows a few response distributions for the 0, 2, and 4 repetition conditions. In the 0-repetition condition (Panel C), all responses come from the guessing state. For the two-repetition condition, half the responses are from guessing and the other half are from detection. For the four-repetition, the mixture is 75% from detection and 25% from guessing. To best see the discrete-state constraint, it is useful to plot these mixtures on the same graph, and Panel F shows the pattern across all three repetition conditions. The mixture from 0 and 4 repetitions are shown in the usual orientation while the mixture from the 2-repetition condition is projected downward, that is, greater mass corresponds to a larger negative deflection. The signature of the mixture pattern, that the modes line up, is transparent in this representation. Figure 3G shows the contrasting case where the confidence ratings come from a typical signal-detection model. Here, the distributions of confidence gradually move toward the old-item as the number of repetitions is increased. It is straightforward to show that the mixture constraint on the distributions of confidence ratings is mathematically equivalent to the above discrete-state representability constraint on ROC plots, and a proof is provided in the supplement.

Province and Rouder (2012) collected confidence ratings across repetition conditions in recognition memory and found that participants by-and-large obeyed this discrete-state representability constraint. Figure 4 shows an example from a selected participant, and this participant is typical in that the modes of the distributions tended to line up. Of the 89 participants that performed the task, the discrete-state conditional independence model could not be rejected for 81 of them (92%). Moreover, for 76 of the 89 participants (85%), a simple discrete-state model outperformed a similarly parameterized normal-distribution signal detection model.

Province and Rouder provide a test of conditional independence, that response distributions from given states are not dependent on stimulus strength. Broder, Kellen, Shutz, and Rohmaler (2013) provide a complementary finding. These researchers assessed whether the conditional distribution of responses change in response to experimenter response demands (changing base rates). Not only do they do so in reasonable ways, the probability of stimulus detection remained invariant to the base rate manipulation. Taken together, the Province and Rouder and Broder et al. invariances provide noteworthy support for the discrete-state account of recognition memory.

Potential Difficulties With Discrete-State Models

In the above development, curved ROCs were accounted for by relaxing the certainty assumption. It seems reasonable to expect that certainty may hold, however, when detection is very good. Consider, for example, a perception experiment where the participant has to identify a letter that is briefly presented and subsequently masked. If the letter is presented sufficiently briefly, then accuracy is not near ceiling and the participant may violate certainty to account for payoffs and other response demands. These violations mean that $a < 1$ and $b > 0$ in the discrete-states model. If the letter is presented longer, say for 1 second, then performance should be at or near ceiling. Performance, however, is only near ceiling when $a \approx 1$ and $b \approx 0$. Consequently, it is reasonable to expect that a and b may depend on duration, and this dependence is in direct violation of conditional independence. If conditional independence is violated, then the discrete-state representability constraint may not hold, and, in fact, the discrete-state model may not be falsifiable. Luce (1963) noted a similar problem with his low-threshold model. Our recommendation is that researchers hold conditional independence provisionally tenable, especially for experiments with sub-ceiling performance in all conditions. It may be that conditional independence fails in the extreme, but understanding where it holds and where it fails will undoubtedly drive better theory development.

Revisiting Signal Detection Models and ROC plots

The basic findings that ROCs are curved and asymmetric have been interpreted according to the standard story as support for the unequal-variance normal distribution model and the dual-process model. In this section we re-examine the standard story for signal-detection models. One of our main questions is whether ROC curves may be used to assess whether processing is mediated by a single process or multiple processes, and after re-examining the standard story, we develop a new property, called *shift representability*, that should prove useful for inferring whether processing is mediated by one or more than one process.

What's Wrong With The Standard Story For Signal-Detection Models?

To assess the wisdom in the standard story, it is useful to separate the core of signal detection models from secondary assumptions that have no psychological content. The core is that decisions are made by comparing a scalar-valued latent-strength value to a criterion. These strength values vary randomly from trial to trial, and the underlying distributions depend on whether the test item is old or new. Added to this core are secondary assumptions, for example, that the underlying distributions are normally distributed. Some parametric assumption seems needed, and the normal is as convenient as any. Many other distributions, however, could do just as well. Figure 5 shows a few alternatives including the gamma and an ex-Gaussian model and their predicted ROC curves. These alternative parametric models make similar but not identical predictions. Because the predictions are not identical, it might seem desirable to test which parametric form is best. We think, however, such tests, even if they could be done, are of little value because they contrast the suitability of one contentless, arbitrary assumption against another.

A more valuable question is whether there exists some signal-detection model that can account for any observed ROC curve. Restated, the question is whether the core of the signal detection model—the notion that there is a scalar latent strength which is compared to criteria—is unfalsifiable. The unfortunate answer is yes. As a matter of logic, there always exists a signal detection model that accounts for any ROC curve. Figure 6 provides a schematic view of the proof. Figure 6A shows an equal-variance normal signal detection model and Figure 6B shows the corresponding ROC curve. Highlighted is a point on strength space at .84, which corresponds to a false alarm rate of .20 and a hit rate of .56. There is a critical relationship between the derivative of the ROC curve at this point and the densities of the signal and noise distributions at the corresponding strength. The derivative is the slope of the tangent line, shown, at the point, and in this case it is 1.4, meaning that at this point hits increase 1.4 times as quickly as false alarms. This slope equals the ratio of the densities at the point; that is, the density of the old-item distribution is 1.4 times as large as that of the new-item distribution. The equality between the derivative and ratio of densities is well known (Egan, 1975), and the proof is provided in the supplement. This equality implies that ROC curves tell us about the ratio of densities rather than the form of distributions. Figure 6C shows a different signal detection model where the new-item distribution is a uniform rather than a normal. Because the uniform has density of 1.0, the slope of the ROC curve in this case corresponds to the density of the old-item distribution. The model in Figure 6C corresponds exactly to the ROC curve in Figure 6B, and the highlighted point in Figure 6C corresponds to the highlighted point in Figure 6B. In particular, the highlighted point in Figure 6C yields a false-alarm rate of .2, and the ratio between the densities is 1.4. The consequence of this equality is that any ROC curve may be accounted for by a uniform new-item distribution and an old-item distribution whose density function is the derivative of the ROC. That is, there always exists parametric forms that can account for any ROC curve, and the core of the signal detection model is unfalsifiable.

The core of the signal-detection model accounts for any single ROC curve. Previously, we showed that the class of discrete-state models with the conditional independence assumption yields constraints across a family of ROC curves. Unfortunately, the same cannot be said about the core of signal detection. The core yields no testable constraint about the relationships among several curves. As a matter of mathematical logic, for any family of curves, say those from different repetition conditions, there always exist a set of new-item and old-item distributions that will perfectly account for these curves. For any data set, the core always remains unfalsifiable. Unfortunately, this important point is not part of the standard story.

The development in Figure 6 suggests that there is a universal signal detection representation framework for all ROC curves where the noise is distributed as a uniform on $(0, 1)$. Figure 7 provides some examples of this universal representation. The first three rows show the cases for the equal-variance normal distribution, the unequal-variance normal distribution, and the gamma distribution[1], respectively. The left column, labeled "Common Representation," shows the usual depiction where the noise is distributed as a normal or as a gamma. The center column, labeled "Universal Representation," shows the case where the noise is uniform. The universal and common representations in a row are fully equivalent in that they predict exactly the same ROC curves. No data set could discriminate between them.

The lack of falsifiability of the core of signal detection means that models not usually considered signal detection models will have universal signal detection representations. The fourth row shows the case for a single high-threshold model (Blackwell, 1953), which is a discrete-state model. Because this model makes ROC predictions, it has an equivalent universal signal-detection representation, which is shown (Banks, 1970; Slotnick & Dodson, 2005). The fifth row shows the universal signal-detection representation of the Yonelinas two-process model. This signal detection model produces ROC predictions that are identical to the Yonelinas model. Changes in familiarity affect the distribution below a strength of 1.0; changes in recollection change the relative weight of strength above and below 1.0 without affecting the shape of either component.

This lack of falsifiability of the core of signal detection stands in contrast to how signal detection is typically conceived. Signal detection is offered as a simple model that captures the workings of a unified process by representing stimuli on a unidimensional scale. In the literature, it serves as a simple contrast for more complex generalizations such as the two-process model or multivariate detection models (Ashby, 1992; DeCarlo, 2003; Rotello, McMillan, & Reeder, 2004; Yonelinas & Parks, 2007). The entirety of the setup, that signal detection is simple and that these generalizations are more complex, depends critically on the assumption that strength or familiarity is distributed as a normal, which is ancillary to the core of signal detection.

Some researchers have argued that the the normal parametric assumption is not arbitrary, but instead is the principled choice (e.g., Wixted & Mickes, 2010). The justification is provided by Swets (1961), who invokes the central-limit theorem. The claim is that the latent strength distributions are the sum of many small, independent and identically distributed deviations, and that this sum is therefore well approximated by a normal. The argument, however, is not compelling because it is unknown what these small, independent addends might be. In fact, we know of no neurophysiological correlates that would bolster the plausibility of the claim. Moreover, as a a counter to such a justification, we offer the case of response times (RT). It seems plausible that response times are the sum of small latencies, and by the same logic, RT therefore should be normally distributed. RTs however are not normal and display a characteristic positive skewness (Luce, 1986). In summary, we find the assertion that the normal is the proper parametric form to be uncompelling, and find it troubling that so much theoretical distinction and assessment is critically based upon it. We provide a more compelling nonparametric set of constraints subsequently.

The lack of falsifiability of the core of signal detection requires researchers to be mindful and deliberate on how they constrain their models to obtain testable predictions. One principled approach is to place process-model constraints on signal detection representations as exemplified by REM (Shiffrin & Steyvers, 1997) and TODAM (Murdock, 1993). These models specify psychologically motivated processing architectures which place constraints on the shape of individual ROC curves, and in some cases, constraints on the relationships between curves. One difficulty, however, is that these models are often overly specific and quite latent given the relatively impoverished nature of ROC data. It often remains unclear how the details of the processing architectures link directly to patterns in the data. Consequently, the assessments of these models have been few, and researchers have provided proof-of-concept demonstrations but not critical tests of experimental control of parameters and mechanisms. In the next section, we provide a different approach to understanding constraint in processing—a focus on the constraint in the relationship among several ROC curves.

Simple and Complex ROC Families of Isosensitivty Curves

The standard story provides a focus on the shapes of individual isosensitivity curves, and, consequently, researchers have repeatedly noted the curvature and asymmetry of these curves. In the previous sections, we noted that these properties are not diagnostic. Consideration of the discrete-state model provides motivation for a focus on the structure in the relationship between curves rather than a focus on any single one. To date, the field has not provided summaries of how ROC curves differ from one another, how such differences vary across conditions, tasks, and domains, and how such differences may constrain theory development.

One approach to understanding the relationships among ROC curves is to focus on parsimony. Figure 8A-C and Figure 9A-B provides examples of families of ROC curves. Each panel shows a family of ROC curves, and each curve in a family comes from a different experimental condition. Figure 8A shows a family where each curve is curved and symmetric. Figure 8B shows a family where each curve is asymmetric to roughly the degree seen in recognition memory data. Figure 8C shows a family where each curve is a straight line. Although these families differ in curvature and symmetry, these properties are not diagnostic. What is diagnostic is the simplicity of the relations between curves. In each panel, changes from one curve to the next can be described simply, and the curves within each plot form a simple, orderly, and parsimonious family. There are elegant relational constraints, such that two curves never cross, exhibited by these families. If such relational simplicity is observed, it is assuredly a strong diagnostic marker that simple models—say those that model condition differences with a single parameter—are warranted. Conversely, such a marker would indicate that complex models—say those that model condition differences with multiple parameters or multiple processes, are overly complex.

The contrasting cases are shown in Figure 9A-B. Here the families themselves are far more complex, and the relationship between any two curves in a family cannot be described easily. If such relational complexity is observed, it is assuredly a strong diagnostic marker that complex models—say those that model condition differences with multiple parameters or multiple processes are indicated, and that simpler models are too simple. To our knowledge, there are no systematic, model-free treatments of the relationships among empirical ROCs across collections of experimental conditions. This lack of exploration into the relational properties among curves contrasts with our large corpus of knowledge about the shapes and asymmetries of individual curves. In many ways the situation is sadly ironic: we know much about the least diagnostic features of ROC curves and little about the most diagnostic features.

Shift Representability of ROC curves

The contrast between the families in Figure 8A-C and Figure 9A-B provides an informal sense of simplicity and complexity across families of ROC curves. Here we formalize the notion, and do so with recourse to the concept of linearizing ROC curves. Readers may be familiar with the concept of $z - ROC$ curves, which are used to linearize ROCs under the assumption that latent strength distributions are normally distributed. To linearize these curves,, a specifc function is applied to both hit and false alarm rates. The resulting values, termed z-hits and z-false alarms, are plotted. Figure 8D shows an example of the transforming function, and Figure 8G shows the results when applied to the hit and false-alarm rates in Figure 8A. The transformed curves are now straight lines and they have a common slope. They differ only in intercept, an easy-to-see, one-parameter difference. The linearization recodes the simplicity of the field of ROC curves in Figure 8A into that in Figure 8D. The story is the same for the families in Figure 8B and Figure 8C; there exists a simple function that linearizes all the curves in the family, shown in Figures 8E and 8F, respectively. When these curves are linearized, the resulting lines all have the same slope and differ only in intercept, as shown in Figures 8H and 8I, respectively.

This linearization strategy may be applied to the complex families as well, and the complexity in the families is reflected in the resulting curves. Figure 9C shows the results of linearizing Figure 9A. The curves in Figure 9A are from the unequal-variance normal-distribution signal-detection model, and the curves come from different combinations of d' and σ. The linearizing function is the same as for the equal-variance normal-distribution, and indeed straight lines result. The slope of these lines reflect the standard-deviation parameter σ; hence, the family may be linearized, but not all members will have common slope. Figure 9D shows the results of attempting to linearizing Figure 9B. In fact, there is no one function that can simultaneously linearize all the lines. The curves in Figure 9B come from the dual-process model, and one can linearize with respect to one process or the other, but not both. The resulting partially-linearized curves in Figure 9D come from applying the linearizing function for the normal component (familiarity).

In summary, linearization is a useful approach to understanding the complexity of relations among ROC curves. If a family of curves can be linearized into a space where only intercept varies, then the variation across the experimental conditions may be explained with recourse to a single parametric change in a single process. Conversely, if curves cannot be linearized into a space where only intercept varies, then the variation across these conditions may more profitably be explained by multiple parameters and perhaps multiple processes. We refer to ROC families that can be linearized into families with common slope as *shift representable,* where the shift is instantiated as changes in intercepts. The reason for this choice in naming will become apparent in the next section.

The Signal-Detection Interpretation of Shift Representability

The concept of shift-representable ROC curves has a signal-detection equivalent. Figure 7 showed the concept of a universal representation—all models of ROC curves have a universal signal-detection representation. These distributions, however, are not required to be shifts of each other. Sometimes, however, the distributions in the universal representation may be transformed into a shift family of distributions. The third column of Figure 7, labeled “Shift Representation,” shows the shift representations when they exist. Examples of models that have a shift signal-detection representation are the equal-variance normal, the gamma-distribution model, and the single high-threshold model. These shifted signal-detection models are equivalent to the universal representations, and they predict the same family of ROC curves. Examples of models that do not have a shift signal-detection representation are the unequal variance normal signal-detection model and the dual-process model.

As might be surmised, shift-representability in ROC curves and shift representability in signal-detection models are equivalent in that one implies the other and vice-versa. If a family of ROC curves is shift representable, then there exists a signal detection model with shifted distributions that predicts the curves. Conversely, shifted distributions in a signal-detection model produce shift-representable ROC curves. There is also a relationship between the shift distribution family and the linearizing function. Let F be the cumulative distribution function of the noise distribution in a shift signal detection model. Then the linearizing function for the resulting ROC curves, denoted θ, is

$$\theta = -F^{-1}(1-p),$$

where p is the hit or false-alarm rate. The proof of this relationship between linearizing functions and signal-detection representability is provided in the supplement.

We recommend that shift-representability be a necessary condition for advocating a signal-detection theory. Most researchers use signal detection to model a single memory process or memory system, and signal detection is meant to capture simple relations across conditions. Unfortunately, signal detection without constraint does not capture the intended simplicity because without constraint the signal-detection theory is unfalsifiable. Shift representability adds the needed constraint, and does so without any parametric assumptions. Shift representability restores signal detection back to what it was intended—a model that captures a simplicity in processing—and does so without recourse to contentless parameter assumptions. This recommendation, while derived from principle, does not sentiment in the extant literature. For example, the UVSD model, which is a preferred signal detection model (e.g., Dube & Rotello, 2012; Wixted et al., 2007), is not shift representable, and should be avoided unless families of curves cannot be suitably linearized.

Assessing Shift Representability

Shift representability provides a needed tool for assessing the complexity of processing. If a family of ROC curves are shift representable, then a simple mechanism is appropriate to describe the data; conversely, if ROCs are not shift representable then more complex mechanisms are indicated. This inferential logic raises the important question of how a researcher may know if a family of ROC curves are shift representable. The crux of the matter is establishing that there does or does not exist a linearizing transform, and this is a difficult problem, the solution to which is necessarily technical. We provide a small summary of the approach here, and provide the details in the supplement. To see if shift representability holds, we linearize each ROC curve separately by applying a function that is composed of a series of straight line segments, called splines. These segments are joined at specific points called knots, and the values of the function at these knot points serve as free parameters. The approach we favor is to estimate the values at these knot points separately for each curve in a family. Shift representability holds only if these values do not vary across different curves in the family.

Conclusions

In this paper, we provide a critique of the standard story for interpreting ROC curves and call for a new focus on the relationships between ROC curves rather than a focus on the curvature or symmetry of isolated curves. We show here in contrast to the standard story that discrete-state models need not predict straight-line ROCs, and that one can always find a signal detection model to account for any family of ROC curves.

Discrete-state models, however, provide a natural constraint on families of ROC curves. In these models, stimulus factors affect the probability of entering one state or another, but not the probability of making a certain response conditional on being in a given state. This conditional independence leads directly to strong and testable constraints on families of ROC curves. The situation is not so sanguine for signal-detection models. The core of these models, that latent strength values are compared to a criterion, is unfalsifiable without recourse to additional assumptions. Traditionally, these assumptions have taken parametric form, say that latent strength is distributed as a normal, and the normal has been used so often that it feels natural. Nonetheless, these parametric assumptions, the normal included, have no psychological content. We provide an alternative constraint, shift representability, that is not parametric. Shift representability captures the intuition that in simple, single-process models, the stimulus variables should affect latent strength in one way, and in this case, there is some representation of latent strengths where the old-item distribution is a shift of the new-item distribution. We argue that shift-representability is the appropriate ROC property to assess if processing is mediated by one or more than one process.

We provide two different constraints on ROC structure, discrete-state representability and shift representability. One salient question for practitioners is whether both properties may be assessed, and, if so, how to interpret joint results. These properties are not exclusive, and ROC curves can meet both properties, one and not the other, or none. Our preference is to give priority to the discrete-state representability property for the following reasons: **1.** Discrete-state representability provides for much stronger constraint on ROC patterns than shift representability. In the discrete-state model, constraint may be observed not only on ROC curves, but the exact position each point must subtend in the ROC plot. In Figure 2, for example, not only is there a prediction for ROC curves (dashed lines) predicted, but there is a prediction for where on these lines the points lie from different payoff conditions. This exactness stands in contrast to the shift-representation predictions. Shift representation implies that given a linearized family, each curve must follow a specific form, but it does not constrain where points lie from different payoff conditions. Discrete-state representability is therefore a more constraining property, and certainly one that is easier to falsify if discrete-state models do not hold. Therefore, we are more impressed if the discrete-state constraint holds than if shift-representability holds. **2.** The notion of all-or-none loss is nowhere near as popular as graceful degradation. If discrete models do hold, then the theoretical consequences are larger as more development would need revision. It is with these two points in mind, that discrete-state models are easier to falsify and stand out of compliance with the graceful degradation theme in the field, that we promote the current set of studies that find support for discrete states including our work (Province & Rouder, 2012) and that from Kellen, Klauer, Broder and colleagues (Kellen, Klauer, & Broder, 2013; Broder, Kellen, Schutz, & Rohrmeier, 2013; Klauer & Kellen, 2010). We think these results should place strong constraints on mnemonic theory. If discrete-state constraints do not hold in a domain, we recommend researchers consider shift representability to adjudicate between a suitably constrained single-process mechanism and a less constrained multiple-process mechanism. In our view, ROC structure supports multiple processes when it is so complex that neither shift representability nor discrete-state representability hold.

With these conceptualizations of constraints, and with the developing statistical tools to assess them, we hope researchers will now be in position to systematically explore the relations among families of ROC curves to better understand the structure of processing.

References

Ashby, G. F. (1992). *Multidimensional models of perception and cognition.* Hillsdale, NJ: Erlbaum.

Attneave, F. (1959). *Applications of information theory to psychology : a summary of basic concepts, methods, and results.* New York: Holt, Rinehart and Winston.

Banks, W. P. (1970). Signal detection theory and human memory. *Psychological Bulletin*, *74*(2), 81 - 99.

Bays, P. M., & Husain, M. (2008). Dynamic Shifts of Limited Working Memory Resources in Human Vision. *Science*, *321*(5890), 851-854. Retrieved from http://www.sciencemag.org/cgi/content/abstract/321/5890/851

Blackwell, H. R. (1953). Psychological thresholds: Experimental studies of methods of measurement. *Bulletin of the Engineering Research Institute of the University of Michigan*, *No. 36*.

Broadbent, D. E. (1966). Two-state threshold model and rating-scale experiments. *Journal of the Acoustical Society of America*, *40*, 244-245.

Broder, A., Kellen, D., Schutz, J., & Rohrmeier, C. (2013). Validating a two-high threshold measurement model for confidence rating data in recognition. *Memory*, 1-29.

Broder, A., & Schutz, 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*(3), 587-606.

DeCarlo, L. T. (2003). An application of signal detection theory with finite mixture distributions to source discrimination. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 767–778.

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.

Egan, J. P. (1975). *Signal detection theory and ROC analysis.* New York: Academic Press.

Erdfelder, E., & Buchner, A. (1998). Process-dissociation measurement models: Threshold theory or detection theory. *Journal of Experimental Psychology: General*, *127*(1), 83–96.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Healy, A. F., & Kubovy, M. (1981). Probability matching and the formation of conservative decision rules in a numerical analog of signal detection. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 344-354.

Howard, M., Bessette-Symons, B., Zhang, Y., & Hoyer, W. (2006). Aging selectively impairs recollection in recognition memory for pictures: Evidence from modeling and ROC curves. *Psychology and Aging*, *21*, 96-106.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513-541.

Kellen, D., Klauer, K., & Broder, A. (2013). Recognition memory models and binary-response rocs: A comparison by minimum description length. *Psychonomic Bulletin & Review*, 1-27.

Kinchla, R. A. (1994). Comments on Batchelder and Riefer's multinomial model for source monitoring. *Psychological Review*, *101*, 166-171.

Klauer, K., & Kellen, D. (2010). Toward a complete decision model of item and source recognition: A discrete-state approach. *Psychonomic Bulletin & Review*, *17*(4), 465-478.

Luce, R. D. (1963). A threshold theory for simple detection experiments. *Psychological Review*, *70*, 61-79.

Luce, R. D. (1986). *Response times.* New York: Oxford University Press.

Ma, W. (2012). Organizing probabilistic models of perception. *Trends in Cognitive Sciences*, *16*(10), 511-518.

Ma, W., Beck, J., Latham, P., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*(11), 1432-1438.

Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, N.J.: Lawrence Erlbaum Associates.

Malmberg, K. J. (2002). On the form of rocs constructed from confidence ratings. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 380-387.

Mandler, G. (1980). Recognizing: The judgment of previous occurence. *Psychological Review*, *87*, 252-271.

Massaro, D. W. (1998). Models for reading letters and words. In D. Scarborough & S. Sternberg (Eds.), *Methods, models, and conceptual issues: An invitation to cognitive science, vol. 4* (p. 301-364). Cambridge, MA: MIT Press.

Massaro, D. W., & Oden, G. C. (1979). Integration of featural information in speech perception. *Psychological Review*, *85*, 172-191.

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.

Murdock, B. B. (1993). Todam2: A model for the storage and retrieval of item, association, and serial order information. *Psychological Review*, *100*, 183-203.

Myers, J. L. (1976). Probability learning and sequence learning. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Approaches to human learning and motivation* (p. 171-205). Hillsdale, NJ: Erlbaum.

Province, J. M., & Rouder, J. N. (2012). Evidence for discrete-state processing in recognition memory. *Proceedings of the National Academy of Sciences*.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59-108.

Rotello, C., McMillan, N. A., & Reeder, J. A. (2004). Sum-difference theory of remembering and knowing: A two-dimensional signal-detection model. *Psychological Review*, *111*, 588-616.

Rouder, J. N. (2004). Modeling the effects of choice-set size on the processing of letters and words. *Psychological Review*, *111*, 80-93.

Rouder, J. N., Tuerlinckx, F., Speckman, P. L., Lu, J., & Gomez, P. (2008). A hierarchical approach for fitting curves to response time measurements. *Psychonomic Bulletin & Review*, *15*(1201-1208).

Schacter, D., & Tulving, E. (1994). What are the memory systems of 1994? In D. Schacter & E. Tulving (Eds.), *Memory systems 1994* (p. 1-38). Cambridge, MA: MIT Press.

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, *15*, 233-250.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM–retrieving effectively from memory. *Psychonomic Bulletin and Review*, *4*, 145-166.

Sims, C., Jacobs, R., & Knill, D. (2012). An ideal observer analysis of visual working memory. *Psychological Review*, *119*(4), 807-830.

Slotnick, S. D., & Dodson, C. S. (2005). Support for a continuous (single-process) model of recognition memory and source memory. *Memory & Cognition*, *33*, 1499-1517.

Swets, J. A. (1961). Is there a sensory threshold? *Science*, *134*, 168-177.

Townsend, J. T. (1971). Theoretical analysis of an alphabetic confusion matrix. *Perception & Psychophysics*, *9*, 40-50.

Townsend, J. T., & Landon, D. E. (1982). An experimental and theoretical investigation of the constant-ratio rule and other models of visual letter confusion. *Journal of Mathematical Psychology*, *25*, 119-162.

Vulcan, N. (2000). An economists perspective on probability matching. *Journal of Economic Surveys*, *14*, 101-118.

Wixted, J., & Mickes, L. (2010). A continuous dual-process model of remember/know judgments. *Psychological Review*, *117*, 1025–1054.

Wixted, J. T., & Mickes, L. (2010). Useful scientific theories are useful: A reply to Rouder, Pratte, and Morey (2010). *Psychonomic Bulletin and Review*, *17*, 436-442.

Yonelinas, A. P. (1994). Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1341-1354.

Yonelinas, A. P., & Parks, C. M. (2007). Receiver operating characteristics (ROCs) in recognition memory: A review. *Psychological Bulletin*, *133*, 800-832.

Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, *20*, 423-428.

Footnotes

[1]The gamma distribution may be parameterized with a scale and shape parameters. We fix the shape of the gamma to 2.0, that is the gamma distributions in this report describe the sum of two exponential distributions. The scale is set to 1.0 for new items and is increased for old items. The degree of this increase serves as a measure of strength.

Figure Captions

Figure 1. The standard story: **A**. The high threshold model expressed as a processing tree. **B.** ROC curves from the high-threshold model are straight lines. **C-D.** The unequal-variance signal detection model and the corresponding ROC predictions, respectively. **E.** Yonelinas' (1994) two-process mixture model expressed as a processing tree. The parameters $\tilde{h}_i$ and $\tilde{f}_i$ are the hit and false alarm rates from an equal-variance signal detection model; $\tilde{h}_i = \Phi(d' - c)$ and $\tilde{f}_i = \Phi(-c)$. **F.** The corresponding ROC predictions.

Figure 2. A discrete-state model without the certainty assumption and the corresponding ROC curves. **A.** The model. **B.** The curve when $d_s = d_n = .7$, $g_0 = 0$, $a_0 = 3$ and $b_0 = -1$ (see Eqs. 9-11); this curve is similar to empirical curves in that it is not straight and has an appropriate degree of asymmetry. **C.-D.** Two examples of the constraint from discrete-state models. Each ROC curve formed by varying the detection parameter must be a linear combination of the curves when $d = 0$ and when $d = 1$. See the text for details.

Figure 3. Conditional independence constraint as seen in distributions of confidence ratings. The words "FROG" and "TABLE" serve as the old and new item for the 2AFC trial. **A.** Distribution of confidence ratings from a guess state. **B.** Distribution of confidence ratings from a detect state. **C.** Distribution of responses from the zero-repetition condition which directly reflects the guess state. **D.** Distribution of responses from the two-repetition condition which is an equal mixture of the guess and detect states. **E.** Responses from the four-repetition condition, which reflects primarily the detect state with some guessing. **F.** Predictions from a discrete-state model for all three conditions. Note the confidence-ratings distribution from the two-repetition condition is projected downward. **G.** Contrasting predictions from a latent-strength model for all three conditions. The confidence-ratings distribution from two-repetition condition

is again projected downward.

Figure 4. Confidence rating distributions from two selected participants from Province and Rouder (2012) exhibit discrete-state representability. The distribution of confidence ratings for the 1-repetition and 2-repetition condition are mixtures of those from the 0-repetition and 4-repetition condition. This type of pattern held broadly in Province and Rouder, and is equivalent in constraint to the discrete-state representability ROC pattern of Figure 2C.

Figure 5. Signal detection models with alternative parametric assumptions. **A.** A gamma-distribution signal-detection model. **B.** An ex-Gaussian distribution signal-detection model. **C.** ROC curves from different models are often similar though not identical.

Figure 6. The unfalsifiability of the core of signal detection. **A-B.** A conventional signal detection model and the corresponding ROC curve. The slope of the ROC curve at the point in Panel B corresponds to the ratio of densities at the corresponding strength value in Panel A. **C.** An equivalent signal detection model which makes the same ROC predictions. The ratio of densities corresponds to the slope of ROC points. Every ROC curve corresponds to some signal detection model representation.

Figure 7. Equivalent signal-detection representations. Each row shows equivalent signal detection representation for popular models. The left column is the common signal detection representation. The center column is the universal representation where the new-item distribution is a uniform. The last column is a shift signal-detection representation. The rows show the case for the equal-variance normal signal-detection model, the unequal-variance normal signal-detection model, the gamma signal-detection model, the high-threshold model, and Yonelinas' dual-process model.

Figure 8. Shift-representable families of ROC curves. **A-C**. Families of ROC curves from the equal-variance normal signal detection model, the gamma signal detection model, and the high-threshold model, respectively. **D-F.** Corresponding linearizing functions. **G-I.** Resulting transformed ROC curves are linear with common slope.

Figure 9. Families of ROC curves that are not shift representable **A-B**. Families of ROC curves from the unequal-variance normal signal detection model and Yonelinas' dual-process model, respectively. **C-D.** Results from applying a zROC transform. The zROC transform can linearize the unequal-variance normal ROCs, but the slopes vary across the family. The zROC transfer can linearize the familiarity component in Yonelinas dual-process model but not the mixture of recollection and familiarity.

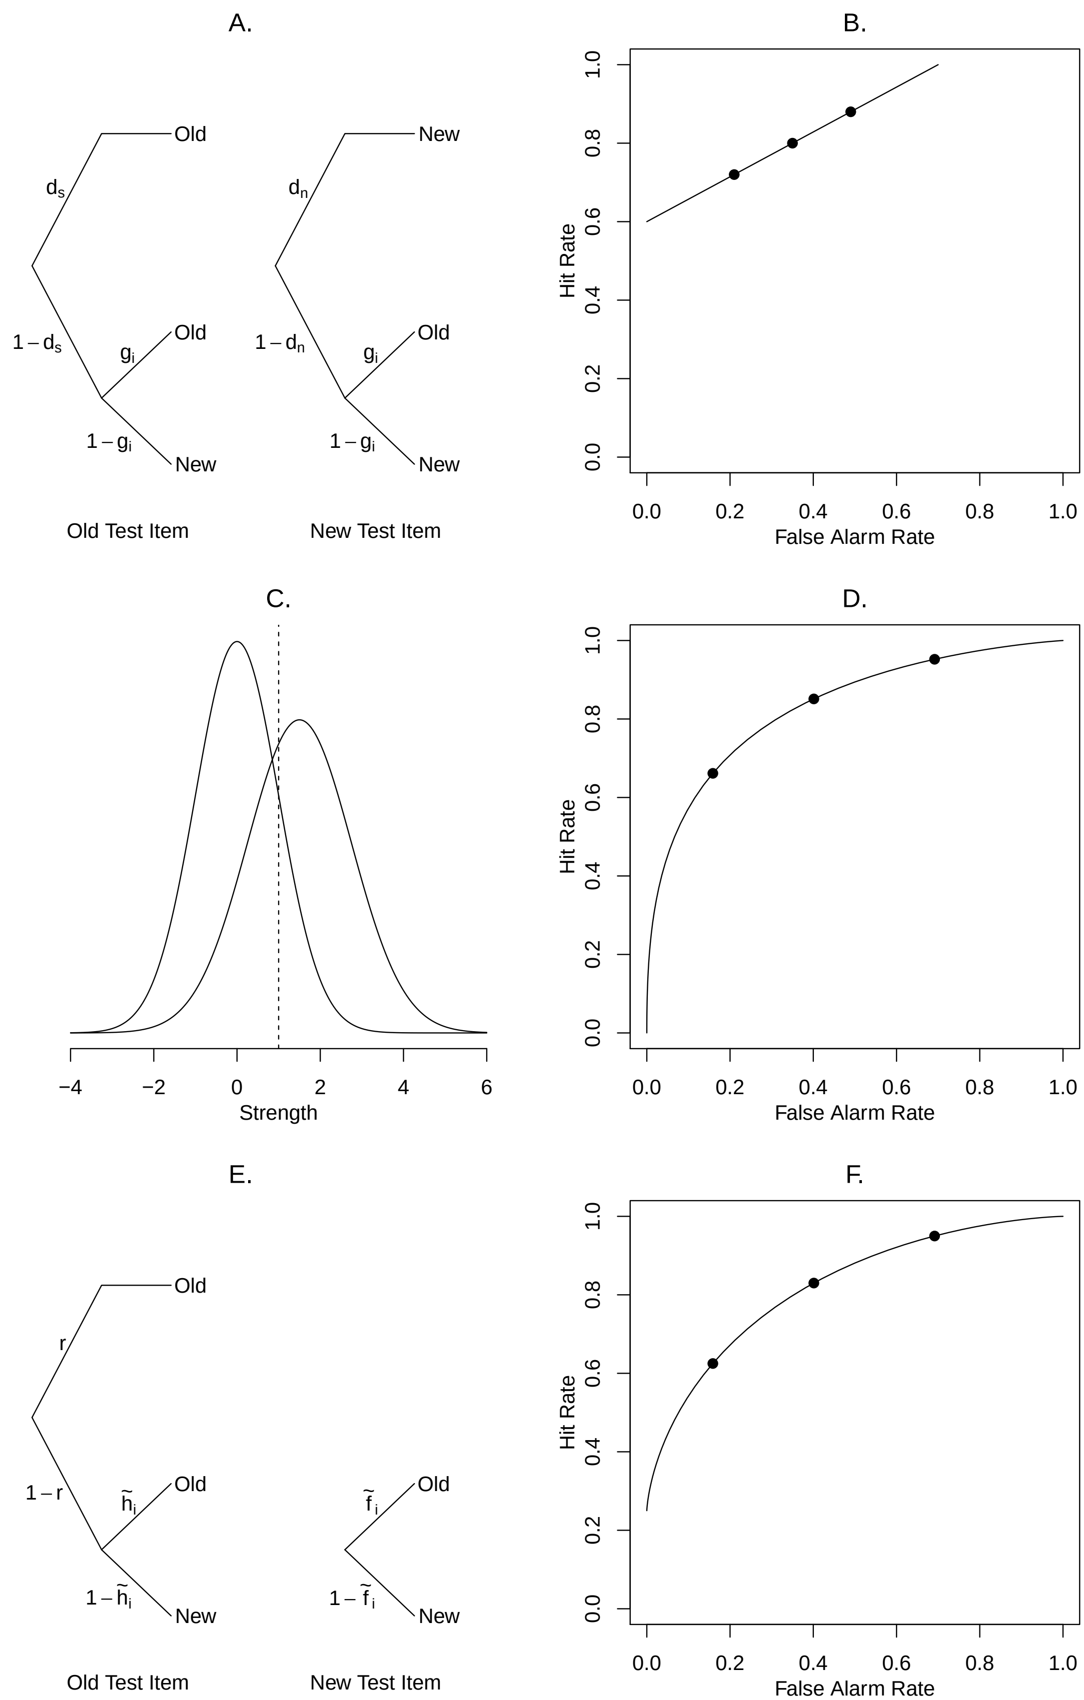
A.
Old
d_s
$1-d_s$
g_i
Old
$1-g_i$
New
Old Test Item
New
d_n
$1-d_n$
g_i
Old
$1-g_i$
New
New Test Item
B.
Hit Rate
False Alarm Rate
C.
Strength
D.
Hit Rate
False Alarm Rate
E.
Old
r
$1-r$
$\tilde{h}_i$
Old
$1-\tilde{h}_i$
New
Old Test Item
$\tilde{f}_i$
Old
$1-\tilde{f}_i$
New
New Test Item
F.
Hit Rate
False Alarm Rate

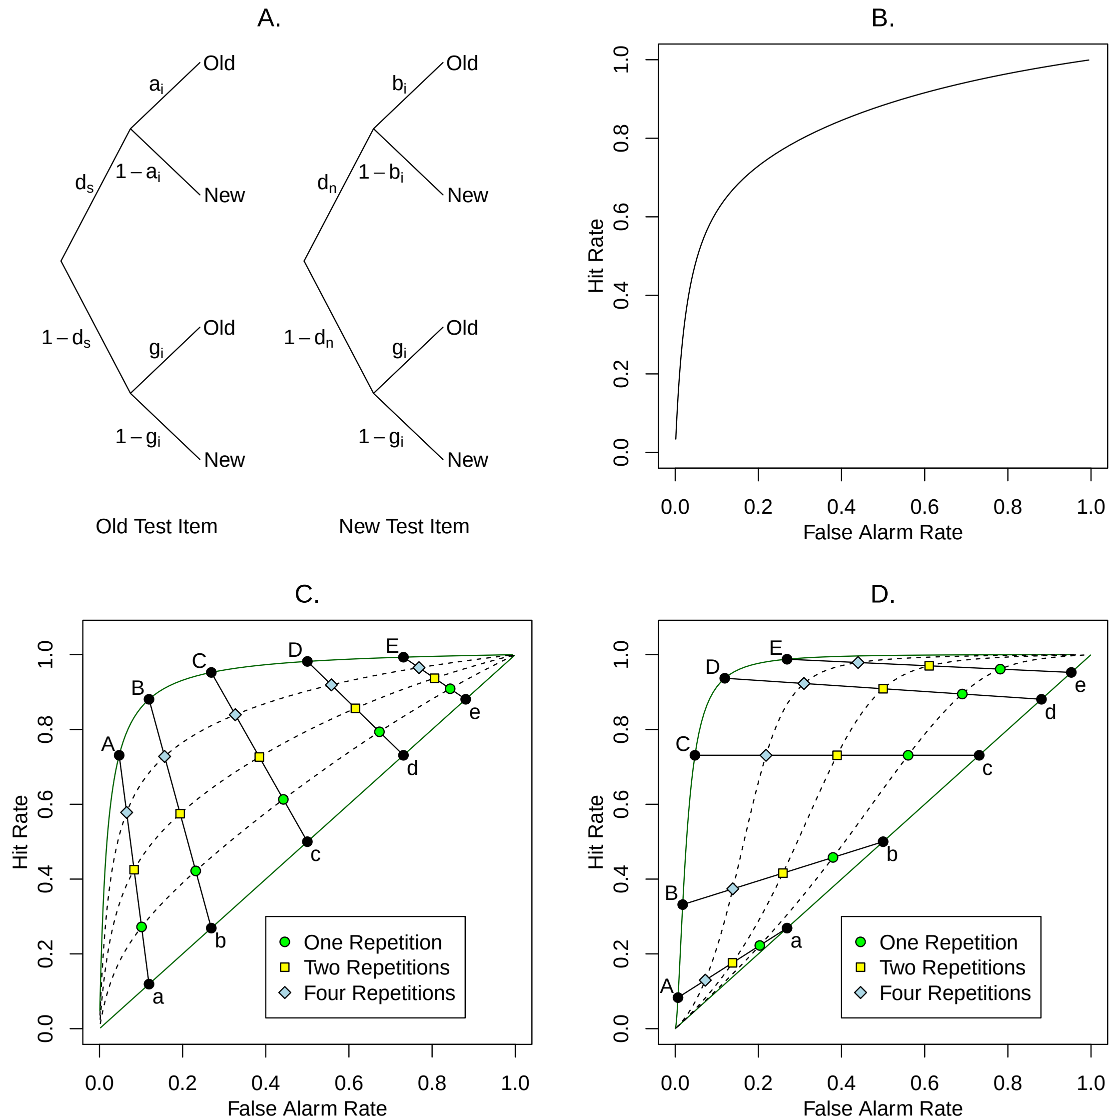
A.
Old
New
Old
New
Old Test Item
Old
New
Old
New
New Test Item
B.
Hit Rate
False Alarm Rate
C.
One Repetition
Two Repetitions
Four Repetitions
Hit Rate
False Alarm Rate
D.
One Repetition
Two Repetitions
Four Repetitions
Hit Rate
False Alarm Rate

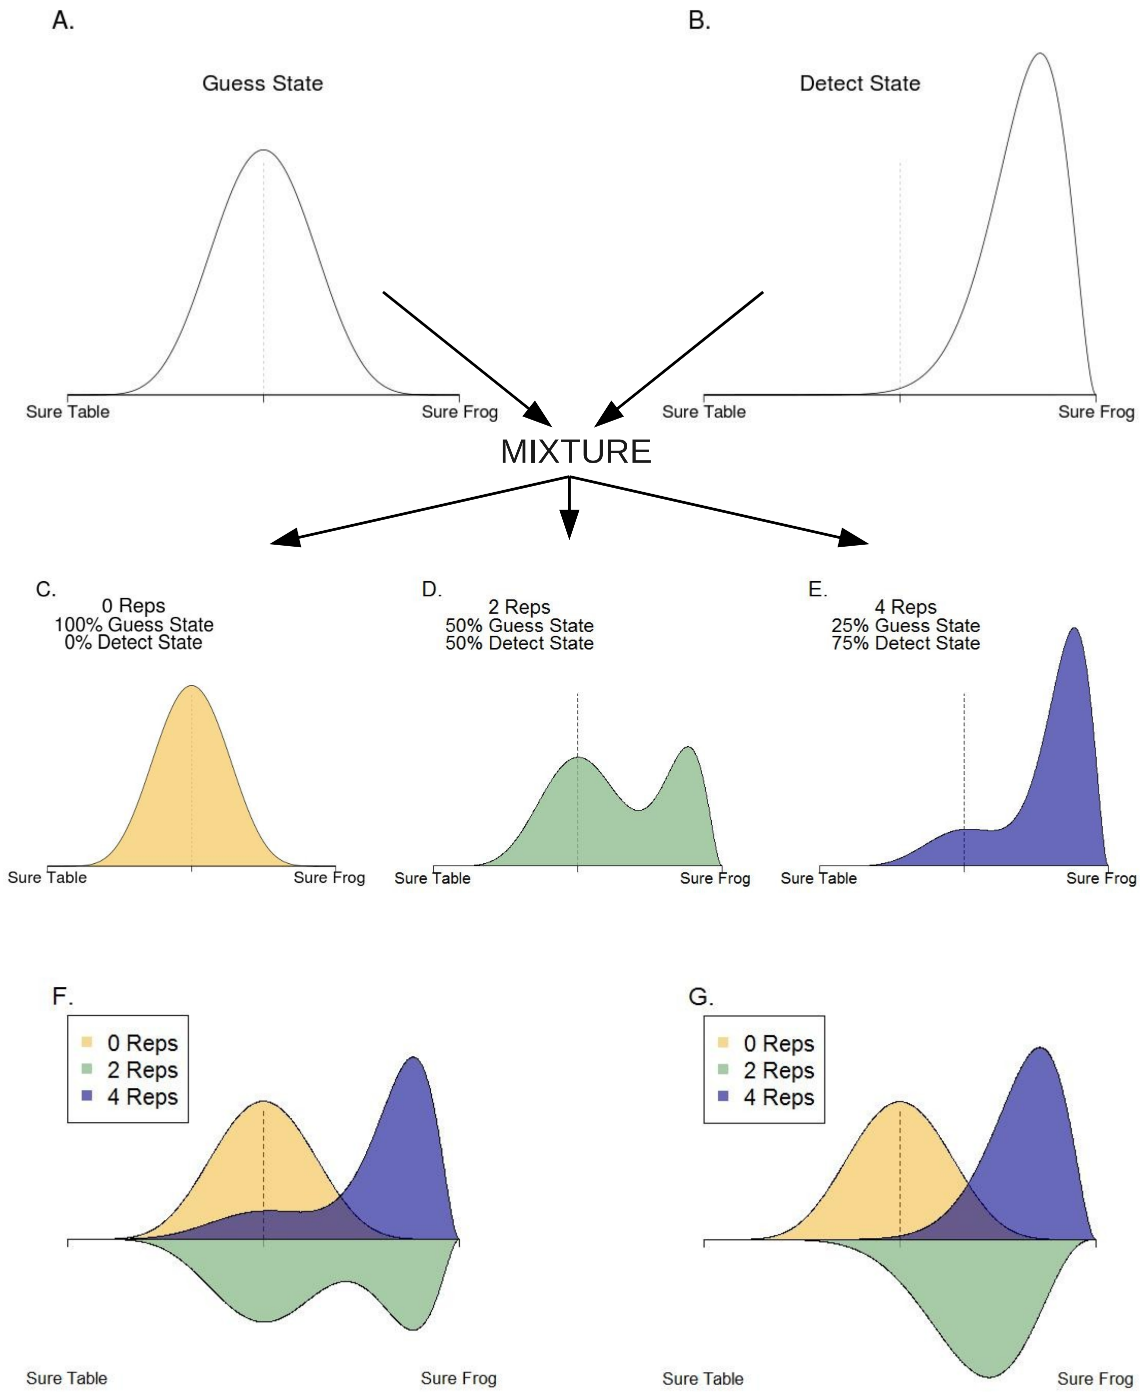
A.
Guess State
Sure Table
Sure Frog
B.
Detect State
Sure Table
Sure Frog
MIXTURE
C.
0 Reps
100% Guess State
0% Detect State
Sure Table
Sure Frog
D.
2 Reps
50% Guess State
50% Detect State
Sure Table
Sure Frog
E.
4 Reps
25% Guess State
75% Detect State
Sure Table
Sure Frog
F.
0 Reps
2 Reps
4 Reps
Sure Table
Sure Frog
G.
0 Reps
2 Reps
4 Reps
Sure Table
Sure Frog

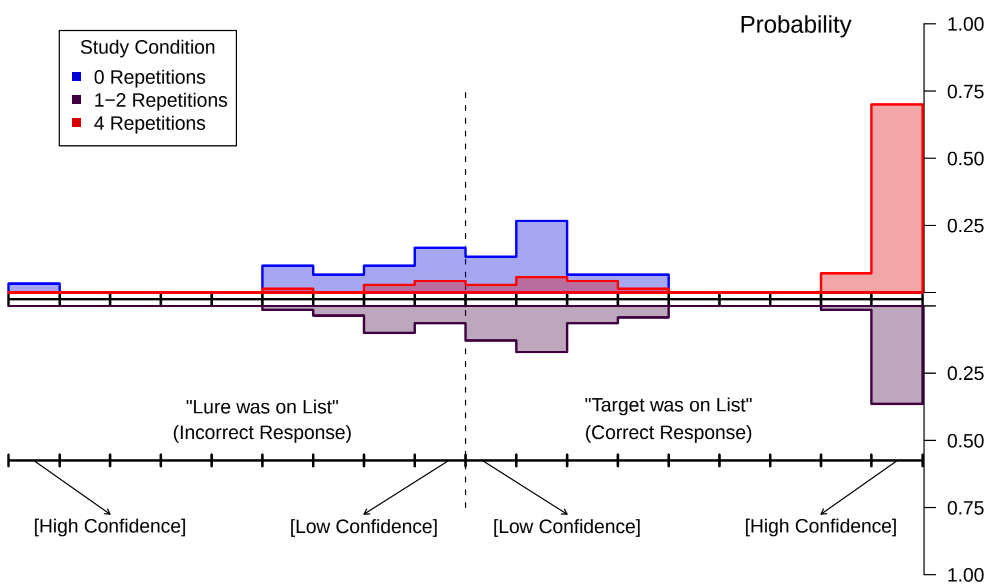
Probability
1.00
0.75
0.50
0.25
0.25
0.50
0.75
1.00
Study Condition
0 Repetitions
1−2 Repetitions
4 Repetitions
"Lure was on List"
(Incorrect Response)
"Target was on List"
(Correct Response)
[High Confidence]
[Low Confidence]
[Low Confidence]
[High Confidence]

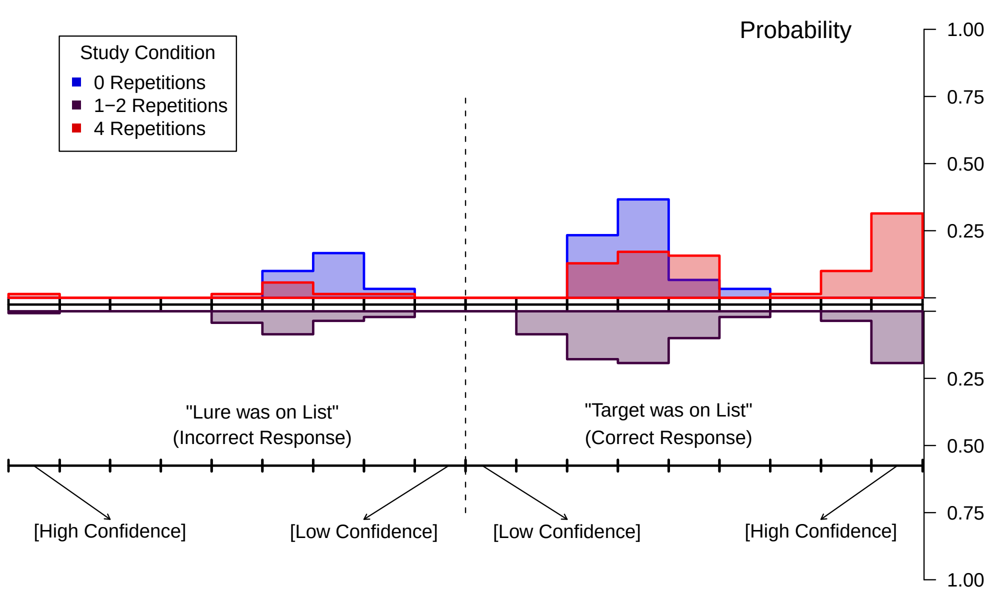
Probability
1.00
0.75
0.50
0.25
0.25
0.50
0.75
1.00
Study Condition
0 Repetitions
1−2 Repetitions
4 Repetitions
"Lure was on List"
(Incorrect Response)
"Target was on List"
(Correct Response)
[High Confidence]
[Low Confidence]
[Low Confidence]
[High Confidence]

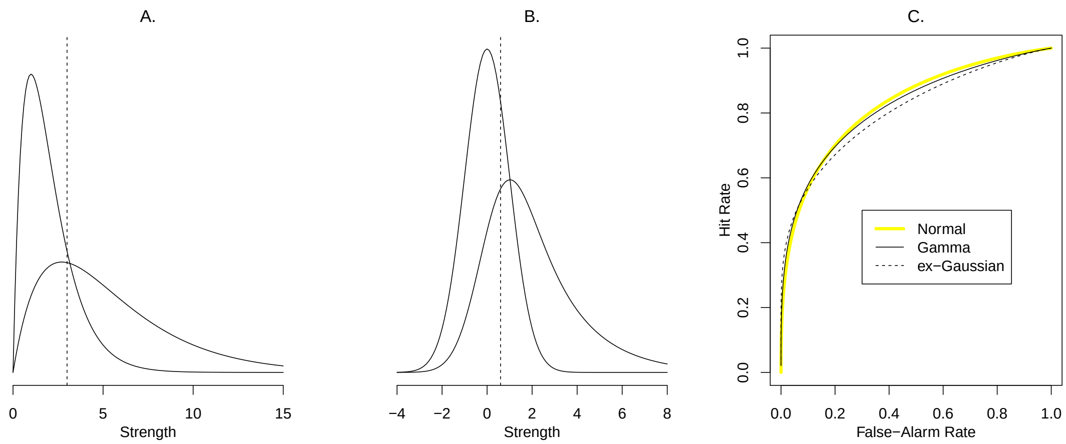
A.
B.
C.
Strength
Strength
0
5
10
15
−4
−2
0
2
4
6
8
Hit Rate
False−Alarm Rate
0.0
0.2
0.4
0.6
0.8
1.0
Normal
Gamma
ex−Gaussian

From ROC Curves To Psychological Theory, Figure 6

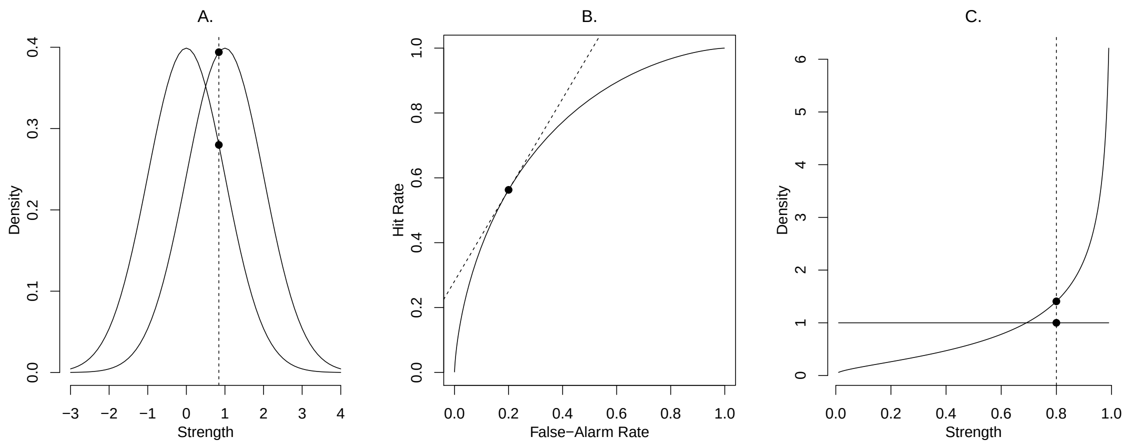

From ROC Curves To Psychological Theory, Figure 7

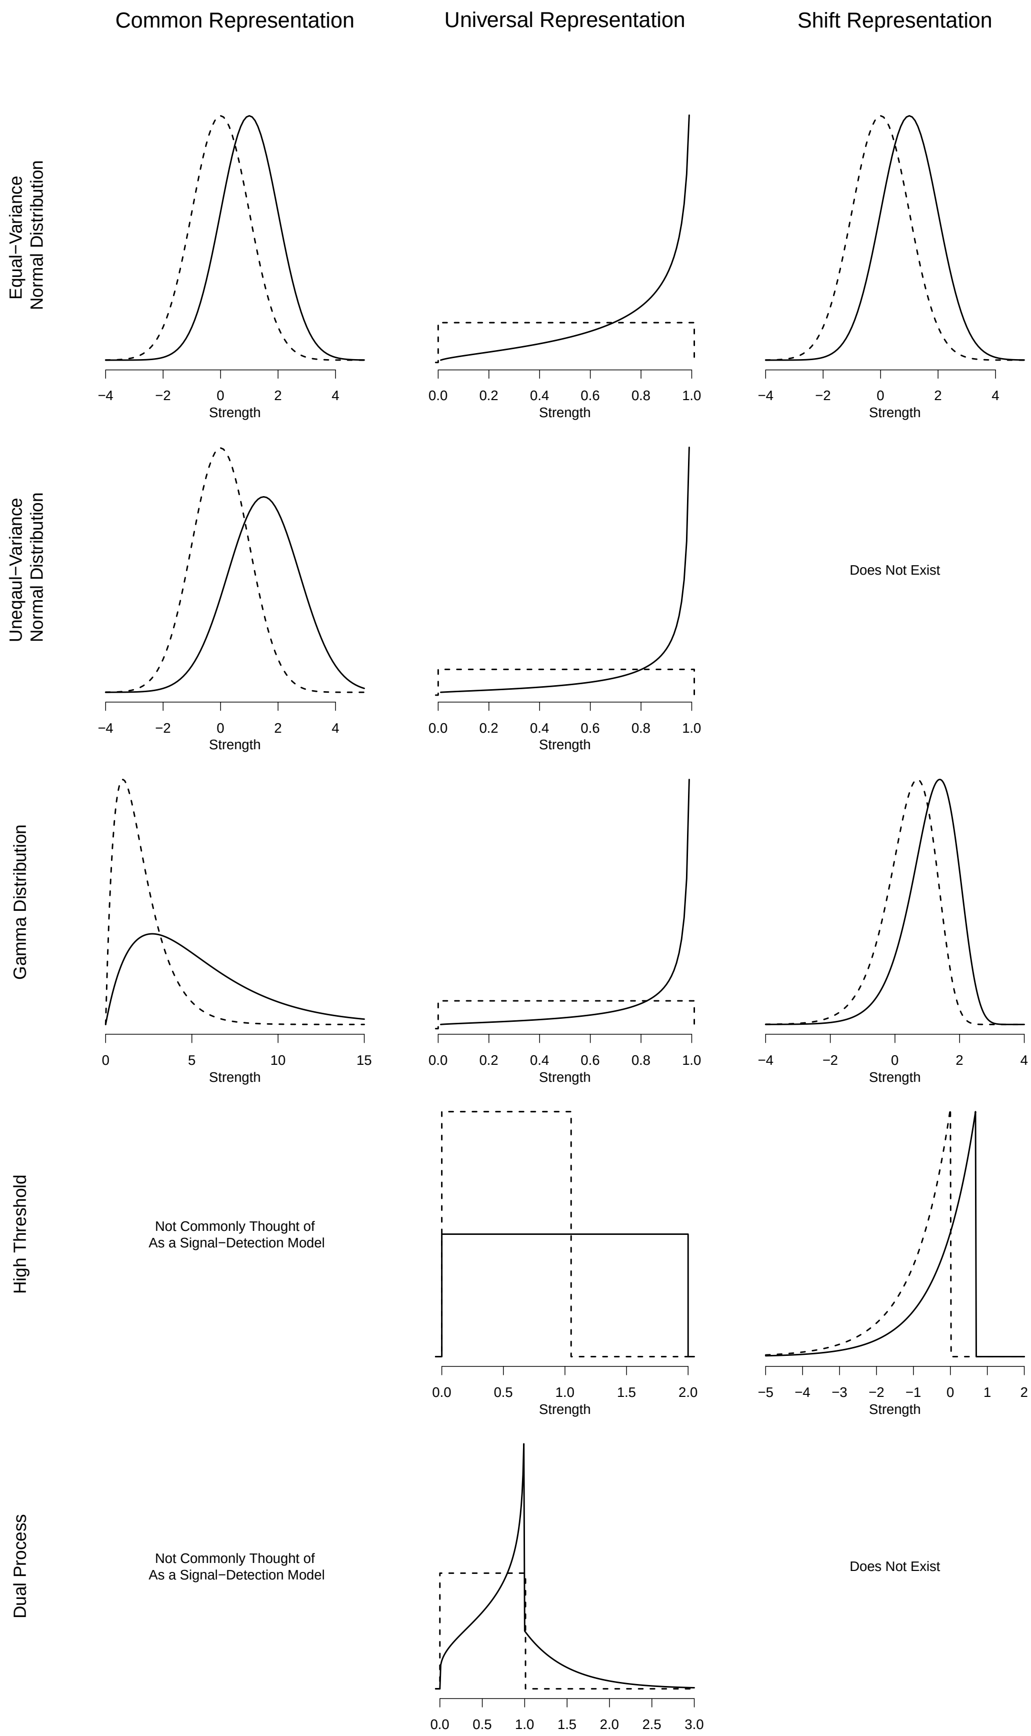

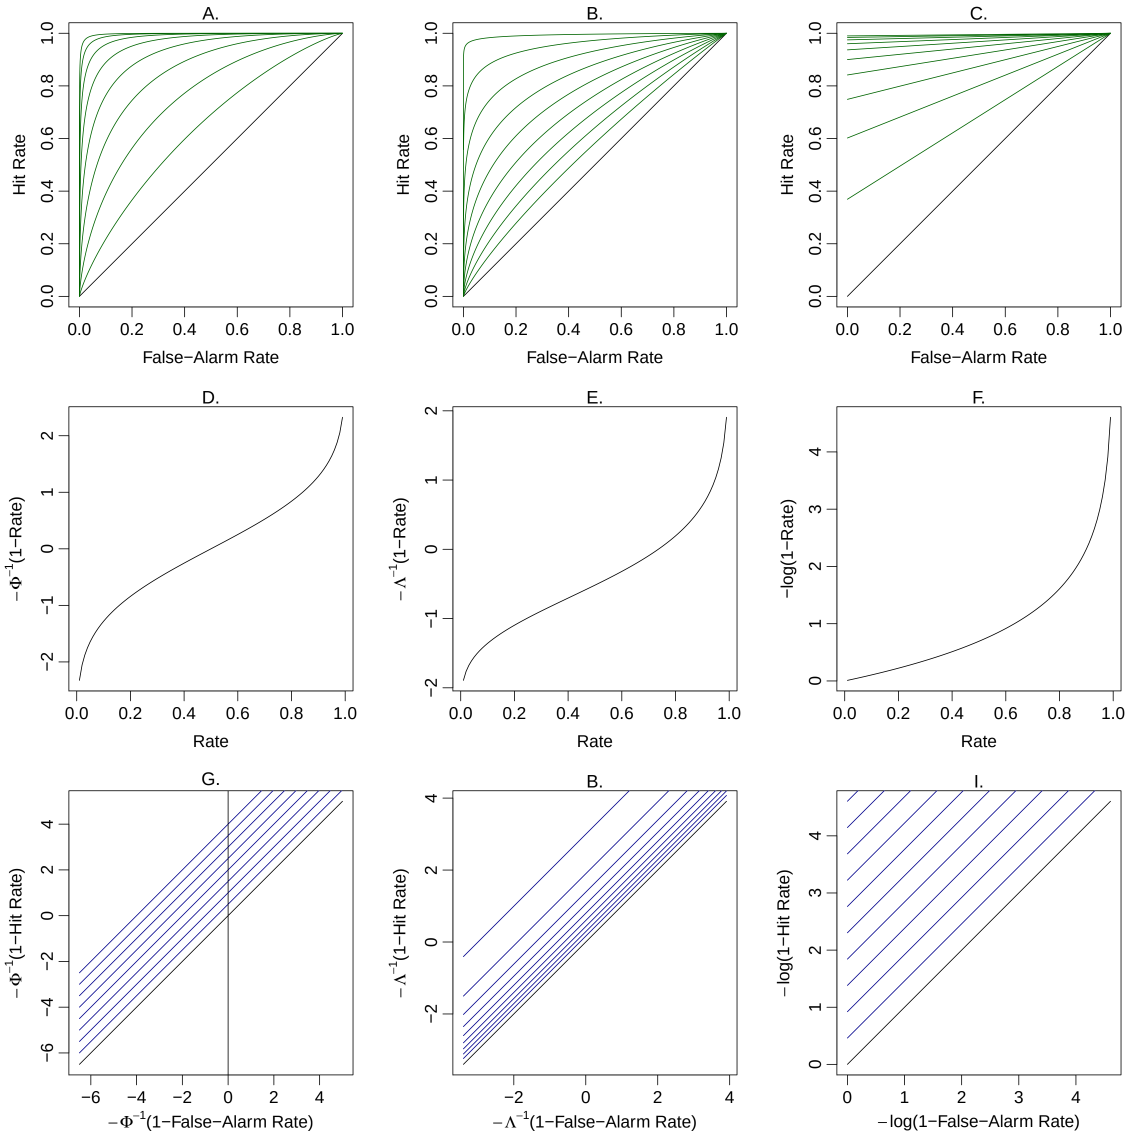
A.
B.
C.
Hit Rate
False−Alarm Rate
D.
E.
F.
−Φ^{−1}(1−Rate)
−Λ^{−1}(1−Rate)
−log(1−Rate)
Rate
G.
B.
I.
−Φ^{−1}(1−Hit Rate)
−Λ^{−1}(1−Hit Rate)
−log(1−Hit Rate)
−Φ^{−1}(1−False−Alarm Rate)
−Λ^{−1}(1−False−Alarm Rate)
−log(1−False−Alarm Rate)

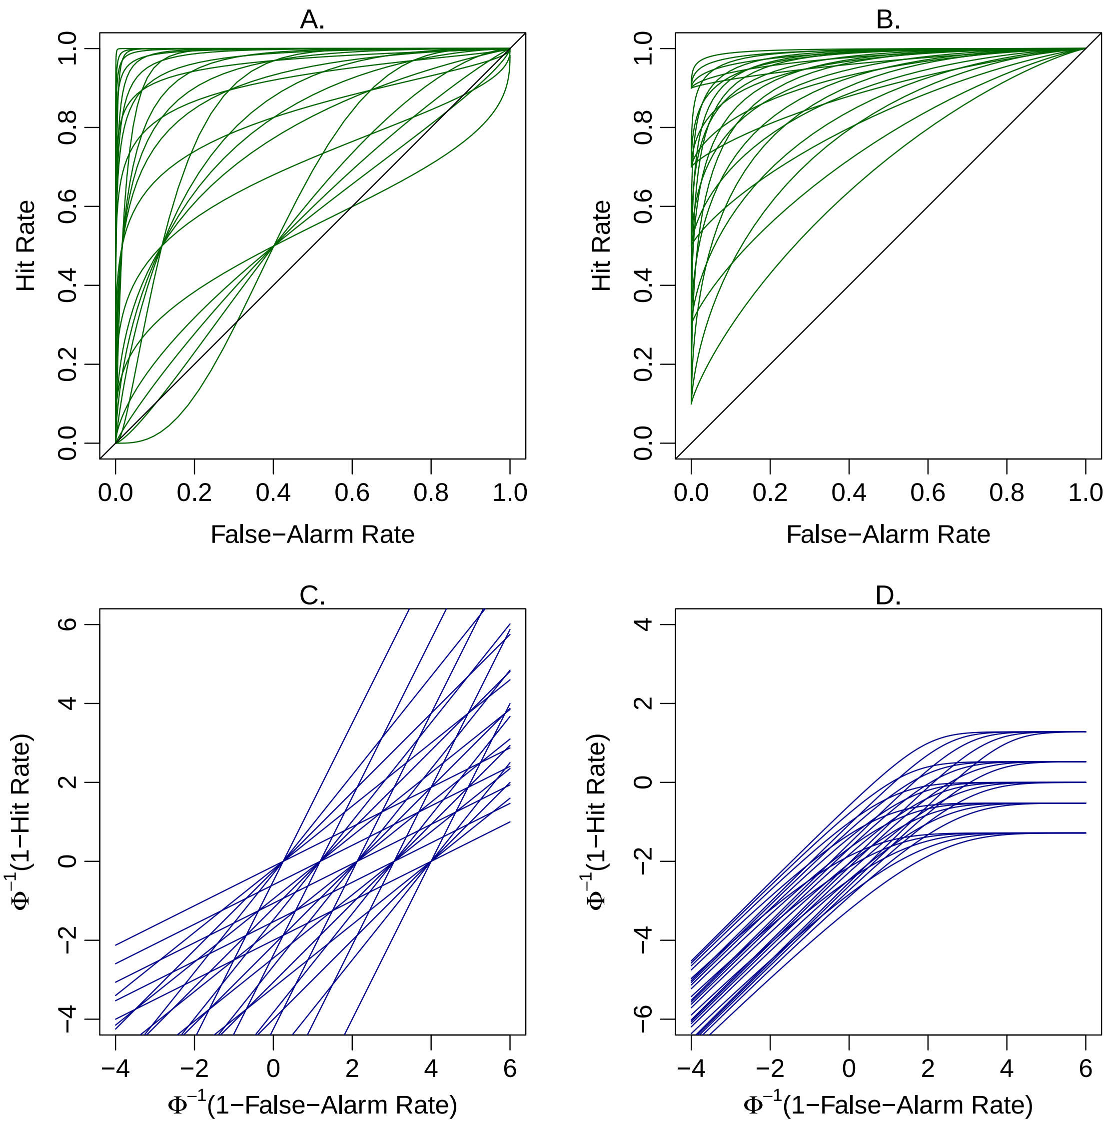
A.
B.
Hit Rate
False−Alarm Rate
C.
D.
Φ^{-1}(1−Hit Rate)
Φ^{-1}(1−False−Alarm Rate)