

Symmetry-based methodology for decision-rule identification in *same–different* experiments

ALEXANDER A. PETROV
Ohio State University, Columbus, Ohio

The standard practice of reducing every *same–different* data set to two numbers (hits and false alarms) is wasteful, because the response pattern to all four stimulus pairs carries information about the decision rule adopted by the observer. We describe eight rules organized in three families: differencing, covert classification, and likelihood ratio. We prove that each family produces a characteristic pattern of (in)equalities among the response probabilities. We propose two simple qualitative tests. Is the performance on stimulus pairs AA and BB statistically indistinguishable? If not, differencing and likelihood-ratio strategies can be rejected. Is the performance on pairs AB and BA indistinguishable? If not, covert classification can be rejected. We present algorithms for fitting two covert-classification models and illustrate the new methodology in a perceptual learning experiment on visual motion-direction discrimination. The standard assumption of symmetric decision criteria was violated.

Judging whether two things are the same or different is a fundamental cognitive operation. Such judgments are a staple of daily life and a building block for higher-level processes. The *same–different* task is an important tool in the psychological laboratory, where it is used with animals, infants (habituation), and patients. It is also widely used in perceptual discrimination studies with adult observers, particularly when the relevant stimulus features are difficult to describe (e.g., wine tasting). Thus, understanding the mechanisms of *same–different* judgments is an important research goal for both theoretical and methodological reasons.

Signal detection theory (SDT; Green & Swets, 1966; Macmillan & Creelman, 2005) is a prominent and fruitful framework for perceptual discrimination. Its appeal is that, given certain assumptions, SDT can separate the perceptual and nonperceptual determinants of the observed performance. The basic experimental design for discrimination involves only two stimuli, referred to as A and B here. Because of perceptual noise, repeated presentations of each stimulus give rise to a whole distribution of internal magnitudes. Because the two stimuli are similar, their respective distributions overlap and have the same standard deviation σ . The normalized distance, $d' = (\mu_B - \mu_A)/\sigma$, between the means of the two magnitude distributions is a measure of the perceptual discriminability of the two stimuli. It can be estimated from the observed proportions of hits (correct classifications of B as “B”) and false alarms (incorrect classifications of A as “B”). Importantly, d' is a sensitivity statistic that characterizes the stimulus pair but not the response bias. The latter is a nonperceptual factor under strategic control that affects the observed proportions but is calculated out of the d' estimate (Green & Swets, 1966).

The SDT framework has been extended to the *same–different* task (e.g., Macmillan & Creelman, 2005; Noreen, 1981; Sorkin, 1962). This task is more complex, because each trial involves two stimulus presentations, usually in succession. The decision space is thus two-dimensional, with x - and y -axes representing the magnitudes of the first and second stimuli, respectively. The fixed *same–different* design involves four stimulus pairs: AA, AB, BA, and BB. In a noise-free system, they would be represented by four points arranged in a square configuration (cf. Figures 1A–1I). The system is noisy, however, and repeated presentations give rise to four bivariate distributions. Those centered on $\langle\mu_A, \mu_A\rangle$ and $\langle\mu_B, \mu_B\rangle$ are labeled “same,” whereas those centered on $\langle\mu_A, \mu_B\rangle$ and $\langle\mu_B, \mu_A\rangle$ are labeled “different.” This is equivalent to a difficult classification problem known as the XOR (exclusive OR) problem. The *same* and *different* regions cannot be separated by a linear boundary (Ashby & Townsend, 1986; Minsky & Papert, 1969; Rumelhart, Hinton, & Williams, 1986). A complex decision rule is necessary—one that involves nonlinear processing and/or multiple criteria.

In this article, we examine eight decision rules (or *strategies*) for the *same–different* task. They belong to three families: differencing, covert classification, and likelihood ratio. Observers who follow a differencing strategy respond on the basis of an explicit representation of the difference between the two stimuli presented on a trial (Sorkin, 1962). Observers who follow a covert-classification strategy label each component stimulus as either “A” or “B” (or, possibly, “don’t know”) and then combine the two labels (Pollack & Pisoni, 1971). Finally, the statistically optimal strategy compares the likelihood that the two magnitudes are drawn from a *different* distribution with

A. A. Petrov, apetrov@alexpetrov.com

the likelihood that they are drawn from a *same* distribution (Dai, Versfeld, & Green, 1996; Noreen, 1981). It is doubtful that human observers have the requisite knowledge and processing power to implement the likelihood-ratio strategy in the general case, but it is an important theoretical benchmark for the other strategies.

The researcher needs to know what strategy is adopted in a given experiment. First of all, without such knowledge one cannot estimate d' . The current practice is to gather the proportions of hits and false alarms according to Equation 1 and then calculate d' according to various strategy-specific formulas (Macmillan & Creelman, 2005). The problem is that different decision rules imply different d' for the same data; the discrepancy can be almost twofold in extreme cases (Macmillan & Creelman, 2005). Second, the response bias is underdetermined too. All indices of *same-different* bias proposed in the literature are strategy specific (Hautus & Collins, 2003; Irwin, Hautus, & Francis, 2001). Third, the decision rule can be of considerable interest in itself. For example, Petrov (2009) investigated whether perceptual learning involves gradual transitions from suboptimal to nearly optimal strategies.

$$\begin{aligned} H &= P(\text{"different"} | AB \text{ or } BA) \\ F &= P(\text{"different"} | AA \text{ or } BB). \end{aligned} \quad (1)$$

There is a pervasive tendency in the methodological literature to analyze the *same-different* data exclusively in terms of hits and false alarms (e.g., Dai et al., 1996; Irwin, Hautus, & Butcher, 1999; Macmillan & Creelman, 2005; Noreen, 1981). This anchors the *same-different* analysis in the better understood *yes-no* analysis (Green & Swets, 1966). However, this approach ignores valuable information available in the *same-different* data. Because Equation 1 reduces everything to two numbers (H and F) and the SDT model has two parameters (d' and bias), no degrees of freedom are left to provide constraints to infer the strategy. Additional information is needed and can be sought in various ways. One method is to obtain *same-different* and *yes-no* responses with the same materials and compare the resulting d' values (see Macmillan & Creelman, 2005, p. 228, for a review). Another method looks for strategy-specific features in the shapes of the isosensitivity and/or isobias curves (Francis & Irwin, 1995; Hautus & Collins, 2003; Irwin & Francis, 1995). A third method adds a random and independent perturbation to each stimulus and examines the resulting pattern of correlations (Berg, 1989; Dai et al., 1996; Wang, Irwin, & Hautus, 2005). Ashby and Gott's (1988) randomization technique adds large amounts of known external noise to the stimuli so that the unknown internal noise becomes relatively insignificant. The decision boundary can then be mapped out directly (Wang et al., 2005). These methods are valuable tools for studying the decision-making aspects of the task. Their complicated experimental designs, however, limit their applicability when the perceptual aspects are of main interest.

In this article, we develop a methodology that does not require special experimental designs. The main idea is to take full advantage of the fact that the *same-different* task

yields four pieces of information rather than just two. We can estimate four probabilities, because there are four distinct stimulus pairs:

$$\begin{aligned} p_{AA} &= P(\text{"different"} | AA) \\ p_{AB} &= P(\text{"different"} | AB) \\ p_{BA} &= P(\text{"different"} | BA) \\ p_{BB} &= P(\text{"different"} | BB). \end{aligned} \quad (2)$$

The traditional analysis assumes that $p_{AB} = p_{BA}$ and $p_{AA} = p_{BB}$. We will demonstrate that these identities can fail empirically and that their violations are very informative theoretically. Different strategies predict different patterns of equalities and inequalities among the four probabilities in Equation 2. All strategies whose predictions are violated in a given data set can be ruled out from further consideration (Platt, 1964; Popper, 1963). In particular, we will prove that all differencing strategies can be ruled out when $p_{AA} \neq p_{BB}$, and most covert-classification strategies can be ruled out when $p_{AB} \neq p_{BA}$. The proof is based on symmetry considerations and allows parameter-free inferences from qualitative patterns in the data. No goodness-of-fit comparisons are required.

The next section introduces eight decision rules that cover the space of plausible theoretical alternatives. Their symmetries and the observable consequences thereof are presented next. This leads to a systematic methodology for analyzing *same-different* data. It identifies four qualitative patterns of (in)equalities among the observed proportions and lists the strategies compatible with each. The methodology is illustrated on data from a perceptual-learning experiment. Mathematical proofs and d' formulas for covert classification are given in the Appendices.

Eight Decision Rules

Representational assumptions. We adopt the standard representational assumptions for the *same-different* task (Dai et al., 1996; Green & Swets, 1966; Laming, 1986; Macmillan & Creelman, 2005; Noreen, 1981): (1) The perceptual effect (or *magnitude*) of the relevant stimulus feature is a unidimensional Gaussian random variable; (2) the two stimuli being discriminated are very similar and thus give rise to magnitude distributions with equal variance; (3) the two successive presentations in a trial are statistically independent; (4) a stimulus has the same perceptual effect, regardless of whether it is presented first or second; and (5) the memorial variability is negligible relative to perceptual variability.

The decision space is two dimensional, with horizontal and vertical axes corresponding to the first and second presentation interval, respectively. The four stimulus pairs give rise to four bivariate Gaussian distributions with centers located on the corners of a square, as is illustrated in Figure 1A. Under these assumptions, the covariance matrices of all four distributions have the same σ^2 along the diagonal and 0 elsewhere. Without loss of generality, we can choose a coordinate system with origin at the center of the square and with unit equal to $\sigma = 1$. The mean perceptual effects of stimuli A and B are thus $\mu_A = -d'/2$ and $\mu_B = +d'/2$. Our goal is to estimate d' —the size of the square.

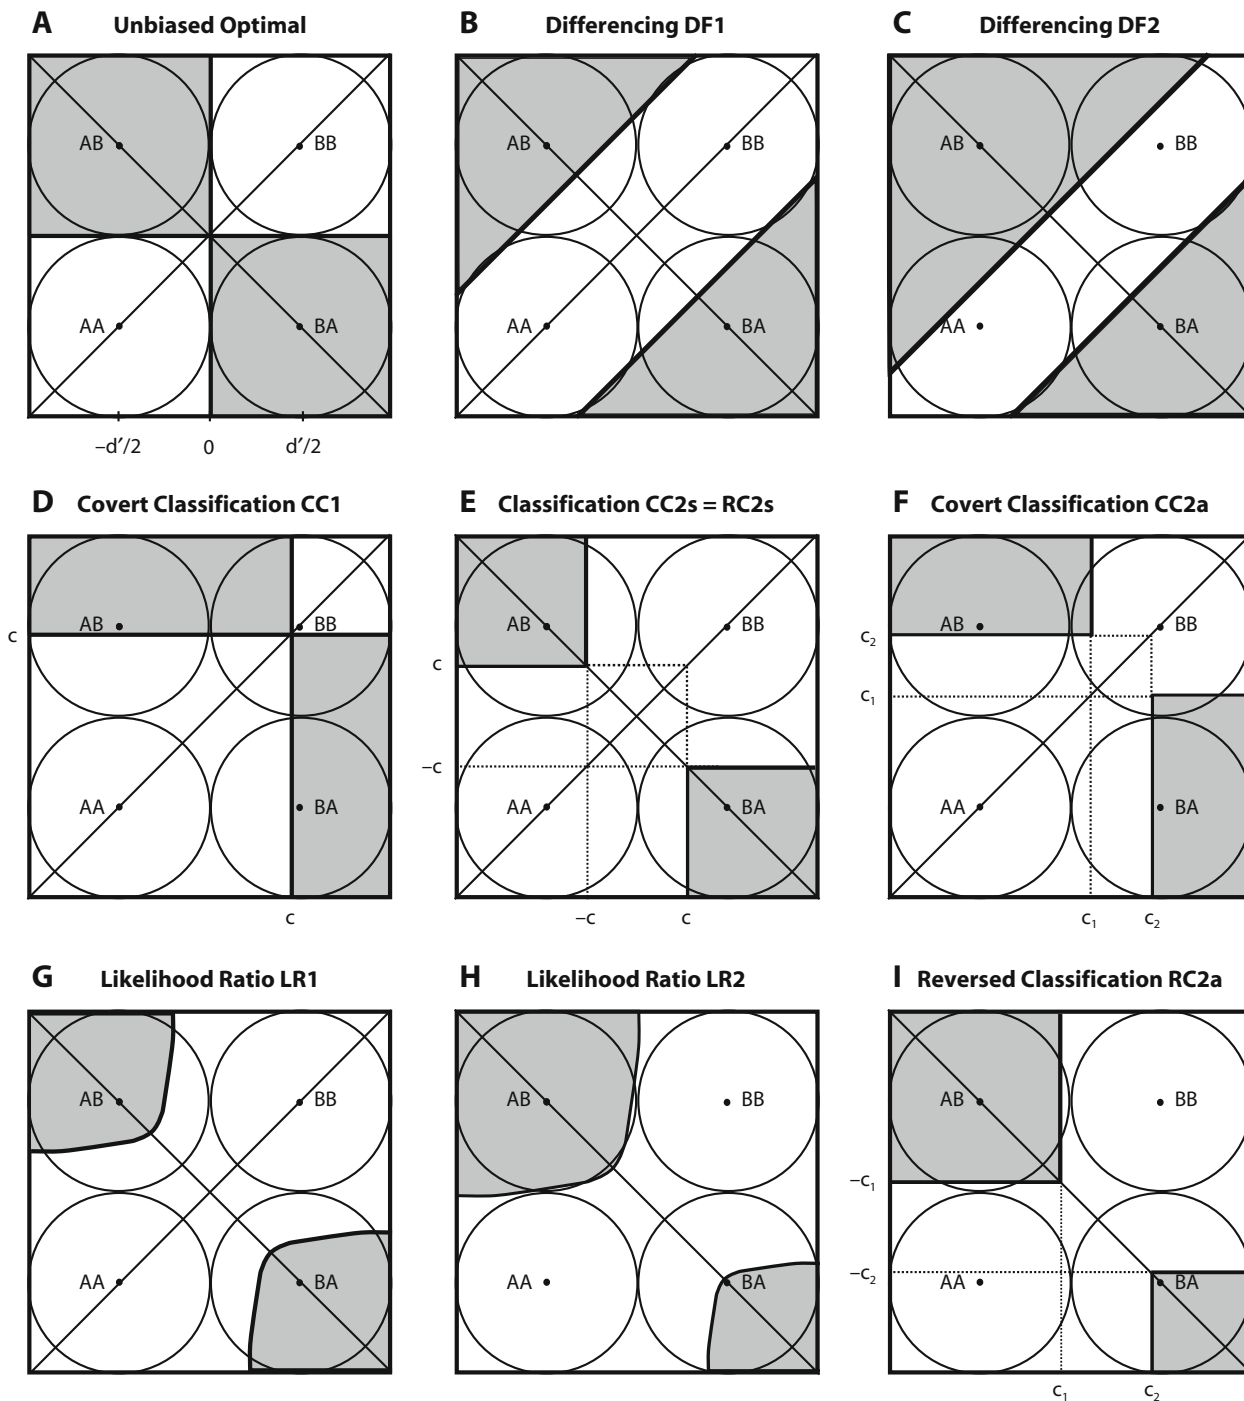


Figure 1. Schematic representation of the decision space. White areas are labeled “same” and shaded areas “different.” The perceptual effects of the first and second stimuli vary along the *x*- and *y*-axes, respectively. The four circles in each panel are contours of equal probability density for the joint effects of stimulus pairs AA, AB, BA, and BB. Panel A illustrates how an unbiased optimal observer partitions the space. Panels B–I illustrate the eight decision rules listed in Table 1 and discussed in the text. The diagonal(s) show the axis (or axes) of symmetry for each decision boundary. *c*, *c*₁, and *c*₂ denote covert-classification criteria.

This is a sensitivity measure that characterizes only the stimulus pair, not the decision rule or the response bias.

To produce a *same* or *different* response, the observer must partition the decision space. Figure 1A shows the strategy of an ideal, unbiased observer. The decision boundary consists

of two perpendicular straight lines, each equidistant from the centers of the respective magnitude distributions. These lines coincide with the axes in our chosen coordinate system. The two shaded quadrants are labeled “different”; the two white quadrants are labeled “same.” Table 1 lists eight

decision rules organized in three families: differencing, covert classification, and likelihood ratio. Figures 1A–II illustrate how each family partitions the decision space.

Differencing strategies. The defining feature of all differencing strategies is that they form an explicit representation of the difference, $\delta = (x_2 - x_1)$, of the stimulus magnitudes and that all subsequent processing is based on this difference. Magnitude pairs that differ by the same amount are treated as equivalent. This collapses the two-dimensional space $\langle x_1, x_2 \rangle$ to a single dimension δ . Geometrically, the differencing operation projects the plane onto the negative diagonal.

The simplest differencing strategy was proposed by Sorkin (1962). We refer to it as DF1 (differencing strategy with one criterion). It compares the absolute value of the difference with a criterion c . The observer responds “same” iff $|\delta| < c$. Equivalently, DF1 can be formulated with a pair of criteria equidistant from the neutral point. The observer responds “same” iff $-c < \delta < c$. Projected back to the two-dimensional space, the criteria become lines parallel to and equidistant from the positive diagonal (Figure 1B). The band that contains the origin is labeled “same”; the regions on either side are labeled “different.”

Exchanging the stimuli in the two presentation intervals reverses the sign of δ . Under DF1, however, this sign does not affect the final response. Thus, DF1 predicts identical response probabilities for stimulus pairs that differ only in the order of presentation. In particular, DF1 predicts $p_{AB} = p_{BA}$ for heterogeneous pairs. Our experimental data violate this identity and thus rule out DF1. The data are consistent with a more general strategy that we refer to as DF2 (differencing strategy with two criteria). It involves two arbitrary criteria that may not be equidistant from the neutral point. The observer responds “same” iff $-c_1 < \delta < c_2$. Figure 1C plots an example in which $p_{AB} > p_{BA}$.

DF1 is a special case of DF2 in which the two criteria satisfy the constraint $c_1 + c_2 = 0$.

Differencing strategies have several things to recommend them. DF1 maximizes accuracy in two-alternative forced choice (2AFC) designs, where the two stimuli are always different and the task is to indicate which one came first (Macmillan & Creelman, 2005). DF1 is also optimal in roving *same-different* designs, where absolute intensities vary but differences are kept constant (Dai et al., 1996). A plausible neural representation of δ has been found in physiological recordings with monkeys perform-

ing roving 2AFC tactile flutter discrimination: The firing rates of many neurons in secondary somatosensory cortex correlate with the stimulus difference (Romo, Hernández, Zainos, Lemus, & Brody, 2002).

On the other hand, the differencing strategy is suboptimal for fixed *same-different* designs, which are our main focus here (Irwin et al., 1999; Noreen, 1981). For example, suppose $d' = 3$ and the criterion is chosen so that the false alarm rate is 1%. Then the hit rate is approximately 50% under the optimal strategy (LR1) but only 32% under DF1. Conversely, a 50% hit rate implies $d' = 3.65$ under DF1 and $d' = 3$ under LR1.

To carry out a differencing strategy, the observer must maintain the first magnitude in working memory and then calculate the difference to the second magnitude. This can be effortful and error prone. To relieve the cognitive load, the observer may opt for a covert-classification strategy instead.

Covert-classification strategies. The defining feature of all covert-classification strategies is that they classify the stimuli individually. Neither memorization nor subtraction of continuous magnitudes is necessary.

The most straightforward covert-classification strategy was proposed by Pollack and Pisoni (1971). We refer to it as CC1 (covert classification with one criterion). The first stimulus is classified as either “A” or “B.” Then the second stimulus is classified *independently and with the same criterion*. The observer responds “same” iff the two labels match. Figure 1D plots the resulting partitioning of the decision space. When the criterion is unbiased, the overt responses are unbiased and optimal (Figure 1A; Noreen, 1981). When the criterion is biased, the overt responses are biased in favor of “same.” This prediction is in qualitative agreement with much *same-different* data (Macmillan & Creelman, 2005). Quantitatively, however, CC1 requires unrealistically high covert biases to account for the observed probabilities of responding “same.”

A preference for *same* responses seems natural with hard-to-discriminate stimuli. On many trials, the stimuli are perceived as similar because neither of them can be (covertly) classified with confidence. This idea can be formalized by a decision rule with two criteria. The first stimulus is labeled “A” if its magnitude is below c_1 , “B” if the magnitude is above c_2 , and “ambiguous” otherwise. Then the second stimulus is classified independently and with the same criteria. The observer responds “different” iff the labels are unambiguous and different. All ambiguous stimuli produce a *same* response. When the two criteria satisfy the constraint $c_1 + c_2 = 0$, we have a covert-classification strategy with two symmetric criteria (CC2s; Figure 1E). Without this constraint, we have a CC strategy with two arbitrary criteria (CC2a; Figure 1F).

Reversed-classification strategies—the final members of this family—formalize the following intuition. Suppose that the first magnitude exceeds some high criterion and is confidently labeled as “B.” To respond “different,” the observer demands that the second magnitude be classified with comparable confidence in the opposite category, demarcated by an equally stringent criterion of the opposite sign. Thus, if the first stimulus is classified as “B” when $x_1 > c_2$, the second stimulus is classified as “A” when

Table 1
Abbreviated Labels of the Decision Strategies
Discussed in the Text

Label	Decision Strategy
DF1	Differencing strategy with one criterion
DF2	Differencing strategy with two criteria
CC1	Covert classification (CC) with one criterion
CC2a	CC with two asymmetric criteria and bias for <i>same</i> responses
CC2s	CC with two symmetric criteria and bias for <i>same</i> responses
RC2a	Reversed-classification with two asymmetric criteria
RC2s	Equivalent to CC2s
LR1	Likelihood-ratio strategy with one criterion
LR2	Likelihood-ratio strategy with two criteria

$x_2 < -c_2$. If the first stimulus is classified as “A” when $x_1 < c_1$, the second stimulus is classified as “B” when $x_2 > -c_1$. With symmetric criteria ($c_1 = -c_2$), the resulting rule coincides with the CC2s rule. With asymmetric criteria, however, RC2a and CC2a partition the decision space differently and predict different patterns among the observed probabilities (Figures 1I and 1F).

Likelihood-ratio strategies. The third family of decision rules involves likelihood ratios. They guarantee statistically optimal performance but require substantial computational power and detailed knowledge of the parameters of the perceptual distributions. This is not psychologically realistic, especially in the beginning of the experimental session. Thus, likelihood-ratio rules are interesting mainly as ideal-observer benchmarks.

Let $L_{d/s}$ denote the likelihood ratio of *different* over *same* responses—that is, the likelihood that $\langle x_1, x_2 \rangle$ is drawn from AB or BA divided by the likelihood that it is drawn from AA or BB. Under uniform presentation frequencies and standard representational assumptions, this ratio is given in Equation 3 (Dai et al., 1996; Noreen, 1981). $L_{d/s}$ defines a saddle-shaped surface that opens up along the positive diagonal and down along the negative diagonal. The likelihood-ratio strategy with one criterion (LR1) is to respond “different” iff $L_{d/s} > c$. The decision boundary consists of two disjoint nonlinear curves. When $c > 1$, there is an overall bias for *same* responses, and the boundary intersects the negative diagonal as is illustrated in Figure 1G. When $c < 1$, there is an overall bias for *different* responses, and the boundary intersects the positive diagonal. When $c = 1$, there is no overall bias, and the boundary consists of two perpendicular straight lines (Figure 1A).

$$L_{d/s}(x_1, x_2) = \frac{e^{x_1 d'} + e^{x_2 d'}}{e^{(x_1 + x_2) d'} + 1}. \quad (3)$$

Data sets in which $p_{AB} \neq p_{BA}$ are inconsistent with the LR1 strategy. The likelihood-ratio strategy with two criteria (LR2) is designed to accommodate such cases. It postulates separate criteria for the two branches of the decision boundary. Figure 1H plots an example in which $p_{AB} > p_{BA}$. The decision process involves two consecutive comparisons. First, adopt criterion c_1 when $x_1 < 0$ and criterion c_2 otherwise. Second, respond “different” iff $L_{d/s}$ exceeds the criterion adopted on the first step.

The curvilinear contours of equal likelihood ratios can be approximated by angle-shaped piecewise-linear contours (Irwin & Hautus, 1997). In our notation, the LR1 strategy is approximated by the RC2s strategy and, equivalently, by the CC2s strategy (cf. Figures 1G and 1E). The LR2 strategy is approximated by the RC2a strategy introduced here for the first time (cf. Figures 1H and 1I). The CC1 strategy is optimal only in the special case of unbiased criterion (Figure 1A).

Symmetries and Their Observable Consequences

So, a participant in a *same-different* experiment can adopt a variety of decision rules organized in three families. Our goal is to develop a framework for making observation-based inferences about these rules. The notion of axial symmetry plays a key role in this framework.

Definition. A curve is symmetric with respect to an axis when the mirror image of each point on the curve also lies on the curve.

Our main leverage comes from the fact that the decision boundaries of different strategies are symmetric with respect to different axes; the positive diagonal (through $\langle \mu_A, \mu_A \rangle$ and $\langle \mu_B, \mu_B \rangle$), the negative diagonal (through $\langle \mu_A, \mu_B \rangle$ and $\langle \mu_B, \mu_A \rangle$), or both. Figures 1A–1I are a visual guide to these symmetries. Each panel plots the axes of symmetry of the corresponding strategy. Given that the perceptual distributions are symmetric too, each decision rule gives rise to a characteristic pattern of equalities among the response probabilities. These patterns are summarized in Theorem 1, which is the main theoretical result in this article.

Theorem 1. *Under the representational assumptions listed in the previous section, the decision rules in Table 1 make the following parameter-free predictions about the response probabilities defined in Equation 2:*

- DF1 implies $p_{AA} = p_{BB}$ and $p_{AB} = p_{BA}$
- DF2 implies $p_{AA} = p_{BB}$
- CC1 implies $p_{AB} = p_{BA}$ and $(p_{AB} - .5)^2 = (p_{AA} - .5)(p_{BB} - .5)$
- CC2s implies $p_{AB} = p_{BA}$ and $p_{AA} = p_{BB}$ and $p_{AB} \geq p_{AA}$
- CC2a implies $p_{AB} = p_{BA}$ and $p_{AB}^2 \geq p_{AA} p_{BB}$
- RC2a implies $p_{AA} = p_{BB}$ and $p_{AB} p_{BA} \geq p_{AA}^2$
- LR1 implies $p_{AA} = p_{BB}$ and $p_{AB} = p_{BA}$
- LR2 implies $p_{AA} = p_{BB}$ when there is an overall bias for “same” responses and $p_{AB} = p_{BA}$ when there is an overall bias for “different” responses.

The proof is straightforward and is given in Appendix A. In a nutshell, symmetric decision boundaries demarcate symmetric response regions, the perceptual distributions have (mutually) symmetric probability density functions, and the integrals of symmetric densities over symmetric regions are equal.

Theorem 1 reveals the implicit regularities of each family. The differencing family is characterized by $p_{AA} = p_{BB}$. Any observer who follows any differencing strategy must show this pattern, regardless of their sensitivity or criteria. The reason is built into the decision rule itself: The difference, $\delta = (x_2 - x_1)$, always has the same (zero) mean when x_1 and x_2 are drawn from distributions with identical means. Thus, the distinction between AA and BB is lost, because each of these pairs projects to the same δ .

The covert-classification family, with the exception of RC2a, is characterized by $p_{AB} = p_{BA}$. Again, the reason is built into the decision rule itself: The XOR rule that combines the covert-classification labels is order invariant. Specifically, the overt response for an “A” followed by a “B” matches that for a “B” followed by an “A.” Therefore, any strategy that applies the same covert-classification criteria on both presentation intervals predicts identical performance for the two heterogeneous stimulus pairs. (Homogeneous pairs break the XOR symmetry because of their nonoverlapping cross-classification possibilities.)

The reversed-classification rule with two asymmetric criteria (RC2a, Figure 1I) is exceptional, because it applies different criteria on the two intervals.

The characterization of the likelihood-ratio family depends on the overall response bias. Typically, observers favor *same* over *different* responses. In those cases, all LR strategies predict $p_{AA} = p_{BB}$. This is because cutting the saddle surface defined by Equation 3 above the neutral point produces decision boundaries that are symmetric with respect to the negative diagonal (Figure 1H). Strategy RC2a, designed to approximate LR2, also follows this pattern. In the complementary (and rare) case, the saddle is cut below the neutral point, and the axis of symmetry flips to the positive diagonal. The observable manifestations are a preponderance of *different* responses and $p_{AB} = p_{BA}$.

Finally, all three families contain a member whose decision boundaries are symmetric with respect to both diagonals. Thus, strategies DF1, CC2s, and LR1 predict both $p_{AA} = p_{BB}$ and $p_{AB} = p_{BA}$.

METHODOLOGY

These predictions can be used as empirical tests to rule out individual strategies or even whole families. The strength of Theorem 1 is that it makes *qualitative* predictions about the contrasting families. When such a prediction is violated in a given data set, an entire family can be ruled out from further consideration. This process is referred to as *strong inference* (Platt, 1964) or *falsification* (Popper, 1963) in philosophy of science. It differs from model fitting, where families compete via their best-fitting representatives. In strong inference, families are ruled out when their qualitative predictions are falsified. The two methodologies are complementary. Strong inference is applied first to narrow the field. No parameter search is necessary. Model fitting, which requires search, can then be applied to the remaining candidates.

We recommend that the analysis of every *same-different* data set begin with two qualitative tests: Is p_{AA} equal to p_{BB} ? Is p_{AB} equal to p_{BA} ? This leads to the following four cases.

Case I: $p_{AA} = p_{BB}$ but $p_{AB} \neq p_{BA}$

When there are statistically significant differences between the probabilities to respond “different” to stimuli AB and BA, strategies CC1, CC2s, CC2a, DF1, and LR1 can be rejected. The likelihood-ratio strategy LR2 can be rejected too when there is an overall preponderance of *different* responses. Otherwise, LR2 is technically compatible with the data. It remains psychologically implausible, however, because it depends on unrealistic amounts of knowledge and processing power.

This leaves only two main contenders: the differencing rule with two criteria (DF2) and the reversed-classification rule with two asymmetric criteria (RC2a). Unfortunately, they imply different d' s. The RC2a model approximates the optimal decision strategy (LR2) and thus needs higher perceptual noise to account for the observed error levels. The DF2 model accounts for them through a combina-

tion of perceptual noise and suboptimal decision making. Thus, d'_{RC2a} is always less than d'_{DF2} on the same data.¹

Case II: $p_{AB} = p_{BA}$ but $p_{AA} \neq p_{BB}$

When there are statistically significant differences between the probabilities to respond “different” to stimuli AA and BB, strategies DF1, DF2, RC2a, CC2s, LR1, and LR2² can be rejected.

This leaves two covert-classification strategies: CC1 and CC2a. They predict identical d' s, because the former is a special case of the latter. CC1 can be identified by its strong prediction $(p_{AB} - .5)^2 = (p_{AA} - .5)(p_{BB} - .5)$, proven in Appendix B.

Case III: $p_{AA} = p_{BB}$ and $p_{AB} = p_{BA}$

This is the case that is implicitly assumed throughout the literature (e.g., Dai et al., 1996; Macmillan & Creelman, 2005; Noreen, 1981). There is a common hit rate, $H = p_{AB} = p_{BA}$, and a common false alarm rate, $FA = p_{AA} = p_{BB}$ (Equation 1). All asymmetric rules (DF2, CC1, CC2a, and RC2a) can be rejected. Two main contenders remain: the differencing strategy with one criterion (DF1) and the covert-classification strategy with two symmetric criteria (CC2s, which coincides with RC2s). The CC2s rule closely approximates the optimal LR1 rule (Irwin & Hautus, 1997) and is much more psychologically plausible. As in Case I, the near-optimal rule implies lower d' than the differencing rule: $d'_{CC2s} < d'_{DF1}$. The discrepancy is most pronounced when the probability of responding “different” is small (Macmillan & Creelman, 2005, p. 224).

Case IV: $p_{AA} \neq p_{BB}$ and $p_{AB} \neq p_{BA}$

When both symmetries are broken, all decision rules in Table 1 must be rejected. The assumptions of independence, stimulus unidimensionality, and/or negligible memory noise are probably violated. There may be strong sequential and/or configural effects that are not captured by the signal detection framework. The observers may be following a memory-based strategy (Cohen & Nosofsky, 2000).

Statistical Test

The present methodology requires a test of statistical difference between two proportions. We recommend the chi-squared test³ because of its versatility. Let n_i denote the number of presentations of stimulus pair i during an experimental block, and let y_i denote the observed number of *different* responses. The proportions $\hat{p}_i = y_i/n_i$ estimate the true response probabilities p_i . The null hypothesis is that these probabilities are equal for the two stimulus pairs under consideration: $p_1 = p_2 = p$. The pooled proportion \hat{p} in Equation 4 allows a more stable estimate of the sampling variance. When the null hypothesis is correct, the χ^2 statistic in Equation 5 has a χ^2 distribution on one degree of freedom (Collett, 2003);

$$\hat{p} = \frac{y_1 + y_2}{n_1 + n_2}, \quad (4)$$

$$\chi^2 = \frac{n_1(\hat{p}_1 - \hat{p})^2 + n_2(\hat{p}_2 - \hat{p})^2}{\hat{p}(1 - \hat{p})}. \quad (5)$$

These equations must be applied to observations from a single experimental block for a single participant. To aggregate the data across N blocks and/or participants, one should calculate N individual X^2 s and then add them up. The sum has a χ^2 distribution on N degrees of freedom. One cannot simply add the raw counts, because when $p_1 < p_2$ for some participants and $p_1 > p_2$ for others, the inequalities cancel out with misleading effect.

Resources for Calculating d' and Bias

Appendix B outlines algorithms for fitting all covert-classification models (CC1, CC2a, and RC2a). Approximate d' formulas for the optimal (LR1) model are given by Macmillan and Creelman (2005, Equation 9.3; see also Dai et al., 1996). An algorithm for fitting the differencing (DF1) model is given by Macmillan and Creelman (2005, Equation 9.8). Open-source MATLAB scripts implementing these algorithms and calculating the corresponding d' are available at <http://alexpetrov.com>.

Various indices of response bias have been developed for the LR1 and DF1 strategies (Hautus & Collins, 2003; Irwin et al., 2001). They all assume that the bias can be described with a single number. Consequently, none of them is applicable in Cases I or II above.

AN EXPERIMENTAL EXAMPLE

We illustrate this methodology in a perceptual learning experiment on visual motion-direction discrimination. Single-interval (*yes-no*) and two-interval (*same-different*) blocks were alternated within subjects for six practice sessions. Only the *same-different* data are reported here; the *yes-no* data and their implications for perceptual learning are reported separately (Petrov, 2009).

Method

Stimuli and Tasks. Each stimulus consisted of 80 black dots that moved coherently along parallel tracks inside a white circular aperture. The direction of motion could take four possible values: -52.5° versus -47.5° from vertical for the negative and 37.5° versus 42.5° for the positive reference direction. The trials were organized in blocks. The reference direction and the task were fixed in each block and varied between blocks. Each block began with four demo trials that explicitly indicated the reference direction with a line drawn on the screen. The reference line was not shown during the experimental trials.

The single-interval task involved a single stimulus presentation per trial. The observer was instructed to indicate whether the direction of motion was counterclockwise or clockwise from the reference direction. In a negative block, for example, the task was to press one button for -52.5° and another button for -47.5° . The two possible stimuli were presented with equal frequencies and in random order.

The *same-different* task involved two consecutive stimulus presentations per trial. The two directions of motion were either identical or differed by 5° . The observer was instructed to press one button for *same* and another button for *different*. The four possible stimulus pairs were presented with equal frequencies and in random order.

Observers. Thirteen students with normal or corrected-to-normal vision participated in the study. They were paid \$6/h plus a bonus contingent on their accuracy.

Design. Each observer completed six sessions: a pretest on Day 1, followed by four practice sessions, followed by a posttest on Day 6. Each session consisted of four blocks of 240 trials each. The task and reference direction were manipulated within subjects.

The single-interval task was tested during Blocks 1 and 3 on Days 1 and 6 and was practiced during Blocks 1 and 2 on Days 2–5. The *same-different* task was tested during Blocks 2 and 4 on Days 1 and 6 and was practiced during Blocks 3 and 4 on Days 2–5.

Seven participants practiced the negative reference direction throughout the experiment, except on Blocks 1 and 2 of the pre- and posttest, which tested transfer to the positive direction. The directions were reversed for the other 6 participants.

Procedure. Single-interval trials began with a short beep followed 500 msec later by a 400-msec motion stimulus. Two-interval trials continued with a 300-msec interstimulus interval followed by a second 400-msec stimulus. The observers typed their response on the computer keyboard. They scored a bonus point for each correct response and lost a point for each error. The bonus was displayed above the fixation dot at all times. There was also auditory feedback (an unpleasant beep) after errors. See Petrov (2009) for additional methodological details.

Results and Discussion

Following our methodology for the *same-different* data, we asked first whether p_{AA} equals p_{BB} . The answer seemed to be *yes* for 11 participants in our sample. None of the 11 individual X^2 values was statistically significant [median $X^2(12) = 14.2$; Equation 5]. The 2 remaining participants violated this equality [$X^2(12) = 32.4$ and 45.5 , $p < .001$]. Next, we asked whether p_{AB} equals p_{BA} . The answer is *no* for the same 11 participants [median $X^2(12) = 63.3$, $p < .001$] and *yes* for the 2 exceptional ones [$X^2(12) = 19.8$ and 14.9 , n.s.].

Thus, the great majority of our observers fall into Case I of our classification: $p_{AA} = p_{BB}$ but $p_{AB} \neq p_{BA}$. Figure 2 illustrates this qualitative pattern. Most of the points on the top panel are within the 95% confidence ellipse for the equality $p_{AA} = p_{BB}$. The situation on the bottom panel is markedly different. The constraint $p_{AB} = p_{BA}$ is clearly violated.

According to Theorem 1, these 11 participants could not have been using any of the following strategies: CC1, CC2s, CC2a, DF1, or LR1. The rejection of the covert-classification (CC) family is particularly important, because it prevents a tempting mistake in the analysis. It would have been tempting to assume that many observers would adopt a covert-classification strategy in our experimental design. The single-interval blocks in the beginning of each session demanded overt classification. The path of least resistance for the subsequent *same-different* blocks with the same stimuli would have been simply to switch to covert classification. The inequality $p_{AB} \neq p_{BA}$ proves that the participants did not adopt this strategy. It also indicates that the standard analytic procedure, which tacitly assumes $p_{AB} = p_{BA} = H$, produces misleading d' estimates for these data!

Our methodology also identifies individual differences. Two participants showed a qualitative pattern that differed from that of the majority: $p_{AB} = p_{BA}$ but $p_{AA} \neq p_{BB}$. This suggests that these 2 observers adopted the covert-classification strategy with asymmetric criteria (CC2a). Note that the standard analytic procedure is not applicable to these data either.

Figure 3 illustrates the individual differences. The top panel shows a representative member of the majority; the bottom panel shows one of the two exceptions. Circles plot hit probabilities (p_{AB} vs. p_{BA} for each block), Xs plot false alarm probabilities (p_{AA} vs. p_{BB}), and lines connect




Figure 2. Probability to respond “different” to stimulus pairs AA versus BB (top) and AB versus BA (bottom). The data are consistent with the constraint $p_{AA} = p_{BB}$ (top) but violate $p_{AB} = p_{BA}$ (bottom). The ellipses are 95% confidence regions. Each scatterplot contains 156 data points (12 blocks \times 13 observers) jittered slightly for visibility. Each probability is estimated from 60 binary responses.

consecutive blocks. Although the overall hit and false alarm rates were very similar for these 2 observers, they followed different response strategies. In the top panel, more circles than crosses lay outside the ellipse, whereas the reverse is true for the bottom panel. The standard analytic procedure is blind to such individual differences.

In summary, the methodology pays off. It pinpointed the decision rule for 2 of the observers and dramatically narrowed the field of possible rules for the rest. All of the observers seemed to use asymmetric decision criteria,

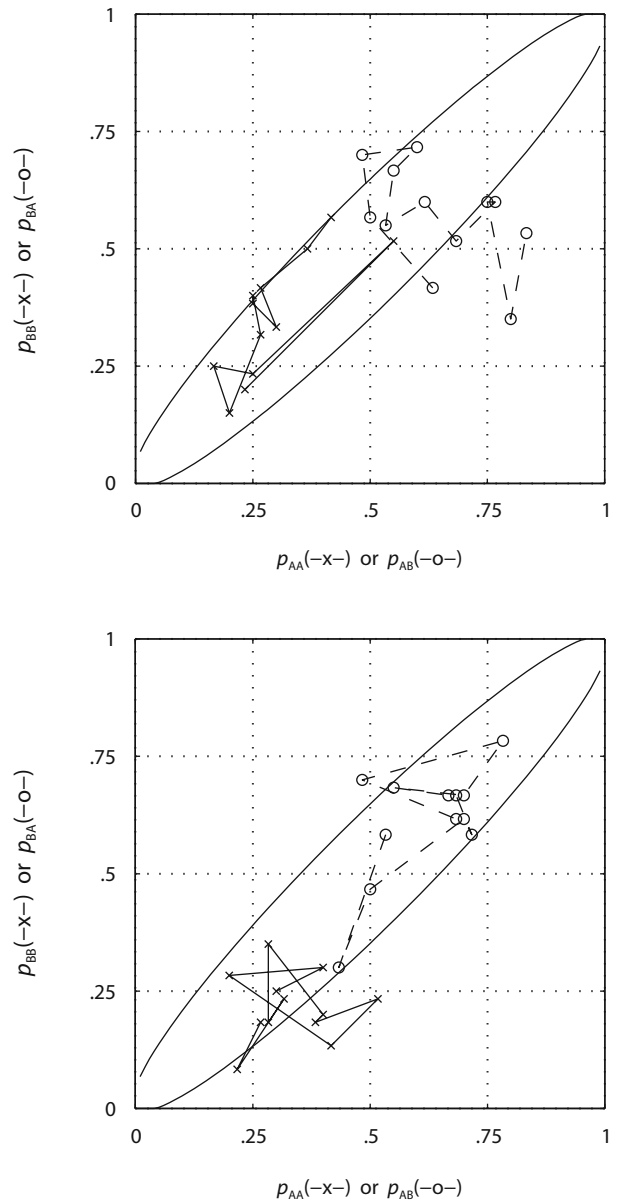


Figure 3. Probability to respond “different” to stimulus pairs AA versus BB (xs) and AB versus BA (Os). Lines connect consecutive experimental blocks. The two panels present two individual observers who follow different response strategies, giving rise to different qualitative patterns: $p_{AA} = p_{BB}, p_{AB} \neq p_{BA}$ (top), as opposed to $p_{AA} \neq p_{BB}, p_{AB} = p_{BA}$ (bottom). The ellipses are 95% confidence regions that the corresponding probabilities are equal.

which violates the tacit assumption of the standard analytic procedure.

Some ambiguity still remains, however. The majority of the data sets are consistent with two qualitatively different strategies: DF2 and RC2. This illustrates an important limitation of the methodology. Although it does narrow the field, it cannot fully disambiguate the decision rule when $p_{AA} = p_{BB}$.

For some observers, the ambiguity can be resolved by complementary data from the single-interval blocks.

There were observers, however, whose strategy remained underdetermined by all available data. Figure 4 illustrates 1 individual of each kind. In the top panel, the *yes-no* learning curve (Xs with error bars) seems to match the *same-different* learning curve assuming DF2 strategy (circles). The comparison must be done at the level of learning curves, because d' improves across sessions. The error bars span approximate 95% confidence intervals for the *yes-no* estimates (Gourevitch & Galanter, 1967, tabulated in Miller, 1996). Note that the differencing strategy (circles) always implies a higher d' than the near-optimal RC2a strategy (triangles). The discrepancy varies across observers depending on biases and other factors. For the observer in the top panel, the two *same-different* curves differ by approximately 15%, and the *yes-no* data seem to favor the DF2 strategy. For the observer in the bottom panel, however, the two *same-different* curves are closer together, and the *yes-no* curve seems equally consistent with either strategy.

GENERAL DISCUSSION

In this article, we developed a simple but powerful idea: Take full advantage of the fact that the *same-different* task yields four probability estimates per condition (Equation 2). We examined eight decision rules organized in three families: differencing, covert classification, and likelihood ratio (Table 1). Each of them partitions the space in a particular way and gives rise to various symmetries (Figure 1) that impose constraints on the overt response probabilities. Theorem 1—our main theoretical result—derives the specific pattern for each strategy. This leads to a systematic methodology for strong inference (Platt, 1964; Popper, 1963) based on two simple qualitative tests. Are the two homogeneous stimulus pairs (AA and BB) labeled “different” with equal probability? If not, all differencing and likelihood-ratio strategies are ruled out. Are the two heterogeneous stimulus pairs (AB and BA) labeled “different” with equal probability? If not, all covert-classification strategies (except RC2a) are ruled out.

We presented experimental evidence that neither of these constraints can be taken for granted. Yet it is standard practice to take them both for granted. Reducing the data to hit and false alarm rates tacitly assumes that $p_{AB} = p_{BA}$ and $p_{AA} = p_{BB}$. None of the 13 observers in our sample (nor 15 others in a follow-up experiment; Petrov, 2009) showed this pattern. The generality of this finding is a matter of further investigation, but it is probably very general. Even if these assumptions are only rarely violated, however, it still behooves the experimenter to verify them in each case. The relevant statistical test is straightforward (Equation 5). We recommend it as a standard preamble to every report of *same-different* data.

The test is so simple that it can be applied to published data, provided that the four response probabilities are reported separately. We analyzed one such report (Wang et al., 2005), which is exemplary in its thoroughness. The chi-squared test of the data in their Table V suggests that 4 of their 6 observers violated $p_{AA} = p_{BB}$. This rules out the very strategies (differencing and likelihood ratio)

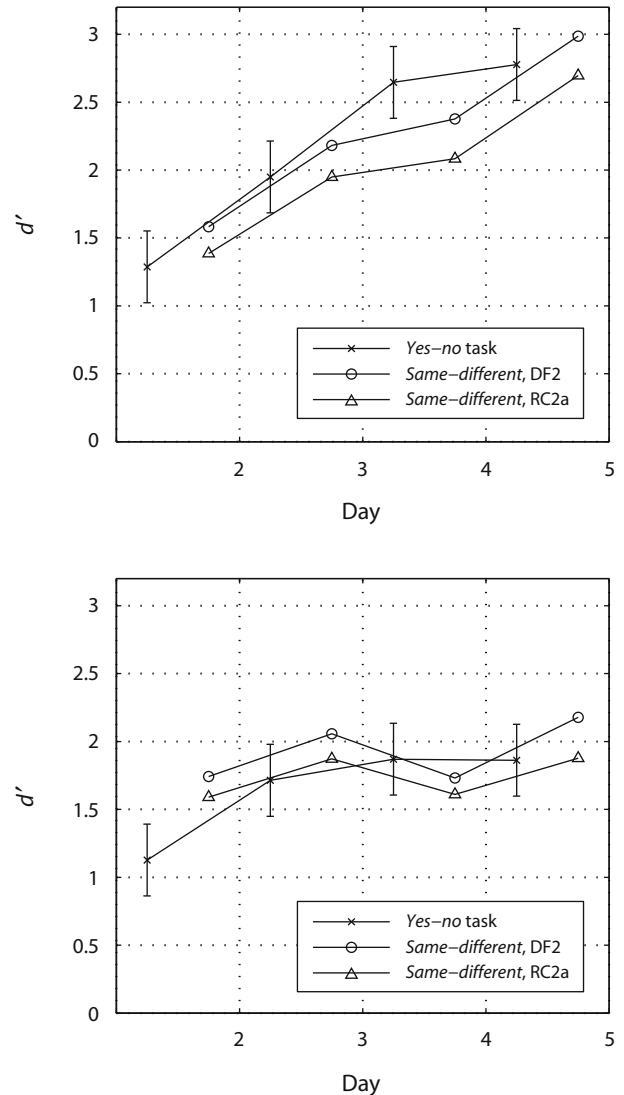


Figure 4. Comparison of the discrimination performance on *yes-no* and *same-different* tasks for two observers. The individual in the top panel seems to follow a differencing (DF2) strategy. The data in the bottom panel are consistent with both differencing and reversed classification (RC2a). Each data point is based on 480 binary observations. The error bars span approximate 95% confidence intervals for the *yes-no* d' estimates. The pre- and posttests on Days 1 and 6 are omitted for simplicity.

that Wang et al. set out to investigate! The 2 remaining observers satisfied both symmetry conditions and apparently followed the DF1 strategy.

Our data also revealed individual differences. It cannot be taken for granted that all observers in a given experiment follow the same decision rule. The symmetry-based methodology can sort the participants into equivalence classes in a principled way, and the data analysis can be tailored accordingly. In Appendix B, we develop algorithms for fitting the covert-classification and reversed-classification models and for calculating the associated d' .

All of the decision rules in Table 1 impose at least one constraint on the observed probabilities. (It is reassuring

that no observer in our sample violated both equalities at the same time.) Therefore, a *same-different* data set has at most 3 *dfs* under these rules. This gives rise to one d' estimate and two criterion estimates. Indices of response bias (Irwin et al., 2001) that reduce all nonperceptual factors to a single number are applicable only in special cases with 2 *dfs*.

Theorem 1 was proven for fixed designs with only two individual stimuli. It can be extended to roving designs that vary the individual stimuli in a range of possible values. The differencing strategy is very natural (and optimal; Dai et al., 1996) in roving designs. Theorem 1 still applies to any subset of data that involves two specific stimuli.

Criterion Noise and Memory Noise

Throughout this article, the boundaries between the response categories were assumed to be relatively stable. This simplifying assumption is not critical, provided the presentation sequence is randomized. Response criteria do shift in psychophysical tasks, and these shifts often correlate with previous stimuli and responses (e.g., Lages & Treisman, 1998; Mueller & Weidemann, 2008). However, the criteria remain fixed for the duration of each particular trial and, with proper randomization, do not correlate with the stimuli on that trial. Thus, the response probabilities still satisfy the equality constraints in Theorem 1 on a trial-by-trial basis. These equalities are preserved when multiple trials with varied criteria are mixed in a block.⁴

It was also assumed that the noise in visual working memory (Pasternak & Greenlee, 2005) is negligible relative to perceptual noise. This simplifying assumption is not critical for the present results either. There is evidence that memory noise can constrain performance in two-interval tasks (e.g., Johnson, 1980; Vogels & Orban, 1986). Nevertheless, the qualitative constraints in Theorem 1 remain approximately true. Memory noise degrades the performance in the differencing and likelihood-ratio models, but the degradation is virtually the same for all four stimulus pairs. We verified this in Monte Carlo simulations at various levels of memory noise. Even when the memory variance was four times greater than the perceptual variance, the discrepancy between the observed probabilities p_{AA} and p_{BB} was insignificant relative to the binomial sampling fluctuations. Thus, the critical prediction $p_{AA} = p_{BB}$ retains its diagnostic value. As for covert classification, its memory load is minimal, because only a verbal label is kept during the interstimulus interval.

Related Research

This article continues a long line of theoretical analyses of the *same-different* task from a signal detection perspective (e.g., Dai et al., 1996; Green & Swets, 1966; Irwin et al., 2001; Macmillan & Creelman, 2005; Noreen, 1981). Our analysis applies to unidimensional stimuli only. The observers were instructed to attend a single attribute, and only this attribute was relevant to the task. The general recognition theory (GRT; Ashby & Townsend, 1986) extends the signal detection framework to multidimensional stimuli, such as colors or faces. GRT offers a sophisticated methodology to explore various kinds of separability and

integrality. A key empirical test is whether performance on one attribute is affected by task-irrelevant variation of another attribute (e.g., Garner, 1974). This manipulation is incompatible with the unidimensional *same-different* task. It can be applied to multidimensional *same-different* judgments (Thomas, 1996), but the decision space becomes very complex (at least four-dimensional). GRT theorists (e.g., Ashby & Gott, 1988) have explored linear and quadratic decision rules in spaces of arbitrary dimensionality.

Cohen and Nosofsky (2000) presented evidence that *same-different* judgments may rely in part on long-term memories for past stimulus pairs. Their extension of the exemplar-based random walk model (Nosofsky & Palmeri, 1997) provides an interesting link to the literature on speeded classification. In roving designs, the model generally predicts asymmetries for both hit and false-alarm probabilities and response times. For fixed designs, the predictions would depend on factors such as pair presentation frequency and equal-variance assumptions (R. Nosofsky, personal communication, June 1, 2009).

On the experimental side, there are a few comparisons of single-interval and two-interval discrimination designs (see Macmillan & Creelman, 2005, p. 228, for a review). The comparison of the method of constant stimuli with two-interval forced choice has been of particular interest (Morgan, Watamaniuk, & McKee, 2000; Nachmias, 2006; Vogels & Orban, 1986). To our knowledge, our experiment is one of very few comparisons involving a *same-different* task and the first comparison of any kind in a learning context. See our empirical report (Petrov, 2009) for a detailed discussion.

Take-Home Message

In closing, we reiterate our recommendation to begin the analysis of every *same-different* dataset with these simple questions: Are stimulus pairs AA and BB labeled “different” with equal probability? What about AB and BA? The statistical test is straightforward, and the answers are very informative. Add them to your toolbox!

AUTHOR NOTE

The author thanks James Todd and Vilislav Boutchaktchiev. Correspondence concerning this article should be addressed to A. A. Petrov, Department of Psychology, 200B Lazenby Hall, Ohio State University, Columbus, OH 43210 (e-mail: apetrov@alexpetrov.com).

REFERENCES

- ASHBY, F. G., & GOTT, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 33-53.
- ASHBY, F. G., & TOWNSEND, J. T. (1986). Varieties of perceptual independence. *Psychological Review*, **93**, 154-179.
- BERG, B. G. (1989). Analysis of weights in multiple observation tasks. *Journal of the Acoustical Society of America*, **86**, 1743-1746.
- COHEN, A. L., & NOSOFSKY, R. M. (2000). An exemplar-retrieval model of speeded same-different judgments. *Journal of Experimental Psychology: Human Perception & Performance*, **26**, 1549-1569.
- COLLETT, D. (2003). *Modeling binary data* (2nd ed.). Boca Raton, FL: Chapman & Hall/CRC.
- DAI, H., VERSFELD, N. J., & GREEN, D. M. (1996). The optimum decision rules in the *same-different* paradigm. *Perception & Psychophysics*, **58**, 1-9.

- FRANCIS, M. A., & IRWIN, R. J. (1995). Decision strategies and visual-field asymmetries in same-different judgments of word meaning. *Memory & Cognition*, **23**, 301-312.
- GARNER, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.
- GOUREVITCH, V., & GALANTER, E. (1967). A significance test for one parameter isosensitivity functions. *Psychometrika*, **32**, 25-33.
- GREEN, D. M., & SWETS, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- HAUTUS, M. J., & COLLINS, S. (2003). An assessment of response bias for the same-different task: Implications for the single-interval task. *Perception & Psychophysics*, **65**, 844-860.
- IRWIN, R. J., & FRANCIS, M. A. (1995). Perception of simple and complex visual stimuli: Decision strategies and hemispheric differences in same-different judgments. *Perception*, **24**, 787-809.
- IRWIN, R. J., & HAUTUS, M. J. (1997). Likelihood-ratio decision strategy for independent observations in the same-different task: An approximation to the decision-theoretic model. *Perception & Psychophysics*, **59**, 313-316.
- IRWIN, R. J., HAUTUS, M. J., & BUTCHER, J. C. (1999). An area theorem for the same-different experiment. *Perception & Psychophysics*, **61**, 766-769.
- IRWIN, R. J., HAUTUS, M. J., & FRANCIS, M. A. (2001). Indices of response bias in the same-different experiment. *Perception & Psychophysics*, **63**, 1091-1100.
- JOHNSON, K. O. (1980). Sensory discrimination: Decision process. *Journal of Neurophysiology*, **43**, 1771-1792.
- KOSTRIKIN, A. I., & MANIN, Y. I. (1989). *Linear algebra and geometry*. New York: Gordon and Breach Science Publishers.
- LAGES, M., & TREISMAN, M. (1998). Spatial frequency discrimination: Visual long-term memory or criterion setting? *Vision Research*, **38**, 557-572.
- LAMING, D. (1986). *Sensory analysis*. London: Academic Press.
- MACMILLAN, N. A., & CREELMAN, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, NJ: Erlbaum.
- MATHWORKS (1999). *MATLAB user's guide*. Natick, MA: The MathWorks, Inc.
- MILLER, J. (1996). The sampling distribution of d' . *Perception & Psychophysics*, **58**, 65-72.
- MINSKY, M. L., & PAPERT, S. A. (1969). *Perceptrons*. Cambridge, MA: MIT Press.
- MORGAN, M. J., WATAMANIUK, S. N. J., & MCKEE, S. P. (2000). The use of an implicit standard for measuring discrimination thresholds. *Vision Research*, **40**, 2341-2349.
- MUELLER, S. T., & WEIDEMANN, C. T. (2008). Decision noise: An explanation for observed violations of signal detection theory. *Psychonomic Bulletin & Review*, **15**, 465-494.
- NACHMIAS, J. (2006). The role of virtual standards in visual discrimination. *Vision Research*, **46**, 2456-2464.
- NOREEN, D. L. (1981). Optimal decision rules for some common psychophysical paradigms. In S. Grossberg (Ed.), *Mathematical psychology and psychophysiology* (pp. 237-280). Providence, RI: American Mathematical Society.
- NOSOFSKY, R. M., & PALMERI, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, **104**, 266-300.
- PASTERNAK, T., & GREENLEE, M. W. (2005). Working memory in primate sensory systems. *Nature Reviews Neuroscience*, **6**, 97-107.
- PETROV, A. A. (2009). *Single-interval and same-different discrimination tasks yield consistent estimates of perceptual learning of visual motion*. Manuscript submitted for publication.
- PLATT, J. R. (1964). Strong inference: Certain systematic methods of scientific thinking may produce much more rapid progress than others. *Science*, **146**, 347-353.
- POLLACK, I., & PISONI, D. B. (1971). On the comparison between identification and discrimination tests in speech recognition. *Psychonomic Science*, **1**, 125-126.
- POPPER, K. (1963). *Conjectures and refutations: The growth of scientific knowledge*. London: Routledge & Kegan Paul.
- ROMO, R., HERNÁNDEZ, A., ZAINOS, A., LEMUS, L., & BRODY, C. D. (2002). Neuronal correlates of decision-making in secondary somatosensory cortex. *Nature Neuroscience*, **5**, 1217-1225.
- RUMELHART, D. E., HINTON, G. E., & WILLIAMS, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Exploration in the microstructure of cognition* (Vol. 1, pp. 318-362). Cambridge, MA: MIT Press.
- SORKIN, R. D. (1962). Extension of the theory of signal detectability to matching paradigms in psychoacoustics. *Journal of the Acoustical Society of America*, **34**, 1745-1751.
- THOMAS, R. D. (1996). Separability and independence of dimensions within the same-different judgment task. *Journal of Mathematical Psychology*, **40**, 318-341.
- VOGELS, R., & ORBAN, G. A. (1986). Decision processes in visual discrimination of line orientation. *Journal of Experimental Psychology: Human Perception & Performance*, **12**, 115-132.
- WANG, M.-F., IRWIN, R. J., & HAUTUS, M. J. (2005). Detection-theoretic analysis of same-different judgments for the amplitude discrimination of acoustic sinusoids. *Journal of the Acoustical Society of America*, **117**, 1305-1313.

NOTES

1. Many people find it counterintuitive that the optimal strategy leads to lower d' estimates. Keep in mind that d' is defined on the basis of a yes-no task with the same stimuli. The optimal strategy maximizes the external same-different performance for given internal distributions. The d' estimation procedure works in reverse—from given data back to the perceptual signal-to-noise ratio.
2. Technically, LR2 is viable when $P(\text{"same"}) < .5$, which is rare.
3. χ^2 depends on the normal approximation to the binomial distribution, which is justified for the block sizes in most experiments. There is also an exact test based on the hypergeometric distribution (Collett, 2003).
4. Incidentally, *inequalities* are not necessarily preserved. For instance, a mixture of trials with $p_{AA} < p_{BB}$ and trials with $p_{AA} > p_{BB}$ can produce $p_{AA} \approx p_{BB}$ at the block level. Criterion variability can thus make an asymmetric decision rule appear symmetric.

APPENDIX A
Definitions and Proofs

In this Appendix, we build the mathematical infrastructure necessary to prove Theorem 1 stated in the main text.

Definition. A region D is symmetric with respect to an axis A when the mirror image $s(\mathbf{x})$ of each point \mathbf{x} in the region also lies in the region:

$$\mathbf{x} \in D \iff s(\mathbf{x}) \in D. \quad (\text{A1})$$

Lemma 1. Let the piecewise curve C partition the plane in two or more regions that can be painted in black and white so that regions bordering on opposite sides of the curve are always painted in opposite colors. Let the boundary C be symmetric with respect to an axis A and have no more than countably many points in common with it. Then the union of all black regions is symmetric with respect to A and so is the union of all white regions.

APPENDIX A (Continued)

Proof outline. An orientation can be assigned to the curve C so that the black partition is always on the left side. To prove that the colored partition is symmetric, it is sufficient to prove that the regions are in a consistent relation relative to the oriented boundary after mirror reflection (Kostrikin & Manin, 1989). Axial symmetry reverses the orientation of the plane, because the determinant of the transformation is $\det(s) = -1$ (Kostrikin & Manin, 1989). We will prove that it also reverses the orientation of the curve C . The black partition now lies on the right side of the reflected curve, thus preserving the original color.

The proof proceeds by induction on the continuous segments of C . Let c be an arbitrary such segment. There are two possible cases illustrated in Figures A1A and A1B.

Case 1. c is symmetric with respect to A . Given that it cannot overlap with A (by assumption), the orientation of c is reversed after reflection (Figure A1A).

Case 2. The reflection of c is some other segment, d . Then the reflection of d is c . Although the orientation of each segment is preserved after reflection, the orientation of the pair $c \cup d$ is reversed (Figure A1B).

The technical condition that C has no more than countably many points in common with A rules out the possibility that c coincides with A . In this problematic case, the partition is still symmetric, but the colors are exchanged after reflection.




Figure A1. Solid lines with bold arrows show the original curve and its assigned orientation. Dashed lines with arrows show the reflected curve. Although the orientation is systematically reversed, the color of each region is preserved after reflection.

Definition. Two probability density functions f_1 and f_2 are mutually symmetric with respect to an axis A when for each point \mathbf{x} :

$$f_1(s(\mathbf{x})) = f_2(\mathbf{x}). \tag{A2}$$

Definition. A probability density function f is symmetric with respect to A when f is mutually symmetric with itself.

Lemma 2. Let the probability density functions f_1 and f_2 be mutually symmetric with respect to an axis A . Let C be a curve that satisfies the conditions of Lemma 1 with respect to the same axis. Let S_1 and S_2 be two stimulus pairs with perceptual effects distributed according to f_1 and f_2 , respectively. Then

$$P(\text{“B”} | S_1) = P(\text{“B”} | S_2), \tag{A3}$$

where $P(\text{“B”} | S_i)$ denotes the probability of covertly classifying S_i into the black region demarcated by C .

Proof. Consider the indicator function $I(\mathbf{x})$ that is 1 when \mathbf{x} lies in the black region and 0 otherwise. By Lemma 1, $I(\mathbf{x}) = I(s(\mathbf{x}))$, because the reflection s preserves the color. Then

$$\begin{aligned} P(\text{“B”} | S_1) &= \int I(\mathbf{x})f_1(\mathbf{x})d\mathbf{x} \\ &= \int I(s(\mathbf{x}))f_1(s(\mathbf{x}))ds(\mathbf{x}) \\ &= \int I(\mathbf{x})f_2(\mathbf{x})d\mathbf{x} = P(\text{“B”} | S_2). \end{aligned} \tag{A4}$$

This uses the mutual symmetry $f_1(s(\mathbf{x})) = f_2(\mathbf{x})$. Also, $ds(\mathbf{x}) = d\mathbf{x}$ because $|\det(s)| = 1$ (Kostrikin & Manin, 1989).

APPENDIX A (Continued)

Lemma 3. Assume perceptual independence, no memory noise, bivariate Gaussian distributions, and equal variance. Let stimulus pairs AA, AB, BA, and BB give rise to perceptual effects distributed according to f_{AA} , f_{AB} , f_{BA} , and f_{BB} , respectively. Then

- f_{AB} and f_{BA} are mutually symmetric w.r.t. the positive diagonal
- f_{AA} is symmetric w.r.t. the positive diagonal
- f_{BB} is symmetric w.r.t. the positive diagonal
- f_{AA} and f_{BB} are mutually symmetric w.r.t. the negative diagonal
- f_{AB} is symmetric w.r.t. the negative diagonal
- f_{BA} is symmetric w.r.t. the negative diagonal

Proof. Without loss of generality, we can choose the coordinate system illustrated in Figure 1A. The mean perceptual effects of stimuli A and B are $\mu_A = -d'/2$ and $\mu_B = +d'/2$. The common variance is $\sigma^2 = 1$. Then the bivariate Gaussian densities are:

$$f_{XY}(x_1, x_2) = \frac{1}{2\pi} e^{-\frac{1}{2}K_{XY}(x_1, x_2)}, \text{ where}$$

$$K_{AA}(x_1, x_2) = (x_1 + d'/2)^2 + (x_2 + d'/2)^2,$$

$$K_{AB}(x_1, x_2) = (x_1 + d'/2)^2 + (x_2 - d'/2)^2,$$

$$K_{BA}(x_1, x_2) = (x_1 - d'/2)^2 + (x_2 + d'/2)^2,$$

and

$$K_{BB}(x_1, x_2) = (x_1 - d'/2)^2 + (x_2 - d'/2)^2. \quad (A5)$$

The positive diagonal defines the mirror reflection $s^+(x_1, x_2) = (x_2, x_1)$, whereas the negative diagonal defines $s^-(x_1, x_2) = (-x_2, -x_1)$. The conclusions of the lemma are verified directly.

Lemma 4. Under standard representational assumptions:

- Any strategy whose decision boundary satisfies the condition of Lemma 1 with respect to the positive diagonal implies $p_{AB} = p_{BA}$. That is, the probability to respond “different” is the same for stimulus pairs AB and BA.
- Any strategy whose boundary satisfies the condition with respect to the negative diagonal implies $p_{AA} = p_{BB}$.

Proof. Follows from Lemmas 2 and 3.

The decision boundaries of all strategies in Table 1 have no more than countably many points in common with either diagonal. Thus, the condition of Lemma 1 reduces to axial symmetry.

Proof of Theorem 1. In light of Lemma 4, all that remains to be proven is the axial symmetry of the decision rules in Table 1:

- *Differencing strategy with two criteria (DF2).* The decision boundary (cf. Figure 1C) consists of two disjoint segments with equations $x_2 - x_1 = c_i$, where c_1 and c_2 are the two criteria. As $s^-(x_2 - x_1) = ((-x_1) - (-x_2)) = (x_2 - x_1)$, each segment is symmetric with respect to the negative diagonal. By Lemma 4, $p_{AA} = p_{BB}$.

- *Differencing strategy with one criterion (DF1).* This is a special case of DF2, in which $c_2 = -c_1$ (cf. Figure 1B). As $s^+(x_2 - x_1) = (x_1 - x_2) = -(x_2 - x_1)$, the two segments are mutually symmetric with respect to the positive diagonal. By Lemma 4, $p_{AB} = p_{BA}$.

- *Likelihood-ratio strategy with one criterion (LRI).* It can be verified by direct substitution that the likelihood ratio $L_{d/s}$ in Equation 3 is preserved by both s^+ and s^- :

$$L_{d/s}(x_1, x_2) = L_{d/s}(x_2, x_1) = L_{d/s}(-x_2, -x_1). \quad (A6)$$

Thus, any decision boundary $L_{d/s} = c$ is symmetric with respect to both diagonals (Figure 1G).

- *Likelihood-ratio strategy with two criteria (LR2).* This strategy applies different criteria for points on different sides of the positive diagonal (Figure 1H). The intent is to break this symmetry and allow for $p_{AB} \neq p_{BA}$. The other symmetry still holds, and $p_{AA} = p_{BB}$.

- *Covert-classification strategy with one (CC1) or two (CC2a) asymmetric criteria.* The covert classification of each individual stimulus is the same, regardless of whether the stimulus was presented first or second. Thus, the classification does not change when the presentation order is reversed, which is equivalent to mirror reflection with respect to the positive diagonal (Figures 1D and 1F). The proof that $p_{AB}^2 \geq p_{AA}p_{BB}$ is given in Appendix B.

- *Reversed-classification strategy with two asymmetric criteria (RC2a).* This decision rule is constructed so that the tips of the L-shaped decision boundaries lie on the negative diagonal (Figure 1I). This introduces a symmetry with respect to this diagonal, and $p_{AB} = p_{BA}$.

- *Covert-classification strategy with two symmetric criteria (CC2s = RC2s).* As a special case of both CC2a and RC2a, this decision boundary is symmetric with respect to both diagonals (Figure 1E).

APPENDIX B

Algorithms for Computing d' for the Covert-Classification Rules**Covert-Classification Rule With One Criterion**

The CC1 decision rule uses a single internal criterion for covert classification (Figure 1D). Let a_A denote the covert probability for classifying stimulus A as an "A." Analogously, $b_A = P(\text{"B"} | A)$, $b_B = P(\text{"B"} | B)$, and $a_B = P(\text{"A"} | B)$. Then the overt probabilities to respond "different" are:

$$\begin{aligned} p_{AA} &= a_A b_A + b_A a_A \\ p_{AB} &= a_A b_B + b_A a_B \\ p_{BA} &= a_B b_A + b_B a_A \\ p_{BB} &= a_B b_B + b_B a_B \end{aligned} \quad (\text{B1})$$

Note that $p_{AB} = p_{BA}$. Moreover, $a_X + b_X = 1$, because there is only one criterion. The second constraint in Theorem 1 follows by direct substitution:

$$(p_{AB} - .5)^2 = (p_{AA} - .5)(p_{BB} - .5). \quad (\text{B2})$$

The two covert probabilities are uniquely determined: $a_A = (1 + \sqrt{2p_{AA} - 1})/2$, $b_B = (1 + \sqrt{2p_{BB} - 1})/2$ (Pollack & Pisoni, 1971). Thus, the d' can be calculated by the standard formula for single-interval designs:

$$d' = z(a_A) - z(a_B) = z(b_B) - z(b_A). \quad (\text{B3})$$

Finally, assuming balanced presentation frequencies, there is an overall bias $P(\text{"same"}) = .5 + (a_A - b_B)^2/2 \geq .5$.

Covert-Classification Rule With Two Criteria

The CC2a decision rule is a generalization of CC1 that uses three covert categories: "A," "B," and "ambiguous." This requires two criteria $c_1 \leq c_2$ (Figure 1F). The covert probabilities are given by the Gaussian cumulative density function Φ :

$$\begin{aligned} a_A &= \Phi(c_1 + d'/2) \\ b_A &= 1 - \Phi(c_2 + d'/2) \\ a_B &= \Phi(c_1 - d'/2) \\ b_B &= 1 - \Phi(c_2 - d'/2) \end{aligned} \quad (\text{B4})$$

The observer responds "different" iff one stimulus is unambiguously classified "A" and the other "B." Thus, Equation B1 still holds. However, a_X and b_X no longer add to 1, and Equation B2 relaxes to:

$$p_{AB}^2 \geq p_{AA} p_{BB}. \quad (\text{B5})$$

Proof. $D = p_{AB}^2 - p_{AA} p_{BB} = (a_A b_B - a_B b_A)^2 \geq 0$.

To reduce the problem to one dimension, the covert probabilities in Equation B1 can be expressed as functions of a single unknown x :

$$\begin{aligned} a_A(x) &= x(p_{AB} + \sqrt{D})/p_{BB} \\ a_B(x) &= x \\ b_A(x) &= (p_{AB} - \sqrt{D})/(2x) \\ b_B(x) &= p_{BB}/(2x) \end{aligned} \quad (\text{B6})$$

From $a_B + b_B \leq 1$ and $a_A + b_A \leq 1$, we get, respectively:

$$\frac{1}{2}(1 - \sqrt{1 - 2p_{BB}}) \leq x \leq \frac{1}{2}(1 + \sqrt{1 - 2p_{BB}})$$

and

$$\frac{p_{BB}(1 - \sqrt{1 - 2p_{AA}})}{2(p_{AB} + \sqrt{D})} \leq x \leq \frac{p_{BB}(1 + \sqrt{1 - 2p_{AA}})}{2(p_{AB} + \sqrt{D})}. \quad (\text{B7})$$

Finally, Equation B3 leads to Equation B8, which can be solved numerically within the bounds in Equation B7.

$$z(a_A(x)) - z(a_B(x)) + z(b_A(x)) - z(b_B(x)) = 0. \quad (\text{B8})$$

Once x is determined, the d' can be calculated from Equations B6 and B3 and the two criteria from Equation B9.

$$\begin{aligned} c_1 &= (z(a_A) + z(a_B))/2 \\ c_2 &= -(z(b_A) + z(b_B))/2 \end{aligned} \quad (\text{B9})$$

Reversed-Classification Rule With Two Criteria

The RC2a decision rule is sensitive to the presentation order. It uses covert criteria $c_1 \leq c_2$ for the first and $-c_2 \leq -c_1$ for the second stimulus (Figure 11). The covert probabilities are analogous to Equation B4 with subscripts indicating the presentation interval: $a_{A1} = P("A" | S_1 = A)$, etc. The symmetry $\Phi(u) = 1 - \Phi(-u)$ implies that $b_{B2} = a_{A1} = a_A$, $a_{B2} = b_{A1} = b_A$, $b_{A2} = a_{B1} = a_B$, and $a_{A2} = b_{B1} = b_B$. Equation B1 thus becomes

$$\begin{aligned}
 p_{AA} &= a_{A1}b_{A2} + b_{A1}a_{A2} = a_Aa_B + b_Ab_B \\
 p_{AB} &= a_{A1}b_{B2} + b_{A1}a_{B2} = a_A^2 + b_A^2 \\
 p_{BA} &= a_{B1}b_{A2} + b_{B1}a_{A2} = a_B^2 + b_B^2 \\
 p_{BB} &= a_{B1}b_{B2} + b_{B1}a_{B2} = a_Ba_A + b_Bb_A.
 \end{aligned}
 \tag{B10}$$

Note that $p_{AA} = p_{BB}$, in agreement with Theorem 1. Equation B5 becomes

$$p_{AB}p_{BA} \geq p_{AA}^2.
 \tag{B11}$$

As before, we express the covert-classification probabilities in terms of a single unknown x :

$$\begin{aligned}
 a_A(x) &= \frac{1}{p_{BB}} \left(xp_{AA} + b_B(x) \sqrt{p_{AB}p_{BA} - p_{AA}^2} \right) \\
 a_B(x) &= x \\
 b_A(x) &= \sqrt{p_{AB} - a_A^2(x)} \\
 b_B(x) &= \sqrt{p_{BA} - x^2}.
 \end{aligned}
 \tag{B12}$$

Inequalities such as $a_A + b_A \leq 1$ and $a_Aa_B \leq p_{AA}$ impose upper bounds on x :

$$\begin{aligned}
 x &\leq \sqrt{p_{BA}} \\
 x &\leq p_{AA} / \sqrt{p_{AB}} \\
 x &\leq \left(1 - \sqrt{2p_{BA} - 1} \right) / 2, & \text{if } p_{BA} > 1/2 \\
 x &\leq \sqrt{\frac{p_{AB}p_{BA} - p_{AA}^2}{p_{AB} + p_{BA} - 2p_{AA}}}, & \text{if } p_{BA} > p_{AA}
 \end{aligned}
 \tag{B13}$$

It is easy to solve Equation B8 within these bounds using off-the-shelf numerical methods. We use the *fzero* function in MATLAB (MathWorks, 1999). The d' can then be calculated from Equations B12 and B3 and the two criteria from Equation B9. MATLAB scripts implementing these calculations are available at <http://alexpetrov.com>.