Psychological Review

The Relative Psychometric Function: A General Analysis Framework for Relating Psychological Processes

Brian Maniscalco, Olenka Graham Castaneda, Brian Odegaard, Jorge Morales, Sivananda Rajananda, Rachel N. Denison, and Megan A. K. Peters

Online First Publication, June 15, 2026. https://dx.doi.org/10.1037/rev0000625

AMERICAN PSYCHOLOGICAL ASSOCIATION

Psychological Review

ISSN: 0033-295X

https://doi.org/10.1037/rev0000625

The Relative Psychometric Function: A General Analysis Framework for Relating Psychological Processes

Brian Maniscalco[1, 2], Olenka Graham Castaneda[1, 2], Brian Odegaard[3], Jorge Morales[4, 5], Sivananda Rajananda[2], Rachel N. Denison[6], and Megan A. K. Peters[1, 2, 7, 8, 9, 10]

[1] Department of Cognitive Sciences, University of California, Irvine
[2] Department of Bioengineering, University of California, Riverside
[3] Department of Psychology, University of Florida
[4] Department of Psychology, Northeastern University
[5] Department of Philosophy and Religion, Northeastern University
[6] Department of Psychological and Brain Sciences, Boston University
[7] Department of Logic and Philosophy of Science, University of California, Irvine
[8] Center for the Theoretical Behavioral Sciences, University of California, Irvine
[9] Center for the Neurobiology of Learning and Memory, University of California, Irvine
[10] Program in Brain, Mind, and Consciousness, Canadian Institute for Advanced Research

Psychophysics seeks to quantitatively characterize relationships between objective properties of the world and subjective properties of perception. However, traditional approaches fall short of identifying quantitative relationships *among* psychological processes. This latter goal is particularly important when certain processes may depend on each other, such as for subjective experience and task performance: Typically, better performance is accompanied by stronger subjective experiences of clarity, vividness, or confidence. But is the relationship between performance and subjective experience fixed, or can it vary, for example, by task or attentional demands? Such questions are key for better understanding psychological processes in general and subjective experience in particular. Here, we develop and showcase a new psychophysical method designed to answer such questions: *relative psychometric function* (RPF) analysis, which characterizes the nonlinear psychometric relationships between psychological processes and how these relationships change under different circumstances (e.g., experimental manipulations). We demonstrate the advantages of RPF analysis using a data set in which human subjects discriminated random dot kinematograms which varied in motion coherence and overall dot density (dots per visual degree) and rated confidence. RPF analysis revealed systematic changes in the relationship between performance and two subjective measures (confidence and metacognitive sensitivity) due to dot density and task design choices. These empirical results showcase how RPF analysis may reveal changes in quantitative relationships between any two psychological measures in future studies: performance, vividness, clarity, reaction time, confidence, and more. To enable the scientific community to use RPF analysis on their data, we also present our open-source RPF toolbox.

Keywords: psychophysics, psychometric functions, relative psychometric function, subjective experience, quantitative psychology

Supplemental materials: https://doi.org/10.1037/rev0000625.supp

Han L. J. Van der Maas served as action editor.
Megan A. K. Peters https://orcid.org/0000-0002-0248-0816
Analyses presented in this article were completed with the open-source relative psychometric function toolbox (Maniscalco & Peters, 2026; https://github.com/CNClaboratory/RPF). The data and analysis scripts from the empirical case study can be found at https://osf.io/62rwk/ (Castaneda et al., 2026). The ideas, analyses, and data presented in this article were shared in part through poster presentations and talks at various international conferences, including the Association for the Scientific Study of Consciousness and the Vision Sciences Society. Previous versions of this article were also posted as a preprint: Maniscalco et al. (2020; https://doi.org/10.31234/osf.io/5qrjn). Some previous versions possessed a different title: The metaperceptual function: Exploring dissociations between confidence and task performance with Type 2 psychometric curves. Please see the version history for all previous versions.

The authors also acknowledge support from the Templeton World Charity Foundation ("An Adversarial Collaboration to Test Predictions of First-Order and Higher-Order Theories of Consciousness," empleton World Charity Foundation Number: 0567, to Megan A. K. Peters and Rachel N. Denison) and the Canadian Institute for Advanced Research (Fellowship in the Brain, Mind, & Consciousness Program, to Megan A. K. Peters). The funders had no role in the design or execution of this project. They thank Karen Tian, Michael Epstein, Angela Shen, and Emil Olsson for helpful discussions in the development of this methodology. Because this research was funded in whole, or in part, by Templeton World Charity Foundation (Grant 0567), for the purpose of open access, the author has applied a CC BY public copyright license to any author accepted manuscript version arising from this submission.

Brian Maniscalco played a lead role in conceptualization, data curation, formal analysis, investigation, methodology, software, validation, visualization, writing–original draft, and writing–review and editing, a supporting

continued

Arguably, the field of quantitative psychology began in the 1860s with Fechner's *Elemente der Psychophysik—Elements of Psychophysics* (Fechner, 1860; Fechner et al., 1966). Fechner's work set the foundation for what is now over 150 years of concerted effort to map objective properties of the world to properties of the mind and brain. Weber, Stevens, and others followed, seeking to establish the functional forms of these relationships: That an observer's "just noticeable difference" in discriminating two stimuli depends on their absolute magnitude (Weber's law), and that the perceived magnitude of a stimulus (brightness, loudness, painfulness) exhibits an exponential relationship to the objective stimulus magnitude (Stevens' power law). These "introductory psychology course" concepts are foundational pillars in the modern study of psychology.

The success of this framework underscores our deep motivation to build models of our minds, but standard psychophysical approaches represent one-to-one mappings between the physical and mental. Ultimately, we wish to understand not only how psychological processes relate to properties of the world but also how psychological processes relate to *each other*. For example, increasing stimulus strength typically leads to faster, more accurate decisions, and increased sense of confidence in those decisions. Likewise, the subjective sense of clarity may also systematically vary with stimulus properties. But what is the relationship among all these psychological variables, and is it fixed across different attentional states, tasks, or individuals? While the relationship linking stimulus magnitude, discriminability, and absolute magnitude estimation has recently been described (Zhou et al., 2024), what about the relationships linking all these other psychological properties to each other?

Characterizing quantitative relationships between psychological variables is also especially important when those relationships themselves may change depending on properties of the world or other psychological processes. Perhaps nowhere is this more evident than in psychophysical studies of subjective experience, where there are clear, empirically observed relationships among stimulus intensity, task performance, and second-order judgments such as confidence (judgment of whether a given discrimination decision is likely to be correct) or subjective visibility, clarity, or vividness (e.g., judgment of the clarity with which one saw a stimulus, regardless of its objective properties). In most instances, these subjective aspects covary with objective performance (Baranski & Petrusic, 1994): A higher probability of correctly identifying a stimulus is typically accompanied by higher confidence, higher clarity ratings, and a higher probability of reporting having seen the stimulus at all. This means that any neural or psychophysical measures of subjective experience are easily confounded by processes driving objective performance (Morales et al., 2022).

A standard approach to disentangling the neural correlates of subjective experience from those underlying objective performance has been to control for these "performance confounds" through either experimental or analytic approaches (Lau, 2008; Lau & Passingham, 2006; Morales et al., 2022; Peters, Kentridge, et al., 2017). One way of implementing this idea is to create multiple experimental conditions in which performance (e.g., percent correct responses or the signal detection theoretic metric d') is held constant (e.g., through subject-specific staircasing) but subjective reports (confidence, vividness, visibility, clarity) vary. However, this performance matching approach has some drawbacks. First, statistical demonstrations that performance is constant across conditions might reflect failures to detect subtle but extant differences in performance, for example, failure to reject a false null hypothesis. This situation might arise simply from situations in which a chosen manipulation had a smaller or noisier effect on performance than it did on subjective reports rather than no effect at all. Second, the effect of an experimental manipulation on a subjective measure at a given matched performance level may strongly depend on the absolute level of performance (see also Morales et al., 2022, for further discussion). Such condition-driven differences in confidence at matched performance have been observed in many different paradigms (Koizumi et al., 2015; Lau & Passingham, 2006; Maniscalco et al., 2016; Odegaard, Chang, et al., 2018; Odegaard, Grimaldi, et al., 2018; Rahnev et al., 2011; Rouault et al., 2018; Samaha et al., 2016; Stolyarova et al., 2019), but the effect size (and sometimes even direction!) can vary across stimulus or task manipulations as well as the (matched) performance level itself.

Notably, both shortcomings of the performance matching approach could be addressed by estimating the full psychometric relationship between task performance and subjective visibility, from chance to ceiling levels of performance. Comparison of such fitted curves across experimental conditions would allow assessment of differences in visibility at matched levels of performance without relying on statistical null effects, while simultaneously revealing how such differences change with performance level.

To date, however, no framework exists for precisely characterizing such nonlinear relationships among psychological variables and how they may change in interesting ways. An impediment to this enterprise has been that these relationships are, by definition, nonlinear linkages between variables measured with error. Fitting any function linking two such variables constitutes a nonlinear "errors-in-variables" problem, where both variables are measured with error and thus standard methods for maximum likelihood estimation or ordinary least squares optimization are not applicable. Whereas methods exist for fitting linear errors-in-variables models (e.g., Deming regression), no method exists for fitting nonlinear

role in project administration, and an equal role in conceptualization. Olenka Graham Castaneda played a supporting role in data curation and writing–review and editing. Brian Odegaard played a supporting role in conceptualization, visualization, and writing–review and editing. Jorge Morales played a supporting role in conceptualization, visualization, and writing–review and editing. Sivananda Rajananda played a supporting role in data curation and writing–review and editing. Rachel N. Denison played a supporting role in conceptualization, funding acquisition, visualization, and writing–review and editing. Megan A. K. Peters played a lead role in conceptualization, formal analysis, funding acquisition, investigation, methodology, project administration, resources, supervision, writing–original draft, and writing–review and editing, a supporting role in software, and an equal role in data curation, validation, and visualization.

Correspondence concerning this article should be addressed to Brian Maniscalco, Department of Cognitive Sciences, University of California, Irvine, 2201 Social and Behavioral Sciences Gateway, Irvine, CA 92697-5100, United States, or Megan A. K. Peters, Department of Logic and Philosophy of Science, University of California, Irvine, 2201 Social and Behavioral Sciences Gateway, Irvine, CA 92697-5100, United States. Email: bmaniscalco@gmail.com or megan.peters@uci.edu

models without using additional information, for example, instrumental variables or repeated observations (see, e.g., Hausman et al., 1995; Z. Huang et al., 2023; T. Li, 2002; Wolter & Fuller, 1982). Moreover, even if this problem were solved, the functional form linking two or more psychological variables is unlikely to be known a priori, requiring us to fall back on nonparametric methods (e.g., rank-based correlations) designed merely to reveal the *presence* and *strength* of a potential relationship, not its shape.

In short, we need a framework to (a) quantitatively characterize the nonlinear relationships among psychological variables measured with error; and (b) quantitatively characterize how much—and in what way—those relationships change with experimental manipulations, neural factors, or individual differences.

Here, we introduce an analytic framework to address these problems. The approach, which we term *relative psychometric function* (RPF) analysis, aims to systematically characterize the relationship between any two psychometric variables and assess how that relationship changes with experimental manipulations, individual differences, and so on. The framework is sufficiently general to be applied to investigation of the relationship between any psychological or neural processes P_1 and P_2 which can be expressed as psychometric functions of a common continuous variable such as stimulus intensity. Investigation of RPFs can significantly expand and deepen our understanding of how various psychological and neural processes interrelate at both phenomenological and mechanistic levels of analysis.

One particularly salient application of the RPF framework is to the study of perceptual metacognition. By investigating subjective measures of perception (e.g., confidence or visibility ratings) as a function of task performance over a wide range of performance levels, we can broaden and deepen our understanding of both in a number of ways. Of particular note is that RPF analysis can be used to control for performance confounds and better isolate changes pertaining specifically to subjective measures, thereby sharpening inferences about the computational and neural mechanisms underlying metacognition and subjective experience. Importantly, as foreshadowed above, RPF analysis controls for performance confounds in a more comprehensive and robust way than existing performance matching approaches. The RPF method as applied to perceptual metacognition also answers recent calls for a “metacognitive” or “introspective” psychophysics (Fleming, 2023; Peters, 2025) and builds upon the “metacognition as a step toward explaining phenomenology” approach introduced by Peters (2022), which calls for research to seek canonical metacognitive computations as a strategy for revealing how subjective experience in general may be generated.

In what follows, we introduce the general principles of relative psychometric function analysis, investigate the mathematical form and behavior of an example RPF, develop and validate interpretable summary statistics, consider model comparison approaches, and discuss its interpretation using a sample data set in which performance and confidence were independently manipulated across a large range of stimulus strengths.

We believe this framework will prove a highly flexible and powerful analysis tool in psychology and neuroscience for studying relationships between various psychological and neural processes. To facilitate this goal, all the methods, data, and analyses presented here are also used to introduce the RPF toolbox (https://github.com/CNClaboratory/RPF, Maniscalco & Peters, 2026)—an open-source resource for the community to apply RPF analyses to any suitable data set. Thus, we make reference to this toolbox throughout and include additional details about implementation on the case study data set presented here in the Supplemental Material.

Methods, Results, and Discussion

Deriving and Interpreting the RPF

Foundations of RPF Analysis

The General Form of the RPF. We define the *relative psychometric function*, or RPF for short, as the function describing the relationship between any two conventional psychometric functions that are expressed in terms of a common independent variable. More formally, suppose we have two psychometric functions

$$P_1 = F_1(x; \theta_1),$$
$$P_2 = F_2(x; \theta_2), \tag{1}$$

where x is stimulus strength and P_1 and P_2 are different measures of performance, such as p(correct) or average confidence.[1,2] To investigate the relationship between these variables, we wish to express P_2 as a function of P_1. Provided that F_1 is invertible such that $x = F_1{}^{-1}(P_1)$, we may do so by writing $P_2 = F_2(F_1{}^{-1}(P_1))$. We may then define the relative psychometric function as

$$R = F_2 \circ F_1^{-1}, \tag{2}$$

such that

$$P_2 = F_2(F_1^{-1}(P_1)) = R(P_1; \theta_1, \theta_2), \tag{3}$$

and write $P_2 = R(P_1)$ for short. Thus, R uses the known psychometric functions F_1 and F_2 of a common independent variable x in order to express P_2 as a function of P_1 (Figure 1).

Since R relates F_2 to F_1 by means of their common independent variable x, it can be seen as a parametric curve whose parameter x is mapped onto the (P_1, P_2) plane via parametric equations (Equation 1).[3] The general form of the RPF given in Equation 3 is the result of implicitizing these parametric equations—that is, eliminating the parameter x to yield a single implicit equation in terms of P_1 and P_2—when F_1 is invertible. Although it is possible to implicitize these equations even if F_1 is not invertible, the resulting equation will not express P_2 as a function of P_1, that is, it will not be possible to construct a relative psychometric *function* where each P_1 value maps onto a unique P_2 value. Since interpretation of such cases would pose conceptual complications, and since such cases are unlikely to affect the vast majority of relevant applications in psychology and neuroscience, in this work, we focus on cases where it is possible to construct an RPF as in Equation 3. As a consequence, our conceptual and formal approaches emphasize analysis of the implicitized RPF (Equation 3) rather than a more general framing in terms of parametric curve analysis, although there is an underlying flavor of the latter insofar as many operations involved in fitting and analyzing

[1] Note that P is intended as shorthand for “performance” and does not necessarily connote a probability.

[2] Here, we use F rather than the conventional ψ to denote psychometric functions for consistency with the RPF toolbox notation, in which it is more convenient to use F.

[3] We thank reviewer Jürgen Heller for pointing out this connection.

Figure 1
The Relative Psychometric Function

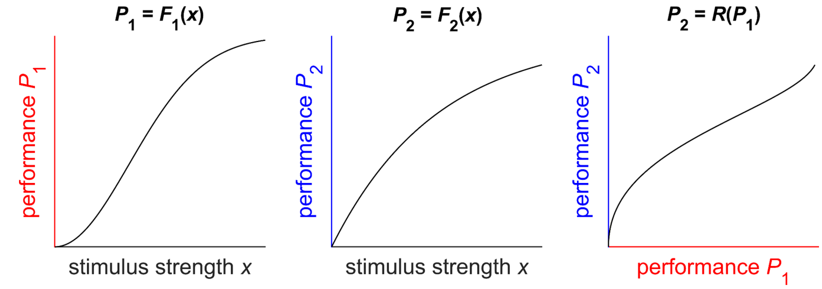

Note. (Left, middle) Conventional psychometric functions characterize the curve relating performance on a given task to stimulus strength. (Right) The *relative psychometric function* characterizes the curve relating performance on one task (P_2) to performance on another (P_1), given knowledge of how both relate to a common stimulus feature *x*. This function thus reveals how different measures of psychological processes relate to each other over a wide range of performance levels. In situations where RPFs differ across experimental conditions, analysis of the RPF can be used to help tease apart the behavior and underlying mechanisms of P_1 and P_2. In the special case where P_1 and P_2 correspond to objective and subjective aspects of perception (see main text), the RPF is a *metaperceptual function* which can be used to isolate subjective aspects of perception from potentially confounding aspects of objective task performance. RPF = relative psychometric function. See the online article for the color version of this figure.

the RPF can proceed via separate analysis of F_1 and F_2 rather than requiring an analytic expression for *R*. It is possible that future work that explores the connection of RPF analysis to parametric curve analysis more deeply may be fruitful. For further discussion, see Supplemental Material S1.

R can be seen as the result of a coordinate transformation of F_2 in which the *x* input is replaced with a P_1 input derived from the mapping $x = F_1^{-1}(P_1)$. Thus, the plot of $P_2 = R(P_1)$ resembles a warped plot of $P_2 = F_2(x)$ in which the *y*-axis values are identical but their distribution along the *x*-axis is warped according to the (likely nonlinear) transformation specified by $P_1 = F_1(x)$ (cf. the three panels of Figure 1 and Supplemental Figures S1 and S2). Over *x* intervals where $F_1(x)$ is shallow, *x* values map onto a small range of P_1 values. This effectively makes the corresponding *x* intervals of the $F_2(x)$ plot contract in their transformation to small P_1 intervals of the $R(P_1)$ plot. Conversely, over *x* intervals where $F_1(x)$ is steep, similar considerations make the corresponding *x* intervals of the $F_2(x)$ plot expand in their transformation to large P_1 intervals of the $R(P_1)$ plot.

Importantly, deriving $P_2 = R(P_1)$ via the relationship of P_1 and P_2 to a common independent variable *x*, where *x* is known exactly rather than measured with error, bypasses difficulties that would arise from attempting to fit a function directly to the (P_1, P_2) data. Since both P_1 and P_2 are variables measured with error and likely have a nonlinear relationship, attempting to fit $P_2 = R(P_1)$ directly requires a nonlinear errors-in-variables model. However, there is no known solution for fitting such models directly, and existing approaches require incorporation of additional data and application of complex analysis methods tailored to specific cases (see, e.g. Hausman et al., 1995; Z. Huang et al., 2023; T. Li, 2002; Wolter & Fuller, 1982).

In the case of analyzing the relationships between psychometric functions in experimental psychology research, the (P_1, P_2) data we might wish to relate are themselves already derived from systematic manipulation of the common independent variable *x*, and thus, the information needed to estimate the relationships of P_1 and P_2 to *x* comes "for free" in the collection of the (P_1, P_2) data. Thus, the approach described in this work is a natural choice for conducting RPF analysis that easily bypasses thorny analysis issues with readily available data.

The Metaperceptual RPF. As discussed in the introduction, one particular application of interest is the case where P_1 and P_2 correspond to objective and subjective measures of perception, respectively. Here, we conceive of objective measures of perception as pertaining to judgments about objective states of the world (e.g., detecting stimulus presence or discriminating stimulus features) and subjective measures as pertaining to judgments about one's own perceptual processing (e.g., assessing confidence in an objective judgment or reporting on the qualities of one's perceptual experience). This characterization of "objective" and "subjective" categories can be seen as a generalization of the classical distinction between Type 1 and Type 2 perceptual tasks, in which the Type 1 task is to classify a stimulus event, and the Type 2 task is to classify one's Type 1 judgment as correct or incorrect (Clarke et al., 1959; Galvin et al., 2003; Maniscalco et al., 2024).

Taking inspiration from the term "psychophysics," we call this special class of RPFs *metaperceptual RPFs* or *metaperceptual functions*. Just as the roots of the word "psychophysical" connote "relationship of perception (psycho-) to stimulus (physical)," so the roots of the word "metaperceptual" connote "relationship of judgments *about* perception (meta-) to perception (perceptual)." We may also use the term *Type 2 psychometric function* to refer to more restricted cases, where the RPF relates Type 2 judgments about Type 1 accuracy (typically confidence ratings) to Type 1 accuracy itself (e.g., as in *p*(correct) or *d*′).

Objective measures of perception include accuracy measures such as p(correct) and the signal detection theory (SDT) measure of sensitivity d' and response bias measures such as p(response) and the SDT measure of criterion c.[4] Subjective measures include ratings of confidence and reports of experiential qualities such as visibility, clarity, intensity, and so on. Subjective measures may also characterize the relationship between subjective and objective judgments, for example, by measuring how well confidence ratings track accuracy as in the SDT measure of metacognitive sensitivity meta-d' (Fleming, 2017; Maniscalco & Lau, 2012, 2014).

Considerations for Fitting the Component Psychometric Functions of the RPF. Psychometric functions can be fitted to probability measures such as p(correct) and p(high confidence) with standard maximum likelihood estimation (MLE) procedures (Kingdom & Prins, 2016). However, MLE fitting of psychometric functions to nonprobabilistic measures requires a different approach. Least square fits maximize likelihood when errors in the fit can be assumed to be normally distributed with constant variance (Burnham & Anderson, 2002), but this assumption may not always hold (e.g., as for d'; see Miller, 1996).

In Supplemental Material S2.1, we derive approaches to achieving MLE psychometric function fits to several variables of central interest for metaperceptual functions: d', meta-d', and mean rating (e.g., for confidence or visibility ratings). These approaches work by relating the variable in question to probabilities for single-trial outcomes and thus only require the standard MLE assumption that outcome probabilities are constant and independent across trials. For cases where specifying or fitting analytic psychometric functions is problematic, we also develop nonparametric RPF analysis methods (see below and Supplemental Material S2.2 for further discussion) and demonstrate that MLE and nonparametric approaches are comparable in their ability to retrieve certain characteristics of the true RPF (see Supplemental Material S2.3).

It is possible to take a radically modular approach to constructing the RPF from its component psychometric functions, in the sense that because the F_1 and F_2 fits can be treated independently, they can be applied to any variable and approached with any fitting method prior to being combined in an RPF. Thus, for example, if F_1 is fitted via MLE, this does not constrain the possibilities for fitting F_2 via MLE or least squares or nonparametric methods. In all cases, the approach described in Equation 3 is sufficiently general to conduct RPF analysis, with the proviso that F_1 must be invertible.

In keeping with the modularity and generality of this approach, in the present work, we are not primarily concerned with considering how to choose specific psychometric functions for different use cases, nor with investigating possible theoretical links or mathematical relationships between the functions chosen for F_1 and F_2 (as might occur, e.g., in the context of characterizing the relationship between curves for d' and p(high confidence) according to a SDT model). Where possible, we treat F_1 and F_2 in a general way without specifying an exact functional form. Where specification of a function is helpful or required, we use the Weibull function as a representative psychometric function for fitting probabilistic data (see following section) and a scaled Weibull function (see Supplemental Material S2.1) as a representative function for fitting nonprobabilistic data. The simplicity of this approach allows us to focus on the more general considerations and principles of constructing, fitting, analyzing, and interpreting RPFs. Naturally, applications of the RPF method to specific research questions may want to use alternative psychometric functions for conventional, practical, or theoretical reasons and may want to consider exploring theoretical links or formal constraints between the functions chosen for F_1 and F_2. The principles of RPF analysis discussed here apply regardless of such choices.

All of these approaches to RPF analysis—that is, modular construction of the RPF via some combination of MLE fitting for probabilistic variables and certain nonprobabilistic variables, least square fitting, and nonparametric analysis—can be readily implemented in the RPF toolbox (Supplemental Material S4; https://github.com/CNClaboratory/RPF, Maniscalco & Peters, 2026).

Probing RPF Behavior: A Case Study Using the Weibull RPF

How should we measure, summarize, and analyze the RPF? Can we summarize its behavior neatly with a small number of parameter values, similar to how conventional psychometric functions are typically analyzed in terms of location and slope parameters? Since the RPF depends on the mathematical forms of the two psychometric functions F_1 and F_2 from which it is composed, the answer to these questions requires specifying the equations for those functions. Here, as a representative example, we consider the behavior of the RPF when P_1 and P_2 are probabilities, for example, p(correct) and p(high confidence), fitted by Weibull functions F_1 and F_2 (Kingdom & Prins, 2016).

The Weibull function for F_1 and F_2 takes on the form

$$P_n = F_n(x) = \gamma_n + (1 - \lambda_n - \gamma_n)\left[1 - e^{-(x/\alpha_n)^{\beta_n}}\right], \tag{4}$$

in this equation,

- n denotes the psychometric function to which all terms pertain, with $n = 1$ and 2 corresponding to F_1 and F_2, respectively
- P_n is performance (here, outcome probability)
- x is stimulus strength
- γ_n is the chance level of responding for P_n
- λ_n is the lapse rate, such that asymptotic performance for P_n is $1 - \lambda_n$
- α_n is the location parameter for $F_n(x)$
- β_n is the slope parameter for $F_n(x)$

Solving Equation 4 for x in the context of F_1 gives

$$x = F_1^{-1}(P_1) = \alpha_1\left(\ln\left(\frac{1 - \lambda_1 - \gamma_1}{1 - \lambda_1 - P_1}\right)\right)^{\frac{1}{\beta_1}}. \tag{5}$$

[4] Sometimes measures like p(response) and c are considered to be *subjective* measures of response bias. Here, we consider them to be *objective* measures of perception insofar as these measures pertain to judgments *about* the world rather than judgments about one's own perceptual processing.

Substituting Equation 5 into the general equation for R in Equation 3 gives

$$P_2 = R_W(P_1) = \gamma_2 + (1 - \lambda_2 - \gamma_2)\left[1 - e^{-\left(\left(\frac{\alpha_2}{\alpha_1}\right)^{-\beta_2}\left(\ln\left(\frac{1-\lambda_1-\gamma_1}{1-\lambda_1-P_1}\right)\right)^{\frac{\beta_2}{\beta_1}}\right)}\right]. \quad (6)$$

We name Equation 6 the *Weibull RPF* (abbreviated R_W), as this is the mathematical form of the RPF in the case where both F_1 and F_2 are Weibulls. We can decompose the Weibull RPF into the following components:

- The F_2 guess rate γ_2 and lapse rate λ_2, which determine the minimum, chance level of performance and maximum, asymptotic level of performance for the R_W just as they do for F_2.
- The *performance ratio* $\frac{1-\lambda_1-\gamma_1}{1-\lambda_1-P_1}$, which characterizes performance P_1 relative to its possible range of values in $[\gamma_1, 1 - \lambda_1]$. When P_1 is at the chance value of γ_1, the performance ratio = 1, and R_W is at the chance level of performance for P_2, that is, γ_2. As P_1 approaches the asymptotic value of $1 - \lambda_1$, the performance ratio approaches infinity, and R_W approaches the asymptotic level of performance for P_2, that is, $1 - \lambda_2$.
- The *relative location* $\alpha_R = \frac{\alpha_2}{\alpha_1}$.
- The *relative slope* $\beta_R = \frac{\beta_2}{\beta_1}$.
- The F_2 slope β_2.

We explore how R_W depends on α_R, β_R, and β_2 in Figure 2. Without loss of generality, we set the scaling parameters, $\gamma_1 = 0.5$, $\gamma_2 = 0$, and $\lambda_1 = \lambda_2 = 0$.

Overall, the behavior of the Weibull RPF is considerably more complicated than the standard Weibull. First, we can observe that R_W does not have a direct analogue to the Weibull's slope parameter β but rather has a variable shape and degree of curvature depending on the relative slope β_R and relative location α_R. When $\beta_R > 1$, R_W is sigmoidal; when $\beta_R < 1$, R_W is inverse sigmoidal; and when $\beta_R = 1$, R_W is concave down, linear, or concave up depending on α_R. The complexities of how R_W's shape changes depending on combinations of parameter values obscures a straightforward and universally applicable interpretation of the function in terms of a slope parameter, and so it is not clear that characterizing R_W in terms of a slope parameter is the best way to insightfully summarize its behavior.

Similarly, R_W does not have a simple analogue to the Weibull's location parameter α. In the Weibull (Equation 4), α acts as a location parameter in the sense that the function takes on 63.2% of its maximal value above chance (i.e., 63.2% of the way between γ and $1 - \lambda$) when $x = \alpha$, since substituting this value of x into the formula entails $1 - e^{-\left(\frac{x}{\alpha}\right)^{\beta}} = 1 - e^{-1} = 0.632$. Thus, α tells us what value of x ("location") yields this threshold function value of 63.2% of the above-chance maximum. From Equation 6, we see that R_W achieves 63.2% of its above-chance maximum when $\frac{\alpha_1}{\alpha_2}\left(\ln\left(\frac{1-\lambda_1-\gamma_1}{1-\lambda_1-P_1}\right)\right)^{\frac{1}{\beta_1}} = 1$. Solving for P_1 in this equation yields R_W's equivalent of the Weibull's location parameter α, which can be expressed as

$$P_1 = \gamma_1 + (1 - \lambda_1 - \gamma_1)\left[1 - e^{-\left(\frac{\alpha_2}{\alpha_1}\right)^{\beta_1}}\right] = F_1(\alpha_2). \quad (7)$$

Thus, R_W takes on its threshold value at the value of P_1 given by $F_1(x)$ evaluated at the location parameter of F_2, that is, at $P_1 = F_1(\alpha_2)$. This result is intuitive in that the RPF derives from a transformation of the input variable of F_2 from x to P_1, while leaving F_2's output P_2 unchanged (Equation 3). Thus, since α_2 is the value of x at which F_2 achieves its threshold value of P_2, R_W must achieve its threshold value of P_2 at whatever value of P_1 that α_2 maps onto in the RPF transformation, which is just $F_1(\alpha_2)$.

Although the value of $F_1(\alpha_2)$ provides a measure of what P_1 value yields R_W's threshold value, its interpretation is more complex than that of α for the Weibull function. Intuitively, lower and higher values of α in the Weibull function roughly correspond to the curve "shifting" or "tilting" left or right on the x-axis.[5] By contrast, the value at which R_W achieves its threshold value is strongly influenced by its curvature, which in turn depends on multiple parameters from F_1 and F_2. For instance, in the lower left plot of Figure 2, the concave down and concave up curves achieve their threshold values at very low and high values of P_1, respectively, due primarily to their differences in curvature. This difference in threshold location cannot be attributed to a shift, tilt, or translation in an otherwise similar curve, as is the case for the Weibull, and thus $F_1(\alpha_2)$ cannot serve the same conceptual role as the Weibull's location parameter α.

Thus, it appears that while there are indeed aspects of the Weibull RPF's behavior that can be summarized with a small number of parameters—β_R and α_R control shape, and $F_1(\alpha_2)$ determines threshold—it is not clear that these parameters provide the same ease of interpretation and leverage for understanding the behavior of the RPF in terms of psychophysical performance as their counterparts α and β do for conventional psychometric functions. Furthermore, the exact mathematical formulation for such parameters depends on the psychometric functions used for F_1 and F_2, entailing that different choices for these functions may lead to different formulations for RPF summary parameters. These difficulties motivate the alternative approaches for comparing RPFs across conditions that we develop and discuss below.

Comparing RPFs Across Conditions

If using parameter values to summarize aspects of RPF behavior is not as straightforward and perhaps not as fruitful as it is for conventional psychometric functions, what alternatives are there for using RPFs to enrich our understanding of psychological processes?

One major goal of RPF analysis would be to investigate how the relationship between two target psychological processes changes across different conditions. For example, analysis of the metaperceptual RPF would be well-suited to address questions on how the relationship between objective and subjective aspects of perceptions are influenced by various factors, such as "Is the relationship between subjective judgments and task accuracy the same in central versus peripheral visual field locations?" (Odegaard,

[5] The slope of the Weibull function is actually controlled by both α and β, although when plotted against log x, α controls function translation and β controls slope (Kingdom & Prins, 2016).

Figure 2

Behavior of the Weibull Relative Psychometric Function (R_W) as a Function of Its Three Main Parameters

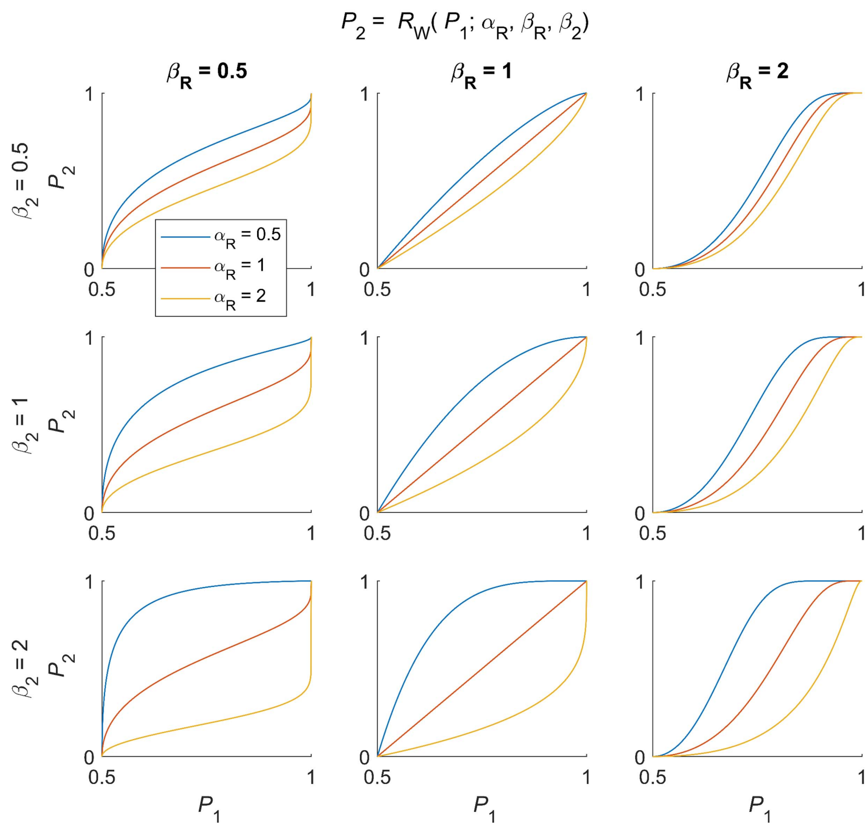

Note. Each plot shows how performance on one task (P_2) depends on performance on another task (P_1) according to the Weibull RPF specified by parameters for relative location $\alpha_R = \frac{\alpha_2}{\alpha_1}$ (separate lines within each plot), relative slope $\beta_R = \frac{\beta_2}{\beta_1}$ (columns), and F_2 slope β_2 (rows), as derived from the parameters of the component Weibull functions F_1 and F_2. RPF = relative psychometric function. See the online article for the color version of this figure.

Chang, et al., 2018; Winter & Peters, 2022) or "Does transcranial magnetic stimulation to a certain region of interest alter the relationship between confidence and task accuracy?" (Peters, Fesi, et al., 2017; Rahnev et al., 2012; Rounis et al., 2010; Ruby et al., 2018).

The behavior of the RPF across conditions sheds light on the relationship between P_1 and P_2. Any across-condition changes in the RPF would indicate a differential effect of condition on P_1 and P_2, such that the changes in P_2 due to condition could not be solely attributed to changes in P_1 or vice versa (or else the RPF would be identical), and would demonstrate that P_1 and P_2 are produced by at least partially separable (i.e., not identical) processes—for example, the processes producing task accuracy and confidence. Alternatively, if P_1 and P_2 differ across conditions, but do so in such a way that preserves the RPF describing their relationship, this would be consistent with the possibility that the changes in P_2 are indeed attributable entirely to changes in P_1 (or vice versa), or that both are products of a single underlying process characterized by a constant RPF.

Below, we consider two approaches to comparing RPFs across conditions: an area under the curve (AUC)-based approach and a model comparison approach.

AUC Approach: Area Under the RPF Curve. As discussed in the Introduction, the motivating example behind this work is performance matching in the consciousness/metacognition literature, in which we seek to find conditions where objective task performance (P_1) is the same but subjective reports of awareness or confidence (P_2) differ (Morales et al., 2022). Notice that the performance matching approach essentially attempts to compare a vertical slice of two RPFs at a particular P_1 value. Thus, a natural generalization of the performance matching approach is to compare two RPFs across a fixed *interval* of P_1 values rather than at a single fixed value. Within a given RPF, summing the values of P_2 across the entire interval of P_1 values amounts to computing the AUC of the RPF, and dividing this AUC by the length of the P_1 interval yields the average value of P_2 over that interval. These AUCs and average P_2 values can then be compared across conditions to assess whether condition affects P_2 over and above any effects it may have on P_1. (Note: AUCs and average P_2 values are suitable to evaluate whether there is any RPF difference between conditions but are not suitable if the user's goal is to characterize the exact shape of the RPF function; see the Benefits and Limitations of the AUC Method section.)

More formally, the RPF AUC is given by

$$\text{AUC} = \int_a^b R(P_1)dP_1. \tag{8}$$

This integral can be computed without an analytic solution, and indeed without specifying an equation for R, by using $x = F_1^{-1}(P_1)$ and $P_2 = F_2(x)$ (Equation 1) to compute the RPF as $P_2 = F_2(F_1^{-1}(P_1))$ (Equation 3) and performing numerical integration.

Normalizing the AUC by the length of the P_1 interval over which it is computed yields the average P_2 value over that interval:

$$\bar{P}_2 = \frac{1}{b - a}\int_a^b R(P_1)dP_1 = \frac{\text{AUC}}{b - a}. \tag{9}$$

$\bar{P}_2$ has an intuitive connection to the performance matching approach. Whereas performance matching seeks to measure the difference between subjective reports at a *fixed* value of task performance, comparing $\bar{P}_2$ for the metaperceptual RPF across conditions gives the *average* difference between subjective reports over a *range* of task performance levels. More generally, $\bar{P}_2$ measured over a given P_1 interval is the most direct analogue in RPF analysis to what a simpler, nonpsychometric design might measure as a point estimate of P_2 corresponding to a given value of P_1.

Equation 9 also implies a useful way of interpreting AUC: It is the average P_2 value over a given P_1 interval multiplied by the width of that interval, that is, $\text{AUC} = (b - a)\bar{P}_2$. These two measures thus give complementary perspectives on the data, with $\bar{P}_2$ being a more direct measure of (average) P_2 over a given P_1 interval, and AUC augmenting this measure with additional information about the P_1 interval over which it was computed. Further consideration of the roles of these measures in RPF analysis requires us to turn our attention to how we should choose the P_1 intervals over which they are computed.

Constraints on Defining Within-Subject P_1 Intervals. Let us begin by considering the P_1 intervals to be used for analyzing RPFs in different within-subject conditions (e.g., attended and unattended) for a single subject. Since our aim is to assess the effect of condition on the RPF as quantified by AUC and $\bar{P}_2$, and since these summarize RPF behavior over a given P_1 interval, it is important that we hold this interval constant to ensure that the same portion of the RPF is being compared across all conditions. But for AUC analysis over this fixed P_1 interval to be possible in every condition, it must be the case that all RPFs being analyzed fully span that interval. In general, this is not guaranteed to be the case unless the fixed P_1 interval is chosen appropriately. For instance, in a grating tilt discrimination task having conditions where the grating is attended or unattended across several levels of grating contrast spanning the full possible range of contrasts from 0 to 1, the fitted psychometric function for p(correct) in the attended condition may range from chance performance of 0.5 at zero contrast to a near-ceiling value (e.g., 0.98) at maximal contrast, whereas the fitted function for the unattended condition may range from chance performance at zero contrast to a level of performance at maximal contrast that is considerably lower than in the attended condition (e.g., 0.8). Thus, although the attended condition RPF spans a P_1 interval of [0.5, 0.98], its AUC can only be compared to that of the unattended condition for a fixed P_1 interval over [0.5, 0.8].

More generally, the full ranges of P_1 values spanned by each RPF being compared jointly determine lower and upper bounds on possible intervals of P_1 values that are common to all RPFs. The lower bound L on the common P_1 interval is given by

$$L = \max_c \ \min_x \ P_{1\ c,x}, \tag{10}$$

where $P_{1c,x}$ denotes the value of P_1 at condition c and stimulus level x. In other words, the lower bound for a common P_1 interval across conditions is the minimal *within-condition* value of P_1 that is maximal *across* conditions (i.e., the smallest P_1 value shared by all conditions). By similar reasoning, the upper bound U is given by

$$U = \min_c \ \max_x \ P_{1\ c,x}, \tag{11}$$

that is, the maximal within-condition value of P_1 that is minimal across conditions (i.e., the largest P_1 value shared by all conditions). For AUCs to be computed with a fixed P_1 interval $[a, b]$ that is common to all conditions, it must be the case that

$$a \geq L, b \leq U, \tag{12}$$

and the widest possible common P_1 interval is given by $[L, U]$.

These considerations are illustrated in Figure 3. The psychometric functions for P_1 in conditions A and B have different values at the minimum and maximum values of x (left panel), which entails that their corresponding RPFs do not span the same range of P_1 values (right panel). Thus, to compare AUC for a fixed P_1 interval, this interval must be restricted to the set of P_1 values that is common to both functions (shaded region). The lower and upper bounds of this common interval are set by the largest across-condition minimum P_1 value and the smallest across-condition maximum P_1 value, respectively, as described in Equations 10 and 11.

In the example of Figure 3, RPF AUC is larger for condition A than for condition B. This indicates that condition influences P_2 over and above its influence on P_1, and that this differential effect holds over a wide range of P_1 values.

It may often be the case that the variable chosen for P_1 is an accuracy measure that must take on a chance value in the absence of relevant stimulus information, for example, $d' = 0$ or $p(\text{correct}) = 0.5$ in a two-choice task. In these cases, L takes on this chance value, allowing for the lower bound of the RPF analysis to extend all the way to P_1's theoretical minimum value. By contrast, even if P_1 has a theoretical maximum (e.g., 1 for a probability), this value may not be achieved even when stimulus strength is maximal (e.g., contrast = 1) or arbitrarily high (for a stimulus feature with no maximum strength, e.g., brightness), as illustrated in the hypothetical example above where p(correct) is well below ceiling under inattention even at maximal stimulus contrast. Thus, in the common case where P_1 is an accuracy measure, U is the limiting factor for interval width.

Considerations for Defining P_1 Intervals Across Subjects. One way to define P_1 intervals across subjects is to simply use a single fixed interval. This approach would seem to lend a certain straightforwardness and simplicity to conducting and interpreting AUC analysis, and is a natural choice if one has a prior reason for being interested in RPF behavior over a specific range of P_1 performance. However, this approach also meets with significant difficulties. One is practical: The width of the fixed P_1 interval applied across subjects is constrained by the largest across-subject value for L and the smallest across-subject value for U (Equations 10 and 11), potentially yielding a narrow P_1

Figure 3

Comparing RPFs With the AUC Approach

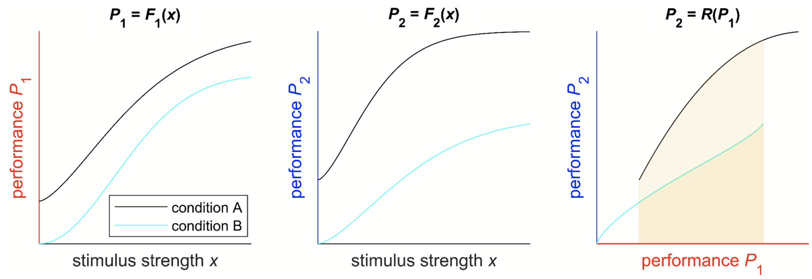

Note. In this illustrative example, psychometric functions for P_1 and P_2 differ across condition and so do their corresponding RPFs. The difference in the RPFs can be quantified by comparing the AUC over the set of P_1 values that both RPFs share in common (shaded region). Here, Condition A has the higher AUC, indicating higher levels of P_2 across the fixed P_1 interval. The AUCs can be divided by the length of this common P_1 interval to yield the average P_2 values over the interval. RPF = relative psychometric function; AUC = area under the curve. See the online article for the color version of this figure.

interval that might fail to adequately capture RPF behavior for many subjects.

The other primary difficulty is that it is not necessarily the case that the same P_1 interval houses comparable portions of the RPF for different subjects. For instance, a given interval might capture the rising slopes of the sigmoidal RPFs in each condition for subject A, whereas the analogous rising slopes for subject B might fall beyond the upper bound of that same interval. In cases like this, a better overall comparison of the RPFs for A and B would arguably be yielded by using different P_1 intervals that capture as much RPF behavior as possible for each subject, rather than using a fixed interval for both. Although this entails nonuniformity of P_1 intervals across subjects, in an important sense, it is the best way possible for the AUC method to achieve a uniform assessment of whole-RPF behavior for each subject.

Indeed, the power of RPF analysis lies in being able to assess the behavior of entire RPFs across broad ranges of P_1, ideally spanning floor to ceiling values if possible. As a general principle, then, it is desirable to use the widest P_1 intervals possible for AUC analysis (assuming one does not have specific interest in analyzing a more restricted interval). This can be accomplished by setting each subject's P_1 interval to the $[L, U]$ bounds computed from that subject's data (Equations 10 and 11). This yields the widest possible interval for each subject and thus maximizes P_1 coverage and corresponding RPF behavior for the analysis as a whole. However, we must also consider issues around the resulting across-subject heterogeneity in P_1 intervals.

Note that using different intervals across subjects is not nearly as problematic as using different intervals for different within-subject conditions (which the approach outlined above strictly prohibits). Doing the latter would confound the effect of condition on RPF behavior (what we want to know) with the effect of condition on what region of the RPF is selected for analysis (an "effect" of no interest that we should eliminate). By contrast, using different intervals across subjects does not introduce confounds with the effect of condition on RPF behavior (assuming a repeated measure design). And while this approach does introduce between-subject variability in the P_1 interval, it does so with the intention of minimizing the more important between-subject variability in how much of each subject's RPFs are excluded from the analysis.

Further considerations around across-subject heterogeneity in P_1 intervals arise in the interpretation of AUC and $\bar{P}_2$.

Considerations for Interpreting AUC and $\bar{P}_2$. When P_1 intervals are chosen to maximize interval width for each subject, we can consider statistical analysis of the $\bar{P}_2$ data to be answering the question, "What is the effect of condition on RPF behavior, as measured by the average P_2 over the widest possible P_1 interval for each subject?" Discoveries of any effects would then indicate that condition influences overall RPF behavior across a broad P_1 range.

As noted above, AUC over a given P_1 interval can be seen as the average P_2 over that interval multiplied by the interval width. On one way of looking, the dependence of AUC on interval width is a nuisance factor that we control for by computing $\bar{P}_2$. However, it is important to keep in mind that while $\bar{P}_2$ controls for the effect of interval width on the AUC computed over a given interval, it does not control for the influence of the *placement* of that interval relative to the RPF; the same RPF will yield different $\bar{P}_2$ depending on what aspects of its behavior are captured by the chosen P_1 interval, and between-subject differences in $\bar{P}_2$ might be partially attributable to differences in RPF coverage afforded by their different P_1 intervals. These concerns are minimized when we maximize the interval width for each subject.

From another point of view, in the context of comparing RPF behavior across the widest possible P_1 intervals for each subject, AUC's dependence on interval width can be seen as informative insofar as it magnifies $\bar{P}_2$ in proportion to the width of the interval over which it was computed, effectively giving greater influence on the outcome of statistical tests to subjects with wider P_1 intervals—that is, subjects for whom the measured $\bar{P}_2$ value characterizes a greater fraction of the RPFs' overall behavior.

Consideration of RPF coverage may be especially important in cases where interval constraints potentially limit probing RPF regions of particular interest. For instance, consider analysis of RPFs relating meta-d' (P_2) to d' (P_1). Since both variables have chance values of zero, and since we expect meta-d' to be monotonically increasing with d', any potential effects of condition on the RPF will likely manifest as influences on its slope (with additional possible influences on its potentially nonlinear, but still monotonically increasing, shape). Such changes are likely to result in changes in AUC and $\bar{P}_2$ that grow with the upper bound of the P_1 interval (with the lower bound always corresponding to the chance value of $d' = 0$). Therefore, subjects for whom the interval's upper bound is smaller will contribute AUC and $\bar{P}_2$ data that, on average, are a poorer reflection of RPF behavior in the P_1 regions where the effect of interest is maximal. In this case, statistical analysis of the AUC data might implicitly take this issue into account by "weighting" each subject's $\bar{P}_2$ data in proportion to their interval width, relative to the uniform "weight" of each data cell when the same statistical test is performed on the $\bar{P}_2$ data.

However, whereas Equation 9 implies a straightforward analytical sense in which AUC is $\bar{P}_2$ scaled by P_1 interval width, caution is warranted in regarding AUC as a "weighted" $\bar{P}_2$ in the context of statistical tests. The manner in which this "weighting" is implicitly achieved by AUC may differ in its formal characterization and interpretation from more principled approaches to weighting data in statistical analyses, and the interpretation and desirability of this implicit "weighting" may vary depending on the specific features of the RPFs and P_1 intervals under consideration as well as the assumptions and logic of the statistical analysis that is employed. A nuanced understanding of the relationship between AUC and $\bar{P}_2$ requires considerations of factors like these on a case-by-case basis.

Nonparametric Computation of AUC. In the foregoing, we have assumed that P_1 and P_2 data are fitted with psychometric functions F_1 and F_2. However, there may be cases where fitting P_1 and/or P_2 encounters difficulties, such as:

- The researcher may be uncertain about the most appropriate functional form to choose for F_1 and/or F_2.
- The researcher may prefer to avoid making parametric assumptions about F_1 and/or F_2.
- For certain dependent variables, it may be unclear, complicated, and/or labor intensive to develop an MLE fitting approach and implement the fitting procedure in analysis code (consider, e.g., the discussion of MLE fitting for d', meta-d' and mean rating as discussed above and in Supplemental Material S2.1).
- When plotted against stimulus strength x, the data to be fitted may be monotonically decreasing (e.g., reaction time data) or nonmonotonic (e.g., when rating confidence in a detection task, confidence may be high for "no" responses at low values of x and "yes" responses at high values of x, yielding a U-shaped function of confidence when collapsed across response type), whereas standard psychometric functions are monotonically increasing with x.
- Limitations and noise in the data may cause technical difficulties with the fitting procedure or may yield fitted parameter values that are implausible or present analysis difficulties (e.g., infinite slope).

These difficulties can be circumvented by computing AUC nonparametrically. The simplest nonparametric approach is to perform linear interpolation between the data points in the plot of P_2 versus P_1 and compute AUC from the resulting trapezoids, analogous to the nonparametric measure of area under the Receiver Operating Characteristic curve, A_g (Pollack & Hsieh, 1969). A hybrid approach can also be applied in which the RPF is constructed from a parametric fit of $P_1 = F_1(x)$ and a nonparametric estimation of $P_2 = F_2(x)$ via interpolation. However, note that a hybrid approach where P_1 data are interpolated and P_2 data are fitted is not viable, since the function yielded by interpolation of P_1 will in general not be monotonic with x and so will not be invertible—unless isotonic regression or similar approaches are employed, as discussed in Supplemental Material S2.2.4—preventing the computation of the RPF as described in Equation 3.

Although most implementations of nonparametric estimation of AUC can handle nonmonotonicity in the P_1 data as a function of x, AUC analysis is still only appropriate under the assumption that the underlying function F_1 that generated those P_1 data is itself monotonic (with nonmonotonicity in the P_1 data being attributed to statistical noise). If F_1 is not monotonic over a given P_1 interval, then the basic logic of using AUC to summarize RPF behavior over that interval breaks down (see Supplemental Material S1).

In Supplemental Material S2.2, we discuss methodological considerations for nonparametric computation of RPF AUC in more detail, and in Supplemental Material S2.3, we present simulations demonstrating that nonparametric methods are comparable to parametric methods at estimating the true AUC of a known generating RPF under data collection conditions typical of those used in psychophysical experiments. These sections also detail the reliability and accuracy of AUC estimation under full parametric psychometric function fitting.

Benefits and Limitations of the AUC Method. Summarizing RPFs with AUC (or $\bar{P}_2$) in this manner has a number of virtues:

1. *Ease of computation*. RPF AUC can be computed via numerical integration based on $F_1(x)$ and $F_2(x)$ without needing to find a closed form expression for $R(P_1)$.
2. *Ease of interpretation*. RPF AUC provides a single, easy to interpret measure (compared to the multiple, complex, interrelated parameters of the Weibull RPF, for example).
3. *Universality*. AUC computations are applicable to any RPF for any P_1 and P_2, regardless of the functional forms of F_1 and F_2 (provided that F_1 is monotonic).
4. *Robustness*. AUC is more robust to measurement error than general psychometric function parameter estimation. For instance, in certain cases, small changes in the data can yield relatively large differences in the fitted parameters without having large effects on the overall shape of the psychometric function, which in turn would lead to only small changes in the RPF AUC. In fact, AUC estimation can even be robust if $F_n(x)$ is constructed from piecewise linear interpolation rather than fitting a function, further simplifying the analysis approach; we explore this possibility in detail in the Supplemental Material S2.3.

The AUC method is most straightforward to interpret in cases where the RPFs do not intersect over the chosen fixed P_1 interval, since in such cases, the values of P_2 in one condition are always higher than in the other for every value of P_1 in the interval. However, if the empirical RPFs *do* intersect in this interval, then the relationship between AUCs across conditions differs on either side of the intersection point, which complicates interpretation of AUC computed over the whole interval. Two possibilities must be considered: (1) The "true" generating RPFs are similar or identical over this interval, and the intersection in the empirical RPFs is due to statistical noise; or (2) the generating RPFs are distinct and do indeed intersect over this interval, as is validly reflected in the empirical RPFs. Since across-condition AUCs have opposite relationships on either side of the intersection point, computing AUC over the entire interval will tend to wash out any across-condition differences. This behavior can be a virtue that accurately reflects the absence of an effect in case (1) but may underestimate or even fail to detect the presence of a true effect in case (2).

For instance, consider an idealized case where over a P_1 interval [0, 1], the empirical RPF in condition A has a constant value of 0.5, and the empirical RPF in condition B is linear with values [0, 1] at the end points of the P_1 interval. In this case, RPF A forms a rectangle with base 1 and height 0.5, and RPF B forms a triangle with base 1 and height 1 that intersects RPF A at $P_1 = 0.5$. Both RPFs have an AUC of 0.5 despite differing considerably in their shape, since A's AUC is larger than B's over $P_1 \in [0, 0.5]$ and the opposite is true over $P_1 \in [0.5, 1]$.

Thus, if the "true" generating RPFs for A and B are similar or identical, and the empirical RPFs A and B differ due to noise, their identical AUCs will accurately reflect the absence of a difference in the generating RPFs. Conversely, if the generating RPFs have forms that are well represented by the empirical RPFs A and B, then computing AUC over the interval [0, 1] will fail to quantify the difference between the generating RPFs due to their intersection over that interval. In cases where the generating RPFs intersect in this way, the model comparison approach described below can still detect the difference between them.

Model Comparison Approach. An alternative approach to comparing RPFs across conditions is to capitalize on the observation that if a functional form for the RPF is available, the parameters of this function can be constrained in such a way as to ensure that fitted RPFs across conditions are identical. The data can then be fitted with two different models, one of which allows parameters to vary freely in such a way that the fitted RPFs can differ across conditions ("varying RPF model"), and one of which constrains parameters in such a way that the fitted RPFs are constrained to be constant across conditions ("fixed RPF model"). Standard model comparison analysis approaches can then be conducted to investigate whether the varying or fixed RPF model provides a better account of the data, taking into account how the greater degrees of freedom in the varying RPF model introduce the possibility of overfitting. This model comparison analysis can be performed, for example, with information theoretic measures such as Akaike information criterion or Bayesian information criterion (Vrieze, 2012) or alternatively with cross validation methods to assess model generalizability (de Rooij & Weeda, 2020).

For instance, consider the functional form of the Weibull RPF R_W discussed above (Equation 6). A trivial way to ensure that R_W is constant across conditions would be to constrain all parameters for F_1 and F_2 to be constant across conditions. However, a more artful approach would be to allow the parameters of F_1 and F_2 to have the maximal degree of freedom possible while still constraining the corresponding RPFs to be constant across conditions. Investigation of Equation 6 shows that this latter goal can be achieved by constraining the following parameter values to be constant across conditions: γ_1, λ_1, β_1, γ_2, λ_2, β_2, and $\frac{\alpha_2}{\alpha_1}$. This set of constraints allows α_1 and α_2 to vary across conditions in the fitting of F_1 and F_2 so long as the resulting parameter values conserve constant values for $\frac{\alpha_2}{\alpha_1}$ across conditions, thus allowing for variation in the fits for F_1 and F_2 while preserving a fixed RPF. The varying RPF model, by contrast, would relax some or all of these constraints and thus allow fitted RPFs to differ across conditions.

Generalizing the above discussion, there are actually multiple ways to define "fixed RPF" and "varying RPF" models, depending on what constraints on across-condition parameter values are imposed over and above the key set of constraints determining whether RPFs can vary across conditions or not. For example, we have just shown how there are two ways to constrain parameters across conditions to yield fixed RPF models for the Weibull RPF, and every way to relax any aspect of these constraints—such as, for example, allowing β_1 but not β_2 to differ across condition, or allowing α_1 and α_2 be completely unconstrained across condition, and so on—would yield a different varying RPF model. An extended model comparison analysis could thus consider a family of models, some of which are fixed RPF models and others of which are varying RPF models. Interpretation of the results of such an analysis could reveal findings such that, for example, the best fitting model constrains $\frac{\alpha_2}{\alpha_1}$ to be constant across conditions but allows β_1 and β_2 to vary, or similar patterns, which could provide a more nuanced understanding of how condition influences specific aspects of RPF behavior and potentially shed light on the underlying mechanisms.

Regardless, the most basic and foundational question would still be whether the empirical RPFs are best characterized by fixed or varying RPF models. If varying RPF models are best supported in the model comparison analysis, this would suggest that RPFs are modulated by condition, and thus that the psychological processes generating P_1 and P_2 are at least partially separable. Conversely, if fixed RPF models are best supported, this would suggest that the observed RPFs are consistent with the possibility that P_1 determines P_2 (or vice versa), or that both are generated by a single underlying process characterized by a constant RPF.

The model comparison approach has several advantages over the AUC approach. Whereas the AUC approach has certain complexities and limitations related to considerations of the P_1 interval, modeling can straightforwardly take into account all aspects of RPF behavior for every condition of every subject. Model comparison can also achieve a more precise and nuanced analysis of RPF behavior; it can detect differences in RPFs even in cases where RPFs intersect in a way that yields similar AUC values, and it can pinpoint which specific aspects of RPF behavior (as controlled by specific RPF parameters) are influenced by condition. Additionally, unlike AUC analysis, model comparison can be applied to analysis of relative psychometric performance even if P_2 cannot be expressed as a function of P_1 (Supplemental Material S1). However, the model comparison approach has the disadvantages that it is more complex and resource intensive to conduct and may require devising an approach that is tailor-made to a specific analysis plan or data set. Additionally, although most aspects of model comparison can proceed via assessing the joint fits of F_1 and F_2 across conditions, characterizing what parameter constraints yield the minimally constrained fixed RPF model may require deriving and analyzing an analytic expression for R (as, e.g., in the above consideration of fixed RPF models for the Weibull RPF).

Empirical Case Study

In this next section, we demonstrate the power and utility of the RPF method by applying it to an empirical data set in which subjects made perceptual decisions about coherent dot motion and rated confidence. Seven levels of motion coherence were presented, allowing construction of psychometric functions for accuracy, confidence, and metacognitive sensitivity. Experimental conditions were contrived so as to attempt to modulate the relationship between confidence and accuracy, naturally inviting an RPF analysis approach.

Experimental Methods

Twenty-one healthy adult human subjects viewed random dot kinematogram (RDK) stimuli which continuously filled the entirety of a computer monitor with random dot motion. In a two-interval forced choice task design, on each trial of the experiment, a circular patch of these dots to the left or right of a central fixation cross briefly displayed coherent motion in a downward direction (Figure 4). The observer's task was first to indicate which side of the display contained the coherent downward motion and then to rate their confidence on a scale of 1–4, using separate keypresses for each report.

We varied three aspects of the task to examine their effects on the relationship between accuracy and confidence. First, we varied motion coherence by randomly selecting the coherence of the downward dot motion on each trial from a list of seven values evenly spaced between 10% and 80% coherence. Second, we varied dot density by setting the density of the dots across the whole display on each trial to one of three levels (low = 1 dot/deg^2, medium = 3 dots/deg^2, high = 9 dots/deg^2). Third, we varied changes in dot density across trials by either setting dot density randomly on each trial (interleaved trial structure) or by holding dot density constant with each block of trials (blocked trial structure; Figure 4).

See Supplemental Material S3 for full details of participants, stimuli, equipment, and experimental design. This study was not preregistered.

Data Analysis

Following previous demonstrations (Koizumi et al., 2015; Odegaard, Grimaldi, et al., 2018; Rollwage et al., 2020; Samaha et al., 2016; Stolyarova et al., 2019), we expected higher dot density conditions to yield higher confidence, even when task performance was similar. We also examined whether metacognitive sensitivity—quantified as meta-d'—would differ across dot density conditions.

To explore these possibilities, we fit psychometric functions to d', mean confidence, and meta-d' as a function of dot motion coherence for each subject in each condition (dot density: high, medium, low; trial structure: blocked, interleaved). We fit scaled Weibull functions for d' and meta-d and modeled mean confidence via the sum of three separate Weibull fits to $p(\text{confidence} \geq y|x)$ data for $y \in \{2, 3, 4\}$, using the methods for MLE fitting of these variables developed in the Supplemental Material S2.1. This allowed us to specify two categories of metaperceptual RPFs, one relating d' to mean confidence and another relating d' to meta-d'. For each of these

Figure 4
Behavioral Task Procedures

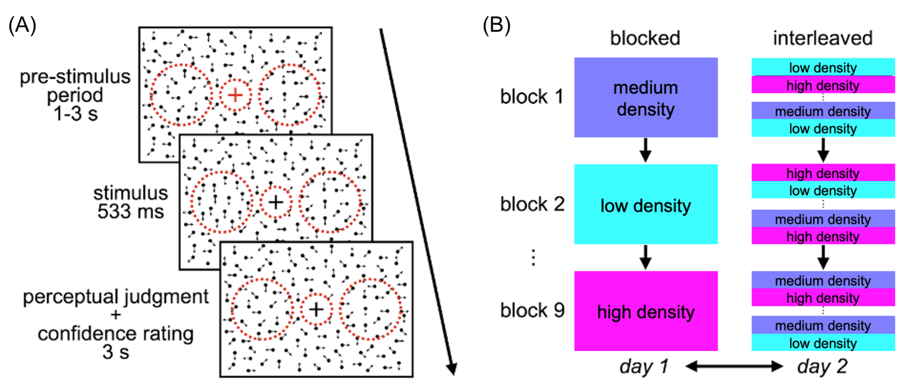

Note. (A) Each trial began with a prestimulus period, during which full-field random dot motion was shown (black arrows illustrate dot motion direction). Subsequently, within one of two circular regions of the screen (indicated here by the red circles to the left and right of fixation—red circles were shown to participants only during preliminary practice trials but not during experimental trials), coherent downward dot motion occurred, followed by a response period in which participants indicated on which side they saw the coherent motion and rated their decision confidence. The central red circle indicates an area around the fixation cross where no dots were presented; this red circle was not shown to participants and is used here for illustration purposes. (B) Participants underwent two trial structure conditions, blocked and interleaved, on two different days of testing. In the blocked condition, dot density was constant across trials within a given block, whereas in the interleaved condition, dot density varied randomly across trials. Blocked versus interleaved days and order of density blocks was counterbalanced across all participants. See the online article for the color version of this figure.

metaperceptual RPFs, we computed AUC and $\bar{P}_2$ separately for each condition of each subject, using the widest possible P_1 interval for each subject (Equations 10 and 11). We assessed the effects of the experimental manipulations on these measures using 3 (Dot Density: High, Medium, Low) × 2 (Trial Structure: Blocked, Interleaved) repeated-measures analyses of variance (ANOVAs). For further details on how we implemented RPF analyses for these data, see Supplemental Material S3.5.

All RPF analyses were performed using our open-source RPF toolbox, available at https://github.com/CNClaboratory/RPF (Maniscalco & Peters, 2026); data and code used specifically for these analyses are also shared openly on the Open Science Framework (https://osf.io/62rwk; Castaneda et al., 2026).

Subsequent to performing RPF analysis, we noticed that one subject exhibited an abnormally narrow P_1 interval due to particularly poor task performance (d') in the interleaved condition, and so we chose to exclude this subject from the main analysis results (see Supplemental Material S3.1 for full discussion). However, we also present full statistical results for all analyses when including this subject's data in Supplemental Material S3.6. The net effect of including this subject is that all statistically significant results get slightly weaker, with all significant main effects remaining significant, but some significant interaction effects falling slightly below the significance threshold. We moderate our conclusions in the discussion below accordingly.

Empirical Results and Discussion

We found that dot density did indeed have the expected effect on the RPF for mean confidence as a function of d', particularly in the interleaved trial structure. We also found evidence that the RPF for meta-d' versus d' varies as a function of trial structure and exhibits sensitivity to dot density in the interleaved trial structure.

In the plotted data (Figures 5 and 6), for illustrative purposes, we show RPF curves for each trial structure and dot density condition that are fitted to the combined group data concatenated across all subjects, rather than an average across the fitted curves for each subject individually. However, we remind the reader that all statistical measures were derived from single-subject fits. In Figure 7, we show group averages of raw AUC and $\bar{P}_2$ as a function of dot density and trial structure for both the mean confidence and meta-d' RPFs.

For mean confidence versus d' (5), using the raw AUC (Figure 7A) in a 3 (Dot Density: High, Medium, Low) × 2 (Trial Structure: Blocked, Interleaved) repeated-measures ANOVA, we found a main effect of Dot Density, $F(2, 38) = 4.169$, $p = .023$, $\eta_p^2 = 0.180$, but not Trial Structure, $F(1, 19) = 0.102$, $p = .752$, $\eta_p^2 = 0.005$, and a significant Trial Structure × Dot Density interaction, $F(2, 38) = 4.867$, $p = .013$, $\eta_p^2 = 0.204$, such that mean confidence increased with increasing dot density in the interleaved but not blocked trial structure. The pattern was similar when we used a second repeated-measures ANOVA to examine the normalized AUC measure $\bar{P}_2$ (Figure 7B), with a main effect of Dot Density, $F(2, 38) = 4.354$, $p = .020$, $\eta_p^2 = 0.186$, but not Trial Structure, $F(1, 19) = 0.173$, $p = .682$, $\eta_p^2 = 0.009$, and a significant Trial Structure × Dot Density interaction, $F(2, 38) = 4.689$, $p = .015$, $\eta_p^2 = 0.198$—again suggestive that mean confidence increased with increasing dot density primarily in the interleaved but not blocked trial structure.

For metacognitive sensitivity (meta-d') versus d' (6), we observed a somewhat different pattern. Using raw AUC (Figure 7C), a repeated-measures ANOVA revealed no main effect of Dot Density, $F(2, 38) = 1.979$, $p = .152$, $\eta_p^2 = 0.094$, but a main effect of Trial Structure, $F(1, 19) = 7.697$, $p = .012$, $\eta_p^2 = 0.288$, and again a Trial Structure × Dot Density interaction, $F(2, 38) = 3.491$, $p = .041$, $\eta_p^2 = 0.155$, such that meta-d' was significantly lower in the interleaved trial structure, particularly for lower dot densities. A final repeated-measures ANOVA on the normalized AUC measure $\bar{P}_2$ (Figure 7D) revealed again no main effect of Dot Density, $F(2, 38) = 1.848$, $p = .171$, $\eta_p^2 = 0.089$, but a main effect of Trial Structure, $F(1, 19) = 9.341$, $p = .007$, $\eta_p^2 = 0.330$, and a marginal interaction between Trial Structure and Dot Density, $F(2, 38) = 2.916$, $p = .066$, $\eta_p^2 = 0.133$—again suggestive that metacognitive sensitivity was lower in the interleaved trial structure, particularly for lower dot densities.

It is notable that RPF analysis for meta-d' versus d' has special resonance with existing analysis approaches that call for meta-d' to be interpreted by way of comparison to d' (Fleming, 2017; Maniscalco & Lau, 2012, 2014), for example, by computing $M_{\text{ratio}} =$ meta-d'/d' as a measure of metacognitive efficiency, with values lower than 1 indicating suboptimal metacognitive sensitivity relative to SDT expectation. In plots of meta-d' versus d', the slope of the function roughly corresponds to M_{ratio} (with this relationship being exact for linear functions with y-intercept = 0). RPF analysis suggests another natural way to link meta-d' versus d' functions with M_{ratio}: dividing the average meta-d' value over a given d' interval (i.e., $\bar{P}_2$) by the average d' value over that interval (i.e., $\frac{b+a}{2}$) yields a measure of average M_{ratio} over the interval. Characterizing RPF analysis results in this way can add extra richness to our understanding. For instance, Figure 7D shows average meta-d' over each subject's d' interval for each condition. Dividing these by the corresponding average d' values for each subject's d' interval and averaging across subjects reveals that in the blocked trial structure, average M_{ratio} is around 0.87 for all dot densities, whereas in the interleaved trial structure, M_{ratio} takes on values of 0.6, 0.66, and 0.86 for low, medium, and high dot densities. These quantities give further context for understanding subjects' overall metacognitive performance and how this is modulated by condition. More generally, these observations highlight how RPF analysis of meta-d' versus d' synergizes with and extends existing frameworks for analyzing metacognitive sensitivity relative to task performance.

These findings are of utility to the community in several ways. First, from a basic science perspective, the observation that a manipulation as simple as the density of an RDK can induce changes in overall mean confidence over and above any effect on task performance capacity is consistent with findings in the literature on the so-called "positive evidence" or "response-congruent evidence" bias in metacognition (Rollwage et al., 2020; Samaha & Denison, 2022). In a number of empirical investigations, it has been shown that higher amounts of absolute magnitude of evidence available to the observer to make a perceptual decision are associated with increased subjective confidence reports; these manipulations of evidence can take the form of contrast or luminance (e.g., Koizumi et al., 2015; Rausch et al., 2017; Samaha et al., 2016, 2019), dot motion coherence (Zylberberg et al., 2012), or even more cognitive type evidence such as facial attractiveness (Ceja et al., 2022). (Note: These demonstrations of "positive evidence" or "response-congruent evidence" bias in metacognition are not without

Figure 5

Results of the Empirical Case Study Showing the Metaperceptual Relative Psychometric Function Relating d′ and Mean Confidence Ratings

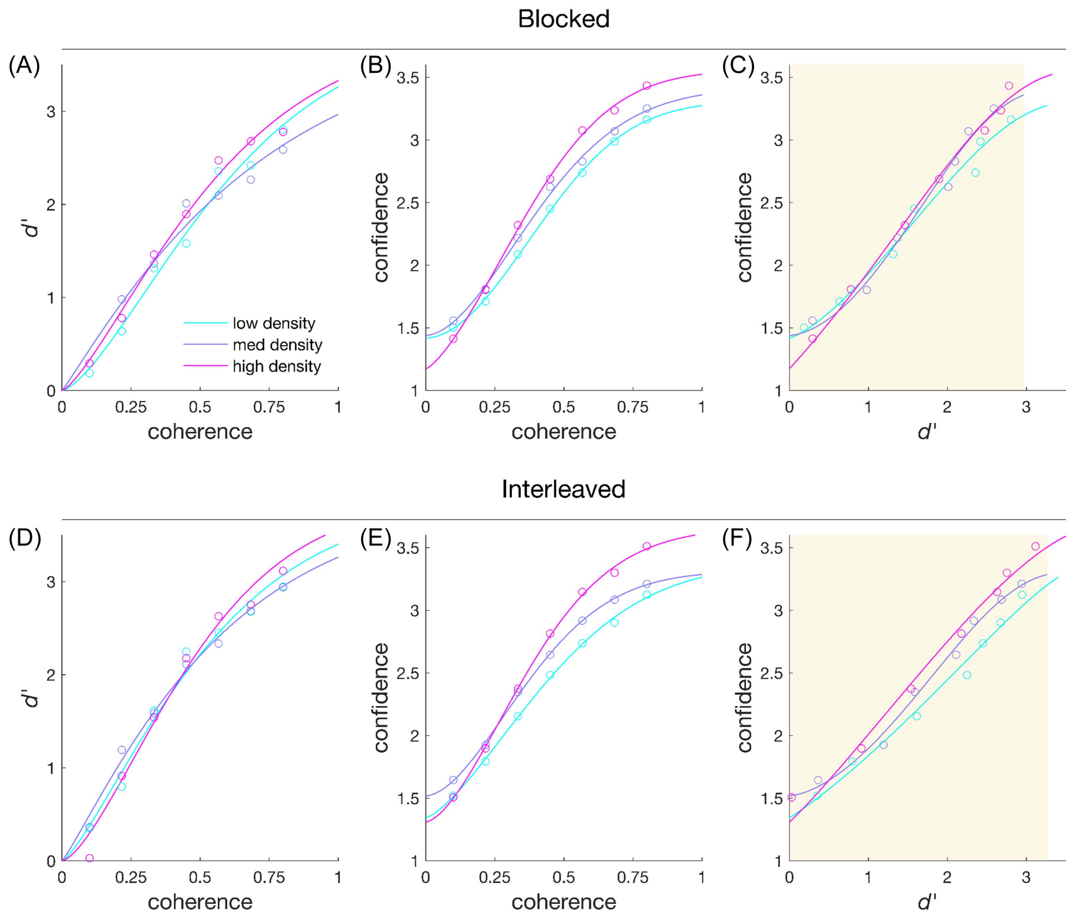

Note. Plots here visualize the statistical effects across subjects (see Figure 7A, 7B) via direct fits to the group-level data. (A–C) show F_1 (d' vs. dot motion coherence), F_2 (mean confidence vs. dot motion coherence), and the RPF R (mean confidence vs. d') for the blocked trials; (D–F) show the same for the interleaved trials. Fitted RPFs for the blocked (C) and interleaved (F) trial structure show visually that the blocked trials resulted in little-to-no apparent differences in the RPF as a function of dot density, while the interleaved trials show RPF separation with higher mean confidence in higher dot density conditions over the same interval of task performance. Shaded regions in (C) and (F) show the d' interval common to all dot density conditions in the group-level data. RPF = relative psychometric function. See the online article for the color version of this figure.

counterexample; see, for example, Locke et al., 2020; Boldt et al., 2025.) Models have been proposed to account for these and similar findings, placing constraints on how confidence might be (neurally) computed in perceptual decisions (e.g., Maniscalco et al., 2016, 2021; Peters, Thesen, et al., 2017).

Here, we add to these previous findings by investigating how trial structure and dot density manipulations influence not only mean confidence levels but also the metacognitive sensitivity exhibited by those confidence ratings (as measured by meta-d'). We found a seemingly robust main effect, whereby meta-d' decreases when different dot densities are interleaved on a trial-by-trial basis, with weaker statistical suggestions that this decrease may be particularly strong for lower dot densities. These findings may be relevant to recent work exploring observers' capacity to update Type 2 criteria with changing uncertainty conditions (Lee et al., 2023; Rahnev & Denison, 2018). However, although lower *overall* metacognitive sensitivity in interleaved trials is consistent with the possibility that interleaved dot densities induce trial-to-trial shifts in Type 2 criteria (which would result in a lower measured value for metacognitive sensitivity), it remains to be explained why this effect might differ across specific dot density levels.

We regard these meta-d' results as tentative, pending subsequent empirical corroboration with higher powered experimental designs. The data set used here was not designed with meta-d' analysis in mind; only 36 trials were presented for each permutation of motion coherence, dot density, and trial structure, and despite the visual appearance of a seemingly impressive interaction between dot density and trial structure (6 and 7), there was only inconsistent statistical support for this interaction across all analyses we conducted. Nonetheless, the present results suggest potentially fertile

Figure 6

Results of the Empirical Case Study Showing the Metaperceptual RPF Relating d′ and Metacognitive Sensitivity, Measured With Meta-d′

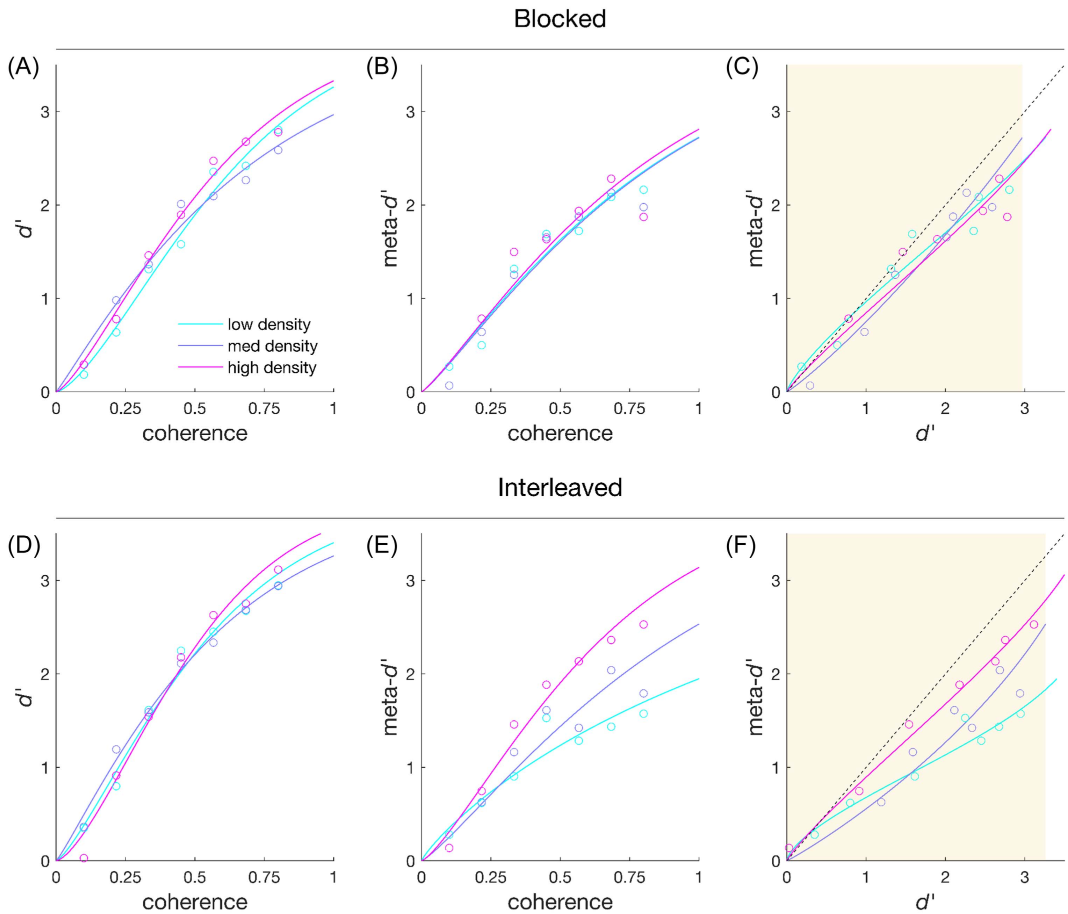

Note. Plots here visualize the statistical effects across subjects (see Figure 7C, 7D) via direct fits to the group-level data. Similar to the plots for mean confidence (Figure 5). (A–C) show F_1 (d' vs. dot motion coherence), F_2 (meta-d' vs. dot motion coherence), and the RPF R (meta-d' vs. d') for the blocked trials; (D–F) show the same for the interleaved trials. Fitted RPFs for the blocked (C) and interleaved (F) trial structure show visually that the blocked trials resulted in no apparent differences in meta-d' as a function of dot density, while the interleaved trials show RPF separation with higher meta-d' in higher dot density conditions over the same interval of task performance. Shaded regions in (C) and (F) show the d' interval common to all dot density conditions in the group-level data. Dashed lines in (C) and (F) show the line of unity where meta-$d' = d'$, corresponding to the expected value of meta-d' under signal detection theory. RPF = relative psychometric function. See the online article for the color version of this figure.

paths for future research. If future research corroborates the tentative findings presented here, such findings would contribute to our understanding of the computational relationship between performance and confidence ratings on a trial-by-trial basis and provide rich opportunities for future computational studies to use these and similar kinds of data to arbitrate among candidate process models giving rise to metacognitive judgments (e.g., Adler & Ma, 2018a, 2018b; Aitchison et al., 2015; Denison et al., 2018; Kiani et al., 2014; Maniscalco et al., 2016; Miyoshi & Lau, 2020; Peters, Fesi, et al., 2017; Peters & Lau, 2015; Peters, Thesen, et al., 2017; Winter & Peters, 2022; Xue et al., 2024b; Zylberberg et al., 2014, 2016).

Our intent with this empirical demonstration was not only to demonstrate that manipulations of dot motion evidence may influence confidence and task performance in distinct ways, which has previously been established. Here, we were also concerned with showing how full RPF analyses can provide benefit over measuring differences in confidence or metacognitive sensitivity at one or two levels of (matched) performance. Importantly, one can see clearly through the RPF approach that the size of the effect on subjective experience depends strongly on the performance level at which the effect is measured: For both mean confidence (5 and 7A, 7B) and metacognitive sensitivity (6 and 7C, 7D), the difference in the subjective report appears to grow as a function of task performance; this occurs with specific quantitative relationships to task performance in both cases. By extension, if a researcher were to try to measure the effect size of a manipulation's influence on mean confidence or meta-d', but were unable to precisely match task performance across conditions or subjects, the effect size of interest

Figure 7

Results of the Empirical Case Study Showing the AUC-Based Analyses Using AUC and $\bar{P}_2$

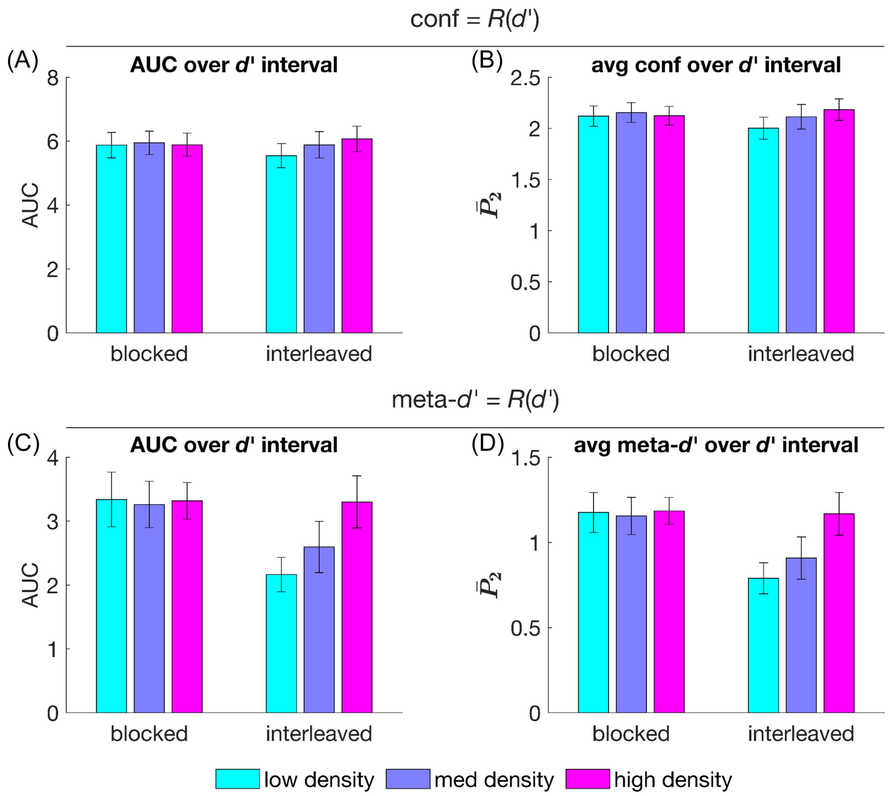

Note. Both raw AUC (A, C) and its normalized variant, $\bar{P}_2$ (B, D) confirm visual inspection of the RPFs in Figures 5 and 6, showing significant main effects of the ANOVA factors dot density (for the average confidence RPF) and Trial Structure (for the meta-*d′* RPF), with both main effects being modulated by Trial Structure × Dot Density interactions. That is, higher dot density led to higher mean confidence over a matched performance interval, especially (or only) in the interleaved trials; and the interleaved trial structure led to lower metacognitive sensitivity (meta-*d′*) than the blocked trial structure over a matched performance interval, especially for lower dot densities. See main text for statistical details. RPF = relative psychometric function; AUC = area under the curve; ANOVA = analyses of variance; avg = average; conf = confidence. See the online article for the color version of this figure.

would be at best poorly estimated, or at worst entirely missed (e.g., at lower levels of *d′*). By measuring and fitting the entire RPF and engaging in the AUC-based analyses presented here, such differences due to nuisance variables can be minimized, revealing a robust and quantitatively precise measure of subjective experience differences independent of task performance.

Importantly, process models of metacognition or subjective experience in general—such as those mentioned above—become much more highly constrained if they must explain behavioral data across the entire RPF in multiple conditions, opening an exciting new set of questions for the community (Fleming, 2023; Peters, 2022, 2025). Recently, numerous process models purporting to capture the relationship between performance and confidence in tasks using RDKs and other stimuli have been proposed, with increasing emphasis on attending to differences in experimental setup, behavioral paradigm, and stimuli to explain seemingly divergent findings (e.g., Fung & Rahnev, 2024; Fung et al., 2025; Xue et al., 2024a). These newer explorations are especially important because most available data sets—and indeed, most studies done in the field—test limited combinations of stimuli and task designs (Rahnev et al., 2020; Shekhar & Rahnev, 2024), leading to calls for expansion into more generalizable approaches (Rahnev et al., 2022) and tasks designed to specifically and efficiently arbitrate among candidate process models (Rong & Peters, 2023). Excitingly, the RPF analysis demonstrated here is not limited to the behavioral metrics we explored (percent correct, mean confidence, and meta-*d′*) but could also be used to highly constrain competing process models which may include evidence accumulation time (Kiani et al., 2014; Maniscalco et al., 2021; Zylberberg et al., 2016), or to evaluate whether and how static versus dynamic integration of prior expectations is integrated into the decision-making process (Khoudary, Bornstein, & Peters, 2025; Khoudary, Peters, & Bornstein, 2025). That our stimuli were dynamic RDKs—well-studied and characterized at most levels of visual processing—can meaningfully set these activities on the right track, and we welcome any follow-up attempts using our data (shared openly on the Open Science Framework; link included in the Data and Code Availability section).

While the data presented here perhaps consist of fewer subjects than would be ideal, they nevertheless demonstrate the robustness of the RPF method to small numbers of subjects or even to few trials per condition. Here, we collected only 36 trials per condition (seven levels of RDK motion coherence, Blocked vs. Interleaved Trial Structure, and three levels of Dot Density). The entire data set was collected across only approximately 2 hr of testing per subject, meaning that even as few as 36 trials per condition can be sufficient for conducting robust and precise RPF analyses comparing across conditions with the AUC-based metrics. (Of course, more trials are better, and following best practices for fitting d' or meta-d' in any data set would suggest at least 100 trials per condition for robust estimates of these metrics.) Future work should seek to confirm and expand the initial findings presented here regarding the effect of dot density and trial structure manipulations on performance, confidence, metacognitive sensitivity, and their interrelations.

Overall, this empirical case study highlights an exciting direction for the study of subjective experience and for use of the RPF analytic approach in general. We believe these results and the analytic approach to be of great value both within the metacognition and subjective experience community (Michel et al., 2019; Rahnev et al., 2022) and beyond.

General Discussion and Future Directions

Summary

In this piece, we have laid out a novel framework for investigating, in general, the quantitative relationship between two psychological processes measured under noisy conditions and how these relationships may vary with any experimental manipulation or intervention that is of interest to the researcher. This framework includes the derivation of the *relative psychometric function* (RPF) under parametric assumptions, including special considerations for fitting customized psychometric functions to nonstandard psychometric variables such as task performance capacity measured with the signal detection theoretic metric d', average confidence ratings, and metacognitive sensitivity (meta-d'; Maniscalco & Lau, 2012, 2014). We also developed and tested a series of metrics and algorithms designed to provide intuitive insight into how the RPF may change across experimental conditions, including the area under the RPF (AUC) method and its normalized variant, $\bar{P}_2$. These metrics provide a clear, precise, and interpretable approach for interpreting variations in RPFs across experimental conditions. And for those who wish to precisely evaluate other relationships among RPFs or how they may be captured by process models (e.g., signal detection or Bayesian decision theoretic, evidence accumulation models, etc.), we also lay out a model comparison approach in which the RPF can be constrained to be equivalent across conditions or free to vary in different ways. This model-based approach can provide important nuance and context to supplement the AUC-based analyses developed here.

We demonstrated the utility of the RPF framework by way of example, showing how the RPF approach can facilitate quantifying precisely how a manipulation of interest impacts subjective processing independent of (or over and above its effect on) objective processing. In this case study on the metaperceptual RPF, we found that our dot density manipulation led to changes in mean confidence that were separable from the influence of this manipulation on task performance, particularly when dot densities were randomly interleaved across trials rather than being organized into blocks of trials. We also found preliminary evidence that interleaving dot density reduces meta-d', particularly for trials with lower dot density.

Although these empirical results are valuable and contribute to the literature on how metacognition behaves, our primary excitement lies in the promise of the RPF framework to study the quantitative relationship between any pair of psychological variables the researcher may desire. Thus, we emphasize that the RPF framework can be used not only to study the relationship between objective processing capacity and subjective experience, but for characterizing the quantitative relationship among any two (likely nonlinearly) related psychological processes—including those for which no functional form relating each process to objective stimulus properties is known or presumed (see Supplemental Material S2.2 and S2.3 for details). As part of developing this highly general analysis framework, we have also developed novel methods for conducting maximum likelihood fitting of psychometric functions to three particular dependent variables that may commonly be of interest for RPF analysis: d', meta-d', and mean confidence (Supplemental Material S2.1).

We have also developed the RPF toolbox as an open-source community resource, available for download and extension from https://github.com/CNClaboratory/RPF. The toolbox supports a full analysis pipeline from raw trial-level data for a single subject to comprehensive RPF analysis results and plots. It is designed to allow for an easy, out-of-the-box analysis pipeline that can be conducted using only a few high-level functions while implicitly handling many of the subtleties and complexities of RPF analysis under the hood, while still allowing for complete control and customizability of the finer details of the analysis where desired. It is highly flexible, including built-in support for computing many dependent variables of interest from trial-level data and various methods for fitting or interpolating the data. For more details, see Supplementary Material S4.

Advantages of the RPF Method Over Standard Performance-Matching for the Study of Subjective Experience

A primary use for RPF analysis is for isolating the neural or computational correlates of subjective aspects of perception from those giving rise to task performance. Since Lau first articulated this need (Lau, 2008; Lau & Passingham, 2006), many groups have sought to control for "performance confounds" by finding one or two levels of matched performance across various experimental manipulations, and then examining how subjective measures differ (Morales et al., 2022; Peters, Kentridge, et al., 2017). However, as mentioned in the introduction, this performance matching approach is limited for two primary reasons: It relies on a statistical null effect (finding conditions where subjective experience differs but performance *fails* to differ), and differences in subjective experience can depend on the level of matched performance selected by the experimenter.

As we have seen, RPF analysis circumvents these challenges by revealing differences in P_2 (e.g., confidence) over an entire matched interval of P_1 (e.g., performance). Importantly, however, we can also relate components of RPF analysis directly back to more traditional performance-matching approaches to facilitate direct comparison with existing literature. For example, we can see that if one measures the entire RPF for each condition of interest, RPF AUC analyses can be tuned to any intervals within the available common P_1 interval across conditions of interest. In the limit, as this interval approaches zero, computing RPF AUC reduces to "reading off" the P_2 (subjective) values given a particular P_1 (performance) value, that is, selecting *exactly* the matched performance level desired through relying on the fitted functions. Doing so avoids the methodological and statistical disadvantages of using staircasing or other methods to discover conditions where subjective measures vary but performance measures *fail* to vary. The freedom to select one or two levels of exactly matched performance also evokes performance-matching studies which have used two or more levels of performance-matched conditions (e.g., hard and easy, Koizumi et al., 2015; see Rahnev et al., 2020, for other potential data sets). RPF AUC analyses could be used to reexamine such data, potentially providing a more principled analytic approach; this might also be possible through the interpolation-based nonparametric approach (described in more detail in Supplemental Material S2.2 and S2.3) even if fitting a parametric RPF is not possible. Thus, RPF analysis provides a natural extension to more traditional performance-matching approaches in a way that facilitates direct comparison to previous empirical and theoretical literature.

Relationship to Other Recent Work Linking Relative and Absolute Judgments

The study of psychophysics has a long and clever history, spanning 150 years of quantitative psychological research. A large literature has developed documenting the relationship between small changes in physical stimulus magnitude and the ability of humans (or nonhuman animals) to discriminate or detect such differences as well as the relationship between physical stimulus magnitude and absolute stimulus magnitude judgments—even of a subjective nature (brightness, loudness, painfulness, and so on). Weber's law, Fechner's law, Stevens' power law—these are all well-known, foundational examples that collectively support quantitative psychology across nearly countless domains of study.

Recently, a unifying framework linking such relative and absolute psychometric judgments—that is, the relationship between questions such as "Was the left light brighter than the right one?" versus "How bright is this light?"—was proposed by Zhou et al. (2024). In this work, the authors combined generalizations of work by Fechner and classic SDT to show how internal noise properties that accompany stimulus representation can explain so-called "power law" intensity percepts. Their unifying framework thus elegantly links both relative and absolute psychophysical judgments to stimulus properties in the environment.

On the surface, Zhou et al.'s (2024) approach may seem related to the RPF approach developed here. However, their goal was to unify multiple psychophysical laws through exploring mappings among stimulus strength, perceived stimulus strength, and noise models. Here, instead, we propose a psychometric framework to quantitatively characterize the relationship between *any* two psychological processes—not just relative and absolute intensity judgments. Specifically, we have through our case study focused on the "metaperceptual" RPF. This form of the RPF is thus not limited to judgments about the subjective evaluation of *stimuli in the world* (absolute magnitude estimation judgments) but is also capable of handling *introspective* or *metacognitive* judgments (judgments about one's own processing capacity or one's own internal experience). In other words, the metaperceptual RPF is sufficiently general so as to evaluate the relationship between the world and first-order internal representations of the world, and between those first-order internal representations and higher order metacognitive or introspective evaluation of them (R. Brown et al., 2019; Overgaard & Mogensen, 2017). Thus, the metaperceptual RPF directly addresses recent calls for a psychophysical introspective research program (Fleming, 2023; Kammerer & Frankish, 2023; Morales, 2024; Peters, 2024) as a targeted technique for understanding phenomenological experience in general (Peters, 2022, 2024), building upon previous research programs seeking to isolate subjective experience for scientific study by holding performance constant (Lau, 2008; Lau & Passingham, 2006; Morales et al., 2022; Peters, Kentridge, et al., 2017). We expect that other field- and question-specific variants of the RPF will emerge, for example, relating confidence judgments to reaction times, clarity assessments to criterion bias, or even extension to triads of variables or more. Future work may wish to combine the approach taken by Zhou et al. with the RPF methodology to expand understanding of various RPFs through considering mappings among the noise present in P_1 and P_2 to noise in the RPF.

Excitingly, RPF analysis for introspective or metacognitive judgments—including in these expanded variable sets—may also facilitate inference about similarities among individuals, stimuli, or tasks by strongly constraining process models—as briefly discussed above. Specifically, direct comparison of entire (multimodal) RPFs provides a sensitive and specific framework for evaluating the similarity of constellations or signatures of psychological processing across many linked behaviors (Taylor et al., 2022). If two tasks, domains, or individuals produce similar RPFs (especially when expanded to multiple variables), one might more easily build a single process model to characterize both, offering a path toward revealing unifying explanations of otherwise seemingly disparate domains.

Final Thoughts

In sum, the RPF framework holds great promise as a foundation for the next generation of psychophysics. To facilitate the exploration and use of this framework across disciplines and psychological areas of study, we encourage interested readers to make use of and extend the open-source RPF toolbox (https://github.com/CNClaboratory/RPF).

References

Adler, W. T., & Ma, W. J. (2018a). Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLOS Computational Biology*, *14*(11), Article e1006572. https://doi.org/10.1371/journal.pcbi.1006572

Adler, W. T., & Ma, W. J. (2018b). Limitations of proposed signatures of Bayesian confidence. *Neural Computation*, *30*(12), 3327–3354. https://doi.org/10.1162/neco_a_01141

Aitchison, L., Bang, D., Bahrami, B., & Latham, P. E. (2015). Doubly Bayesian analysis of confidence in perceptual decision-making. *PLOS Computational Biology*, *11*(10), Article e1004519. https://doi.org/10.1371/journal.pcbi.1004519

Baranski, J. V., & Petrusic, W. M. (1994). The calibration and resolution of confidence in perceptual judgments. *Perception & Psychophysics*, *55*(4), 412–428. https://doi.org/10.3758/BF03205299

Boldt, A., Sun, Y., & Desender, K. (2025). How disconfirmatory evidence shapes confidence in decision-making. *Communications Psychology*, *3*(1), 150. https://doi.org/10.1038/s44271-025-00325-3

Brown, G. S., & White, K. G. (2005). The optimal correction for estimating extreme discriminability. *Behavior Research Methods*, *37*(3), 436–449. https://doi.org/10.3758/BF03192712

Brown, R., Lau, H., & LeDoux, J. E. (2019). Understanding the higher-order approach to consciousness. *Trends in Cognitive Sciences*, *23*(9), 754–768. https://doi.org/10.1016/j.tics.2019.06.009

Burnham, K. P., & Anderson, D. (2002). *Model selection and multi-model inference: A practical information-theoretic approach*. Springer International Publishing.

Castaneda, O. G., Peters, M. A. K., & Maniscalco, B. (2026, January 19). *The relative psychometric function: A general analysis framework for relating psychological processes* [Data set]. Open Science Framework. https://doi.org/10.17605/OSF.IO/62RWK

Ceja, V., Ezzeldine, Y., & Peters, M. A. K. (2022). Models of confidence to facilitate engaging task designs. *2022 Conference on Cognitive Computational Neuroscience*. https://2022.ccneuro.org/view_paperd568.html?PaperNum=1150

Clarke, F., Birdsall, T., & Tanner, W. P. (1959). Two types of ROC curves and definitions of parameters. *The Journal of the Acoustical Society of America*, *31*, 629–630. https://doi.org/10.1121/1.1907764

de Rooij, M., & Weeda, W. (2020). Cross-validation: A method every psychologist should know. *Advances in Methods and Practices in Psychological Science*, *3*(2), 248–263. https://doi.org/10.1177/2515245919898466

Denison, R. N., Adler, W. T., Carrasco, M., & Ma, W. J. (2018). Humans incorporate attention-dependent uncertainty into perceptual decisions and confidence. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(43), 11090–11095. https://doi.org/10.1073/pnas.1717720115

Fechner, G. T. (1860). *Elemente der psychophysik* [Elements of psychophysics]. Breitkopf und Härtel. https://www.goodreads.com/book/show/35465569-elements-of-psychophysics

Fechner, G. T., Howes, D. H., & Boring, E. G. (1966). *Elements of psychophysics* (Vol. 1). Holt, Rinehart and Winston New York.

Fernández, A., & Carrasco, M. (2020). Extinguishing exogenous attention via transcranial magnetic stimulation. *Current Biology*, *30*(20), 4078–4084.e3. https://doi.org/10.1016/j.cub.2020.07.068

Fleming, S. M. (2017). HMeta-d: Hierarchical Bayesian estimation of metacognitive efficiency from confidence ratings. *Neuroscience of Consciousness*, *2017*(1), Article nix007. https://doi.org/10.1093/nc/nix007

Fleming, S. M. (2023). Metacognitive psychophysics in humans, animals, and AI. *Journal of Consciousness Studies*, *30*(9–10), 113–128. https://doi.org/10.53765/20512201.30.9.113

Fung, H., & Rahnev, D. (2024). Different stimulus manipulations produce dissociable confidence–accuracy relationships. *Journal of Vision*, *24*(10), Article 399. https://doi.org/10.1167/jov.24.10.399

Fung, H., Shekhar, M., Rausch, M., & Rahnev, D. (2025). *Similarities and differences in the effects of different stimulus manipulations on accuracy and confidence*. https://osf.io/b7kru_v1/

Galvin, S. J., Podd, J. V., Drga, V., & Whitmore, J. (2003). Type 2 tasks in the theory of signal detectability: Discrimination between correct and incorrect decisions. *Psychonomic Bulletin & Review*, *10*(4), 843–876. https://doi.org/10.3758/BF03196546

Hanning, N. M., Fernández, A., & Carrasco, M. (2023). Dissociable roles of human frontal eye fields and early visual cortex in presaccadic attention. *Nature Communications*, *14*(1), Article 5381. https://doi.org/10.1038/s41467-023-40678-z

Hausman, J. A., Newey, W. K., & Powell, J. L. (1995). Nonlinear errors in variables estimation of some Engel curves. *Journal of Econometrics*, *65*(1), 205–233. https://doi.org/10.1016/0304-4076(94)01602-V

Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d′. *Behavior Research Methods, Instruments, & Computers*, *27*(1), 46–51. https://doi.org/10.3758/BF03203619

Herrmann, K., Montaser-Kouhsari, L., Carrasco, M., & Heeger, D. J. (2010). When size matters: Attention affects performance by contrast or response gain. *Nature Neuroscience*, *13*(12), 1554–1559. https://doi.org/10.1038/nn.2669

Huang, L., & Dobkins, K. R. (2005). Attentional effects on contrast discrimination in humans: Evidence for both contrast gain and response gain. *Vision Research*, *45*(9), 1201–1212. https://doi.org/10.1016/j.visres.2004.10.024

Huang, Z., Meng, S., & Ye, Z. (2023). Effective estimation of nonlinear errors-in-variables models. *Communications in Statistics: Simulation and Computation*, *54*(7), 2345–2363. https://doi.org/10.1080/03610918.2023.2219434

Jigo, M., & Carrasco, M. (2020). Differential impact of exogenous and endogenous attention on the contrast sensitivity function across eccentricity. *Journal of Vision*, *20*(6), Article 11. https://doi.org/10.1167/jov.20.6.11

Kammerer, F., & Frankish, K. (2023). What forms could introspective systems take? A research programme. *Journal of Consciousness Studies*, *30*(9–10), 13–48. https://doi.org/10.53765/20512201.30.9.013

Khoudary, A., Bornstein, A. M., & Peters, M. A. K. (2025). Investigating implicit and explicit expectations in perceptual decision making. *Proceedings of the Cognitive Science Society*, *47*. https://escholarship.org/uc/item/7dm8x2wj

Khoudary, A., Peters, M. A. K., & Bornstein, A. M. (2025). Reasoning goals and representational decisions in computational cognitive neuroscience: Lessons from the drift diffusion model. *The European Journal of Neuroscience*, *61*(7), Article e70098. https://doi.org/10.1111/ejn.70098

Kiani, R., Corthell, L., & Shadlen, M. N. (2014). Choice certainty is informed by both evidence and decision time. *Neuron*, *84*(6), 1329–1342. https://doi.org/10.1016/j.neuron.2014.12.015

Kingdom, F. A. A., & Prins, N. (2016). *Psychophysics: A practical introduction* (2nd ed.). Academic Press.

Koizumi, A., Maniscalco, B., & Lau, H. (2015). Does perceptual confidence facilitate cognitive control? *Attention, Perception & Psychophysics*, *77*(4), 1295–1306. https://doi.org/10.3758/s13414-015-0843-3

Lau, H. (2008). Are we studying consciousness yet? In L. Weiskrantz & M. Davies (Eds.), *Frontiers of consciousness* (pp. 2008–2245). Oxford University Press.

Lau, H., & Passingham, R. E. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proceedings of the National Academy of Sciences*, *103*(49), 18763–18768. https://doi.org/10.1073/pnas.0607716103

Lee, J. L., Denison, R., & Ma, W. J. (2023). Challenging the fixed-criterion model of perceptual decision-making. *Neuroscience of Consciousness*, *2023*(1), Article niad010. https://doi.org/10.1093/nc/niad010

Li, H.-H., Pan, J., & Carrasco, M. (2021). Different computations underlie overt presaccadic and covert spatial attention. *Nature Human Behaviour*, *5*(10), 1418–1431. https://doi.org/10.1038/s41562-021-01099-4

Li, T. (2002). Robust and consistent estimation of nonlinear errors-in-variables models. *Journal of Econometrics*, *110*(1), 1–26. https://doi.org/10.1016/S0304-4076(02)00120-3

Ling, S., & Carrasco, M. (2006). Sustained and transient covert attention enhance the signal via different contrast response functions. *Vision Research*, *46*(8–9), 1210–1220. https://doi.org/10.1016/j.visres.2005.05.008

Locke, S. M., Mamassian, P., & Landy, M. S. (2020). Performance monitoring for sensorimotor confidence: A visuomotor tracking study. *Cognition*, *205*, Article 104396. https://doi.org/10.1016/j.cognition.2020.104396

Macmillan, N. A., & Creelman, C. D. (2004). *Detection theory: A user's guide*. Taylor & Francis.

Maniscalco, B., Castaneda, O. G., Odegaard, B., Morales, J., Rajananda, S., Denison, R., & Peters, M. A. K. (2020). The metaperceptual function: Exploring dissociations between confidence and task performance with type 2 psychometric curves. *Earliest post: 2020, July 6*. https://doi.org/10.31234/osf.io/5qrjn

Maniscalco, B., Charles, L., & Peters, M. A. K. (2024). Optimal metacognitive decision strategies in signal detection theory. *Psychonomic Bulletin & Review*, *32*(3), 1041–1069. https://doi.org/10.3758/s13423-024-02510-7

Maniscalco, B., & Lau, H. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition*, *21*(1), 422–430. https://doi.org/10.1016/j.concog.2011.09.021

Maniscalco, B., & Lau, H. (2014). Signal detection theory analysis of type 1 and type 2 data: Meta-d′, response-specific meta-d′, and the unequal variance SDT model. In S. M. Fleming & C. D. Frith (Eds.), *The cognitive neuroscience of metacognition* (pp. 25–66). Springer International Publishing. https://doi.org/10.1007/978-3-642-45190-4_3

Maniscalco, B., Odegaard, B., Grimaldi, P., Cho, S. H., Basso, M. A., Lau, H., & Peters, M. A. K. (2021). Tuned normalization in perceptual decision-making circuits can explain seemingly suboptimal confidence behavior. *PLOS Computational Biology*, *17*(3), Article e1008779. https://doi.org/10.1371/journal.pcbi.1008779

Maniscalco, B., & Peters, M. A. K. (2026). *RPF toolbox* [Code]. Github. https://github.com/CNClaboratory/RPF

Maniscalco, B., Peters, M. A. K., & Lau, H. (2016). Heuristic use of perceptual evidence leads to dissociation between performance and metacognitive sensitivity. *Attention, Perception & Psychophysics*, *78*(3), 923–937. https://doi.org/10.3758/s13414-016-1059-x

Michel, M., Beck, D., Block, N., Blumenfeld, H., Brown, R., Carmel, D., Carrasco, M., Chirimuuta, M., Chun, M., Cleeremans, A., Dehaene, S., Fleming, S. M., Frith, C., Haggard, P., He, B. J., Heyes, C., Goodale, M. A., Irvine, L., Kawato, M., … Yoshida, M. (2019). Opportunities and challenges for a maturing science of consciousness. *Nature Human Behaviour*, *3*(2), 104–107. https://doi.org/10.1038/s41562-019-0531-8

Miller, J. (1996). The sampling distribution of d′. *Perception & Psychophysics*, *58*(1), 65–72. https://doi.org/10.3758/BF03205476

Miyoshi, K., & Lau, H. (2020). A decision-congruent heuristic gives superior metacognitive sensitivity under realistic variance assumptions. *Psychological Review*, *127*(5), 655–671. https://doi.org/10.1037/rev0000184

Morales, J. (2024). Introspection is signal detection. *The British Journal for the Philosophy of Science*, *75*(1), 99–126. https://doi.org/10.1086/715184

Morales, J., Odegaard, B., & Maniscalco, B. (2022). The neural substrates of conscious perception without performance confounds. In F. De Brigard & W. Sinnott-Armstrong (Eds.), *Neuroscience and philosophy* (pp. 285–323). MIT Press.

Odegaard, B., Chang, M. Y., Lau, H., & Cheung, S.-H. (2018). Inflation versus filling-in: Why we feel we see more than we actually do in peripheral vision. *Philosophical Transactions of the Royal Society of London: Series B, Biological Sciences*, *373*(1755), Article 20170345. https://doi.org/10.1098/rstb.2017.0345

Odegaard, B., Grimaldi, P., Cho, S. H., Peters, M. A. K., Lau, H., & Basso, M. A. (2018). Superior colliculus neuronal ensemble activity signals optimal rather than subjective confidence. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(7), E1588–E1597. https://doi.org/10.1073/pnas.1711628115

Overgaard, M., & Mogensen, J. (2017). An integrative view on consciousness and introspection. *Review of Philosophy and Psychology*, *8*(1), 129–141. https://doi.org/10.1007/s13164-016-0303-6

Peters, M. A. K. (2022). Towards characterizing the canonical computations generating phenomenal experience. *Neuroscience and Biobehavioral Reviews*, *142*, Article 104903. https://doi.org/10.1016/j.neubiorev.2022.104903

Peters, M. A. K. (2024). Introspective psychophysics for the study of subjective experience. *Cerebral Cortex*, *35*(1), 49–57. https://doi.org/10.1093/cercor/bhae455

Peters, M. A. K. (2025). Introspective psychophysics for the study of subjective experience. *Cerebral Cortex*, *35*(1), 49–57. https://doi.org/10.1093/cercor/bhae455

Peters, M. A. K., Fesi, J., Amendi, N., Knotts, J. D., Lau, H., & Ro, T. (2017). Transcranial magnetic stimulation to visual cortex induces suboptimal introspection. *Cortex: A Journal Devoted to the Study of the Nervous System and Behavior*, *93*, 119–132. https://doi.org/10.1016/j.cortex.2017.05.017

Peters, M. A. K., Kentridge, R. W., Phillips, I., & Block, N. (2017). Does unconscious perception really exist? Continuing the ASSC20 debate. *Neuroscience of Consciousness*, *2017*(1), Article nix015. https://doi.org/10.1093/nc/nix015

Peters, M. A. K., & Lau, H. (2015). Human observers have optimal introspective access to perceptual processes even for visually masked stimuli. *eLife*, *4*, Article e09651. https://doi.org/10.7554/eLife.09651

Peters, M. A. K., Thesen, T., Ko, Y. D., Maniscalco, B., Carlson, C., Davidson, M., Doyle, W., Kuzniecky, R., Devinsky, O., Halgren, E., & Lau, H. (2017). Perceptual confidence neglects decision-incongruent evidence in the brain. *Nature Human Behaviour*, *1*(7), Article 0139. https://doi.org/10.1038/s41562-017-0139

Pollack, I., & Hsieh, R. (1969). Sampling variability of the area under the ROC-curve and of d'e. *Psychological Bulletin*, *71*(3), 161–173. https://doi.org/10.1037/h0026862

Prins, N., & Kingdom, F. A. A. (2018). Applying the model-comparison approach to test specific research hypotheses in psychophysical research using the Palamedes toolbox. *Frontiers in Psychology*, *9*, Article 1250. https://doi.org/10.3389/fpsyg.2018.01250

Rahnev, D., Balsdon, T., Charles, L., de Gardelle, V., Denison, R., Desender, K., Faivre, N., Filevich, E., Fleming, S. M., Jehee, J., Lau, H., Lee, A. L. F., Locke, S. M., Mamassian, P., Odegaard, B., Peters, M., Reyes, G., Rouault, M., Sackur, J. … Zylberberg, A. (2022). Consensus goals in the field of visual metacognition. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, *17*(6), 1746–1765. https://doi.org/10.1177/17456916221075615

Rahnev, D., & Denison, R. N. (2018). Suboptimality in perceptual decision making. *The Behavioral and Brain Sciences*, *41*, Article e223. https://doi.org/10.1017/s0140525x18000936

Rahnev, D., Desender, K., Lee, A. L. F., Adler, W. T., Aguilar-Lleyda, D., Akdoğan, B., Arbuzova, P., Atlas, L. Y., Balcı, F., Bang, J. W., Bègue, I., Birney, D. P., Brady, T. F., Calder-Travis, J., Chetverikov, A., Clark, T. K., Davranche, K., Denison, R. N., Dildine, T. C., … Zylberberg, A. (2020). The confidence database. *Nature Human Behaviour*, *4*(3), 317–325. https://doi.org/10.1038/s41562-019-0813-1

Rahnev, D., Maniscalco, B., Graves, T., Huang, E., de Lange, F. P., & Lau, H. (2011). Attention induces conservative subjective biases in visual perception. *Nature Neuroscience*, *14*(12), 1513–1515. https://doi.org/10.1038/nn.2948

Rahnev, D., Maniscalco, B., Luber, B., Lau, H., & Lisanby, S. H. (2012). Direct injection of noise to the visual cortex decreases accuracy but increases decision confidence. *Journal of Neurophysiology*, *107*, 1556–1563. https://doi.org/10.1152/jn.00985.2011

Rausch, M., Hellmann, S., & Zehetleitner, M. (2017). Confidence in masked orientation judgments is informed by both evidence and visibility. *Attention, Perception & Psychophysics*, *80*(1), 134–154. https://doi.org/10.3758/s13414-017-1431-5

Rollwage, M., Loosen, A., Hauser, T. U., Moran, R., Dolan, R. J., & Fleming, S. M. (2020). Confidence drives a neural confirmation bias.

Nature Communications, *11*(1), Article 2634. https://doi.org/10.1038/s41467-020-16278-6

Rong, Y., & Peters, M. A. K. (2023). Toward "computational-rationality" approaches to arbitrating models of cognition: A case study using perceptual metacognition. *Open Mind*, *7*, 652–674. https://doi.org/10.1162/opmi_a_00100

Rouault, M., Seow, T., Gillan, C. M., & Fleming, S. M. (2018). Psychiatric symptom dimensions are associated with dissociable shifts in metacognition but not task performance. *Biological Psychiatry*, *84*(6), 443–451. https://doi.org/10.1016/j.biopsych.2017.12.017

Rounis, E., Maniscalco, B., Rothwell, J. C., Passingham, R. E., & Lau, H. (2010). Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. *Cognitive Neuroscience*, *1*(3), 165–175. https://doi.org/10.1080/17588921003632529

Ruby, E., Maniscalco, B., & Peters, M. A. K. (2018). On a "failed" attempt to manipulate visual metacognition with transcranial magnetic stimulation to prefrontal cortex. *Consciousness and Cognition*, *62*, 34–41. https://doi.org/10.1016/j.concog.2018.04.009

Samaha, J., Barrett, J. J., Sheldon, A. D., LaRocque, J. J., & Postle, B. R. (2016). Dissociating perceptual confidence from discrimination accuracy reveals no influence of metacognitive awareness on working memory. *Frontiers in Psychology*, *7*, Article 851. https://doi.org/10.3389/fpsyg.2016.00851

Samaha, J., & Denison, R. (2022). The positive evidence bias in perceptual confidence is unlikely post-decisional. *Neuroscience of Consciousness*, *2022*(1), Article niac010. https://doi.org/10.1093/nc/niac010

Samaha, J., Switzky, M., & Postle, B. R. (2019). Confidence boosts serial dependence in orientation estimation. *Journal of Vision*, *19*(4), Article 25. https://doi.org/10.1167/19.4.25

Shekhar, M., & Rahnev, D. (2024). How do humans give confidence? A comprehensive comparison of process models of perceptual metacognition. *Journal of Experimental Psychology: General*, *153*(3), 656–688. https://doi.org/10.1037/xge0001524

Stolyarova, A., Rakhshan, M., Hart, E. E., O'Dell, T. J., Peters, M. A. K., Lau, H., Soltani, A., & Izquierdo, A. (2019). Contributions of anterior cingulate cortex and basolateral amygdala to decision confidence and learning under uncertainty. *Nature Communications*, *10*(1), Article 4704. https://doi.org/10.1038/s41467-019-12725-1

Taylor, A. H., Bastos, A. P. M., Brown, R. L., & Allen, C. (2022). The signature-testing approach to mapping biological and artificial intelligences. *Trends in Cognitive Sciences*, *26*(9), 738–750. https://doi.org/10.1016/j.tics.2022.06.002

Vrieze, S. I. (2012). Model selection and psychological theory: A discussion of the differences between the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). *Psychological Methods*, *17*(2), 228–243. https://doi.org/10.1037/a0027127

Winter, C. J., & Peters, M. A. K. (2022). Variance misperception under skewed empirical noise statistics explains overconfidence in the visual periphery. *Attention, Perception & Psychophysics*, *84*(1), 161–178. https://doi.org/10.3758/s13414-021-02358-2

Wolter, K. M., & Fuller, W. A. (1982). Estimation of nonlinear errors-in-variables models. *Annals of Statistics*, *10*(2), 539–548. https://doi.org/10.1214/aos/1176345794

Xue, K., Shekhar, M., & Rahnev, D. (2024a). A novel behavioral paradigm reveals the nature of confidence computation in multi-alternative perceptual decision making. *Research Square*, Article rs.3.rs-5510856. https://doi.org/10.21203/rs.3.rs-5510856/v1

Xue, K., Shekhar, M., & Rahnev, D. (2024b). Challenging the Bayesian confidence hypothesis in perceptual decision-making. *Proceedings of the National Academy of Sciences of the United States of America*, *121*(48), Article e2410487121. https://doi.org/10.1073/pnas.2410487121

Zhou, J., Duong, L. R., & Simoncelli, E. P. (2024). A unified framework for perceived magnitude and discriminability of sensory stimuli. *Proceedings of the National Academy of Sciences of the United States of America*, *121*(25), Article e2312293121. https://doi.org/10.1073/pnas.2312293121

Zylberberg, A., Barttfeld, P., & Sigman, M. (2012). The construction of confidence in a perceptual decision. *Frontiers in Integrative Neuroscience*, *6*, Article 79. https://doi.org/10.3389/fnint.2012.00079

Zylberberg, A., Fetsch, C. R., Shadlen, M. N., & Frank, M. J. (2016). The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision. *eLife*, *5*, Article e17688. https://doi.org/10.7554/eLife.17688

Zylberberg, A., Roelfsema, P. R., & Sigman, M. (2014). Variance misperception explains illusions of confidence in simple perceptual decisions. *Consciousness and Cognition*, *27C*, 246–253. https://doi.org/10.1016/j.concog.2014.05.012

Received November 29, 2024
Revision received February 6, 2026
Accepted February 18, 2026 ■