Manuscript Title: A Mathematical Model of Person Judgment. Part 1: Cue Emergence

Running Head: A Mathematical Model of Person Judgment: Cue Emergence

Version Date: 18.03.2023

Fresh vs. Resubmission

If this is the first submission round, use "Fresh". If this is a resubmission, use "R" and the round (e.g., 1, 2, 3).

Status: R1

Paper Category and Type

Please insert an "X" into one of the 14 gray-shaded cells to indicate the primary category and, if applicable, type of the paper. White spaces cannot be selected. Only one "X" should be indicated.

		Primary Paper Categories				
		Theory	Methodology	Empirical	Applied	Comment
No Specific Paper Type		X				
Specific Paper Types (if applicable)	State of the Art Review					
	Tutorial					
	Projects & Data					
	Insights & Ideas					
	Cumulative Blitz Report					
	Registered Report					
	Replication					
	Meta-analysis					

Invited Paper

Choose "Yes" if an editor has specifically invited this paper. If not, please choose "No".

No

OPTIONAL: Paper Collection (Theme Bundle or Controversy Exchange)

If this paper should be considered as a part of a collection, then please state the intended collection. Paper collections can be Theme Bundles or Controversy Exchanges. Please indicate which one applies.

Type:

Title:

OPTIONAL: Authors of the Paper

If you wish to remain anonymous in the peer-review process, please leave the space below blank or as is.

Daniel Leising, Technische Universität Dresden;
Rene L. Schilling, Technische Universität Dresden

Note: You also need to upload a completed Contributorship Template to the online submission system. That form will not be sent out for peer-review.

English Abstract

Max. 150 words; no formatting – plain text only

We present the first part of a fully parametrized mathematical model of person judgment, in an attempt to streamline and better organize theory in this research area. This first part focuses on the emergence of substantive information (“cues”) about target persons that perceivers may use in forming judgments. It incorporates time as a continuous variable, and accounts for the effects of targets, situations and their interactions. It also accounts for randomness overall, between-target differences in randomness, and overlap between observation intervals. We discuss how all of this influences the similarity of targets’ substance levels in any two observation intervals, enabling predictions regarding retest reliability, inter-rater agreement, cross-situational consistency, and predictive validity. We also explain how the model relates to existing theories, and to key concepts from the person judgment literature.

English Keywords

Max. 8 keywords (all in lower-case, except for names), delimited by semicolons

personality; judgment; cue; modelling; formalization; mathematics; theory

OPTIONAL: Abstract in another language

Max. 150 words; no formatting – plain text only; leave blank or as is if not relevant

Abstract goes here

OPTIONAL: Keywords in another language (same as Abstract above)

Max. 8 keywords, delimited by semicolons; leave blank or as is if not relevant

keyword; keyword; keyword; keyword; keyword; keyword; keyword; keyword

Language: State the language of the additional abstract and keywords here

Key Insights

3-5 bullet points as summaries, with max. 10 words each

- Model formalization may help integrate previously unconnected strands of theory
- Formal theories are more exact and parsimonious than narrative ones
- We demonstrate this by formally modelling the emergence of “cues”
- A single formula enables a broad range of predictions
- The model makes the enormous complexity of the topic visible

Relevance Statement

Max. 250 characters, including spaces

Please do not just summarize the paper, but highlight its relevance, interestingness, and unique merits.

Formalization of psychological theories is often called for, but examples are rare in personality science. We present a fully parametrized model of cue emergence, and use it to highlight advantages of formal models as compared to “narrative” ones.

OPTIONAL: Plain-Language Summary

500-750 words. Leave blank or as is if not relevant.

A non-technical summary for a lay audience aims to increase the dissemination of your findings beyond disciplines and academia more generally. It should avoid jargon as much as possible and communicate main findings and their practical and policy-related implications in clear and simple language.

Please use the following sub-headings (that can be modified for theoretical and review articles) to structure the non-technical summary:

1. What is the study's background?
2. Why was this study done?
3. What did the researchers do and find?
4. What do these findings mean?

URLs to Supplements

Please insert working web addresses where supplements can be accessed.

Type	URL
Pre-registration	n/a
Data	n/a
Codebook	n/a
Documents, stimuli, images, etc.	n/a
Scripts, code, etc.	n/a
Tests, assessment instruments	n/a
Additional analyses, tables, etc.	n/a
Preprints	Technical Supplement: https://psyarxiv.com/myua4
Other supplements	n/a

Note: Insert URLs if applicable (e.g., in empirical articles). If not applicable, simply leave blank.

If certain supplements are referenced or linked to in the manuscript (e.g., data, codebook, scripts, etc.), then please use the exact same URL(s) as inserted into the table above. In other words, any and each URL included in the manuscript also needs to be found in this table.

Supplements are subject to pre-publication peer-review. If you wish to remain anonymous, please ensure that no identifying information is contained in the URLs or their contents.

Before a paper has been conditionally accepted, any accessible repository can be used to store the supplements. Thus, at the time of submitting the paper, no supplements need to be uploaded to the online submission system; linking to them suffices for the time being. However, after a paper has been conditionally accepted, all supplements need to be uploaded in a zip-folder (along with other documents and forms) to the journal submission system. For the time being, PsychOpen staff of the publisher will then upload your supplements to PsychOpen's PsychArchives repository (<https://www.psycharchives.org/>) and automatically add the new respective URLs into the published paper. Once your paper has been conditionally accepted, you will receive more information on this process in the editorial decision letter.

A Mathematical Model of Person Judgment

Part 1: Cue Emergence

Daniel Leising & René L. Schilling

Technische Universität Dresden

Revised Version After First Round of Peer-Review

March 18th 2023

Corresponding author:

Daniel Leising,

Technische Universität Dresden

01062 Dresden, Germany

Daniel.Leising@tu-dresden.de

Introduction

In this paper, we present the first part of an integrative mathematical model of person judgment. Our first major aim is to help streamline and better organize the vast literature in this field, by connecting previously unconnected strands of theory with one another, and by formalizing some key concepts that, so far, have mainly featured in the literature in „narrative“ (i.e. natural language) form. Our second major aim is to demonstrate that, despite being fairly challenging, such formalization is actually possible and bears great potential for making personality research a more cumulative science. Among the gains to be expected from stricter formalization are a higher level of exactness and a lower level of redundancy.

In terms of content, most of the parameters in our model are not new. What is new is the strict mathematical form in which we present them. In fact, we believe that many of our colleagues in the field at least implicitly subscribe to a model that is quite similar to what we present here. What we hope to be able to add is a greater level of explicitness, conceptual clarity and parsimony. We present this model in two different forms: The present paper introduces and explains the core components of the model in a manner that should be accessible to most readers with a background in personality research. Throughout it, we try to elucidate how the model's parameters relate to key concepts from the personality literature. In addition, we also offer a full mathematical account as an online Technical Supplement (Leising & Schilling, 2023). This latter version is more abstract, goes into greater mathematical detail, and only occasionally touches on specific concepts from the personality literature. Its purpose is to demonstrate that the analysis overall is mathematically sound and that all components of the model are connected within an exact, coherent framework, leaving no room for ambiguities in terms of meaning. This is important because conceptual ambiguities are one of the main problems that psychological theory is grappling with to this day (Frankenhuis, Panchanathan & Smaldino, 2022; Glöckner & Betsch, 2011).

To make it as easy as possible to read the paper and the Technical Supplement side-by-side, we reference mathematical formulas by the same consecutive numbering in both documents. However, most of the more specific formulas from the Technical Supplement do *not* appear in the present paper.

Most psychologists are only just beginning to discover and fully appreciate the advantages of formal scientific models over the “narrative” ones that still dominate the field: Formal models are able to capture greater complexity with greater exactness. They are better specified and thus better falsifiable. They enable a clearer organization of how research findings relate to each another, and an easier identification of contradictions, circularity, redundancy, specification gaps, and compatibility between different models (Borgstede & Eggert, 2023a, 2023b; Leising, Thielmann, Glöckner, Gärtner & Schönbrodt, 2022a; Rodgers, 2010). For all of these reasons, formal models at least have the potential to be more conducive to cumulative knowledge growth than narrative ones, which is why a stricter formalization of psychological theory has been called for many times (e.g., Glöckner & Betsch, 2011; Oberauer & Lewandowski, 2019; Robinaugh, Haslbeck, Ryan, Fried & Waldorp, 2021; Smaldino, 2020).

However, formal modelling has not yet been established as a routine scientific practice in personality psychology, which we think is due to three factors: First, a lack of training. Most psychologists (including the first author of the present paper) were never properly trained in mathematical modelling and thus lack the necessary expertise to do the hands-on work of developing a full mathematical model (Rodgers, 2010). Second, due its greater demands in terms of exactness, mathematical modelling is significantly more difficult and time-consuming as compared to

the more common “narrative” approach. The present paper, for example, reflects the outcome of several years of intense collaborative work. Third, proper formal modelling is not only not incentivized in psychology today, but there are actually “good” strategic reasons for shying away from it (Frankenhuis et al., 2022). For example, the more exact a model becomes, the more easily it will be refuted.

Despite these impediments, there *have* of been several prior attempts at formalizing theoretical thinking on person judgment processes. Among the most ambitious and successful of these are David Kenny’s Weighted Average (1994) and PERSON (2004, 2019) models. Brunswik’s (1956) Lens Model has also been highly influential as a conceptual framework in this field (Back & Nestler, 2016; Funder, 1995). This paper presents the first part of a model that is supposed to jointly capture most of the key components from these earlier models and integrate them in a coherent, formalized fashion.

Note, however, that this first part of the model only deals with the substantive bases of person judgments which are called “cues” in the Lens Model. However, the cues in our model do *not* necessarily relate to some joint, underlying cause variable. They simply represent “the facts” about targets that perceivers *may* or may not base their judgments on. What perceivers actually do with that information will be the subject of a follow-up paper. Note that our model is supposed to capture other-judgments as well as self-judgments. So, a cue may either be something that is observable from the outside (e.g. smiling), or something that is only perceivable by the target (e.g., felt pain).

We account for the emergence of cues in a way that is compatible with Classical Test Theory (CTT; Novick, 1966; Lord & Novick, 1968). Our model goes beyond CTT, however, in that it decomposes the unobservable “true” (or better: noise-free) score into three sources of variance: target, situation, and their interaction. Thus, it captures the key components of psychology’s person-situation-debate (Kenrick & Funder, 1988; Mischel, 1968; Vazire & Sherman, 2017) and also overlaps significantly with Latent-State-Trait Theory (LST; Hintz, Geiser & Shiffman, 2019; Steyer, Ferring & Schmitt, 1992; Steyer, Schmitt & Eid, 1999; Steyer, Mayer, Geiser & Cole, 2015). A brief model comparison is provided further below.

With the current presentation, we hope to achieve a variety of goals: (a) show that an integration of key concepts from the person perception literature by way of formalization is *possible*, (b) show that such an integration is also *worthwhile*, because it helps reduce theoretical ambiguity and redundancy, while at the same time enabling more precise predictions across a wide range of phenomena. We also highlight the crucial difference between the “substance” that is being described and the ways in which perceivers may describe it (e.g., by means of person-descriptive terms from the natural language), because this difference is still being overlooked by too many contributions in the field.

We would like to encourage our readers to use our parametrization in their own work, in order to foster the emergence of a more cumulative personality science (in which different researchers describe the same things using the same labels). We also hope to be able to (c) showcase the enormous amount of complexity permeating this field of study, which remains hidden as long as proper formalization is avoided, and thus (d) illustrate a likely reason why proper formalization has been avoided so far by most researchers working on these issues. Ultimately, all of this is intended to help increase the field’s appreciation for this crucial yet neglected kind of work.

Going Through the Model Components Step by Step

Basic Considerations

This first part of the model accounts for the emergence of cues on which perceivers may base their judgments of targets. We translate the concept of a cue into the *average level of some substance variable for a given interval*. By “substance”, we simply mean an actual characteristic that the target has – at least temporarily - irrespective of whether or how that characteristic is ever perceived or used in making a judgment.

In line with CTT, we model a target’s level on a substance variable as comprising a systematic component (θ) and random noise (W). The systematic component is modelled as depending on the target person (τ), the situational circumstances (σ) at the given point in time (t), and the interaction between target and situation ($\tau \times \sigma$). We also assume that targets may differ in how much random noise their substance levels typically contain ($\gamma(\tau)$).

As a general rule, we use a tilde (as in $\tilde{\theta}$) to indicate all kinds of operations (i.e., averaging, but also taking variances and covariances) in which the variation “is in” targets (τ) (i.e., the targets are the source of variation). We use a bar (as in $\bar{S}$) to symbolize averaging over time (t). In most of what follows, such averaging is bound to a specific interval (as in $\bar{S}(\tau, I, \omega)$), which stands for a target’s observable (but noisy) substance level in the time-interval I . In fact, this is the primary outcome of our modelling effort – a cue. The ω in the argument represents the presence of random influences.

The different types of operations may also be combined, as in $\tilde{\bar{\theta}}(I)$ which stands for the average target’s noise-free substance level across interval I . Once a parameter is subjected to such an operation, the component of its argument across which the operation takes place necessarily vanishes from its argument. For example, the argument of $\tilde{\bar{\theta}}(I)$ does not contain τ anymore because it is the outcome of averaging $\bar{\theta}(\tau, I)$ across τ .

The Model’s Components

In this section, we introduce the model’s components and their relationships with one another step by step. Table 1 in the Appendix lists all parameters along with their respective meanings. To facilitate comprehension, we introduce a specific example: Action Unit 12 (AU12) activation. In Ekman and Friesen’s (1978) Facial Action Coding System, activation of AU 12 (mainly reflecting activation of the *Musculus Zygomaticus Major*) causes the “cheek dimples” that are commonly interpreted as “smiling”. In the present context, however, it is very important to keep in mind that we are only interested in the actual, substantive level of activation of this part of the face – whether this activation is actually perceived and possibly even interpreted by anyone is irrelevant for now.

In accordance with CTT, we distinguish between the observable level (S) of a substance variable and that variable’s “true” level (θ). The former always contains at least some perturbations caused by random noise whereas the latter does not. Because the label “true” sometimes leads to misunderstandings regarding the meaning of the θ parameter, we will call it “noise-free” instead.

Figure 1. Two targets' (τ and τ') noise-free substance levels (blue: $\theta(\tau, \sigma(t))$; red: $\theta(\tau', \sigma(t))$) in five intervals (I_1 to I_5), displayed as colored solid lines. For simplicity, each interval reflects the influence of just one situation ($\sigma(t)$). The average target's noise-free level in each interval ($\tilde{\theta}(\sigma(t))$) is represented by the solid black lines. Random changes in substance levels from moment to moment (blue: $dS(\tau, \sigma(t), \omega)$; red: $dS(\tau', \sigma(t), \omega)$) are represented by the wiggly scatter around the noise-free levels.

Figure 1 is to illustrate some of the key concepts. We look at the AU12 activations of two target persons as they pass through a number of different situations. The first target (τ) is Tessa, the second target (τ') is Trudy. Both Tessa and Trudy are university professors. In Figure 1, Tessa's AU12 activation levels are displayed in blue whereas Trudy's are displayed in red. The time (t) axis is divided into five intervals (I_1 to I_5) schematically representing five situations. For simplicity, we assume that the situation remains the same within each of these intervals (i.e., we have "one-situation intervals"). The first interval is the morning where both Tessa and Trudy get ready to go to work. The second interval is a lecture that each of them is giving. The third situation is a funeral that each of them attends. The fourth interval is another lecture. The fifth interval is a lunch break. Note that the model treats the targets' substance levels as being independent. For example, the two targets may have identical schedules, but still live and work in different cities.

The solid lines in Figure 1 represent the noise-free substance levels of the two targets, and of the average target, in each of the five intervals. They are also captured by parameters $\theta(\tau, \sigma(t))$ for Tessa (blue), $\theta(\tau', \sigma(t))$ for Trudy (red), and $\tilde{\theta}(\sigma(t))$ for the average target (black). Note that all of these expressions depend directly on $\sigma(t)$, and not on the respective interval. This is because the only thing that matters about the intervals are the respective situational circumstances applying to them. For example, Interval 2 and 4 are identical in terms of situation ("lecture"), and thus produce the exact same noise-free scores for each individual target, and for the average target.

In Figure 1 we see that Trudy (τ' , red) has a higher noise-free level of AU12 activation than Tessa (τ , blue) in the morning (I_1). This reverses, however, when they both give their first lecture (I_2). Obviously, lectures are situations in which Tessa's AU12 gets into overdrive, whereas Trudy's level of AU12 activation remains basically unchanged compared to when she was at home in the morning. The funeral (I_3) brings a drastic change: Trudy's AU12 activation goes down a bit, but not remotely as much as Tessa's, whose face now shows almost no AE 12 activation anymore. During the second

lecture (I_4), the pattern is the exact same as in Interval 2. Finally, the pattern during the lunch break (I_5) largely resembles that from the first interval, only with both Tessa and Trudy showing slightly lower AU12 activation.

One of the key tenets of personality psychology is that target persons may differ from one another “in general”, that is, across some kind of longer, relevant time frame (e.g., their entire lifespan). This time frame is represented by another parameter (I_{Max}) which does not appear in Figure 1. It stands for the reality that one is trying to model, and from which one is drawing more limited time - samples when gathering data. Most important, I_{Max} is assumed to contain the typical setup of situations, including their typical lengths, of the reality that one is interested in. Each target has an individual time-average on the substance variable for this hypothetical interval ($\bar{\theta}(\tau, I_{Max})$ for Tessa and $\bar{\theta}(\tau', I_{Max})$ for Trudy) and there is also an average across targets (“Grand Mean”) for this interval ($\bar{\theta}(I_{Max})$). To avoid clutter, these parameters are not included in Figure 1. They would be displayed as straight lines lying parallel to the x-axis, with the Grand Mean lying exactly in between the individual means of the two targets. These parameters will come into play when we describe the observable cue level as a sum of various “effects” (see next section).

Up to this point, we were only talking about the targets’ *noise-free* substance levels, which may never actually be observed. In Figure 1, the wiggly scatter around the noise-free levels represents random changes of the targets’ substance levels from moment to moment. The model parameters for these would be $dS(\tau, \sigma(t), \omega)$ for Tessa (blue) and $dS(\tau', \sigma(t), \omega)$ for Trudy (red). The arguments here contain an ω , which stands for the presence of a random influence. The use of dS in this context implies maximum unpredictability (see below). Note that, despite the noise being equally unpredictable for the two targets, its *variance* is obviously larger for Trudy than it is for Tessa, as indicated by the different amplitude of the respective scatters. In our model, the target-specific variance of the noise is captured by the parameter $\gamma^2(\tau)$ for Tessa (blue) and $\gamma^2(\tau')$ for Trudy (red).

The key explanandum of part 1 of our model are the targets’ average observable levels on the substance variable for a given interval (“cues”). In the model, they are represented by the parameter $\bar{S}(\tau, I, \omega)$ (see next section). Both Tessa and Trudy would have an average observable level of AU 12 activation for each of the five intervals (i.e., ten such levels altogether). To avoid clutter, we do not display these average levels separately in Figure 1 but introduce them in detail in the next section. Just as in CTT, the targets’ average substance levels are assumed to be the sum of their respective noise-free levels and random noise.

Putting the Components Together

The full mathematical account of our model may be found in the Technical Supplement (Leising & Schilling, 2023). In the following, we will only highlight those parts of it that are of immediate importance for personality research and that connect most easily to established – but often ill-defined – concepts in this field.

Let us repeat:

$\tilde{\theta}(I_{Max})$ is the noise-free Grand Mean of the substance variable for interval I_{Max} .

$\tilde{\theta}(\sigma(t))$ is the noise-free mean of the substance variable across targets for the situation at time t .

$\bar{\theta}(\tau, I_{Max})$ is the noise-free time-mean of the substance variable for target τ and interval I_{Max} .

$\theta(\tau, \sigma(t))$ is the noise-free level of the substance variable for target τ at time t .

The respective parameters for the other target (τ') would be defined accordingly. We now introduce the following differences (see formulas 2 and 3 in the TS), which will become important later on, when we start interpreting covariances:

A target-effect ($\Delta\bar{\theta}(\tau, I_{Max}) = \bar{\theta}(\tau, I_{Max}) - \tilde{\theta}(I_{Max})$) is the difference between a given target's noise-free time-mean across I_{Max} and the respective average of all targets' noise-free time-means. In our example, it would be (e.g.) the difference between Tessa's noise-free AU12 activation across I_{Max} and the average of Tessa's and Trudy's noise-free AU12 activation across I_{Max} . There would be exactly two such target-effects in our example, one for Tessa and one for Trudy. Generally speaking, some targets do certain things more often than the average target does them. This is a basic assumption underlying much of trait-focused personality psychology.

A situation-effect ($\Delta\tilde{\theta}(\sigma(t), I_{Max}) = \tilde{\theta}(\sigma(t)) - \tilde{\theta}(I_{Max})$) is the difference between the average target's noise-free substance level for the situation present at time t and the average target's noise-free substance level across I_{Max} . In our example, this would be the average target's noise-free AU12 activation in a given situation (e.g., lecture) and the average target's noise-free AU12 activation across I_{Max} . There would be exactly five such situation-effects in our example, one for each of the one-situation intervals (I_1 to I_5). Generally speaking, some situations evoke a reaction from the average person that is different from how the average person generally reacts. This is a basic assumption underlying much of experimental social psychology.

An interaction-effect ($\Delta\theta(\tau \times \sigma(t)) = \theta(\tau, \sigma(t)) - \Delta\bar{\theta}(\tau, I_{Max}) - \Delta\tilde{\theta}(\sigma(t), I_{Max}) - \tilde{\theta}(I_{Max})$) is what remains when subtracting the relevant target-effect, the relevant situation-effect and the grand mean from a given targets's momentary noise-free substance level at time t . In our example, interaction-effects are responsible for the fact that $\theta(\tau, \sigma(t))$ and $\theta(\tau', \sigma(t))$ are not always equally far apart from $\tilde{\theta}(\sigma(t))$, and that their rank-order sometimes even reverses (e.g., during a lecture, Tessa's AU12 activation is higher than Trudy's, whereas at the funeral the opposite is the case). Interaction-effects such as these are at the core of personality theories in which personality is conceptualized in terms of the overall pattern of how a target reacts to specific situations (most notable: Mischel & Shoda, 1995). Note that, for brevity, we will call this type of effect "interaction-effect" from here on. In our example, there would be exactly (2 targets by 5 situations =) ten such effects.

The Microscopic Model. Formula 1 in the Technical Supplement defines what we call the “microscopic model”. The term implies that this formula contains all the relevant components in their most basic form and may not be reduced any further.

$$(1) \quad dS(\tau, t, \omega) = \theta(\tau, \sigma(t)) dt + dW^\tau(t, \omega)$$

Here, the $dS(\tau, t, \omega)$ stands for the change rate in a target’s substance level at time t and the $dW^\tau(t, \omega)$ stands for the random component in that change rate. To model this random noise, we borrow from the field of physics in which the use of *Brownian Motion* is the standard solution for this issue (Schilling, 2021). The basic idea therein is to have *maximum unpredictability* of the noise, which means that the noise at any point in time may not predict the noise at any neighboring point in time, regardless of how small the time difference between these two points is. By integrating across an interval comprising several points (s) in time (between 0 and t), one obtains a concrete value for the current level of noise. That value is:

$$(9) \quad W^\tau(t, \omega) = \int_0^t dW^\tau(s, \omega)$$

Figure 2 in the Technical Supplement further illustrates what the concept of Brownian Motion is about. Most of this is omitted here. In the present context, one only needs to know that $W^\tau(t, \omega)$ has an expectancy value of zero (exactly as in CTT) and a target-specific variance of $\gamma^2(\tau)t$. Note that our model does account for differences in how noisy the substance levels of individual targets are. This is the reason why there is a τ superscript accompanying parameter W . In our example (see Figure 1), Trudy obviously has a larger W parameter than Tessa.

The variance ($\gamma^2(\tau)t$) of W increases with time, because the longer such a random process goes on, the more room there will be for the sum of all random changes to deviate from zero, in either direction. Note, however, that this concerns the variance of the *cumulative* noise only – the noise at any given point in time has the same instantaneous variance everywhere. Also, the increase of the cumulative noise will be kept in check when we average across intervals in the next step.

Observable Substance Levels (Cues). In analyzing actual person judgment data, one will never have to use the microscopic model because person judgments in reality are never based on infinitely small observation intervals. Rather, the substantive base for any person judgment is an average, observable substance level for a given interval (I) with length $|I|$. By integrating formula 7 over this interval, one obtains the observable cue level that is of key concern in the present paper:

$$(10) \quad \bar{S}(\tau, I, \omega) = \frac{1}{|I|} \int_I \theta(\tau, \sigma(t)) dt + \frac{1}{|I|} \int_I dW^\tau(t, \omega)$$

The bar above the S on the left-hand side means that this is a mean of variable S over time, specifically interval I . Note that the same formula also applies when abandoning the assumption that the situation stays the same across the entire interval. This will become relevant further below.

Due to the assumed omnipresence of random noise, a target’s observable substance levels ($\bar{S}(\tau, I, \omega)$) will differ somewhat from their noise-free substance levels for the same interval ($\bar{\theta}(\tau, I)$). However, the longer the observation interval becomes, the more this noise will cancel out via integration, and the more the targets’ observable substance levels will resemble their noise-free

ones. In our example, a student watching Trudy give a lecture for ten minutes would obtain a less noisy observation of her typical A12 activation in such situations than would a student observing her for 1 minute only.

Substantive Similarity. In person judgment research, we are often interested in how similar judgments are, and this similarity is typically assessed in terms of correlation coefficients. This is the case when studying (e.g.) inter-rater agreement, the factor structure of a set of items, the internal consistency of a psychometric scale, or the predictive validity of some kind of assessment (e.g., of academic potential) with regard to some other kind of assessment (e.g., of academic achievement). The common basis for all of this is the substantive similarity of two different sets of observations, that is, the extent to which the substance levels of the targets in two observation intervals (I and I') resemble each other. We will assess this similarity in terms of a correlation, as well. But before we discuss the respective formula, we first have to talk about *interval overlap*.

Kenny (1994) introduced overlap as a component of his Weighted Average Model. He uses his model to explain consensus, that is, the level of agreement between different perceivers judging the same targets on the same item. Accordingly, Kenny conceptualizes overlap as the extent to which different perceivers base their judgments on the same vs. different information about the targets. Obviously, when this overlap becomes larger, the substantive base for the perceivers' judgments is shared to a greater extent.

Kenny addresses the topic in terms of the proportion of "acts" that two perceivers have seen a target engage in, as opposed to acts that only one of them has observed. However, this approach entails the problem of defining what an "act" is (e.g., when it begins and when it ends). In contrast, we define overlap more directly in terms of *observation intervals*, thus circumventing the need for demarcating distinct acts. Note that the term "observation interval" does not imply that any observation actually takes place. We are still dealing with substance only.

Observation intervals (I and I') may have any length, cut across any number of different situations (i.e., we now lift the "one situation per interval" constraint) and also have any degree of overlap with one another. Figure 2 below is to illustrate this. The figure uses the exact same time axis as Figure 1 above. The one-situation-intervals (I_1 to I_5) are also the same as before. However, this time the Y-axis does not represent the targets' substance levels (as in Figure 1) but the respective situational circumstances that are present for each of these intervals (M = Morning, L = Lecture, F = Funeral, B = Break). These are on a nominal scale and could thus be rearranged freely along the y-axis.

Figure 2. Upper half: the function $\sigma(t)$ giving the situation at time t , with the same “one-situation intervals” (I_1 to I_5) as in Figure 1. M = Morning, L = Lecture, F = Funeral, B = Break. Lower half: two partially overlapping observation intervals (I and I') cutting across the one-situation intervals from the upper half, and their overlapping interval ($I \cap I'$).

The lower half of Figure 2 displays two observation intervals (I and I') that cut freely across the one-situation intervals used in the upper half: Observation interval I starts approximately in the middle of interval I_1 and ranges from there into not quite the middle of interval I_3 . Observation interval I' starts approximately in the middle of interval I_2 and ranges from there to a point close to the beginning of interval I_5 . The overlapping interval created by this will be called $I \cap I'$ from here on. In our model, all observation intervals (e.g., I and I') are contained in interval I_{Max} and may have any degree of overlap with one another (e.g., one may be completely included in the other). This is explained in more detail in the Technical Supplement (formulas 61 and 62). For our more applied purposes here, we only need the ratio of the lengths of the observation interval ($|I \cap I'|$) and the product of the lengths of the two original intervals ($|I| \cdot |I'|$). This is our universal measure of overlap and features formula 73, which describes substantive similarity as a covariance.

$$(73) \quad \text{Cov}(\bar{S}(I), \bar{S}(I')) = \text{cov}(\bar{\theta}(I), \bar{\theta}(I')) + \left[\frac{1}{N} \sum_{\tau=1}^N \gamma^2(\tau) \right] \frac{|I \cap I'|}{|I| \cdot |I'|}$$

To understand the implications of this formula, one needs to understand the difference between $\text{Cov}(\bar{S}(I), \bar{S}(I'))$ and $\text{cov}(\bar{\theta}(I), \bar{\theta}(I'))$ (see Technical Supplement, Section 2). Both of these covariances assess the similarity of the targets' substance levels in intervals I and I' . However, the first covariance reflects both the systematic and the random variation in the data whereas the second covariance “sees” only the systematic variation. As long as there is no overlap ($|I \cap I'| = 0$),

the first covariance is identical to the second. However, to the extent that there is overlap, the average target's random noise variance ($\frac{1}{N} \sum_{\tau=1}^N \gamma^2(\tau)$) will *add* to the first covariance because it features in both intervals (the N in the formula is the number of targets we average across).

To ensure comparability (e.g., across studies investigating different substance variables), we recommend using a normalized version of formula 73 in which everything is divided by the common variance of the noisy substance levels in the two intervals. This leads to formula 74:

$$(74) \quad \text{Corr}(\bar{S}(I), \bar{S}(I')) = \frac{\widetilde{\text{cov}}(\bar{\theta}(I), \bar{\theta}(I')) + \left[\frac{1}{N} \sum_{\tau=1}^N \gamma^2(\tau) \right] \frac{|I \cap I'|}{|I| \cdot |I'|}}{\sqrt{\widetilde{\text{var}}(\bar{S}(I))} \sqrt{\widetilde{\text{var}}(\bar{S}(I'))}}$$

As long as there is no overlap, the denominator on the right-hand side will be larger than the numerator to the extent that targets' substance levels do contain noise. This is because the denominator "sees" this random variation whereas the numerator does not. However, with longer intervals, this difference between numerator and denominator will become smaller, because the noise will cancel out more. The potential for this to happen, however, will decrease with overlap. When overlap is perfect, substantive similarity will be one.

Up to this point, we only discussed the role that noise may play in substantive similarity, given different levels of overlap. We will now devote some more attention to the component of the two formulas that represents all the systematic influences ($\widetilde{\text{cov}}(\bar{\theta}(I), \bar{\theta}(I'))$). To be clear this is the covariance of $\bar{\theta}(\tau, I)$ and $\bar{\theta}(\tau, I')$, taken in τ . In order to understand the role of this component, it is helpful to split these parameters as follows:

$$(75) \quad \bar{\theta}(\tau, I) = \underbrace{\bar{\theta}(\tau, I) - \bar{\theta}(\tau, I_{\max})}_{=\delta\bar{\theta}(\tau, I, I_{\max})} + \underbrace{\bar{\theta}(\tau, I_{\max}) - \bar{\theta}(I_{\max})}_{=\Delta\bar{\theta}(\tau, I_{\max})} + \bar{\theta}(I_{\max})$$

Recall that $\bar{\theta}(I_{\max})$ is the Grand Mean of the substance variable, and that $\Delta\bar{\theta}(\tau, I_{\max})$ is the target-effect, respectively. We now use $\delta\bar{\theta}(\tau, I, I_{\max})$ to denote the target-specific time-average of all situation- and interaction-effects in interval I . Using this, we obtain the following decomposition of the covariance that interests us the most (to preserve readability, we temporarily simplify such that $\Delta\bar{\theta}(I_{\max})$ becomes $\Delta\bar{\theta}$, $\delta\bar{\theta}(I, I_{\max})$ becomes $\delta\bar{\theta}(I)$, and $\delta\bar{\theta}(I', I_{\max})$ becomes $\delta\bar{\theta}(I')$. Note that τ vanishes from the argument(s) of the covariance, due our taking this covariance in τ (as signaled by our use of the tilde).

$$(76) \quad \widetilde{\text{cov}}(\bar{\theta}(I), \bar{\theta}(I'))$$

$$(77) \quad = \widetilde{\text{cov}}(\Delta\bar{\theta} + \delta\bar{\theta}(I), \Delta\bar{\theta} + \delta\bar{\theta}(I'))$$

$$(78) \quad = \widetilde{\text{cov}}(\Delta\bar{\theta}, \Delta\bar{\theta}) + \widetilde{\text{cov}}(\delta\bar{\theta}(I), \delta\bar{\theta}(I')) + \widetilde{\text{cov}}(\Delta\bar{\theta}, \delta\bar{\theta}(I')) + \widetilde{\text{cov}}(\delta\bar{\theta}(I), \Delta\bar{\theta})$$

$$(79) \quad = \widetilde{\text{var}}(\Delta\bar{\theta}) + \widetilde{\text{cov}}(\delta\bar{\theta}(I), \delta\bar{\theta}(I')) + \widetilde{\text{cov}}(\Delta\bar{\theta}, \delta\bar{\theta}(I')) + \widetilde{\text{cov}}(\delta\bar{\theta}(I), \Delta\bar{\theta}).$$

Thus, the overall target-covariance of the noise-free scores in the two observation intervals ($c\overline{ov}(\bar{\theta}(I), \bar{\theta}(I'))$) reflects several independent influences: The first contribution comes from the variance of the target-effects ($\overline{var}(\Delta\bar{\theta})$). The more Trudy and Tessa generally differ in their AU12 activation levels (across I_{Max}), the greater substantive similarity will be.

The second contribution comes from the target-covariance of the time-averages of the sum of all situation- and interaction-effects in the two intervals ($c\overline{ov}(\delta\bar{\theta}(I), \delta\bar{\theta}(I'))$). Note, however, that the situation-effects are the same for all targets. Thus, they do affect $\delta\bar{\theta}(\tau, I)$ and $\delta\bar{\theta}(\tau, I')$, but not their covariance in τ . Accordingly, $c\overline{ov}(\delta\bar{\theta}(I), \delta\bar{\theta}(I'))$ will only reflect the interaction-effects. The sums of all interaction-effects in each interval (I and I') will covary more strongly the more similar the respective situational line-ups in the two intervals are (situational or σ -similarity; Sherman, Nave & Funder, 2010; see part 5 of the Technical Supplement). For example, if interval I includes the first lecture and interval I' includes the second lecture, substantive similarity should increase. However, substantive similarity will also be high when the situations in the two intervals are different but produce similar outcomes in terms of the targets' substance levels.

The third and fourth contributions come from the covariances between the target-effects and the interaction-effects in each interval: $c\overline{ov}(\Delta\bar{\theta}, \delta\bar{\theta}(I'))$ and $c\overline{ov}(\delta\bar{\theta}(I), \Delta\bar{\theta})$. This means that substantive similarity will be higher to the extent that the situations present in the observation intervals *amplify* the existing overall differences between the targets (i.e., target-effects), via target by situation interaction.

Applicability to Typical Research Questions

Formula 74 is the primary outcome of our current analyses. It enables predictions of the correlation between substance levels in any two intervals. Note that the formula is not only useful for predicting inter-rater agreement, but any correlation between observations of the same substance variable in the same targets. This will be discussed next. Note that everything that follows in this section is derived from the same model, and particularly from formula 74. With this, we hope to be able to demonstrate how much precision and parsimony may be gained by way of proper formalization of the theoretical concepts one employs.

Retest reliability. A proper test of retest reliability would require repeated assessments of the same substance variable in the same targets, under identical situational circumstances, with the same interval lengths, but with no overlap between intervals. Formula 74 tells us that under these circumstances, substantive similarity will equal the proportion of systematic variance in the overall variance, which is how reliability is commonly defined. Note that the systematic variance includes not only the variance of the target-effects but also the three other covariances from formula 79. In any case, longer intervals will mean higher reliability, depending on how noisy the data is to begin with.

Cross-situational consistency. Cross-situational consistency is of particular interest to many personality psychologists, because “personality” may be conceptualized as the extent to which differences between people persist across varying circumstances. This view basically equates personality with target-effects. However, estimates of cross-situational consistency will always be confounded with situation-similarity – the more the circumstances under which one observes targets resemble each other, more the covariance of the respective interaction-effects – and the covariances between these and the target effects – will contribute to consistency.

Inter-rater agreement. One major reason why different judgments of the same targets may agree with one another is that the different judges saw the targets behave similarly. According to formula 74, there are several possible reasons for this to happen, all of which are not mutually exclusive: First, target-variance may be large. Second, there may be overlap, and thus: shared noise. Third, there may be situation similarity. Fourth, the situations in which the targets are observed may amplify the existing target-variance.

Validity. The term “validity” is used in a fairly unsystematic fashion throughout the psychometrics literature (e.g., Zachar & Jablensky, 2015). Here, we use it to denote associations between the targets’ observable substance levels in some (“predictor”) interval (I) and their observable substance levels in some other (“criterion”) interval (I') which is the interval that is of greater interest. Typically, the latter interval is larger, situationally more diverse, and located later on the timeline, but both intervals are samples from I_{Max} . With this research question, permitting any overlap between I and I' would be undesirable because it would introduce “predictor-criterion-contamination”.

For example, we may use Trudy’s and Tessa’s average AU12 activation levels in the morning (I_1) to predict their average AU12 activation levels in the three following intervals (I_2 to I_4), when they have to deal with other people. This would be a classic example of *predictive validity*, with the measured substance variable being the same in both intervals (so formula 74 applies).

We would expect this type of predictive validity to be larger to the extent that there is target-variance. We would also expect it to increase with the lengths of both intervals, because this should cancel out more of the noise. To maximize validity, the situational setup of the predictor and criterion interval should be as similar as possible (or at least be reasonably expected to elicit similar substance levels). This is true regardless of whether one wants to validly predict the targets’ behavior in a very circumscribed set of situations, or the stable differences between them more broadly. In the latter case, it will be helpful to either use a diverse situational setup (so interaction-effects can be expected to cancel out), or a setup that elicits few interaction-effects to begin with.

Comparison with Previous Models

Kenny’s (1994) Weighted Average Model

As we stated above, our model covers much of the same ground that is also covered by David Kenny’s (1994) Weighted Average Model (WAM; see also the closely related PERSON model by Kenny, 2004, 2019). Specifically, versions of the acquaintance (n), overlap (q) and consistency ($r1$) parameters from Kenny’s WAM feature in our model, as well: Using our terminology, one would represent acquaintance as $|I|$, overlap as $\frac{|I \cap I'|}{|I| \cdot |I'|}$ and cross-situational consistency in terms of Formula 74, with the additional specification that both intervals contain just one situation, and these situations are different from one another. The remaining WAM parameters (see Kenny, 1994; pages 63-65) are not addressed in the current article, because they all involve contributions attributable specifically to perceivers. These will be treated in a follow-up article.

In several respects, however, the model that we present here goes beyond WAM in terms of precision and/or generality and/or scope. We only have room here to highlight some of these:

First, we incorporate time as a continuous variable. This enables us to account for both overlap and a core aspect of “acquaintance” (i.e., interval length) objectively and in a more precise fashion than is possible when the targets’ continuous response streams have to be segmented into separate “acts”. Second, our model (and specifically the part of it on which the present paper focuses)

incorporates a fully formalized account of what may be going on at the level of the substance, independent of whether this substance is ever used by perceivers to inform their judgments. It therefore lends itself to establish a more direct connection to the branch of personality research that uses *objective* (i.e., non-judgment) data (e.g., Stachl et al., 2021), and even to making exact predictions in that realm. Third, our model does not distinguish between “categorical” information (e.g., biological sex, height, skin color, age) and “behavioral” information. Rather, it accounts for all different types of substantive information about targets the same way. What is called “categorical information” in Kenny’s WAM model would just be substance variables that are relatively stable across time (i.e., variables with little noise, situation effects, and interaction-effects). Fourth, our model explicitly incorporates target-effects, situation-effects, and interaction-effects. Thus, it establishes a direct connection to key concepts from the so-called “person-situation debate”; Mischel, 1968; Kenrick & Funder, 1988). Some of these connections are discussed in more detail in the next section. Fifth, our model also expressly accounts for the influence and the properties of random noise, thus establishing direct links both to Classical Test Theory as well as to mainstream modelling practices in other fields of science (e.g., physics).

Latent State-Trait Theory

The present model also has significant overlap with Latent-State-Trait-Theory (LST; Steyer, Ferring & Schmitt, 1992; Steyer, Schmitt & Eid, 1999; Steyer, Mayer, Geiser & Cole, 2015). Most important, both the current model and LST focus on explaining the level of a variable that is measured at some occasion and assume that person-effects, situation-effects, interaction-effects and random noise contribute to this measured level. In fact, what is called the “latent state” in LST is largely identical in meaning to the $\theta(\tau, \sigma(t))$ parameter in our model.

Whereas earlier models within the LST framework did not systematically disentangle situation- and interaction-effects, a recent model by Geiser et al. (2015) does precisely that. Even more recent models (e.g., Koch, Voelkle & Driver, 2023) do consider time as a continuous factor, but only between measurements, not within the measurement interval itself. While our model was developed independent of LST – we started from Kenny’s (1994) WAM and Brunswik’s (1956) Lens Model – the parallels are obvious and may actually be viewed as encouraging, because independent strands of theory development converge so well.

There are also a few important differences, however: First, the model we present does systematically decompose the (“latent”) noise-free substance levels into target, situation and interaction effects while also accounting for time as a continuous variable. We are not aware of any LST model so far that combines both of these properties. However, in our model, it is the length of the observation intervals rather than the spacing between them that matters. The spacing between observations will only become relevant once it is assumed that the systematic components of what is being measured (especially target- and interaction-effects) may vary over time. The present model may of course be amended this way, however.

Second, our model allows for any composition of an observation interval in terms of situations whereas the LST models that we are aware of only handle (repeated) assessments of targets in the same situation as opposed to another situation. This feature enables our model to capture situation similarity as a gradual phenomenon (see below). Third, our model expressly covers overlap. Fourth, our model does account for target-specific levels of noisiness.

Perhaps even more important, in the present paper we emphasize that all of this applies to the substance on which judgments may be based, rather than to those judgments. It seems important to point this out once more because LST models are sometimes used as if the difference between a

direct assessment of substance levels and judgments delivered by human perceivers were irrelevant. It is not, however, because there are numerous additional sources of co-variation in the latter kind of data (e.g., acquiescence, extreme responding, attitudes, content-based perceiver-effects) that may be misinterpreted in terms of substantive phenomena (Baird et al., 2017; Rau et al., 2021; Heynicke, Rau, Leising, Wessels & Wiedenroth, 2022; Wetzel, Lüdtke, Zettler & Böhnke, 2016). Recent versions of LST have begun to consider the influence of measurement methods (e.g., self- vs. other-ratings) more systematically, however (e.g., Hintz, Geiser & Shiffman, 2018).

Brunswik's Lens Model

The present paper describes the part of our model that captures the emergence of a target's average substance level for a given interval ($\bar{S}(\tau, I, \omega)$). As we pointed out above, this level plays the same role conceptually as a "cue" does in Brunswik's (1956) Lens Model. In a follow-up paper, we will talk about how perceivers may use such average substance levels in judging targets.

Probably the main purpose of Brunswik's model is to describe how ("proximal") cues may be expressions of ("distal") causes, and how judgments based on cues may be used to infer such causes. The cue levels whose emergence we model here are perfectly suited for such analyses. Specifically, if the targets' levels on different substance variables are found to correlate with one another, this may be due to variation in a third variable that causally affects each of these substance variables. Before this possibility may be seriously considered, however, one has to rule out the alternative possibility that the levels of the different substance variables influence each other directly.

The correlations between the targets' cue levels and their levels of an underlying cause variable – if it exists – are often called "cue validities" in analyses based on the Lens Model. What the present paper contributes to the understanding of cue validities is that there may be two relevant sources of attenuation involved in this regard: random noise, and target by situation interactions. The influence of both is likely to be reduced with longer observation intervals.

Relationships to Key Concepts in the Personality Literature

In the remainder of this paper, we will briefly touch on how some key concepts from the personality literature may either be expressed in terms of the model we introduced so far, or added to it as extensions. Our main goal in this is to showcase how formalization allows for a more precise specification of theoretical concepts that, so far, have mostly been described in terms of the natural language.

Situation characteristics. Our treatment of situations so far has been fairly simplistic, as we represented them with a single, nominally scaled variable ($\sigma(t)$). For many research purposes, this may actually be sufficient. For example, many studies in social psychology use experimental designs that, mathematically, have the same properties. However, in certain cases it may be necessary to represent situations *multidimensionally* (e.g., Rauthmann et al., 2014) in terms of several properties ($\sigma_1(t)$, $\sigma_2(t)$ etc.) that may vary *continuously* (or even randomly) over time, and with varying degrees of independence from one another.

Person-, situation- and interaction-effects. Obviously, our model is very closely connected to key concepts from the "person-situation debate" in psychology (Mischel, 1968; Kenrick & Funder, 1988). In fact, the $\Delta\bar{\theta}(\tau, I_{max})$ parameter *is* the person-effect (better: "target-effect", because perceivers, which will come to interest us only later, are persons, too), the $\Delta\tilde{\theta}(\sigma(t), I_{Max})$ parameter *is* the situation effect, and the $\Delta\theta(\tau \times \sigma(t), I_{Max})$ parameter *is* the target by situation interaction-effect. Comparing the proportions of the overall substance variance that are accounted for by each of these parameters will be directly informative regarding the relative influence of persons,

situations, and their interactions. This, however, requires controlling for random error, which may only be separated from the other parameters by having the same targets go through the same setup of situations a second time. Note that, whereas our model is applicable to both experimental and naturalistic designs, it does not explicitly account for differences in the sets of situations that individual targets typically encounter (“situation contact”; Rauthmann, Sherman, Nave & Funder, 2015).

Situation Similarity. Our model does capture the idea that the situations in which targets are observed may be more or less similar to one another, and that this similarity may affect the substantive similarity of the targets’ behavior (Sherman, Nave & Funder, 2010). As we conceptualize it, situation similarity is a property of a pair of observation intervals that are being compared in this regard. If one describes the situation using a single nominal scale (as we did in the present paper), situation similarity may simply be assessed as the percentage of two observation intervals in which sigma has the same value. More complex assessments of situation characteristics are conceivable, however (see above), and would enable more complex assessments of situation similarity.

Intraindividual variability. For decades, psychologists have been interested in how stable people’s behavior and experiences are across time and situations (Vazire & Sherman, 2017), including the extent to which there are systematic *differences* between targets in this stability (e.g., Fleeson, 2001). Note again that this concerns the targets’ momentary *substance* levels, which are the primary focus of the present paper. Research has often used person judgments as a stand-in for these substance levels, but that approach may fall short of its purpose because the variability of interest may not be disentangled from the perceivers’ variability in responding to items (e.g., Baird, Lucas & Donnellan, 2017). Our model also clarifies that there are two possible sources of substantive intra-individual variability (i.e., target-specific noise and interaction-effects) that may and should probably be studied separately, because otherwise one would lump together two components with entirely different meanings.

Personality. Our model connects to several different ways in which the term “personality” is commonly used. The first approach is to simply think of personality in terms of stable differences between targets. Using this approach, a target’s personality may be described for every individual substance variable, as the target-effect ($\Delta\bar{\theta}(\tau, I_{Max})$). The second approach focuses on interaction-effects instead. Here, a target’s personality is described in terms of the pattern of if-then relationships between specific situational circumstances and the target’s respective response tendencies. This is the approach proposed by Mischel and Shoda (1995), whose concept of a “personality signature” basically comprises a given target’s set of interaction-effects across a number of different situations. Finally, the term “personality” may be used to denote measurable or entirely hypothetical substance variables that causally account for co-variation among cues. This “reflective” idea of personality (Schmittmann et al., 2013) informs most personality researchers’ use of Brunswik’s (1956) Lens Model, and is also consistent with Funder’s (1995) conceptualization. Note that, in contrast, Kenny’s (1994, 2004, 2019) approach is purely “formative” in nature, in that it is only concerned with *how perceivers use* the information that they receive about targets. In his PERSON model (Kenny, 2004), personality (P) is the judgment that the average perceiver would make if he or she had access to all the relevant cue information.

Strong vs. weak persons / situations. The person-situation literature prominently features the idea that both situations and targets may be distinguished from one another in how “strong” vs. “weak” they are (e.g., Schmitt et al., 2013). Note that this terminology is not meant to imply any evaluation (e.g., *strong* being somehow “better” than *weak*), but only to compare situations with one another, and targets with one another, in terms of how pervasive their respective influences are on

substance levels. There is no binding convention as to how these terms are to be used. We suggest the following (Leising & Müller-Plath, 2009): A target (τ) would be considered “strong” (“weak”) to the extent that – averaging across time – her noise-free, momentary substance levels ($\theta(\tau, \sigma(t))$) resemble her noise-free substance level overall ($\bar{\theta}(\tau, I_{max})$) more (less) than they resemble the average target’s noise-free substance level at the same point in time ($\bar{\theta}(\sigma(t), I_{Max})$). Using this approach, Trudy (τ' , red) is a stronger target than Tessa (τ , blue) (see Figure 1). Strong (weak) situations may be defined analogously.

Good vs. bad targets. “Good targets” are targets that are easy to judge (Funder, 1995). According to our model, there are actually *two* ways in which one may be a good target. First, being a good target could mean that there is little random variation in a target’s behavior stream. Using this understanding, Tessa (τ , blue) would be a better target than Trudy (τ' , red) because Tessa’s $\gamma^2(\tau)$ parameter is smaller and thus her observable substance levels *in each individual situation* will resemble her noise-free levels *in those situations* more (see Figure 1). Second, being a good target may also be interpreted in terms of how easy it is to infer a target’s noise-free substance level *overall*. Using this understanding, Trudy would be the better target because her noise-free substance levels in individual situations deviate less strongly from her overall substance level. Note that the latter understanding (being true to one’s overall level) is basically identical to what it means to be a “strong person” (see above).

Good vs. bad Information. Good Information is one of the key moderators of judgment accuracy in Funder’s (1995) *Realistic Accuracy Model*. However, Funder himself states that the term covers at least two largely distinct concepts: First, information will become “better” with longer observation intervals. This will be the case because longer intervals will result in more cancelling out of random noise, but only to the extent that those substance levels actually *are* noisy. Longer intervals (which likely, but not necessarily, go along with a more diverse setup of situations) may also be advantageous because interaction-effects will cancel out more. Second, some situations may be more “diagnostic” than others (Borkenau, Mauer, Riemann, Spinath & Angleitner, 2004) in that the targets’ substance levels in them predicts the same targets’ substance levels in the criterion interval particularly well.

Acquaintance. The extent to which perceivers *know* their targets has played a prominent role in the personality judgment literature (e.g., as a predictor of inter-rater agreement; Biesanz, West & Millevoi, 2007; Funder & Colvin, 1988). Interval length ($|I|$) clearly represents an important aspect of this concept, as greater acquaintance will usually imply having spent *more time* observing a target. Given some level of random-noise, this should lead to greater reliability and validity. However, the acquaintance concept is ambiguous, because being better acquainted with a target may also mean (a) being able to observe the target in different (e.g., more intimate) situations, and (b) being able to communicate more with the target (e.g., about his or her subjective experiences) (Kenny, 1994). These two aspects are *not* part of the model as we present it here, but may of course be added.

Limitations

The model as we present it here uses a number of constraints, in order to keep complexity within manageable bounds: First, we assume that the noise-free substance components are stable over time. That is, target-effects do not change, situation-effects do not change, and interaction-effects do not change, either. In reality, however, it is quite likely that changes in all of these parameters do occur. This may be accounted for by adding more parameters to the model (e.g., ones assessing overall developmental trends with age, and/or ones assessing decreasing predictability).

There are now variants of LST that do account for this additional layer of complexity (Koch et al., 2023).

Second, our model (like CTT and LST) operates on the assumption that substance variables have unlimited variability without lower or upper bounds. In reality, however, most the respective continua do have such bounds and this may have important consequences for model accuracy. Most important, target and situation effects will come to limit each others' influences under these more realistic assumptions (Blum & Schmitt, 2018; Schmitt et al., 2013). For example, when a situation already induces very high levels of anxiety, target-effect can do little to increase that level even further. The model may be amended to reflect these considerations, but this will make subsequent derivations more complicated. The question of how much model complexity is actually needed in this regard is ultimately an empirical one.

Third, in the present paper, we introduced our model using a single substance variable. However, the model may just as well be applied to several substance variables at once, since S may well be a vector, rather than a scalar. This is relevant, for example, with regard to the analysis of whole *profiles* of personality judgments (i.e., judgments of targets by perceivers on many different items; Biesanz, 2010; Furr, 2008) and will be addressed in a follow-up paper. It is also relevant with regard to studies making predictions of one substance variable (in interval I') based on another substance variable (in interval I).

Conclusion and Outlook

We hope that the presentation so far has shown just how much complexity needs to be accounted for when trying to conceptualize person judgment processes in a strict, formal way. Note again, however, that the current paper only covers cue emergence and does not even begin to address the ways in which perceivers may *use* cues in judging people. Despite this somewhat daunting complexity, we also hope that the presentation has shown how much clarity and parsimony may be gained by attempting a stricter formalization: A single model comprising a relatively limited set of parameters may go a long way in expressing all sorts of important but previously unconnected theoretical ideas, and connecting them in a rather seamless manner. The same model may enable a large number of exact predictions regarding a broad range of important phenomena. Of course, it is not necessary to operationalize each and every parameter in every empirical study. But a model like the one presented here may nevertheless be useful because it explicates all the factors that are likely to have an influence on the data. This enables researchers to plan their studies more systematically (e.g., by letting the factors they do not measure vary randomly), which may in turn contribute to better validity. In order to motivate more psychologists to get involved with this crucial but highly demanding kind of academic work (i.e., modelling), it may be necessary to find ways of incentivizing it better (e.g., Leising et al., 2022a, 2022b). It will also be helpful to turn to colleagues who already possess the required expertise, invite them to collaborate, and learn from them.

Appendix

Table 1: Overview of the parameters

Parameter	Meaning	Ref.
t	time	
τ	target	
π	perceiver	
$\sigma(t)$	situation at time t	Fig. 1
ω	unsystematic randomness	
$I_{\max}$	standard interval	Fig. 1
I	interval of observation	Fig. 1
$ I $	length of the interval I	
$\theta(\tau, \sigma(t))$	noiseless score of target τ at time t	(1), (2)
$\bar{\theta}(\sigma(t))$	the average target's noiseless score for the situation $\sigma(t)$ at time t	
$\bar{\theta}(\tau, I)$	noiseless cue level of target τ for the interval I	(17), (36)
$\bar{\bar{\theta}}(I) = \bar{\bar{\theta}}(I)$	grand mean of the substance variable across I	(29)
$\Delta\bar{\theta}(\sigma(t), I)$	$\bar{\theta}(\sigma(t)) - \bar{\bar{\theta}}(I)$; situation effect if $I = I_{\max}$	(33), (2)
$\Delta\bar{\theta}(\tau, I)$	$\bar{\theta}(\tau, I) - \bar{\bar{\theta}}(I)$; target effect if $I = I_{\max}$	(33), (2)
$\Delta\theta(\tau \times \sigma(t), I)$	$\theta(\tau, \sigma(t)) - \Delta\bar{\theta}(\tau, I) - \Delta\bar{\theta}(\sigma(t), I) - \bar{\bar{\theta}}(I)$ target by situation interaction effect if $I = I_{\max}$	(34), (3)
$W^\tau(t, \omega)$	cumulative noise level up to time t for target τ , modelled by a Brownian motion with mean zero and variance $\gamma^2(\tau)t$	
$\gamma(\tau)\sqrt{t}$	standard deviation of the noise at time t for target τ	Fig. 2
$dS(\tau, t, \omega)$	rate of the (noisy) substance level of target τ at time t	(1)
$dW^\tau(t, \omega)$	rate of the random error for target τ at time t ; "white noise"	(1)
$\bar{S}(\tau, I, \omega)$	observable cue level of target τ for the interval I	(10), (35)
$\widetilde{}, \widetilde{\text{var}}, \widetilde{\text{cov}}$	mean, variance and covariance taken only in targets (τ)	Sec. 2
$\mathbb{E}, \text{Var}, \text{Cov}$	mean value, variance and covariance taken only in the noise (ω)	Sec. 2
$\mathbb{V}\widetilde{\text{ar}}, \mathbb{C}\widetilde{\text{ov}}$	variance and covariance taken in targets (τ) and noise (ω) simultaneously	Sec. 2

References

1. Back, M. D., & Nestler, S. (2016). Accuracy of judging personality. In J. A. Hall, M. S. Mast, & T. V. West (Eds.), *The social psychology of perceiving others accurately* (pp. 98–124). Cambridge University Press. <https://doi.org/10.1017/CBO9781316181959.005>
2. Baird, B.M., Lucas, R. E., Donnellan, M. B. (2017). The role of response styles in the assessment of intraindividual personality variability. *Journal of Research in Personality*, 69, 170-179. <https://doi.org/10.1016/j.jrp.2016.06.015>
3. Biesanz, J. C. (2010). The Social Accuracy Model of interpersonal perception: Assessing individual differences in perceptive and expressive accuracy. *Multivariate Behavioral Research*, 45, 853–885. <https://doi.org/10.1080/00273171.2010.519262>
4. Biesanz, J. C., West, S. G., & Millevoi, A. (2007). What do you learn about someone over time? The relationship between length of acquaintance and consensus and self-other agreement in judgments of personality. *Journal of Personality and Social Psychology*, 92(1), 119–135. <https://doi.org/10.1037/0022-3514.92.1.119>
5. Blum, G.S. & Schmitt, M. (2018). The Nonlinear Interaction of Person and Situation (NIPS) Model and its values for a psychology of situations. In: Funder, D. C., Rauthmann, J. F., & Sherman, R. A. *Oxford Handbook of Psychological Situations*. Oxford, UK: Oxford University Press.
6. Borgstede, M., & Eggert, F. (2023a). Squaring the circle: From latent variables to theory-based measurement. *Theory & Psychology*, 33(1), 118–137. <https://doi.org/10.1177/09593543221127985>
7. Borgstede, M., & Eggert, F. (2023b). Meaningful measurement requires substantive formal theory. *Theory & Psychology*, 33(1), 153–159. <https://doi.org/10.1177/09593543221139811>
8. Borkenau, P., Mauer, N., Riemann, R., Spinath, F. M., & Angleitner, A. (2004). Thin Slices of Behavior as Cues of Personality and Intelligence. *Journal of Personality and Social Psychology*, 86(4), 599–614. <https://doi.org/10.1037/0022-3514.86.4.599>
9. Brunswik, E. (1956). *Perception and the representative design of psychological experiments*. Berkeley, CA: University of California Press.
10. Ekman, P., Friesen, W. V. (1978). *Facial Action Coding System: A technique for the measurement of facial movement*. Palo Alto, CA: Consulting Psychologists Press.
11. Fleeson, W. (2001). Toward a structure- and process-integrated view of personality: Traits as density distributions of states. *Journal of Personality and Social Psychology*, 80(6), 1011–1027. <https://doi.org/10.1037/0022-3514.80.6.1011>

12. Frankenhuys, W., Panchanathan, K., & Smaldino, P. E. (2022, July 23). Strategic ambiguity in the social sciences. <https://doi.org/10.31222/osf.io/kep5b>
13. Funder, D. C. (1995). On the accuracy of personality judgment: A realistic approach. *Psychological Review*, 102, 652-670. <https://doi.org/10.1037/0033-295X.102.4.652>
14. Funder, D. C., & Colvin, C. R. (1988). Friends and strangers: Acquaintanceship, agreement, and the accuracy of personality judgment. *Journal of Personality and Social Psychology*, 55(1), 149-158. <https://doi.org/10.1037/0022-3514.55.1.149>
15. Furr, M. R. (2008). A framework for profile similarity: Integrating similarity, normativeness, and distinctiveness. *Journal of Personality*, 76(5), 1267-1316. <https://doi.org/10.1111/j.1467-6494.2008.00521.x>
16. Glöckner, A., & Betsch, T. (2011). The empirical content of theories in judgment and decision making: shortcomings and remedies. *Judgment and Decision Making*, 6(8), 711-721.
17. Heynicke, M., Rau, R., Leising, D., Wessels, N., & Wiedenroth, A. (2022). Perceiver effects in person perception reflect acquiescence, positivity, and trait-specific content: evidence from a large-scale replication study. *Social Psychological and Personality Science*, 13(4), 839–848. <https://doi.org/10.1177/19485506211039101>
18. Hintz, F., Geiser, C., & Shiffman, S. (2019). A latent state–trait model for analyzing states, traits, situations, method effects, and their interactions. *Journal of Personality*, 87(3), 434–454. <https://doi.org/10.1111/jopy.12400>
19. Kenny, D. A. (1994). *Interpersonal perception. A social relations analysis*. New York, NY: The Guilford Press.
20. Kenny, D. A. (2019). *Interpersonal perception. The foundation of social relationships*. New York, NY: The Guilford Press.
21. Kenny, D. A. (2004). PERSON: A general model of interpersonal perception. *Personality and Social Psychology Review*, 8(3), 265-280. https://doi.org/10.1207/s15327957pspr0803_3
22. Kenrick, D. T., & Funder, D. C. (1988). Profiting from controversy: Lessons from the person-situation debate. *American Psychologist*, 43(1), 23–34. <https://doi.org/10.1037/0003-066X.43.1.23>
23. Koch, T., Voelkle, M. C., & Driver, C. C. (2023) Analyzing longitudinal multirater data with individually varying time intervals. *Structural Equation Modeling: A Multidisciplinary Journal*, 30(1), 86-104. <https://www.tandfonline.com/doi/full/10.1080/10705511.2022.2096612>

24. Leising, D., & Müller-Plath, G. (2009). Person-situation integration in research on personality problems. *Journal of Research in Personality*, 43(2), 218–227.
<https://doi.org/10.1016/j.jrp.2009.01.017>
25. Leising, D., & Schilling, R. (2021, October 4). Mathematical Model of Person Judgment: Part 1 (Cue Emergence). <https://doi.org/10.31234/osf.io/myua4>
26. Leising, D., Thielmann, I., Glöckner, A., Gärtner, A., & Schönbrodt, F. D. (2022a). Ten steps toward a better personality science - how quality may be rewarded more in research evaluation. *Personality Science*, 3, 1-44. <https://doi.org/10.5964/ps.6029>
27. Leising, D., Thielmann, I., Glöckner, A., Gärtner, A., & Schönbrodt, F. (2022b). Ten steps toward a better personality science – a rejoinder to the comments. *Personality Science*, 3, 1-15.
<https://doi.org/10.5964/ps.7961>
28. Lord, F. M. & Novick, M. R. (1968). *Statistical theories of mental test scores*. Reading MA: Addison-Wesley Publishing Company.
29. Mischel, W. (1968). *Personality and assessment*. New York, NY: Wiley.
30. Mischel, W., & Shoda, Y. (1995). A cognitive-affective system theory of personality: Reconceptualizing situations, dispositions, dynamics, and invariance in personality structure. *Psychological Review*, 102(2), 246–268. <https://doi.org/10.1037/0033-295X.102.2.246>
31. Novick, M.R. (1966). The axioms and principal results of classical test theory. *Journal of Mathematical Psychology*, 3(1), 1-18. [https://doi.org/10.1016/0022-2496\(66\)90002-2](https://doi.org/10.1016/0022-2496(66)90002-2)
32. Oberauer, K., & Lewandowsky, S. (2019). Addressing the theory crisis in psychology. *Psychonomic Bulletin and Review*, 26(5), 1596-1618. <https://doi.org/10.3758/s13423-019-01645-2>
33. Rau, R., Carlson, E. N., Back, M. D., Barranti, M., Gebauer, J. E., Human, L. J., Leising, D., & Nestler, S. (2021). What is the structure of perceiver effects? On the importance of global positivity and trait-specificity across personality domains and judgment contexts. *Journal of Personality and Social Psychology*, 120(3), 745–764. doi: 10.1037/pspp0000278
34. Rauthmann, J. F., Gallardo-Pujol, D., Guillaume, E. M., Todd, E., Nave, C. S., Sherman, R. A., Ziegler, M., Jones, A. B., & Funder, D. C. (2014). The Situational Eight DIAMONDS: A taxonomy of major dimensions of situation characteristics. *Journal of Personality and Social Psychology*, 107(4), 677–718. <https://doi.org/10.1037/a0037250>
35. Rauthmann, J. F., Sherman, R. A., & Funder, D. C. (2015). Principles of situation research: Towards a better understanding of psychological situations. *European Journal of Personality*, 29, 363–381.
<http://dx.doi.org/10.1002/per.1994>

36. Robinaugh, D.J., Haslbeck, J.M.B., Ryan, O., Fried, E.I., & Waldorp, L.J. (2021). Invisible hands and fine calipers: a call to use formal theory as a toolkit for theory construction. *Perspectives on Psychological Science*, 16(4), 725-743. <https://doi.org/10.1177/1745691620974697>
37. Rodgers, J. L. (2010). The epistemology of mathematical and statistical modeling: A quiet methodological revolution. *American Psychologist*, 65(1), 1–12. <https://doi.org/10.1037/a0018326>
38. Schilling, R. L. (2021). *Brownian Motion* (3rd ed.). Berlin: De Gruyter.
39. Schmitt, M., Gollwitzer, M., Baumert, A., Blum, G., Gschwendner, T., Hofmann, W., & Rothmund, T. (2013). Proposal of a Nonlinear Interaction of Person and Situation (NIPS) model. *Frontiers in Psychology*, 4, 499. <https://doi.org/10.3389/fpsyg.2013.00499>
40. Sherman, R.A., Nave, C.S., & Funder, D.C. (2010). Situational similarity and personality predict behavioral consistency. *Journal of Personality and Social Psychology*, 99(2), 330-343. <https://doi.apa.org/doiLanding?doi=10.1037%2Fa0019796>
41. Smaldino, P. E. (2020). How to translate a verbal theory into a formal model. *Social Psychology*, 51(4), 207–218. <https://doi.org/10.1027/1864-9335/a000425>
42. Stachl, C., Boyd, R. L., Horstmann, K. T., Khambatta, P., Matz, S. C., & Harari, G. M. (2021). Computational Personality Assessment. *Personality Science*, 2(1), 1-22. <https://doi.org/10.5964/ps.6115>
43. Steyer, R., Ferring, D. & Schmitt, M. J. (1992). States and traits in psychological assessment. *European Journal of Psychological Assessment*, 8, 79-98.
44. Steyer, R., Mayer, A., Geiser, C., & Cole, D. A. (2015). A theory of states and traits – revised. *Annual Review of Clinical Psychology*, 11(1), 71-98. <https://doi.org/10.1146/annurev-clinpsy-032813-153719>
45. Steyer, R., Schmitt, M., & Eid, M. (1999). Latent state–trait theory and research in personality and individual differences. *European Journal of Personality*, 13(5, Spec Issue), 389–408. [https://doi.org/10.1002/\(SICI\)1099-0984\(199909/10\)13:5<389::AID-PER361>3.0.CO;2-A](https://doi.org/10.1002/(SICI)1099-0984(199909/10)13:5<389::AID-PER361>3.0.CO;2-A)
46. Vazire, S., & Sherman, R. A. (2017). Introduction to the special issue on within-person variability in personality. *Journal of Research in Personality*, 69, 1-3 . <https://doi.org/10.1016/j.jrp.2017.07.004>
47. Wetzel, E., Lüdtke, O., Zettler, I., & Böhnke, J. R. (2016). The stability of extreme response style and acquiescence over 8 years. *Assessment*, 23(3), 279–291. <https://doi.org/10.1177/1073191115583714>
48. Zachar, P., & Jablensky, A. (2015). Introduction: The concept of validation on psychiatry and psychology. In P. Zachar, D. S. Stoyanov, M. Aragona & A. Jablensky (Eds.), *Alternative perspectives on psychiatric validation* (pp. 3-24). Oxford, UK: Oxford University Press.