The Problem of Coordination and the Pursuit of Structural Constraints in Psychology

David Kellen

Syracuse University

Clinton P. Davis-Stober

University of Missouri

John C. Dunn

University of Western Australia

Michael L. Kalish

Syracuse University

Author Note

We thank Henrik Singmann, Elliot A. Ludvig, Mikhail S. Spektor, Richard D. Morey, and two anonymous reviewers for their valuable comments. Correspondence should be sent to David Kellen (davekellen@gmail.com).

Abstract

Paul Meehl's famous critique laid out in detail many of the problematic practices and conceptual confusions that stand in the way of meaningful theoretical progress in psychological science. By integrating many of Meehl's points, we argue that one of the reasons for the slow progress in psychology is the failure to acknowledge the *problem of coordination*. This problem arises whenever we attempt to measure quantities that are not directly observable, but can be inferred from observable variables. The solution to this problem is far from trivial, as demonstrated by a historical analysis of thermometry. The key challenge is the specification of a functional relationship between theoretical concepts and observations. As we demonstrate, empirical means alone will not allow us to determine this relationship. In the case of psychology, the problem of coordination has dramatic implications in the sense that it severely constrains our ability to make meaningful theoretical claims. We discuss several examples and outline some of the solutions that are currently available.

Keywords: theory, measurement, scaling, quantitative modeling, order-constrained inference

The Problem of Coordination and the Pursuit of Structural Constraints in Psychology

In 1978, Paul E. Meehl (1920-2003) offered a scathing criticism of psychological science. According to Meehl, psychologists were busy occupying themselves with theories that were both “scientifically unimpressive and technologically worthless” (p. 806). The consequence of such an activity is an impediment of cumulative theoretic progress, with entire research communities trapped in vicious cycles in which *theories never die but simply fade away* (see also Newell, 1973). Behind this unfortunate state of affairs, Meehl argued, was psychologists’ tendency to overlook basic considerations regarding the falsifiability of theories and the inappropriate use of null-hypothesis testing.

The goal of the present paper is to relate Meehl’s critique of psychology’s theory-testing practices to the ‘*problem of coordination*’ which scientists, historians, and philosophers have discussed for well over a century (e.g., Chang, 2004; Mach, 1896/1986; Reichenbach, 1957; Tal, 2017; van Fraassen, 2008).¹ We argue that by not addressing this problem, psychologists have compromised their ability to assess the relative merits of competing theories, resulting in a proliferation of theoretical concepts or phenomena for which there is little or no actual evidence. Relying on historical and philosophical analyses of thermometry (Chang, 2004; Mach, 1896/1986; Sherry, 2011), we make the case that the answer to the problem of coordination involves a careful and systematic *joint* development of theoretical models and experimental knowledge. Finally, we will discuss readily-available testing approaches that sidestep the problem of coordination.

The Falsification of Theories in Psychology

Let $\mathcal{T}$ denote the theoretical construct under investigation. For example, $\mathcal{T}$ could be a statement about whether a particular activity is governed by a single- or dual cognitive process. Let $\mathcal{A}$ denote the auxiliary assumptions, such that when considered jointly with $\mathcal{T}$ gives rise to a set of predicted outcomes $\mathcal{O}$. The assumptions in $\mathcal{A}$ may include common

¹ Chang (2004) refers to the problem of coordination as the ‘*problem of nomic measurement*.’

statistical assumptions (e.g., independence of responses) but may also include other elements regarding how constructs in $\mathcal{T}$ relate to observations, such as linearity assumptions among independent variables (see Kellen, 2019). The interplay between these concepts lies at the heart of our critique.

The falsifiability of any given theory $\mathcal{T}$, along with auxiliary assumptions $\mathcal{A}$, presupposes the ability to differentiate between the set $\mathcal{O}$ of outcomes deemed permissible and the complementary set $\bar{\mathcal{O}}$ of those that are not. *Modus Tollens* can then be invoked to falsify the conjunction $\mathcal{T} \& \mathcal{A}$:

If $\mathcal{T} \& \mathcal{A}$ is true, then $\mathcal{O}$.

We observe $\bar{\mathcal{O}}$.

Therefore, $\mathcal{T} \& \mathcal{A}$ is false.

The falsifiability of $\mathcal{T} \& \mathcal{A}$ can be low due to the small size of $\bar{\mathcal{O}}$ relative to $\mathcal{O}$. For example, consider a theory stating that two population means, from a continuous dependent variable, are not equal. Such a theory is vacuous given that $\bar{\mathcal{O}}$ is a single point on a continuum! Careless consideration of $\bar{\mathcal{O}}$ can lead to theories which are unlikely (or impossible) to be falsified. However, a relatively large $\bar{\mathcal{O}}$ doesn't necessarily mean that $\mathcal{T}$ is now easily falsifiable. After all, the falsification of the conjunction $\mathcal{T} \& \mathcal{A}$ can be attributed to a failure of one or more of the auxiliary assumptions in $\mathcal{A}$ (Duhem, 1951; Quine, 1963). Borrowing language from Lakatos (1976), $\mathcal{A}$ effectively serves as a “protective belt” over $\mathcal{T}$, saving it from falsification. This situation leads researchers to engage in an iterative process in which $\mathcal{A}$ is scrutinized and amended, before making any determination on the merits of $\mathcal{T}$ (Lakatos, 1976; Meehl, 1990).

Alternatively, one can try to make a case for $\mathcal{T}$ by appealing to the falsification of a complementary theory $\bar{\mathcal{T}}$ using a modified logical argument:²

² Please note $\bar{\mathcal{T}}$ does not need to be the complement, in a set-theoretic sense, of $\mathcal{T}$.

If $\bar{\mathcal{T}} \& \mathcal{A}$ is true, then $\bar{\mathcal{O}}$.

We observe $\mathcal{O}$.

Therefore, $\bar{\mathcal{T}} \& \mathcal{A}$ is false.

Therefore, either $\mathcal{T}$ is true or $\mathcal{A}$ is false.

At the center of Meehl's (1978) critique is the fact that these important considerations are often ignored or misunderstood by psychologists, who merrily entertain vague theories without "*sufficient conceptual power (especially mathematical development) to yield the kinds of strong refuters expected by Popperians, Bayesians, and unphilosophical scientists in developed fields like chemistry.*" (p. 829). To make matters worse, the kind of testing psychologists often engage in involves a degenerate form of the modified logical argument given above. Specifically, they test *null hypotheses* that are trivially false and whose alternatives have little connection with any target theory:

... if you have enough cases and your measures are not totally unreliable, the null hypothesis will always be falsified, regardless of the truth of the substantive theory. (p. 822)

All sorts of competing theories are around, including my grandmother's common sense, to explain the nonnull statistical difference. (p. 824)

Meehl (1978) contrasted this problematic practice with the kind of testing found in the "hard" sciences, where the alternative hypothesis stands in close relation with a substantive candidate theory:

... the logical distance, the difference in meaning or content, so to say, between the alternative hypothesis and substantive theory $\mathcal{T}$ is so small that only a logician would be concerned to distinguish them. (p. 824)

The fact that Meehl’s critique is now over forty years-old presents itself as an opportunity to revisit some of its main points. At first blush, the fact that we encounter theoretical tours de force making a number of precise predictions (e.g., Cox & Shiffrin, 2017) suggests that things have improved considerably. Our point of contention here is that some of the progress in psychology as a whole is only apparent, given that it is predicated on a misunderstanding of the distinction between theory and auxiliary assumptions. More specifically, some elements of $\mathcal{A}$, whose specific purpose is to bridge the “deductive gap” between theoretical and observational statements, are assumed to belong to $\mathcal{T}$ and/or $\bar{\mathcal{T}}$ without proper justification. Consequently, these elements will not be scrutinized and refined by researchers, as envisioned by Lakatos (1976). Instead, they will be left untouched, as they are (illegitimately) seen as part of the theories’ “hard cores”.

One consequence of such misunderstandings is the spurious rejection of viable theoretical accounts and the latent-variable structures they propose. For example, Stephens, Dunn, and Hayes (2018) showed that previous rejections of single-process theories of syllogistic reasoning (i.e., $\bar{\mathcal{T}}$), taken as supporting a dual-process account (i.e., $\mathcal{T}$), hinge on auxiliary assumptions (e.g., a linear relation between latent processes and performance) that are simply taken for granted. When these assumptions are relaxed, it can be shown that the data at large are successfully captured by a single-process account (i.e., the different dependent variables can be described by single latent-variable).³ The problem identified by Stephens et al. (2018) is that previous attempts to test these theories illegitimately considered certain aspects of elements of $\mathcal{A}$ as part of $\mathcal{T}$ and/or $\bar{\mathcal{T}}$, which in turn results in a minimization of $\bar{\mathcal{O}}$. What this means is that single-process theories are being set up to fail, the end result being the false idea that a successful characterization of the data requires the involvement of two or more processes (i.e., latent variables).

Another consequence is the overstatement of support for certain theories. The

³ A more general analysis that includes other research domains such as category learning is given by Stephens, Matzke, and Hayes (2019).

empirical success of a conjunction $\mathcal{T} \& \mathcal{A}$ can be quite impressive when $\mathcal{O}$ is small. However, it is important to disentangle the contribution of the different elements in $\mathcal{T}$ and $\mathcal{A}$ to the size of $\mathcal{O}$. Otherwise, one might erroneously attribute the success to the theoretical statements in $\mathcal{T}$ when in fact most of the leg work is being accomplished by $\mathcal{A}$. One example of such a situation was recently identified by Jones and Dzhafarov (2014), who showed that the long-celebrated family of diffusion and ballistic-accumulator models, which are used to obtain precise joint descriptions of response frequencies and latencies, is not falsifiable until auxiliary parametric assumptions are introduced (e.g., the growth-rate variability between trials follows a Gaussian distribution). In other words, the empirical success of $\mathcal{T}$ alone is a sure thing.

The Problem of Coordination

To better understand the challenge of establishing and justifying a precise relationship between theoretical and observational statements, it is useful to frame our discussion within the context of measurement. In a nutshell, measurements are statements in which two quantities are placed in relation to each other according to an established set of rules. The problem of coordination refers to the circular relation that exists between the meaning of an unobservable quantity and its measurement:

1. Let X be a postulated quantity that is not directly observable.
2. Let Y be a directly-observable quantity that is connected to X by a coordination function $f(\cdot)$, such that $Y = f(X)$.
3. In order to measure X through Y , we must know the coordination function $f(\cdot)$ that maps the former into the latter. The problem is that this function is both *unknown* and *unknowable*. It cannot be established empirically, as that would require knowing joint instances of Y and X , the latter being the unobservable quantity that we were trying to measure in the first place. Therefore, the coordination function has to be

defined by a theory instead of being *discovered* through empirical means.

The problem of coordination is endemic to measurement in all of science, not just psychology. It is therefore instructive to consider the way in which the problem was conceived, approached, and provisionally settled in a context where measurement seems intuitively easy: the measurement of temperature. According to Mach (1896/1986), earlier efforts in scale construction in thermometry often framed these scales as attempts to approximate some platonic idea of temperature. In Mach's view, such a framing is misconceived, as it overlooks the fact that any conception of thermal states can only exist by virtue of an arbitrary definition that coordinates them with empirical facts. In other words, the two questions, "*What counts as a measurement of X?*" and "*What is X?*", cannot be addressed independently of each other (van Fraassen 2008, Chap. 5).

Although temperature is one of the physical magnitudes that people are most familiar with, it turns out that its measurement is far from trivial. In a body of work that spans over two-hundred years, we see that the process that led to the thermometers we know today involved a number of development stages, each associated with specific challenges (Chang, 2004; Mach, 1896/1986; Sherry, 2011). Similarly, people are deeply familiar with psychological concepts such as memory, attention, or intelligence due to the role these play in our language (Maraun, 1998). Despite this familiarity, their measurement so far has proven to be frustratingly difficult (e.g., Borsboom, Molenbergh, & Heerden, 2004; Maraun, 1998; Michell, 1999; Slaney, 2017).

Snapshots from the development of temperature scales

Initial attempts to measure temperature led to the development of ingenious instruments known as *thermoscopes*. These instruments were based on the observation that most liquids expand with heat, which meant that their registered volume in a sealed container such as a glass tube could be used to determine whether the temperature of *A* is less than, equal to, or greater than that of *B*. In other words, the thermoscope provides us

with an *ordinal scale* of temperature. Note that the development of thermoscopes hinges on the assumption that the relationship between volume and temperature is *monotonically increasing*.⁴ A function $f(\cdot)$ is monotonically increasing if $X_i \leq X_j \iff f(X_i) \leq f(X_j)$, $\forall i, j$.

From the use of thermoscopes, we learned that different substances left in the same environment long enough will end up at the same temperature (i.e., they will reach thermal equilibrium), even though they may feel differently. This is the so-called *zeroth law of thermodynamics* (Reif, 1965). The zeroth law enables thermoscopes to become thermometers by being calibrated against each other, as it allows the establishment of *fixed points* that can be used to set an origin as well as the units of the scale. But determining fixed points turned out to be extremely challenging, as exemplified by the discrepant measurements of the boiling point of water (e.g., 112.20 °C / 233.96 °F) when using different experimental apparatuses.⁵ The solution found by a commission appointed by the Royal Society of London in 1776 was to *define* the boiling point of water as the value recorded when exposing a thermoscope to the steam emerging from the water. The rationale was that measurements obtained under this definition showed little variation across different experimental setups.

Unfortunately, the availability of fixed points did not provide a solution to the problem of coordination: To suppose that the points on the scale measure temperature, rather than just volume, is to assume a *linear* coordination, such that $f(X) = \alpha X + \beta$, with α and β being free parameters. In fact, this assumption is rejected by the disagreements observed between thermometers using different liquids (e.g., water, alcohol, mercury, olive oil). These disagreements also showed that different substances have distinct coordination functions. This insight motivated the work of Henri Victor Regnault in the mid 1800s, who evaluated the merits of different thermometers by means of *consistency*

⁴ Interestingly, this relationship is nonmonotonic when the liquid is water.

⁵ It is worth noting that these discrepancies were not due to unreliable equipment or poor experimental design. In fact, many of the reported results have been successfully replicated by Chang (2008).

testing. Regnault's position was that, if there is an attribute that takes on some value, and we can specify different ways through which that value could be measured, then these measurements should all agree. Results showed a high degree of consistency between thermometers filled with gases such as hydrogen and air, namely linear relationships. Importantly, cases with poor agreement (e.g., sulfuric-acid gas) were deemed unsuitable as thermometric substances. Using the terminology introduced earlier, Regnault treated any observed inconsistencies to $\mathcal{A}$, but not $\bar{\mathcal{T}}$.

Regnault's identification of linear relations between the pressures of the gases didn't change the fact that their respective relations with temperature remained unknown. In other words, the problem of coordination stood unresolved. However, this shouldn't come as a surprise: Any attempt to measure X requires us to make a statement about what X is, something only a theory can ultimately provide. In the case of thermometry, the provisional theoretical solution came in the form of the kinetic theory of heat, which *redefines* temperature as the average kinetic energy of the particles of an ideal gas. Based on this theory, it became possible to establish a linear relationship between the volume of (near-) ideal gases and temperature, and subsequently use these results to precisely calibrate thermometers based on other substances such as mercury (e.g., Beattie, Blaisdell, Kaye, Gerry & Johnson, 1941).

But even though the problem of coordination requires a theoretical response, this does not imply that experimental studies like the ones described above are in some way secondary. As discussed by Chang (2004, Chap. 5), the circularity inherent to the problem of coordination requires researchers to engage in 'epistemic iterations', a process in which successive stages of experimental knowledge and theoretical understanding, each building on the preceding stages: From noticing that our feelings of warmth are (somewhat imperfectly) tracked by the volume of substances, to the stabilization and refinement of experimental procedures. Each development makes the subsequent possible, furthering our scientific goals. Importantly, these goals are not necessarily reducible to the pursuit of

Figure 1. Analogy between thermometry and psychological measurement.

some kind of realist aspirations (e.g., Kellen, 2019). For instance, note how the definition of temperature offered by kinetic theory is nothing more than an *abstraction* within the theory’s logical space that happens to characterize the outcomes of a stabilized procedure. The average kinetic energy isn’t something that exists, in the same sense that the “average person” doesn’t exist (see Van Fraassen, 2008, Chap. 5).

The Problem of Coordination in Psychology

As in thermometry, so in psychology: latent variables give rise to observed variables by means of a coordination function. We can draw an analogy between the measurement of the temperature of a substance with the measurement of a psychological attribute of a person. In measuring temperature, a scientist chooses a procedure $\mathcal{P}$ that yields an observation Y that is theoretically related to the unknown temperature X . In the case of a thermometer, Y is the volume of the enclosed thermometric substance. Formally, $Y = f_{\mathcal{P}}(X)$. The process is the same for psychological attributes. A scientist chooses a procedure $\mathcal{P}'$ that yields an observation Y' that is theoretically related to the value X' of the attribute. Again, $Y' = f_{\mathcal{P}'}(X')$. Figure 1 illustrates this analogy. For a concrete example, take the case of *memory*, a capacity that can be defined in broad strokes as the ability to remember. The quantity, or accuracy, of remembering can be easily recorded using a variety of experimental procedures such as free recall or single-item recognition.

Drawing a closer analogy, these procedures include a study phase analogous to *heating*, in the sense that better studied items become more memorable (hotter). This increase is tracked by the responses given in the test phase (e.g., hit rates or recall rates), the same way that changes in the temperature of substance are tracked by its volume.

Looking back at our historical discussion of temperature, one might expect to find similar efforts in the development of measures of memory. But one would be disappointed. Nor do we find the kind of virtuous circularities that Chang (2004) alludes to when discussing the process of epistemic iterations. There are three issues that remain unresolved:

1. *Whether different measurement procedures measure the same or different attributes.*

For example, many researchers might feel that apparently dissimilar procedures, such as cued recall and single-item recognition, may measure different attributes (i.e., types of memory), but fewer might hold that this is the case for more similar procedures, such as recognition memory applied to different classes of items (e.g., different random word lists; or pictures of faces vs. pictures of houses). This issue was at the center of the implicit/explicit memory debates that once dominated memory research (e.g., Schacter, Chiu, & Ochsner, 1993).

2. *Even if two procedures are taken to measure the same attribute, they may have different coordination functions.* This is equivalent to having a thermometer whose measurements hinge on the thermometric substance inside its vessel (e.g., water, mercury, air). In memory research, when very different procedures are compared, such as cued recall and recognition, no one would suppose that a score of say 70% correct responses on each task can be taken as measuring the same strength of memory. However, when more similar procedures are compared, such claims are frequently made (as we will see below).

3. *Even if two procedures are taken to have the same coordination function, this function*

remains unknown. We encountered this issue earlier when referring to Regnault’s consistency tests. In psychology, this issue leads to the interpretability problems discussed by Loftus (1978) and more recently by Wagenmakers, Krypotos, Criss, & Iverson (2012) and Garcia-Marques, Garcia-Marques, and Brauer (2014), in which the understanding of data in terms of interactions and main-effects obtained with ANOVA-type decompositions is shown to collapse under alternative coordination functions.

Operationalism about coordination

One of the reasons why the development of measures of psychological attributes differs from what is found in other domains such as thermometry is the fact that psychologists have by and large adopted (even if tacitly) a peculiar view of measurement known as *operationalism* (Bridgman, 1927). According to it, measurement is simply viewed as “*any precisely specified operation that yields a number*” (Dingle, 1950, p. 11).

Operationalism was popularized in psychology by S. S. Stevens when he distinguished four different types of scales (nominal, ordinal, interval, and ratio) on the basis of the operations that produced them (Stevens, 1946). It is this operationalist view of psychological measurement that justifies the assertion that a sum-score of numerical ratings constitutes some measure of “something” (e.g., of attitudes, intelligence; for critiques, see Green, 2001; Koch, 1992; Leahey, 1980; Michell, 1999).

It is also this operationalist view that drives researchers to make strong theoretical claims based on the fact that two different measures do not change in exactly the same way across experimental conditions. For instance, the observation that performance in one memory task is affected by an experimental manipulation — whereas performance in a second task isn’t — has been interpreted by many as evidence for the existence of separate memory systems (for a critical review, see Newell, Dunn & Kalish 2011). The argument being that, if the measures behave differently, it must be because they measure different

things. *Note how this line of reasoning runs completely counter to Regnault's:* In this particular example, we are dealing with the conjunction $\bar{\mathcal{T}} \& \mathcal{A}$ where $\bar{\mathcal{T}}$ is the theory that there is a single attribute measured by both tasks, and part of $\mathcal{A}$ is the assumption that both tasks have the same linear coordination function. When both hold, we should expect both measures to register exactly the same changes in performance across conditions. But if they do not (i.e., an interaction is observed), it follows that either $\bar{\mathcal{T}}$ or $\mathcal{A}$ or both are at fault. Whereas Regnault would interpret such a result as a failure of $\mathcal{A}$, a psychologist operating under the tenets of operationalism would choose to reject $\bar{\mathcal{T}}$.

The reason behind this disagreement is that operationalism ignores the problem of coordination and effectively enforces a complete disconnect with natural reality: We are no longer concerned with the ability of Y capturing some unobservable X — quantity Y has become its own measure. Also, it follows that any concept of validity has to go out the window, as we are no longer able to evaluate the consistency of measurements coming from different instruments: They measure different things simply because they involve different operations — a vicious circularity.⁶ Under operationalism, many of the achievements observed in thermometry would have not been possible. In fact, operationalism would legitimize absurd claims, such as that temperature is a multi-valued magnitude, on the grounds that the application of distinct thermometric instruments (e.g., pyrometers, thermocouples, and thermometers) yield different numbers prior to calibration.

Case Study: Face-Inversion Effect

For a more concrete example on how the choice of coordination affects theory testing, let us consider the case of the *face-inversion effect* (Yin, 1969), which is illustrated in Figure 2: According to the face-inversion effect, people's ability to recognize faces is more affected by inversion than other pictures of mono-oriented objects such as houses. One

⁶ In reaction to these problems, one might propose that different operationalizations are somehow capturing distinct facets of larger '*super-concept*'. Unfortunately, such a move offers no real solution: If the super-concept has its own operationalization, then the other measures are redundant. However, if it does not have an operationalization, then the super-concept has to be deemed meaningless.

Figure 2. Top Panel: Illustration of the paradigm used by Yin (1969), in which pictures of faces and houses are studied upright or inverted, and later tested in a two-alternative forced-choice recognition task. Bottom Panel: The observed accuracy, and a characterization in which inverted/upright faces and houses have the same memory strength.

common interpretation of the face-inversion effect is that there is “something special” about the way we process facial stimuli (for discussions, see Dunn & Kalish, 2018, Chap. 2; Loftus, Oberg, & Dillon, 2004).

In a typical experiment, participants study items that may be either pictures of faces or houses and either upright or inverted. This conforms to a 2×2 factorial design. It is argued that if there is nothing “special” about faces then the effect of inversion should be the same for both faces and houses. Using the analogy with temperature, while upright faces and houses may have different ‘temperatures’ after being ‘heated’ by study (because faces are easier to study – absorb ‘heat’ more readily – than houses), the ‘cooling’ that comes from inverting them should lower their final temperatures by the same amount. Let

X be the memory strength of upright faces and let Δ_{house} be the difference in memory strength for houses and let $\Delta_{inverted}$ be the (negative) change in strength due to inversion. Then

$$\begin{aligned} X_{\text{face,upright}} &= X, \\ X_{\text{face,inverted}} &= X + \Delta_{\text{inverted}}, \\ X_{\text{house,upright}} &= X + \Delta_{\text{house}}, \\ X_{\text{house,inverted}} &= X + \Delta_{\text{inverted}} + \Delta_{\text{house}}. \end{aligned}$$

It follows that

$$X_{\text{face,upright}} - X_{\text{face,inverted}} = X_{\text{house,upright}} - X_{\text{house,inverted}}.$$

Now, let $f(\cdot)$ be a coordination function. If it is linear, then the *equality* above holds. That is,

$$f(X_{\text{face,upright}}) - f(X_{\text{face,inverted}}) = f(X_{\text{house,upright}}) - f(X_{\text{house,inverted}}).$$

The accuracy data in Figure 2 show that $f(X_{\text{face, upright}}) - f(X_{\text{face, inverted}})$ is larger than $f(X_{\text{house, upright}}) - f(X_{\text{house, inverted}})$, indicating that the effect of inversion on memory is greater for faces. This result, which would be captured by interaction effects in a linear model (e.g., via ANOVA or regression), is used to support the claim that there is something special about the way we process and remember faces.

But as discussed above, assuming a single linear coordination for both faces and houses lacks justification. One possible remedy is to retreat to the more modest idea that the relationship between accuracy and memory strength is monotonically increasing. This move is equivalent to admitting that we only have a ‘memory thermoscope’ available. To begin with, we may suppose that faces and houses have the same monotonically increasing coordination function, $f(\cdot)$. In this case, the equality of differences is replaced by the

following two implications, which are comprised of *inequalities*:

$$f(X_{\text{face,upright}}) - f(X_{\text{face,inverted}}) \geq 0 \iff f(X_{\text{house,upright}}) - f(X_{\text{house,inverted}}) \geq 0,$$

and

$$f(X_{\text{face,upright}}) - f(X_{\text{house,upright}}) \geq 0 \iff f(X_{\text{face,inverted}}) - f(X_{\text{house,inverted}}) \geq 0.$$

Because the second of these is violated by the data shown in Figure 2, we could conclude that there is ‘something special’ about faces. However, this conclusion depends crucially on the assumption of a common coordination function. This too may be relaxed by assuming that faces and houses have different coordination functions. In this case, it is perfectly possible for the memory strength of studied houses and faces to be the same, and for them to be equally affected by inversion. Under this assumption, the potential violation of $\bar{\mathcal{T}}$ is attributed to $\mathcal{A}$ and we can no longer stand by our earlier statement that there is something special about our memory for faces simply because they show a larger inversion effect in a recognition task. It is analogous to saying that there is something special about the temperature of water relative to that of mercury simply because these substances expand differently when exposed to the same heat, effectively ignoring the problem of thermoscope/thermometer calibration (also, we know of no attempts to establish fixed points). Finally, note that a similar case can be made regarding the comparison of different groups of individuals: For instance, Laguesse, Dormal, Biervoye, Kuefner and Rossion (2012) compared the size of the face inversion effect in two groups of participants, finding a larger effect for one group than the other which they interpreted as demonstrating that the groups differed in their sensitivity to inversion. While such an interpretation is not illegitimate, it once again depends upon an unjustified assumption of a linear coordination function (for a detailed discussion, see Dunn & Kalish, 2018).

The face-inversion effect illustrates how the problem of coordination can severely

affect the conclusions that can be drawn from observed patterns of data. A theory $\bar{\mathcal{T}}$ is proposed with goal of serving as a kind of null hypothesis – a statement about the world that the researcher seeks evidence against. This may take the form of proposing that there is only one kind of memory, or that the effects of inverting an image on memory for that image are the same for faces and houses, or that the effect differs across groups of people (e.g., younger and older adults). As noted earlier, such a point hypothesis is vacuous because it rules out almost nothing — $\bar{\mathcal{O}}$ corresponds to a single point. By coupling $\bar{\mathcal{T}}$ with a linear coordination function, the minimal size of $\bar{\mathcal{O}}$ is thereby maintained and, as more data are collected and the point null is rejected, the theory of interest $\mathcal{T}$ is apparently supported. But, by considering more general and more plausible candidates for coordination functions, the outcome predicted by $\bar{\mathcal{O}}$ is no longer a single point and so is less readily falsified. This response is nothing more than the textbook prescription that researchers should examine all aspects of their experimental setups (included in $\mathcal{A}$) before drawing theoretically significant conclusions from the observed data. Finally, the face-inversion effect demonstrates that the way psychologists engage with attributes such as memory is completely at odds with the experimental and theoretical developments found in the case of thermometry, a situation that partly due to the persisting influence of operationalism. It is also at odds with Meehl’s (1978) call for a more careful consideration of the link between testing outcomes and theoretical statements.

What Can Be Done?

The testability of a theory is co-determined by the structure of the latent variables it postulates (e.g., their involvement in different dependent variables), as well as the coordination functions imposed (Dunn & Anderson, 2018). The latter determine which transformations can and cannot effect on the mapping of the latent structure onto observations. This relationship between latent structure and coordinations shows that the problem of coordination is not limited to measurement — it is also a problem for theory

testing. At this point, it is not clear whether we can overcome the many problems of coordination found in psychology. On one hand, the replication of some of the achievements found in thermometry — the establishment of fixed points and calibration of scales — seems extremely unlikely.⁷ On the other hand, some of the attributes psychologists are interested in (e.g., intelligence, anxiety, dominance) have extremely complicated ‘grammars’, which can seriously compromise their measurability (for a discussion, see Maraun, 1998). But if the history of psychological measurement tells us anything, is that powerful advances can be achieved when rigorous thinkers are willing to put some intellectual muscle into the enterprise (Rozeboom, 1966).

In any case, it is the responsibility of researchers to specify the nature of their assumed coordination functions. If the interpretation of the empirical results offered by the researcher depends upon a specific coordination function, such as linear, then it is only reasonable to expect that this be made explicit, as any assumption would be. However, in so doing, it should become apparent that some hyper-specific coordinations (such as linear) are only rarely justifiable.⁸ For this reason, researchers may propose coordinations with fewer unjustified commitments in alternative or complementary methods that assume different coordinations (for a recent example, see Kellen, Steiner, Davis-Stober, & Pappas, 2020). Depending on the proposed coordination function(s), the outcome $\bar{O}$ associated with $\bar{T}$ will change accordingly. Point hypotheses involving relationships of equality or additivity will not survive any departure from linearity – other implications will have to be worked out.⁹

Restricting ourselves to the assumption that coordination functions are monotonic

⁷ Such achievements will most likely require a paradigm-shift in the way we conceive and implement our research programs. For instance, towards something akin to Duncan Luce’s psychophysical program (for an overview, see Steingrimsson (2016)).

⁸ Researchers intent on using linear coordination functions in the absence of a comprehensive theory might be wise to study the history of thermometry more closely, and judge whether the body of psychological research that they are building on includes analogous achievements (e.g., determination of fixed points, calibrations).

⁹ Meehl (1978, 1990) makes a similar point when arguing for the use of interval predictions.

seems to be a reasonable option in most contexts.¹⁰ Fortunately, there are a number of readily-available methods that only require the assumption of monotonicity. For instance, *Signed-Difference Analysis* (Dunn & Anderson, 2018; Dunn & James, 2003) can be used to identify the structural properties of a theory’s latent variables that hold when assuming monotonic coordinations. These structural properties are observable in the *directions* (+, −, and 0) in which the observable variables can jointly change across conditions in a given experimental design. These differences are described by *sign vectors*. Note that one special case of signed-difference analysis is *State-Trace Analysis* (Bamber, 1979; Dunn & Kalish, 2018), which focuses on the question of whether two dependent variables can be described by a single latent variable with monotonic coordinations. Stephens et al. (2018) applied signed-difference analysis to a corpus of studies used to compare single- and dual-process models of syllogistic reasoning. For example, participants we requested — under deductive or inductive instructions — to judge syllogisms that varied dichotomously in terms of their both *validity* and *causal consistency*. The endorsement rates coming out these studies can be boiled down to eighty-one sign vectors. Each element of the vector corresponding to the sign of the difference in endorsement rates between causal-consistency conditions. For instance, consider the sign vector:

$$\begin{pmatrix} + & - & - & + \\ \text{Deduction} & \text{Deduction} & \text{Induction} & \text{Induction} \\ \text{Valid} & \text{Invalid} & \text{Valid} & \text{Invalid} \end{pmatrix}$$

Under monotonic coordinations, different single-process and dual-process theories create different partitions of sign vectors into $\mathcal{O}$ and $\bar{\mathcal{O}}$. For instance, the sign vector described above cannot be captured by any single-process theory and most dual-process theories. The inconsistencies between the observed differences and each theory’s partition of sign vectors were tested statistically using the order-constrained inference method proposed by

¹⁰ There are of course circumstances in which monotonicity is not defensible option (for a discussion, see Dunn, Kalish, & Newell, 2014).

Kalish, Dunn, Burdakov, and Sysoev (2016).¹¹ Results showed that most theories are rejected, including all testable dual-process theories.

One concern often raised when discussing the use of weaker coordination assumptions is that theories become harder to test. Our immediate response to such concerns is to point out that the ‘increased testability’ that one might be reluctant to forfeit was obtained through illegitimate means. Having said that, it is incorrect to assume that one cannot devise strong tests in the absence of stronger coordination assumptions. This can be achieved through the use of richer experimental designs.¹² For example, Kellen, Winiger, Dunn, and Singmann (2019) and McCausland, Davis-Stober, Marley, Park, and Brown (2020) tested (and upheld) the postulates of Signal Detection Theory and Random Utility Theory, using experimental designs for which the $\mathcal{O}$ s of both theories are minuscule.

Another reason why researchers might feel reluctant to engage with weaker coordinations is the fact the *equations* used in traditional ANOVA-type decompositions (main effects and interactions) are not suited to handle the now-predicted *inequalities*. In response, we raise three points: First, many of the psychological theories that researchers engage with make predictions at the *ordinal* level (i.e., they predict one or more inequalities). In fact, we would argue that the ANOVA-type language imposed by many standard statistical methods has been an hindrance to psychologists in the sense that it very often interferes with our ability to think clearly and speak plainly about theoretical predictions and their encounter with data (e.g., Hatz, Park, McCarty, McCarthy, & Davis-Stober, in press; Rouder, Haaf, Davis-Stober, & Hilgard, 2019). Second, the joint test of the order constraints postulated by a theory allows for powerful omnibus tests that wouldn’t be possible with a piecemeal approach where multiple main-effect and interaction

¹¹ For other methods, see Davis-Stober (2009), Heck and Davis-Stober (2019), and Regenwetter and Cavagnaro (2019).

¹² Also, knowledge of $\bar{\mathcal{O}}$ under monotonic coordinations can be used to tailor experiments to the sole purpose of observing some of its elements (i.e., devise a critical test; see Kellen, 2019; Kellen et al., 2020). For example, the best-performing theory encountered by Stephens et al. (2017) is unable to predict the sign vector (+ − − +) discussed above. This fact led them to conduct a follow-up study with sole purpose of trying to observe data consistent with this vector.

tests would be necessary (for an excellent example, see Iverson, 2006). Third, the resources needed to apply order-constrained inference methods are now readily available (see Heck & Davis-Stober, 2019; Kalish et al., 2017; Regenwetter & Cavagnaro, 2019).

Conclusion

According to Lakatos (1976), auxiliary assumptions $\mathcal{A}$ serve as a ‘protective belt’ around a theory. In a research program that is theoretically and empirically progressive by continuously resolving previous anomalous findings and confirming novel predictions, these auxiliary assumptions are constantly being analyzed and refined. Meehl (1978) despaired of psychologists attempting to test theories against null-hypotheses that are trivially false. The consequence of such practice is the development of degenerative research programs in which the rise and fall of theories is more reflective of fads and fashion than meaningful theoretical development.

In the present work, we argue that some of the problematic practices criticized by Meehl are still present, despite apparent progress. We attribute this persistence to a continued neglect towards the problem of coordination, which leads to spurious support towards/against certain theories. Using the history of thermometry as reference, we showed that the linear coordinations typically employed are often implausible and never justified, giving researchers a false notion of precision, falsifiability, and empirical support (Chang, 2004; Sherry, 2011). This neglect corrupts the Lakatosian process of scientific development, as it prevents researchers from investigating the impact of their assumed coordinations by incorrectly attributing their failures to the theories. As a first step, researchers should restrict themselves to only assuming monotonic coordinations that do not necessarily generalize across procedures. Fortunately, there is a rich toolbox of methods that can operate under such minimal structural constraints. The use of such methods, along with more plausible coordination functions, will result in less falsifiable $\bar{\mathcal{T}}$. This in turn will render theories $\mathcal{T}$ appropriately falsifiable, removing an important impediment to

cumulative theoretical progress in psychology.

References

- Bamber, D. (1979). State-trace analysis: A method of testing simple theories of causation. *Journal of Mathematical Psychology*, *19*, 137–181.
- Beattie, J. A., Blaisdell, B. E., Kaye, J., Gerry, H. T., & Johnson, C. A. (1941). An experimental study of the absolute temperature scale viii. the thermal expansion and compressibility of vitreous silica and the thermal dilation of mercury. *Proceedings of the American Academy of Arts and Sciences*, *74*, 371–388.
- Birnbaum, M. H. (2011). Testing mixture models of transitive preference: Comment on Regenwetter, Dana, and Davis-Stober (2011). *Psychological Review*, *118*, 675–683.
- Borsboom, D., Mellenbergh, G. J., & Van Heerden, J. (2004). The concept of validity. *Psychological Review*, *111*, 1061–1071.
- Bridgman, P. W. (1927). *The logic of modern physics*. New York: Macmillan.
- Chang, H. (2004). *Inventing temperature: Measurement and scientific progress*. Oxford, UK: Oxford University Press.
- Chang, H. (2008). The myth of the boiling point. *Science Progress*, *91*, 219–240.
- Cox, G. E., & Shiffrin, R. M. (2017). A dynamic approach to recognition memory. *Psychological Review*, *124*, 795–860.
- Davis-Stober, C. P. (2009). Analysis of multinomial models under inequality constraints: Applications to measurement theory. *Journal of Mathematical Psychology*, *53*, 1–13.
- Dingle, H. (1950). A theory of measurement. *The British Journal for the Philosophy of Science*, *1*, 5–26.
- Duhem, P. M. M. (1954). *The aim and structure of physical theory*. Princeton, NJ: Princeton University Press.
- Dunn, J. C., & Anderson, L. (2018). Signed difference analysis: Testing for structure under monotonicity. *Journal of Mathematical Psychology*, *85*, 36–54.
- Dunn, J. C., & James, R. N. (2003). Signed difference analysis: Theory and application. *Journal of Mathematical Psychology*, *47*, 389–416.

- Dunn, J. C., & Kalish, M. L. (2018). *State-trace analysis*. New York: Springer.
- Dunn, J. C., Kalish, M. L., & Newell, B. R. (2014). State-trace analysis can be an appropriate tool for assessing the number of cognitive systems: A reply to Ashby (2014). *Psychonomic Bulletin & Review*, *21*, 947–954.
- Garcia-Marques, L., Garcia-Marques, T., & Brauer, M. (2014). Buy three but get only two: The smallest effect in a 2×2 ANOVA is always uninterpretable. *Psychonomic Bulletin & Review*, *21*, 1415–1430.
- Green, C. D. (2001). Operationism again: What did Bridgman say? What did Bridgman need? *Theory & Psychology*, *11*, 45–51.
- Hatz, L. E., Park, S., McCarty, K., McCarthy, D., & Davis-Stober, C. P. (in press). Young adults make rational sexual decisions. *Psychological Science*.
- Heck, D. W., & Davis-Stober, C. P. (2019). Multinomial models with linear inequality constraints: Overview and improvements of computational methods for bayesian inference. *Journal of Mathematical Psychology*, *91*, 70–87.
- Iverson, G. J. (2006). An essay on inequalities and order-restricted inference. *Journal of Mathematical Psychology*, *50*, 215–219.
- Jones, M., & Dzhamfarov, E. N. (2014). Unfalsifiability and mutual translatability of major modeling schemes for choice reaction time. *Psychological Review*, *121*, 1–32.
- Kalish, M. L., Dunn, J. C., Burdakov, O. P., & Sysoev, O. (2016). A statistical test of the equality of latent orders. *Journal of Mathematical Psychology*, *70*, 1–11.
- Kellen, D. (2019). A model hierarchy for psychological science. *Computational Brain & Behavior*, *2*, 160–165.
- Kellen, D., Steiner, M. D., Davis-Stober, C. P., & Pappas, N. R. (2020). Modeling choice paradoxes under risk: From prospect theories to sampling-based accounts. *Cognitive Psychology*, *118*, 101258.
- Kellen, D., Winiger, S., Dunn, J. C., & Singmann, H. (2019). *Testing the Foundations of Signal Detection Theory in Recognition Memory* (preprint). PsyArXiv. doi:

10.31234/osf.io/p5rj9

- Koch, S. (1992). Psychology's Bridgman vs Bridgman's Bridgman: An essay in reconstruction. *Theory & Psychology*, 2, 261–290.
- Laguesse, R., Dormal, G., Biervoye, A., Kuefner, D., & Rossion, B. (2012). Extensive visual training in adulthood significantly reduces the face inversion effect. *Journal of Vision*, 12, 14.
- Lakatos, I. (1976). Falsification and the methodology of scientific research programmes. In S. G. Harding (Ed.), *Can theories be refuted? Essays on the Duhem-Quine thesis* (pp. 205–259). Springer.
- Leahey, T. H. (1980). The myth of operationism. *The Journal of Mind and Behavior*, 1, 127–143.
- Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, 6, 312–319.
- Loftus, G. R., Oberg, M. A., & Dillon, A. M. (2004). Linear theory, dimensional theory, and the face-inversion effect. *Psychological Review*, 111, 835–863.
- Mach, E. (1896/1986). *Principles of the theory of heat historically and critically elucidated*. Dordrecht, Holland: D. Reidel Publishing Company.
- Maraun, M. D. (1998). Measurement as a normative practice: Implications of Wittgenstein's philosophy for measurement in psychology. *Theory & Psychology*, 8, 435–461.
- McCausland, W. J., Davis-Stober, C., Marley, A., Park, S., & Brown, N. (2020). Testing the random utility hypothesis directly. *The Economic Journal*, 130, 183–207.
- Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*, 46, 806–834.
- Meehl, P. E. (1990). Appraising and amending theories: The strategy of Lakatosian defense and two principles that warrant it. *Psychological Inquiry*, 1, 108–141.
- Michell, J. (1999). *Measurement in psychology: A critical history of a methodological*

- concept* (Vol. 53). Cambridge University Press.
- Newell, B. R., Dunn, J. C., & Kalish, M. (2011). Systems of category learning: Fact or fantasy? In *Psychology of learning and motivation* (Vol. 54, pp. 167–215). Elsevier.
- Quine, W. V. O. (1963). Two dogmas of empiricism. In *From a logical point of view* (p. 20-46). New York: Harper & Row.
- Regenwetter, M., & Cavagnaro, D. R. (2018). Tutorial on removing the shackles of regression analysis: How to stay true to your theory of binary response probabilities. *Psychological Methods*, 24, 135-152.
- Reichenbach, H. (1958). *The philosophy of space and time*. New York: Dover.
- Reif, F. (1965). *Fundamentals of thermal and statistical physics*. New York: McGraw-Hill.
- Rouder, J. N., Haaf, J. M., Davis-Stober, C. P., & Hilgard, J. (2019). Beyond overall effects: A bayesian approach to finding constraints in meta-analysis. *Psychological Methods*, 24, 606-621.
- Rozeboom, W. W. (1966). Scaling theory and the nature of measurement. *Synthese*, 16, 170–233.
- Schacter, D. L., Chiu, C.-Y. P., & Ochsner, K. N. (1993). Implicit memory: A selective review. *Annual Review of Neuroscience*, 16, 159–182.
- Sherry, D. (2011). Thermoscopes, thermometers, and the foundations of measurement. *Studies in History and Philosophy of Science Part A*, 42, 509–524.
- Slaney, K. (2017). *Validating psychological constructs: Historical, philosophical, and practical dimensions*. New York: Springer.
- Steingrimsson, R. (2016). Subjective intensity: Behavioral laws, numerical representations, and behavioral predictions in Luce’s model of global psychophysics. *Journal of Mathematical Psychology*, 75, 205–217.
- Stephens, R. G., Dunn, J. C., & Hayes, B. K. (2018). Are there two processes in reasoning? the dimensionality of inductive and deductive inferences. *Psychological Review*, 125, 218.

- Stephens, R. G., Matzke, D., & Hayes, B. K. (2019). Disappearing dissociations in experimental psychology: Using state-trace analysis to test for multiple processes. *Journal of Mathematical Psychology*, *90*, 3–22.
- Stevens, S. S. (1946). On the theory of scales of measurement. *Science*, *103*, 677–680.
- Tal, E. (2017). Measurement in science. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Fall 2017 ed.). <https://plato.stanford.edu/archives/fall2017/entries/measurement-science/>.
- Van Fraassen, B. C. (2008). *Scientific representation: Paradoxes of perspective*. Oxford, UK.: Oxford University Press.
- Wagenmakers, E.-J., Krypotos, A.-M., Criss, A. H., & Iverson, G. (2012). On the interpretation of removable interactions: A survey of the field 33 years after Loftus. *Memory & cognition*, *40*, 145–160.
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*, 141.