A Model Hierarchy for Psychological Science

David Kellen

Syracuse University

Author Note

David Kellen, Department of Psychology. I thank Mike L. Kalish, Henrik Singmann, and Andrea L. Turnbull for valuable comments on an earlier draft. Correspondence should be sent to David Kellen (davekellen@gmail.com).

Abstract

Lee et al. (2019) provided a comprehensive list of recommendations for modelers that aims at improving the robustness of their results. Drawing from the literature on philosophy of science, the present commentary argues for a broader view of modeling that considers the different roles that they play in our scientific practices. Following Suppes (1966), I propose a model hierarchy and discuss the distinct issues that arise at each of its levels. The benefit of a hierarchy of this kind is that it can aid researchers in better understanding the different challenges that they face.

Keywords: cognitive modeling, robustness, model comparison

A Model Hierarchy for Psychological Science

There is much to like in Lee et al.'s (2019) recommendations. However, I cannot avoid thinking that something important is missing. If I had to express my concern in one sentence, it would be that '*methodology is downstream of conceptual analysis*'. In order to discuss practices in modeling and attribute them some utility, we first need to have a clear understanding of the different roles that models can play in our scientific practices. In psychological science, the term 'model' is often understood in a 'statistical sense', denoting a parametric function that yields probability distributions over a well-defined sample space. In this commentary, I will rely on a broader understanding found in the philosophy of science literature, in which a model is *any representation of some 'target system'* (for reviews, see Frigg & Hartmann, 2018; Morgan & Morrison, 1999).

A good starting point for discussing the roles that models can take is Patrick Suppes' 1966 seminal paper *Models of Data*. In this short piece, Suppes argues that there is no direct contact between a researcher's theory and the world. Instead, their relationship is mediated by multiple models that interact with each other in a somewhat hierarchical fashion:

... a whole hierarchy of models stands between the model of the basic theory and the complete experimental experience. Moreover, for each level of the hierarchy there is a theory in its own right. Theory at one level is given empirical meaning by making formal connections with theory at a lower level. (...) It is precisely the fundamental problem of scientific method to state the principles of scientific methodology that are to be used to answer these questions — questions of measurement, of goodness of fit, of parameter estimation, of identifiability, and the like. (Suppes, 1966, pp. 260)

Figure 1 illustrates one possible Suppes-type hierarchy (for alternative formulations, see Giere, 2004, 2010; Mayo, 1996). At the top and bottom levels of the hierarchy we have the

researcher and the world, with the former holding theories and hypotheses about the latter. There is a deductive gap between the two, however. We cannot deduce a given observation O from a theory or hypothesis alone. To do so, we also need to introduce a number of auxiliary assumptions. The implication is that no set of observations (O or not- O) is by itself sufficient to change our beliefs about any scientific theory or hypothesis (Duhem, 1954). This is where the models in the hierarchy come in. By providing *representations* of the theory, the data, and the experiment, they allow the researcher to bridge the deductive gap between her hypotheses and the world. On one hand, models enable the application of an abstract theory to some real system (e.g., observable behavior in a given experiment). On the other, they allow us to represent this real system in a more abstract and idealized manner that is amenable to theoretical description (Frigg & Hartmann, 2018; Morgan & Morrison, 1999).¹ Finally, note that the arrows illustrate the crosstalk between the different levels of the hierarchy - they denote the presence of joint constraint and the fact that learning occurs via an iterative process. All levels of the hierarchy are to be continuously revisited through ‘loosening’ and ‘tightening’ activities, with new ideas being generated and tested (for discussions, see Box, 1976; Fiedler, 2018).

The differentiation of models established by the proposed hierarchy is helpful in two ways. First, it highlights the *nature of representation*: To represent something is to engage in an *intentional act* of relating two objects according to an established set of *goals* and *criteria*. It involves somebody (the researcher) bestowing the representing object (the model) a given function. Cases of misrepresentation are those in which the representing object fails to serve this function and live up to the established criteria (for overviews, see Regt, Henk, Sabina, & Eigner, 2009; Van Fraassen, 2008).² Second, the proposed

¹ When discussing representations, it is often important distinguish between idealization and abstraction. According to Cartwright (1983), idealization is a deliberate distortion that facilitates the representation of a given system (e.g., imposing a given functional form), whereas abstraction involves the omission of some of those system’s properties (e.g., omitting some form of internal noise or variability).

² A subway map deviates in many ways from the actual subway system. But because its goal is to tell you how to go places, we can only refer to it as a misrepresentation if it fails to take you from point A to B.

differentiation makes the problem of appraisal and error checking more tractable, allowing researchers to investigate the possibility of assumption violations separately at each level. We can then refer to our modeling efforts as *robust* when are shown to withstand our best efforts to dismiss them (Mayo, 1996).

In what follows, I will discuss the three different modeling levels included in the proposed hierarchy.

Theoretical Models

Ours is a world in which people manifest a certain number of *capacities*: They see and hear, have beliefs and preferences, remember things, and engage in all kinds of actions based on a variety of reasons. In light of such capacities, theory development consists of attempts to explain them by making some reference to their causal pre-conditions (e.g., the underlying mechanisms; see Cummins, 2000; Trigg & Kalish, 2011; Weiskopf, 2011). Behavioral results (e.g., ‘effects’), even those manifesting law-like regularities, take on an instrumental role in the sense that their value lies in their ability to inform us on the nature of these pre-conditions. For example, the value of *list-strength effects* observed in memory tasks is given by their ability to discard the notion that memory judgments are based on a ‘global match’ between a target item and the presumptively stored memory traces (e.g., Shiffrin, Ratcliff, & Clark, 1990).³ Failure to establish a clear and theory-informed relationship between capacities and behavioral results can lead researchers astray. One drastic example of this can be found in the literature on syllogistic reasoning where, for over thirty years, researchers mistook ANOVA interaction terms as direct measures of change in reasoning ability and developed many different models around these

³ This does not exclude the possibility of certain effects ‘gaining a life of their own’ (Hacking, 1983; Mayo, 1996). Consider the case of the fan effect, which was originally used in the development of ACT-R (Anderson, 1974). Since then the fan effect has grown beyond this role, to the point that it has even been used as a measurement tool in the study of group and individual differences (e.g., Cantor & Engle, 1993; for a discussion, see Garcia-Marques & Ferreira, 2011).

presumed changes (for a discussion, see Rotello, Heit, & Dube, 2015).⁴

Differences in researchers' goals and criteria are responsible for the many classes of theoretical models found in the literature. For instance, we often make the distinction between *process models* (e.g., Decision Field Theory; Busemeyer & Townsend, 1993) and *measurement models* (e.g., Multinomial Processing Tree models; Batchelder, 2010). In some cases, theoretical models are developed with the intent of only capturing the general patterns in the data rather than provide a fine-grained account (e.g., Navarro, 2019; Shiffrin & Nobel, 1997). Some models are to be applied to a wide range of conditions (e.g., Newell, 1990), whereas others are designed with a single experimental paradigm in mind (Batchelder, 2010). As a result of the many lamentations on the current state of psychological science (e.g., Coyne, 2016; Pashler & Wagenmakers, 2012), one might be tempted to think that this diversity among theoretical models is yet another sign of conceptual confusion or immaturity of our field. In fact, it's the opposite. All of these different types of models provide complementary contributions with regard to experimental knowledge, measurement, and theory development that altogether have given us the leverage to make significant developments in many different domains (for discussions, see Garcia-Marques & Ferreira, 2011; Weiskopf, 2011). The overwhelming success of these modeling enterprises becomes especially clear when viewed through the lens that progress in science is measured by its 'problem-solving' abilities rather than some unattainable notion of truth (Laudan, 1977).

Given this diversity in modeling, it is important to keep in mind that the way we evaluate models is not orthogonal to our goals, *it is determined by them*. If we wish to find a model that strikes the 'best' balance between fit and parsimony, then criteria such as Bayes Factors or Normalized Maximum Likelihood are natural choices. However, if our

⁴ Importantly, these interaction effects are 'statistically robust' in the sense that they have been widely replicated. But the only thing this robustness achieved was the perpetuation of a misunderstanding. This suggests that our definition of robust modeling needs to be broader, something along the lines of 'modeling that is unlikely to fool ourselves and others'.

goal is to construct increasingly more encompassing or detailed accounts, other criteria such as minimizing deviance might be more reasonable. After all, a bias towards parsimony will not lead us to select models that are necessarily ‘closer’ to the data-generating truth. In fact, it can even lead us to prefer parsimonious models that don’t fit the data over more complex models that do (Gelman & Rubin, 1999).

This discussion brings me to Lee et al.’s (2019) treatment of model evaluation criteria. They argue that the many different evaluation criteria available can lead to opposite conclusions when applied to the same research question, models, and data (see their Figure 2). Their recommendation then is a preregistration of evaluation criteria, which in my view misses the point. The fact that two or more models can be fit to the same data does not imply that we are presented with an interesting model-comparison exercise. For all we know, the differences in model predictions might be minor and by and large due to their ancillary parametric assumptions. Perhaps the data are too simple to warrant all the different processes included in a model.⁵ What we need is a careful justification of why the data coming from a given study is interesting for purposes of model comparison. What are the diverging predictions and how do they connect with the theoretical claims made in each model? Are the divergences qualitative? Quantitative? How are they affected by ancillary assumptions? What can be inferred from different types of model misfits? In short, *robust modeling first and foremost requires diagnostic data*. The greater the diagnosticity, the less the differences between model-evaluation criteria will matter. In some cases, researchers are going to have to go bigger, and consider more complete accounts (e.g., Molloy, Galdo, Bahg, Liu, & Turner, 2019). In other cases, they might have to go smaller, focusing on specific portions of data in which critical model features are manifested with little to no interference from ancillary assumptions (e.g., Birnbaum, 2008; Kellen & Klauer, 2015).

Requiring researchers to motivate the use of their data also minimizes other

⁵ Take a look at Figure 2 in Lee et al. (2019). One can argue that there is no real disagreement between criteria simply because the performance differences are negligible. Also both models seem to be producing somewhat similar patterns.

problematic issues mentioned by Lee et al. (2019) such as HARKing (Hypothesizing After Results Are Known): First, having researchers justify their design makes it hard for them to convincingly reframe it around some specific set of results. Second, in some cases where critical tests are involved, the history behind an hypothesis is of little importance. Take Lee et al.’s *Example 1: Expected Utility Theory*, a model that is unable to handle people’s choices across a wide range of well-defined scenarios (Allais, 1953; Birnbaum, 2008). Because these failures can be traced back to the model’s axioms, it doesn’t matter if researchers derived them a priori or stumbled upon when inspecting their results.⁶

Models of Data

Researchers often overlook the fact that what we refer to as ‘data’ never corresponds to some collection of ‘raw’ observations, but to some canonical representation — a data model (Suppes, 1966). This data model mediates the relation between theoretical models and the observations made. The adoption of a given data model can be motivated by many factors, such as the properties of the theoretical models one intends to adopt, the auxiliary assumptions one is prepared to assume (e.g., which type of summaries or aggregations do I find acceptable?), or mere tractability (for discussions, see Harris, 2003; Mayo, 1996).

Evaluating the assumptions of the data model is a necessary step to ensure the robustness of our inferences. For instance, researchers often construct data points by summarizing multiple observations. Whether such summaries are reasonable depends on the theoretical models being considered. One example can be found in Navarro (2019), who discussed the modeling of a learning experiment in which responses across multiple trials were aggregated into accuracy scores. By aggregating, one is tacitly assuming that there is

⁶ One reviewer asked about the possibility of HARKing when researchers are free to determine which assumptions are central to the model and which are ancillary. I don’t see a problem in allowing researchers to have complete freedom in the way that they set up their models. The important question is whether they can make a compelling case for the choices they made. Take the case of utility models: Could anybody justify a critical test that *required* the utility function to be linear? I cannot see how, especially without any reference to focused studies testing this assumption (e.g., Kirby, 2011), or the use of an experimental design that would make such an assumption plausible (e.g., Birnbaum, 2008).

no information of interest being lost, such as in cases where observations are independent and identically distributed (iid). When such assumptions are determined to be false, we need to focus our attention on the risks that we are exposed to. In the case of Navarro, the theoretical-modeling efforts would be at serious risk had she intended to capture the exact shape of the data (e.g., Heathcote, Brown, & Mewhort, 2000). However, her goal was to merely capture the qualitative trends, placing her in a much less vulnerable situation (for a related discussion, see Kellen & Klauer, in press).⁷

Sometimes multiple data models are possible, each with its own advantages and shortcomings. Consider the study of transitivity in preferential choice, in which different data and theoretical models have been used (Birnbaum, 2011; Regenwetter, Dana, and Davis-Stober, 2010). Transitivity corresponds to the notion that preferences are based on a subjective representation of options that yields a (partial) rank order. If one prefers option A over B and B over C , then one also prefers A over C . One possible data model, used by Regenwetter et al., constructs choice proportions based on choices recorded across several option-pair replications. An advantage of this data model is that it allows for transitivity to be tested under conditions in which it is very unlikely to hold a priori. The disadvantage is the need for multiple option-pair replications along with the assumption that choices are iid. An alternative data model was proposed by Birnbaum (2011), in which people's *choice patterns across option pairs* are considered instead. An attractive feature of this data model is that it does not require multiple replications or iid assumptions. However, the severity of the testing that can be conducted is comparatively lower (see Regenwetter, Dana, Davis-Stober, & Guo, 2011).

⁷ In some cases, we have a great deal of background knowledge that outright questions the use of certain data models. For example, it is well known in the response-time literature that response-time means are very limited in their ability to characterize the underlying distribution of response times (e.g., Balota & Yap, 2011). Evans, Hawkins, Boehm, Wagenmakers, and Brown (2017) report a comparison between two rival evidence-accumulation models in which the outcome completely depends on whether response-time means or the entire distributions are considered.

Models of the Experiment

Finally, we turn to the level of the hierarchy that has received the least attention among philosophers of science and theoreticians (for notable exceptions, see Galison, 1987; Hacking, 1983; Mayo, 1996). The experiment model establishes the relationship between the theoretical model and the different conditions in an experimental design. According to Suppes (1966), the questions associated with this level are concerned with parameter identifiability, the precision of estimation in general (e.g., Spektor & Kellen, 2018), and selective influence (e.g., Rae, Heathcote, Donkin, Averell, & Brown, 2014). I would argue that this level should also include our assumptions regarding the way individuals engage with tasks. Take the case of forced-choice tasks: Whenever we model participants' responses, we often take for granted that they are based on relative judgments rather than absolute ones (but see Starns, Chen, & Staub, 2017). This understanding of what constitutes a model of the experiment and its demarcation from theoretical models is analogous to the distinction between 'task theory' and 'cognitive architecture' found in the ACT-R literature (for a discussion, see Cooper, 2007).

Based on this understanding, one concern that I have relates with the degree of abstraction that is often applied (see Footnote 1). Many different experiments can be perceived as structurally equivalent if we only focus on some of its properties. This equivalence invites researchers to assume that the same experiment model holds for all of them. However, this equivalence only exists from the researcher's perspective, not the individuals who take part in an experiment. Lee et al.'s (2019) *Example 2: Context Effects in Decision Making* provides a perfect example, as it is a domain in which researchers have translated consumer-choice problems into perceptual-judgment tasks that preserve their basic structure (Trueblood, Brown, Heathcote, & Busemeyer, 2013). A context effect is said to occur when the probability of someone choosing one alternative over the other is influenced by the presence of a third alternative. For instance, the probability of choosing a certain variety of apples over oranges can be increased by introducing an inferior variety of

apples, an *attraction effect*.

Spektor, Kellen, and Hotaling (2018) attempted to further explore the attraction effect using Trueblood et al.’s (2013) perceptual task. Much to their surprise, they were consistently obtaining a strong effect in the opposite direction — a *repulsion effect*. After several experiments, Spektor et al. were able to find some the causes behind this reversal. Chief among these was the arrangement of the stimuli on the screen. Whereas Trueblood et al. had the stimuli displayed horizontally and somewhat close, Spektor et al. used a ‘triangular’ arrangement and placed them somewhat farther apart. This dependency on presentation format outright rejects the experiment model as participants are engaging with the task in a way that is sensitive to experimental-design choices thought to be innocuous. Rejections of this kind can compromise the robustness of theoretical modeling efforts, given that the resulting characterizations are likely to hold only under a very specific set of not-so-interesting circumstances.

Conclusion

Lee et al.’s (2019) discussion is largely predicated on a statistical understanding of models. This is evidenced by their comparison with traditional data analyses, in which they argue that the main difference between the two pretty much lies on whether substantive interpretations can be made. The goal of this commentary is to show that this understanding overlooks the numerous roles that models can play in our scientific practices (Frigg & Hartmann, 2018; Morgan & Morrison, 1999). Hierarchies like the one proposed by Suppes (1966) provide a way to make sense of these roles and understand the different types of errors that one can commit at each level. The end result is a more complete view of what robust modeling means.

References

- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, *21*, 503-546.
- Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive psychology*, *6*, 451-474.
- Balota, D. A., & Yap, M. J. (2011). Moving beyond the mean in studies of mental chronometry: The power of response time distributional analyses. *Current Directions in Psychological Science*, *20*, 160-166.
- Batchelder, W. H. (2010). Cognitive psychometrics: Using multinomial processing tree models as measurement tools. In S. Embretson (Ed.), *Measuring psychological constructs: Advances in model-based approaches*. American Psychological Association.
- Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review*, *115*, 463-501.
- Birnbaum, M. H. (2011). Testing mixture models of transitive preference: Comment on Regenwetter, Dana, and Davis-Stober (2011). *Psychological Review*.
- Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, *71*, 791-799.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, *100*, 432-459.
- Cantor, J., & Engle, R. W. (1993). Working-memory capacity as long-term memory activation: An individual-differences approach. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 1101-1114.
- Cartwright, N. (1983). *How the laws of physics lie*. Oxford, UK: Clarendon Press.
- Cooper, R. P. (2007). The role of falsification in the development of cognitive architectures: Insights from a lakatosian analysis. *Cognitive science*, *31*, 509-533.

- Coyne, J. C. (2016). Replication initiatives will not salvage the trustworthiness of psychology. *BMC Psychology*, 4, 28.
- Cummins, R. (2000). “how does it work?” versus “what are the laws?”: Two conceptions of psychological explanation. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 117–144). Cambridge, Mass.: MIT Press.
- Duhem, P. M. M. (1954). *The aim and structure of physical theory*. Princeton, NJ: Princeton University Press.
- Evans, N. J., Hawkins, G. E., Boehm, U., Wagenmakers, E.-J., & Brown, S. D. (2017). The computations that support simple decision-making: A comparison between the diffusion and urgency-gating models. *Scientific Reports*, 7, 16433.
- Fiedler, K. (2018). The creative cycle and the growth of psychological science. *Perspectives on Psychological Science*, 13(4), 433–438.
- Frigg, R., & Hartmann, S. (2018). Models in science. *Stanford Encyclopedia of Philosophy*.
- Galison, P. (1987). *How experiments end*. Chicago, IL: University of Chicago Press.
- Garcia-Marques, L., & Ferreira, M. . B. (2011). Friends and foes of theory construction in psychological science: Vague dichotomies, unified theories of cognition, and the new experimentalism. *Perspectives on Psychological Science*, 6, 192–201.
- Gelman, A., & Rubin, D. B. (1999). Evaluating and using statistical methods in the social sciences: A discussion of “a critique of the Bayesian Information Criterion for model selection”. *Sociological Methods & Research*, 27, 403–410.
- Giere, R. N. (2004). How models are used to represent reality. *Philosophy of Science*, 71, 742–752.
- Giere, R. N. (2010). An agent-based conception of models and scientific representation. *Synthese*, 172, 269–281.
- Hacking, I. (1983). *Representing and intervening: Introductory topics in the philosophy of natural science*. Cambridge, UK: Cambridge University Press.
- Harris, T. (2003). Data models and the acquisition and manipulation of data. *Philosophy*

- of Science*, 70, 1508–1517.
- Heathcote, A., Brown, S., & Mewhort, D. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic bulletin & review*, 7, 185–207.
- Kellen, D., & Klauer, K. C. (2015). Signal detection and threshold modeling of confidence-rating ROCs: A critical test with minimal assumptions. *Psychological Review*, 122, 542–557.
- Kellen, D., & Klauer, K. C. (in press). Theories of the Wason selection task: A critical assessment of boundaries and benchmarks. *Computational Brain and Behavior*.
- Kirby, K. N. (2011). An empirical assessment of the form of utility functions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37, 461–476.
- Laudan, L. (1977). *Progress and its problems: Towards a theory of scientific growth*. Univ of California Press.
- Lee, M. D., Criss, A. H., Devezzer, B., Donkin, C., Etz, A., Leite, F., . . . others (2019). Robust modeling in cognitive science. *Computational Brain and Behavior*.
- Mayo, D. G. (1996). *Error and the growth of experimental knowledge*. University of Chicago Press.
- Molloy, M. F., Galdo, M., Bahg, G., Liu, Q., & Turner, B. M. (2019). What’s in a response time?: On the importance of response time measures in constraining models of context effects. *Decision*, 6, 171–200.
- Morgan, M. S., Morrison, M., & Skinner, Q. (1999). *Models as mediators: Perspectives on natural and social science*. Cambridge, UK: Cambridge University Press.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, Mass.: Harvard University Press.
- Pashler, H., & Wagenmakers, E. J. (2012). Editors’ introduction to the special section on replicability in psychological science: A crisis of confidence? *Perspectives on Psychological Science*, 7, 528–530.
- Rae, B., Heathcote, A., Donkin, C., Averell, L., & Brown, S. (2014). The hare and the

- tortoise: Emphasizing speed can change the evidence used to make decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40, 1226-1243.
- Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2011). Transitivity of preferences. *Psychological Review*, 118, 42-56.
- Regenwetter, M., Dana, J., Davis-Stober, C. P., & Guo, Y. (2011). Parsimonious testing of transitive or intransitive preferences: Reply to Birnbaum (2011). *Psychological Review*, 118, 684-688.
- Rotello, C. M., Heit, E., & Dube, C. (2015). When more data steer us wrong: Replications with the wrong dependent measure perpetuate erroneous conclusions. *Psychonomic Bulletin & Review*, 22, 944-954.
- Shiffrin, R. M., & Nobel, P. A. (1997). The art of model development and testing. *Behavior Research Methods, Instruments, & Computers*, 29, 6-14.
- Shiffrin, R. M., Ratcliff, R., & Clark, S. E. (1990). List-strength effect: II. theoretical mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 179-195.
- Spektor, M. S., & Kellen, D. (2018). The relative merit of empirical priors in non-identifiable and sloppy models: Applications to models of learning and decision-making. *Psychonomic Bulletin & Review*, 25, 2047-2068.
- Spektor, M. S., Kellen, D., & Hotaling, J. M. (2018). When the good looks bad: An experimental exploration of the repulsion effect. *Psychological Science*, 29, 1309-1320.
- Starns, J. J., Chen, T., & Staub, A. (2017). Eye movements in forced-choice recognition: Absolute judgments can preclude relative judgments. *Journal of Memory and Language*, 93, 55-66.
- Suppes, P. (1966). Models of data. In *Studies in logic and the foundations of mathematics* (Vol. 44, pp. 252-261). Elsevier.
- Trigg, J., & Kalish, M. (2011). Explaining how the mind works: On the relation between

- cognitive science and philosophy. *Topics in cognitive science*, 3, 399–424.
- Trueblood, J. S., Brown, S. D., Heathcote, A., & Busemeyer, J. R. (2013). Not just for consumers: Context effects are fundamental to decision making. *Psychological science*, 24, 901–908.
- Van Fraassen, B. C. (2008). *Scientific representation: Paradoxes of perspective*. Oxford, UK.: Oxford University Press.
- Weiskopf, D. A. (2011). Models and mechanisms in psychological explanation. *Synthese*, 183, 313–338.

Figure 1. A Suppes-type hierarchy. The arrows reflect the cross-talk existing between the different levels.