

## State-trace analysis can be an appropriate tool for assessing the number of cognitive systems: A reply to Ashby (2014)

John C. Dunn · Michael L. Kalish · Ben R. Newell

Published online: 1 May 2014  
© Psychonomic Society, Inc. 2014

**Abstract** Ashby (2014) has argued that state-trace analysis (STA) is not an appropriate tool for assessing the number of cognitive systems, because it fails in its primary goal of distinguishing single-parameter and multiple-parameter models. We show that this is based on a misunderstanding of the logic of STA, which depends solely on nearly universal assumptions about psychological measurement and clearly supersedes inferences based on functional dissociation and the analysis of interactions in analyses of variance. We demonstrate that STA can be used to draw inferences concerning the number of latent variables mediating the effects of a set of independent variables on a set of dependent variables. We suggest that STA is an appropriate tool to use when making arguments about the number of cognitive systems that must be posited to explain behavior. However, no statistical or inferential procedure is able to provide definitive answers to questions about the number of cognitive systems, simply because the concept of a “system” is not defined in an appropriate way.

**Keywords** State-trace analysis · Multiple systems · Math modeling · Model evaluation · Categorization

A fundamental aim of cognitive psychology is to characterize the mental processes that enable people to perform different tasks, such as learning to assign objects to categories.

---

J. C. Dunn (✉)

School of Psychology, University of Adelaide, North Terrace,  
Adelaide 5005, Australia  
e-mail: john.c.dunn@adelaide.edu.au

M. L. Kalish

Department of Psychology, Syracuse University, Syracuse, NY, USA

B. R. Newell

School of Psychology, University of New South Wales, Sydney,  
Australia

Fulfillment of this aim is made difficult by the fact that the processes themselves are not open to view and cannot be isolated for individual study—their nature and very existence must always be inferred. As a result, the logic of this inference will always be subject to debate and criticism. Ashby (2014; hereafter, simply “Ashby”) has recently contributed to this debate through a critique of the inferential logic of state-trace analysis (STA). He argues that STA fails to provide an appropriate inferential tool for assessing the number of cognitive systems involved in a given task. His argument has three main parts: (1) that STA frequently fails in its primary goal of distinguishing single-parameter and multiple-parameter models; (2) that it equally fails to distinguish single-system and multiple-system models, each of which may have multiple parameters; and (3) that STA lacks the statistical power to reject single-parameter outcomes, and thus that any failure to reject can be attributed to a Type II error.

We have employed STA to address claims that category learning relies on multiple cognitive systems (Dunn, Newell, & Kalish, 2012; Newell, Dunn, & Kalish, 2010), and have also separately argued that STA provides evidence that can be used to reason about the number of cognitive systems involved in categorization (Newell & Dunn, 2008; Newell, Dunn, & Kalish, 2011). Perhaps surprisingly, we are also on record as having argued, in agreement with Ashby, that STA cannot be used directly to distinguish between single-system and multiple-system models, each of which may have multiple parameters. We have suggested that “state-trace analysis does not offer a principled means of distinguishing between an interpretation of the data in terms of multiple parameters of a single system . . . or in terms of parameters of multiple systems” (Newell et al., 2011, p. 198). This agreement extends to recommendations as to how such interpretations should be distinguished. Ashby acknowledges that “more traditional dissociation logic is also flawed” (p. 9), and proposes that single-system and multiple-system accounts can only be

distinguished via a converging-operations approach in which “it is vital to consider *all* the available data” (p. 9). In the same vein, we earlier suggested that “distinguishing between these interpretations requires additional criteria such as the nature of experimental effects, their internal logic, and their respective abilities to account for the data” (Newell et al., 2011, p. 198), and called for a theory based on explicit mathematical formulations that precisely define its constructs and allows for its quantitative evaluation.

Despite these points of agreement, we do not agree with the final conclusion that STA is an inappropriate tool for assessing the number of cognitive systems. In addressing Ashby’s three-part argument, we will take the opportunity to describe more fully the logic of STA and how it can be used to weigh different claims concerning the numbers and natures of psychological processes, latent variables, or parameters underlying performance on different cognitive tasks. The following sections address each part of Ashby’s argument. In addition, we will briefly discuss the implications of our analysis for the “more traditional dissociation logic” and distinguish between the empirical question of how many latent variables underlie a particular data set and the philosophical question of how many cognitive systems are involved in a given task.

To foreshadow, our reply to Ashby’s objections are that (1) STA succeeds in distinguishing single-parameter and multiple-parameter models whenever the exceptionally mild and nearly universal assumption of monotonicity of measurements is met, in contrast to inferences based on function dissociation or the analysis of interactions in analyses of variance (ANOVAs), which do not; (2) the attempt to distinguish single-system and multiple-system models cannot be made purely statistically, but only rhetorically; (3) although the statistical power of STA is a topic of current research, significant results have already been obtained in relevant experiments.

### What is state-trace analysis?

STA was introduced by Bamber (1979) as a general method to determine properties of the latent structure that mediates the effects of two or more independent variables on two or more dependent variables. Since that time, it has undergone two lines of development. The first is associated with the work of Geoffrey Loftus and colleagues (e.g., Busey, Tunnicliff, Loftus, & Loftus, 2000; Loftus & Irwin, 1998; Loftus, Oberg, & Dillon, 2004) and is based on the original description of STA by Bamber and on an earlier article by Loftus (1978) that emphasized the critical concept of monotonicity in psychological measurement. Both of these articles highlighted the distinction between the latent variables (e.g., memory strength, degree of learning, or visual acuity) that are the targets of

psychological theory and the manifest dependent variables (e.g., proportions correct, hit rates, or response times) that are presumed to measure them in some way. Loftus proposed that the nature of this relationship is fundamentally unknown, and that it is therefore a mistake to interpret a change in a manifest variable as being equivalent to a change of equal magnitude in the latent variable that it measures. He further proposed that the best that might be hoped for is that the relationship between the two kinds of variables is at least monotonic. That is, if the value of the latent variable increases between two conditions, the value of the manifest variable should not decrease (or should not increase, if the relationship is inversely monotonic). The second line of development is associated with the work of Dunn and colleagues (Dunn, 2008; Dunn & James, 2003; Dunn & Kirsner, 1988, 2003; Newell & Dunn, 2008) and is concerned with the implications of Loftus’s argument for monotonicity in the logic of inferences based on functional dissociation. The conclusion drawn from this work was that functional dissociation is neither necessary nor sufficient to infer the existence of more than one latent variable, and that dissociation logic is superseded by the logic of STA.

As is discussed by Ashby, STA can be used to distinguish two kinds of models: a single-parameter model, in which the effects of two or more independent variables on two (or potentially more) dependent variables are mediated by exactly one latent variable; and a multiple-parameter model, in which the effects are mediated by more than one latent variable. In the first case, a plot of the observed values of one dependent variable against those of the other (called the *state-trace plot*) is confined to a one-dimensional curve in two-dimensional outcome space (the space of all possible values of the two dependent variables). In the second case, the plot is not so confined. This is a simple consequence of the underlying mathematics.

Unfortunately, the distinction between a one-dimensional curve and two-dimensional space is a mathematical idealization that is difficult to make in practice. Because only a relatively small number of points are typically sampled in an experiment, it is impossible to determine, without additional assumptions, whether or not they fall on a one-dimensional curve. As Ashby shows, without any constraints on the forms of the relationships between a single latent variable and two dependent variables, it is possible to generate a great variety of curves in outcome space. For this reason, and following Loftus (1978), STA has assumed that each dependent variable is a monotonic function of the underlying latent variable. As a consequence, the one-dimensional curve, and any points that fall on it, is constrained to be monotonically increasing (or decreasing) in outcome space. This makes it possible—in principle, at least—to distinguish a one-dimensional from a two-dimensional state trace.

### Is monotonicity a reasonable assumption?

Ashby has proposed that in many situations the monotonicity assumption is unlikely to be true, rendering STA irrelevant. This will occur if the measure of performance (the manifest variable) peaks at some intermediate value of a latent variable (e.g., response criterion, bias, learning rate, or relative attention). We make the following three comments on this argument.

First, monotonicity represents a very modest assumption in comparison to the much stronger assumption that a latent variable and the manifest variable that measures it are either identical or linearly related. As was recently discussed by Wagenmakers, Krypotos, Criss, and Iverson (2012), this view is still widely (if covertly) held by most researchers. The assumption of monotonicity follows from the construct validity of the measures: that changes in the latent variable (be it strength of association, or memory, or acuity of perception, or intensity of sensation) are never accompanied by changes in the manifest variable (percent correct on a test, recognition accuracy, detection, or rating) in the opposite direction. Ashby's claim appears to be that little, if anything, can be known about the relationship between a latent and a manifest variable. But this claim goes too far, because if it were true, no inference about the number and nature of psychological processes would be possible. Any experimental effect would be meaningless—if performance on a task increased between two conditions, it would be just as possible for the underlying latent variable (learning rate, memory strength, etc.) to have increased, decreased, or remained unchanged.

Second, the monotonicity assumption is just that—an assumption. Confronted by a nonmonotonic state trace, the researcher may choose to argue for the existence of multiple latent variables or for a violation of monotonicity. One of these options may be more reasonable than the other. Violations of monotonicity represent failures of operational definition. If a latent variable is nonmonotonically related to some performance measure, then some way to measure it needs to be found. For example, Ashby highlights the fact that response bias in a signal detection task is nonmonotonically related to the proportion of correct responses. However, it is monotonically related to both hit rate and false alarm rate, which is another reason to prefer these measures. On the other hand, if the state trace is monotonic, it can be concluded that the data are consistent with a single-parameter model.

Finally, it may be the case that in a particular domain it is more reasonable to assume that each manifest variable is a (different) nonmonotonic function of a single latent variable. However, it may also be reasonable to assume that the relevant functions are not completely unconstrained but have a predictable form, as is shown in the examples offered by Ashby. In these examples, each manifest variable reaches a single maximum at a different value of the latent variable. This

assumption can also serve as a constraint on the form of the resulting one-dimensional curve, to enable it to be distinguished from a two-dimensional alternative. Because it is less constraining than monotonicity, the relevant conclusion will be more difficult to achieve, and more data points may be required, but it may well be possible. The assumption of monotonicity is not integral to STA, it simply facilitates its application.

### What can be inferred from the state-trace plot?

The second part of Ashby's argument is that even if it is possible to identify a given state trace as being unambiguously one-dimensional or two-dimensional, nothing follows from this, since both outcomes are consistent with any number of cognitive systems. He demonstrates this by generating different state-trace plots from two different category-learning models—the single-system GCM and the dual-system COVIS model. As we discussed previously, since STA is sensitive to the number of latent variables (or model parameters), it does not directly identify whether these are packaged into one, two, or more “systems.” From the point of view of STA, both GCM and COVIS are simply multiple-parameter models that, all things considered, should both predict a high-dimensional state trace. In contrast, Ashby was able to generate a one-dimensional state trace from each model, as is shown in his Table 1, and concluded that this demonstrated that nothing useful can be inferred from the dimensionality of the state trace.

We contend, contra Ashby, that the dimensionality of the state trace is highly informative, even when the choice is between two different multiple-parameter models. To see this, it is necessary to understand how two multiple-parameter models are able to produce a one-dimensional state trace. Figure 1 illustrates the three ways in which this can be done. In each case, it is supposed that there are two dependent variables, labeled  $x$  and  $y$ , each of which is a different multivariate function of a set of parameters, represented here by  $u$  and  $v$ , which, in turn, are each a function of a set of independent variables, represented by  $a$  and  $b$ . Whether these parameters are grouped into different “systems” is immaterial for the following argument.

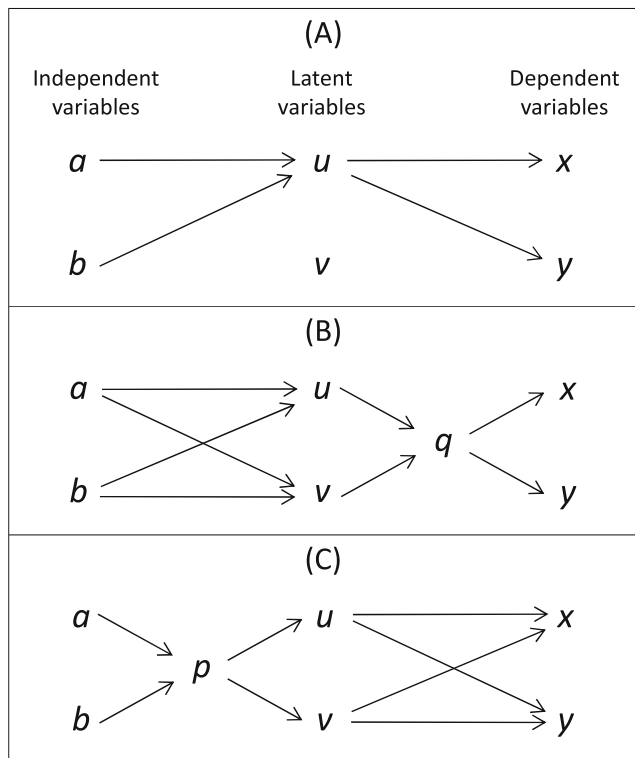
Figure 1A shows the first way of producing a one-dimensional state trace from a multiparameter model. In this instance, the two independent variables affect only one latent variable, which, in turn, affects both dependent variables. The resulting state trace is necessarily a one-dimensional curve. In addition, if each dependent variable is a monotonic function of the latent variable, then the state trace is also monotonic (either increasing or decreasing).<sup>1</sup>

<sup>1</sup> An example of such a state trace is shown in Ashby's Table 1.

**Table 1** Decision table for state-trace analysis

Observed State-Trace	Actual Number of Latent Variables Affected	
	One	More Than One
One-dimensional	One latent variable	Functional dependence
Two-dimensional	Failure of monotonicity	Multiple latent variables

Figure 1b shows the second way of producing a one-dimensional state trace from a multiparameter model. In this case, although each latent variable is a multivariate function of both independent variables, and each dependent variable is a multivariate function of both latent variables, the latter functions have the property that they are *dependent*. This means that each dependent variable is essentially the *same* function of the two latent variables. Put another way, the nature of the functional relationships means that the latent variables do not *differentially influence* the two dependent variables. This is illustrated in Fig. 1b by the appearance of a *virtual latent variable*,  $q$ , that is a function of the set of (actual) latent variables. If each dependent variable is, in turn, a monotonic function of  $q$ , then the resulting state trace will be a monotonic curve.<sup>2</sup> An example of this form of functional dependence,



**Fig. 1** Three different ways in which a multiparameter model may generate a one-dimensional state trace: (a) The independent variables affect only one latent variable. (b) The independent variables affect several latent variables, but these do not differentially affect the dependent variables. (c) The independent variables fail to differentially affect several latent variables

drawn from signal detection theory, is given in Appendix A. For further discussion of functional dependence in this context, see Dunn and Kirsner (1988), Kadlec and Van Rooij (2003), and Pratte and Rouder (2012).

Figure 1c shows the third way of producing a one-dimensional state trace from a multiparameter model. It also involves functional dependence, but in this case between the independent and latent variables. In Fig. 1c, the two independent variables do not differentially influence the set of latent variables. Each latent variable is essentially the same function of the two independent variables, which can therefore be replaced by a single *virtual independent variable*,  $p$ . The resulting state trace is necessarily a one-dimensional curve. If, in addition, the functions that relate  $p$  to each latent variable are also monotonic, and the functions that relate each latent variable to each dependent variable (holding the other latent variables constant) are monotonic, then the state trace will also be a monotonic curve. An example of this form of functional dependence, also drawn from signal detection theory, is given in Appendix B.

What can we conclude from all of this? First, in agreement with Ashby, it is clearly possible for a multiple-parameter model to generate a one-dimensional state trace. Second, rather than this telling us nothing, the result can be highly informative: It reveals that manipulation of the independent variables has not differentially influenced the relevant latent variables. This indicates either (1) that the effects of the independent variables on the dependent variables are mediated by a single latent variable (Fig. 1a); or (2) that the effects are mediated by several latent variables, but they are functionally dependent—either the latent variables do not differentially influence the dependent variables (Fig. 1b) or the independent variables do not differentially influence the latent variables (Fig. 1c). In each case, this tells us something about the model in question. This is the logic that Dunn et al. (2012) used in their study of the effects of number of trials, feedback type, feedback delay, and mask type on learning rule-based and non-rule-based category structures. They applied STA and observed a statistically significant two-dimensional state trace when number of trials and feedback delay were manipulated under minimal (yes/no) feedback and using a confusable mask. When the same independent variables were manipulated under full feedback, or when a less confusable mask was used, a two-dimensional state trace was not observed.<sup>3</sup> This pattern of results was unexpected, and they concluded that “the present results pose a challenge for all current models of categorization” (Dunn et al., 2012, p. 855).

Table 1 shows the decision table for STA, summarizing the points discussed in the current and previous sections. The rows correspond to classifications of the observed state-trace plot as either one-dimensional (i.e., monotonically increasing

<sup>2</sup> An example of such a state trace is shown in Ashby’s Table 1.

<sup>3</sup> Subject to qualification that it may be the result of a Type II error.

or decreasing) or two-dimensional (assuming two dependent variables). The columns correspond to the true state of affairs—whether the experimental manipulations affect one latent variable or more than one. The cells of the table correspond to the inferences that can be drawn at the conjunction of each row and column. If a one-dimensional state trace is observed, then one of two inferences is possible: Either only one latent variable has been affected, or more than one latent variable has been affected but the latent variables are functionally dependent. Depending on the context, these two alternatives may not be equally plausible. It should be apparent from the previous discussion that functional dependence may require a very special set of relationships between variables that, although it is possible, may be highly improbable. The experimenter then has to decide which interpretation of the state-trace plot is the more reasonable one, in terms of the features of the experiment, the theoretical background, and so on. If it is possible that the independent variables have not differentially influenced the latent variables, one strategy that is open to the experimenter is to examine the effects of other sets of independent variables. If these also produce a one-dimensional state trace, it becomes more difficult to attribute this to fortuitous relationships between some combinations of independent variables on the multiple latent variables. More likely, only one latent variable is actually involved.

A similar dichotomy exists if the state trace is observed to be two-dimensional. In this case, two inferences are again possible: Either two or more latent variables have been affected, or only one latent variable has been affected but there has been a failure of monotonicity. And, as in the one-dimensional case, the experimenter must decide which interpretation of the state-trace plot is the more reasonable one. In some contexts it may be more reasonable to question monotonicity, but if this were so, it is unclear why STA (or any other approach, such as ANOVA, which depends on even stronger assumptions) would be attempted in the first place.

In summary, STA is not a magic bullet. It cannot be used to decide unequivocally whether the effects of a set of independent variables on two or more dependent variables are mediated by one or more than one latent variable. It provides one useful source of evidence, but this must be weighed against other evidence and evaluated accordingly.

### Is a one-dimensional state trace a Type II error?

Ashby notes that several studies have reported one-dimensional (i.e., monotonic) state-trace plots in the category-learning literature. Because all models of category learning, whether single-system or multiple-system, find it necessary to invoke more than one parameter, he attributes all apparent claims of a one-dimensional state trace to a classical Type II error—failure to reject the null hypothesis. However, as we explained above, the

fact that multiple parameters may be in play does not guarantee that the state trace will itself be multidimensional. It will be one-dimensional if only one parameter has been affected, or if several parameters are affected but they are functionally dependent. In either case, the result provides evidence concerning the architecture of the system in question.

A Type II error for STA occurs when the true state trace is multidimensional but monotonicity cannot be rejected. Although the question of the statistical power of any test of the dimensionality of the state trace is currently open, if a Type II error is suspected, then it may be better to consider alternative experimental designs rather than relying on traditional designs and statistical analyses (Prince, Brown, & Heathcote, 2012). Needless to say, because many functional dissociations rely on accepting the null hypothesis (thereby asserting that a factor has no effect on a variable), they are likely to generate inherently high Type II error rates. Inference based on ANOVA also has its problems, in relation to the Type I error rate. A Type I error for STA occurs when the true state trace is one-dimensional but the statistical test concludes that it is multidimensional. Because of its reliance on a linear relationship between latent and manifest variables, a significant ANOVA interaction is often found even when the true state trace is one-dimensional, thus leading to an inherently high Type I error rate.

The claim by Ashby that failures to reject a one-dimensional state trace in relevant experiments in the category-learning literature can be attributed to a Type II error is refuted by experience. The detection of a two-dimensional state trace by Dunn et al. (2012, Exp. 1) makes it impossible for STA to always return a Type II error.

### What can be inferred from functional dissociations?

In his article, Ashby acknowledges that his arguments against STA also imply that the “more traditional dissociation logic is also flawed” (p. 9), but he also wants to infer the existence of multiple cognitive systems from such dissociations, as long as there are a large number of them. Thus, although he acknowledges that “even though a careful examination of each dissociation in isolation would likely show that that one result, by itself, was, at best, only weakly diagnostic,” he also argues that if a “multiple-systems model predicts ten new empirical dissociations a priori and . . . all ten are empirically supported” and if “no single-system model is known that can account for all these results,” then “collectively, these ten dissociations should be interpreted as strong support for multiple systems” (p. 9).

We have argued elsewhere that STA supersedes the logic of dissociation (Newell & Dunn, 2008). This is because many dissociations are consistent with a one-dimensional state trace, as was originally shown by Dunn and Kirsner (1988) and again noted by Ashby (p. 5). Briefly, under a model in which two dependent variables are monotonically increasing

functions of a single latent variable, STA tests exclusively for crossover interactions (or negative associations) in which one dependent variable increases and the other decreases between two experimental conditions. A dissociation, which relies on the absence of an effect, sits on the boundary of a positive association, in which both dependent variables increase (or decrease), and a crossover interaction. If the null difference is nudged in one direction, it is consistent with a one-dimensional state trace, but if sufficiently nudged the other way, it is consistent with a two-dimensional state trace. Because it sits on the fence in this way, a dissociation has no evidential value.

These considerations motivated the study by Newell et al. (2010), in which they applied STA in order to determine whether a dissociation reported by Zeithamova and Maddox (2006) was or was not associated with a two-dimensional state trace. Zeithamova and Maddox had found that the addition of a working memory load affected learning a rule-based category structure but had little effect on learning a non-rule-based structure, and they argued that this was consistent with the view that there were two different category-learning systems: One system, principally involved in learning rule-based structures, is affected by working memory load, and another system, principally involved in learning non-rule-based structures, is unaffected by load. In contrast, Newell et al. argued that a dissociation can only be viewed as providing evidence in support of a multiple-systems model if it is also inconsistent with a one-dimensional state trace. They found that the relevant dissociation was in fact consistent with a one-dimensional state trace when participants who failed to learn any structure were excluded. If these participants were included, the state trace became more obviously two-dimensional, but this was a necessary consequence of the differential inclusion of nonlearners across the two category structures.

It follows from this that ten or 100 dissociations, each consistent with a one-dimensional state trace, provide no support for the existence of multiple systems. In contrast, if at least one of these dissociations corresponds to a two-dimensional state trace, then this does offer support for the existence of multiple latent variables that may (or may not) be packaged into different “systems.”

### What is a system?

Ashby criticizes STA for not directly identifying whether one or multiple cognitive systems exist. But this is an impossible task for any statistical procedure or inferential logic, because the concept of a “system” is itself not well defined. The same term means different things to different people. Sherry and Schacter (1987) have suggested that systems may have evolved because they perform incompatible functions, but other theorists have stressed their interdependence (e.g., Kim

& Baxter, 2001; Klein, Cosmides, Tooby, & Chance, 2002), and others have emphasized functional differences based on anatomical features of the brain (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Manns & Eichenbaum, 2006). The question of how many systems are present is therefore not a strictly empirical question, at least not in the same way as the question of how many latent variables are present. This is because, as we have argued here, the question of how many latent variables are present can be decided empirically: Different hypotheses make predictions that directly interface with the data. Questions about how many systems are present turn just as much on philosophical judgments or rhetorical moves as they do on empirical judgments. We therefore do not see that placing questions about parameters and questions about systems in opposition to one another, as Ashby does, is useful. Any warning that STA will not solve philosophical problems is like having a warning label on your car that reminds you that the car will not take you to Mars.<sup>4</sup>

### Conclusion

In summary, we have argued that Ashby’s critique is unconvincing because the premises on which his argument rests fail to support his conclusion. We believe that there should be universal agreement that when a set of dependent variables are monotonically related to the latent variable they purport to measure (which is generally what “measure” means), then STA is a useful tool to draw inferences concerning the minimum number of latent variables that mediate the effects of a set of independent variables on that set of dependent variables. However, like all tools, STA must be used carefully, and should not be applied to tasks for which it is not suited. Specifically, no statistical or inferential procedure is able to provide definitive answers to questions about the number of cognitive systems, simply because the concept of a “system” is not defined in an appropriate way. We have also shown that STA logically supersedes reliance on functional dissociation and provides a principled foundation for theorizing about the complexity of the mechanisms of cognition. To the extent that this kind of theorizing is relevant to assessing the number of cognitive systems that exist, STA is an appropriate tool upon which to base such arguments.

**Author note** The authors acknowledge the support of the Australian Research Council (through Grant Nos. DP130101535, DP110100751, DP0878630, and DP0877510) and the support of the National Science Foundation (Award No. 1256959). We also thank our colleagues Andrew Heathcote, Scott Brown, Brett Hayes, Michael Lee, Richard Morey, Melissa Prince, and Rachel Stephens for many stimulating discussions, and two anonymous reviewers for their helpful comments on an earlier draft of the manuscript.

<sup>4</sup> We thank Richard Morey for this insightful comment.

## Appendix A: Functional dependence between latent and dependent variables

The first form of functional dependence affects the relationship between a set of latent variables and a set of dependent variables. To illustrate it, we draw on a hypothetical example from the theory of signal detection. As is well known, this theory is used to account for the proportions of hits and false alarms in a detection task. It has two parameters: discriminability  $d'$  and a response criterion  $c$  (corresponding to the latent parameters  $u$  and  $v$  in Fig. 1b). In one parameterization, the false alarm rate is a monotonic function<sup>5</sup> of  $c$ , whereas the hit rate is a (potentially different) monotonic function of  $c$  and  $d'$  (for a complete discussion, see Macmillan & Creelman, 2005). Imagine that an experiment is performed in which both  $d'$  and  $c$  are differentially affected. If the hit rate is plotted against the false alarm rate for each condition, the resulting state trace will be two-dimensional. However, suppose that we are not interested in the dependent variables, hit rate and false alarm rate, but are interested in two other commonly used measures of performance—overall proportion correct and the difference between the hit and false alarm rates. In this case, the state trace will be a one-dimensional monotonically increasing curve. The reason for this is that the functions that relate the model parameters to these dependent variables are functionally dependent (unlike the functions relating the same parameters to hit rate and false alarm rate). Let  $q$  be the difference between hit rate and the false alarm rate. Then proportion correct is, by definition, equal to  $(1 + q)/2$ . In other words, proportion correct and the difference between hit and false alarm rates are both functions of the same derived parameter (or virtual latent variable),  $q$ , and all information relating to the two underlying parameters is lost.

## Appendix B: Functional dependence between independent and latent variables

The second form of functional dependence affects the relationship between a set of independent variables and a set of latent variables. To illustrate this, we again draw on a hypothetical example from the theory of signal detection. Imagine that an experiment is conducted in which detection of a faint light is tested under four conditions defined by the orthogonal combination of two independent variables: target contrast (low vs. high) and retinal location (central vs. peripheral). Four experienced observers are tested, and each is assigned to a different combination of target contrast and retinal location. Each observer is unbiased and minimizes their overall error rate by placing their decision criterion at the midpoint

between the signal and noise distributions on the perceptual continuum—assuming, for the sake of argument, that these distributions are equal-variance Gaussian. The two dependent variables are hit rate (a monotonic function of  $d'$  and  $c$ ) and the false alarm rate (a monotonic function of  $c$ ). We also suppose that  $d'$  is some (complex) multivariate function of target contrast and retinal location. But, because each observer places his or her decision criterion optimally, it turns out that  $c$  is a function of  $d'$ . Specifically, because it is assumed that the target and noise distributions are Gaussian with equal variances,  $c = \frac{1}{2}d'$ . Both parameters can therefore be considered functions of a single virtual independent variable,  $p$ , where  $d' = p$  and  $c = \frac{1}{2}p$ . The result is that the state-trace plot (of hit rate against false alarm rate) is necessarily a one-dimensional (and, in this case, monotonically decreasing) curve. Because the two independent variables do not differentially influence the two latent variables, all information concerning the separate identities of these latent variables is lost.

## References

- Ashby, F. G. (2014). Is state-trace analysis an appropriate tool for assessing the number of cognitive systems? *Psychonomic Bulletin & Review*. doi:10.3758/s13423-013-0578-x
- Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, 105, 442–481. doi:10.1037/0033-295X.105.3.442
- Bamber, D. (1979). State-trace analysis: A method of testing simple theories of causation. *Journal of Mathematical Psychology*, 19, 137–181. doi:10.1016/0022-2496(79)90016-6
- Busey, T., Tunnicliff, J., Loftus, G., & Loftus, E. (2000). Accounts of the confidence–accuracy relation in recognition memory. *Psychonomic Bulletin & Review*, 7, 26–48. doi:10.3758/bf03210724
- Dunn, J. C. (2008). The dimensionality of the remember–know task: A state-trace analysis. *Psychological Review*, 115, 426–446. doi:10.1037/0033-295X.115.2.426
- Dunn, J. C., & James, R. N. (2003). Signed difference analysis: Theory and application. *Journal of Mathematical Psychology*, 47, 389–416. doi:10.1016/S0022-2496(03)00049-X
- Dunn, J. C., & Kirsner, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review*, 95, 91–101. doi:10.1037/0033-295X.95.1.91
- Dunn, J. C., & Kirsner, K. (2003). What can we infer from double dissociations? *Cortex*, 39, 1–7. doi:10.1016/S0010-9452(08)70070-4
- Dunn, J. C., Newell, B. R., & Kalish, M. L. (2012). The effect of feedback delay and feedback type on perceptual category learning: The limits of multiple systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38, 840–859. doi:10.1037/a0027867
- Kadlec, H., & Van Rooij, I. (2003). Beyond existence: Inferences about mental processes from reversed association. *Cortex*, 39, 183–187. doi:10.1016/S0010-9452(08)70098-4
- Kim, J. J., & Baxter, M. G. (2001). Multiple brain-memory systems: The whole does not equal the sum of its parts. *Trends in Neurosciences*, 24, 324–330. doi:10.1016/S0166-2236(00)01818-X
- Klein, S. B., Cosmides, L., Tooby, J., & Chance, S. (2002). Decisions and the evolution of memory: Multiple systems, multiple functions.

<sup>5</sup> A<sup>1</sup> The monotonic function corresponds to a cumulative distribution function that is often assumed to be Gaussian.

- Psychological Review*, 109, 306–329. doi:10.1037/0033-295X.109.2.306
- Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, 6, 312–319. doi:10.3758/bf03197461
- Loftus, G. R., & Irwin, D. E. (1998). On the relations among different measures of visible and informational persistence. *Cognitive Psychology*, 35, 135–199. doi:10.1006/cogp.1998.0678
- Loftus, G. R., Oberg, M. A., & Dillon, A. M. (2004). Linear theory, dimensional theory, and the face-inversion effect. *Psychological Review*, 111, 835–863. doi:10.1037/0033-295X.111.4.835
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide*. Mahwah: Erlbaum.
- Manns, J. R., & Eichenbaum, H. (2006). Evolution of declarative memory. *Hippocampus*, 16, 795–808. doi:10.1002/hipo.20205
- Newell, B. R., & Dunn, J. C. (2008). Dimensions in data: Testing psychological models using state-trace analysis. *Trends in Cognitive Sciences*, 12, 285–290. doi:10.1016/j.tics.2008.04.009
- Newell, B. R., Dunn, J. C., & Kalish, M. (2010). The dimensionality of perceptual category learning: A state-trace analysis. *Memory and Cognition*, 38, 563–581. doi:10.3758/MC.38.5.563
- Newell, B. R., Dunn, J. C., & Kalish, M. (2011). Systems of category learning: Fact or fantasy? In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 54, pp. 167–215). San Diego: Academic Press.
- Pratte, M. S., & Rouder, J. N. (2012). Assessing the dissociability of recollection and familiarity in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38, 1591–1607. doi:10.1037/a0028144
- Prince, M., Brown, S., & Heathcote, A. (2012). The design and analysis of state-trace experiments. *Psychological Methods*, 17, 78–99. doi:10.1037/a0025809
- Sherry, D. F., & Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological Review*, 94, 439–454. doi:10.1037/0033-295x.94.4.439
- Wagenmakers, E.-J., Krypotos, A.-M., Criss, A. H., & Iverson, G. J. (2012). On the interpretation of removable interactions: A survey of the field 33 years after Loftus. *Memory & Cognition*, 40, 145–160. doi:10.3758/s13421-011-0158-0
- Zeithamova, D., & Maddox, W. T. (2006). Dual-task interference in perceptual category learning. *Memory & Cognition*, 34, 387–398. doi:10.3758/BF03193416