Theories of the Wason Selection Task: A Critical Assessment of Boundaries and
Benchmarks

David Kellen
Syracuse University

Karl Christoph Klauer
Albert-Ludwigs-Universität Freiburg

Author Note

David Kellen, Department of Psychology. Karl Christoph Klauer, Department of Social Psychology and Methodology. Correspondence should be sent to David Kellen (davekellen@gmail.com). David Kellen received support from the Swiss National Science Foundation Grant 100014_165591. Karl Christoph Klauer was supported by DFG Reinhart-Koselleck grant DFG Kl 614/39-1. Data and analysis scripts can be found at https://osf.io/mvpbh/?view_only=df692cae0cda4b92bdcadb72c823fee2.

Abstract

The Wason selection task is one of the most prominent paradigms in the psychology of reasoning, with hundreds of published investigations in the last fifty odd years. But despite its central role in reasoning research, there has been little to no attempt to make sense of the data in a way that allows us to discard potential theoretical accounts. In fact, theories have been allowed to proliferate without any comprehensive evaluation of their relative performance. In an attempt to address this problem, Ragni, Kola, and Johnson-Laird (2018) reported a meta-analysis of 228 experiments using the Wason selection task. This data corpus was used to evaluate sixteen different theories on the basis of three predictions: 1) the occurrence of *canonical selections*, 2) dependencies in selections, and 3) the effect of counter-example salience. Ragni et al. argued that all three effects cull the number of candidate theories down to only two, which are subsequently compared in a model-selection analysis. The present paper argues against the diagnostic value attributed to some of these predictions. Moreover, we revisit Ragni et al.'s model-selection analysis and show that the model they propose is non-identifiable and often fails to account for the data. Altogether, the problems discussed here suggest that we are still far from a much-needed theoretical winnowing.

Keywords: hypothesis testing, mental models, reasoning, selection task

Theories of the Wason Selection Task: A Critical Assessment of Boundaries and
Benchmarks

The report of my death has been greatly exaggerated.

— (famous misquote of) Mark Twain

Since its introduction by Peter Wason over fifty years ago, the card selection task has been a staple task in the study of human reasoning and hypothesis-testing behavior (Wason, 1960). In a typical selection task, participants see a set of four cards (see Figure 1), each having a number on one side, and a letter on the other. Their visible sides show, for example, the letters and numbers *E*, *K*, 6, and 3. Participants are then given a rule such as ‘*If there is a vowel on one side of a card, then there is an even number on the other side.*’, and their task is to indicate which cards they would need to turn (if any) in order to test whether rule is true or false. In more abstract terms, the rule given corresponds to an indicative conditional ‘*If p, then q*’ with each of the four cards having a logical relationship with the two propositions *p* and *q* in the rule (see also Figure 1). When the rule is understood as establishing that the antecedent being true (i.e., *p*) is *sufficient*, but not *necessary* for the consequent to be true (i.e., *q*), it follows that the logically correct selection corresponds to the two cards showing *E* and 3 — selection $p\bar{q}$.¹ However, the results reported by Wason (1960) as well as many other follow-up studies show that out of all sixteen possible selection patterns (from no cards selected to $p\bar{p}q\bar{q}$), only a small percentage of respondents, usually no more than 10%, makes the logically correct selection, with most of them selecting *p* or *pq* instead.

Although much has been learned since Wason (1960), we are still far from fully understanding the underlying reasoning processes, as showcased by the large and diverse set of candidate theories that are still considered viable (for reviews, see Evans, 2017; Evans & Over, 2004; Oaksford & Chater, 2003). In a recent review, Ragni, Kola, and Johnson-Laird’s (2018; henceforth RKJ) listed sixteen candidate theories, which we report in Table 1. These theories differ dramatically in terms of the underlying characterization: whereas some only make a vague reference to the processes involved (e.g., Relevance

¹ We refer to selections of cards by strings such as pq that concatenate the symbols for the selected cards.

Theory; Sperber, Cara, & Girotto, 1995), others invoke quite detailed processing architectures (e.g., Parallel-Distributed-Processing Model; Leighton & Dawson, 2001).

The existence of so many theories indicates the need for a comprehensive evaluation of their merits in order to reduce the number of viable candidates. Chief among these merits is each theory’s ability to accommodate the different behavioral regularities found in the literature. RKJ attempt to achieve this reduction: First, they use a large data corpus comprised of 228 different studies to establish empirical support for three predictions.

- **Prediction 1:** The preponderance of so-called canonical selection responses (p , pq , $p\bar{q}$, and $pq\bar{q}$).
- **Prediction 2:** The presence of dependencies in the card selections.
- **Prediction 3:** The effect of counter-example salience in card selection.

We elaborate on these predictions below. According to RKJ, these predictions are highly diagnostic, as each of them can by itself exclude over half of the candidate models (for predictions 1-3, the exclusion percentages are 56%, 75%, and 63%; see Table 1). Taken together, RKJ were able to reduce the number of candidate theories from sixteen to only two — a *model theory* (MT) based on an algorithm originally proposed by Johnson-Laird and Wason (1970), and the *inference-guessing theory* (IGT) developed by Klauer, Stahl, and Erdfelder (2007). The two models, which are members of the Multinomial Processing Tree class (MPT class; Riefer & Batchelder, 1988), are illustrated in Figures 2 and 3, respectively.

MT (illustrated in Figure 2) assumes that people’s selections depend on the intuitions they have based on the meaning of the hypothesis, and how these can be used to generate counterexamples. Given the rule ‘*If p, then q*’, people are assumed to list the items included in it, resulting in the list pq (with probability c) or p (with probability $1 - c$), depending on whether the rule is seen as implying its converse ‘*If q, then p*’ or not. With probability $1 - e$, the person has no further insights regarding the rule and therefore selects

all of the items included in the mental list (p or pq). With probability e , an explicit model of the rule is constructed and with probability $(1 - f)$, some degree of insight is reached: For a mental list p , item q is included due its potential for confirming the rule, resulting in the selection pattern pq . When the original list was pq , the construction of an explicit model of the rule leads to the inclusion of $\bar{q}$ given its refutation potential, resulting in selection $pq\bar{q}$. Finally, with probability f , full insight is reached so that the selection ultimately considered only consists of elements providing potential counter-examples to the rule — $p\bar{q}$.

In contrast with MT, the IGT does not postulate a specific algorithm that all participants are expected to perform. Instead, the model assumes that people interpret the rule in many different ways, with some participants merely guessing. The model is purposely vague about the nature of the underlying reasoning process, as its goal is to describe the relative frequencies of all 16 possible selection patterns based on the relative frequencies of different interpretations of the rule and the availability of a few simple inferences such as modus ponens. The structure of the model is illustrated in Figure 3. The parameters in the model reflect the different possible interpretations of the rule and different ways to apply a number of simple inferences (for a detailed discussion, see Klauer et al., 2007): c = conditionality *versus* biconditionality, x = bidirectionality *versus* case distinction, d = forward inference *versus* backward inference, s = perceived sufficiency *versus* necessity, and i = irreversible *versus* reversible reasoning. Please note that the order in which the parameters occur in the model is a theoretically plausible one, but it does not play a role in model performance (one can build statistically-equivalent models using different orders).

RKJ reported a model-selection analysis in which MT and a *simplified version* of the IGT were fitted to the canonical selections found in the data corpus and compared using goodness-of-fit scores and the Bayesian Information Criterion (BIC; Kass & Raftery, 1995). The model comparisons conducted indicated that MT outperforms the simplified IGT even

in terms of goodness-of-fit scores that do not penalize IGT for its greater number of free parameters.

We completely agree with RKJ that a reduction of the number of theories of the Wason selection task (or their integration) is long overdue (Klauer et al., 2007). We also appreciate the focus on specific patterns in the data as an approach to theory testing (i.e., a critical-test approach; see Birnbaum, 2008; Kellen & Klauer, 2014, 2015). However, three aspects of their work raise important concerns that deserve pause and discussion. First, we note that RKJ’s evaluation hinges on *aggregate data*, with each participant contributing a single response. This situation limits the type of theories that can be tested without the need for questionable assumptions (e.g., the false assumption that between- and within-subject variability are exchangeable). It also raises questions about the diagnostic value of some of the canonical responses and the observation of selection dependencies. Second, we point out that focusing on only four of sixteen selection patterns, labeled canonical, is problematic for a number of different reasons. Finally, we revisit RKJ’s MPT analysis and model-selection procedure in which they compare MT and IGT. Here we will demonstrate that even the simplified version of IGT considered by RKJ is not testable when considering the canonical selections alone. This fact indicates that some of the model-fit results reported by RKJ are *impossible*, whereas others are illegitimate. Turning to the MT, our analyses will show that it is non-identifiable and that it fails to describe the data at a non-negligible rates (i.e., above the nominal 5% rate). Altogether, the issues discussed here raise serious questions on the validity of RKJ’s conclusions.

Modeling Aspirations are Constrained by the Nature of the Data

RKJ’s evaluation left them with two viable candidates, their proposed MT model and Klauer et al.’s (2007) IGT. One of RKJ’s arguments for preferring the former over the latter is that MT provides a specific reasoning algorithm, whereas the IGT is intentionally vague on the nature of the underlying processes. However, one can argue that the IGT’s

vagueness and agnosticism are advantageous features *given the nature of the data at hand*. Specifically, we are dealing with datasets aggregating a *single* response per person. It is not unreasonable to assume that different people engage with the task in qualitatively different ways (e.g., use a matching heuristic; Evans, 1972; see also Newstead, Handley, Harley, Wright, & Farrelly, 2004; Stanovich & West, 1998), and that some of them rely on some form of guessing. Given this level of heterogeneity, it makes sense to rely on a model that can encompass many different types of processes (including guessing) and that does not make strong, algorithm-level assumptions. In fact, one can argue that the development of an algorithm-level model based on aggregate data is questionable to the extent that we are conflating between-subject and within-subject differences. In other words, we are committing an ecological fallacy (Molenaar, 2004; see also Regenwetter & Robinson, 2017).

Before moving on, let us take the opportunity to dispel some potential misunderstandings: We are not denouncing any specific level of theory development (e.g., Marr, 1982), or minimizing their scientific value. Nor are we trying to make a general case against the use of aggregate data for purposes of theory development and testing. What we are saying is that the level of theoretical work that we wish to engage in is going to be constrained by the nature of the data available. Specifically, when discussing theory predictions, we need to determine whether these are preserved under aggregation, and if so, under which conditions/assumptions — one cannot merely take them for granted. To be clear, there are many instances in which predictions at the individual level are preserved under aggregation (e.g., Heck & Erdfelder, 2017; Kellen, Singmann, & Batchelder, 2018). However, certain properties such as independence and functional form, which are critical when attempting to characterize processes in a fine-grained manner, are well known to be violated (Estes, 1956; Regenwetter & Robinson, 2017).

Dependencies in Aggregate Data do not Imply Dependencies in Individual Selections

The use of aggregate data introduces several restrictions on what can be learned from the data. Consider RKJ's main Prediction 2, the dependence of selections of evidence. The claim here is that an individual does not select cards independently of each other. But even if each individual selects each card independently of other cards, the aggregate data are still likely to exhibit dependencies because we are dealing with an heterogeneous sample. In other words, RKJ's main Prediction 2 is in fact non-diagnostic when tested on aggregate data, because dependent card selections in the aggregate data do not imply that any individuals' selection of evidence is dependent. Thus, RKJ's falsely argue that the inability to describe these dependencies is grounds for rejecting nine out of the sixteen candidate theories. The problem can be stated as follows: Assume a 2×2 table reporting the relative frequency with which cards q (rows) and $\bar{q}$ (columns) were selected. As shown in Table 2, a table in which the independence of selecting q and $\bar{q}$ is rejected can be comprised of a *mixture* of two tables in which independence holds. Please note that the reported mixture is only one possibility among many. We also refrain from making any claims regarding the nature of the different mixture components. What this demonstration shows is simply that the observation of dependencies at the aggregate level cannot be used as an argument against models assuming response independence. Spurious violations of independence at the aggregate level are well known in the social sciences. In fact, they form the basis of data-analytic methods such as latent-class analysis (Clogg, 1995; Lazarsfeld & Henry, 1968), in which response dependencies are characterized by a mixture of distributions in which response independence is satisfied.

Finally, one should keep in mind that the distortions introduced by data aggregation cannot be downplayed through appeals to the robustness and/or magnitude of the observed effects, nor can they be mitigated by attempts to increase statistical power. The fact that RKJ (and others before them) report robust dependencies at the level aggregate

data do not make the case for individual-level dependencies any more plausible, in the same way that looking at larger electorates would not enable us to dismiss the possibility of *Condorcet paradoxes* (Regenwetter & Robinson, 2017). Simply put, the problems associated with aggregation are *logical, not statistical*. Interpreting statistical robustness as a proxy for logical soundness is a common, well-documented misunderstanding among researchers that carries serious risks (for reviews, see Regenwetter & Robinson, 2017; Rotello, Heit, & Dube, 2015).

Problems with Focusing on Only Four Selection Patterns

RKJ focus their analysis on only four selection patterns, termed canonical, out of sixteen possible patterns. But what makes the canonical patterns canonical? There is consensus that the selections of p and pq stand out due to their frequency. Historically, much of the interest in the Wason selection task stems from the fact that a further pattern, $p\bar{q}$, actually occurs quite infrequently in standard versions of the selection task. This is interesting because from certain normative standpoints, this pattern is the correct choice. Much work actually went into finding conditions under which this pattern occurred more frequently (e.g., Cheng & Holyoak, 1985; Platt & Griggs, 1993), and it is probably due to this research interest that many studies now exist in which the pattern occurs quite frequently (see RKJ's Table 2). The malleability of the frequency with which the $p\bar{q}$ pattern occurs and its normative status certainly justify considering this pattern.

On the other hand, similar research effort has not been devoted to changing the frequency with which other relatively rare patterns occur, leaving it an open question whether the unstudied patterns might not turn out to be equally important. In fact, IGT also predicts that patterns such as q , $\bar{p}\bar{q}$, and other non-canonical patterns should arise more frequently under certain conditions. Klauer et al. (2007) found that the frequency of some of these patterns, notably the selection of q alone, was in fact systematically affected by a number of theoretically-motivated experimental manipulations. In Evans and Lynch's

(1973) negations paradigm, dramatic and robust effects can be found on the selection of the $\bar{p}q$ and $\bar{p}\bar{q}$ patterns among others (e.g., Stahl, Klauer, & Erdfelder, 2008), whereas Gigerenzer and colleagues report dramatic effects on the selection of $\bar{p}q$ as a consequence of what they call perspective change in thematic versions of the selection task (e.g., Gigerenzer & Hug, 1992). MT needs to claim that each occurrence of these other patterns is “idiosyncratic and rare” (RKJ, p. 784), which seems unsatisfactory given the just-mentioned systematic effects. Conversely, from RKJ’s Table 2, it is clear that the partial insight pattern $pq\bar{q}$, which RKJ consider canonical, is not more frequent than q or $p\bar{p}q\bar{q}$. And thus, to the extent to which inclusion of the partial insight pattern is based on RKJ’s main Prediction 1 (the preponderance of the canonical patterns), we should either consider the prediction falsified or include these two patterns among the canonical ones.

In contrast, we would argue that excluding an average of 26% responses/participants selecting non-canonical patterns is risky and likely to distort the assessment of the canonical patterns. One reason is that whatever processes (such as, for example, guessing) cause the 26% non-canonical selections are likely to contribute to the recorded selections of canonical patterns as well. In consequence, the frequencies of the canonical patterns themselves in all likelihood comprise a mixture of selections caused by the processes that also cause the selection of non-canonical patterns and of selections caused by other processes such as those described by MT or the inference part of the IGT model. This means that:

- a) The 26% of selections due to processes that MT cannot describe is in all likelihood an underestimation of their predominance.
- b) We need a means to correct for selections of canonical patterns due to the “wrong” processes that also underlie the selection of non-canonical patterns before we can validly evaluate the absolute or relative performance of candidate models.
- c) One correction approach consists of establishing a “guessing” probability distribution

over all sixteen card selection patterns. This approach was adopted in the original proposal of the IGT model (Klauer et al., 2007).

The problem we are facing here is similar to the one memory researchers encountered a couple of decades ago when discussing the role of processes or systems operating at different levels of consciousness. Data coming from any given memory task are in all likelihood produced by a mixture of processes, which means that one cannot legitimately take on a ‘process-pure view’ when interpreting them (for a critique, see Reingold & Toth, 1996). But whereas memory researchers have extended experimental designs in an attempt to estimate these mixtures (e.g., Buchner, Erdfelder, & Vaterrodt-Plünecke, 1995; Jacoby, 1991), RKJ’s narrow focus on canonical selections limits our ability to do so.

Evaluation Criteria Need to be Charitable and Consistent

The importance of guessing processes goes beyond contributing to a more accurate characterization of the processes underlying selection behavior. They also play a critical role when evaluating the merit of competing theories, especially when doing so on the basis of aggregate data. As proposed by Klauer et al. (2007) and acknowledged by RKJ, it is very likely that participants engage the Wason task in different ways and that some people’s selections are driven by guessing. Indeed, this notion stands out as one of the IGT model’s defining features. However, it would be incorrect to infer from this distinctive feature that any of the other candidate theories is in any way antagonistic to it. In light of this concordance, it is charitable to extend the possibility of guessing to all competing theories.

Such an extension is not inconsequential, as it enables theories to successfully handle the occurrence of some of the canonical selections, something which they would not be able to do otherwise. For instance, the selection $pq\bar{q}$ would become available to theories such as Oaksford and Chater’s (1994) *optimal data selection* account.² Note that the IGT, which

² RKJ’s review is often not explicit on which exact canonical patterns are not predicted by each theory (see their Supplemental Materials), but from our reading, most of them appear to fail to predict only $pq\bar{q}$.

RKJ consider to be one of the two viable theories, relies on guessing to accommodate this specific selection. What this example shows is that one cannot dismiss theories based on the occurrence of so-called canonical selections, unless we can demonstrate that this limitation cannot be overcome by means of a guessing process. Otherwise, the models are not being evaluated according to the same criteria.³

In their discussion, RKJ appear to minimize the role of guessing, for instance by stating that guessing-based selections should occur no more often than a selection made at random ($\frac{1}{16} \approx 6\%$).⁴ The problem with this statement is that it hinges on a view of guessing that is much more constrained than the ones found in the literature. For example, according to Klauer et al.'s (2007) IGT, selections are *independent* under guessing, such that the probability of all sixteen possible selection patterns is governed by four parameters only, θ_p , $\theta_{\bar{p}}$, θ_q , and $\theta_{\bar{q}}$, each corresponding to the probability that a given card is selected in guessing. But aside from independence, no other constraint is imposed, especially in terms of the values that each θ parameter can take: They can differ from each other, and can be anywhere within the unit interval. Also, the values can vary according to the task context or the type of cues provided, among other factors (see Klauer et al., 2007, p. 682). A uniform distribution over the selection patterns as alluded to by RKJ would only hold if all θ parameters were fixed to .50.

As a counter-argument, one can raise the issue that the characterization of the inference processes provided by the model is dependent on the flexibility of the non-inference processes being postulated (e.g., Riesterer & Ragni, 2018). One can then argue that this dependency can compromise the entire modeling enterprise. We completely agree with the first statement. In fact, this point is well documented in other domains such as recognition memory, where the estimation of *stimulus-state mappings* is dependent on

³ To ensure consistency, one should identify one specific guessing process and implement it in all of the other models.

⁴ Interestingly, RKJ's own criteria would determine the partial insight pattern $pq\bar{q}$, which occurs in 4.7% of all selections across the 228 studies, guessing-based as predicted by IGT.

the proposed *state-response mappings* (for discussions, see Klauer & Kellen, 2010; Kellen, Singmann, & Klauer, 2014; Kellen, Singmann, Vogt, & Klauer, 2015). One insight coming from this literature is the existence of complexity tradeoffs between mappings. For example, when attempting to fit data, the adoption of overly simplified state-response mappings often introduces the need for overly complex stimulus-state mappings and vice-versa. However, these two types of complexity in models should not be treated alike. Theoretical discussions are typically centered on the stimulus-state mappings, which means that the search for a parsimonious account will generally involve somewhat flexible state-response mappings. In any case, there are ways to address concerns with undue flexibility: As long established in the MPT literature, the reasonableness of a given set of postulated processes can be ascertained via *validity studies* in which one attempts to selectively influence them (Batchelder & Riefer, 1999; Erdfelder et al., 2009). The goal of the experiments reported by Klauer et al. (2007) was to achieve exactly that for the IGT.

Revisiting RKJ's MPT Model Analysis

RKJ attempted to compare MT with IGT by fitting both models to the four canonical response frequencies. One challenge they faced was the fact that any set of canonical responses only provides three degrees of freedom (as proportions have to sum to 1), which required a series of simplifications of the IGT. The resulting simplified IGT is illustrated in Figure 4. The results reported indicated that the MT outperformed the simplified IGT in terms of goodness of fit as well as in terms of the BIC.

Several aspects of RKJ's analysis are worth discussing and revisiting. First, the introduction of any kind of simplification is *not* necessitated by the author's focus on the set of canonical responses. After all, an inspection of any model's ability to describe some set of responses does not require the model to *not* consider other responses. In most cases, the inclusion of additional data makes comparisons crisper and dispels confounds (e.g., Fific, 2014). Moreover, as already mentioned, although the canonical patterns other than

the partial insight pattern tend to be selected frequently, 26% of the participants selected non-canonical patterns, with a range of 0% to 95% across studies.

Simplifying the IGT restricts the number of pathways by which different responses can occur, and thus risks distorting and/or limiting the model’s ability to account for data at large. Importantly, note that the restrictions go way beyond the removal of simple guessing processes: For instance, the simplified model allows for the possibility of sufficiency and necessity interpretations in the case of forward inferences. In contrast, a backward inference necessarily leads to the $p\bar{q}$ selection in the simplified model, whereas in the original model this is only one of four possible patterns stemming from backward inferences. Specifically, the simplified model tacitly assumes that backward inferences adopt a necessity interpretation and to be applied to the visible and invisible sides, whereas these assumptions are not made for forward inferences in the simplified model.⁵

Second, the simplified IGT is not correctly specified. RKJ’s simplified IGT model incorrectly assumes that a forward inference with a necessity interpretation yields a $p\bar{q}$ selection, when it is clear from the full model (see Figure 3) that such a selection is deemed impossible under those conditions (only $\bar{p}$ and $\bar{p}q$ are possible). One way to remove this discrepancy consists of removing the sufficiency/necessity branch by fixing $s = 1$ (see Figure 4). Interestingly it turns out that, when only dealing with the four canonical responses, this restriction can be imposed without any loss of generality (we will return to this point below; see Footnote 8). In any case, it should be clear that the restriction suggested here is motivated on purely *pragmatic grounds*. We are not arguing against the sufficiency/necessity branches of IGT. Finally, note also that restricting guessing to the selection of the partial insight pattern in the simplified IGT model removes the corrective function of the guessing part of the original IGT model. Namely, to provide a baseline for

⁵ In their defense, RKJ state that “*Its proponents might well object that we have pruned it too much, but we had to eliminate six of its parameters, because only the canonical selections are relevant*” (p. 789). But again, the focus on a specific set of responses does not impose the need to ignore any other responses (they could simply have an ancillary role in parameter estimation). The need only arises because RKJ *decided* to focus on a limited set of responses.

the frequency of selecting each of the 16 patterns that does not need to be accounted for in terms of inference processes.

Third, RKJ appear to express their motivation for simplifying the IGT when stating that ‘*standard algorithms cannot fit models with more parameters than categories of results*’ (p. 789). This statement is, however, inaccurate: There is no such impediment, the only issue being the inability to obtain a *unique* set of best-fitting parameter estimates (i.e., the model is not identifiable; Bamber & van Santen, 2000). Typically, models with parameters equal or larger than the number of degrees of freedom are not testable because the range of predictions completely covers the space of possible outcomes. However, it is possible to find exceptions in the literature: For instance, Regenwetter and Davis-Stober (2012) fitted and tested a model with 540 parameters to data that only provided 20 degrees of freedom. Their model was testable because it only covered $\frac{1}{2000}$ of the space of possible outcomes.

Fourth, the models fitted by RKJ are over/saturated and non-identifiable: MT has three parameters and the simplified IGT four, which respectively equal and exceed the three degrees of freedom provided by the canonical response frequencies.⁶ The fact that the simplified IGT is oversaturated means that its parameters are not identifiable, which *completely disallows* any attempt to quantify its flexibility by means of parameter counting, as done by BIC (for a thorough discussion, see Moran, 2016).⁷ But this is not the only issue invalidating RKJ’s model-selection procedure. A more general question concerns the testability of the two candidate models. As previously discussed, over/saturated models are usually not testable. If this is the case, then no meaningful selection is possible given that they will always perfectly fit the data and any suitable model-selection statistic will always yield a tie. Alternatively, if only one model is testable, it is questionable

⁶ Please note that this state of affairs would not change if we also included a fifth response category concerning the occurrence of non-canonical selections. Although it would introduce a degree of freedom, one would be forced (irrespective of the model) to ‘spend it’ on an additional parameter describing the probability of non-canonical selections.

⁷ Note that BIC corresponds to a Laplace approximation of a model’s marginal likelihood (Kass & Raftery, 1995), which requires the model’s log-likelihood function to have an unique global maxima (which in turn implies that all parameters are identifiable).

whether the entire model-selection procedure can go beyond evaluating how well this model fares. RKJ do not provide any discussion or clarification on this matter, but it turns out that the simplified IGT is not testable whereas the MT is.⁸ Including untestable models in a model-comparison procedure creates a scenario in which the results are *completely dependent* on the performance of their testable counterparts. Any ‘dismissal’ of untestable models would be entirely based on a preference for more parsimonious accounts rather than on some failure to describe the data in any way. Such a scenario is particularly problematic when comparing MT with IGT because the latter’s lack of testability is entirely owed to RKJ’s decision to consider only a small subset of the observable selection patterns.

MT is also non-identifiable: When solving the system of equations describing the tree structure of MT, we see that there is no single solution:

$$\begin{aligned} e &= \frac{2 - 2p - p\bar{q} \pm \sqrt{(2p + p\bar{q} - 2)^2 + 4(p(pq + 1) + p\bar{q} - 1)}}{2} \\ f &= \frac{pq}{e} \\ c &= \frac{pq\bar{q}}{e(1 - f)} \end{aligned}$$

For example, consider the vector [.363, .191, .394, .052] comprised of, in order, the probabilities associated with p , pq , $p\bar{q}$, and $pq\bar{q}$, found across all 104 ‘abstract’ datasets.

This vector is perfectly described by MT with parameters $c = 0.134$, $e = .581$, and

⁸ To see this, simply note that when we fix IGT parameter s to 1, the resulting three-parameter model is a simple reparametrization of a fully unconstrained four-outcome multinomial distribution (Purdy & Batchelder, 2009). Importantly, note that the case of the IGT provides a concrete example of why it is non-sensical to apply BIC in the case of non-identifiable models: Allowing s to be freely estimated would not change the flexibility of the IGT (it is already maximal). However, the BIC penalty would increase, as if somehow model complexity increased (with penalty severity increasing along with sample size). Turning to MT, its testability can be demonstrated by showing that it cannot perfectly fit some datasets. We will later show that the data corpus gathered by RKJ includes many such examples. This raises another puzzle with RKJ’s model analyses. RKJ report root mean square error (RMSE) as a measure of model fit. That the simplified IGT is untestable implies that it can perfectly fit each pattern of canonical selections so that RMSE must be zero for each dataset. Yet in their Table 6, all RMSE values for IGT were positive and in each case larger than RMSE values for MT, which is simply impossible. Another ironic implication is that computing BIC with an appropriate parameterization of the IGT model (with three parameters) necessarily leads to BIC values for IGT that are smaller or at best equal to MT’s BIC values, reversing RKJ’s model-selection results.

$f = .329$. However, a perfect fit is also achieved with parameters $c^* = .482$, $e^* = .300$, and $f^* = .637$. This lack of identifiability compromises the ability to interpret parameters as theoretically-meaningful characterizations of the data, as more than one interpretation within the same model is almost always available.⁹

One likely reason why many of the issues surrounding both IGT and MT did not come up before is the fact that RKJ did not engage in the kind of empirical parameter-validation analysis that traditionally accompanies the proposal of any new MPT model (Batchelder & Riefer, 1999; Erdfelder et al., 2009). It has long been established in the MPT literature that simply demonstrating that a model can fit the data of interest is not enough. It is also necessary to demonstrate that the characterization it provides is a reasonable one. Accordingly, Klauer et al.'s (2007) proposal of the IGT is accompanied by a series of experiments that successfully targeted the different processes it postulates via selective-influence studies. The purpose of a set of systematic parameter-validation studies is twofold: (a) To show that the parameters associated with substantively interpreted mental processes in fact behave as one would predict if they measure such processes with construct validity and (b) to show that the postulated parameters are indeed needed to account for the entire pattern of effects obtained in the parameter-validation studies. At a minimum, one would expect an interpretation of these previous studies under MT (in particular those affecting non-canonical selections such as q), as well as a new set of validation studies tailored to MT's processes.

⁹ This is an atypical case of non-identifiability in the sense that there are exactly two solutions instead of an entire range of solutions (i.e., the surface of the likelihood function has two peaks instead of having one or more ridges; see Spektor & Kellen, 2018). Also important is the fact that this non-identifiability case is easy to overlook: One common approach for checking the identifiability of model parameters is to verify whether the Jacobian matrix has full rank, which can be done by showing that its determinant is non-zero (Bamber & van Santen, 2000). However, full rank is a necessary but not sufficient condition for parameter identifiability (for a discussion, see Schmittmann, Dolan, Raijmakers, and Batchelder, 2010). In order to show how this approach can fail, we derived the determinant of a 3×3 submatrix of MT's Jacobian concerning three of the canonical selections (note that the four canonical selections only provide three independent probabilities). The determinant of this submatrix is $ec + e^2(f(1 - c) - 1)$, which is non-zero (i.e., Jacobian has full rank) everywhere except when $c = \frac{(1-f)e}{1-fe}$, $e = 0$, $e = \frac{-c}{f(1-c)-1}$, and $f = \frac{1}{1-c} - \frac{c}{e(1-c)}$. Outside of these special cases, MT's Jacobian has full rank, incorrectly suggesting that MT's parameters are identifiable.

Fitting MT and IGT to the Canonical Selections

Given the concerns associated with data aggregation, we refitted MT to each single study with MPTinR (Singmann & Kellen, 2013) using the maximum-likelihood method. MT being a saturated model means that it imposes only inequality restrictions on the predicted probabilities for the canonical selections. The observed relative frequencies either satisfy these inequality restrictions, in which case MT fits the data perfectly, or they do not, in which case there will be misfit.¹⁰ The results summarized in Table 3 show that MT does not perfectly fit sixty-four (28%) datasets. Among these datasets, 45% exhibited statistically significant misfit ($p < .05$), for an overall rejection rate of 13%. This result is somewhat higher than the expected 5% when assuming that the model captures the data-generating processes. The occurrence of statistically-significant misfits did not occur equally often across different kinds of selection tasks, ranging from 10% to 23%, but this could be due the differences in sample sizes and its impact in terms of statistical power.

We inspected the model-fit residuals in order to see whether there was any stable pattern in the misfits obtained across studies (Kellen & Singmann, 2016). As shown in Figure 5, indeed there was: We found that in the cases where fit is not perfect, MT *always* underestimates the probability of responses p and $pq\bar{q}$ and overestimates pq . This pattern indicates that MT is failing to describe data in a systematic way. Subsequent simulations showed that this pattern of residuals occurs is to be expected when MT corresponds to the data-generating process. We further investigated the range of predictions that MT can make. Using a grid search, we found that the MT's predicted selection probabilities all fall into sixteen of the 24 possible *selection probability inequalities* (e.g., $p \geq pq \geq p\bar{q} \geq pq\bar{q}$).

¹⁰ Because the model is saturated, no degrees of freedom are thus left for conducting traditional null-hypothesis tests (Bamber & Van Santen, 2000; Davis-Stober, 2009). Another problem is that many of the individual-study datasets have very few responses overall, which can be problematic when attempting to rely on tests based on asymptotic results. As a solution, we computed p -values using an adapted double-bootstrap procedure (van de Schoot, Hooijink, & Deković, 2010). The adaptation consisted of not taking non-parametric bootstrap samples data vector $\mathbf{x}$ in the first step, but sample from a Dirichlet distribution with concentration parameter vector $\mathbf{x} + \alpha$ where $\alpha = [1, 1, 1, 1]$. This adaptation enabled responses associated with zero cells to be sampled (with some small probability).

This means that there are some general inequality patterns that MT cannot predict at all, and among those it can predict, only some values are permitted.¹¹ Overall, MT appears to have trouble accounting for cases in which pq selections occur very seldom and both p and $p\bar{q}$ or just $p\bar{q}$ occur fairly often.¹² ¹³

Discussion

RKJ's work represents a much-needed effort towards reducing the number of theories attempting to characterize selections in the Wason task. But despite its merits, their approach has important limitations that question the validity of the conclusions reached. One key problem is the fact that, in the face of person heterogeneity, there is only so much that one can learn from aggregate data. These limitations are reflected in models such as the IGT, which places its focus on the prevalence of different types of reasoning processes among individuals, rather than on fine-grained characterizations of each process. The presence of person heterogeneity raises important questions regarding the diagnostic value of some of RKJ's predictions:

- Prediction 1 (*preponderance of canonical selections*): There is no clear rationale for considering only certain selection patterns canonical while discounting others as “idiosyncratic and rare” (RKJ, p. 784). Historically, the canonical selection patterns

¹¹ The inequalities among selection probabilities can be conveniently represented by a string indicating the rank for each selection. The ‘forbidden’ ranks (from highest to lowest) for $[p, pq, p\bar{q}, pq\bar{q}]$ are 1423, 1432, 1342, 3412, 3421, 2341, 2431, and 2413.

¹² For example, among the studies MT fares poorly, one involves an abstract rule and requires participants to justify each selection (Stanovich & West, 2008). Among the participants making a canonical selection, 48% and 21% selected p and $p\bar{q}$ respectively, whereas only 16% chose pq . Another study involved the so-called *age-drinking problem*, in which the rule to be tested is “*If a person is drinking beer then the person must be over 21 years of age.*” (Stanovich & West, 1998). Ninety-one percent and 2% of the participants selected $p\bar{q}$ and pq respectively, among those making a canonical selection.

¹³ One anonymous reviewer suggested that we are making a conceptual error here, in which we conflate models and theories. The argument is that the range of predictions covered by the MPT model entitled ‘MT’ does not cover all the predictions under the umbrella of ‘Model Theory’ (Johnson-Laird & Wason, 1970), such that any failure of the former does not necessarily carry over to the latter. We are happy to concede this point, noting however that it is not obvious (at least to us) what the exact differences are, nor is the testability of Model Theory clear to us.

have often been reported, but historical contingency is not a scientifically valid argument. In fact, the partial insight pattern, considered canonical, is no less frequently observed than at least two other patterns, and there is evidence that the frequency of other selection patterns can be systematically affected by experimental manipulations. Moreover, more than a quarter of the participants make non-canonical selections, with much heterogeneity in the percentage of non-canonical selections across studies, leaving a quarter of the responses as well as the considerable variance in the overall selection of non-canonical patterns unaccounted for. Finally, even if a theory's reasoning processes do not predict some of the canonical selections, one could and perhaps should attribute their occurrence to non-reasoning processes such as guessing in a charitable extension of the theory.

- Prediction 2 (*dependency of individual selections*): Observed selection dependencies might be due to the mixture of independent selections. This means that the inference from dependencies in the aggregate data to dependencies in individual selections is a non-sequitur.

Moreover, many additional issues were discussed with respect to RKJ's analyses contrasting MT and a simplified version of IGT. The simplified IGT is misspecified and not an appropriate simplification of the IGT. The goodness-of-fit results reported in terms of root mean square error, RMSE, are impossible as reported and must by necessity favor IGT rather than, as reported, MT. Use of BIC to select between the two models is invalidated because the simplified IGT as stated is overparameterized and MT is non-identifiable. The simplified IGT is furthermore consistent with any possible dataset for the canonical selections, implying that even if a more appropriate selection index favored one or the other model, this preference would depend solely on MT's performance in fitting the data. This is especially problematic because the untestability of IGT is due to RKJ's decision to focus on only the canonical selections. In a reanalysis of MT's performance in fitting the canonical selections, MT was rejected at higher rates than expected. Finally, no

model-validation analysis of MT was provided.

In light of these results, we think it is premature to defend a reduction of candidate theories to just MT and IGT, let alone a preference for MT. RKJ's case for MT can be framed in terms of a cost-benefit analysis: Roughly speaking, the costs consist of 1) accepting that ignoring all but the canonical selections is reasonable. 2) that aggregation biases are not problematic, 3) that one does not have to assume a consistent and charitable set of evaluation criteria across models. The benefit is that one ends up with a model that assumes an algorithm-level account, but is unable to make predictions regarding non-canonical selections, and fails to provide a unique characterization of underlying processes. Altogether, the case for MT requires too much and returns too little.

But despite all of our criticisms, it would be unreasonable to simply dismiss some of the predictions used by RKJ: Although they are not diagnostic at the level of aggregate data, they could be used to motivate the development of research programs that are more centered around person-level data and the heterogeneity found in them (e.g., Oberauer, Geiger, Fischer, & Weidenfeld, 2007; Stanovich & West, 2000; Stenning & van Lambalgen, 2008; Trippas, Handley, & Verde, 2013). This would be a timely move given the current ability to collect data from a large and diverse groups of individuals and the availability of statistical methods that can merge them in ways that nevertheless preserve their idiosyncrasies (for a discussion, see Trippas et al., 2018). Skovgaard-Olsen, Kellen, Hahn, and Klauer (in press) recently adopted such an approach, using individuals' probability judgments to classify them in terms of their interpretation of indicative conditionals (e.g., "If A, then C"), and used these classifications to test a number of predictions associated with the different interpretations. It should be noted that Skovgaard-Olsen's et al.'s approach shares important similarities with the single-case methods proposed by Stenning and van Lambalgen (2008), which have been applied to the Wason task.

Finally, consider famous observations such as that clocks on satellites orbiting earth run differently from clocks on earth, that a beam of laser light passing through two narrow

parallel slits positioned close together throws dark and light bands on a screen behind the slits, or that human blood-group chimeras (humans with two blood types) exist. On the one hand, these observations lend strong support to models in their respective fields that are more complex than traditional ones. On the other hand, these observations, being very specialized and/or rare, would in all likelihood not lead to a preference for the more complex models in a model-selection exercise that included all relevant data on, respectively, time measurement with clocks, light diffraction, and human genetics, because the simpler models, being able to describe most of the datasets well, would probably strike the better overall compromise between parsimony and fit. For such reasons, it may be wise to complement formal model-selection criteria by substantive evaluation criteria as argued by Klauer et al. (2007) for the IGT. In particular, there are many more reliable findings in the Wason selection-task paradigm (see Section “Problems with focusing on only four selection patterns” for examples) than the one considered in RKJ’s main Prediction 3 (salience of counterexamples increases selections of potential falsifications). Perhaps it would be fruitful to compile a list of the reliable benchmark effects that can be observed in the Wason selection task as has been done, for example, in working-memory research (Oberauer et al., 2018) and to evaluate (a) whether different empirically validated models provide plausible accounts of the different effects in terms of being able to reproduce them and (b) if so, whether the effects are mapped in a meaningful way on the different model parameters.

References

- Bamber, D., & Van Santen, J. P. (2000). How to assess a model's testability and identifiability. *Journal of Mathematical Psychology*, *44*, 20–40.
- Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review*, *6*, 57–86.
- Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review*, *115*, 463–501.
- Buchner, A., Erdfelder, E., & Vaterrodt-Plünnecke, B. (1995). Toward unbiased measurement of conscious and unconscious memory processes within the process dissociation framework. *Journal of Experimental Psychology: General*, *124*, 137–160.
- Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, *17*, 391–416.
- Clogg, C. C. (1995). Latent class models. In G. Arminger, C. C. Clogg, & E. Soberl (Eds.), *Handbook of statistical modeling for the social and behavioral sciences* (p. 311–359). New York: Plenum.
- Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? studies with the Wason selection task. *Cognition*, *31*, 187–276.
- Davis-Stober, C. P. (2009). Analysis of multinomial models under inequality constraints: Applications to measurement theory. *Journal of Mathematical Psychology*, *53*, 1–13.
- Eliasmith, C. (2005). Cognition with neurons: A large-scale, biologically realistic model of the Wason task. In *Proceedings of the 27th annual meeting of the cognitive science society* (p. 624–629).
- Erdfelder, E., Auer, T.-S., Hilbig, B. E., Aßfalg, A., Moshagen, M., & Nadarevic, L. (2009). Multinomial processing tree models: A review of the literature. *Zeitschrift für Psychologie/Journal of Psychology*, *217*, 108–124.
- Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin*, *53*, 134–140.

- Evans, J. S. B. T. (1972). Interpretation and matching bias in a reasoning task. *Quarterly Journal of Experimental Psychology*, *24*, 193–199.
- Evans, J. S. B. T. (1977). Toward a statistical theory of reasoning. *Quarterly Journal of Experimental Psychology*, *29*, 621–635.
- Evans, J. S. B. T. (1984). Heuristic and analytic processes in reasoning. *British Journal of Psychology*, *75*, 451–468.
- Evans, J. S. B. T. (2017). A brief history of the Wason selection task. In N. Galbraith, E. Lucas, & D. Over (Eds.), *The thinking mind: A Festschrift for Ken Manktelow* (pp. 15–28). New York, NY: Routledge.
- Evans, J. S. B. T., & Lynch, J. S. (1973). Matching bias in the selection task. *British Journal of Psychology*, *64*, 391–397. doi: 10.1111/j.2044-8295.1973.tb01365.x
- Evans, J. S. B. T., & Over, D. E. (2004). *If*. Oxford, England: Oxford University Press.
- Fific, M. (2014). Double jeopardy in inferring cognitive processes. *Frontiers in Psychology*, *5*, 1130.
- Gigerenzer, G., & Hug, K. (1992). Domain-specific reasoning: Social contracts, cheating, and perspective change. *Cognition*, *43*, 127–171. doi: 10.1016/0010-0277(92)90060-U
- Hattori, M. (2002). A quantitative model of optimal data selection in wason’s selection task. *Quarterly Journal of Experimental Psychology*, *55*, 1241–1272.
- Heck, D. W., & Erdfelder, E. (2017). Linking process and measurement models of recognition-based decisions. *Psychological Review*, *124*, 442–471.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513–541.
- Johnson-Laird, P. N., & Wason, P. C. (1970). A theoretical analysis of insight into a reasoning task. *Cognitive Psychology*, *1*, 134–148.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*, 773–795.
- Kellen, D., & Klauer, K. C. (2014). Discrete-state and continuous models of recognition

- memory: Testing core properties under minimal assumptions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40, 1795-1804.
- 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., & Singmann, H. (2016). ROC residuals in signal-detection models of recognition memory. *Psychonomic Bulletin & Review*, 23, 253-264.
- Kellen, D., Singmann, H., & Klauer, K. C. (2014). Modeling source-memory overdistribution. *Journal of Memory and Language*, 76, 216-236.
- Kellen, D., Singmann, H., Vogt, J., & Klauer, K. C. (2017). Further evidence for discrete-state mediation in recognition memory. *Experimental Psychology*, 62, 40-53.
- Kellen, H., D. and Singmann, & Batchelder, W. H. (2018). Classic-probability accounts of mirrored (quantum-like) order effects in human judgments. *Decision*, 5.
- Kirby, K. N. (1994). Probabilities and utilities of fictional outcomes in wason's four-card selection task. *Cognition*, 51, 1-28.
- Klauer, K. C., & Kellen, D. (2010). Toward a complete decision model of item and source recognition: A discrete-state approach. *Psychonomic Bulletin & Review*, 17, 465-478.
- Klauer, K. C., Stahl, C., & Erdfelder, E. (2007). The abstract selection task: New data and an almost comprehensive model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 680-703.
- Krauth, J. (1982). Formulation and experimental verification of models in propositional reasoning. *The Quarterly Journal of Experimental Psychology*, 34, 285-298.
- Lazarsfeld, P. F., & Henry, N. W. (1968). *Latent structure analysis*. Boston: Houghton Mifflin.
- Leighton, J. P., & Dawson, M. R. (2001). A parallel distributed processing model of Wason's selection task. *Cognitive Systems Research*, 2, 207-231.

- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco, CA: WH Freeman.
- Molenaar, P. C. (2004). A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology, this time forever. *Measurement*, 2, 201–218.
- Moran, R. (2016). Thou shalt identify! The identifiability of two high-threshold models in confidence-rating recognition (and super-recognition) paradigms. *Journal of Mathematical Psychology*, 73, 1–11.
- Newstead, S. E., Handley, S. J., Harley, C., Wright, H., & Farrelly, D. (2004). Individual differences in deductive reasoning. *Quarterly Journal of Experimental Psychology*, 57, 33–60.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101, 608–631.
- Oaksford, M., & Chater, N. (2003). Optimal data selection: Revision, review, and reevaluation. *Psychonomic Bulletin & Review*, 10, 289–318.
- Oberauer, K., Geiger, S. M., Fischer, K., & Weidenfeld, A. (2007). Two meanings of “if”? individual differences in the interpretation of conditionals. *The Quarterly Journal of Experimental Psychology*, 60, 790–819.
- Oberauer, K., Lewandowsky, S., Awh, E., Brown, G. D. A., Conway, A., Cowan, N., . . . Ward, G. (2018). Benchmarks for models of short-term and working memory. *Psychological Bulletin*, 144, 885–958.
- Platt, R. D., & Griggs, R. A. (1993). Facilitation in the abstract selection task: The effects of attentional and instructional factors. *The Quarterly Journal of Experimental Psychology*, 46, 591–613. doi: 10.1080/14640749308401029
- Purdy, B. P., & Batchelder, W. H. (2009). A context-free language for binary multinomial processing tree models. *Journal of Mathematical Psychology*, 53, 547–561.
- Ragni, M., Kola, I., & Johnson-Laird, P. N. (2018). On selecting evidence to test hypotheses: A theory of selection tasks. *Psychological Bulletin*, 144, 779–796.

- Regenwetter, M., & Davis-Stober, C. P. (2012). Behavioral variability of choices versus structural inconsistency of preferences. *Psychological Review*, *119*, 408-416.
- Regenwetter, M., & Robinson, M. M. (2017). The construct-behavior gap in behavioral decision research: A challenge beyond replicability. *Psychological Review*, *124*, 533-550.
- Reingold, E. M., & Toth, J. P. (1996). Process dissociations versus task dissociations: A controversy in progress. In G. Underwood (Ed.), *Implicit cognition* (p. 159-202). Oxford, UK: Oxford University Press.
- Riesterer, N., & Ragni, M. (2018). Implications of guessing types in multinomial processing tree models: Conditional reasoning as an example. *Paper presented at the 16th International Conference on Cognitive Modeling (Madison, WI)*.
- Rips, L. J. (1994). *The psychology of proof*. Cambridge, Mass.: MIT Press.
- Rotello, C. M., Heit, E., & Dubé, C. (2015). When more data steer us wrong: Replications with the wrong dependent measure perpetuate erroneous conclusions. *Psychonomic Bulletin & Review*, *22*, 944-954. doi: 10.3758/s13423-014-0759-2
- Schmittmann, V. D., Dolan, C. V., Raijmakers, M. E., & Batchelder, W. H. (2010). Parameter identification in multinomial processing tree models. *Behavior Research Methods*, *42*, 836-846.
- Singmann, H., & Kellen, D. (2013). MPTinR: Analysis of multinomial processing tree models in R. *Behavior Research Methods*, *45*, 560-575.
- Skovgaard-Olsen, N., Kellen, D., Hahn, U., & Klauer, K. C. (in press). Norm conflicts and conditionals. *Psychological Review*.
- Smith, J. B., & Batchelder, W. H. (2008). Assessing individual differences in categorical data. *Psychonomic Bulletin & Review*, *15*, 713-731.
- 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.

Sperber, D., Cara, F., & Girotto, V. (1995). Relevance theory explains the selection task.

Cognition, 57(1), 31–95.

Stahl, C., Klauer, K. C., & Erdfelder, E. (2008). Matching bias is not eliminated by explicit

negations. *Thinking and Reasoning*, 14, 281–303. doi: 10.1080/13546780802116807

Stanovich, K. E., & West, R. F. (1998). Individual differences in rational thought. *Journal*

of Experimental Psychology: General, 127, 161–188.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications

for the rationality debate? *Behavioral and Brain Sciences*, 23, 645–665.

Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases

and cognitive ability. *Journal of Personality and Social Psychology*, 94, 672–695.

Stenning, K., & van Lambalgen, M. (2008). *Human reasoning and cognitive science*.

Boston, Mass.: MIT Press.

Trippas, D., Handley, S. J., & Verde, M. F. (2013). The SDT model of belief bias:

Complexity, time, and cognitive ability mediate the effects of believability. *Journal of*

Experimental Psychology: Learning, Memory, and Cognition, 39, 1393–1402.

Trippas, D., Kellen, D., Singmann, H., Pennycook, G., Koehler, D. J., Fugelsang, J. A., &

Dubé, C. (2018). Characterizing belief bias in syllogistic reasoning: A hierarchical

Bayesian meta-analysis of ROC data. *Psychonomic Bulletin & Review*, 25,

2141–2174.

Van De Schoot, R., Hoijtink, H., & Deković, M. (2010). Testing inequality constrained

hypotheses in SEM models. *Structural Equation Modeling*, 17, 443–463.

Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task.

Quarterly Journal of Experimental Psychology, 12, 129–140.

Wason, P. C. (1966a). Reasoning. In B. Foss (Ed.), *New Horizons in Psychology*.

Harmondsworth, England: Penguin Books.

Wason, P. C. (1966b). Reasoning. In B. Foss (Ed.), *New Horizons in Psychology*.

Harmondsworth, England: Penguin Books.

Table 1

Theories of the Wason Selection Task and their ability to accommodate the three predictions established by Ragni et al. (2018).

Theory	Source	Dependent Selections	Canonical Selections	Salient Counter examples
Defective truth tables	Wason (1966)	✗	✓	✗
Model theory (MT)	Johnson-Laird and Wason (1970)	✓	✓	✓
PSYCOP	Rips (1994)	✗	✗	✗
Relevance	Sperber, Cara, and Girotto (1995)	✓	✗	✓
Multiple interpretations	Stenning and van Lambalgen (2008)	✓	✓	✗
Pragmatic schemas	Cheng and Holyoak (1985)	✓	✗	✗
Innate modules	Cosmides (1989)	✗	✗	✓
Stochastic theory	Evans (1977)	✗	✗	✗
Matching & verifying	Krauth (1982)	✗	✗	✗
Heuristic-analytic	Evans (1984)	✓	✗	✗
Inference-guessing (IGT)	Klauer, Stahl, and Erdfelder (2007)	✓	✓	✓
Probability & utility	Kirby (1994)	✗	✗	✗
Optimal information gain (v1)	Oaksford and Chater (1994)	✓	✗	✓
Optimal information gain (v2)	Hattori (2002)	✗	✗	✓
Parallel distributed processes	Leighton and Dawson (2001)	✗	✗	✗
Neurons	Eliasmith (2005)	✗	✗	✗

Note. This table is based on Table 5 reported by Ragni et al. (2018). The two theories that can account for all three predictions are in bold.

Table 2

Testing Independence in the Selection of q and $\bar{q}$

Observed			Mixture Component 1 (30%)			Mixture Component 2 (70%)		
	No $\bar{q}$	Yes $\bar{q}$		No $\bar{q}$	Yes $\bar{q}$		No $\bar{q}$	Yes $\bar{q}$
No q	.293	.244	No q	.000	.000	No q	.419	.349
Yes q	.343	.119	Yes q	.850	.150	Yes q	.126	.105

Note. The observed data consists of the proportions observed when aggregating all datasets reported by RKJ for which all sixteen possible selection patterns were available ($N=7991$). The hypothesis of selection independence was rejected ($G^2(df = 1) = 37.52$, $p < .0001$). The percentages in parentheses correspond to the contribution of each component to the observed table. In both mixture components, the selection probabilities of q and $\bar{q}$ are independent.

Table 3
MT Fit Results

Task	Studies (N)	Median Study Size	Proportion Misfits ($G^2 > 0$)	Proportion $p < .05$
Abstract	104	16	.21	.10
Everyday	44	26	.45	.23
Deontic	80	18	.28	.11
Overall	228	18	.28	.13

Figure 1. Illustration of the Wason selection task and the logical relationship between each card and the propositions p and q found in the rule ‘*If p , then q* ’.

Figure 2. Illustration of the MT model proposed by Ragni et al. (2018). Parameters c , e , and f correspond to probabilities (for details, see the main text).

Figure 3. Illustration of the IGT model proposed by Klauer et al. (2007). Parameters c , d , x , s , and i correspond to probabilities. Note that for space reasons, only the ‘inference’ branches are illustrated (i.e., guessing branches are omitted).

Figure 4. Illustration of the simplified IGT model proposed by Ragni et al. (2018) to describe the canonical selections. According to Ragni et al., the path $r \times c \times d \times (1 - s)$ should lead to selection $p\bar{q}$. This expectation is incorrect, as can be seen in the full model in Figure 3. The expected selections are $\bar{p}$ and $\bar{p}q$.

Figure 5. Fit residuals when fitting MT to the four canonical selections