[In press in *Cognition*, 2020]

A test of two processes: The effect of training on deductive and inductive reasoning

Rachel G. Stephens^a, John C. Dunn^b, Brett K. Hayes^c & Michael L. Kalish^d

a. School of Psychology, University of Adelaide

Adelaide SA 5005, Australia

rachel.stephens@adelaide.edu.au

b. School of Psychological Science, University of Western Australia

Perth WA 6009, Australia

john.dunn@uwa.edu.au

c. School of Psychology, University of New South Wales

Sydney NSW 2052, Australia

b.hayes@unsw.edu.au

d. Department of Psychology, Syracuse University

Syracuse NY 13244, USA

mlkalish@syr.edu

Address for correspondence:

Rachel Stephens, School of Psychology, University of Adelaide, Adelaide SA 5005,
Australia.

Email: rachel.stephens@adelaide.edu.au Phone: +61 8 8313 2817

Abstract

Dual-process theories posit that separate kinds of intuitive (Type 1) and reflective (Type 2) processes contribute to reasoning. Under this view, inductive judgments are more heavily influenced by Type 1 processing, and deductive judgments are more strongly influenced by Type 2 processing. Alternatively, single-process theories propose that both types of judgments are based on a common form of assessment. The competing accounts were respectively instantiated as two-dimensional and one-dimensional signal detection models, and their predictions were tested against specifically targeted novel data using signed difference analysis. In two experiments, participants evaluated valid and invalid arguments, under induction or deduction instructions. Arguments varied in believability and type of conditional argument structure. Additionally, we used logic training to strengthen Type 2 processing in deduction (Experiments 1 & 2) and belief training to strengthen Type 1 processing in induction (Experiment 2). The logic training successfully improved validity-discrimination, and differential effects on induction and deduction judgments were evident in Experiment 2. While such effects are consistent with popular dual-process accounts, crucially, a one-dimensional model successfully accounted for the results. We also demonstrate that the one-dimensional model is psychologically interpretable, with the model parameters varying sensibly across conditions. We argue that single-process accounts have been prematurely discounted, and formal modeling approaches are important for theoretical progress in the reasoning field.

Keywords: Inductive and deductive reasoning; dual-process theories; single-process theories; signed difference analysis; signal detection theory; training

1. Introduction

A widespread view is that there are two types of processes in high-level cognition (see Evans & Stanovich, 2013; Melnikoff & Bargh, 2018), as epitomized by the well-known *Star Trek* characters, Captain Kirk and Mr. Spock. Kirk reasons via gut-feelings and intuitions, while Spock generally applies cold analytical thinking and logic. For a given problem, it seems that people can reason either like Kirk or like Spock. In the lab, researchers have studied this using an argument evaluation task (e.g., Evans, Handley, Harper, & Johnson-Laird, 1999; Rips, 2001; Rotello & Heit, 2009). In this task, participants consider arguments such as:

If the US cuts fuel emissions then global warming will be reduced. (1)

The US did not cut fuel emissions.

Global warming was not reduced.

Some participants are given *induction* reasoning instructions, in which they are asked to judge whether the conclusion below the line is *plausible* based on the premises above the line.¹ Others are given *deduction* reasoning instructions in which they judge whether the conclusion *necessarily follows* from the premises. For Argument (1), under induction instructions people may reason more like Kirk and use their prior beliefs about fuel emissions and global warming to decide that the conclusion is plausible. In contrast, under deduction instructions, if people correctly apply Spock-like logic, the conclusion would be deemed not necessarily true (the argument structure is *denying the antecedent*, which is logically invalid). Though these might appear to be different ways of drawing inferences or conclusions, a key question is whether they reflect the operation of qualitatively different cognitive processes.

¹ Note that our focus is on “inductive reasoning” in the sense of assessing novel predictions (i.e., uncertain conclusions) in light of existing knowledge (i.e., given premises) (Hayes & Heit, 2017), as opposed to other definitions such as generalizing from specific exemplars to broader categories.

Popular dual-process theories propose that there are distinct “Type 1” and “Type 2” processes in human reasoning, judgment and decision making. Such views have been highly influential, with programs based on these theories now advocated in education and assessment (Gillard, Van Dooren, Schaeken, & Verschaffel, 2009; Stanovich, 2016), medical diagnosis (Croskerry, Singhal, & Mamede, 2013), and managerial decision making (Dane & Pratt, 2007), and the concept is being taken up in industry to try to avoid reasoning errors (see Melnikoff & Bargh, 2018). Type 1 processing is generally assumed to be intuitive: It is autonomous, does not require working memory, tends to be fast, and tends to produce responses biased by background knowledge. In contrast, Type 2 processing is seen as reflective: It involves effortful hypothetical thinking, requires working memory, tends to be slow, and tends to produce normative responses (see Evans & Stanovich, 2013). Some theorists propose that the two kinds of processes operate in parallel (e.g., Handley & Trippas, 2015; Sloman, 1996, 2014), while others suggest that Type 1 processing generates intuitive default responses, which may or may not be altered by subsequent high-effort Type 2 processing (e.g., De Neys, 2012; Evans, 2007, 2008; Kahneman & Frederick, 2002). Regardless of the particular version that is preferred, according to dual-process theories, when people consider a reasoning problem such as Argument (1), they could access distinct assessments of argument strength based on Type 1 or Type 2 processes. It is often assumed that induction judgments are particularly dependent on Type 1 processes, while deduction judgments are more dependent on Type 2 processes (Evans, Handley, & Bacon, 2009; Evans & Over, 2013; Rotello & Heit, 2009; Singmann & Klauer, 2011; Verschueren, Schaeken, & d’Ydewalle, 2005).

In contrast, single-process theories propose that a common core process underlies responding in various reasoning, judgment and decision making tasks (cf. Keren, 2013; Keren & Schul, 2009; Kruglanski, 2013; Kruglanski & Gigerenzer, 2011; Osman, 2004, 2013). Under this view, both induction and deduction judgments for reasoning problems such as Argument

(1) are based on a common assessment of subjective argument strength (Rips, 2001). One possibility is that this strength-assessment may be produced by generating and testing mental models of the premises and conclusions (Johnson-Laird, 1994). Another is that it is based on the perceived conditional probability of the conclusion given the premises (Lassiter & Goodman, 2015; Oaksford & Chater, 2001, 2007).

Dual-process accounts are often framed as verbal models, and a key form of empirical support for them is the existence of functional dissociations (Evans, 2008; Evans & Stanovich, 2013) – including important demonstrations that particular factors affect induction judgments more than deduction judgments, or vice versa (for reviews see e.g., Heit, Rotello, & Hayes, 2012; Stephens, Dunn, & Hayes, 2018). In many studies demonstrating such dissociations, arguments are presented like those in Table 1, which vary in both logical validity and prior believability (based on background knowledge), and participants are asked to evaluate the arguments according to deduction or induction instructions. Factors such as consistency with background causal knowledge, argument length, and premise-conclusion similarity have a greater effect on induction judgments (e.g., Handley, Newstead, & Trippas, 2011; Heit & Rotello, 2010; Rips, 2001; Rotello & Heit, 2009; Singmann & Klauer, 2011), while argument validity, working memory load, and cognitive ability have a stronger impact on deduction judgments (e.g., Evans, Handley, Neilens, & Over, 2010; Heit & Rotello, 2010; Howarth, Handley, & Walsh, 2016; Rotello & Heit, 2009).

Table 1

Example Causal Conditional Arguments used in Experiments 1 and 2

Argument type	Validity	P/C	Believable	Unbelievable
Affirmation	Valid (<i>Modus ponens</i>)	P1	If a company advertises during the Super Bowl then the company's sales will increase.	If contraception is cheaper then there will be more pregnancies.
		P2	The company advertised during the Super Bowl.	Contraception was cheaper.
		C	The company's sales increased.	There were more pregnancies.
	Invalid (<i>Affirming the consequent</i>)	P1	If a company advertises during the Super Bowl then the company's sales will increase.	If contraception is cheaper then there will be more pregnancies.
		P2	The company's sales increased.	There were more pregnancies.
		C	The company advertised during the Super Bowl.	Contraception was cheaper.
Denial	Valid (<i>Modus tollens</i>)	P1	If a company advertises during the Super Bowl then the company's sales will increase.	If contraception is cheaper then there will be more pregnancies.
		P2	The company's sales did not increase.	There were not more pregnancies.
		C	The company did not advertise during the Super Bowl.	Contraception was not cheaper.
	Invalid (<i>denying the antecedent</i>)	P1	If a company advertises during the Super Bowl then the company's sales will increase.	If contraception is cheaper then there will be more pregnancies.
		P2	The company did not advertise during the Super Bowl.	Contraception was not cheaper.
		C	The company's sales did not increase.	There were not more pregnancies.

Note. P1 = premise 1, P2 = premise 2, C = conclusion.

However, as many have argued (e.g., Dunn & Kirsner, 1988; Newell & Dunn, 2008; Prince, Brown, & Heathcote, 2012; Stephens et al., 2018; Stephens, Matzke, & Hayes, 2019), although such dissociations are consistent with dual-process accounts, they do not provide compelling evidence for the existence of more than one underlying mechanism for assessing arguments. For this reason, important tests of competing single-process and dual-process

theories in reasoning have instead been based on the more rigorous logic of reversed association or state-trace analysis (Bamber, 1979; Dunn & Kalish, 2018; Dunn & Kirsner, 1988), or signed difference analysis (Dunn & Anderson, 2018; Dunn & James, 2003). State-trace analysis tests for evidence against simple single-process models that postulate one key underlying mechanism or latent variable, while signed difference analysis allows the testing of more detailed, formal single- and dual-process models.

In the sections that follow, we review evidence based on state-trace analysis of reasoning data and on formal reasoning models tested via signed difference analysis (i.e., Hayes, Stephens, Ngo, & Dunn, 2018; Rips, 2001; Singmann & Klauer, 2011; Stephens et al., 2018). This review shows that although the simplest single-process models can be rejected, a more complex version based on the signal detection framework can successfully account for induction and deduction judgments – this is a viable model against dual-process competitors. We then discuss how the successful single-process model can be more explicitly tested, and present two new experiments that do so via manipulations including training participants in deductive logic. We use signed difference analysis to test the model in its most general form. To foreshadow, despite the popularity of dual-process accounts, a single-process model can also explain both extant and new data. We also test a more specific (Gaussian) version of the single-process model to show that the model parameters are psychologically interpretable.

1.1. Tests of the Simplest Single-Process Model Using State-Trace Analysis

According to one of the simplest possible single-process models, induction and deduction judgments are governed by a single underlying latent variable, corresponding to a common psychological dimension of argument strength (Rips, 2001). A key implication of this account is that endorsement rates for arguments assessed under induction versus deduction instructions must be monotonically related; the model “forbids” an ordinal pattern of difference between two conditions in which induction endorsements increase while deduction

endorsements decrease, or vice versa (i.e., contributing to a reversed association; see Dunn & Kirsner, 1988). This prediction holds without having to commit to strong (and potentially false) assumptions about exactly how the single latent variable maps onto observed endorsement rates – simply that the mapping functions are monotonic. However, there has been some – albeit inconsistent – evidence against a monotonic relationship between induction and deduction judgments.

Rips (2001) initially demonstrated a monotonicity violation when there was conflict between the validity and believability of arguments. Rips found that valid-unbelievable arguments received higher deduction endorsements than invalid-believable arguments, but this pattern was reversed for induction judgments. However, Stephens et al. (2018) applied state-trace analysis to these data (see Dunn & Kalish, 2018), which involved plotting induction endorsement rates against deduction endorsement rates, and examining the fit of a model that assumes a single latent variable and thus a monotonic relationship (using an appropriate conjoint monotonic regression [CMR] procedure; Kalish, Dunn, Burdakov, & Sysoev, 2016). They found that the model closely approximated the data, so there was no clear evidence for more than one kind of argument strength assessment. Similarly, Hayes et al. (2018) applied state-trace analysis to three experiments where people made deduction or induction judgments about arguments varying in validity and believability. Notably, these experiments tested several factors relevant to important theoretical distinctions between Type 1 and 2 processing: working memory load, working memory capacity, and decision time (e.g., De Neys, 2012; Evans & Stanovich, 2013; Handley & Trippas, 2015). Although there were clear dissociations between argument endorsement rates in induction and deduction, Hayes et al. (2018) found the data could be explained by variations in a single latent variable (again, based on CMR tests).

However, Singmann and Klauer (2011) found larger violations of monotonicity across deduction and induction tasks. In two experiments they manipulated validity, believability

(plausible, implausible and neutral) and argument type. Argument types were “affirmation” causal conditional arguments that contrast the valid *modus ponens* (if A then B, A, therefore B) with the invalid *affirming the consequent* (if A then B, B, therefore A), and “denial” causal conditional arguments that contrast the valid *modus tollens* (if A then B, not B, therefore not A) with the invalid *denying the antecedent* (if A then B, not A, therefore not B). Stephens et al. (2018) applied state-trace analysis to data from both experiments, and this time found reliable departures from monotonicity, rejecting the hypothesis that induction and deduction judgments were driven by a single latent variable. This result is consistent with the view that there are two distinct psychological dimensions of argument strength, as predicted by dual-process accounts. However, by themselves, demonstrations of reversed associations or monotonicity violations are silent on exactly what the multiple latent variables may be. For this reason, it is important to specify and test formal models that instantiate the competing theories, with latent variables defined by the model parameters. More complex single-process models might – and indeed, can – account for non-monotonic data like those from the Singmann and Klauer (2011) experiments (Stephens et al., 2018).

1.2. Signal Detection Models of Reasoning

Signal detection theory offers a useful framework for formulating and testing single- and dual-process accounts of the argument evaluation task (Heit & Rotello, 2010; Rotello & Heit, 2009; Rotello, Heit, & Kelly, 2019; Stephens et al., 2018). As shown schematically in Figure 1, under this framework, arguments are assumed to vary along continuous dimension(s) of subjective argument strength. On the one hand, dual-process models are *two-dimensional* (2D) – they assume that induction and deduction judgments are based on two different strength dimensions, one based primarily on the output of Type 1 processing and the other based primarily on the output of Type 2 processing, respectively (Figure 1b). On the other hand, single-process models are *one-dimensional* (1D) – they assume only a single strength

dimension, such that both induction and deduction judgments are based on a common assessment of argument strength (Figure 1a). Both model classes assume that valid and invalid arguments form distinct distributions in their 2D or 1D space, with the relative distance between them reflecting the extent to which participants distinguish the two argument types. This distance is captured respectively by two discriminability parameters for the 2D model (d_D for deduction and d_I for induction), or a single discriminability parameter (d) for the 1D model. The models also assume that in the argument evaluation task, participants set a decision threshold or criterion, endorsing only those arguments that sit above the criterion in strength. Figure 1 shows the most general model variants, labelled the *independent-1D* and *independent-2D* models by Stephens et al. (2018) – these variants have distinct, independent criterion parameters for deduction and induction judgments (c_D and c_I , respectively). Notably, this “single-process” model actually has three parameters or latent variables, while the corresponding “dual-process” model has four. However, simpler 1D and 2D model variants can also be tested, such as those that include only one shared criterion parameter for both deduction and induction (i.e., *dependent-1D* and *-2D* models), or those that fix the criterion across different experiment conditions (e.g., factorial combinations of believability and affirmation/denial argument type) – hence there is no criterion parameter (i.e., *fixed criterion-1D* and *-2D* models).

Figure 1. Two signal detection models of the argument evaluation task. a) A three-parameter single-process model, with discriminability parameter, d , and decision criteria c_I for induction judgments and c_D for deduction judgments. b) A four-parameter dual-process model, which adds distinct discriminability parameters for induction and deduction judgments, d_I and d_D , respectively. See the online article for the color version of this figure. [2-column fitting image]

1.3. Tests of Reasoning Models Using Signed Difference Analysis

The distributions of argument strength in Figure 1 are shown as different Gaussian and Gamma distributions, but we use these for illustrative purposes only. As we and others have noted (e.g., Dunn, 2008; Loftus, 1978; Rouder, Pratte, & Morey, 2010; Stephens et al., 2018), the true forms of such response distributions are unknown. Hence, it is prudent to test rival signal detection models of reasoning using an approach that makes only minimal assumptions about the relationship between changes in model parameters and the observed induction and deduction endorsement rates. The approach that we adopt, signed difference analysis (SDA;

Dunn & Anderson, 2018; Dunn & James, 2003), simply assumes that this relationship is monotonic; for example, if some manipulation produces a positive shift in the model parameters we will see an increase (or at a minimum, no decrease) in argument endorsements. SDA is a natural extension of state-trace analysis. While the latter involves testing for evidence against a model with a single latent variable or parameter in two-dimensional data space (e.g., induction vs. deduction endorsement rates), SDA can be used to test for evidence against more complex models, in higher dimensions.

In the current application of SDA, the data space is four dimensional, based on induction and deduction endorsement rates for valid and invalid arguments (Stephens et al., 2018), as shown along the x-axis in Figure 2a. Endorsement rates for deduction-valid, deduction-invalid, indication-valid and induction-invalid form the “dependent variables”, and SDA involves testing observed *ordinal patterns of difference* between “conditions” across these four dependent variables. For example, the “hypothetical” data in Figure 2a are based on the results observed by Evans et al. (2009) from an argument evaluation task with causal conditional arguments (the non-speeded conditions). In the Figure, Condition 2 roughly corresponds to a condition with affirmation arguments (modus ponens and affirming the consequent) and low-believability, and Condition 1 roughly corresponds to a condition with denial arguments (modus tollens and denying the antecedent) and high-believability. The ordinal pattern of difference here is that deduction-valid and induction-valid endorsements are higher in Condition 2 than in Condition 1, plus deduction-invalid and induction-invalid endorsements are lower in Condition 2 than in Condition 1. Note that this pattern is consistent with more accurate validity discrimination in Condition 2 relative to Condition 1 (i.e., endorsing more valid arguments and rejecting more invalid arguments).

Figure 2. Hypothetical examples of two different possible ordinal patterns of difference between conditions, captured by signed difference vectors: a) $\pm(+, -, +, -)$; b) $\pm(+, -, -, +)$.

[2-column fitting image]

Crucially, each signal detection model can be shown to predict some ordinal data patterns but not others. Observation of ordinal patterns that are “forbidden” by a given 1D or 2D model would falsify that model. These forbidden patterns can be formally identified (see Stephens et al., 2018) – but are also often consistent with the contrasting intuitive assumptions of each reasoning model. For example, all 1D models assume that induction and deduction judgments share a single discriminability parameter. Hence, across two experimental conditions, validity discrimination should never be found to both increase for deduction and decrease for induction (or vice versa). A hypothetical example of this qualitative forbidden pattern is illustrated in Figure 2b (and will be discussed in more detail shortly); relative to Condition 2, Condition 1 suggests more accurate validity discrimination for deduction (i.e., a higher endorsement rate for valid arguments and a lower endorsement rate for invalid arguments), but less accurate discrimination for induction (i.e., a lower endorsement rate for valid arguments and a higher endorsement rate for invalid arguments).

Stephens et al. (2018) used SDA to test a set of 16 1D and 2D models (the independent-1D and independent-2D models described above, plus variants with more restricted criteria parameters) against a large database of 26 experiments with induction and deduction endorsement rates for valid and invalid arguments (including the data from Rips, 2001, and Singmann & Klauer, 2011). These experiments included a wide range of factors such as variation in the number of premises, causal consistency or believability, the similarity between categories in the premise and conclusion statements, the cognitive ability of participants, time pressure, and the pleasantness of the content. Stephens et al. (2018) also conducted a new experiment involving a manipulation of perceived base-rates, to test the prediction of the independent-1D model (vs. the competing dependent-2D model) that induction and deduction response criteria are indeed independent and thus can be pushed in opposite directions.

Applying an SDA extension of the CMR statistical test to these datasets (see Experiment 1 Results below for further details on the CMR test), Stephens et al. (2018) found that all models were ruled out except the independent-1D model and the independent-2D model. Given that the independent-1D model has three parameters while the 2D variant is a saturated model with four parameters, the former may be preferred on the grounds of parsimony. At a minimum, the success of the independent-1D model shows that the distinction between Type 1 and Type 2 processing, or any similar distinction that assumes two separate assessments of argument strength, is not required to account for extant data from the argument evaluation task. This evidence of a viable single-process model is especially impressive given that the Stephens et al. (2018) database included variation in factors such as cognitive ability, causal consistency, argument structure, and time pressure that have all been claimed by previous researchers to differentially affect Type 1 and Type 2 processing.

However, an important limitation of the Stephens et al. (2018) archival analysis is that none of the reasoning experiments examined were specifically designed to test the competing

predictions of the independent-1D and independent-2D models. Therefore, the primary goal of the current experiments was to perform a more targeted test of the independent-1D model, searching for the critical evidence that would reject it in favor of the more complex independent-2D model.

1.4. New Tests of the Independent-1D Model

In order to perform a targeted SDA test of the independent-1D model, its permitted and forbidden ordinal data patterns must be understood. In SDA, the relevant ordinal differences between conditions can be captured by *signed difference vectors* with four elements corresponding to the dependent variables: in this case, (deduction-valid, deduction-invalid, induction-valid, induction-invalid). Each element can take the value +, −, or 0, although it is rare to observe mean differences of exactly zero. Figure 2 presents two hypothetical examples of different possible signed difference vectors. Figure 2a shows an example of $\pm(+, -, +, -)$, in which relative to condition 1, the condition 2 means are higher for deduction-valid and induction-valid, but lower for deduction-invalid and induction-invalid. Figure 2b shows an example of $\pm(+, -, -, +)$, in which relative to condition 2, the condition 1 means are higher for deduction-valid and induction-invalid, but lower for deduction-invalid and induction-valid. In total, there are 40 possible signed difference vectors (see Stephens et al., 2018).

Different models specify different forbidden and permitted data patterns. Stephens et al. (2018) identified that the independent-1D model has one forbidden sign vector, $\pm(+, -, -, +)$, as shown in Figure 2b. This is the pattern we considered earlier – it is consistent with separate discriminability parameters for deduction and induction judgments, because validity discrimination rates can shift in opposite directions. For example, in the Figure, relative to Condition 2, in Condition 1 deduction judgments show higher discrimination of valid versus invalid arguments (i.e., a higher endorsement rate for valid arguments and a lower endorsement rate for invalid arguments), while the induction judgments show the opposite pattern. This

pattern is consistent with the saturated independent-2D model and if observed, it would count as compelling evidence against the independent-1D model.

If there are distinct Type 1 and Type 2 processes that differentially affect induction and deduction judgments, it should be possible to observe the critical pattern forbidden by the independent-1D model, $\pm(+, -, -, +)$. How might an experiment produce this complex pattern? It is most likely to be observed between two conditions formed by a *combination* of experimental factors, such as untrained versus trained judgments, believability, and type of argument form. The first part of the critical pattern (i.e., the first $(+, -, \dots)$ for deduction-valid, deduction-invalid) requires an experimental manipulation that improves validity discrimination, more so for deduction than for induction. This might be achieved by training participants in deductive logic.

The potential for a training manipulation to improve validity discrimination is highlighted by the fact that such discrimination is often poor for the typical untrained undergraduate participant, for many types of logical arguments. For example, across all conditions included in the Stephens et al. (2018) database, mean deduction endorsement rates were .80 ($SD=.14$) for valid arguments and .34 ($SD=.27$) for invalid arguments, which suggests room for improvement towards the normative values of 1 and 0, respectively. A training manipulation also has theoretical importance for tests of single- and dual-process reasoning models; in particular, it is possible that participants' general lack of logic training masked the observation of distinct Type 1 and 2 processing in the Stephens et al. (2018) database. If participants do not know how to correctly assess validity for some argument forms, perhaps Type 2 processing has simply not been able to exert an influence on responses that is sufficiently distinguishable from that of Type 1 processing. In addition to teaching participants how to assess validity correctly, logic training may also further clarify the distinction between

the induction and deduction tasks. Thus post-training there may be a stronger influence of Type 1 processing in induction and Type 2 processing in deduction.

For these reasons, we include logic training as a key factor in our experiments. In Experiment 1, we train both induction and deduction participant groups on how to distinguish valid and invalid arguments. This confirms that the logic training procedure is successful, and examines people's judgments when only logic-based reasoning is trained. In Experiment 2, to further magnify differences between deduction and induction judgments after training, we also train both groups on how to distinguish arguments with believable and unbelievable content. Thus, we train people in both "logic-based" reasoning and "belief-based" reasoning, in principle creating the ideal conditions for them to apply different types of processing in induction and deduction.

The second part of the critical $\pm(+, -, -, +)$ pattern, (i.e., the latter $(-, +)$ for induction-valid, induction-invalid) might be more likely to appear if – for untrained participants – there is one argument type for which valid arguments are endorsed more often than invalid arguments, and another argument type for which the opposite occurs. This kind of responding has been observed before, for causal conditional arguments with content about real-world events and believable versus unbelievable variants, similar to those shown in Table 1. For affirmation arguments, participants typically endorse valid arguments (modus ponens) more than invalid arguments (affirming the consequent), but for denial arguments (modus tollens vs. denying the antecedent), this effect is often reversed. For instance, these effects were found by Evans et al. (2009), Singmann and Klauer (2011), Trippas et al. (2014), and the new experiment by Stephens et al. (2018). To illustrate, as mentioned earlier, the "hypothetical" results in Figure 2a are based on the results observed by Evans et al. (2009; these data were included in the Stephens et al., 2018 database), with Condition 2 roughly corresponding to their (non-speeded) affirmation, low-believability condition, and Condition 1 roughly corresponding to

their (non-speeded) denial, high-believability condition. It is suggested that participants rely primarily on the believability of the conditional when assessing the validity of modus tollens (Evans et al., 2010; Singmann & Klauer, 2011).

Considering the pattern of responses in Figure 2a, note that if in a new Condition 3, validity discrimination can be “corrected” for deduction judgments in Condition 1 (but induction judgments are not altered), this would produce the (idealized) data pattern in Figure 2b that is forbidden by the independent-1D model. Training may produce this kind of effect, as discussed above. To be clear, to create the ordinal $\pm(+, -, -, +)$ pattern, training needs to have a *larger* impact on deduction than induction – not necessarily *no* impact on induction. Note that the hypothetical conditions in Figure 2b would then correspond to a combination of different experiment factors: trained, denial, high-believability responses (Condition 1) versus untrained, affirmation, low-believability responses (Condition 2). Therefore, in addition to training, our experiments also include the factors of argument type (affirmation versus denial causal conditional arguments) and believability (low, high, and neutral, with the latter included simply to increase opportunities for observing the critical SDA forbidden pattern).

Although we do not know of any previous studies of the effects of logic training (and/or belief training) on both induction and deduction judgments, some have investigated the effects of logic training within the laboratory on deductive reasoning. Some training procedures have been shown to improve accuracy for various deduction tasks, including evaluating categorical syllogisms (Prowse Turner & Thompson, 2009), generating conclusions from conditional (and other) premise structures (Klauer, Meiser, & Naumer, 2000; Klauer, Stegmaier, & Meiser, 1997) and the related Wason selection task (e.g., Cheng, Holyoak, Nisbett, & Oliver, 1986; Klaczynski & Laipple, 1993). Although it is unclear exactly which training components are essential, explanation of valid and invalid inferences with concrete examples, plus practice with

immediate feedback appear to be beneficial. Therefore, we include these features in our logic training procedure.

1.5. Further Exploratory Model Testing

Our primary goal is to use signed difference analysis to test the independent-1D and -2D models in their most general form, with minimal distributional assumptions. As we have argued, this SDA testing is important because the “true” distributional forms are unknown. The approach is also powerful because if SDA rules out the independent-1D model, all possible variants of the model with stronger distributional assumptions are also ruled out. However, given that we end up retaining both the independent-1D and -2D models, our subsequent goal is to assess them further, in a more exploratory fashion with stronger distributional assumptions (i.e., Gaussian distributions). This allows us to examine the models’ best-fitting parameter values for the experimental conditions. First, we investigate whether the two discriminability parameters of the independent-2D model are correlated. If so, this model is mimicking the independent-1D model and thus its extra complexity is unwarranted. Second, we examine the interpretability of the discriminability and two criterion parameters of the independent-1D model; how do they vary across experimental conditions?

2. Experiment 1

The primary aim of Experiment 1 was to test for the ordinal pattern forbidden by the independent-1D model, using signed difference analysis. The dependent variables for SDA were induction and deduction endorsement rates (manipulated between groups), elicited for a common set of valid and invalid arguments. Three within-participants factors were the type of causal conditional argument (denial vs. affirmation), the believability of the content (unbelievable, believable, or neutral) and logic training (pre- vs. post-training block). Participants completed three tasks: a pre-training block of 36 trials, then the logic training and test, then a post-training block of another 36 trials.

2.1. Materials and Methods

2.1.1. Participants. There were 94 undergraduates at Syracuse University, USA, who participated for course credit (informed consent was obtained). Mean age was 19.2 years ($SD = 1.1$), and there were 45 males (47.9%). Four participants reported having received some prior formal logic training (e.g., a philosophy class). However, inspection of their pre-training responses in the first block suggested that the degree to which their decisions matched argument-validity was well below ceiling (64-72% of trials), so they were not excluded. Participants were randomly assigned to the instruction conditions: 46 for induction and 48 for deduction.

2.1.2. Stimuli. The conditional statements used for the arguments in the two main experiment blocks referred to real world events and were based on those used by Evans et al. (2009), but adapted for participants in the USA (see examples in Table 1). The full stimulus sets are available in the supplementary materials. To create believable and unbelievable argument content, a list of 104 conditional statements was rated by a separate group of 27 participants, drawn from the same population as the main study. These participants were asked to rate their certainty that each statement is true, from "definitely false" to "definitely true" (0-100). From this, 24 high believability ($M = 73.9$, $SD = 5.0$) and 24 low believability statements ($M = 24.3$, $SD = 4.9$) were selected for Experiment 1 (see Appendix A). These 48 statements were split into two equivalent sets (signaled in Appendix A), so that each statement in the first set was very closely matched in believability rating to a corresponding statement in the second set. Thus we created matched pairs of statements that would permit yoking across the pre- and post-training blocks. To create arguments with *neutral* believability, another 24 conditional statements were constructed that included a nonsense word, such as "If the surgeon is scomp then the patient will recover quickly" (see Appendix A for the full list). These neutral

statements were also arbitrarily assigned to the two sets, resulting in a total of 36 statements per set (i.e., 12 neutral, 12 believable, 12 unbelievable).

The final collection of 72 statements was used to create the four kinds of conditional inference problems (modus ponens, accepting the consequent, modus tollens and denying the antecedent). For each participant, to form the 36 pre-training and 36 post-training trials, the 36 matched pairs of statements were each randomly allocated to a problem kind, generating three arguments per experiment cell, as defined by the four problem kinds and the three believability levels. These matched pairs were presented such that one randomly appeared in the pre-training block, and the other appeared in the post-training block, thus yoking combinations of believability-rating and problem-kind across blocks.

2.1.3. Procedure. Participants completed the reasoning task individually, within a one-hour laboratory session. Software for the experiments was programmed in Matlab using the Psychophysics Toolbox extensions (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007). During the pre- and post-training blocks, each argument was presented one-by-one, with a line separating the conclusion from the premises. Participants were required to evaluate each argument according to either deductive or inductive criteria, as described in the instructions below. Responses were made by clicking either the Valid or Invalid button (for deduction), or the Strong or Weak button (for induction), presented underneath a given argument. Participants then rated their confidence on a scale that appeared, ranging from 50 (Guessing) to 100 (Certain) in increments of 10 (the confidence ratings will not be considered for the present analyses).

At the start of the experiment, participants in the deduction condition were given the following instructions:

In this task, you will see some arguments and be asked to make judgements about the VALIDITY of the arguments.

For example, you will see an argument like this:

If it snows more then the temperature will increase

It snowed more

The temperature increased

and will be asked to judge how valid this argument is.

'Valid' arguments are those for which the sentence below the line is NECESSARILY TRUE, given that the information above the line is true.

'Invalid' arguments are those for which the sentence below the line is NOT NECESSARILY TRUE, given that the information above the line is true.

For each argument, you should assume that the information above the line is true.

Induction instructions were identical, except they referred to argument strength rather than validity, and defined the response options as:

'Strong' arguments are those for which the sentence below the line is PLAUSIBLE, assuming that the information above the line is true.

'Weak' arguments are those for which the sentence below the line is IMPLAUSIBLE, assuming that the information above the line is true.

We asked participants to take their time, and told them they would receive immediate feedback from the computer program at the end of the study. We also forewarned that some of the arguments would include a nonsense word (i.e., neutral items), and for these they should make their best guess. Throughout the study, for each block a trial counter was presented at the top left corner of the screen. The 36 arguments within each block were presented in random order for each participant. The pre- and post-training blocks were identical, except for the argument content and trial order.

Between these two blocks, participants received thorough automated logic training on the causal conditional arguments, introduced as explaining “how a logic expert might evaluate the kinds of arguments you have seen so far”. The training was identical for induction and deduction conditions. Participants first read through some materials at their own pace, presented across several screens, between which they could browse freely (available as supplementary material). The materials clarified how to interpret the conditional statement using an Euler diagram, and explained how to use the diagram to assess validity for each of the four kinds of conditional arguments. Participants then completed a short test to help them apply the training, in which they made valid/invalid judgments. The test involved 12 arguments generated from three new neutral conditional statements, repeated for each of the four kinds of argument structure. Feedback was provided after each response, which also reminded participants *why* the argument was valid or invalid. At the end of the test, accuracy (i.e., percentage of correct “valid” or “invalid” responses for valid or invalid items, respectively) was reported for each of the four argument kinds, along with a brief reminder of the actual validity of each kind.

Before the post-training block, participants were reminded of the induction or deduction instructions given at the start of the experiment. To give induction participants “permission” to again consider argument content if they wanted to, the instructions were prefaced with:

We have now taught you how logicians determine whether the form of an argument is valid or invalid. However, we are interested in how people generally judge whether an argument is weak or strong. In real life, the content or meaning of arguments can also be important.

2.2. Results and Discussion

Mean endorsement rates for each condition are shown in Figure 3, plotted according to the design for the signed difference analysis, with the four SDA dependent variables on the x-

axis, and the 12 conditions as different lines (raw data from both experiments are available as supplementary material). Prior to training, as has been observed previously (e.g., Evans et al., 2009; Stephens et al., 2018; Trippas et al., 2014), for affirmation arguments, valid arguments were endorsed more often than invalid arguments, but for denial arguments, this pattern was generally inverted. Notably, inspection of Figure 3 suggests that the logic training had a substantial impact, but the figure does not reveal any ordinal patterns corresponding to the vector forbidden by the independent-1D model (cf. Figure 2b).

Given the different response patterns for affirmation and denial arguments, we first analyzed the endorsement rates separately for these two argument types, using more customary statistical tests to examine the impact of instructions, validity, believability and training block. The goals here were to examine whether we replicated typical effects of these four factors, and to see whether there were dissociation effects similar to those that have traditionally been used to support dual-process theories. Second, we carried out the critical SDA model testing. Although no cases of the forbidden vector are apparent in Figure 3, with affirmation and denial conditions displayed in separate plots, SDA involves a comprehensive and rigorous test for this data pattern by examining the ordinal patterns for *all* pairwise combinations of experimental conditions (including between pairs of conditions involving affirmation versus denial arguments).

Before proceeding to the analyses of data from the main experiment blocks, we confirmed that both induction and deduction participants engaged with the logic training materials. Indeed, accuracy (i.e., correct endorsement of valid arguments and correct rejection of invalid arguments) for the 12 test trials during the training phase was quite high, and similar for both instruction groups, $t(92) = .13$, $p = .89$, with mean accuracy of .78 ($SD = .20$) for the induction group and 0.78 ($SD = .22$) for the deduction group.

Figure 3. Mean endorsement rates from Experiment 1, broken down by training block (pre- and post-training) and argument type (affirmation and denial). X-coordinates are slightly perturbed to aid visibility. Error bars show *SEM*. [2-column fitting image]

2.2.1. Tests of experimental effects. Because the data in each cell were based on binary responses to three arguments for each participant, for the initial analysis we used generalized linear mixed models (GLMM), with a binomial family and logit link function, and random intercepts for participants, implemented in R (Bates, Maechler, Bolker, & Walker, 2015; to avoid non-convergence, we used the adaptive Gauss-Hermite optimizer, based on 10 integration points). To examine the effect of each factor (induction/deduction reasoning instructions, validity, believability, and pre/post-training block), plus two- and three-way

interactions between them, we applied likelihood ratio tests that each compared a “reduced” or nested model *without* the factor (or interaction) of interest against a “full” model *with* the factor. As examples, to test the effect of induction/deduction reasoning instructions, the reduced model did not include any fixed effects, while the full model included instructions as a predictor. Similarly, to test for an interaction between reasoning instructions and training, the reduced model included only those two factors, while the full model also added in the interaction between them. Finally, to test for a three-way interaction such as between instructions, training, and validity, the reduced model included those three factors and the three possible two-way interactions between them, while the full model added in the three-way interaction.

2.2.1.1. Validity and believability effects. For both affirmation and denial arguments, the GLMM analyses revealed a significant effect of validity; respectively, $\chi^2(1) = 967.88, p < .001$ and $\chi^2(1) = 116.28, p < .001$, with valid arguments endorsed more often than invalid arguments overall (but for denial arguments this effect is largely driven by the post-training block – see interaction effects below). There was also a significant effect of believability for each argument type; $\chi^2(2) = 73.52, p < .001$, and $\chi^2(2) = 44.83, p < .001$. Relative to endorsement rates for neutral arguments, believable arguments were not significantly different, $ps > .05$, but unbelievable arguments were endorsed less often, $ps < .001$ (see Figure 3, although note that the figure presents the mean proportion of endorsements, which is different to the data structure used for GLMM). There was also a significant interaction between validity and believability for affirmation arguments, $\chi^2(2) = 15.13, p < .001$, but not for denial arguments, $p = .19$.

The lower endorsement rate for unbelievable than believable arguments is consistent with standard belief bias effects (e.g., Dube, Rotello, & Heit, 2010; Evans, Barston, & Pollard, 1983; Newstead, Pollard, Evans, & Allen, 1992; Trippas, Handley, & Verde, 2013). One possible explanation for the similar endorsement rates between neutral and believable

arguments is that the nonsense words in the neutral conditional statements were assumed to mean something plausible. For example, “scomp” in “If the surgeon is scomp then the patient will recover quickly” might be interpreted as something like “skilled”, “experienced”, or “competent”.

2.2.1.2. Instructions, validity and believability effects. More importantly for our current interest in instructions, Experiment 1 did *not* replicate the effects sometimes observed that 1) validity has a larger effect on deduction than induction and 2) believability has a larger effect on induction than deduction (e.g., Hayes et al., 2018; Rips, 2001; Rotello & Heit, 2009). For both affirmation and denial arguments, the instructions $\times$ validity and instructions $\times$ believability effects showed $p > .05$; argument believability and validity had a similar influence on induction and deduction judgments.

2.2.1.3. Training effects. Turning to the novel training manipulation, a striking feature of Figure 3 is that the logic training had a substantial effect on argument evaluations. Overall, endorsement rates reduced after training (affirmations: $\chi^2(1) = 50.63, p < .001$; denials: $\chi^2(1) = 64.20, p < .001$). This reduction was larger for the deduction than the induction group, for both argument types (affirmations: $\chi^2(1) = 5.01, p = .03$; denials: $\chi^2(1) = 6.52, p = .01$). Otherwise, training had similar effects on both induction and deduction groups.

Crucially, for both affirmation and denial arguments, there was a significant interaction between validity and block, $\chi^2(1) = 166.66, p < .001$ and $\chi^2(1) = 301.43, p < .001$, with training generally “correcting” the validity effects. As shown in the figure, after training, valid arguments were often endorsed and invalid arguments were rarely endorsed. This pattern was reversed for believability effects – after training, argument believability had little effect on endorsements for both argument types, $\chi^2(2) = 49.36, p < .001$ and $\chi^2(2) = 17.00, p < .001$. However, the most interesting three-way interactions (i.e., testing for differential training effects on induction vs. deduction) of instructions $\times$ block $\times$ validity and instructions $\times$ block

× believability were not statistically significant, $p > .05$. All other likelihood ratio tests were not statistically significant, $p > .05$.

2.2.1.4. Summary. We replicated standard effects of validity and belief bias, and the training had a substantial impact, generally heightening validity discrimination and reducing believability effects. Although we specifically trained participants on deductive logic, both induction and deduction groups applied the training to their judgments. Overall, endorsement rates were reduced more for deduction than for induction, but for both affirmation and denial arguments (considered separately), training altered validity (and believability) effects in the same ordinal direction for deduction and induction judgments.

2.2.2. Signed difference analysis. To test the independent-1D model against the data (including pairwise comparisons across all affirmation and denial conditions), we applied the conjoint monotonic regression statistical procedure developed by Kalish et al. (2016) for fitting a one-parameter model to two outcome variables (i.e., state-trace analysis), and extended by Stephens et al. (2018) for higher-dimensional dataspace and models (i.e., SDA). In signed difference terms, state-trace analysis fits a one-dimensional model in which the signed difference between any two conditions cannot be equal to $\pm(+, -)$ which is the forbidden vector for this model. Model goodness-of-fit is measured by the sum of weighted least-squared differences between the observed and predicted means across conditions. A p -value associated with the observed fit value is obtained through a bootstrap resampling procedure – for details see Dunn and Kalish (2018). This is a test of the hypothesis that the data are consistent with the permitted signed difference pattern(s) derived from a given model. The extension to SDA is straightforward. In this case, the data are fit by a given model under the constraint that the signed difference between any two conditions cannot equal any of the forbidden signed difference vectors specified by that model. As with state-trace analysis, a least-squares measure of model fit is calculated via a custom optimization routine, then the corresponding p -value is

estimated by bootstrap resampling of the data (using 10,000 samples in the current tests). If the null hypothesis is retained, the model successfully accounts for the data.

The independent-1D model fit the Experiment 1 data perfectly ($\text{fit} = 0$, $p = 1$); there were no instances of the forbidden signed difference vector, $\pm(+, -, -, +)$, across any pairwise comparisons between the 12 conditions shown in Figure 3. Figure 4 shows an example of two conditions in which the hypothetical pattern in Figure 2b plausibly might have occurred, across a conjunction of the training, argument type and believability factors, as discussed at the outset. Note that the pre-training, affirmation, unbelievable condition replicated the general results observed by Evans et al. (2009; non-speeded task), which were approximately those shown for Condition 2 in Figure 2b. We had proposed that if logic training had a large effect on the deduction group but a limited effect on the induction group, endorsement rates like those shown for Condition 1 in Figure 2b might have been observed, particularly for the post-training, denial, believable condition. However, instead the logic training produced similar response rates for both induction and deduction judgments, as can be seen in Figure 4, and the critical opposing shifts in validity discrimination rates for deduction versus induction were not observed.

Figure 4. Mean endorsement rates from two key conditions based on the conjunction of all independent variables in Experiment 1, showing the pattern $\pm(+, -, +, -)$. X-coordinates are slightly perturbed to aid visibility. Error bars show *SEM*. [single-column fitting image]

2.3. Conclusions

Experiment 1 replicated standard effects of validity and believability on endorsement rates for affirmation and denial arguments, and successfully improved validity discrimination after a short, automated training session. A key aim was to apply SDA to test for the ordinal pattern forbidden by the independent-1D model, but this pattern was not observed. As found by Stephens et al. (2018), although consistent with the saturated 2D model, the data can also be accounted for by the independent-1D model, which has distinct criteria for induction and deduction judgments and a single discriminability parameter.

However, evidence against the 1D model depends upon observing opposing shifts in validity discrimination rates for deduction and induction judgments, between two conditions. We had proposed that this result may have been observed across a combination of factors (argument type, believability and training block) if the logic training had more strongly affected the deduction than the induction group. Instead, we found that logic training had similar effects

on both groups. This occurred despite our post-training instructions to induction participants, indicating that they may also consider the content or meaning of the arguments when evaluating plausibility. However, even with these instructions, there may have been strong demand characteristics for induction participants to use the logic training in the final block, leading to similar responses for induction and deduction groups. To reduce this possibility and perform a stronger test of the models, in Experiment 2 we added some additional “content training” after the logic training, for both induction and deduction groups. The goal was to train participants in both logic-based and belief-based reasoning. This should help to further clarify the distinction between the induction and deduction tasks, and create the ideal conditions for heightened Type 1 processing for induction judgments and heightened Type 2 processing for deduction judgments.

3. Experiment 2

The primary aim of Experiment 2 was to further test for the ordinal pattern forbidden by the independent-1D model. The design of Experiment 2 was identical to Experiment 1 except that between the logic training and the final post-training block, all participants also completed some “content training”, which asked them to consider whether argument content was “sensible”. Changes to the design are noted below. Our intention was that after both training tasks, deduction participants would know how to base their judgments on validity only, and induction participants would understand that both validity and believability can be important for the plausibility or strength of an argument.

3.1. Materials and Methods

3.1.1. Participants. There were 87 students at Syracuse University, who participated for course credit or received US\$10 (informed consent was obtained). Mean age was 22.2 years ($SD = 6.8$), and there were 36 males (41.4%). Twenty-one participants reported having received some prior formal logic training (e.g., a philosophy class, set theory in a mathematics). For 15

of these participants, pre-training validity discrimination was well below ceiling (50-72%), so they were not excluded. The remaining six participants had high pre-training validity discrimination of 81-100%. If these six participants were excluded, the same conclusions were drawn from the GLMM and SDA tests, so we report the fully inclusive results from all 87 participants here. Participants were randomly assigned to the instruction conditions: 43 for induction and 44 for deduction.

3.1.2. Procedure. In the content/belief training task, we explained to participants that during the logic training we had taught them how “logicians” determine validity, however in real life, the content or meaning of arguments is important when evaluating their plausibility. We discussed an example (modus ponens) argument that was valid yet implausible, and commented that, “Although the form of this argument is logically valid, you probably agree that it is NOT plausible - based on what you know about the content. Therefore, while we can apply logic to assess the validity of the abstract FORM of an argument, the content of the argument can affect how convincing it is.”

Next, participants were told they would practice thinking about the content of an argument. They then judged whether 16 arguments were “sensible” or “not sensible”, focusing only on content, and were given feedback after each response: “Correct/Incorrect. Most people think this content is (NOT) sensible”. The first eight arguments all included a bi-conditional statement (“If and only if water gets on the wood shavings then the shavings will be wet” vs. “will ignite”), presented within structures equivalent to the four kinds of conditional arguments. The arguments were introduced as all being logically valid. The second eight arguments included conditional statements with “maybe” (“If the ice-cream is out in the sun then maybe the ice-cream will melt” vs. “will freeze”), also presented within the four structures but introduced as logically invalid. For both sets, participants were instructed to focus on content.

Before the final post-training block began, deduction participants were instructed as follows [with induction variants in square brackets]:

We have now taught you how to correctly determine whether the form of an argument is valid or invalid. We have also asked you to consider whether the content makes an argument sensible or not sensible.

When judging logical validity, content is ignored, but in real life the content or meaning will affect how convincing an argument is.

We are now interested in your judgments about whether an argument is valid or invalid [weak or strong].

In this final task, again you will see some arguments and be asked to make judgments about the VALIDITY [STRENGTH] of the arguments.

Participants were then reminded of the deduction or induction instructions.

3.2. Results and Discussion

Mean endorsement rates for each condition are shown in Figure 5. The pre-training endorsement rates were similar to those in Experiment 1, although notably the pre-training (inverted) validity effects appeared smaller for denial arguments in Experiment 2. Figure 5 suggests that training had a less extreme impact than in Experiment 1, and there are no obvious cases of the ordinal data pattern that would disconfirm the 1D-independent model. Again however, CMR tests are needed for a more rigorous examination of data patterns across all pairwise combinations of conditions. We applied the same GLMM then SDA model tests to the data as for Experiment 1.

Similar to Experiment 1, we first confirmed that accuracy for the 12 test trials during logic training was again quite high, and similar for both instruction groups, $t(85) = .04, p = .97$, with mean accuracy of .82 ($SD = .21$) for the induction group and 0.82 ($SD = .16$) for the deduction group. Accuracy for the 16 test trials during content training was also high, although

slightly higher for the induction group (mean = .89, $SD = .14$) than the deduction group (mean = .82, $SD = .18$), $t(85) = 2.01$, $p = .048$.

Figure 5. Mean endorsement rates from Experiment 2, broken down by training block (pre- and post-training) and argument type (affirmation and denial). X-coordinates are slightly perturbed to aid visibility. Error bars show *SEM*. [2-column fitting image]

3.2.1. Tests of experimental effects. The same GLMM procedure was applied, as per Experiment 1.

3.2.1.1. Validity and believability effects. For both affirmation and denial arguments, the GLMM analyses revealed a significant effect of validity; respectively, $\chi^2(1) = 999.55$, $p < .001$ and $\chi^2(1) = 158.07$, $p < .001$, with valid arguments endorsed more often than invalid

arguments overall. There was also a significant effect of believability (i.e., belief bias) for each argument type, $\chi^2(2) = 73.39, p < .001$ and $\chi^2(2) = 58.52, p < .001$. As in the previous study, Figure 5 suggests that, overall, unbelievable arguments were endorsed less often than believable and neutral arguments. Follow-up tests confirmed that relative to neutral arguments, unbelievable arguments were endorsed less often, $ps < .001$, and believable arguments were not significantly different for affirmations ($p = .12$), but endorsed more often for denials, $p = .049$. As in Experiment 1, there was also a significant interaction between validity and believability for affirmation arguments, $\chi^2(2) = 46.43, p < .001$, but not for denial arguments, $p = .16$.

3.2.1.2. Instructions, validity and believability effects. Turning to instructions effects, in contrast to Experiment 1, for denial arguments (but not affirmation), there was a main effect of reasoning instructions, such that induction endorsement rates were generally lower than deduction, $\chi^2(1) = 3.85, p = .05$. More interestingly, validity and believability now had some differential effects on induction and deduction judgments overall (as found by e.g., Hayes et al., 2018; Rips, 2001). The validity effect was larger for the deduction than the induction group for affirmation arguments, $\chi^2(1) = 5.88, p = .02$, but not for denial arguments, $p = .55$. Conversely, believability effects were larger for induction than for deduction, for denial arguments, $\chi^2(2) = 6.23, p = .04$, but not for affirmation arguments, $p = .07$.

3.2.1.3. Training effects. As in Experiment 1, Figure 5 suggests that the logic training was very effective, although with the additional content training, the effects of argument believability were more evident in the final block of Experiment 2. Overall, endorsement rates reduced after training, for both affirmation and denial arguments, $\chi^2(1) = 99.58, p < .001$ and $\chi^2(1) = 58.05, p < .001$, respectively. As per Experiment 1, this reduction was larger for the deduction than the induction group, for both argument types, $\chi^2(1) = 7.22, p = .01$ and $\chi^2(1) = 28.13, p < .001$. Also, for both affirmation and denial arguments, training generally improved

validity discrimination rates, $\chi^2(1) = 76.78, p < .001$ and $\chi^2(1) = 174.96, p < .001$, such that in the final block, valid arguments were often endorsed and invalid arguments were relatively rarely endorsed. However, unlike Experiment 1, the interaction between training block and believability was not significant for affirmation nor denial arguments, $p = .41$ and $p = .57$; overall, believability effects remained after training.

In contrast to Experiment 1, training had some differential effects on induction and deduction groups; this is the kind of dissociation evidence that has been used to support dual-process accounts. For affirmation (but not denial) arguments, the three-way interactions of instructions $\times$ block $\times$ validity and instructions $\times$ block $\times$ believability were statistically significant; respectively $\chi^2(1) = 35.39, p < .001$ and $\chi^2(2) = 6.97, p = .03$ (see Figure 5a and 5c). Relative to the pre-training block, during post-training the deduction group showed a larger increase in validity discrimination than the induction group. In contrast, believability effects increased for induction but decreased for deduction. All other likelihood ratio tests were not statistically significant, $p > .05$.

3.2.1.4. Summary of key findings. These results indicate that with the content/belief training added, logic training no longer overwhelmed induction endorsement rates in the final block. Potential demand characteristics that may have been present in Experiment 1 were successfully reduced; we were able to encourage the induction group to consider both validity and believability after training. Under a dual-process account, the post-training results are consistent with heightened Type 1 reasoning in induction compared to deduction, and heightened Type 2 reasoning in deduction compared to induction. However, for both affirmation and denial arguments, training still generally altered validity discrimination in the same ordinal direction for deduction and induction judgments. Again, SDA needs to be applied to test for the $\pm(+, -, -, +)$ pattern across all conditions – including between pairs of conditions involving affirmation versus denial arguments.

3.2.2. Signed difference analysis. Despite the apparent dissociations found in the GLMM analyses, as with Experiment 1, the CMR test found that the independent-1D model fit the Experiment 2 data perfectly ($\text{fit} = 0, p = 1$). There were no instances of the forbidden signed difference vector, $\pm(+, -, -, +)$, across any pairwise comparisons between the 12 conditions shown in Figure 5. As considered in Experiment 1, Figure 6 shows an example of two conditions in which the forbidden pattern might have occurred, across a conjunction of the training, argument type and believability factors. Again, the pre-training, affirmation, unbelievable condition replicated the general results observed by Evans et al. (2009; non-speeded task). However, rather than more selectively affecting the deduction group, logic training also had sizeable effects on the induction group. Most notably, the critical opposing shifts in validity discrimination rates for deduction versus induction were not observed.

We also tested the independent-1D model against the data from both experiments combined, forming a total of 24 conditions across an additional between-participants factor of Experiment 1 versus 2. Again, the 1D model could not be rejected ($\text{fit} = .22, p = .69$).

Figure 6. Mean endorsement rates from two key conditions in Experiment 2, showing the pattern $\pm(+, +, -, -)$, or almost $\pm(+, +, 0, 0)$. X-coordinates are slightly perturbed to aid visibility. Error bars show *SEM*. [single-column fitting image]

3.2.3. Alternative interpretation of training effects. A possible alternative interpretation of our results is that during logic training, people simply learnt to recognize valid versus invalid “patterns” or perceptual features of the arguments, and then after training engaged neither Type 1 nor Type 2 reasoning processes. However, while such a pattern-based strategy may make sense for deduction judgments, there is little reason for people to use it for induction judgments where it is nominally irrelevant and relatively effortless Type 1 processing may be applied. Furthermore, believability effects were reduced but not eliminated after training, suggesting that people still considered the content of the arguments, not just surface features.

3.3. Conclusions

Experiment 2 added content training and compared to Experiment 1, successfully induced more differential patterns of endorsement between deduction and induction judgments. However, the ordinal pattern forbidden by the independent-1D model was still not observed;

there was no evidence that demands two independent discriminability parameters (i.e., a dual-process reasoning model). Combining both experiments into one larger dataset also could not reject the 1D model. Therefore, extending the results of Stephens et al. (2018), we find that although consistent with the independent-2D model, induction and deduction judgments can be accounted for by the independent-1D model.

4. Examination of the Models with Distributional Assumptions

The SDA tests have shown that Experiments 1 and 2 do not reject the independent-1D model – in its most general form, with minimal distributional assumptions. Therefore, the 1D model remains as a viable alternative to the saturated independent-2D model – both models can account for the data. While the 1D variant may be preferred on the basis of parsimony, other considerations are also important in model comparison, such as how the parameters vary across conditions and whether they do so in a way that means they are psychologically interpretable. To assess the two competing models on this kind of parameter behavior, stronger distributional assumptions must be made than the monotonicity assumption of SDA, so that we can examine each model's best-fitting parameter values for the experimental conditions. Unfortunately, the “true” distributional forms of the independent-1D and -2D models are unknown, so an important goal for future work is to test systematically a variety of possible distributional forms across a wide range of data from the argument evaluation task. But as an initial examination, here we assume that the invalid and valid distributions for induction and deduction judgments are Gaussian, as has been assumed before for the argument evaluation task by Rotello and Heit (2009; Heit & Rotello, 2010). We focus on exploring Gaussian versions of the models for the Experiment 2 data, where there were larger differences between induction and deduction judgments, which would often be interpreted as evidence for dual-processes (but see Appendix B for the Experiment 1 results, based on the same procedure).

In fitting the Gaussian-independent-1D and -2D models, for all 12 conditions (training block $\times$ believability $\times$ affirmation/denial argument type) the mean for the invalid-deduction and invalid-induction distributions was set to 0 (this defines reference points for estimating the distance between valid and invalid distributions). In line with the 1D model tested by SDA, each of the four distributions (valid and invalid arguments for induction and deduction) was permitted to have a different standard deviation (SD), but each value was fixed across the 12 conditions. Thus for the Gaussian-independent-1D model, the invalid-deduction SD was set to 1 (which sets the scale for the distribution widths), and the three other SD s were estimated from the data, concurrently along with the other parameters. This 1D model then had three additional free parameters (d , c_D , c_I) to account for deduction-valid, deduction-invalid, induction-valid, induction-invalid responses in each condition. In contrast, the Gaussian-independent-2D model was already saturated, so all four SD s were set to 1. This 2D model had four free parameters (d_D , d_I , c_D , c_I) to account for responses in each condition.

First, we considered the Gaussian-independent-2D model. This saturated model necessarily fits the data perfectly. The best-fitting parameter values are shown in Table 2. An important question is whether the induction and deduction discriminability parameters are highly correlated across the 12 conditions (i.e., training block $\times$ believability $\times$ argument type). If so, then this model is mimicking a more parsimonious Gaussian-independent-1D model. Indeed, we found a high correlation between d_D and d_I , $r = .894$. This suggests that the extra discriminability parameter of the independent-2D model is redundant, and offers further evidence that the independent-1D model should be preferred for these data.

Second, we considered the Gaussian-independent-1D model. This model fits the data well; there was a very high correlation between the observed and predicted endorsement rates, $r = .997$, and the null hypothesis that the model fits the data could not be rejected, $G^2(9) = .65$, $p = 1.00$. The best-fitting parameter estimates, and diagrams of the criteria and argument

strength distributions are shown in Figures 7 and 8. They show that the Gaussian-independent-1D model parameters are psychologically interpretable, varying in sensible directions across the training and believability factors. For example, qualitatively, it can be seen that the training both increases discriminability and shifts the response thresholds to be more conservative, reducing participants' general bias to often endorse arguments. Also, for affirmation arguments, the greater post-training believability effects for induction participants compared to deduction participants (i.e., Figure 5) are related to a more conservative criterion setting for the induction group when responding to unbelievable items (see Figure 7e). In contrast, for denial arguments, induction and deduction groups had similar post-training response criteria for each level of argument believability, both adopting a stricter criterion for unbelievable arguments (see Figure 8e) than for neutral or believable arguments.

Table 2

The Independent-2D Model's Best-fitting Parameter Values for Experiment 2

Argument type	Block	Believability	d_I	d_D	c_I	c_D
Affirmation	Pre-training	Neutral	1.22	0.89	-0.11	-0.49
		Unbelievable	1.04	0.81	0.31	-0.02
		Believable	1.94	1.30	-0.05	-0.47
	Post-training	Neutral	1.91	2.81	0.54	1.03
		Unbelievable	1.30	2.27	0.95	1.06
		Believable	2.87	2.88	0.45	0.88
	Pre-training	Neutral	-0.02	0.00	-0.19	-0.63
		Unbelievable	-0.04	-0.06	0.21	-0.37
		Believable	0.06	-0.09	-0.35	-0.67
Denial	Post-training	Neutral	1.44	1.37	0.66	0.67
		Unbelievable	1.17	1.06	0.86	0.75
		Believable	1.04	0.93	0.29	0.39

Note. d_I and d_D are discriminability parameters and c_I and c_D are criterion parameters, for induction and deduction, respectively. The four standard deviation parameter values (i.e., for invalid-deduction, invalid-induction, valid-deduction and valid-induction) are fixed to 1 across conditions.

Figure 7. Illustration of the estimated parameter values for a Gaussian instantiation of the independent-1D model, for the affirmation arguments in Experiment 2. Aff = affirmation arguments; d = discriminability parameter; c_I = induction criterion parameter; c_D = deduction criterion parameter; IV = invalid; V = valid; D = deduction; I = induction. The standard deviation parameter values are fixed across conditions: invalid-deduction is set to 1; invalid-induction = 1.36; valid-deduction = 0.99; valid-induction = 0.84. See the online article for the color version of this figure. [2-column fitting image]

Figure 8. Illustration of the estimated parameter values for a Gaussian instantiation of the independent-1D model, for the denial arguments in Experiment 2. Den = denial arguments; d = discriminability parameter; c_I = induction criterion parameter; c_D = deduction criterion parameter; IV = invalid; V = valid; D = deduction; I = induction. The standard deviation parameter values are the same as for affirmation arguments. See the online article for the color version of this figure. [2-column fitting image]

5. General Discussion

Across two experiments, we performed a targeted test for evidence against a “single-process” independent-1D model, which had successfully accounted for a large body of existing data from the argument evaluation task (Stephens et al., 2018). Participants made induction or deduction judgments about valid and invalid arguments, which varied in argument believability and conditional argument form. Additionally, we used training to heighten Type 1 versus Type 2 processing, under a dual-process account. In Experiment 1 we used logic training to strengthen the influence of Type 2 processing in deduction. In Experiment 2 we added content or belief training to also strengthen the influence of Type 1 processing in induction. Combinations of the experimental factors showed differential effects on induction and deduction judgments, particularly in Experiment 2, which are often interpreted as evidence for dual-process accounts via dissociation logic (see e.g., Evans & Stanovich, 2013). However, we instantiated competing single- and dual-process accounts as formal signal detection models, and tested them in their most general form using signed difference analysis. The key result was that – in agreement with Stephens et al. (2018) – no compelling evidence was found against the independent-1D model. Indeed, this model could fit each of our two experiments perfectly, and it could not be rejected by combining both experiments into a larger dataset. This model assumes a common assessment of argument strength and distinct decision criteria for induction and deduction judgments. The success of this 1D model indicates that there is no need to invoke the dual-process core idea of distinct assessments of argument strength for induction and deduction, based on different weightings of Type 1 and Type 2 output.

We also showed that the independent-1D model has psychologically interpretable parameters that can help to unpack the effects of experimental manipulations. To do this, stronger distributional assumptions needed to be made. As an initial examination, we assumed Gaussian distributions, which is a standard assumption in the signal detection framework. The

Gaussian-independent-1D model suggested that, for instance, logic training both increased the discrimination of valid and invalid arguments, and reduced a general bias towards endorsing arguments. It also suggested that apparent dissociations between induction and deduction judgments may be driven by differences in response criteria. Additionally, an examination of the competing Gaussian-independent-2D model suggested that the two discriminability parameters were highly correlated, and thus superfluous.

Crucially, these parameter interpretations and assessments are contingent on the particular models and distributional assumptions. A major advantage of applying formal modeling to reasoning data is that the underlying assumptions are made explicit, and therefore can be scrutinized. Further theoretical development in this area will involve testing the effects of adding assumptions (e.g., about the form of the distributions linking latent parameters to observed responding) to our signal detection models. Our initial examination of a Gaussian-independent-1D model found that this model fit the current data very well. However, further tests of this model are needed using data based on other argument forms (e.g., categorical syllogisms, disjunctions) and other experimental manipulations. Other possible distributional assumptions could also be tested (e.g., gamma, exponential, equal-variance Gaussian). These more-specific variants of the independent-1D model could be assessed both in terms of relative fit, and in terms of parameter estimates across theoretically important experimental manipulations: Do decision thresholds and the discriminability parameter change in sensible and consistent ways? This is an important direction for future work. We note that Rotello and Heit (2009; Heit & Rotello, 2010) have previously rejected a Gaussian-1D signal detection model of induction and deduction judgments, although their data could be fit by the independent-1D model when tested using SDA (Stephens et al., 2018). This suggests that the Gaussian distributional form may not always be the most appropriate one for explaining reasoning data.

5.1. Convergence with Other Recent Evidence

We see our results as converging with other evidence that has challenged the classic dual-process view that belief-based judgments are typically based upon fast, automatic Type 1 processing, while logical reasoning depends upon slow, effortful Type 2 processing. For example, in a variant of the argument evaluation task, Handley et al. (2011) asked participants to make deduction judgments or to judge simply whether the argument conclusions were believable or unbelievable. They found that believability judgments were made *more slowly* and *less accurately* than deduction judgments. Similar effects have been found with base-rate reasoning tasks (e.g., Pennycook, Trippas, Thompson, & Handley, 2013).

These kinds of patterns have led to the proposal of more complex dual-process accounts that blur the classic distinction between Type 1 and 2 processing. For example, the parallel competitive model of Handley and Trippas (2015) assumes that structural problem features and background knowledge are activated simultaneously, and *both* are potentially reliant on Type 1 and Type 2 processing. Similarly, the logical intuition model of De Neys (2012) assumes that deliberate Type 2 processing may be triggered after initial intuitive Type 1 processing, but only if conflict is detected between *intuitively-based* logical and heuristic responses. Relatedly, Thompson, Pennycook, Trippas, and Evans (2018) recently found that high-capacity reasoners – even under a response deadline – were more accurate for deduction (or statistics) judgments than belief judgments when the two cues conflicted, whereas the reverse was true for low-capacity reasoners. This was interpreted as evidence that high-capacity reasoners have “better” intuitions that can include logical form or probabilities. However, instead of increasing the complexity of dual-process accounts and the similarities between Type 1 and 2 processing, evidence of “exceptions” to the predictions of more traditional dual-process theories may be consistent with a more parsimonious single-process account, in which the Type 1/2 distinction is unnecessary.

5.2. Further Theoretical Development

The success of the independent-ID model demonstrates the viability of a “single-process” account of reasoning. However, signal detection models might be regarded as *measurement models*; they specify the relationship between latent variables or parameters (discriminability and decision criteria) and observed outcomes (argument endorsement rates) but do not detail how the values of the parameters are set. Hence, there is still an important open question about *how* the value of the independent-ID model’s discriminability parameter is determined – what common mechanism is used to assess argument strength in deduction and induction tasks? One possibility is that it involves an assessment of the conditional probability of an argument’s conclusion based on the premises, combined with relevant background knowledge (Oaksford & Chater, 2007). Alternately, it could reflect an assessment of an argument’s general capacity for changing degrees of belief (i.e., the “force” of an argument; see Hahn & Oaksford, 2007), or some combination of both. For conditional arguments like those used in the current experiments (with the first premise as “If *antecedent* then *consequent*”), assessment of conditional probability may be moderated by the perceived relevance of the antecedent to the consequent (see Skovgaard-Olsen, Singmann, & Klauer, 2016; 2017). For example, unbelievable-valid arguments could be seen as negatively relevant (e.g., “If contraception is cheaper then there will be more pregnancies”). Thus, the “single dimension” of argument strength in the independent-ID model may be based on multiple types of information. Crucially, however, this model assumes that the same types of information underlie assessments of argument strength in induction and deduction.²

Another idea is that subjective argument strength may be based on an assessment of coherence or comprehensibility, driven by more general mechanisms for linguistic or discourse processing (cf. Krzyżanowska, Collins, & Hahn, 2017; Morsanyi & Handley, 2012; Trippas,

² The authors thank Mike Oaksford for suggesting this possibility.

Handley, Verde, & Morsanyi, 2016). The key idea is that in everyday communication, both logical form and content-believability contribute to the coherence of discourse. Thus, rather than being an extra skill outside of the ordinary, understanding or producing discourse regularly requires the ability to determine argument validity, not just believability. This is consistent with evidence that the so-called “primitive” Trobriand Islanders made arguments based on valid deductive inferences during disputes over land-rights (Hutchins, 1981). Along these lines, perhaps core components of the Singmann, Klauer, and Beller (2016) dual-source model could be applied to help connect these ideas with future extensions of the independent-1D model. According to the dual-source model, judgments in argument evaluation tasks are driven by a weighted integration of both logical form and argument content.

5.3. Alternative Tests of the Independent-1D Model

In the current studies we attempted to establish conditions that would provide a strong test of the independent-1D reasoning model. If dual-process accounts are correct then under these conditions, we should be able to see evidence of people using separate dimensions to evaluate argument strength in induction and deduction (leading to the signed difference vector forbidden by the independent-1D model). However, our signed difference analysis found no evidence of such results, leading us to retain the 1D model. Of course, proponents of dual-process accounts might be able to suggest other experimental designs that are more likely to demonstrate that two independent argument strength assessments (corresponding to Type 1 and Type 2 processing) are needed. One possibility is that alternative *independent* variables are needed to reveal the different processing types. That is, perhaps our independent variables of logic training, believability and conditional argument type did not differentially affect Type 1 and 2 processing. Therefore, future experiments could stick with induction and deduction judgments as the core dependent variables or tasks, but manipulate other combinations of independent variables, such as the emotional content of arguments, argument difficulty,

working memory, cognitive ability, or decision time. Note however, that Hayes et al. (2018) examined working memory load, working memory capacity, and decision time for induction and deduction judgments about simple arguments, and still found no evidence to support multiple processes based on state-trace analysis.

Another possibility is that alternative *dependent* variables or tasks (other than deduction vs. induction tasks) are needed to reveal two distinct kinds of processing. However, Stephens et al. (2019) used state-trace analysis to re-analyze studies cited by Evans and Stanovich (2013) as core experimental evidence for dual-process reasoning accounts. This included using alternative dependent variables such as reasoning judgments under time pressure versus no time pressure, or reasoning by individuals with high versus low working memory or cognitive ability. Much of the evidence cited by Evans and Stanovich (2013) did not support multiple processes according to state-trace analysis, although Stephens et al. (2019) identified a small number of studies deserving of further investigation. Additionally, Hayes, Wei, Dunn, and Stephens (2019) tested 1D and 2D signal detection models using SDA against an argument evaluation task, but this time comparing deduction judgments against – arguably more intuitive – *liking* judgments (i.e., participants simply rated how much they liked conclusions for valid and invalid arguments, see Morsanyi & Handley, 2012; Trippas et al., 2016). Hayes et al. also manipulated the critical independent variables of working memory capacity and working memory load while participant made the liking/deduction judgments, but again found no evidence to reject the independent-1D model.

5.4. Training Deductive Reasoning

Beyond testing the competing single- and dual-process accounts, our experiments also show that people can be trained in just a few minutes to quite accurately assess logical validity. This training can be automated (e.g., unlike the face-to-face training of Prowse Turner & Thompson, 2009), although its success may depend upon problem difficulty (cf. Klauer et al.,

1997). This is promising for educational applications aimed at improving people's reasoning – especially given some previous work suggesting that even an entire first-year philosophy logic course improved deductive reasoning only moderately (Leighton, 2006) or barely at all (Cheng et al., 1986). Important avenues for future research include isolating the components of training that are most beneficial and/or efficient (see e.g., Klauer et al., 2000), and whether training generalizes to other, non-trained, argument forms. Our training procedure may be useful for future theoretical lab work. As Evans (2010, p.323) argued, "...we will learn more about rationality by studying people's ability to apply rule-based reasoning that they have been taught, rather than by constantly studying the performance of naïve participants on novel problems."

5.5. Implications for Other Domains and Tasks

Our findings have broad implications across many areas of psychology. Dual-process or dual-systems theories are becoming increasingly popular, according to citation rates (Pennycook, 2018). Examples include declarative versus non-declarative memory and learning (e.g., McLaren et al., 2014; Schacter & Tulving, 1994; Squire, 1992), explicit rule-based versus procedural category learning (Ashby, Alfonso-Reese, Turken, & Waldron, 1998), holistic and featural processing of faces (Tanaka & Gordon, 2011), visuo-spatial versus phonological working memory (Baddeley, 2012), automatic versus explicit processing of social information (Evans, 2008; Smith & DeCoster, 2000), and distinct modes of thinking for utilitarian versus deontological moral judgments (Paxton & Greene, 2010). However, important evidence for such accounts is often based on functional dissociations. Moving beyond this kind of evidence, we have added to the growing body of claims that single-process accounts have been prematurely discounted (e.g., in recognition memory: Dunn, 2008; Hayes, Dunn, Joubert, & Taylor, 2017; face perception: Loftus, Oberg, & Dillon, 2004; category learning: Newell, Dunn, & Kalish, 2010; Stephens & Kalish, 2018). We have illustrated a useful approach for

formally instantiating the competing theories, identifying critical ordinal data patterns that can distinguish them, and then performing targeted experiments to see whether single-process models really can be rejected.

5.6. Conclusions

Dual-process theories of high-level cognition have been very influential, and recent work has focused on refining them in light of data that increasingly challenges a clear distinction between Type 1 and 2 processing (e.g., De Neys, 2012; Handley & Trippas, 2015; Thompson et al., 2018). Now is thus an opportune moment to pause and reconsider whether such a distinction need be maintained. The current studies illustrate how formal modeling can be used to guide and perform targeted tests of competing single- and dual-process accounts of reasoning. One-dimensional signal detection models can account for a surprisingly wide range of data, supporting the view that a single-process account is a viable alternative to (increasingly complex) dual-process accounts. Thus, it seems that – as they themselves come to appreciate – Kirk and Spock are not as sharply different as they first appear.

Acknowledgements

Funding: This work was supported by Australian Research Council Discovery Grant DP150101094 and DP190102160 to authors BKH and JCD, and Australian Research Council Discovery Grant DP130101535 to JCD and MLK. Declarations of interest: none. The authors thank Eric Moskowitz, Rebecca Leonard and Tennazha Bradley for their assistance with data collection, and Nicole Cruz for assistance with statistical analyses. Materials and data are available as supplementary files, and at <https://osf.io/cwf8x/>.

References

- 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
- Baddeley, A. (2012). Working memory: Theories, models, and controversies. *Annual Review of Psychology*, 63, 1-29. doi:10.1146/annurev-psych-120710-100422
- 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
- Bates, D., Maechler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67, 1-48. doi:10.18637/jss.v067.i01
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433-436.
- Cheng, P. W., Holyoak, K. J., Nisbett, R. E., & Oliver, L. M. (1986). Pragmatic versus syntactic approaches to training deductive reasoning. *Cognitive Psychology*, 18(3), 293-328. doi:10.1016/0010-0285(86)90002-2
- Croskerry, P., Singhal, G., & Mamede, S. (2013). Cognitive debiasing 1: Origins of bias and theory of debiasing. *British Medical Journal: Quality and Safety*, 22, ii58-ii64. doi:10.1136/bmjqs-2012-001712
- Dane, E., & Pratt, M. G. (2007). Exploring intuition and its role in managerial decision making. *Academy of Management Review*, 32(1), 33-54. doi:10.5465/amr.2007.23463682
- De Neys, W. (2012). Bias and conflict: A case for logical intuitions. *Perspectives on Psychological Science*, 7(1), 28-38. doi:10.1177/1745691611429354

- Dube, C., Rotello, C. M., & Heit, E. (2010). Assessing the belief bias effect with ROCs: It's a response bias effect. *Psychological Review*, 117(3), 831-863. doi:10.1037/a0019634
- 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., & Anderson, L. (2018). Signed difference analysis: Testing for structure under monotonicity. *Journal of Mathematical Psychology*, 85, 36-54.
doi:10.1016/j.jmp.2018.07.002
- Dunn, J. C., & James, R. N. (2003). Signed difference analysis: Theory and application. *Journal of Mathematical Psychology*, 47(4), 389-416. doi:10.1016/S0022-2496(03)00049-X
- Dunn, J. C., & Kalish, M. L. (2018). *State-trace analysis*: Springer.
- Dunn, J. C., & Kirsner, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review*, 95(1), 91-101.
doi:10.1037/0033-295X.95.1.91
- Evans, J. St. B. T. (2007). On the resolution of conflict in dual process theories of reasoning. *Thinking & Reasoning*, 13(4), 321-339. doi:10.1080/13546780601008825
- Evans, J. St. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255-278.
doi:10.1146/annurev.psych.59.103006.093629
- Evans, J. St. B. T. (2010). Intuition and reasoning: A dual-process perspective. *Psychological Inquiry*, 21, 313-326. doi:10.1080/1047840X.2010.521057
- Evans, J. St. B. T., Barston, J. L., & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory & Cognition*, 11, 295-306.
doi:10.3758/BF03196976

- Evans, J. St. B. T., Handley, S. J., & Bacon, A. M. (2009). Reasoning under time pressure: A study of causal conditional inference. *Experimental Psychology*, 56(2), 77-83.
doi:10.1027/1618-3169.56.2.77
- Evans, J. St. B. T., Handley, S. J., Harper, C. N. J., & Johnson-Laird, P. N. (1999). Reasoning about necessity and possibility: A test of the mental model theory of deduction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(6), 1495-1513. doi:10.1037/0278-7393.25.6.1495
- Evans, J. St. B. T., Handley, S. J., Neilens, H., & Over, D. E. (2010). The influence of cognitive ability and instructional set on causal conditional inference. *The Quarterly Journal of Experimental Psychology*, 63(5), 892-909.
doi:10.1080/17470210903111821
- Evans, J. St. B. T., & Over, D. E. (2013). Reasoning to and from belief: Deduction and induction are still distinct. *Thinking & Reasoning*, 19(3-4), 267-283.
doi:10.1080/13546783.2012.745450
- Evans, J. St. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, 8(3), 223-241.
doi:10.1177/1745691612460685
- Gillard, E., Van Dooren, W., Schaeken, W., & Verschaffel, L. (2009). Dual processes in the psychology of mathematics education and cognitive psychology. *Human Development*, 52, 95-108. doi:10.1159/000202728
- Hahn, U., & Oaksford, M. (2007). The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological Review*, 114, 704-732.
doi:10.1037/0033-295X.114.3.704

- Handley, S. J., Newstead, S. E., & Trippas, D. (2011). Logic, beliefs, and instruction: A test of the default interventionist account of belief bias. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(1), 28-43. doi:10.1037/a0021098
- Handley, S. J., & Trippas, D. (2015). Dual processes and the interplay between knowledge and structure: A new parallel processing model. *Psychology of Learning and Motivation*, 62, 33-58. doi:10.1016/bs.plm.2014.09.002
- Hayes, B. K., Dunn, J. C., Joubert, A., & Taylor, R. (2017). Comparing single- and dual-process models of memory development. *Developmental Science*, 20, e12469. doi:10.1111/desc.12469
- Hayes, B. K., & Heit, E. (2017). Inductive reasoning 2.0. *Wiley Interdisciplinary Reviews: Cognitive Science*, 9(3), e1459. doi:10.1002/wcs.1459
- Hayes, B. K., Stephens, R. G., Ngo, J., & Dunn, J. C. (2018). The dimensionality of reasoning: Inductive and deductive inference can be explained by a single process. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 44(9), 1333-1351. doi:10.1037/xlm0000527
- Hayes, B. K., Wei, P., Dunn, J. C., & Stephens, R. G. (2019). Why is logic so likeable? A single-process account of argument evaluation with logic and liking judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Advance online publication. doi:10.1037/xlm0000753
- Heit, E., & Rotello, C. M. (2010). Relations between inductive reasoning and deductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(3), 805-812. doi:10.1037/a0018784
- Heit, E., Rotello, C. M., & Hayes, B. K. (2012). Relations between memory and reasoning. *Psychology of Learning and Motivation*, 57, 57-101. doi:10.1016/B978-0-12-394293-7.00002-9

- Howarth, S., Handley, S. J., & Walsh, C. (2016). The logic-bias effect: The role of effortful processing in the resolution of belief–logic conflict. *Memory & Cognition*, 44(2), 330-349. doi:10.3758/s13421-015-0555-x
- Hutchins, E. (1981). Reasoning in Trobriand discourse. In R. Casson (Ed.), *Language, culture, and cognition: Anthropological perspectives* (pp. 481-489). New York: MacMillan.
- Johnson-Laird, P. N. (1994). Mental models and probabilistic thinking. *Cognition*, 50(1-3), 189-209. doi:10.1016/0010-0277(94)90028-0
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49-81). Cambridge: Cambridge University Press.
- Kalish, M. L., Dunn, J. C., Burdakov, O. P., & Sysoev, O. (2016). A statistical test of the equality of latent orders. *Journal of Mathematical Psychology*, 70, 1-11. doi:10.1016/j.jmp.2015.10.004
- Keren, G. (2013). A tale of two systems: A scientific advance or a theoretical stone soup? Commentary on Evans & Stanovich (2013). *Perspectives on Psychological Science*, 8(3), 257-262. doi:10.1177/1745691613483474
- Keren, G., & Schul, Y. (2009). Two is not always better than one: A critical evaluation of two-system theories. *Perspectives on Psychological Science*, 4(6), 533-550. doi:10.1111/j.1745-6924.2009.01164.x
- Klaczynski, P. A., & Laipple, J. S. (1993). Role of content domain, logic training, and IQ in rule acquisition and transfer. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 19(3), 653-672. doi:10.1037/0278-7393.19.3.653

- Klauer, K. C., Meiser, T., & Naumer, B. (2000). Training propositional reasoning. *The Quarterly Journal of Experimental Psychology Section A*, 53(3), 868-895.
doi:10.1080/713755911
- Klauer, K. C., Stegmaier, R., & Meiser, T. (1997). Working memory involvement in propositional and spatial reasoning. *Thinking & Reasoning*, 3(1), 9-47.
doi:10.1080/135467897394419
- Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? *European Conference on Visual Perception - Abstract Supplement*.
- Kruglanski, A. W. (2013). Only one? The default interventionist perspective as a unimodel—Commentary on Evans & Stanovich (2013). *Perspectives on Psychological Science*, 8(3), 242-247. doi:10.1177/1745691613483477
- Kruglanski, A. W., & Gigerenzer, G. (2011). Intuitive and deliberate judgments are based on common principles. *Psychological Review*, 118(1), 97-109. doi:10.1037/a0020762
- Krzyżanowska, K., Collins, P. J., & Hahn, U. (2017). Between a conditional's antecedent and its consequent: Discourse coherence vs. probabilistic relevance. *Cognition*, 164, 199-205. doi:10.1016/j.cognition.2017.03.009
- Lassiter, D., & Goodman, N. D. (2015). How many kinds of reasoning? Inference, probability, and natural language semantics. *Cognition*, 136, 123-134.
doi:10.1016/j.cognition.2014.10.016
- Leighton, J. P. (2006). Teaching and assessing deductive reasoning skills. *The Journal of Experimental Education*, 74(2), 109-136. doi:10.3200/JEXE.74.2.107-136
- Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, 6(3), 312-319.
doi:10.3758/BF03197461

- 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
- McLaren, I. P. L., Forrest, C. L. D., McLaren, R. P., Jones, F. W., Aitken, M. R. F., & Mackintosh, N. J. (2014). Associations and propositions: The case for a dual-process account of learning in humans. *Neurobiology of Learning and Memory*, *108*, 185-195. doi:10.1016/j.nlm.2013.09.014
- Melnikoff, D. E., & Bargh, J. A. (2018). The mythical number two. *Trends in Cognitive Sciences*, *22*(4), 280-293. doi:10.1016/j.tics.2018.02.001
- Morsanyi, K., & Handley, S. J. (2012). Logic feels so good - I like it! Evidence for intuitive detection of logicity in syllogistic reasoning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *38*(3), 596-616. doi:10.1037/a0026099
- 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. L. (2010). The dimensionality of perceptual category learning: A state-trace analysis. *Memory & Cognition*, *38*, 563-581. doi:10.3758/MC.38.5.563
- Newstead, S. E., Pollard, P., Evans, J. St. B. T., & Allen, J. L. (1992). The source of belief bias effects in syllogistic reasoning. *Cognition*, *45*, 257-284. doi:10.1016/0010-0277(92)90019-E
- Oaksford, M., & Chater, N. (2001). The probabilistic approach to human reasoning. *Trends in Cognitive Sciences*, *5*(8), 349-357. doi:10.1016/S1364-6613(00)01699-5
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford, England: Oxford University Press.

- Osman, M. (2004). An evaluation of dual-process theories of reasoning. *Psychonomic Bulletin & Review*, 11(6), 988-1010. doi:10.3758/bf03196730
- Osman, M. (2013). A case study: Dual-process theories of higher cognition—Commentary on Evans & Stanovich (2013). *Perspectives on Psychological Science*, 8(3), 248-252. doi:10.1177/1745691613483475
- Paxton, J. M., & Greene, J. D. (2010). Moral reasoning: Hints and allegations. *Topics in Cognitive Science*, 2, 511-527. doi:10.1111/j.1756-8765.2010.01096.x
- Pennycook, G. (2018). A perspective on the theoretical foundation of dual process models. In W. De Neys (Ed.), *Dual Process Theory 2.0*. New York, NY: Psychology Press.
- Pennycook, G., Trippas, D., Thompson, V. A., & Handley, S. J. (2013). Base rates: both neglected and intuitive. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 40, 544-554. doi:10.1037/a0034887
- Prince, M., Brown, S., & Heathcote, A. (2012). The design and analysis of state-trace experiments. *Psychological Methods*, 17, 78-99. doi:10.1037/a0025809
- Prowse Turner, J. A., & Thompson, V. A. (2009). The role of training, alternative models, and logical necessity in determining confidence in syllogistic reasoning. *Thinking & Reasoning*, 15(1), 69-100. doi:10.1080/13546780802619248
- Rips, L. J. (2001). Two kinds of reasoning. *Psychological Science*, 12(2), 129-134. doi:10.1111/1467-9280.00322
- Rotello, C. M., & Heit, E. (2009). Modeling the effects of argument length and validity on inductive and deductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(5), 1317-1330. doi:10.1037/a0016648
- Rotello, C. M., Heit, E., & Kelly, L. J. (2019). Do modals identify better models? A comparison of signal detection and probabilistic models of inductive reasoning. *Cognitive Psychology*, 112, 1-24.

- Rouder, J. N., Pratte, M. S., & Morey, R. D. (2010). Latent mnemonic strengths are latent: a comment on Mickes, Wixted, and Wais (2007). *Psychonomic Bulletin & Review*, 17(3), 427-435. doi:10.3758/pbr.17.3.427
- Schacter, D. L., & Tulving, E. (1994). *Memory systems 1994*. Cambridge, MA: MIT Press.
- Singmann, H., & Klauer, K. C. (2011). Deductive and inductive conditional inferences: Two modes of reasoning. *Thinking & Reasoning*, 17(3), 247-281. doi:10.1080/13546783.2011.572718
- Singmann, H., Klauer, K. C., & Beller, S. (2016). Probabilistic conditional reasoning: Disentangling form and content with the dual-source model. *Cognitive Psychology*, 88, 61-87. doi:10.1016/j.cogpsych.2016.06.005
- Skovgaard-Olsen, N., Singmann, H., & Klauer, K. C. (2016). The relevance effect and conditionals. *150*, 26-36. doi:10.1016/j.cognition.2015.12.017
- Skovgaard-Olsen, N., Singmann, H., & Klauer, K. C. (2017). Relevance and reason relations. *Cognitive Science*, 41, 1202-1215. doi:10.1111/cogs.12462
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119(1), 3-22. doi:10.1037/0033-2909.119.1.3
- Sloman, S. A. (2014). Two systems of reasoning, an update. In J. W. Sherman, B. Gawronski, & Y. Trope (Eds.), *Dual process theories of the social mind*. New York, NY: Guilford Press.
- Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, 4(2), 108-131. doi:10.1207/S15327957PSPR0402_01
- Squire, L. R. (1992). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. *Journal of Cognitive Neuroscience*, 4, 232-243. doi:10.1162/jocn.1992.4.3.232

- Stanovich, K. E. (2016). The comprehensive assessment of rational thinking. *Educational Psychologist, 51*(1), 23-34. doi:10.1080/00461520.2015.1125787
- Stephens, R. G., Dunn, J. C., & Hayes, B. K. (2018). Are there two processes in reasoning? The dimensionality of inductive and deductive inferences. *Psychological Review, 125*(2), 218-244. doi:10.1037/rev0000088
- Stephens, R. G., & Kalish, M. L. (2018). The effect of feedback delay on perceptual category learning and item memory: Further limits of multiple systems. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 44*(9), 1397-1413. doi:10.1037/xlm0000528
- Stephens, R. G., Matzke, D., & Hayes, B. K. (2019). Disappearing dissociations in experimental psychology: Using state-trace analysis to test for multiple processes. *Journal of Mathematical Psychology, 90*, 3-22. doi:10.1016/j.jmp.2018.11.003
- Tanaka, J. W., & Gordon, I. (2011). Features, configuration, and holistic face processing. In A. J. Calder, G. Rhodes, M. H. Johnson, & J. V. Haxby (Eds.), *The Oxford handbook of face perception* (pp. 177-194). New York, NY: Oxford University Press.
- Thompson, V. A., Pennycook, G., Trippas, D., & Evans, J. St. B. T. (2018). Do Smart People Have Better Intuitions? *Journal of Experimental Psychology: General, 147*(7), 945-961. doi:10.1037/xge0000457
- 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, & Cognition, 39*, 1393-1402. doi:10.1037/a0032398
- Trippas, D., Handley, S. J., Verde, M. F., & Morsanyi, K. (2016). Logic brightens my day: Evidence for implicit sensitivity to logical validity. *Journal of Experimental*

Psychology: Learning, Memory, & Cognition, 42(9), 1448-1457.

doi:10.1037/xlm0000248

Trippas, D., Verde, M. F., Handley, S. J., Roser, M. E., McNair, N. A., & Evans, J. St. B. T.

(2014). Modeling causal conditional reasoning data using SDT: Caveats and new insights. *Frontiers in Psychology*, 5, 217. doi:10.3389/fpsyg.2014.00217

Verschueren, N., Schaeken, W., & d'Ydewalle, G. (2005). A dual-process specification of

causal conditional reasoning. *Thinking & Reasoning*, 11(3), 239-278.

doi:10.1080/13546780442000178

Appendix A

Stimuli For Both Experiments

Table A1

Believable Conditional Statements and Certainty Ratings

Conditional Statement	<i>M</i>	<i>SD</i>	Set
If the quarterback wins the Heisman Trophy then the quarterback will receive more endorsements	82.0	15.1	1
If people recycle more of their trash then the amount of waste sent to landfills will decrease	81.9	17.9	2
If car ownership increases then traffic congestion will get worse	81.8	17.3	2
If a Beyonce song is used as a movie theme song then sales of the song will increase	80.1	12.9	1
If a company advertises during the Super Bowl then the company's sales will increase	78.8	17.8	1
If Adidas get more superstars to wear their shoes then Adidas sales will increase	78.4	21.3	2
If Sony release a PlayStation 5 then Sony's profits will rise	77.8	17.3	2
If there are more student scholarships then university entries will increase	77.6	21.7	1
If ozone depletion continues then the Arctic ice will melt	76.9	23.0	1
If more people use protective sun cream then cases of skin cancer will be reduced	75.1	18.8	2
If people buy from local businesses more often then there will be more local jobs	74.0	15.9	2
If money is invested in downtown public spaces then there will be more people socializing downtown	73.7	17.6	1
If it becomes compulsory to vote during presidential elections then voter turnout rates will improve	72.5	20.7	1
If Kim Kardashian wears new Gucci clothes then Gucci's sales will increase	72.0	23.2	2
If genetic research continues then a cure for cancer will be found	70.1	14.4	2
If oil prices continue to rise then US petrol prices will rise	69.9	17.8	1
If fertility treatment improves then the world population will rise	69.8	16.9	1
If the US cuts fuel emissions then global warming will be reduced	69.7	25.0	2
If student fees are increased then applications for university places will drop	69.3	20.8	2
If the cost of fruit and vegetables is subsidized then people will eat more healthily	69.3	20.8	1
If elementary school class sizes are reduced then national literacy will improve	69.0	23.2	1
If jungle deforestation continues then Gorillas will become extinct	68.1	23.4	2
If nurses' salaries are improved the recruitment of nurses will increase	67.8	20.9	2
If funding for public broadcasting is lowered then program quality will be reduced	67.5	21.4	1

Note. "Set" indicates the allocation into stimulus set 1 or 2, for yoking across the pre- and post-training blocks.

Table A2

Unbelievable Conditional Statements and Certainty Ratings

Conditional Statement	<i>M</i>	<i>SD</i>	Set
If Allied forces leave Iraq then it will become a democracy	32.8	17.3	1
If the minimum wage is increased then product prices will fall	32.0	21.1	2
If more new houses are built then there will be more homeless people	31.6	19.0	2
If there is less traffic then more children will walk to school	30.9	23.2	1
If the legal drinking age is reduced then there will be a lower incidence of road traffic accidents	30.7	21.1	1
If the US relaxes its immigration policy then there will be more jobs for Americans	29.1	21.7	2
If high schools become more academically rigorous then fewer people will apply to university	28.4	22.3	2
If state-funded preschools are made available then fewer couples will have children	27.2	14.9	1
If the food stamp program is abolished then low-income families will have more food	25.0	26.2	1
If the US restricts imports from China then clothing prices will fall	24.9	19.8	2
If fast food is taxed then childhood obesity will increase	24.0	17.6	2
If quarantine laws are strengthened then mad cow disease will spread to the US	23.1	17.4	1
If people are educated about cruelty to farm animals then there will be fewer vegetarians	23.1	24.5	1
If health insurance is cheaper then there will be more sick people	23.0	27.5	2
If further terrorist attacks occur in the US then President Obama will resign	22.0	19.2	2
If science funding is reduced then there will be more technological advances in the US	20.7	22.7	1
If another airplane disappears during flight then the number of airplane passengers will increase	20.6	18.9	1
If more music is downloaded illegally then the music industry will make more money	20.3	23.8	2
If dogs are trained better then there will be more cases of dog bites	19.5	21.8	2
If contraception is cheaper then there will be more pregnancies	19.4	20.4	1
If more green cards are granted then the US population will decrease	19.3	20.4	1
If customers refuse to tip then service staff will earn a lot of money	18.8	18.8	2
If students lack internet access then it is easier to complete assignments	18.6	24.9	2
If Skype begins charging a fee then there will be more Skype users	17.6	23.0	1

Table A3

Neutral Conditional Statements

Conditional Statement	Set
If the spy novel is horb then the critics will be impressed	1
If the workers are vust then the business will succeed	1
If the video game is mord then the game will be more fun	1
If the wine includes rool then the wine will taste very fruity	1
If the cake is peze then the cake will be gluten-free	1
If the coffee has sede then the coffee will be bitter	1
If the party is bete then the party will be dull	1
If the animal eats thoe then the animal will be fitter	1
If the chairs are dith then the chairs will be more comfortable	1
If the detective is quemp then the offender will be found sooner	1
If the handbag has a brumpt then the handbag will have a patterned lining	1
If the hotel is fring then the hotel rooms will be clean	1
If an insect eats clemp then the insect will become darker	2
If artwork is glat then the art will be more likely to be shown in a gallery	2
If the project is glab then the project will finish within-budget	2
If the tree is zash then the tree will be healthy	2
If the basketball team is maft then the team will score many goals	2
If a film is skad then the film will be a blockbuster	2
If the business is fasp then the office will shift location	2
If the surgeon is scomp then the patient will recover quickly	2
If the restaurant is kend then the food will be excellent	2
If the Mayor is lesp then the city will be well run	2
If the actor is cran then the role will be well performed	2
If the new song is floke then the band will top the music charts	2

Appendix B

Signal Detection Model Fitting Results for Experiment 1

The saturated Gaussian-independent-2D model necessarily fits the Experiment 1 data perfectly. The best-fitting parameter values are shown in Table B1. As in Experiment 2, we found a high correlation between d_D and d_I , $r = .959$, suggesting that the extra discriminability parameter is redundant. This further supports that the independent-1D model should be preferred for these data.

The Gaussian-independent-1D model fits the Experiment 1 data well. There was a very high correlation between the observed and predicted endorsement rates, $r = .998$, and the null hypothesis that the model fits the data could not be rejected, $G^2(9) = -4.08$, $p = 1.00$. The best-fitting parameter estimates are illustrated in Figures B1 and B2. Again, the Gaussian-independent-1D model parameters are interpretable. For example, the logic training led to both increased discriminability and more conservative response thresholds. The similar response ratings across induction and deduction groups (i.e., Figure 3) are reflected by similar criterion settings for each group.

Table B1

The independent-2D model's best-fitting parameter values for Experiment 1.

Argument type	Block	Believability	d_I	d_D	c_I	c_D
Affirmation	Pre-training	Neutral	1.27	0.79	-0.09	-0.51
		Unbelievable	0.76	0.60	0.43	0.21
		Believable	1.31	0.97	-0.26	-0.41
	Post-training	Neutral	2.18	2.15	0.66	0.76
		Unbelievable	1.95	2.07	0.76	0.89
		Believable	2.12	2.40	0.66	0.86
Denial	Pre-training	Neutral	-0.32	0.08	-0.43	-0.43
		Unbelievable	-0.35	-0.12	-0.04	-0.07
		Believable	-0.60	-0.36	-0.62	-0.81
	Post-training	Neutral	1.25	1.43	0.78	0.84
		Unbelievable	1.24	1.01	0.83	0.67
		Believable	1.24	1.33	0.58	0.70

Note. d_I and d_D are discriminability parameters and c_I and c_D are criterion parameters, for induction and deduction, respectively. The four standard deviation parameter values are fixed to 1 across conditions.

Figure B1. Illustration of the estimated parameter values for the Gaussian-independent-1D model, for the affirmation arguments in Experiment 1. Aff = affirmation arguments; d = discriminability parameter; c_I = induction criterion parameter; c_D = deduction criterion parameter; IV = invalid; V = valid; D = deduction; I = induction. The standard deviation parameter values are fixed across conditions; invalid-deduction is set to 1, and the others were estimated to be as follows: invalid-induction = 1.05; valid-deduction = 0.58; valid-induction = 0.51. See the online article for the color version of this figure. [2-column fitting image]

Figure B2. Illustration of the estimated parameter values for the Gaussian-independent-1D model, for the denial arguments in Experiment 1. Den = denial arguments; d = discriminability parameter; c_I = induction criterion parameter; c_D = deduction criterion parameter; IV = invalid; V = valid; D = deduction; I = induction. The standard deviation parameter values are the same as for affirmation arguments. See the online article for the color version of this figure.

[2-column fitting image]