Not All Syllogisms Are Created Equal: Varying Premise Believability Reveals Differences Between Conditional and Categorical Syllogisms

by

Stephanie Solcz

A thesis

presented to the University of Waterloo

in fulfilment of the

thesis requirement for the degree of

Doctor of Philosophy

in

Psychology

Waterloo, Ontario, Canada, 2011

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Deductive reasoning is a fundamental cognitive skill, and consequently has been the focus of much research over the past several decades. In the realm of syllogistic reasoning—judging the validity of a conclusion given two premises—a robust finding is the *belief bias effect*: broadly, the tendency for reasoners to judge as valid more believable than unbelievable conclusions. How the content believability of conclusions influences syllogistic reasoning has been the subject of hundreds of experiments and has informed several theories of deductive reasoning; however, how the content of *premises* influences the reasoning processes has been largely overlooked. In this thesis, I present 5 experiments that examine how premise content influences reasoning about *categorical* (i.e., statements with the words 'some' and 'not') and *conditional* (i.e., 'if/then' statements) syllogisms, which tend to be treated as interchangeable in deductive reasoning literature. It is demonstrated that premise content influences reasoning in these two types of syllogisms in fundamentally different ways. Specifically, Experiment 1 replicates and extends previous findings and demonstrates that for *conditional* syllogisms, belief bias results when premises are both believable and unbelievable; however, reasoners are more likely to judge that a conclusion is valid when it follows from believable than from unbelievable premises. Conversely, belief bias for *categorical syllogisms* results only when premises are believable; conclusion believability does not influence conclusion endorsement when premises are unbelievable.

Based on these preliminary findings, I propose a theory that categorical and conditional syllogisms differ in the extent to which reasoners initially assume the premises to be true, and that this difference influences when in the reasoning process reasoners evaluate the believability of premises. Specifically, I propose that reasoners automatically assume that conditional, but not

categorical, premises are true. It is proposed that, because the word “if” in conditional statements elicits hypothetical thinking, conditional premises are assumed to be true for the duration of the reasoning process. Subsequent to reasoning, premises can be “disbelieved” in a time-consuming process, and initial judgments about the conclusion may be altered, with a bias to respond that conclusions following from believable premises are valid. On the other hand, because categorical premises are phrased as factual propositions, reasoners initially judge the believability of categorical premises prior to reasoning about the conclusion. Unbelievable premises trigger the reasoner to disregard content from the rest of the syllogism, perhaps because the reasoner believes that the information in the problem will not be helpful in solving the problem.

This theory is tested and supported by four additional experiments. Experiment 2 demonstrates that reasoners take longer to reason about conditional syllogisms with unbelievable than believable premises, consistent with the theory that unbelievable premises are “disbelieved” in a time-consuming process. Further, participants demonstrate belief bias for categorical syllogisms with unbelievable premises when they are instructed to assume that premises are true (Experiment 3) or when the word ‘if’ precedes the categorical premises (Experiment 4). Finally, Experiment 5 uses eye-tracking to demonstrate that premise believability influences *post-conclusion* premise looking durations for conditional syllogisms and *pre-conclusion* premise looking durations for categorical syllogisms. This finding supports the hypothesis that reasoners evaluate the believability of conditional premises *after* reasoning about the conclusion but that they evaluate the believability of categorical premises *before* reasoning about the conclusion. Further, Experiment 5 reveals that participants have poorer memory for the content of categorical syllogisms with unbelievable than believable premises, but memory did not differ for conditional syllogisms with believable and unbelievable premises. This suggests that unbelievable premise content in categorical syllogism is suppressed or ignored.

These results and the theory of premise evaluation that I propose are discussed in the context of contemporary theories of deductive reasoning.

Acknowledgments

I would like to thank, first and foremost, Jonathan Fugelsang, my supervisor and friend. His endless academic and personal support was greatly appreciated throughout the ups and downs of the past five years.

Second, thank you to committee members Derek Koehler and Colin MacLeod for their helpful suggestions on this thesis, and to members of the Reasoning and Decision Making Lab, for their academic support and friendship.

Finally, thank you to the many family members and friends who helped to make the past five years incredibly fun and as stress-free as possible.

Table of Contents

List of Figures

Deductive reasoning is a fundamental cognitive ability that has been considered to be one of the cornerstones of logical thought since the time of Aristotle. According to Piaget, demonstrations of deductive reasoning signalled the development of logical thinking abilities and the attainment of formal operations, the final stage of cognitive development (Inhelder & Piaget, 1958). Certainly, the ability to reason deductively is necessary in our daily lives to accurately derive conclusions from sets of information. For example, if a doctor prescribes you medication and tells you not to take aspirin on days that you take the medication, it is important to accurately reason that if you take the medication today, you cannot take aspirin today. Now, imagine the case that you have a headache, for which you typically take aspirin, on a day that you need to take your newly prescribed medication. To avoid dangerous drug interactions, it is imperative to be able to ignore your belief that aspirin is acceptable to take when you have a headache. This essential component of reasoning – reasoning only with relevant information and disregarding beliefs – often proves difficult.

In the laboratory, deductive reasoning is typically studied by presenting participants with syllogisms with two premises and a conclusion. Participants are instructed to determine whether the conclusion necessarily follows from the information contained in the premises, or whether the conclusion is valid or invalid given the premises. Syllogisms may be categorical or conditional in nature. Categorical syllogisms contain two premises and a conclusion that represent relations among different classes of entities. In abstract terms, a typical categorical syllogism can be represented as below:

No A are B

Some C are B

Therefore, some C are not A (Valid)

or

Therefore, some A are not C (Invalid)

Syllogisms may also contain conditional premises and conclusions, which depict hypothetical relations in "If/Then" statements. Conditional syllogisms always have a major premise that is conditional in nature, and a minor premise which may or may not be another conditional. An abstract conditional syllogism with two conditional premises (also known as a hypothetical syllogism, sorite, or syllogism with double conditionals) is represented below:

If A, then B

If B, then C

Therefore, If A, then C (Valid)

or

Therefore, If C, then A (Invalid)

Most studies of reasoning use syllogisms that are categorical in nature or they use problems with a major conditional premise and a minor categorical premise (e.g., *If p, then q; p; therefore*, *q*). Conditional problems of this nature lend themselves to distinct theories of reasoning which will not be further addressed here. For the purposes of this dissertation, conditional reasoning will refer to syllogisms with two conditional premises, unless otherwise noted, so that direct comparisons may be made between categorical and conditional syllogisms.

The Belief Bias Effect

For both categorical and conditional syllogisms, accurate deductive reasoning requires that people evaluate the conclusion using *only* logical relations contained in the premises. However, the tendency for people to take into account their prior knowledge and beliefs results in a common fallacy known as *belief bias* (Evans, Barston, & Pollard, 1983). This extremely robust phenomenon indicates that reasoners are more likely to endorse a conclusion as valid when it is also believable, or consistent with their knowledge about the world. Typically, the effects of conclusion believability are stronger on invalid than on valid conclusions.

The belief bias effect is one of the most widely replicated findings in the deductive reasoning literature, and points strongly to the influence of conclusion content on the ability to reason logically. The effect is found for both categorical (Evans et al., 1983; Klauer, Munch, & Naumer, 2000) and conditional syllogisms (Santamaria, Gargia-Madruga, & Johnson-Laird, 1998; Torrens, Thompson, & Cramer, 1999; Thompson, 1996; Evans & Over, 2004). Most modern explanations of belief bias can account for the effect in both categorical and conditional syllogisms and therefore do not necessarily distinguish between mechanisms underlying the two types of syllogisms. A brief review of two major theories of belief bias demonstrates this.

Theories of Belief Bias

First, the mental-models account of deductive reasoning, pioneered by Johnson-Laird (Johnson-Laird & Byrne, 1991; Johnson-Laird & Bara, 1984), proposes that reasoners begin by constructing models of information contained in the premises. Mental models can be thought of as iconic, diagram-like spatial representations of premises (Johnson-Laird, 2001; Knauff, Mulack, Kassubek, Salih, & Greenlee, 2002). According to the Mental Models account, deductive reasoning occurs in three stages (Johnson-Laird & Byrne, 1991): First, reasoners use their general knowledge to construct a model of the information that the premises describe; second, reasoners attempt to use this model to arrive at a conclusion about the information in the premises; and third, reasoners attempt to search for alternative models that falsify this conclusion. The number of models possible given a set of premises varies, and a conclusion is valid if it is consistent with all possible models.

The Mental Models account of syllogistic reasoning has been applied to categorical (Oakhill, Johnson-Laird, & Garnham, 1989) and conditional (Santamaria et al., 1998; Torrens et al., 1999) syllogistic reasoning. To illustrate, take the example of the following categorical premises:

No cigarettes are inexpensive

Some addictive things are inexpensive

According to the Mental Models theory, reasoners first construct a model of the premises using symbolic tokens. In standard Mental Models notation (Johnson-Laird & Byrne, 1991), each horizontal line represents a single model of the premises. Square brackets indicate that all members of the set are represented exhaustively, that is, no members of the set may occur elsewhere in the model; ellipses after a model indicate that more models may be added to the set in addition to the initial models (i.e., the set has not been entirely “fleshed out”). In these additional fleshed out models, ¬ symbolizes the negation of a token. The first premise, *No cigarettes are expensive,* would be represented as below:

[cigarettes]

[cigarettes]

[inexpensive]

[inexpensive]

. . .

The above model indicates that cigarettes and inexpensive things exist as tokens that are mutually exclusive. The second premise, *Some addictive things are inexpensive*, would be represented as below:

addictive things inexpensive

addictive things inexpensive

. . .

In the second step of the reasoning process, reasoners must combine these models to arrive at an integrated model of both premises, and then derive a conclusion. The integrated model in this example would be:

[cigarettes]

[cigarettes]

[inexpensive] addictive things

[inexpensive] addictive things

This integrated models yields the valid conclusion, *Some addictive things are not cigarettes*. This model also yields the conclusion, *No cigarettes are addictive things*. If the models are fully fleshed out, the reasoner will arrive at the following integrated model:

[cigarettes]		
[cigarettes]		addictive things
	[inexpensive]	addictive things
	[inexpensive]	addictive things

This fleshed out model falsifies the conclusion that *no cigarettes are addictive things*.

Now, take the following conditional syllogism:

If an animal is a bird, then it has a beak

If an animal has a beak, then it has feathers

Again, reasoners would begin by forming models of the premises. The first premise yields the following *fleshed out* model (Johnson-Laird, 2008, personal communication):

[bird]	[beak]
[bird]	[beak]
[¬bird]	[beak]
[¬bird]	[¬beak]
	. . .

The second premise yields the following model:

[beak]	[feathers]
[beak]	[feathers]
[¬beak]	[feathers]
[¬beak]	[¬feathers]
. . .	

The integrated model becomes:

[bird]	[beak]	[feathers]
[¬bird]	[beak]	[feathers]
[¬bird]	[¬beak]	[feathers]
[¬bird]	[¬beak]	[¬feathers]

The model of the premises yields only one valid conclusion: *If an animal is a bird, then it has feathers.* Note that it is not necessary to flesh out the model to refute this conclusion; this is a one-model syllogism.

Errors in deduction arise when reasoners fail to account for all models of the premises (Johnson-Laird & Byrne, 1991). Specifically, for belief bias, the believability of the conclusion affects the model generation step of the reasoning process. First, reasoners construct an initial model of the premises, rejecting conclusions that are invalid given this model. Second, reasoners consider the believability of the conclusion: If the conclusion is believable, then reasoners are more likely to prematurely halt the search for models that falsify the conclusion and are more likely to accept the conclusion as valid. However, if the conclusion is unbelievable, reasoners are more motivated to search through all possible models, which explains why effects of conclusion validity are greater when conclusions are unbelievable (Newstead, Pollard, Evans, Allen, 1992; Oakhill et al., 1989). The Mental Models theory of reasoning has been called upon to explain belief bias resulting from both categorical syllogistic reasoning (Oakhill et al., 1989) and conditional syllogistic reasoning (Santamaria et al., 1998; Torrens et al., 1999).

A second influential class of reasoning theories (though not necessarily exclusive from mental model accounts, Evans, 2003; Byrne & Johnson-Laird, 2009), is known as Dual Process accounts of reasoning (see Evans, 2003, 2008 for reviews). Several variations of Dual Process accounts have been proposed (e.g., Evans, 1989; Sloman, 1996; Stanovich, 1999), although all of them share similar features. In general, Dual Process theorists posit that there are two systems that may be engaged during reasoning (the names of these two systems vary across theories; the terminology of Stanovich, 1999, is used here): System 1 is thought to be unconscious, automatic, fast, and to provide a response based on heuristics or associative processes; System 2 is thought to be conscious, deliberative, slow, and to provide a response based on logical analysis of the problem. System 2 is also thought to suppress and override System 1 in cases where System 1 provides a faulty response (Stanovich, 1999). In the case of belief bias, System 1 is thought to provide a response based on the believability of the conclusion (i.e., “valid” if the conclusion is believable). This response will be correct in cases where the validity and the believability of the conclusion are congruent. However, where the validity and believability of the conclusion conflict (e.g., a valid, yet unbelievable conclusion), if System 2 is engaged, there is a higher likelihood of arriving at the correct answer based on logical analysis of the syllogism.

Dual Process theories differ in how they account for the time course of the reasoning process, falling into two categories. First, Default Interventionist Models (Kahneman & Frederick, 2002, 2005; Evans, 2006) posit that System 1 provides an automatic, default response based on the believability of the conclusion, and that System 2 may subsequently override the response provided by System 1 and provide a response based on logic if there is sufficient time, cognitive resources, and motivation. On the other hand, Parallel Process Models (Sloman, 1996, 2002) propose that System 1 and System 2 are always activated in parallel and that both systems provide a response. Because the two systems are activated in parallel, reasoners are often aware when there is a conflict between the validity and believability of the conclusion. Parallel Process Models have garnered some support from response latency analyses, which show that reasoners typically spend longer on problems in which the validity and believability of the conclusion conflict (Stupple & Ball, 2008, De Neys, 2007).

Reasoning with False Premises

It is clear that the content of the conclusion can greatly affect the outcome of deductive reasoning: Several experiments have replicated the belief bias phenomenon and much effort has been expended explaining this phenomenon. It is curious, then, that relatively little attention has been paid to how the content of the premises within a syllogism affect reasoning. Indeed, two thirds of a syllogism is comprised of information in premises, and the influential theory of Mental Models itself postulates that one begins the reasoning process by modelling the premises. In most studies of conditional and categorical syllogistic reasoning, participants are instructed to assume that premises are true for the purposes of the experiment. However, there are several instances in daily life where we must reason about information that is contrary to our beliefs. For example, scientists must often reason about data that are contrary to their theoretical beliefs, and consumers should be able to reason about unlikely product claims.

Relatively few experiments have examined how people reason about believable versus unbelievable information in premises. This gap in the literature is not without reason: Syllogisms do not readily lend themselves to varying the believability of premises and conclusions. As an illustration, it is logically impossible to create a valid syllogism with two believable premises and an unbelievable conclusion. Nevertheless, there are experiments which have successfully isolated some effects of premise believability on deductive reasoning.

In the realm of conditional reasoning with a major conditional premise and minor categorical premise, George (1995, 1997, 1999) has presented evidence that when major conditional premises are made uncertain by the introduction of qualifying terms such as "probably," reasoners are less likely to accept valid conclusions, as they likely transfer this uncertainty to the conclusion. Further, Markovits and colleagues (Markovits, Saelen, & Forgues, 2009; Markovits & Schroyens, 2007) have shown that under instructions to assume that premises are true, participants accept unbelievable conclusions more readily when the major premise is false than when it is true or plausible. These researchers suggest that when the premise is false, reasoners inhibit information about the items contained in the premises, making it more palatable to accept an unbelievable conclusion. Moreover, the Mental Models account addresses the influence of premise believability on conditional reasoning (Johnson-Laird & Byrne, 2002). According to the *principle of semantic modulation*, "the meanings of the antecedent and consequent, and coreferential links between these two clauses, can add information to models, prevent the construction of otherwise feasible models of the core meaning, and aid the process of constructing fully explicit models" (Johnson-Laird & Byrne, 2002, p. 13). Essentially, the meaning of the premises is proposed to facilitate or to block the formation of models. For example, the proposition, *if it is a game, then it is not soccer*, theoretically yields the following models:

Game	¬Soccer
¬Game	¬Soccer
¬Game	Soccer

However, rational individuals would agree that it is not possible for soccer not to be a game. This knowledge would block the formation of the third model of the premises. Thus, models that are believable are more readily formed than models that are false.

The experiments discussed above used problems with a major conditional premise and a minor categorical premise. Thompson (1996) examined deductive reasoning for conditional syllogisms with two conditional premises, analogous to standard categorical syllogisms (e.g., "If an animal is a bird, then it has a beak; If an animal has a beak, then it flies"). She provided participants with two or three sets of premises and corresponding valid and invalid conclusions. Premise believability varied between participants: For each participant, all premises were either believable or unbelievable. For each set of premises, participants were asked to select, from a list of four possible conclusions, all those that were valid given the premises. Her results indicated that when premises were believable, participants tended to endorse more conclusions as valid than when premises were neutral or unbelievable. Belief bias resulted regardless of premise believability: For both believable and unbelievable premises, participants were more likely to endorse believable conclusions as valid. Unfortunately, Thompson was not able to examine all combinations of believability of premises and conclusions, but her data clearly show that believable premises caused participants to endorse more conclusions relative to unbelievable premises. Further, because belief bias resulted for both believable and unbelievable premises, Thompson concluded that premises are considered separately from conclusions, stating, "given the absence of an interaction between validity and belief, [...] premise believability acts independently of any logical analysis that is performed" (p. 318).

This effect has not been successfully replicated, however. Torrens, Thompson, and Cramer (1999) presented participants with similar syllogisms with two conditional premises. Although they report that participants endorsed more believable than unbelievable conclusions (i.e., the belief bias effect), they did not find that participants endorsed more conclusions following from believable than from unbelievable premises. This failure to replicate Thompson's (1996) finding may be due to differences in methodologies: In Thompson's original experiment, premise believability was manipulated between subjects, whereas Torrens, Thompson, and Cramer manipulated premise believability within subjects. That this effect has not been successfully replicated warrants further study.

Studies of premise believability in categorical syllogistic reasoning are extremely lacking; in fact, I am aware of only one experiment that systematically varied premise believability to determine its effects on reasoning processes for categorical syllogisms. Cherubini, Garnham, Oakhill, and Morley (1998; Experiments 3 and 4) presented participants with nine pairs of categorical premises in which either one or both were false, and asked participants to generate conclusions following from the premises. Interestingly, they found that when both premises were false, belief bias was not evident; that is, participants were not more likely to generate believable than unbelievable conclusions. The authors proposed a modified

Mental Models account, in which information about the terms in the premises is brought to mind before the premises are modeled. If this information is consistent with one's previous knowledge, then one is likely to use this information when deriving a conclusion; however, if it is inconsistent, it will not be used. Again, this research suffers because of the logical structure of syllogisms: The investigators were not able to examine how participants respond to syllogisms with believable premises and unbelievable, valid conclusions.

Conditional versus Categorical Syllogistic Reasoning

This review of past research on premise believability brings to light an interesting discrepancy between categorical syllogisms and conditional syllogisms with double conditional premises. In her study of conditional syllogisms, Thompson (1996) found that premise believability did not influence the belief bias effect; rather, believable premises simply increased the number of "valid" responses given by participants. On the other hand, Cherubini et al. (1998) found that when premises were unbelievable, the belief bias effect was completely eliminated. It is compelling that differential effects of premise believability were found for these two types of syllogisms, given that existing theories of deductive reasoning would not predict differences. Indeed, the two types of syllogisms have identical task demands (i.e., deducing whether a conclusion is valid given two premises) and both elicit the belief bias effect, so without taking differential effects of premise believability into account, there is no reason to assume that reasoners treat these syllogisms differently. It seems, then, that varying premise believability could be a fruitful tool for examining potential differences in how reasoners treat conditional and categorical syllogisms.

This question is a crucial one, given that conditional and categorical syllogisms have been used in the literature to inform the same theories of deductive reasoning and belief bias.

The majority of experiments examining deductive reasoning use categorical syllogisms: Such experiments have been used to support or to falsify various reasoning theories (e.g., Newstead et al., 1993; Klauer, Munch, & Naumer, 2000). However, experiments using conditional syllogisms have also played an influential role in deductive reasoning theories. For example, as discussed above, following Thompson's (1996) experiment using conditional syllogisms, she concluded that premises are considered independently of the conclusion during deductive reasoning, and suggested that current theories of reasoning need to account for this fact. Santamaria and colleagues (Santamaria et al., 1998) report that, just as with categorical syllogisms, participants drew a valid conclusion from double conditional premises more often when the conclusion was believable than when it was unbelievable, and they drew believable valid conclusions quicker than unbelievable conclusions. They interpret these findings in terms of the same Mental Models theory that accounts for results of experiments using categorical syllogisms.

On this same point, Torrens and colleagues (Torrens et al., 1999) used conditional syllogisms in a large study examining individual differences in the belief bias effect, including how skilled participants were at considering different models of premises. They found that intelligence and general deductive reasoning ability was not related to the extent to which participants were influenced by conclusion believability, however, performance on the model generation task strongly predicted belief bias. They concluded that, "the results provide clear support for the Mental Models Theory [...] assumption that the search for alternative models is implicated as an integral part of the belief bias effect" (p.22). Clearly, it is important to question whether these findings contribute to theories of deductive reasoning in general, or theories of reasoning about specific syllogisms. If reasoning about categorical and conditional reasoning differ along important dimensions, such as how the belief bias effect is influenced by premise content, then it becomes imperative for theories of deductive reasoning to be tailored to—or at least to accommodate—different forms of reasoning.

Indeed, other researchers have just recently begun to question whether categorical and conditional syllogisms share similar underlying mechanisms. Reverberi et al. (2010) employed fMRI technology to demonstrate that this may not be the case. They presented participants with 72 categorical syllogisms and 60 conditional syllogisms with two conditional premises (note that they did not control for premise believability) and identified brain regions active during premise integration, which they defined as the point at which the second premise was read. Their results showed that for both conditional and categorical syllogisms, there was activation in the left prefrontal cortex and left basal ganglia. There was additional activation during the integration of categorical premises in the occipital medial gyrus, lateral parietal lobe, and precuneus. These results provide some initial evidence that conditional and categorical syllogistic reasoning may not be governed by the same cognitive mechanisms and, at the very least, that considering categorical premises involves some extra processing over considering conditional premises.

Further, in a study of categorical syllogistic reasoning, Evans and Curtis-Holmes (2005) demonstrated that under time pressure, reasoners made more responses that were consistent with the believability of the conclusion. This finding was taken as evidence that syllogistic reasoning is consistent with Default Interventionist theories of reasoning, because reducing the time that one has to reason through a problem reduces the probability that System 2 will have enough time to override the default answer provided by System 1. However, Evans, Handley, and Bacon (2009) applied a similar methodology to conditional problems (these problems had one major conditional premise and one minor categorical premise), and did not find that time pressure influenced belief bias. They concluded that “these findings confirm our suggestion that belief effect in conditional inference may operate quite differently from the belief bias observed in syllogistic reasoning" (p. 82).

Rationale for this Dissertation

Evidence is recently coming to light that categorical and conditional syllogistic reasoning may rely on different reasoning mechanisms. Although the two types of syllogisms elicit seemingly indistinguishable belief bias effects, it may be the case that varying the content of premises, a manipulation which has largely been overlooked, provides a tool to elucidate differences underlying reasoning about these two types of syllogisms. Previous research examining effects of premise believability on reasoning has produced inconsistent findings and is fraught with methodological difficulties, including those imposed by the logical structure of syllogisms.

This dissertation addresses these issues, beginning with successful replication of the discrepant findings of previous researchers (Experiment 1). In Experiment 2, response latencies provide evidence that reasoners are less likely to take the believability of the conclusion into account for categorical syllogisms with unbelievable than with believable premises, and are slower to reason about conditional syllogisms with unbelievable than with believable premises. Experiments 3 and 4 demonstrate that conditional and categorical premises differ in the extent to which reasoners assume the premises to be true; specifically, instructions to assume that premises are true (Experiment 3) and the word "if" (Experiment 4) both appear to trigger hypothetical thinking. Finally, Experiment 5 uses the somewhat novel approaches of eye-tracking to support the proposed time course of reasoning. It is argued that reasoners initially assume that conditional premises are true, and then reject them after the conclusion has been reasoned about, whereas reasoners evaluate the believability of categorical syllogisms prior to reasoning about the conclusion. Further, Experiment 5 explores memory for the content of syllogisms, and supports the position that reasoners disregard content of categorical syllogisms with unbelievable premises. Overall, I will use premise believability as a tool to extract differences underlying categorical and conditional syllogistic reasoning, and in doing so, will develop a theory of how reasoners consider the premises of these two types of syllogisms.

Experiment 1

As noted in the introduction, there is some evidence that varying premise believability results in differential reasoning outcomes for categorical and conditional syllogisms (Thompson, 1996; Cherubini et al., 1998). However, these two investigations may not be directly comparable: They used different methodologies (Thompson used a between-subjects conclusion selection task, whereas Cherubini et al. used a within-subjects generation task) and of course came from different laboratories. This preliminary experiment was designed to replicate the results from previous experiments using a typical conclusion validation procedure with a within-subjects design so that conditional and categorical reasoning could be directly compared.

This experiment also attempted to address the major shortcomings of previous experiments imposed by the logical structure of syllogisms. Previous experiments were unable to fully examine effects of premise believability because it is not possible to create syllogisms with certain combinations of premises and conclusions. Experiment 1 attempted to overcome this limitation so that a fully orthogonal examination of premise believability, conclusion validity, and conclusion believability could be carried out. For most premises and conclusions, statements were used that were true or false by definition. However, for the categorical and conditional syllogisms where this was not possible, some exceptions were made. For categorical syllogisms with believable premises and valid conclusions, conclusions could not be false by definition. However, conclusions could be used that are unsound, or not fully believable. These statements have opposite forms which are true. For example, the converse of the conclusion "Some dogs are not cats" is "All dogs are not cats," or "No dogs are cats," and is true by definition. Also, for conditional syllogisms, some premises were necessitated which were not always true or false, but were true or false in most cases. For example, although the premise "If a vehicle has wheels, then it has a steering wheel," is true in most but not all circumstances, it was nonetheless used as a believable premise.

Thus, Experiment 1 was designed to fully investigate how varying premise believability differentially affects reasoning for categorical and conditional syllogisms. It was hypothesised that, given the results of Thompson (1996), for conditional syllogisms, varying premise believability would not influence the belief bias effect. It was predicted, however, that participants would endorse as valid more conclusions following from believable than from unbelievable premises. Following the results of Cherubini et al., (1998), it was predicted that for categorical syllogisms, the belief bias effect would be found when premises were believable, but not unbelievable.

Method

Participants

Participants were 69 undergraduate students (30 male, 39 female, mean age = 19.84 yrs, SD = 1.60 yrs) enrolled in psychology courses at the University of Waterloo who participated in exchange for course credit. Participants reported never having taken a course in formal logic.

Materials

Syllogisms were constructed that varied orthogonally across four dimensions (see Appendix A for a complete list of syllogisms used). First, syllogisms were either categorical or conditional in nature; second, premises were either believable or unbelievable; third, conclusions were either believable or unbelievable; and fourth, conclusions were either valid or invalid. Thirty-two syllogisms were constructed, with two syllogisms in each cell. Conditional syllogisms took the form (with words replacing the letters in the actual task), *If A, then B;*

If B, then C;

Therefore, if A, then C (Valid)

or

Therefore, If C, then A (Invalid).

Categorical syllogisms took the form,

No A are B;

Some C are B;

Therefore some C are not A (Valid)

or

Therefore, Some A are not C (Invalid)

Pilot testing for the stimuli was conducted with an independent group of 40 undergraduate students, who rated the believability of premises and conclusions (see pilot data in Appendix A). In general, statements that were rated highly believable or unbelievable were used to construct the syllogisms. For most problems, statements that were true or false by definition comprised the syllogisms, with a few necessary exceptions. First, using typical true or false statements, it is not possible to create valid syllogisms (either categorical or conditional) with believable premises and an unbelievable conclusion. Second, it is not possible to create invalid, categorical syllogisms with unbelievable premises and an unbelievable conclusion. As such, to maintain the orthogonal design of the experiment, some small liberties were taken with these problems. Specifically, either premises were used that were believable, but not necessarily true in all cases (e.g., "If a vehicle has wheels, then it has a steering wheel"), or conclusions were considered unbelievable if their converse was true (cf. Cherubini et al., 1998). For example, the converse of the statement used as an unbelievable conclusion, "Some reptiles are not rats" is true (i.e., "No reptiles are rats" or "all reptiles are not rats").

Procedure

Each participant reasoned through all 32 syllogisms, which were presented to participants in paper booklets with four problems per page. Stimuli were presented in one of eight random orders. Participants first read the following instructions:

> This is an experiment to test people's reasoning ability. This booklet contains 32 problems. Each contains two statements followed by a conclusion. You are asked to determine if the conclusions may be logically deduced from the statements. If you judge that the conclusion necessarily follows from the statements (i.e., that it is a valid conclusion based on the information given in the two statements), you should answer by circling VALID, otherwise INVALID. Please take your time and ensure that you have the right answer. Solve the problems in order, without flipping forward or backward in the booklet."

Participants were tested in small groups of two to eight individuals. They were given as much time as they needed to complete the task.

Results

The proportion of times that participants indicated conclusions were valid are depicted in Figure 1 for Conditional Syllogisms with Believable and Unbelievable Premises, and in Figure 2 for Categorical Syllogisms with Believable and Unbelievable Premises. Endorsement rates (i.e., number of times participants indicated that a conclusion was valid) were submitted to a 2 (Syllogism Type: Conditional, Categorical) x 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVA.

Results will be discussed, first, in terms of the overall belief bias effect collapsed across all problems. Second, the three-way interaction between Syllogism Type, Premise Believability, and Conclusion Believability will be examined to determine whether effects of Conclusion Believability vary as a function of Syllogism Type and Premise Believability. Third, the Syllogism Type x Premise Believability interaction will be examined to determine whether participants endorsed more conclusions following from Believable than from Unbelievable Premises for Conditional Syllogisms compared to Categorical Syllogisms. Other results that are not directly relevant to the theory addressed here are presented in non-bold font in Appendix C for this and all following experiments.

Overall Belief Bias Effect. Collapsed across all problems, the typical belief bias effect was found: There were main effects of Conclusion Believability, $F(1,68) = 48.44$, $MSE = .515$, $p < .001$, $\eta^2 = .416$, such that participants endorsed more believable than unbelievable conclusions, and Conclusion Validity, $F(1,68) = 92.03$, $MSE = .862$, $p < .001$, $\eta^2 = .575$, such that participants endorsed more valid than invalid conclusions.. Further, there was an interaction found between Conclusion Validity and Believability, such that effects of Conclusion Believability were larger for invalid than for valid conclusions, $F(1,68) = 10.71$, $MSE = .368$, $p <.01$, $\eta^2 = .136$.

Effects of Conclusion Believability as a Function of Syllogism Type and Premise Believability. There was a marginally significant three-way interaction between Syllogism Type, Premise Believability, and Conclusion Believability, $F(1,68) = 3.84$, $MSE = .341$, $p = .054$, $\eta^2 = .053$. To examine the locus of this interaction, conclusion endorsement rates for Categorical and Conditional Syllogisms were submitted to two separate 2 (Premise Believability: Believable,

Conditional Syllogisms with Believable Premises

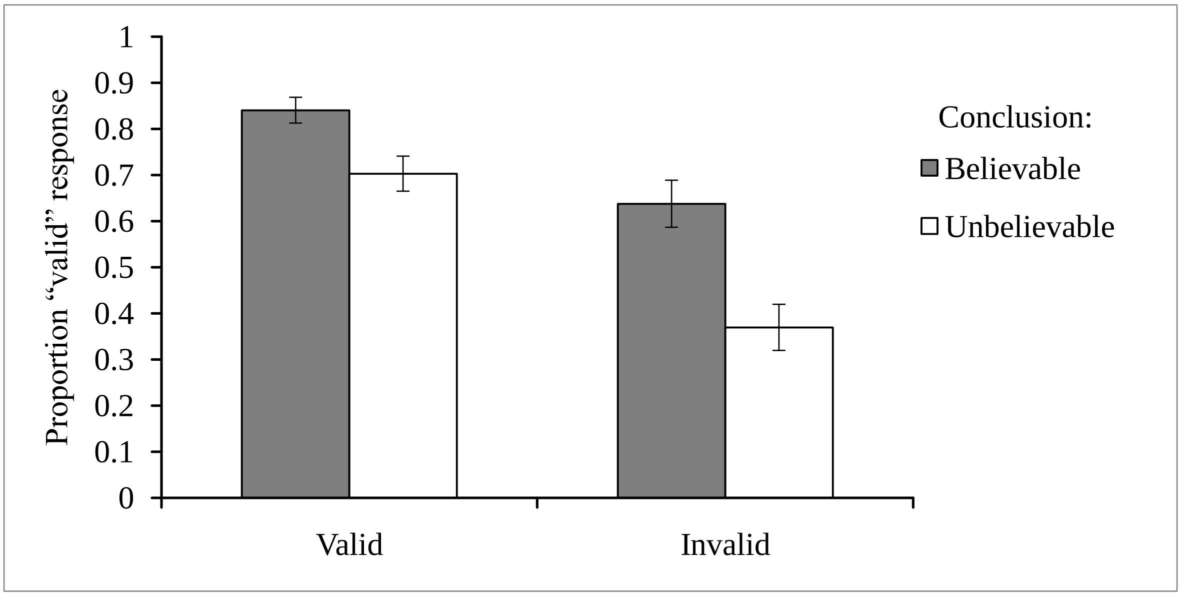

Conditional Syllogisms with Unbelievable Premises

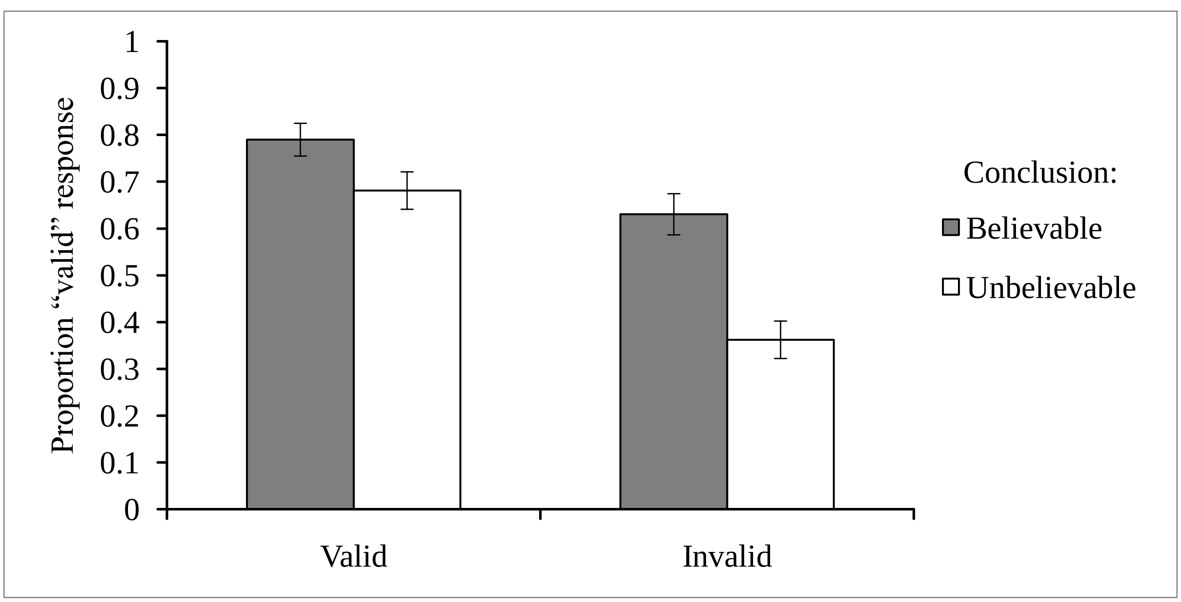

Figure 1. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable Premises and Unbelievable Premises. *The error bars for this figure (and for all remaining figures) represent the standard error for each condition.*

Categorical Syllogisms with Believable Premises

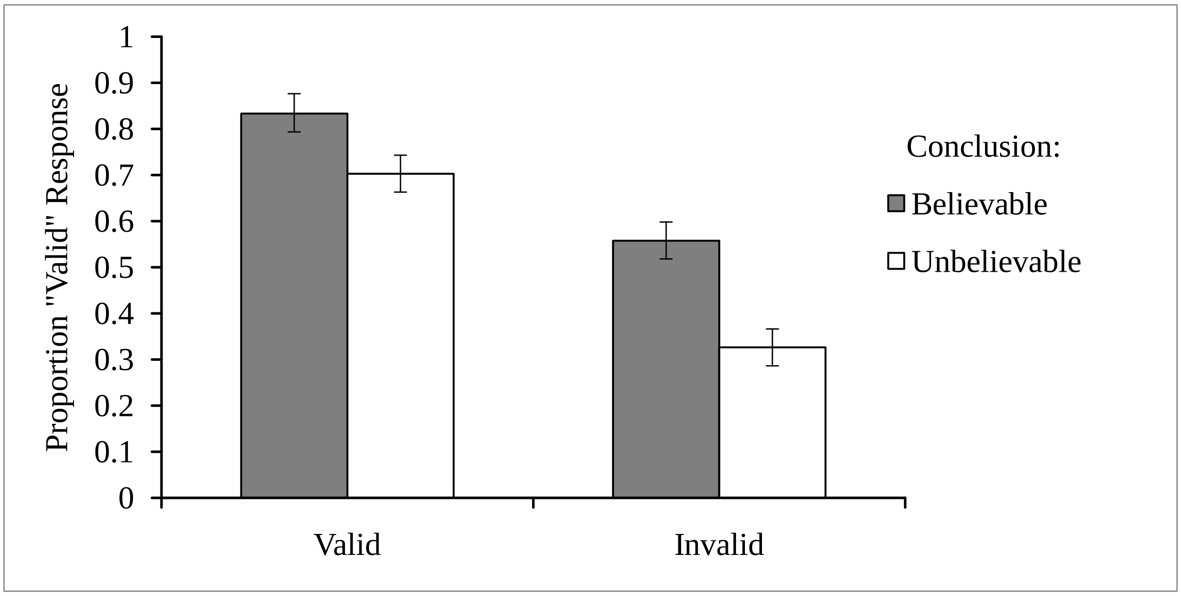

Categorical Syllogisms with Unbelievable Premises

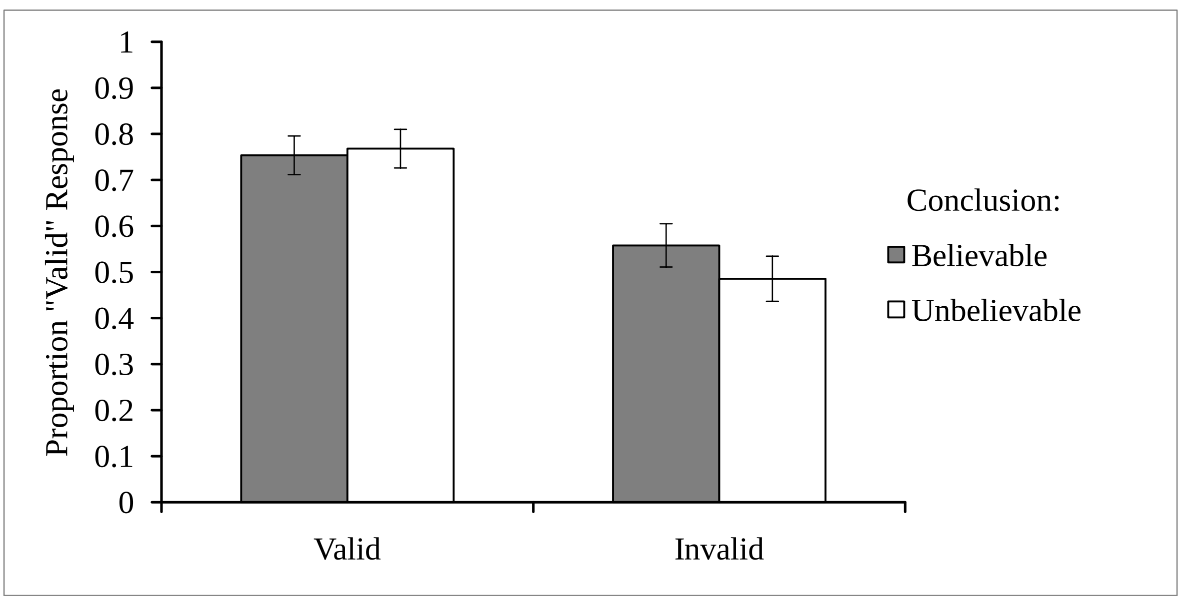

Figure 2. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable Premises and Unbelievable Premises.

Unbelievable) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVAs. Importantly, for Categorical Syllogisms, Premise Believability interacted with Conclusion Believability, F(1,68) = 8.26, *MSE* = .387, p <.01, η^2 = .108. Paired samples t-tests revealed that when premises were believable, participants were more likely to endorse believable conclusions, $t(68) = 4.57$, p <.001, which is not surprising, given the robustness of the belief bias effect. When premises were unbelievable, however, Conclusion Believability played no role in how frequently participants endorsed the conclusion, $t(68) = 0.73$, p =.465. In contrast to that found for Categorical Syllogisms, there was no significant interaction between Premise Believability and Conclusion Believability for Conditional Syllogisms, F(1,68) = 0.12, *MSE* = .250, p = .734, η^2 = .002.

Conclusion Endorsement for Categorical and Conditional Syllogisms as a Function of Premise Believability. The Syllogism Type x Premise Believability interaction was not significant, F(1,68) = 2.13, *MSE* = .435, p = .149, η^2 = .030. Here, participants did not endorse more conclusions following believable than unbelievable premises for either Categorical Syllogisms, F(1,68) = 1.38, *MSE* = .526, p = .245, η^2 = .020, or Conditional Syllogisms, F(1,68) = 0.61, *MSE* = .430, p = .439, η^2 = .009.

Discussion

This experiment was designed to replicate and extend the results of Thompson (1996) and of Cherubini et al. (1998), who found different consequences of varying premise believability for conditional and categorical syllogisms. In general, their results were replicated using a completely matched and orthogonal design. For categorical syllogisms, belief bias was found only when premises were believable; that is, both conclusion validity and believability affected conclusion endorsement, and there was a greater effect of conclusion believability on invalid conclusions. When premises were unbelievable, only conclusion validity affected conclusion endorsement – there was no belief bias effect. Further, there was no main effect of premise believability; participants endorsed the same overall number of conclusions following from believable premises as from unbelievable premises. These results, using a different task (i.e., validation rather than generation), replicated the results of Cherubini et al. (1998); namely, when premises of categorical syllogisms are unbelievable, no belief bias effect is found. That this effect was replicated using a different task and different stimuli points to its generalizability and robustness.

Thompson (1996) found that for conditional syllogisms, belief bias occurred for both believable and unbelievable premises; however, participants endorsed a greater number of conclusions following from believable than from unbelievable conclusions. The results of the current experiment partially replicate Thompson's results, in that belief bias was found regardless of premise believability for conditional syllogisms; that is, participants endorsed more valid than invalid conclusions, more believable than unbelievable conclusions, and the effects of conclusion believability were greater on invalid than on valid conclusions. Although the data show a small trend in the direction of participants endorsing more conclusions that were preceded by believable than by unbelievable premises, this difference was not significant. It is not clear whether Thompson's finding is replicable, or whether the effect was not replicated here due to the difference in methodology between this experiment and Thompson's. Specifically, the current experiment varied premise believability within subjects, whereas Thompson varied premise believability between subjects. Given that others using a within-subjects design (Torrens et al., 1999) failed to find effects of premise believability with conditional syllogisms, it may be that these effects are only evident when participants see either believable or unbelievable premises. Further, the current experiment used a conclusion validation paradigm, whereas Thompson used a conclusion selection task. For now, however, it is interesting that varying premise believability interacted with conclusion believability for categorical syllogisms but not for conditional syllogisms.

This experiment also used a novel approach of taking some small creative liberties with stimuli so that premise believability, syllogism type, conclusion validity and conclusion believability could be examined in a completely orthogonal design. The syllogisms in question, specifically categorical and conditional syllogisms with believable premises and valid, unbelievable conclusions demonstrated what would be predicted from the belief bias effect: Participants endorsed more believable than unbelievable conclusions, even when the unbelievable conclusions were not false by definition. Thus, the stimuli seem to have been successful at simulating strictly believable and unbelievable statements and can be used in paradigms to address effects of premise believability on deductive reasoning.

Along with the results of Thompson (1996) and Cherubini et al., (1998), this experiment provides significant evidence that premise believability differentially affects how people reason through conditional and categorical syllogisms. This in turn suggests that premises are considered differently for categorical and conditional syllogisms. How then might premises be conceptualized? The reasoning literature itself does not provide much evidence to answer this question. The Mental Models account (Johnson-Laird & Byrne, 1991), which purports that the first step in reasoning is a model of premises, claims that the content of the premises can affect how premises are modelled. For example, reasoners will not model as readily what is contrary to their knowledge about the world (Byrne & Johnson-Laird, 2009), and the semantic content of premises can modulate how models are constructed (with unbelievable models being sacrificed for what is believable; Quelhas, Johnson-Laird, & Juhos, 2010). However, even the Mental Models account places greatest importance on the believability of the conclusion, claiming that the believability of the *conclusion* affects how reasoners conceptualize the premises, rather than the other way around.

It is evident that information contained in premises influences reasoning, but to determine how, we may have to consult literatures other than the reasoning literature. Philosophers have long debated how everyday information is conceptualized. Rene Descartes (1644/1984) claimed that we take in information without initially judging whether it is true or false; our belief assessment can be suspended until we rationally analyze the veracity of the information. In contrast, Baruch Spinoza (1677/1982) disagreed, claiming that information must be accepted as true for it to be comprehended. This acceptance happens automatically, and the information can be discredited or "disbelieved" in a process that takes time and cognitive effort. Some modern theorists agree with Spinoza's claim. Johnson-Laird (1988) articulately asserted that to comprehend information, one must "imagine how the world should be, granted its truth" (p.110).

A particularly strong proponent of Spinoza's philosophy on automatic belief is Gilbert (1991). Along with colleagues, he published several experiments showing that when people are provided with information, they are likely to believe the information, even when they are told the information is false, when their processing of the information was interrupted by concurrent cognitive load (Gilbert, Tafarodi, & Malone, 1993), or when processing time was cut short (Gilbert, Krull, & Malone, 1990). Other researchers have found that people are able to evaluate true statements faster than false statements (Just & Carpenter, 1976). These results suggest that the default is to believe information, even when it is signalled to be false, and that disbelieving the information may only happen given sufficient time and resources. Note, however, that the findings are not entirely consistent: Hasson, Simmons, and Todorov (2005) found that even under cognitive load, participants remembered false statements when they provided useful information. Thus, the context and content of statements may modulate how they are conceptualized in terms of their veracity.

Thus, there is evidence that, at least in some contexts, people initially automatically accept statements to be true and then spend time at a later point disbelieving or discrediting them. Applying this theory to the results obtained in Experiment 1, it may be possible to explain differential effects of premise believability for categorical and conditional syllogisms. Take first, conditional premises, the theorized time course of which is depicted in Figure 3. Regardless of premise believability, participants reason about conclusions in a similar fashion, as evidenced by belief bias resulting for both believable and unbelievable premises. It may be the case that reasoners automatically assume that believable and unbelievable conditional premises are true for the purposes of reasoning about the conclusion. Subsequently, reasoners may discredit unbelievable premises, which then leads them to reject conclusions following from unbelievable premises. This would explain Thompson's (1996) finding that participants endorsed more conclusions following from believable than from unbelievable premises.

Consider now categorical syllogisms, for which the theorized reasoning process is outlined in Figure 4. In Experiment 1, there was no evidence that participants took into account conclusion believability when premises were unbelievable; that is, premise believability seems to have determined the course of the reasoning process. Perhaps, contrary to reasoning about conditional syllogisms, reasoners do not accept categorical premises to be true, but rather initially evaluate their believability *prior* to reasoning about the conclusion. Then, when premises are believable, believability information from the rest of the problem is taken into account, and when premises are unbelievable, believability information in the conclusion is disregarded.

In summary, I have outlined a preliminary theory of premise conceptualization. This theory states that reasoners assume that conditional premises are true prior to reasoning, and then discredit them *after* or somewhere outside of reasoning about the conclusion (i,e., in an independent process). In sharp contrast, categorical premises are evaluated *prior* to reasoning about the conclusion, and premise believability determines the course of the reasoning process. This theory will be tested in the following experiments in this dissertation.

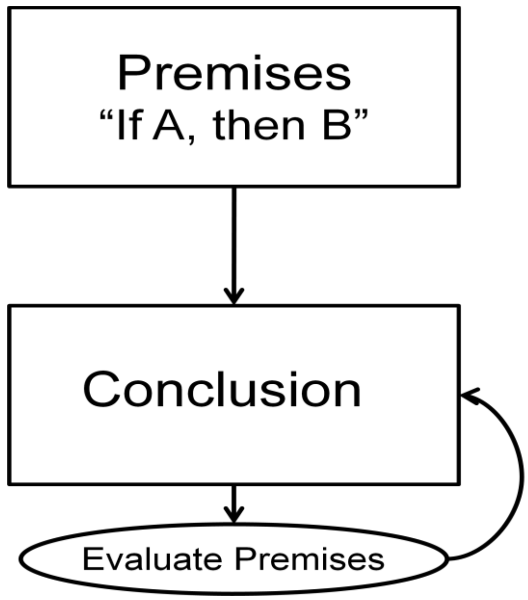

Figure 3. The theorized reasoning process for Conditional Syllogisms. *It is hypothesized that reasoners evaluate the believability of the premises after reasoning about the conclusion, and then alter responses to conclusions accordingly.*

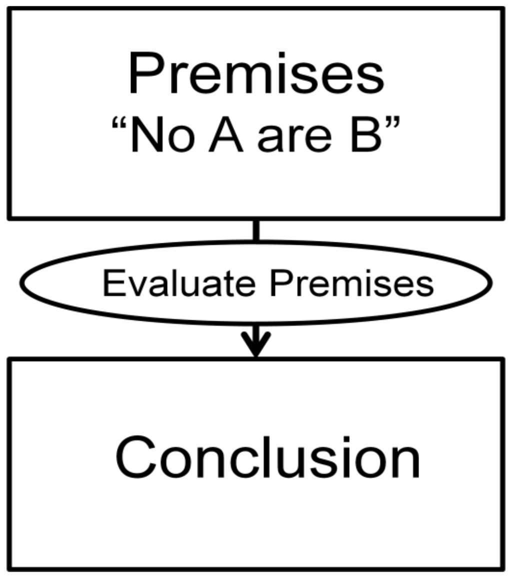

Figure 4. The theorized reasoning process for Categorical Syllogisms. *It is hypothesized that reasoners evaluate the believability of the premises immediately, prior to reasoning about the conclusion.*

Experiment 2

The goal of Experiment 2 was to replicate Experiment 1and extend the findings through the measurement of response latencies. Although using response latencies as evidence for underlying reasoning mechanisms has been successfully used by some researchers (e.g., Thompson, Striemer, Reikoff, Gunter, & Cambell, 2003; Ball, Phillips, Wade, & Quayle, 2006), examining them remains an underused tool, given that most theories of deductive reasoning and belief bias make specific predictions regarding response latencies. The current experiment will use response latencies to evaluate predictions made by the proposed theory of premise evaluation in categorical and conditional reasoning.

Previous experiments examining response latencies in syllogistic reasoning have shown somewhat inconsistent results. Ball et al. (2006) report that participants spent longer on categorical syllogisms in which the validity and believability of the conclusion conflicted than on syllogisms in which the validity and believability of the conclusion were congruent. Conversely, Thompson et al. (2003) did not find this interaction. Rather, these researchers found that, for categorical syllogisms, participants took longer to respond to believable than unbelievable conclusions, and invalid than valid conclusions. Although these results are in opposition, they nonetheless reveal effects of conclusion believability on response latency.

A secondary goal of Experiment 2 was to address the possibility that differential belief bias effects are the result of different numbers of models possible rather than the type of syllogism. In Experiment 1, the categorical syllogisms used were multiple-model syllogisms, whereas the conditional syllogisms used were single-model syllogisms. As such, it is feasible that, given the ease with which single-model syllogisms are conceptualized, reasoners are not influenced by premise believability and demonstrate similar belief bias effects regardless of whether the premises are believable or unbelievable. Conversely, given the complexity of multiple-model syllogisms, it is possible that reasoners are more likely to be influenced by premise believability in a way that influences the belief bias effect. Experiment 2 was designed to eliminate this potential confound by varying the number of models possible in the stimuli.

Given the theory of premise evaluation outlined above, several predictions can be drawn. First, if participants initially accept all conditional premises to be true, and then discredit unbelievable premises, then participants should take longer reasoning about conditional syllogisms with unbelievable than with believable premises. Second, if participants take the believability of the conclusion into account when conditional premises are believable or unbelievable, there should be different response latencies for believable and unbelievable conclusions. Third, if reasoners evaluate categorical premises prior to reasoning, rather than assuming premises are true and then rejecting false premises in an additional process, there should be no difference in response latencies for believable and unbelievable premises. Finally, if reasoners do not take believability information of the conclusion into account when premises are unbelievable for categorical syllogisms, then there should be no effect of conclusion believability on response latencies.

Method

Participants

Participants were 40 students (16 male, 24 female, mean age = 20.15 yrs, SD = 2.47 yrs) enrolled in psychology courses at the University of Waterloo who participated in exchange for course credit. None reported having taken a course in logic.

Materials

Thirty-two syllogisms were used that varied across the same dimensions as those used in

Experiment 1 but that incorporated an additional dimension: number of models required given the premises. To vary difficulty within categorical syllogisms, the number of possible models of the premises was varied. Single-model categorical premises took the form “All A are B; All B are C;” whereas multi-model premises took the form “No A are B; Some C are B.” Single-model conditional premises were of the form, “If A, then B; If B, then C.” Unfortunately, a reasonable multi-model conditional syllogism is not readily apparent (Johnson-Laird, personal communication, 2008). Thus, to increase the difficulty of conditional syllogisms to approximate that of categorical multi-model syllogisms, negation was introduced into the premises, which took the form, “If A, then not B; If not B; then C.” There was one syllogism in each Type x Models x Premise Believability x Conclusion Validity x Conclusion Believability cell. Eight Categorical, Multi-model syllogisms and eight Single-model Conditional syllogisms were selected from Experiment 1 for use in the current experiment. Eight Single-model Categorical syllogisms and eight Difficult Conditional syllogisms were created for this experiment. Syllogisms were comprised of statements that were true or false by definition, with the exception of the following syllogisms: Categorical, Multi- and Single-model syllogisms with Believable Premises and a Valid, Unbelievable Conclusion; Categorical, Multi-model syllogisms with Unbelievable Premises and an Invalid, Unbelievable Conclusion; and Conditional, Single-model and Difficult syllogisms with Believable Premises and a Valid, Unbelievable Conclusion. Twenty undergraduates who did not participate in this experiment rated the believability of the statements used within the syllogisms in an online pilot experiment (these data are presented in Appendix B).

Procedure

All participants were tested individually. Syllogisms were presented to participants on a computer using E-Prime v1.2 (Schneider, Eschman, & Zuccolotto, 2001) software which recorded participants' responses and response latencies to each syllogism (from the onset of the problem on the screen to the participant's key press). Syllogisms and instructions were presented in black, 18 point Courier New font against a white background. The following instructions were presented to participants on the computer screen and verbally by the researcher:

> This is an experiment to test people's reasoning ability. This task contains 32 problems. Each problem contains two statements followed by a conclusion. You are asked to determine if the conclusions may be logically deduced from the statements. If you judge that the conclusion necessarily follows from the statements (i.e., that it is a valid conclusion based on the information given in the two statements), you should answer by pressing the 'z' key. Otherwise, press the 'm' key to indicate that the conclusion is invalid. Please take your time and ensure that you have the right answer before answering.

Participants had the opportunity to ask the experimenter any questions about the task after hearing the instructions. Participants then completed two practice syllogisms and the 32 syllogisms created for this experiment. Each syllogism was presented individually and in a different random order for each participant. Syllogisms remained on the screen until the participant made a response. When the participant pressed a key indicating their judgment about the conclusion, the next syllogism appeared on the screen after a 1000 millisecond delay. After responding to all 32 syllogisms, participants were debriefed.

Results

Response Data

Proportion of conclusion endorsement is depicted for Conditional Syllogisms with Believable and Unbelievable Premises and Categorical Syllogisms with Believable and Unbelievable Premises in Figures 5 and 6, respectively. Endorsement rates were submitted to a 2 (Type: Conditional, Categorical) x 2 (Models: Single, Multi-model) x 2 (Premise Believability:

Conditional Syllogisms with Believable Premises

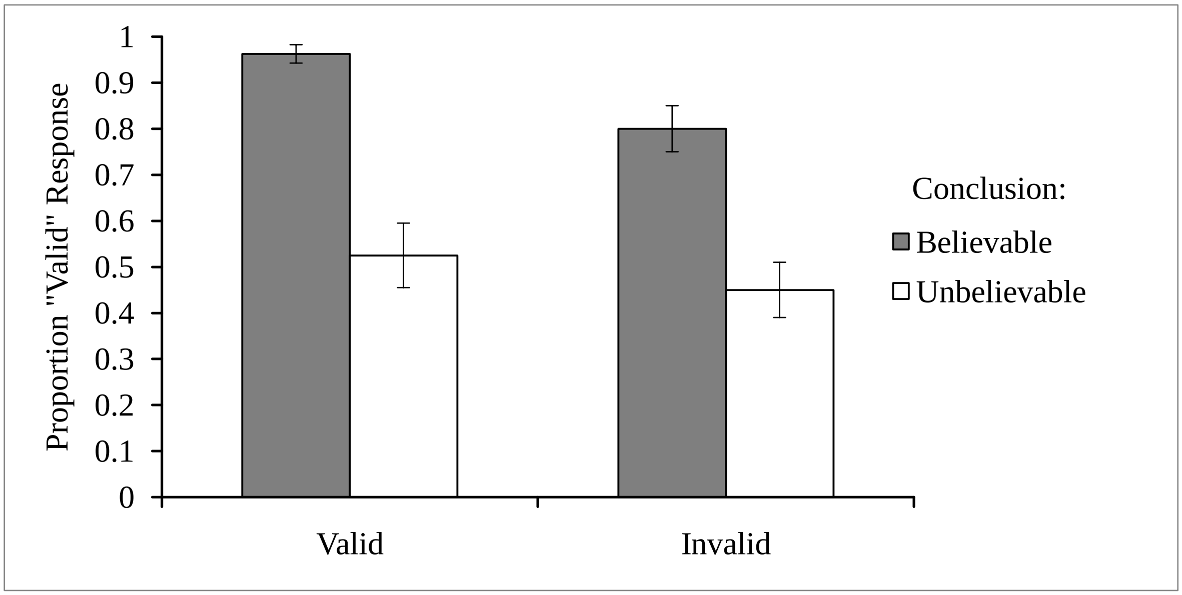

Conditional Syllogisms with Unbelievable Premises

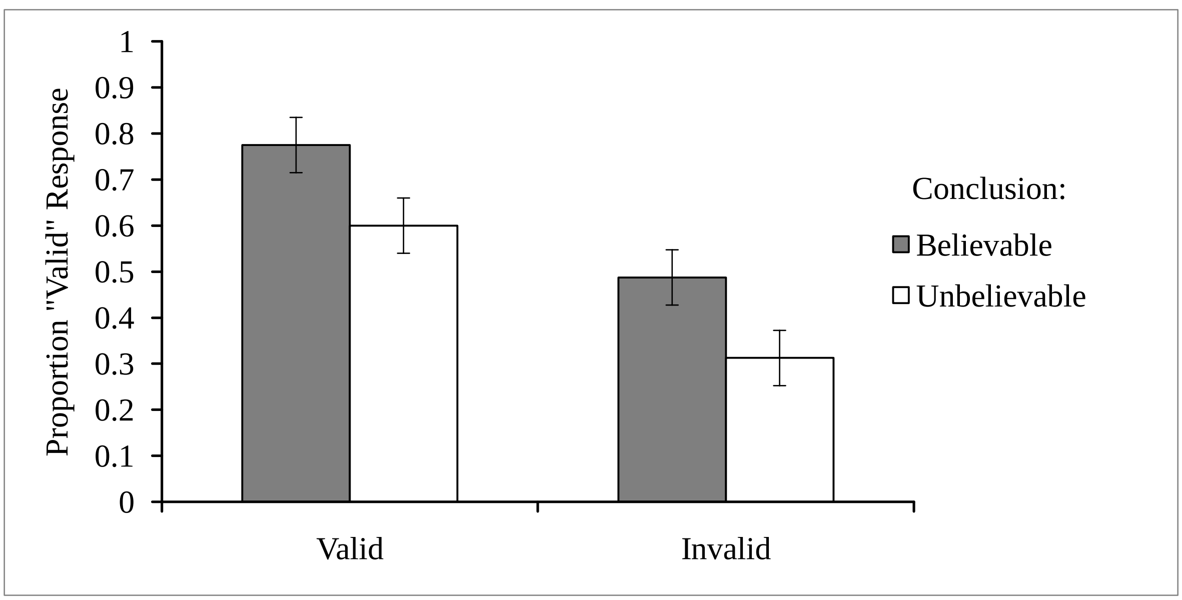

Figure 5. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises.

Categorical Syllogisms with Believable Premises

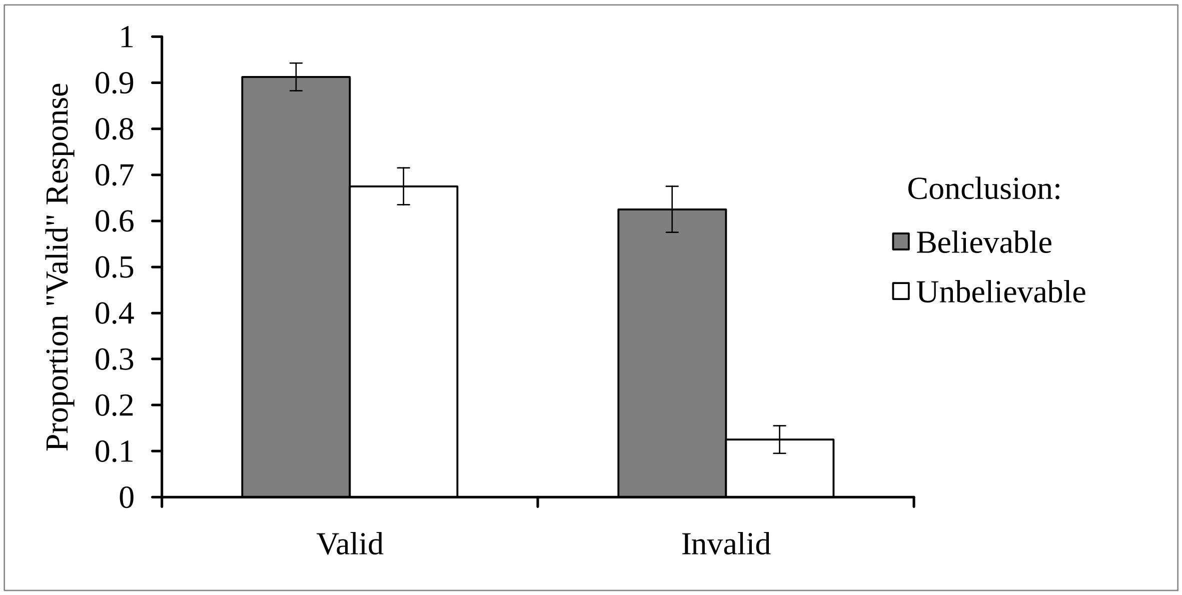

Categorical Syllogisms with Unbelievable Premises

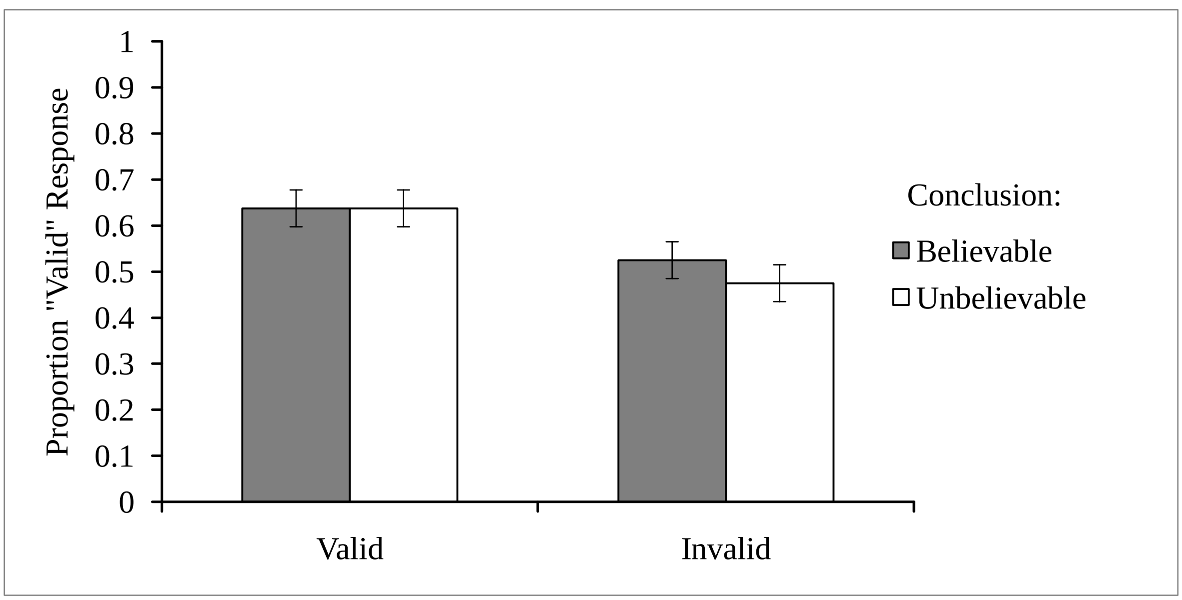

Figure 6. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises.

Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) x 2 (Models: Multiple, Single) repeated measures ANOVA.

Before other results are addressed, general effects of the number of models will be reported to ensure that this manipulation was effective. Then, as in Experiment 1 (and for all subsequent experiments), results will be discussed first in terms of overall belief bias. Second, the three-way interaction between Syllogism Type, Premise Believability, and Conclusion Believability will be examined to determine whether effects of Conclusion Believability vary as a function of Syllogism Type and Premise Believability. Next, the interaction between Syllogism Type and Premise Believability will be discussed to determine whether participants endorsed more conclusions following from Believable than Unbelievable Premises for Conditional compared to Categorical Syllogisms. Other results can be found in Appendix C.

Number of Models Manipulation. There was a Models x Conclusion Validity interaction, $F(1,39) = 5.21$, $MSE = .117$, $p < .05$, $\eta^2 = .118$, indicating that effects of Conclusion Validity were greater for one-model than for multi-model syllogisms. In other words, participants were less successful at distinguishing valid from invalid conclusion for multi-model than for single model syllogisms. There was no Type x Models x Conclusion Validity interaction, $F(1,39) = 0.63$, $MSE = .125$, $p = .408$, $\eta^2 = .018$, indicating that effects of Conclusion Validity were weaker for multi-model syllogisms for both Categorical and Conditional Syllogisms, although difficulty was manipulated for conditional syllogisms in lieu of number of models. Thus, difficulty was successfully manipulated and was manipulated equally across syllogism type.

Overall Belief Bias Effect. Overall, the main effects typically observed in studies of the belief bias were found: There were main effects of Conclusion Believability, $F(1,39) = 54.39$, *MSE* = .341, $p < .001$, $\eta^2 = .582$, and Validity, $F(1,39) = 37.23$, *MSE* = .498, $p < .001$, $\eta^2 = .488$. The interaction between Conclusion Validity and Believability typical of the belief bias effect was not significant when collapsed across Syllogism Type, $F(1,39) = 1.56$, *MSE* = .162, $p = .218$, $\eta^2 = .039$. However, there was a three-way interaction between Type, Conclusion Validity, and Conclusion Believability, $F(1,39) = 6.55$, *MSE* = .122, $p < .05$, $\eta^2 = .114$. To unpack this three-way interaction, endorsement rates for Categorical and Conditional Syllogisms were submitted to separate 2 (Conclusion Validity) x 2 (Conclusion Believability) repeated measures ANOVAs. The typical interaction between Conclusion Validity and Believability was significant for Categorical Syllogisms, $F(1,39) = 5.28$, *MSE* = .740, $p < .05$, $\eta^2 = .119$, but was not significant for Conditional Syllogisms, $F(1,39) = 0.77$, *MSE* = .396, $p = .385$, $\eta^2 = .019$.

Effects of Conclusion Believability as a Function of Syllogism Type and Premise Believability. Crucially, there was a three-way interaction between Type, Premise Believability, and Conclusion Believability, $F(1,39) = 24.25$, *MSE* = .167, $p < .001$, $\eta^2 = .383$. Separate two-way repeated measures ANOVAs computed for Categorical and Conditional syllogisms revealed that there was a significant interaction between Premise Believability and Conclusion Believability for Categorical Syllogisms, $F(1,39) = 33.76$, *MSE* = .140, $p < .001$, $\eta^2 = .464$. Paired-sample *t*-tests were computed to address this interaction. Importantly, Conclusion Believability *only* affected conclusion endorsement when premises were believable, $t(39) = 7.17$, $p < .001$, but not when they were unbelievable, $t(39) = 0.46$, $p = .648$. For Conditional Syllogisms, Premise Believability also interacted with Conclusion Believability, $F(1,39) = 12.93$, *MSE* = .148, $p = .001$, $\eta^2 = .249$, however the form of this interaction was such that Conclusion Believability had a greater effect on endorsement rates when premises were Believable, $t(39) = 7.34$, $p < .001$ rather than when they were Unbelievable, $t(39) = 3.82$, $p < .001$.

The four-way interaction between Type, Premise Believability, Conclusion Believability, and Models (i.e., number of models required for categorical syllogisms) was non-significant, $F(1,39) = 1.35$, $MSE = .113$, $p = .252$, $\eta^2 = .034$, indicating that the pattern of conclusion believability results described above does not vary as a function of the number of possible models required for the syllogism.

Conclusion Endorsement for Categorical and Conditional Syllogisms as a Function of Premise Believability. Contrary to the results of Experiment 1, Syllogism Type interacted with Premise Believability, $F(1,39) = 19.02$, $MSE = .066$, $p < .001$, $\eta^2 = .328$. For Conditional Syllogisms, more conclusions following from Believable Premises were endorsed than those following from Unbelievable Premises, $t(39) = 3.65$, $p = .001$. There was no overall effect of Premise Believability on conclusion endorsement for Categorical Syllogisms, $t(39) = 0.41$, $p = .685$. Again, this interaction was not qualified by the number of models required, $F(1,39) = 0.70$, $MSE = .018$, $p = .408$, $\eta^2 = .018$.

Response Latency Data

Response Latencies[1] as a function of Conclusion Validity and Believability are shown for Conditional Syllogisms with Believable and Unbelievable Premises and Categorical Syllogisms with Believable and Unbelievable Premises in Figures 7 and 8, respectively.

Response Latencies were submitted to a 2 (Type: Conditional, Categorical) x 2 (Models: Single, Multi-model) 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVA.

---

[1] **An outlier analysis was carried out and response latencies that were outside of three standard deviations of the mean for each cell were removed. Because the pattern of data remained identical to when all scores were included, and to remain consistent with other experiments examining response latencies and deductive reasoning (e.g., Thompson et al., 2003), all scores were included for the analyses of this and subsequent experiments.**

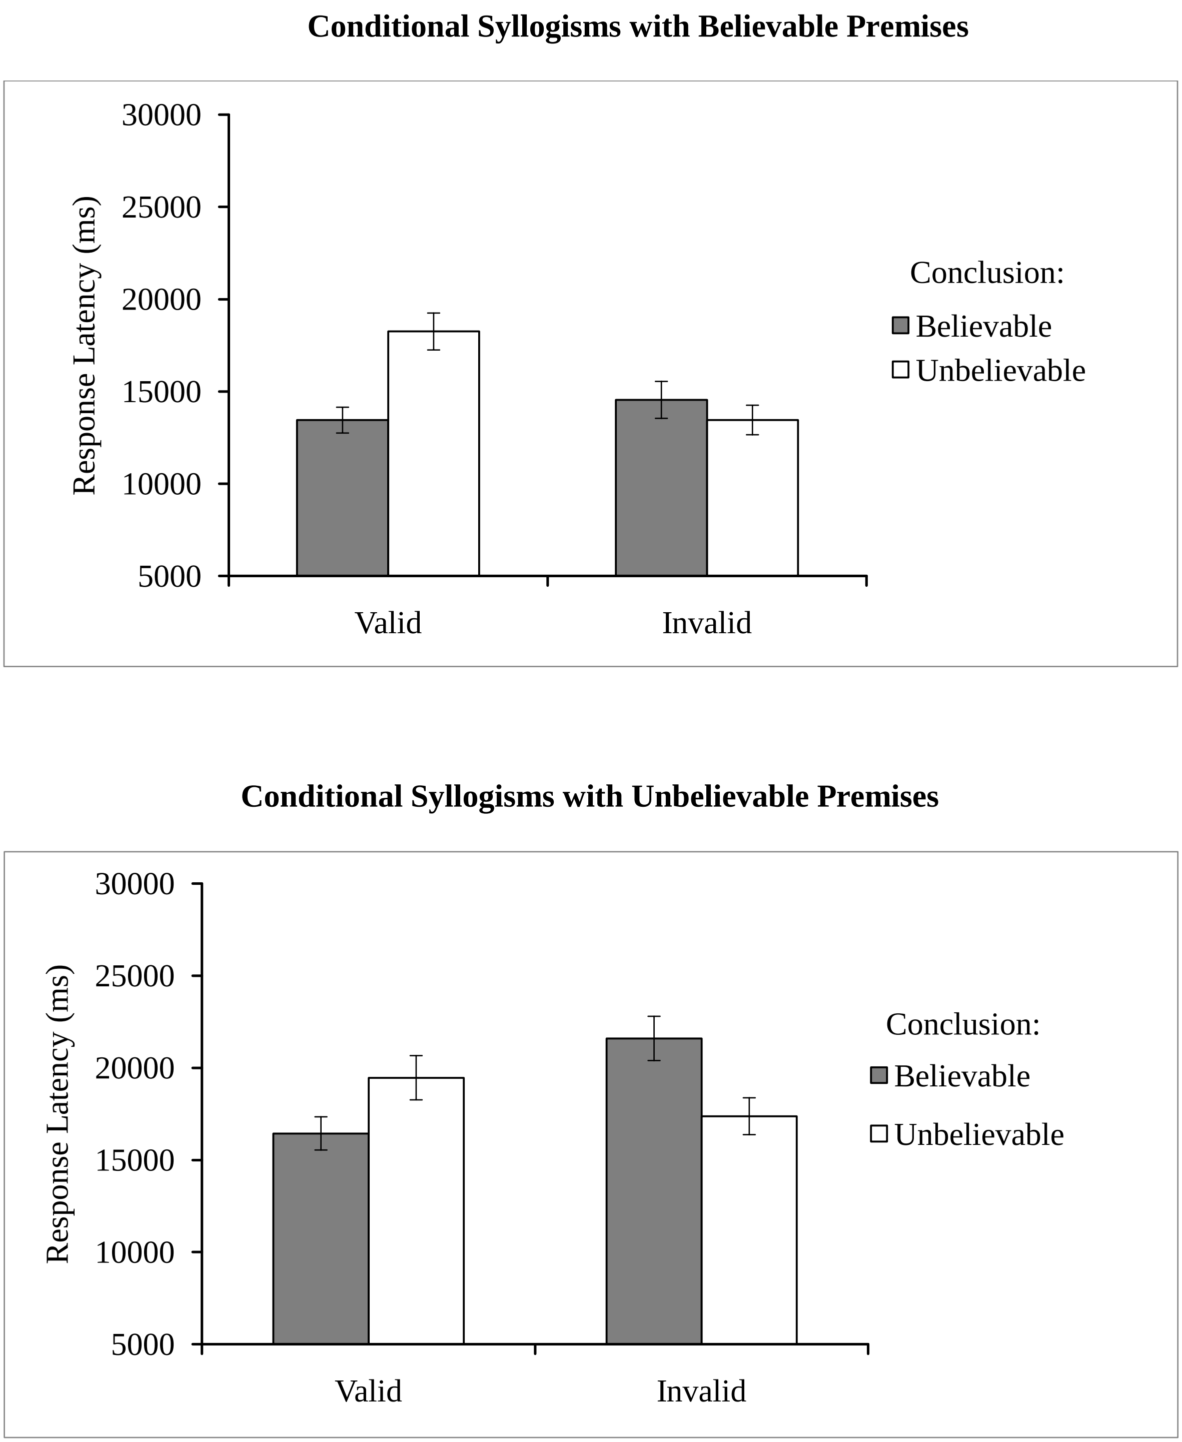

Figure 7. Response latencies as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises.

Categorical Syllogisms with Believable Premises

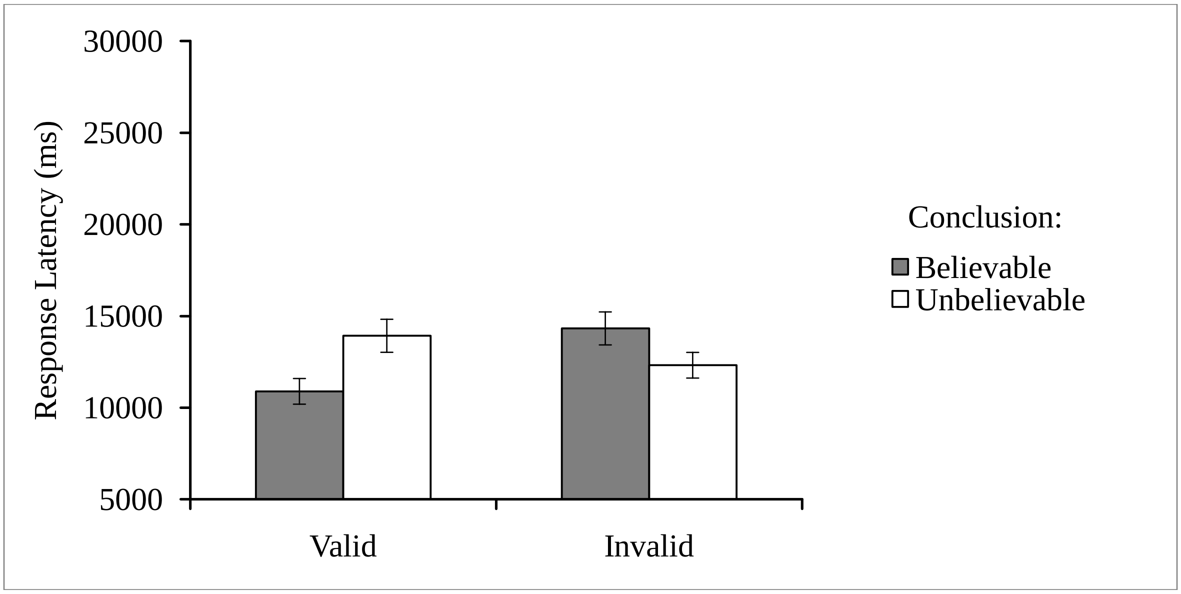

Categorical Syllogisms with Unbelievable Premises

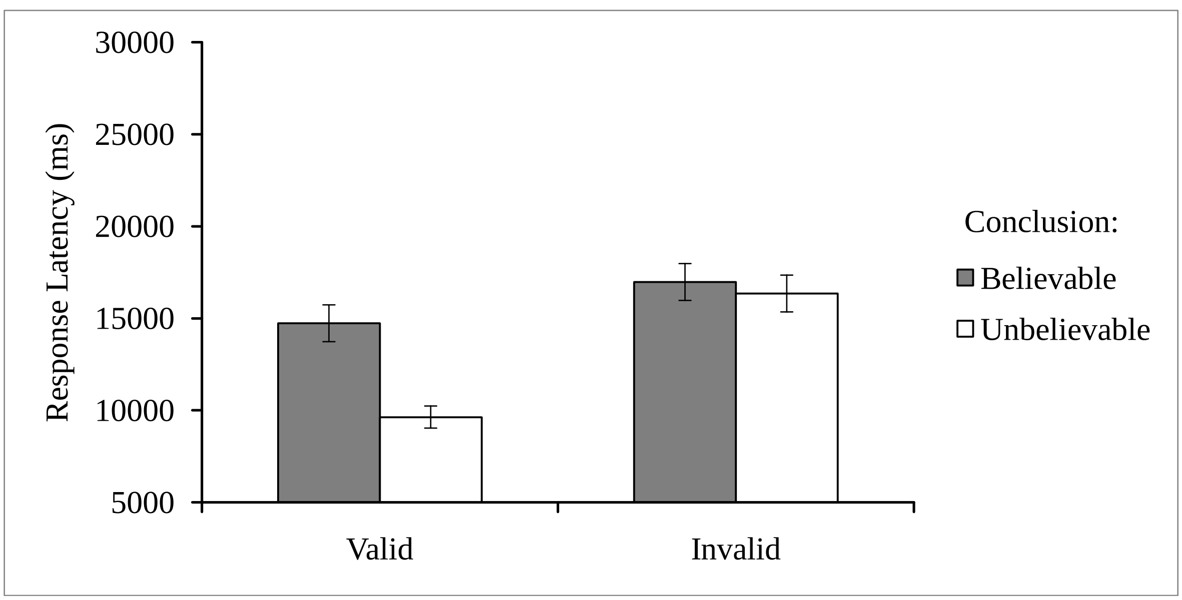

Figure 8. Response latencies as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises.

For this and all subsequent experiments, response latency data will be discussed, first, in terms of interactions between Conclusion Validity and Conclusion Believability and other factors to determine for which syllogisms Conclusion Believability affected response latency. Second, the two-way Syllogism Type by Premise Believability interaction will be examined to determine whether participants spent longer reasoning about Conditional Syllogisms with Believable than Unbelievable Premises compared to Categorical Syllogisms. Other effects are reported in Appendix C.

Effects of Conclusion Believability. There was a significant three-way interaction between Syllogism Type, Conclusion Validity, and Conclusion Believability, $F(1,39) = 4.66$, $MSE = 8.753 \times 10^7$, $p < .05$, $\eta^2 = .107$, and a four-way interaction between Syllogism Type, Premise Believability, Conclusion Validity, and Conclusion Believability, $F(1,39) = 12.09$, $MSE = 1.103 \times 10^8$, $p = .001$, $\eta^2 = .237$. To address these interactions, separate 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVAs were conducted with Categorical and Conditional response latencies. For Categorical Syllogisms, there was a highly significant interaction between Premise Believability, Conclusion Validity, and Conclusion Believability, $F(1,39) = 19.92$, $MSE = 8.831 \times 10^7$, $p < .001$, $\eta^2 = .338$. Separate 2 (Conclusion Validity: Valid, Invalid) x 2(Conclusion Believability: Believable, Unbelievable) repeated measures ANOVAs were conducted for Believable and Unbelievable Premises. Here, there was a significant interaction between Conclusion Validity and Conclusion Believability when premises were believable, $F(1,39) = 20.45$, $MSE = 2.995 \times 10^7$, $p < .001$, $\eta^2 = .334$, such that response latencies were greater when the believability and validity of the conclusion conflicted. This interaction was also significant when premises were unbelievable, $F(1,39) = 4.93$, $MSE =$

5.996 x10^7, p = <.05, η^2 = .112. The nature of this interaction, however, differed from that of believable premises. Specifically, for unbelievable premises, participants were actually *faster* to reason about conflict problems.

For Conditional Syllogisms, Conclusion Validity and Believability interacted, F(1,39) = 14.43, MSE = 9.040 x10^7, p < .001, η^2 = .270. Here, participants spent longer reasoning about problems in which the Validity and Believability of the conclusion conflicted. However, unlike the Categorical Syllogisms, there was no significant three-way interaction between Premise Believability, Conclusion Validity, and Conclusion Believability, F(1,39) = 1.08, MSE = 8.735 x10^7, p =.306, η^2 = .027. There was no interaction between Number of Models, Type, Premise Believability, Conclusion Validity, and Conclusion Believability, F(1,39) = 0.07, MSE = 1.019 x10^8, p =.798, η^2 = .002, indicating that the effects discussed above did not vary as a function of the number of models.

Effects of Premise Believability for Conditional Syllogisms compared to Categorical Syllogisms. There was a significant interaction between Syllogism Type and Premise Believability, F(1,39) = 5.66, MSE = 1.334 x10^8, p <.05, η^2 = .127. Paired-samples t-tests indicated that participants took significantly longer to reason about Conditional Syllogisms with Unbelievable Premises than with Believable Premises, t(39) = 4.25, p <.001. The difference in response latencies between Believable and Unbelievable Premises for Categorical Syllogisms, however, fell short of significance, t(39) = 1.89, p = .066. There was no three-way interaction between Number of Models, Type, and Premise Believability, F(1,39) = 0.27, MSE = 6.160 x10^7, p =.602, η^2 = .007, indicating that the effects outlined above were not affected by number of models.

Discussion

In general, the predictions of the theory of premise evaluation outlined in the introduction were supported. For conditional syllogisms, participants took longer to reason about syllogisms with believable than with unbelievable premises. Further, for both believable and unbelievable premises, there is evidence that participants took conclusion believability into account. Specifically, for both premise types, Conclusion Validity interacted with Conclusion Believability, thus replicating and extending the findings of Ball et al. (2006) that participants spend longer on problems in which the validity and believability of the conclusion conflicted. These results are consistent with the theory that reasoners initially assume that conditional premises are true, and reason about the conclusion under this assumption. After, or otherwise independent of, reasoning about the conclusion, participants may discredit unbelievable premises, which explains the extra time spent on problems with unbelievable premises.

For categorical premises, the trend for participants to spend longer on syllogisms with unbelievable premises fell short of significance. Interestingly, there was no effect of conclusion believability for syllogisms with unbelievable premises in the predicted direction. Indeed, conclusion validity and believability interacted for both believable and unbelievable premises, but the nature of this interaction differed. For believable premises, the interaction resembled that of conditional syllogisms: participants spent longer reasoning about syllogisms for which the believability and validity of the conclusion conflicted. For unbelievable premises, however, participants were *faster* to respond to these conflict problems, particularly those with valid, unbelievable conclusions. If participants took the believability of the conclusion into account, they would not be expected to be faster on conflict problems, given that responses elicited by the validity and believability of the conclusion are in opposition. Thus, the finding that participants were faster on conflict problems is strong evidence that conclusion believability was not accounted for during the reasoning process when premises were unbelievable. In fact, it may be the case that conclusion believability was ignored in such a way that *facilitated* reasoning on conflict problems.

Overall, these results suggest that, for categorical syllogisms, participants evaluate the believability of premises when they first encounter them. Because participants do not accept unbelievable premises and then disbelieve them (that is, there is one step of evaluation for both believable and unbelievable premises), they do not spend significantly longer reasoning about problems with unbelievable premises. Further, when premises are unbelievable, participants do not spend more time reasoning about conflict problems, as would be predicted if conclusion believability was taken into account. This suggests that there is something about unbelievable premises that triggers participants to ignore believability information in the rest of the problem; perhaps it is the case that unbelievable information in the premises acts as a signal that believability information in the rest of the problem is unreliable and should be ignored.

Crucially, the endorsement rates from the current experiment replicated those of Experiment 1. Here, conditional syllogisms elicited the main effects typical of belief bias regardless of premise believability. That is, participants endorsed more valid than invalid conclusions and more believable than unbelievable conclusions. The typically found interaction between conclusion validity and believability, however, was not significant. Unlike in Experiment 1, the current experiment replicated the results of Thompson (1996): Participants were more likely to judge that a conclusion was valid when it was preceded by believable than by unbelievable premises. This finding is also in line with previous studies that have found that reasoners are less likely to endorse, or are less confident in, conclusions that follow from uncertain conditionals (e.g., George, 1995, 1997, 1999). Again, these results suggest that believable and unbelievable premises are initially treated in the same fashion, however, there is an overall bias for participants to respond “valid” for conclusions following from believable premises and “invalid” for conclusions following from unbelievable premises. If reasoners evaluate premises after the reasoning process, it may be the case that they are then biased to endorse conclusions following from sound premises.

For categorical syllogisms, the typical belief bias results were only found for believable premises. When premises are unbelievable, the belief bias effect is suppressed, and only conclusion validity affects conclusion endorsement. Again, this suggests that reasoners somehow disregard believability information in the conclusion when premises are not believable.

The effects described above were not dependent upon the number of models required for syllogisms. Thus, effects of conclusion believability were absent for categorical syllogisms with unbelievable premises for both single and multi-model syllogisms. That the believability of the premises influenced belief bias when only one model of the premises was required speaks to how powerful premise believability can be in affecting how individuals reason about even simple problems. Thus far, there is evidence that reasoners treat the premises of categorical and conditional syllogisms in fundamentally different ways. Conditional premises seem to be accepted as true prior to the reasoning process, whereas categorical premises seem to be evaluated prior to reasoning about the conclusion. Experiments 3 and 4 will examine *why* this is the case.

Experiment 3

Experiments 1 and 2 provide strong evidence for differential effects of premise believability on conditional and categorical syllogisms. Following Experiment 1, I had proposed a theory for why this is the case: specifically, that reasoners initially assume that conditional statements are true, whereas they evaluate the believability of categorical statements as soon as they are encoded. The following three experiments will test various aspects of this explanation.

Most experiments investigating deductive reasoning instruct participants to “assume that premises are true.” This is for good reason: Although logically a valid conclusion can follow from unbelievable premises, pragmatically this is not the case (Evans, 2005; Politzer & Bourmaud, 2002). Experiments 1 and 2 omitted this instruction so that effects of premise believability could be fully explored. Even when participants were not instructed to assume that premises were true, there was still some evidence of logical reasoning: Conclusion validity always affected reasoning such that participants were more likely to endorse valid than invalid conclusions.

The theory is that reasoners assume by default that conditional premises are true prior to reasoning about the conclusion, but evaluate the truthfulness of categorical premises before reasoning about the conclusion. If this is indeed the case, then emphasizing the importance of assuming that premises are true to participants should remove the evaluative step prior to conclusion reasoning for categorical syllogisms, thus making categorical syllogisms resemble conditional syllogisms in both conclusion endorsement rates and response latency patterns. To determine whether conditional and categorical syllogisms differ in the point at which premises are evaluated, an experimental group of participants was provided with augmented instructions that emphasized that they should always assume premises to be true. Relative to the control group, it was predicted, first, that there should be no differences with respect to conditional syllogisms, because participants presumably already assume that premise are true; second, that categorical syllogisms should show belief bias whether premises are believable *or* unbelievable; and finally, that there should be evidence of conclusion believability affecting response latencies (in the form of longer response latencies for conflict problems) for categorical syllogisms with both believable and unbelievable premises.

Method

Participants

Participants were 85 students (39 male, 46 female, mean age = 19.48 yrs, SD = 1.89 yrs) enrolled in psychology courses at the University of Waterloo who participated in exchange for course credit. None had reported having taken a course in logic.

Materials

The same 32 syllogisms used in Experiment 2 were used in the current experiment. Again, syllogisms were either Conditional or Categorical in Type, had either Believable or Unbelievable Premises, Valid or Invalid Conclusions, and Believable or Unbelievable Conclusions.

Procedure

Participants were assigned to one of two groups. Participants in the control group (n = 40) received instructions identical to those in Experiment 2. Participants in the experimental group (n = 45) received instructions that emphasized the importance of assuming premises to be true (note the text in bold font below; this sentence was not in bold in the actually presented instructions):

> This is an experiment to test people's reasoning ability. This task contains 32 problems. Each problem contains two statements followed by a conclusion. You are asked to

> determine if the conclusions may be logically deduced from the statements. **It is very important that you assume that the two statements are true**. If you judge that the conclusion necessarily follows from the statements (i.e., that it is a valid conclusion based on the information given in the two statements), you should answer by pressing the 'z' key. Otherwise, press the 'm' key to indicate that the conclusion is invalid. Please take your time and ensure that you have the right answer before answering.

Other than this instructional manipulation for the experimental group, participants in both groups were tested in a procedure that was identical to that in Experiment 2.

Results

Response Data

The proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability are shown for Conditional Syllogisms with Believable and Unbelievable Premises and Categorical Syllogisms with Believable and Unbelievable Premises in Figures 9 and 10, respectively, for the Control Group and Figures 11 and 12, respectively, for the Experimental Group. Conclusion endorsement rates were submitted to a 2 (Type: Categorical, Conditional) x 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability) repeated measures ANOVA with Condition (Control, Experimental) as a between subjects variable. Because the number of models required for syllogisms did not affect any findings of interest in Experiment 2, number of models was not considered as a factor here or in any of the following experiments.

Overall Belief Bias Effect. Overall, the typical belief bias effect was found: There were main effects of Conclusion Believability, $F(1,83) = 74.84$, $MSE = .814$, $p < .001$, $\eta^2 = .474$, and Validity, $F(1,83) = 102.15$, $MSE = 1.113$, $p < .001$, $\eta^2 = .552$, and an interaction between Conclusion Validity and Believability, $F(1,83) = 11.68$, $MSE = .351$, $p = .001$, $\eta^2 = .123$.

Effects of Conclusion Believability as a Function of Syllogism Type and Premise Believability. There was a Type x Premise Believability x Conclusion Believability interaction,

Conditional Syllogisms with Believable Premises – Control Group

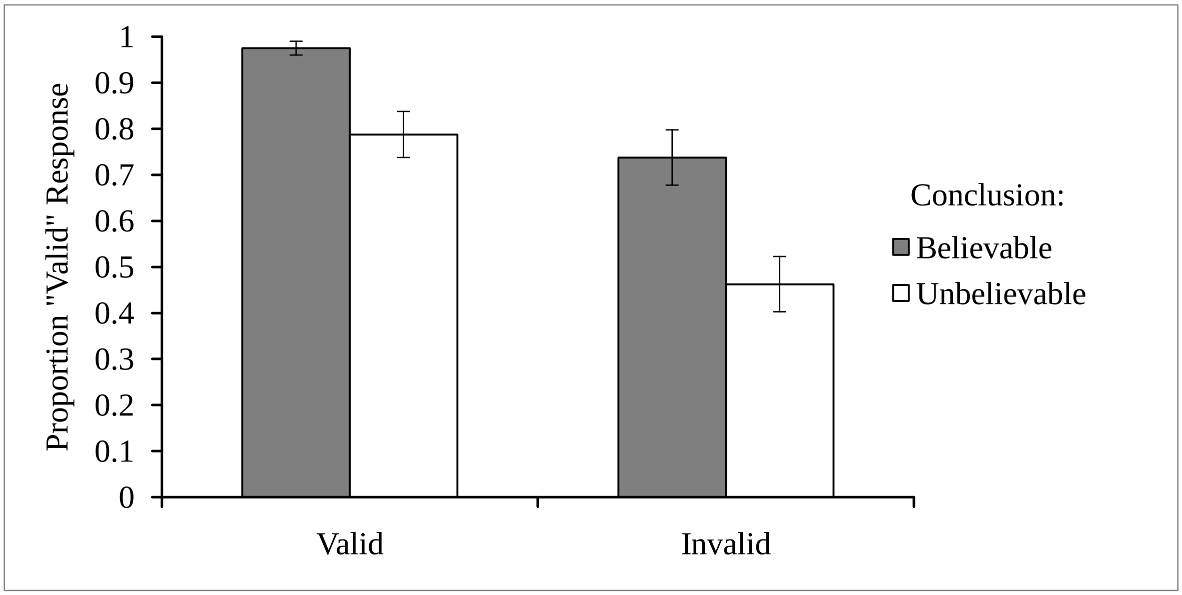

Conditional Syllogisms with Unbelievable Premises – Control Group

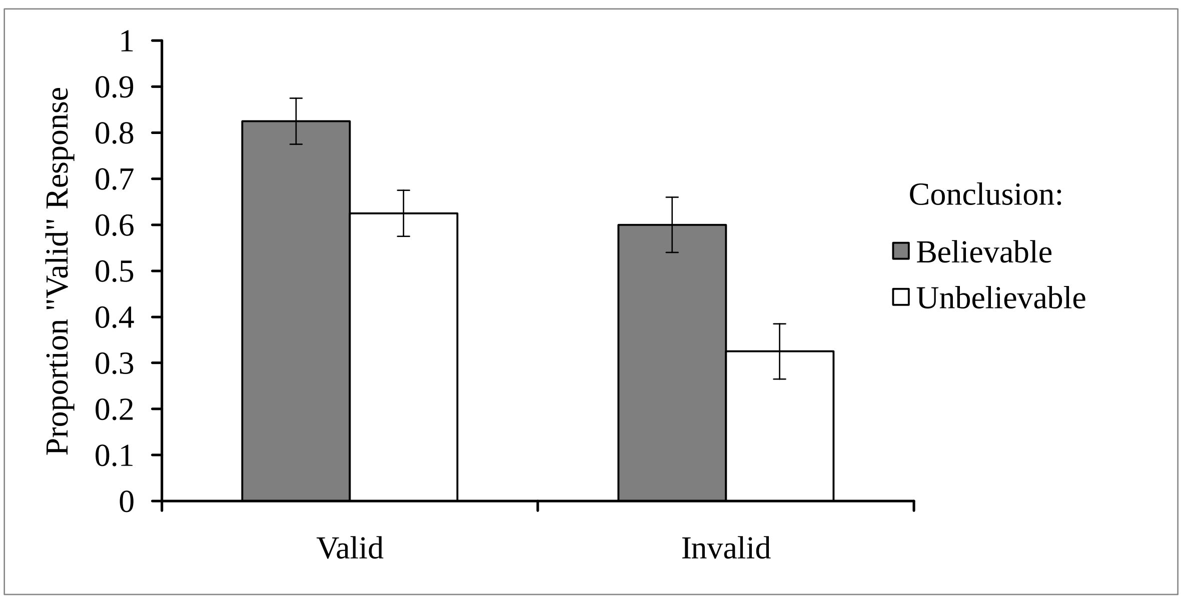

Figure 9. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises *for the* Control Group.

Categorical Syllogisms with Believable Premises – Control Group

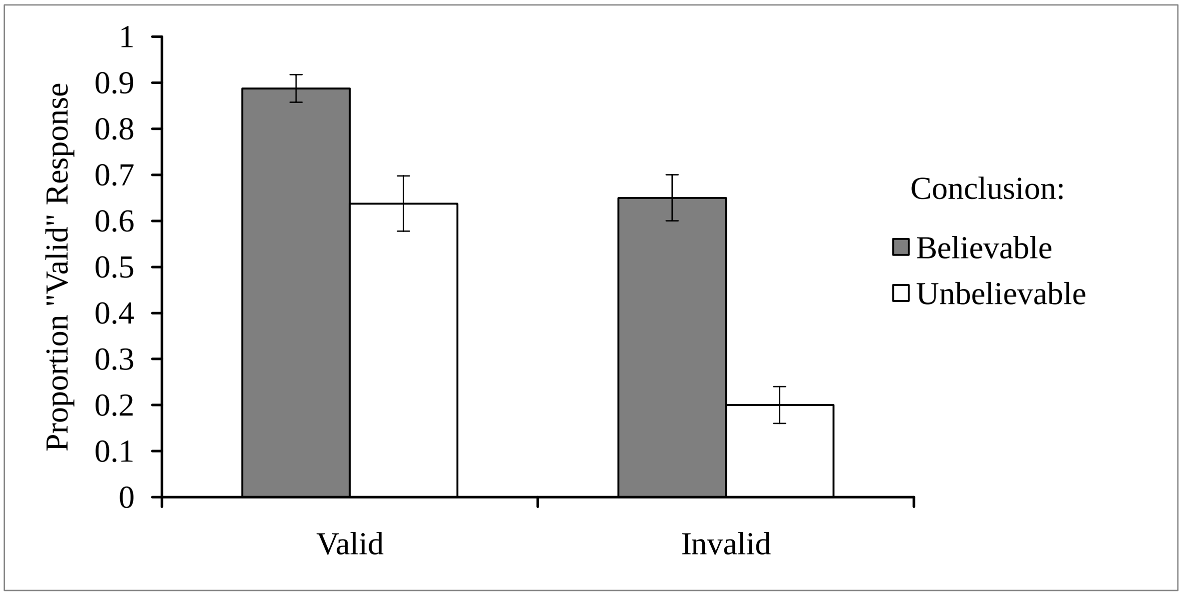

Categorical Syllogisms with Unbelievable Premises – Control Group

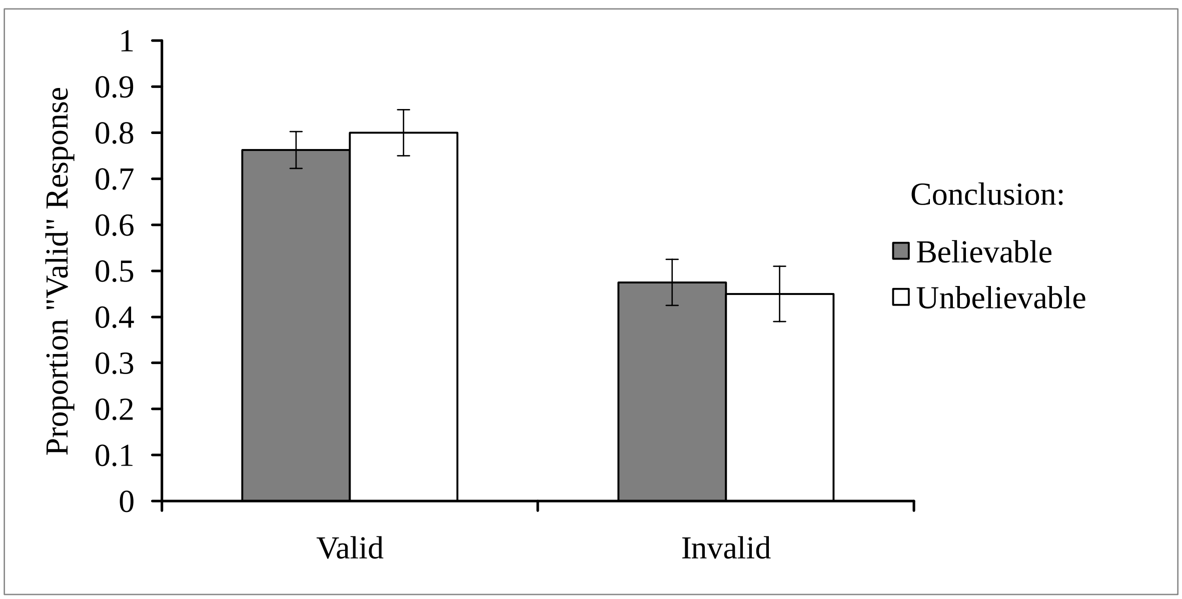

Figure 10. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises *for the* Control Group.

Conditional Syllogisms with Believable Premises – Experimental Group

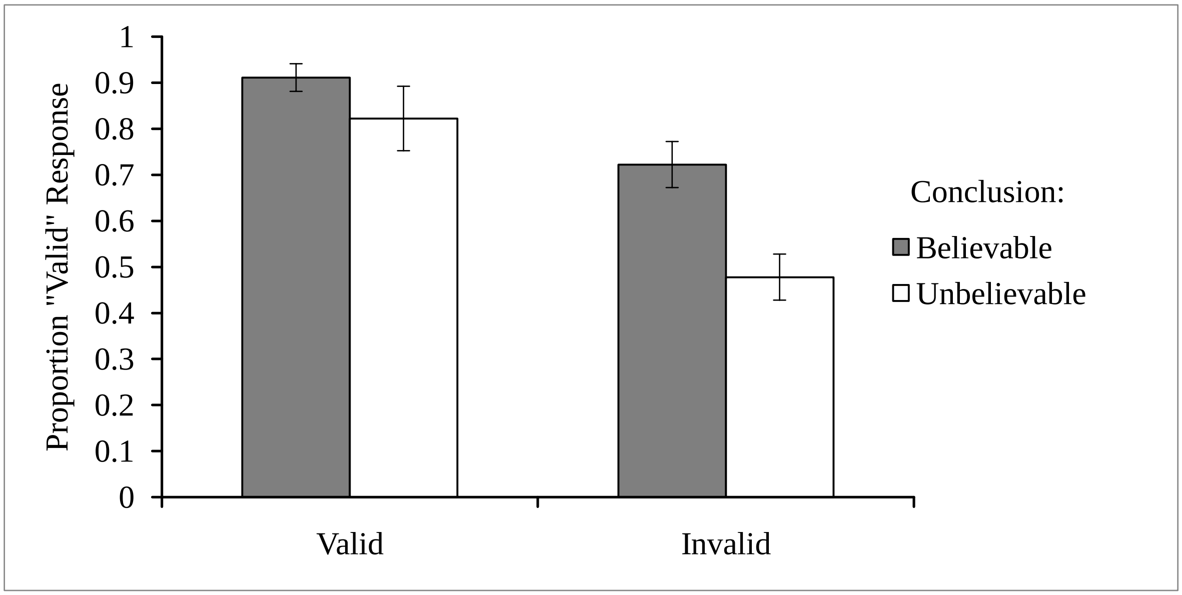

Conditional Syllogisms with Unbelievable Premises – Experimental Group

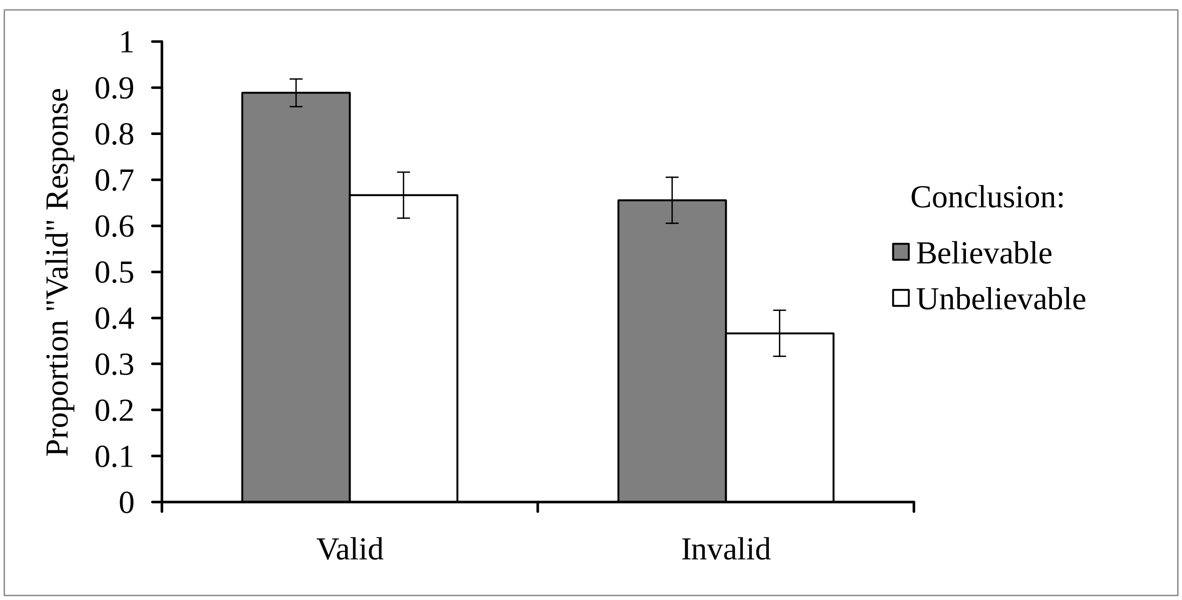

Figure 11. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises *for the* Experimental Group.

Categorical Syllogisms with Believable Premises – Experimental Group

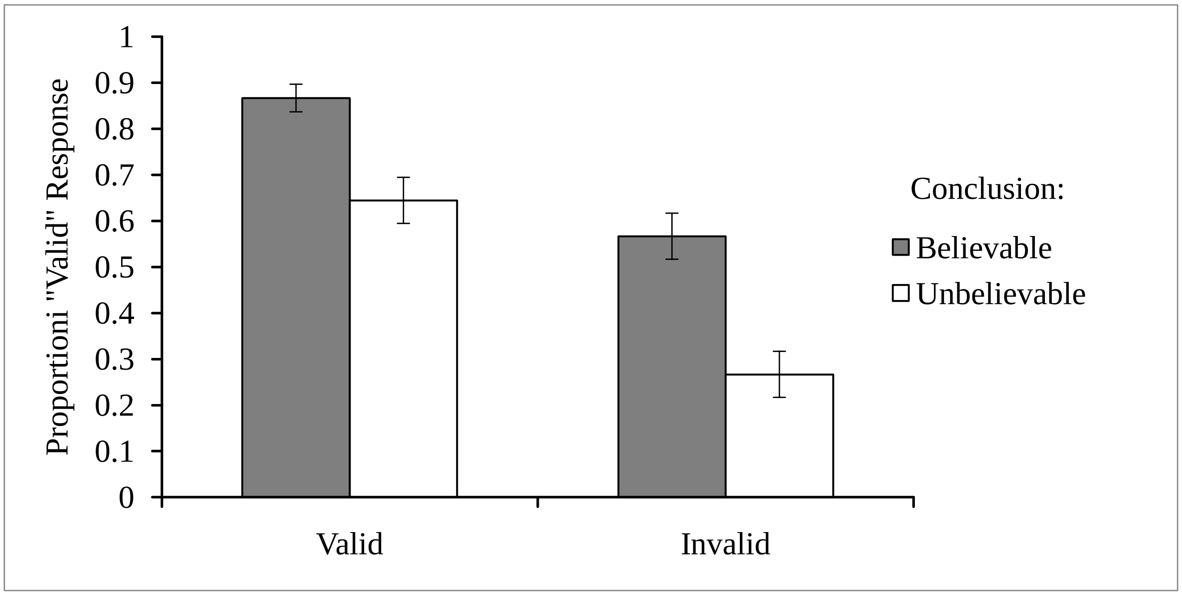

Categorical Syllogisms with Unbelievable Premises – Experimental Group

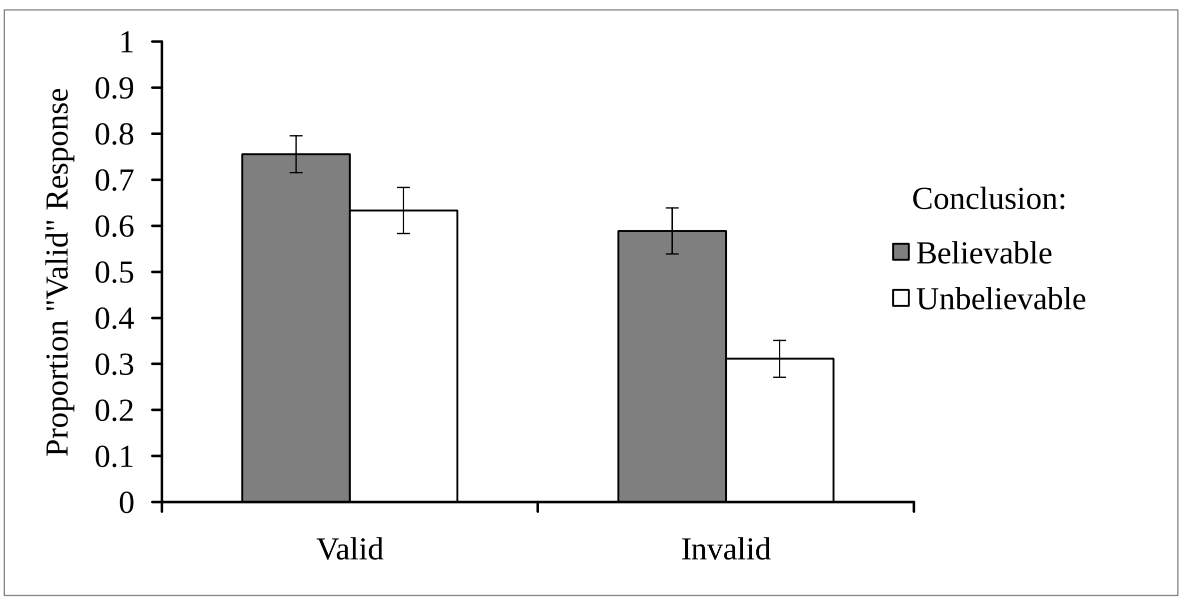

Figure 12. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises *for the* Experimental Group.

$F(1,83) = 21.43$, $MSE = .260$, $p < .001$, $\eta^2 = .205$, which was further qualified by a marginally significant interaction between Type, Premise Believability, Conclusions Believability, and Condition, $F(1,83) = 3.69$, $MSE = .260$, $p = .058$, $\eta^2 = .043$. To further explore this interaction, endorsement rates from the Control Group and Experimental Group were submitted to separate 2 (Type: Categorical, Conditional) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVAs. For the Control Group, there was a highly significant three-way interaction between Type, Premise Believability, and Conclusion Believability, $F(1,83) = 20.01$, $MSE = .263$, $p < .001$, $\eta^2 = .339$. These analyses revealed that for Categorical Syllogisms, Premise Believability and Conclusion Believability interacted, $F(1,39) = 37.74$, $MSE = .524$, $p < .001$, $\eta^2 = .498$, such that Conclusion Believability played a role only when Premises were Believable, $t(39) = 6.23$, $p <.001$, but not when they were Unbelievable, $t(39) = 0.15$, $p = .884$. There was no such interaction for Conditional Syllogisms, $F(1,39) = 0.01$, $MSE = .673$, $p = .924$, $\eta^2 = <.001$. For the Experimental Group, there was no interaction between Premise Believability and Conclusion Believability for either Categorical, $F(1,44) = 0.94$, $MSE = .718$, $p < .338$, $\eta^2 = .021$, or Conditional, $F(1,44) = 3.78$, $MSE = .337$, $p = .058$, $\eta^2 = .079$, Syllogisms.

Conclusion Endorsement for Categorical and Conditional Syllogisms as a Function of Premise Believability. There was a significant Type x Premise Believability interaction, $F(1,83) = 19.40$, $MSE = .273$, $p < .001$, $\eta^2 = .189$, which was further qualified by a marginally significant interaction between Type, Premise Believability, and Condition, $F(1,83) = 3.10$, $MSE = .273$, $p = .082$, $\eta^2 = .036$. To further explore these effects, endorsement rates for the Experimental and Control groups were submitted to 2 (Type: Conditional, Categorical) x 2 (Premise Believability: Believable, Unbelievable) repeated measures ANOVAs. For the Control group, Type interacted with Premise Believability, $F(1,39) = 17.61$, $MSE = 1.113$, $p < .001$, $\eta^2 = .311$. Paired-samples *t*-tests revealed that participants endorsed more conclusions following from Believable Premises than from Unbelievable Premises for Conditional Syllogisms, $t(39) = 5.56$, $p < .001$. For Categorical Syllogisms, however, there was no difference in conclusion endorsement depending on Premise Believability, $t(39) = 0.92$, $p = .362$. For the Experimental Group, Type and Premise Believability did not interact, $F(1,44) = 3.78$, $MSE = 1.073$, $p = .058$, $\eta^2 = .079$.

Response Latency Data

Response Latencies as a function of Conclusion Validity and Believability are shown for Conditional Syllogisms with Believable and Unbelievable Premises and Categorical Syllogisms with Believable and Unbelievable Premises in Figures 13 and 14, respectively, for the Control Group, and Figures 15 and 16, respectively, for the Experimental Group.

Response latencies were submitted to a 2 (Type: Categorical, Conditional) x 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability) repeated measures ANOVA with Condition (Control, Experimental) as a between subjects variable.

Effects of Conclusion Believability. There was a significant four-way interaction between Type, Premise Believability, Conclusion Validity, and Conclusion Believability, $F(1,83) = 9.41$, $MSE = 1.128 \times 10^8$, $p < .01$, $\eta^2 = .102$. To unpack this four-way interaction, separate 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability) repeated measures ANOVAs were conducted for Categorical and Conditional Syllogisms. For Categorical Syllogisms, there was a significant three-way interaction between Premise Believability, Conclusion Validity, and Conclusion Believability, $F(1,83) = 11.73$, $MSE = 2.260 \times 10^7$, $p = .001$, $\eta^2 = .124$. Two-way ANOVAs revealed that when

Conditional Syllogisms with Believable Premises – Control Group

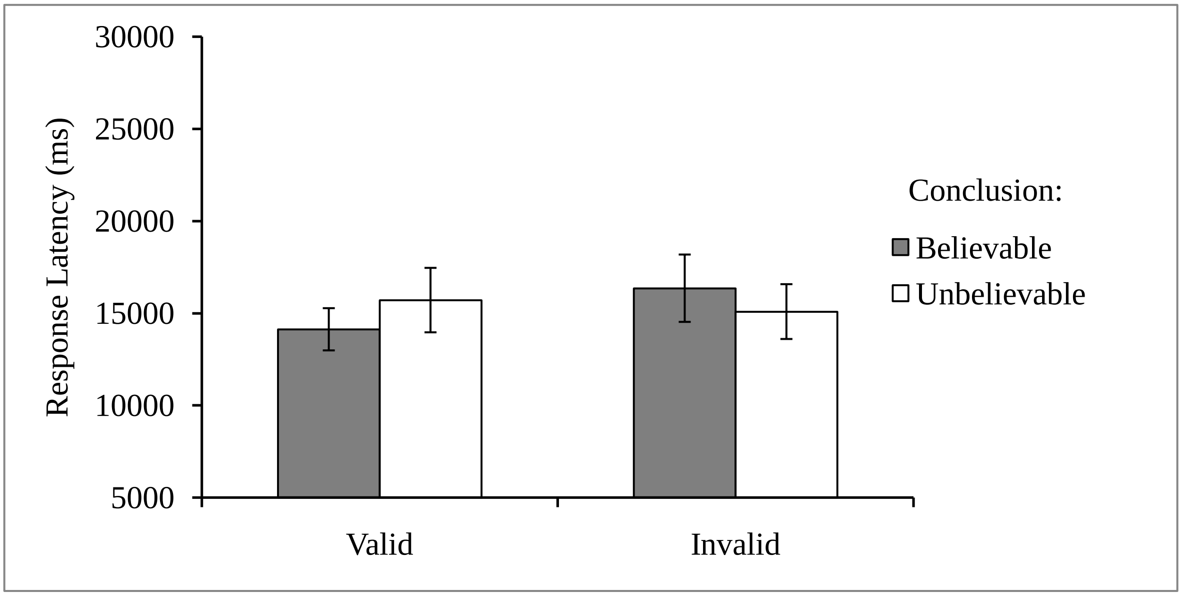

Conditional Syllogisms with Unbelievable Premises – Control Group

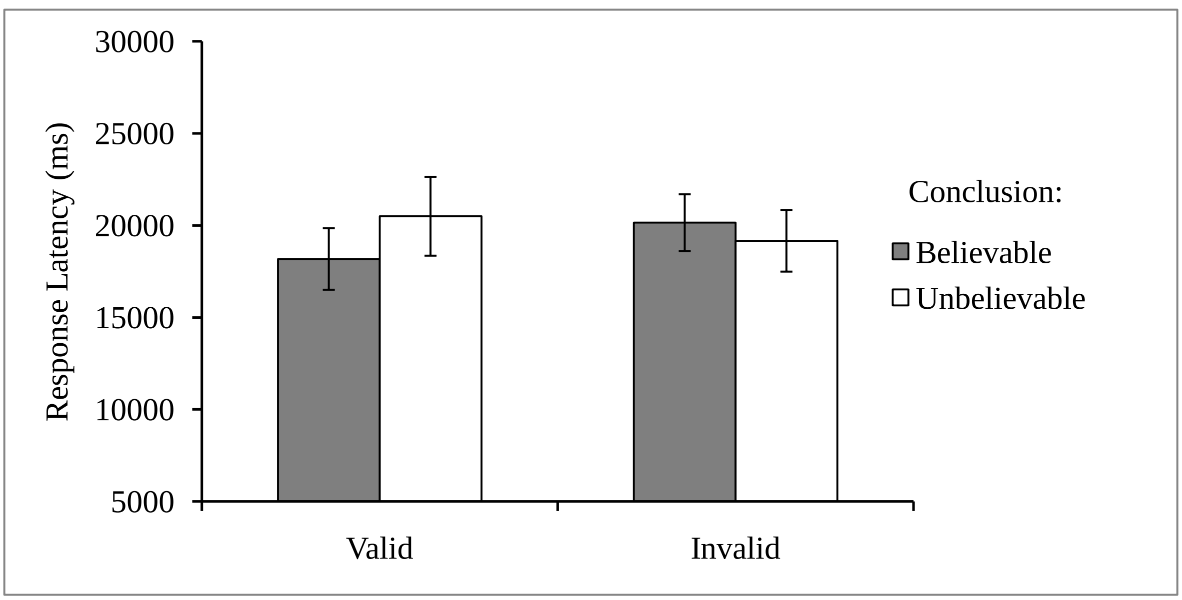

Figure 13. Response latencies as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises *for the* Control Group.

Categorical Syllogisms with Believable Premises – Control Group

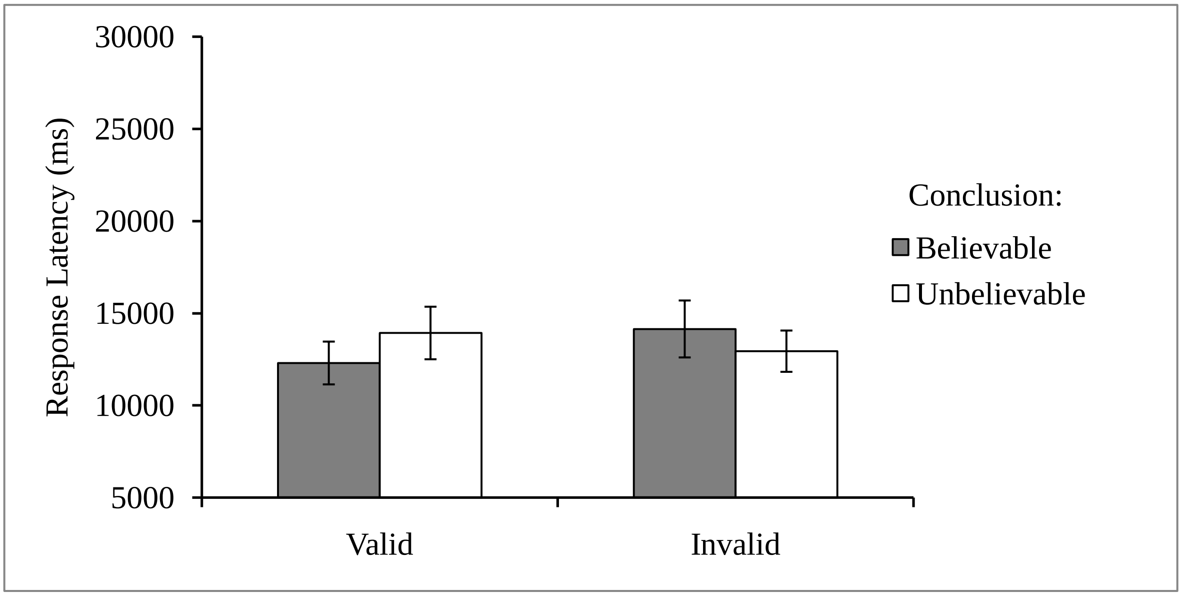

Categorical Syllogisms with Unbelievable Premises – Control Group

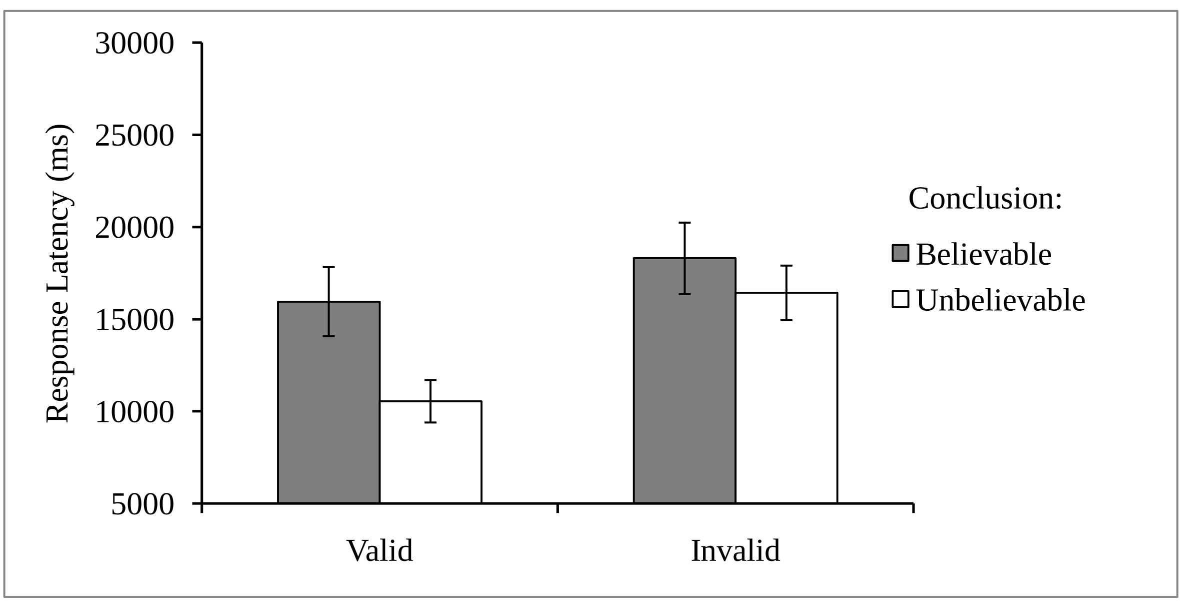

Figure 14. Response latencies as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises *for the* Control Group.

Conditional Syllogisms with Believable Premises – Experimental Group

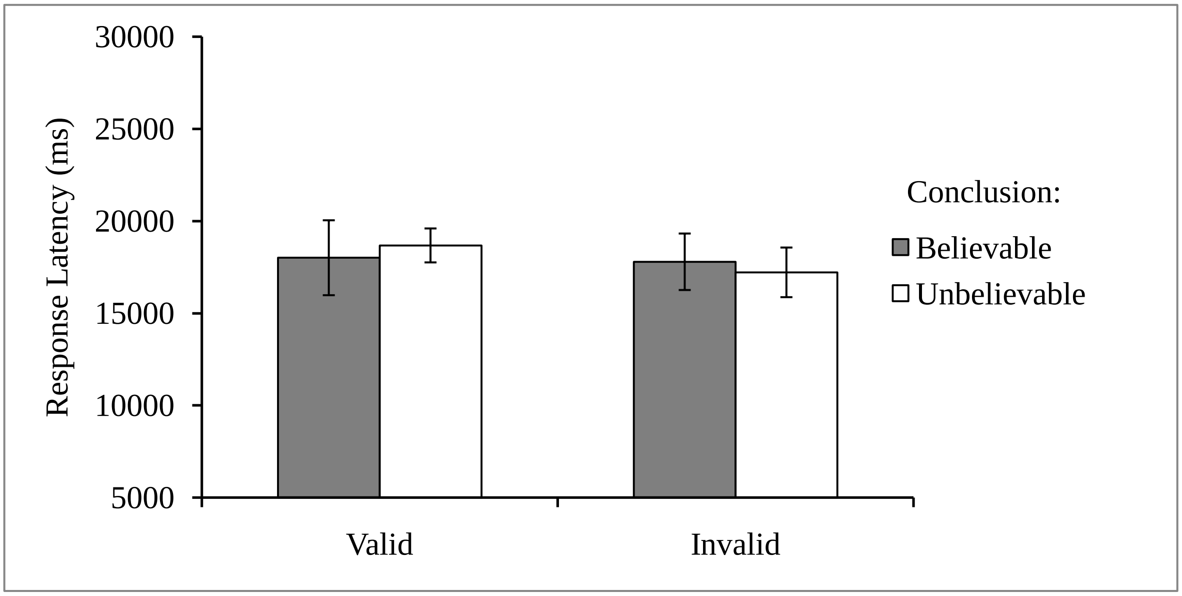

Conditional Syllogisms with Unbelievable Premises – Experimental Group

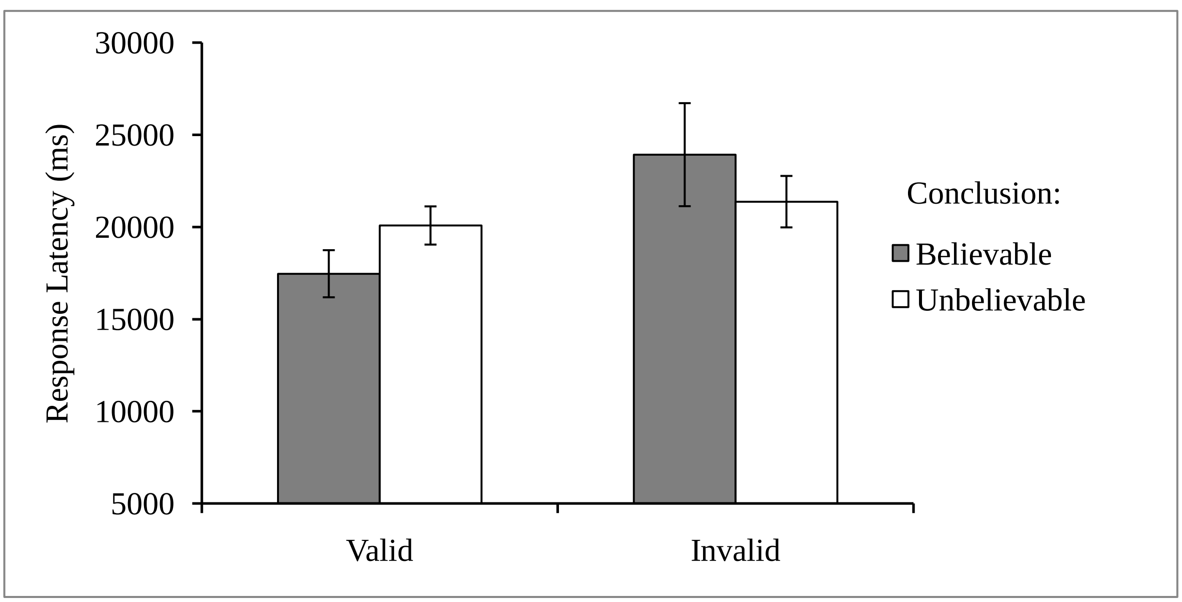

Figure 15. Response latencies as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises *for the* Experimental Group.

Categorical Syllogisms with Believable Premises – Experimental Group

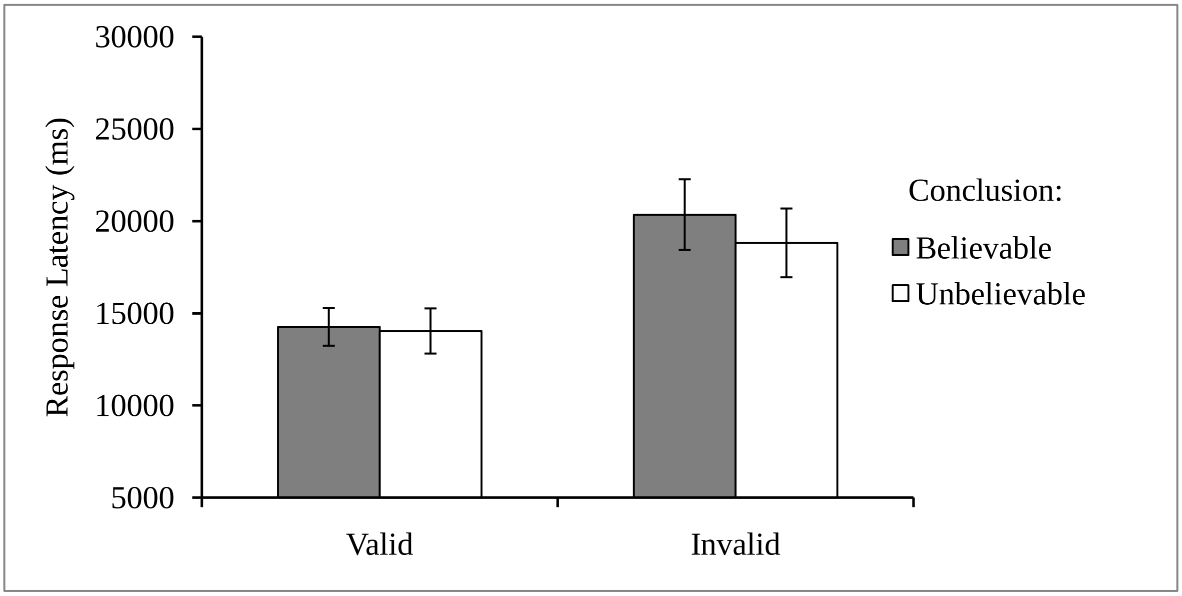

Categorical Syllogisms with Unbelievable Premises – Experimental Group

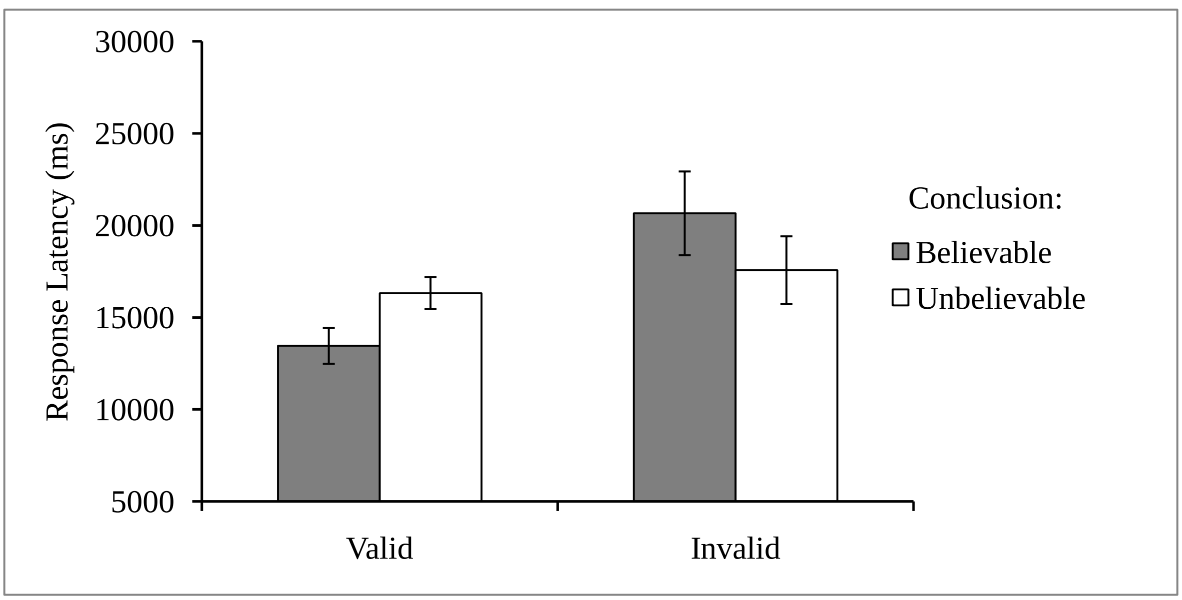

Figure 16. Response latencies as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises *for the* Experimental Group.

premises were believable, there was a significant interaction between Conclusion Validity and Believability, $F(1,83) = 7.12$, $MSE = 5.701 \times 10^7$, $p < .01$, $\eta^2 = .079$, such that participants took longer to reason about conflict problems than non-conflict problems; however, this interaction was only marginally significant when premises were unbelievable, $F(1,83) = 3.22$, $MSE = 7.758 \times 10^7$, $p = .077$, $\eta^2 = .037$. For Conditional Syllogisms, the interaction between Premise Believability, Conclusion Validity, and Conclusion Believability was not significant, $F(1,83) = 0.58$, $MSE = 8.671 \times 10^7$, $p = .447$, $\eta^2 = .007$. The two-way interaction between Conclusion Validity and Conclusion Believability for Conditional Syllogisms was significant, $F(1,83) = 8.63$, $MSE = 4.838 \times 10^7$, $p < .01$, $\eta^2 = .094$, indicating that regardless of premise believability, participants spent longer reasoning about conflict problems than non-conflict problems.

Effects of Premise Believability for Conditional Syllogisms compared to Categorical Syllogisms. There was a Type x Premise Believability interaction, $F(1,83) = 8.09$, $MSE = 1.453 \times 10^8$, $p < .01$, $\eta^2 = .089$. Paired-samples *t*-tests revealed that participants reasoned longer about Conditional Syllogisms with Unbelievable than Believable Premises, $t(84) = 4.84$, $p < .001$. There was no difference in response latencies for syllogisms with Believable and Unbelievable Premises for Categorical Syllogisms, $t(84) = 1.24$, $p = .218$.

Discussion

In general, the results were consistent with the predictions of my proposed theory of premise evaluation. When participants were told to assume that premises were true, belief bias resulted for categorical syllogisms regardless of whether premises were believable or unbelievable, and further, there was a trend for participants to spend longer reasoning about conflict problems when premises were unbelievable. The control group, who did not receive instructions to assume that premises were true, replicated the findings of Experiments 1 and 2:

There was no belief bias when categorical premises were unbelievable, and no evidence that participants took conclusion believability into account when premises were unbelievable.

Conditional syllogisms, of which participants are presumed to assume premises are true as their default, showed the same effects regardless of instruction and replicated the findings of Experiments 1 and 2. Here, belief bias was found regardless of premise believability (participants endorsed more valid than invalid and more believable than unbelievable conclusions, and there were significant interactions between conclusion validity and believability). Thompson's (1996) finding was replicated, such that participants endorsed more conclusions following from believable than unbelievable premises. Further, as in Experiment 2, participants spent longer reasoning about syllogisms with unbelievable premises.

Overall, these findings support the theory that reasoners automatically assume conditional, but not categorical, premises to be true. In the case of conditional syllogisms, this assumption leads reasoners to treat believable and unbelievable conditional premises identically for the purposes of reasoning about the conclusion. Prior to conclusion reasoning, reasoners may evaluate premises and alter responses to be consistent with premise believability. In contrast, for categorical syllogisms, it seems as though rather than assuming that the statements are true, reasoners evaluate the believability of premises prior to considering the conclusion, and then only take into account believability of the conclusion when premises are believable.

It is proposed that, for conditional syllogisms, reasoners assume by default that premises are true, and subsequent to the reasoning process, reject false premises and change responses accordingly. This hypothesis can account for why reasoners spend longer reasoning about syllogisms with unbelievable than believable premises: Evaluating and disbelieving information that was previously accepted takes time (Gilbert, 1991). If instructions to assume that premises are true causes participants to treat categorical premises like conditional ones, then we should also see main effects of premise believability on responses and response latencies for categorical syllogisms (that is, more "valid" responses and shorter responses latencies for believable than unbelievable syllogisms). However, this was not the case: There were equal overall endorsement rates and response latencies for believable and unbelievable categorical premises. Although it is not entirely clear why this is the case, it points to the fact that, even with instructions to assume that premises are true, reasoners nonetheless persist in treating categorical syllogisms differently from conditional syllogisms. It may be that when statements are categorical, reasoners are not motivated to evaluate them and to discredit false ones, perhaps because, despite instructions to assume that premises are true, reasoners could not fully inhibit evaluating categorical premises prior to reasoning about the conclusion. Further, although we see effects of conclusion believability on endorsement rates when categorical premises were unbelievable, the evidence is not as strong for response latencies, arguably a more sensitive measure. That participants were only marginally slower on conflict problems for categorical syllogisms with unbelievable premises is more evidence that participants may not have completely suppressed the initial evaluative component for categorical syllogisms.

Although under instructions to "assume that premises are true" categorical syllogisms do not fully resemble conditional syllogisms, there is still striking evidence that these instructions alter how people reason about categorical syllogisms. Experiment 4 will examine another factor unique to conditional but not categorical syllogisms that may account for differences observed when reasoning about these two syllogism types.

Experiment 4

Experiment 3 provided evidence that reasoners assume from the outset that conditional, but not categorical, premises are true. Why might this be the case? The conditional and categorical statements used in the current experiments are similar in that they make explicit the relations among different classes of items. However, these two types of propositions differ in one crucial way: Conditional syllogisms contain the powerful word '*if.'* The study of this small word itself has a rich background in philosophy and psychology alike. In philosophy literature, there is some debate as to the meaning of the word, with some claiming that *if* introduces a categorical proposition (e.g., Lewis, 1973; Goodman, 1947) and others claiming that the word introduces a suppositional, hypothetical statement (e.g., Mackie, 1973; Barnett, 2006, 2010). These latter philosophers propose that 'if' invites the reader or listener to *suppose* that something is true. Some psychologists agree; Evans and Over (2004) state that "'if' is one of the most important and interesting words in the human language. It is used to express hypothetical thought, which is an essential part of human reasoning and decision making" (p.1) and that "'if' must have, in some sense, a hypothetical or suppositional evaluation..." (p.171). Further, according to a recent theory of conditional reasoning, known as the Suppositional Theory of Conditionals (Evans, Over, & Handley, 2005), "conditionals cue a mental simulation in which people suppose the antecedent (*if* statement) to be true and then assess their degree of belief in the consequent (*then* statement)" (Hadjichristidis et al., 2007, p. 2052).

Perhaps it is this small word that is responsible for the differences in how people treat conditional and categorical premises. It is plausible that the word 'if' itself invites the reader to think hypothetically and to assume that the following proposition is true. On the other hand, categorical statements are presented in a more factual manner, and may signal the reader or listener to accept the statement only if it is true. Experiment 4 will determine whether this is in fact the case using a simple manipulation: adding the word 'if' prior to categorical statements. If this word signals to reasoners that they should assume that premises that follow 'if' are true, even momentarily, then the endorsement patterns and response latencies of categorical syllogisms should resemble those for conditional syllogisms. That is, the word 'if' should halt the execution of evaluative processing of categorical premises that is otherwise hypothesized to occur prior to the reasoning process.

The following predictions can be made. If 'if' causes reasoners to assume that conditional premises are true, then first, adding 'if' to categorical syllogisms should produce a belief bias effect even given unbelievable premises; and second, participants should spend longer reasoning about conflict than non-conflict problems for categorical premises regardless of premise believability and should spend longer reasoning about problems with unbelievable than believable premises.

Method

Participants

Participants were 90 students (31 male, 59 female, mean age = 19.55 yrs, SD = 1.96 yrs) enrolled in psychology courses at the University of Waterloo who participated in exchange for course credit. None had reported having taken a course in logic.

Materials

Stimuli comprised the same syllogisms used in Experiments 2 and 3. For the Control Group, syllogisms were presented exactly as they were in previous experiments. For the Experimental Group, categorical syllogisms were altered slightly by adding the word 'if' to the beginning of premises and the word 'Then' to the beginning of conclusions, as depicted below:

If no A are B, and

If some C are B,

Then, some C are not A

Conditional syllogisms were unchanged from Experiments 2 and 3.

Procedure

Participants were assigned to one of two groups. Both groups received instructions identical to those in Experiment 2 and were tested under the same procedure, with the exception that participants in the Experimental Group received categorical syllogisms augmented with the word ‘if’.

Results

Response Data

Proportions of conclusions endorsed as valid as a function of Conclusion Validity and Believability are shown for Conditional Syllogisms with Believable and Unbelievable Premises and Categorical Syllogisms with Believable and Unbelievable Premises in Figures 17 and 18, respectively, for the Control Group, and Figures 19 and 20, respectively, for the Experimental Group. Conclusion endorsement rates were submitted to a 2 (Type: Categorical, Conditional) x 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability) repeated measures ANOVA with Condition (Control, Experimental) as a between subjects variable.

Overall Belief Bias Effect. Again, as with Experiments 1, 2, and 3, the belief bias effect was found. Specifically, there were main effects of Conclusion Believability, $F(1,88) = 110.02$, $MSE = 1.104$, $p <.001$, $\eta^2 = .556$, and Validity, $F(1,88) = 109.87$, $MSE = .591$, $p <.001$, $\eta^2 = .555$,

Conditional Syllogisms with Believable Premises – Control Group

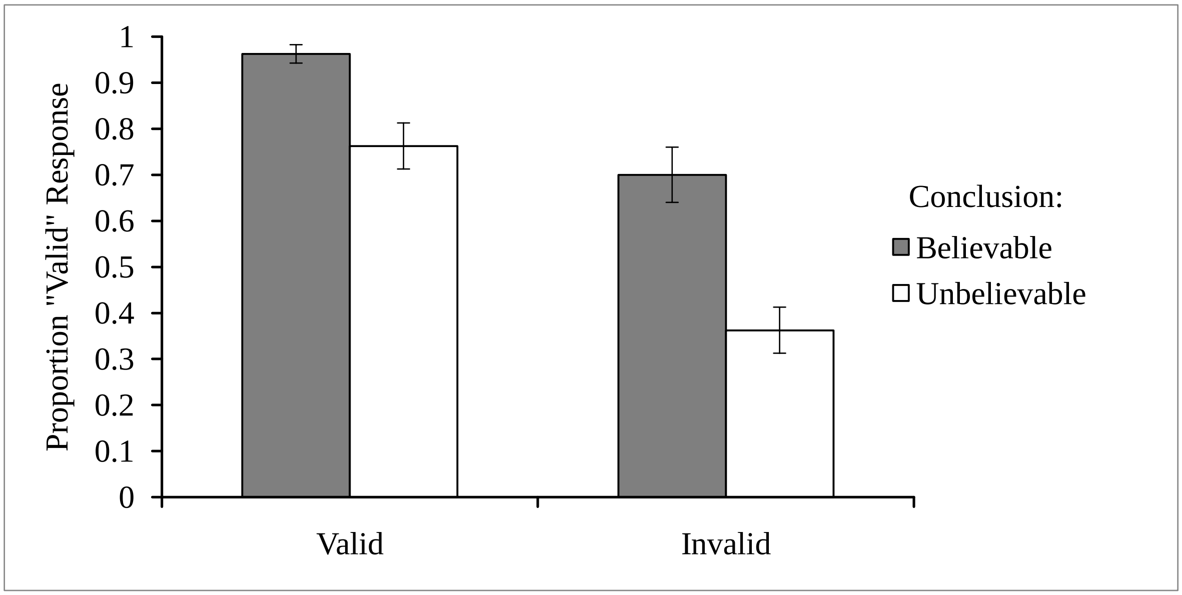

Conditional Syllogisms with Unbelievable Premises – Control Group

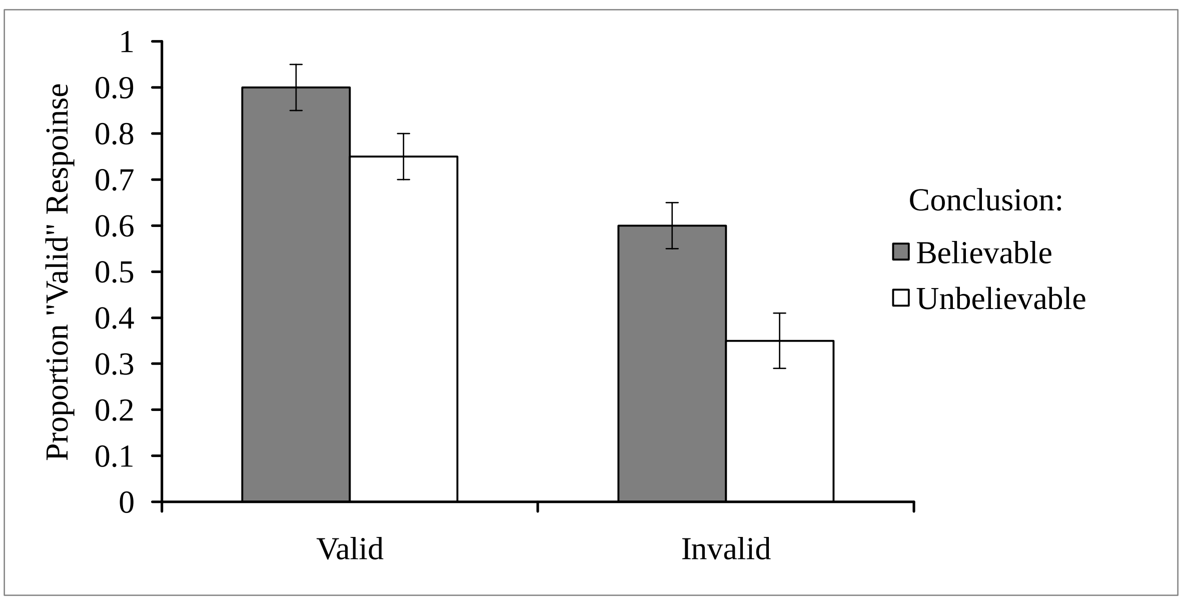

Figure 17. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises *for the* Control Group.

Categorical Syllogisms with Believable Premises – Control Group

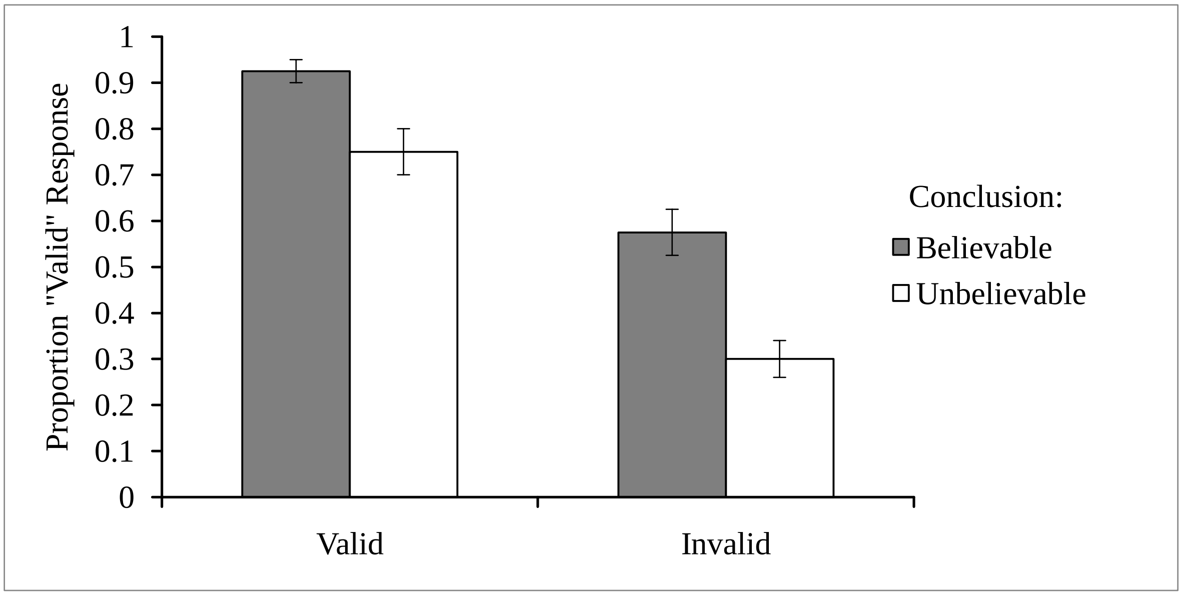

Categorical Syllogisms with Unbelievable Premises – Control Group

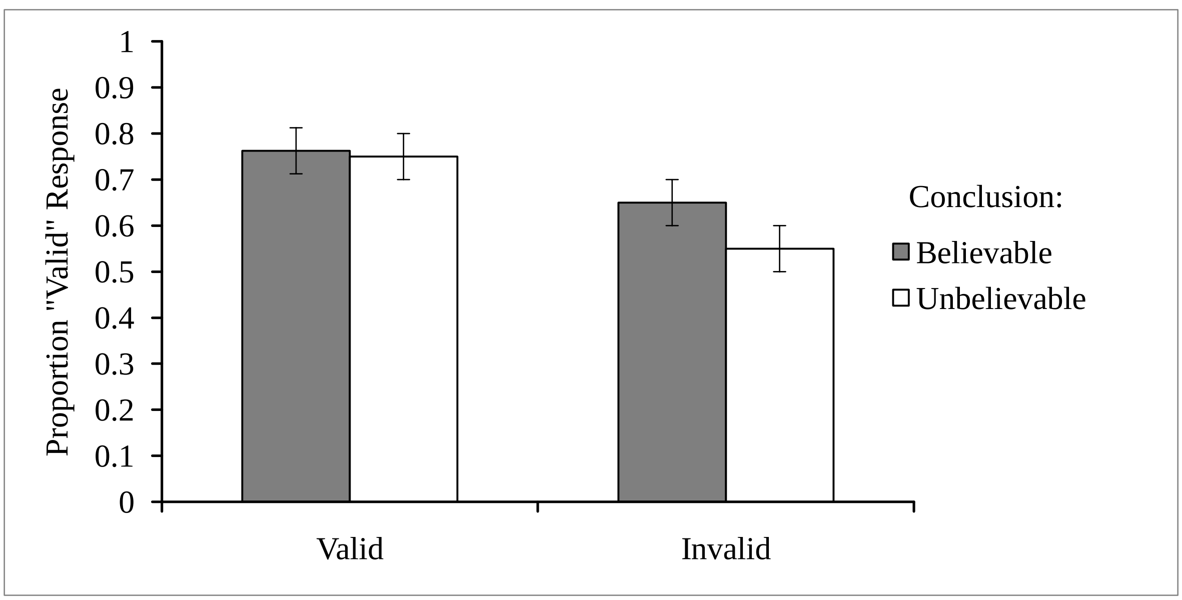

Figure 18. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises *for the* Control Group.

Conditional Syllogisms with Believable Premises – Experimental Group

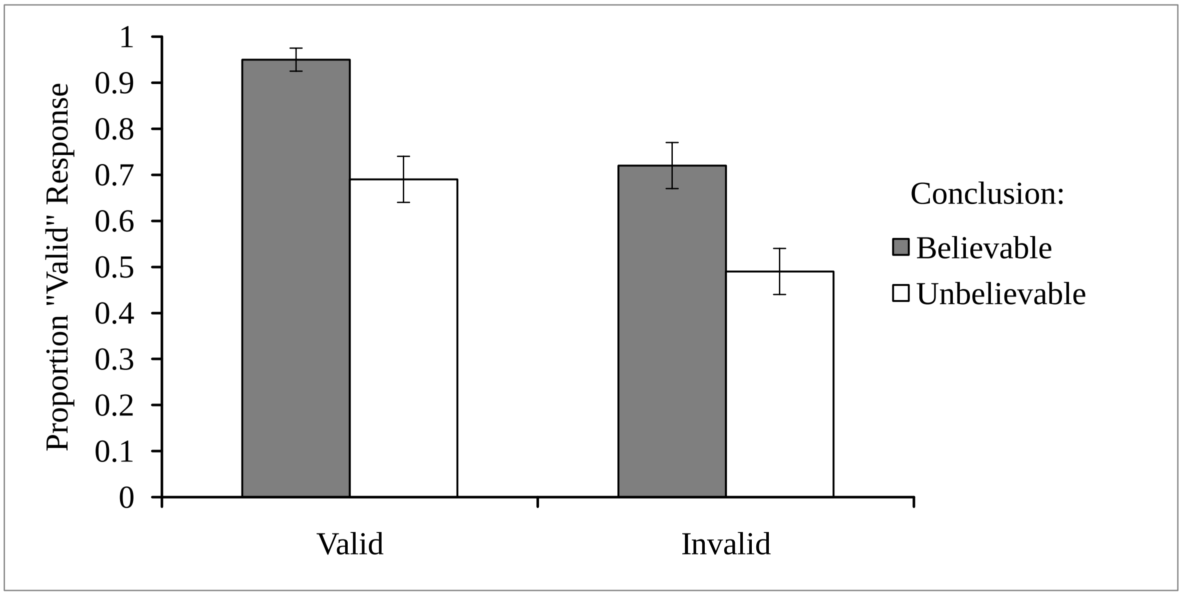

Conditional Syllogisms with Unbelievable Premises – Experimental Group

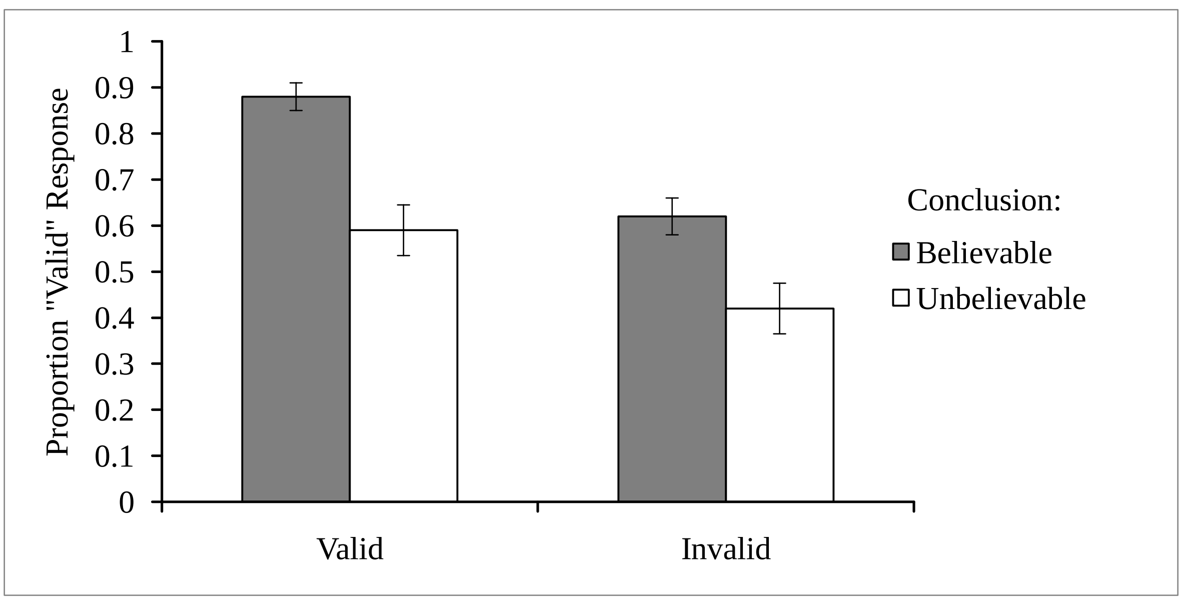

Figure 19. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises *for the* Experimental Group.

Categorical Syllogisms with Believable Premises – Experimental Group

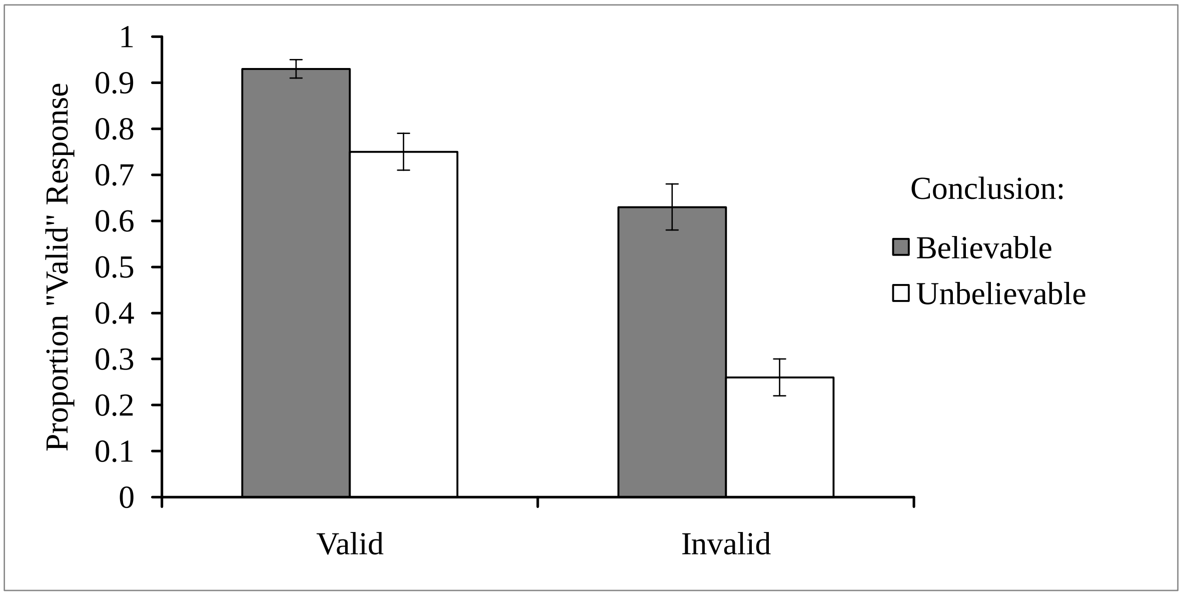

Categorical Syllogisms with Unbelievable Premises – Experimental Group

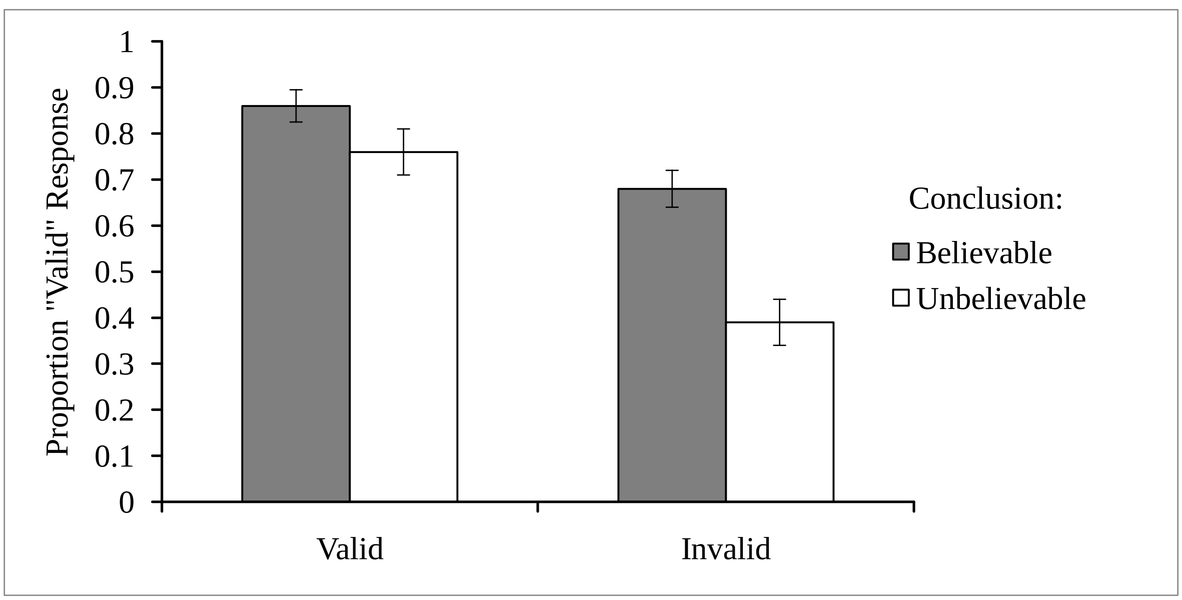

Figure 20. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises *for the* Experimental Group.

and an interaction between Conclusion Validity and Believability, $F(1,88) = 10.32$, $MSE = .253$, $p <.001$, $\eta^2 = .105$.

Effects of Conclusion Believability as a Function of Syllogism Type and Premise Believability. The four-way Type x Premise Believability x Conclusion Believability x Condition interaction was not significant, $F(1,88) = 1.64$, $MSE = .274$, $p = .203$, $\eta^2 = .018$. However, because the influence of condition on the three-way interaction between Syllogism Type, Premise Believability, and Conclusion Believability is essential to the hypotheses, this three-way interaction was examined separately for each condition. For the Control Condition, there was a significant three-way interaction between these three factors, $F(1,39) = 12.66$, $MSE = .218$, $p = .001$, $\eta^2 = .245$. To characterize this three-way interaction, conclusion endorsement rates for participants in the Control Condition were isolated by Syllogism Type and were subjected to two separate 2(Premise Believability: Believable, Unbelievable) x 2(Conclusion Believability: Believable, Unbelievable) repeated measures ANOVAs. Here, there was no interaction between Premise Believability and Conclusion Believability for Conditional Syllogisms, $F(1,39) = 1.51$, $MSE = .500$, $p = .226$, $\eta^2 = .037$. However, similar to previous experiments, Premise Believability and Conclusion Believability did interact for Categorical Syllogisms, $F(1,39) = 8.19$, $MSE = .556$, $p <.01$, $\eta^2 = .174$. Paired-samples *t*-tests revealed that participants endorsed more Believable than Unbelievable Conclusions following from Believable Premises, $t(39) = 4.46$, $p <.001$, whereas Conclusion Believability did not affect endorsement of conclusions following Unbelievable Premises, $t(39) = 1.36$, $p =.183$. In short, in the control condition, the key finding reported in the previous experiments was replicated: Premise believability affects the magnitude of belief bias for categorical but not for conditional syllogisms.

There was no interaction between Type, Premise Believability, and Conclusion Believability for the Experimental Group, $F(1,49) = 2.25$, $MSE = .319$, $p = .140$, $\eta^2 = .044$, suggesting that Premise Believability did not differentially affect the magnitude of the belief bias for Categorical and Conditional Syllogisms.

Conclusion Endorsement for Categorical and Conditional Syllogisms as a Function of Premise Believability. Although there was a Type x Premise Believability interaction, $F(1,88) = 13.94$, $MSE = .261$, $p <.001$, $\eta^2 = .137$, the Type x Premise Believability x Condition interaction was not significant, $F(1,88) = 0.26$, $MSE = .261$, $p = .613$, $\eta^2 = .003$. Overall, premise Believability did not affect overall conclusion endorsement for Categorical Syllogisms, $t(89) = 1.54$, $p = .127$. Conversely, participants endorsed more conclusions following Believable than Unbelievable Premises for Conditional Syllogisms, $t(89) = 3.41$, $p =.001$.

Response Latency Data

Response Latencies as a function of Conclusion Validity and Believability are shown for Conditional Syllogisms with Believable and Unbelievable Premises and Categorical Syllogisms with Believable and Unbelievable Premises in Figures 21 and 22, respectively, for the Control Group, and Figures 23 and 24, respectively, for the Experimental Group. Response Latencies were submitted to a 2 (Type: Categorical, Conditional) x 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability) repeated measures ANOVA with Condition (Control, Experimental) as a between subjects variable.

Effects of Conclusion Believability. Although the five-way interaction between Type, Premise Believability, Conclusion Validity, Conclusion Believability, and Condition was not significant, $F(1,88) = 1.86$, $MSE = 5.678 \times 10^7$, $p = .176$, $\eta^2 = .021$, several 2 (Conclusion

Conditional Syllogisms with Believable Premises – Control Group

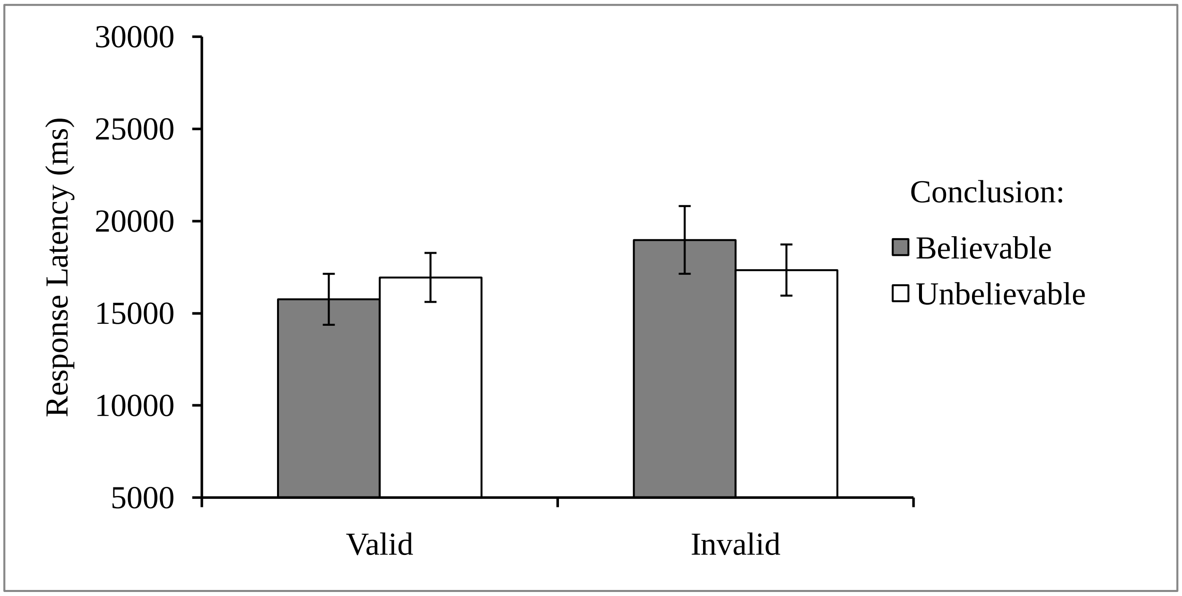

Conditional Syllogisms with Unbelievable Premises – Control Group

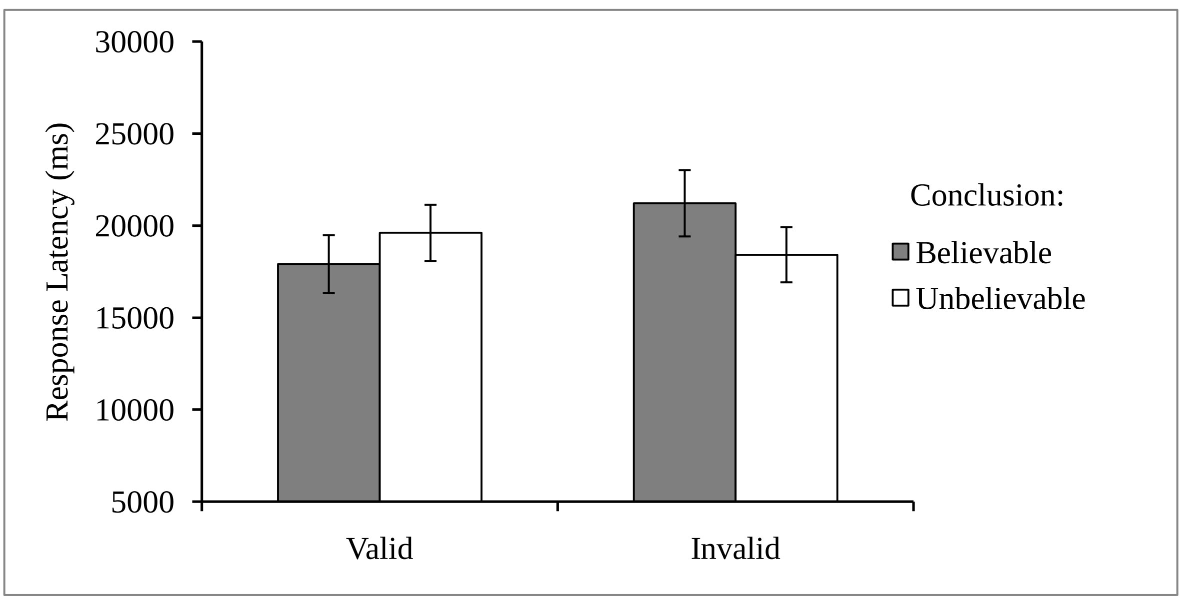

Figure 21. Response latencies as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises *for the* Control Group.

Categorical Syllogisms with Believable Premises – Control Group

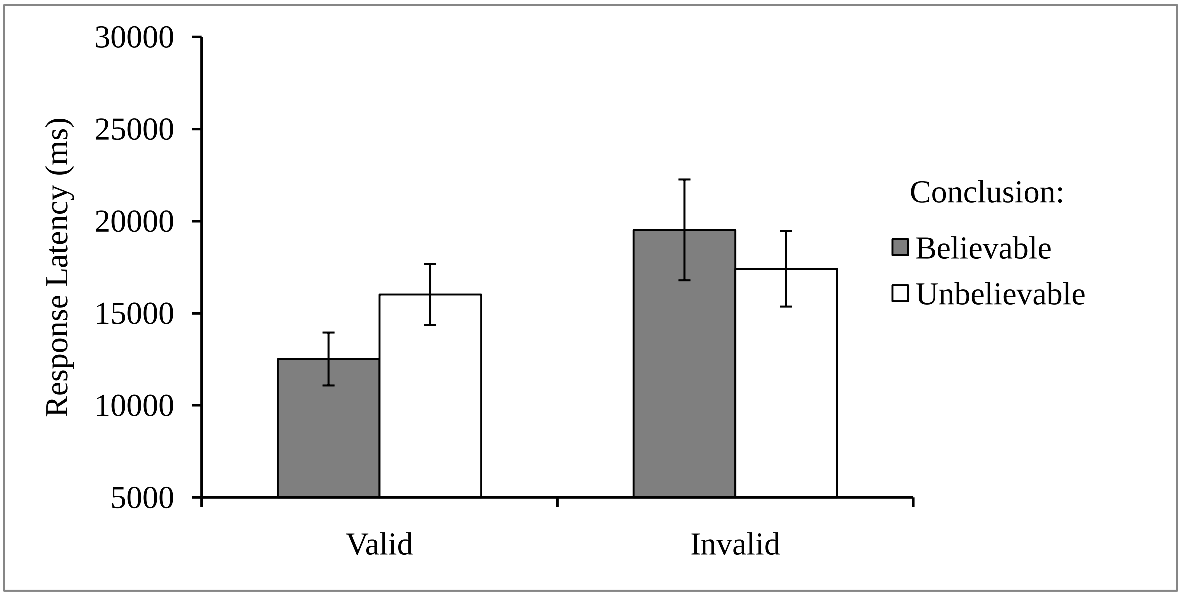

Categorical Syllogisms with Unbelievable Premises – Control Group

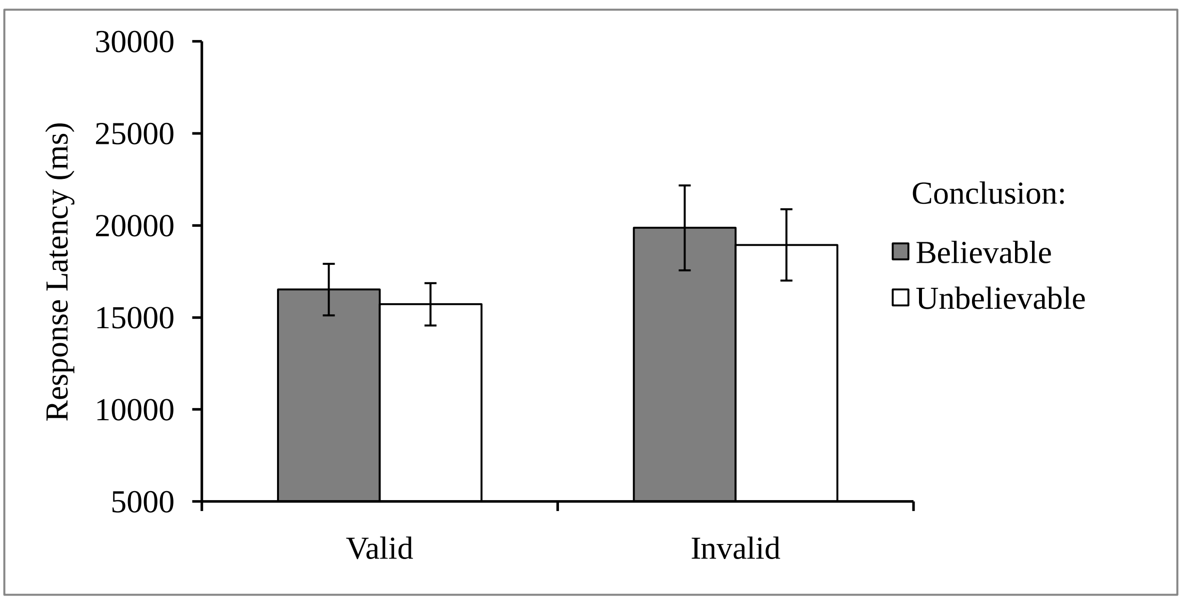

Figure 22. Response latencies as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises *for the* Control Group.

Conditional Syllogisms with Believable Premises – Experimental Group

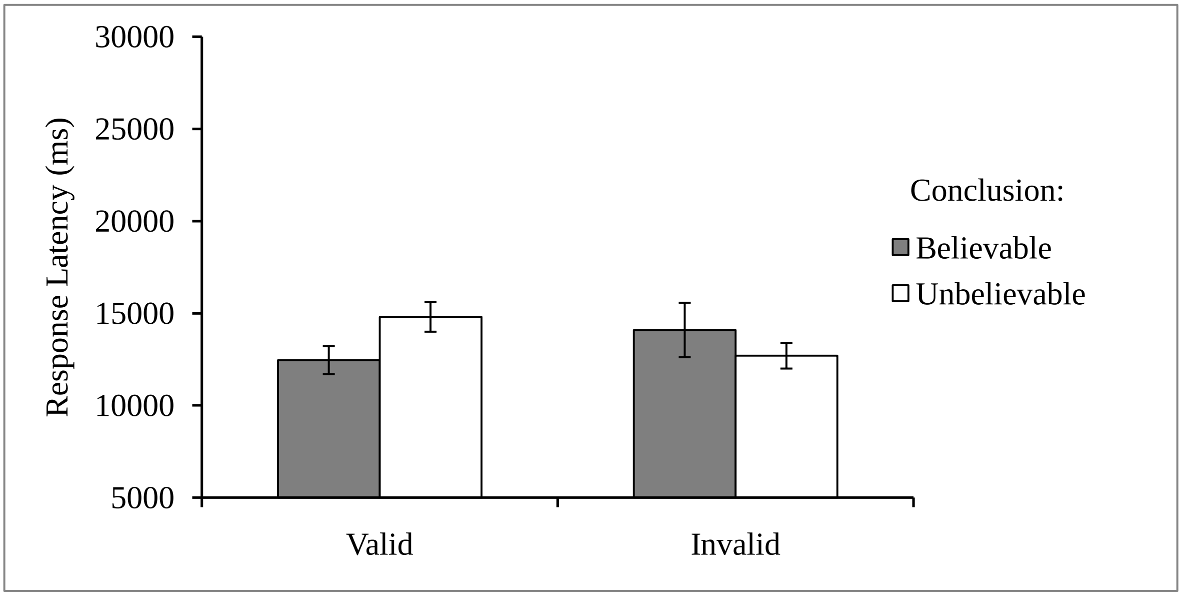

Conditional Syllogisms with Unbelievable Premises – Experimental Group

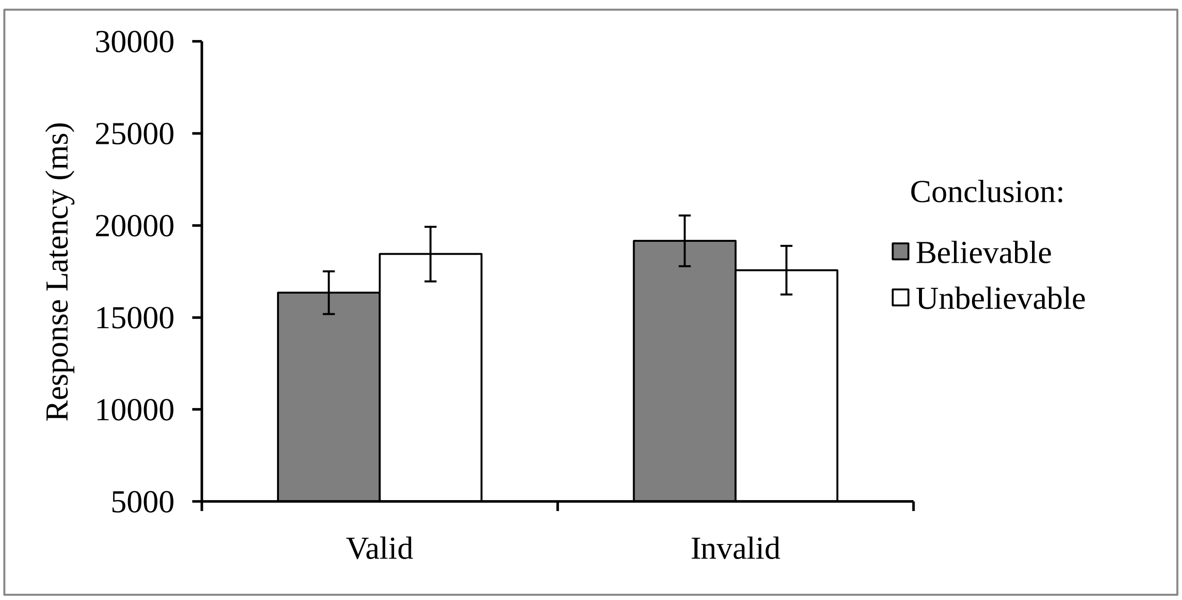

Figure 23. Response latencies as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises *for the* Experimental Group.

Categorical Syllogisms with Believable Premises – Experimental Group

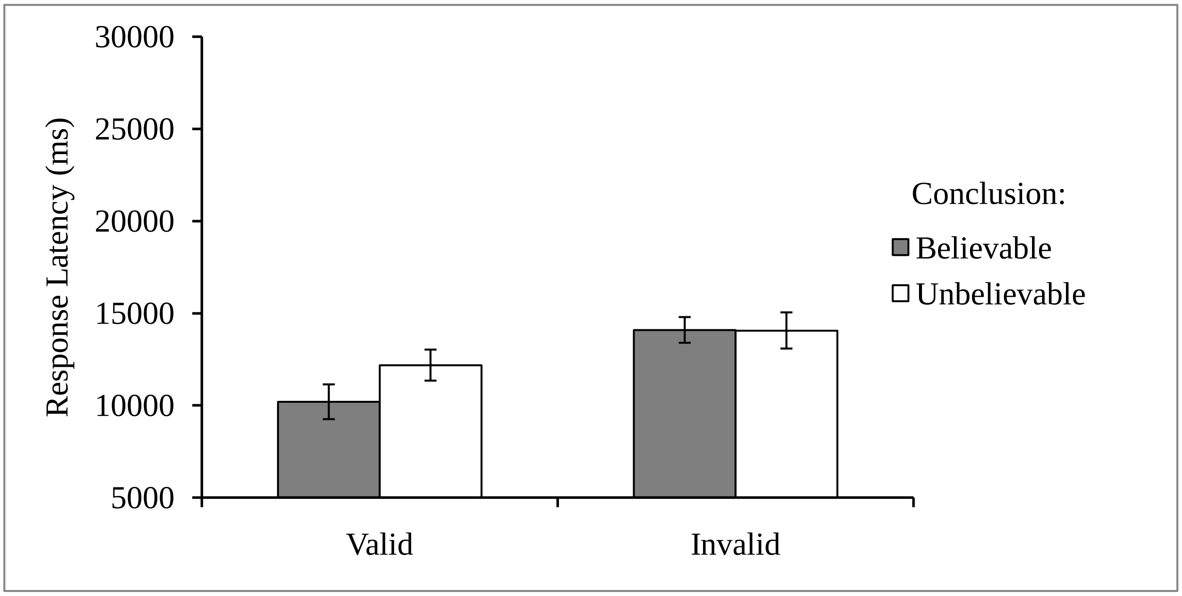

Categorical Syllogisms with Unbelievable Premises – Experimental Group

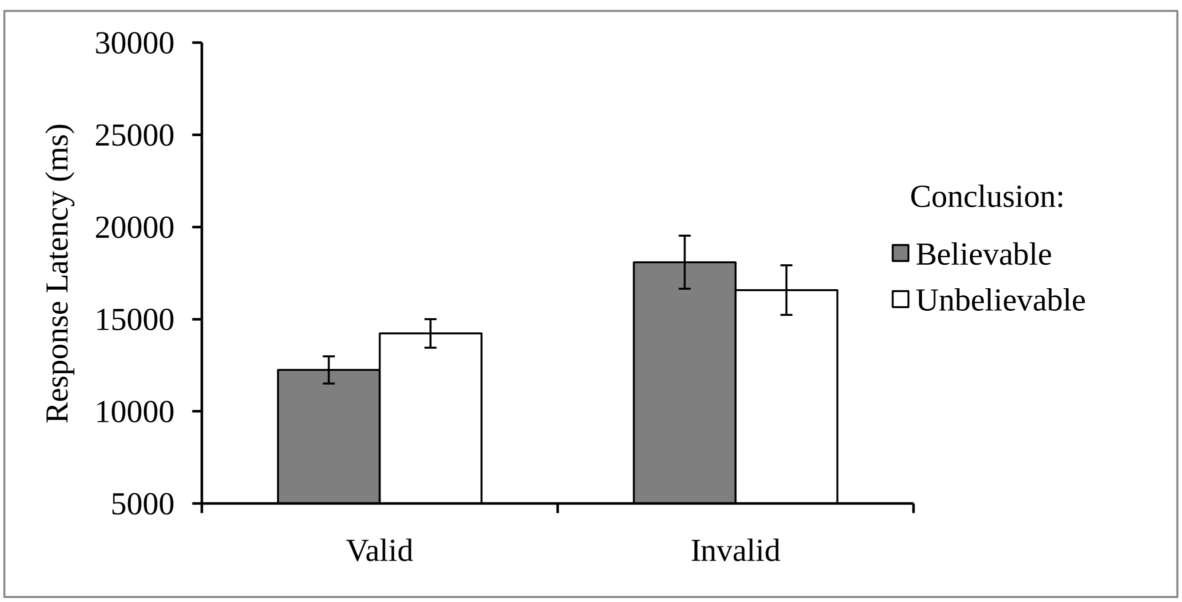

Figure 24. Response latencies as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises *for the* Experimental Group.

Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVAs were conducted to determine whether Conclusion Believability differentially influenced response latencies for Categorical and Conditional Syllogisms with Believable and Unbelievable Premises in the Experimental and Control Conditions.

For the Control Condition, Conclusion Validity and Conclusion Believability interacted for Categorical Syllogisms with Believable Premises, $F(1,39) = 7.16$, $MSE = 4.406 \times 10^7$, $p < .05$, $\eta^2 = .155$. However, this interaction fell short of significance for Conditional Syllogisms with Believable Premises, $F(1,39) = 1.54$, $MSE = 5.178 \times 10^7$, $p = .223$, $\eta^2 = .038$, and for Unbelievable Premises $F(1,39) = 2.42$, $MSE = 8.365 \times 10^7$, $p = .128$, $\eta^2 = .059$. There was also no such interaction for Categorical Syllogisms with Unbelievable Premises, $F(1,39) = 0.003$, $MSE = 5.273 \times 10^7$, $p = .957$, $\eta^2 < .001$, suggesting that conflicting Conclusion Validity and Conclusion Believability did not differentially interfere with reasoning.

For the Experimental Condition, there was a marginally significant interaction between Conclusion Validity and Conclusion Believability for Conditional Syllogisms with Believable Premises, $F(1,49) = 3.93$, $MSE = 4.441 \times 10^7$, $p = .053$, $\eta^2 = .074$; this interaction for Unbelievable Premises fell short of significance, $F(1,49) = 2.22$, $MSE = 7.666 \times 10^7$, $p = .142$, $\eta^2 = .043$. The interaction between Conclusion Validity and Conclusion Believability for Categorical Syllogisms with Believable Premises was not significant, $F(1,49) = 2.31$, $MSE = 2.097 \times 10^7$, $p = .135$, $\eta^2 = .045$, however, there was a main effect of Conclusion Believability, $F(1,49) = 14.28$, $MSE = 2.869 \times 10^7$, $p < .001$, $\eta^2 = .226$. Thus, even though the interaction between Conclusion Validity and Believability was not significant, Conclusion Believability did influence response latencies. Crucially, there was a significant interaction between Conclusion

Validity and Conclusion Believability for Categorical Syllogisms with Unbelievable premises, $F(1,49) = 4.46$, $MSE = 3.383 \times 10^7$, $p < .05$, $\eta^2 = .083$.

Effects of Premise Believability for Conditional Syllogisms compared to Categorical Syllogisms. Neither the Syllogism Type x Premise Believability x Condition interaction, $F(1,88) = 0.40$, $MSE = 6.133 \times 10^7$, $p = .526$, $\eta^2 = .005$, nor Syllogism Type x Premise Believability Interaction, $F(1,88) = 1.98$, $MSE = 6.133 \times 10^7$, $p = .163$, $\eta^2 = .022$, were significant.

Discussion

These results partially support predictions made by the theory of premise evaluation set out earlier. First and most central, adding 'if' to categorical premises elicited a belief bias effect even when premises were unbelievable. Second, participants were slower to reason about problems in which the validity and believability of the conclusion conflicted for categorical syllogisms with unbelievable premises. Taken along with the results of Experiment 3, these results support the theory that 'if' leads reasoners to assume that premises are true and to treat believable and unbelievable premises similarly.

As in Experiment 3, the experimental manipulation did not cause data for categorical syllogisms to fully resemble conditional syllogisms: Participants shown categorical syllogisms with 'if' did not endorse more conclusions following from believable than unbelievable premises and did not spend longer reasoning about problems with unbelievable premises. However, note that although reasoners did accept more conclusions following from believable than unbelievable conditional premises in this experiment, participants did not spend longer reasoning about conditional syllogisms with believable than unbelievable premises, in contrast to Experiments 2 and 3. Thus, reasoners may not be reliably motivated to discredit false information that was once believed, regardless of the format of the premise. Further, although 'if' was added to categorical statements, they are still essentially different from categorical statements. Conditional syllogisms contain three full ‘If/Then’ statements, whereas categorical syllogisms, taken as a whole, contain only one entire ‘If/Then’ statement. This essential difference is unavoidable without altering the type of syllogism and may be the reason why results of categorical syllogisms, under instructions to assume that premises are true or with the addition of ‘if,’ do not fully resemble those of conditional syllogisms. Nonetheless, given the limitations imposed by the logical structure of categorical syllogisms, there is compelling evidence that the word ‘if’ is the key difference between categorical and conditional syllogisms that leads reasoners to assume that conditional premises are true but to evaluate the believability of categorical premises. It should also be noted that the conclusions drawn from Experiments 3 and 4 are based in a large part on marginally significant higher-order interactions. This is not surprising, given the robustness of the belief bias effect itself and how difficult it is to manipulate this effect. That being said, future research should strive to develop more powerful methods of modulating bias when reasoning with these two types of syllogisms.

Experiment 5

Thus far, there is evidence that conditional and categorical syllogisms differ in how premises are considered. I have hypothesized, specifically, that for conditional syllogisms, reasoners automatically assume that premises are true and then, at least sometimes, disbelieve false premises either *after* or otherwise independently of reasoning about the conclusion. For categorical syllogisms, reasoners evaluate the believability of premises right away, and then seem to disregard believability information in the conclusion when premises are unbelievable. Although previous experiments have investigated how premises are evaluated, there remain some unanswered questions: First, what is the time course of reasoning, particularly for conditional syllogisms? Do reasoners evaluate premise believability after the conclusion has been reasoned about, or does this happen in parallel with reasoning about the conclusion? Second, why is belief bias not evident for categorical syllogisms with unbelievable premises? The final experiment of this dissertation was designed to answer these questions using two different yet converging approaches, eye-tracking and a memory task.

Eye-Tracking

Eye-movement monitoring has been used for decades in cognitive psychology research (Rayner, 1998) under the assumption that fixations and eye movements act as a proxy for where one is directing his or her attention. This approach has been applied to various reasoning modalities, including insight problem solving (Grant & Spivey, 2003), analogical reasoning (Bethell-Fox, Lohman, & Snow, 1984), the Wason Card Selection Task (Ball, Lucas, Miles, & Gale, 2003), and spatial reasoning (Korner & Gilchrist, 2004). Use of eye-tracking methodology in the study of deductive reasoning is a relatively new approach to disentangling reasoning processes. Researchers who have employed it (e.g., Ball et al., 2006; Espino, Santamaria, Meseguer, & Carreiras, 2005) report success in using the technology to support or refute theories of reasoning. Particularly, Ball et al. (2006) examined premise inspection times for categorical syllogisms before and after the conclusion was attended to, and reported that pre-conclusion premise inspection times did not differ according to the validity or believability of the conclusion. However, they found that participants looked significantly longer at premises after looking at the conclusion when the validity and believability of the conclusion conflicted. They used these findings to support theories of deductive reasoning that posit that the believability and validity of the conclusion itself interact to determine how reasoning progresses.

The success of prior experiments using eye-tracking technology in supporting or refuting reasoning theories points to the utility of monitoring eye movements during reasoning, yet no eye-tracking experiments to date have examined how eye-movements differ for categorical and conditional syllogistic reasoning and for believable and unbelievable premises. The current experiment will do just that, by comparing premise and conclusion inspection times for categorical and conditional syllogisms with believable and unbelievable premises. Although the previous experiments in this thesis have provided global response latencies, up until this time, it has not been possible to determine how much of this time participants spent on the premises and how much they spent on the conclusion. Eye-tracking can be used to classify response time based on the component of the problem to which participants are attending.

According to the Mental Models account, reasoners begin the reasoning process by forming a model of the premises. Next, they determine whether the conclusion is valid or invalid given this model. Here, the key assumption will be made that once the conclusion is attended to, the process of reasoning about the conclusion begins. Under this assumption, other measures of interest include how often reasoners refer back to the premises after attending to the conclusion (which can be taken as an index of the extent to which participants consider the premises after beginning to reason about the conclusion), and how much time is spent attending to premises before and after the conclusion has been attended to (or how much time is dedicated to the premises pre- and post-reasoning).

Given the theory of premise evaluation that I propose in this dissertation, the following predictions can be drawn. First, if reasoners assume that conditional premises are true and then disbelieve unbelievable premises after reasoning, then they should refer back to the premises after the conclusion has been looked at more often than they do for categorical syllogisms, which reasoners are thought to have evaluated prior to the reasoning process. Further, categorical premises should be attended to longer, and any effects of premise believability should be evident, before the conclusion is attended to rather than after, because this is when premise evaluation is thought to take place. The converse should be true for conditional syllogisms, if evaluation occurs after reasoning. Finally, if reasoners disregard conclusion believability for categorical syllogisms with unbelievable premises, there should be no evidence of conclusion believability affecting eye movements.

Memory for Syllogisms

Examining memory for content that is reasoned about is a novel approach in deductive reasoning, however, its use is warranted to address a key question: Do reasoners ignore believability information in categorical syllogisms with unbelievable premises? Take the following syllogism:

No animals are dogs

Some lions are dogs

Therefore, some lions are not animals

This categorical syllogism has two unbelievable premises, and according to Experiments 1 through 4, reasoners do not show belief bias for this type of syllogism. It is not yet known why this is the case. It is possible that, when given unbelievable premises, which they are thought to evaluate prior to reasoning, reasoners realize that the content of the problem will not help them reason about the conclusion, and may in fact hinder them. Reasoners may then abstract the logical form from the syllogism and disregard the unbelievable content, representing it like so:

No A are B

Some C are B

Therefore, some C are not A

Reasoning about this syllogism as opposed to the above syllogism, if done correctly, yields the same "valid" response without being hindered by the unbelievable content of the above syllogism. Thus, when confronted with a categorical syllogism with premises that are recognized to be unbelievable, reasoners may treat the syllogism as an abstract problem. If this is the case, participants should show reduced memory for the content of these problems relative to categorical syllogisms with believable premises and conditional syllogisms with believable and unbelievable premises. In these situations, reasoners would not abstract the logical form, and therefore would devote more processing to the content, making it more memorable. This hypothesis is also tested in this experiment.

Method

Participants

Participants were 39 students (16 male, 23 female, mean age = 20.89 yrs, SD = 2.12 yrs) enrolled in undergraduate Psychology courses at the University of Waterloo who participated in exchange for bonus course credit. All participants reported that they had never taken a course in logic and reported normal or corrected-to-normal vision.

Stimulus Displays

Each display contained one syllogism in black 28 point Calibri font against a white background. The three components of the syllogism (i.e., each premise and the conclusion) were on separate horizontal lines separated vertically by 3.8 cm, and the entire syllogism was centred both horizontally and vertically on the display. Displays were viewed from a distance of 61 cm. Each statement measured between 9.5 cm to 32.0 cm horizontally, depending on the length of the statement, and 0.8 cm vertically, corresponding to a visual angle of 0.75º.

Eye-Tracking Apparatus

An ASL Eye-Trac6 Desktop Model eye-tracking system monitored eye-movements. A chin rest was positioned 61 cm away from the display screen and camera to ensure that participants' head movements were kept to a minimum. A camera situated beneath the display screen recorded pupil reflections of the most accurately calibrated eye. Prior to beginning the experiment, the system was calibrated for each participant using a nine-point calibration procedure.

Two display screens were used: a Hanns-G PC compatible screen set at a resolution of 1024 x 768 pixels screen presented stimuli to participants, and a second monitor presented the displays and gaze position to the experimenter in real-time, allowing the experimenter to assess accuracy. Stimuli were programmed and presented to participants using E-Prime v1.2 (Schneider et al., 2001) software.

Procedure

The experimental session began with orienting participants to the eye-tracking apparatus and calibrating the eye-tracking system. Participants then proceeded through the reasoning task in a procedure that was identical to that of Experiment 2.

The eye-tracking portion of the experiment ended at the completion of the reasoning task. Participants then completed several measures of cognitive abilities, which will not be discussed further. Approximately one hour after the completion of the reasoning task, participants were given a free recall memory test in which they were asked to recall as many of the content words from the syllogisms as they could remember. Participants were given a sheet to write down these words with the following instructions:

> Think back to first task from this experiment. You were asked to determine whether a conclusion was valid or invalid based on given information. Your job in this task is to write down all of the topics that you can remember from those problems.
> For example, if you remember reasoning about the following problem:
>
> No spoons are knives
> Some silver things are knives
> Therefore, some spoons are not silver things
>
> You could write down *spoons*, *knives*, or *silver things.* (you do not need to write down words like *no* or *are*).
> Write down as many key words as you can remember.

Participants were given as much time as they needed to write down (free recall) as many words as they could remember. After this task, participants were debriefed.

Results

Response Data

The proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability are shown for Conditional Syllogisms with Believable and Unbelievable Premises and Categorical Syllogisms with Believable and Unbelievable Premises in Figures 25 and 26, respectively. Conclusion endorsement rates were submitted to a 2 (Type: Conditional, Categorical) x 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVA.

Overall Belief Bias Effect. The main effects characteristic of the belief bias effect were found: Participants endorsed more Believable than Unbelievable Conclusions, $F(1,38) = 48.76$, $MSE = .581$, $p < .001$, $\eta^2 = .562$, and more Valid than Invalid Conclusions, $F(1,38) = 7.74$, $MSE = 1.715$, $p < .01$, $\eta^2 = .169$. The typical interaction between Conclusion Validity and Believability was non-significant, $F(1,38) = 0.60$, $MSE = .452$, $p = .444$, $\eta^2 = .016$.

Effects of Conclusion Believability as a Function of Syllogism Type and Premise Believability. The critical interaction between Type, Premise Believability, and Conclusion Believability was significant, $F(1,38) = 4.95$, $MSE = .272$, $p < .05$, $\eta^2 = .115$. To explore this interaction, endorsement rates for Categorical and Conditional Syllogisms were submitted to two separate 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVAs. For Conditional Syllogisms, Premise Believability did not interact with Conclusion Believability, $F(1,37) = 0.34$, $MSE = .348$, p .563, $\eta^2 = .009$. For Categorical Syllogisms, there was a significant Premise Believability x Conclusion Believability interaction, $F(1,37) = 15.95$, $MSE = .253$, $p < .001$, $\eta^2 = .301$. Paired-samples *t*-tests revealed that participants endorsed more Believable than Unbelievable Conclusions when categorical premises were believable, $t(37) = 5.87$, $p < .001$. Conversely, Conclusion Believability did not affect conclusion endorsement when premises were unbelievable, $t(37) = 1.18$, $p = .244$.

Conditional Syllogisms with Believable Premises

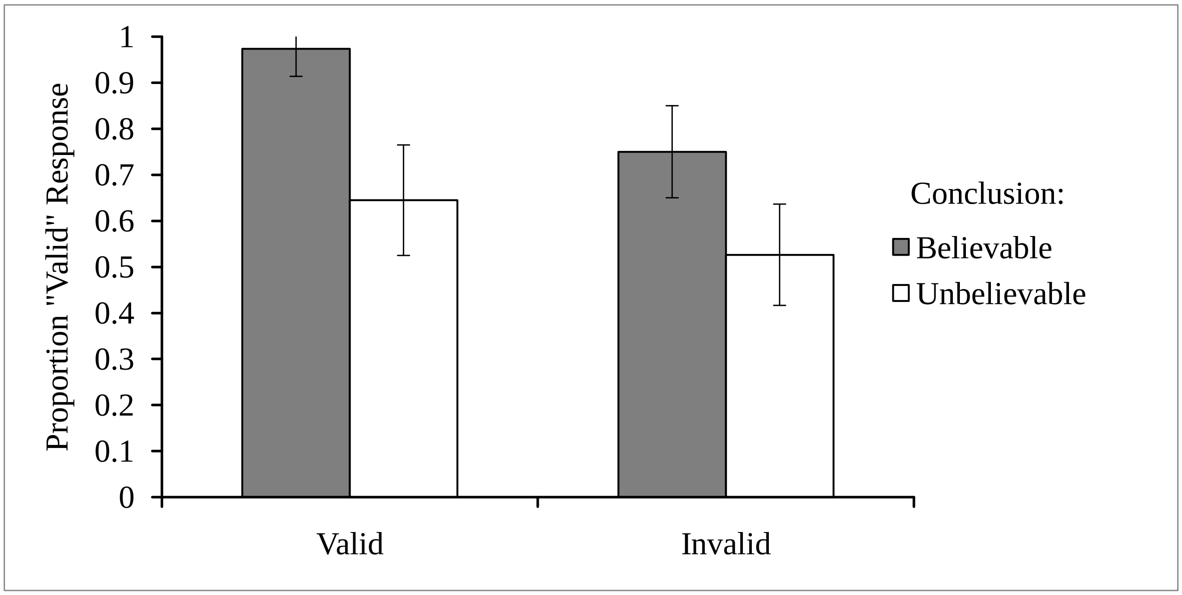

Conditional Syllogisms with Unbelievable Premises

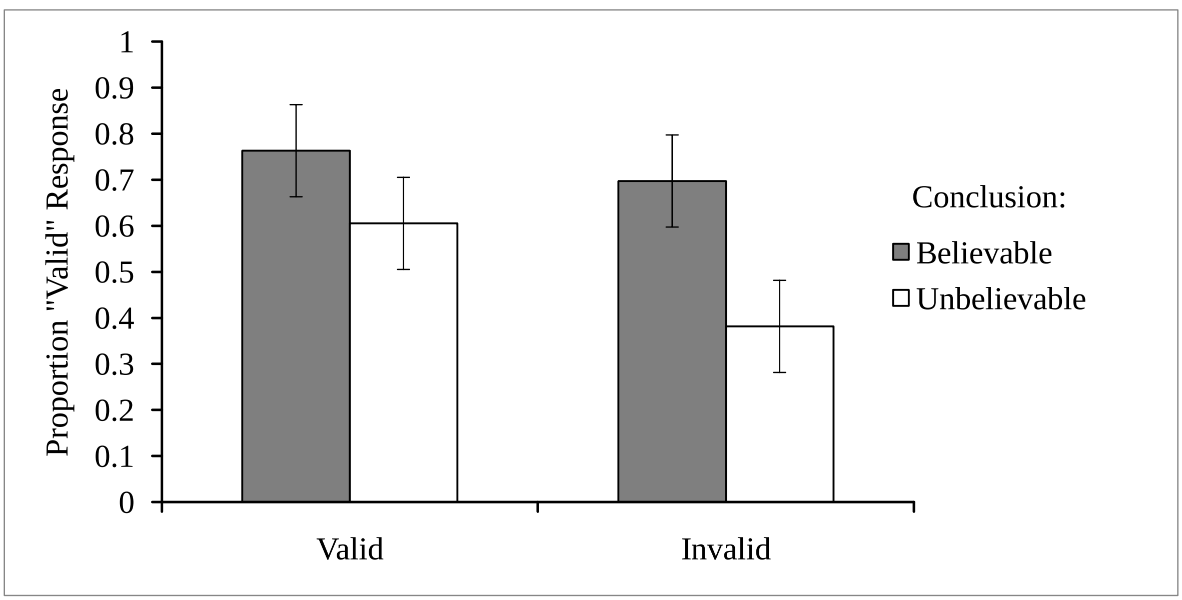

Figure 25. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises.

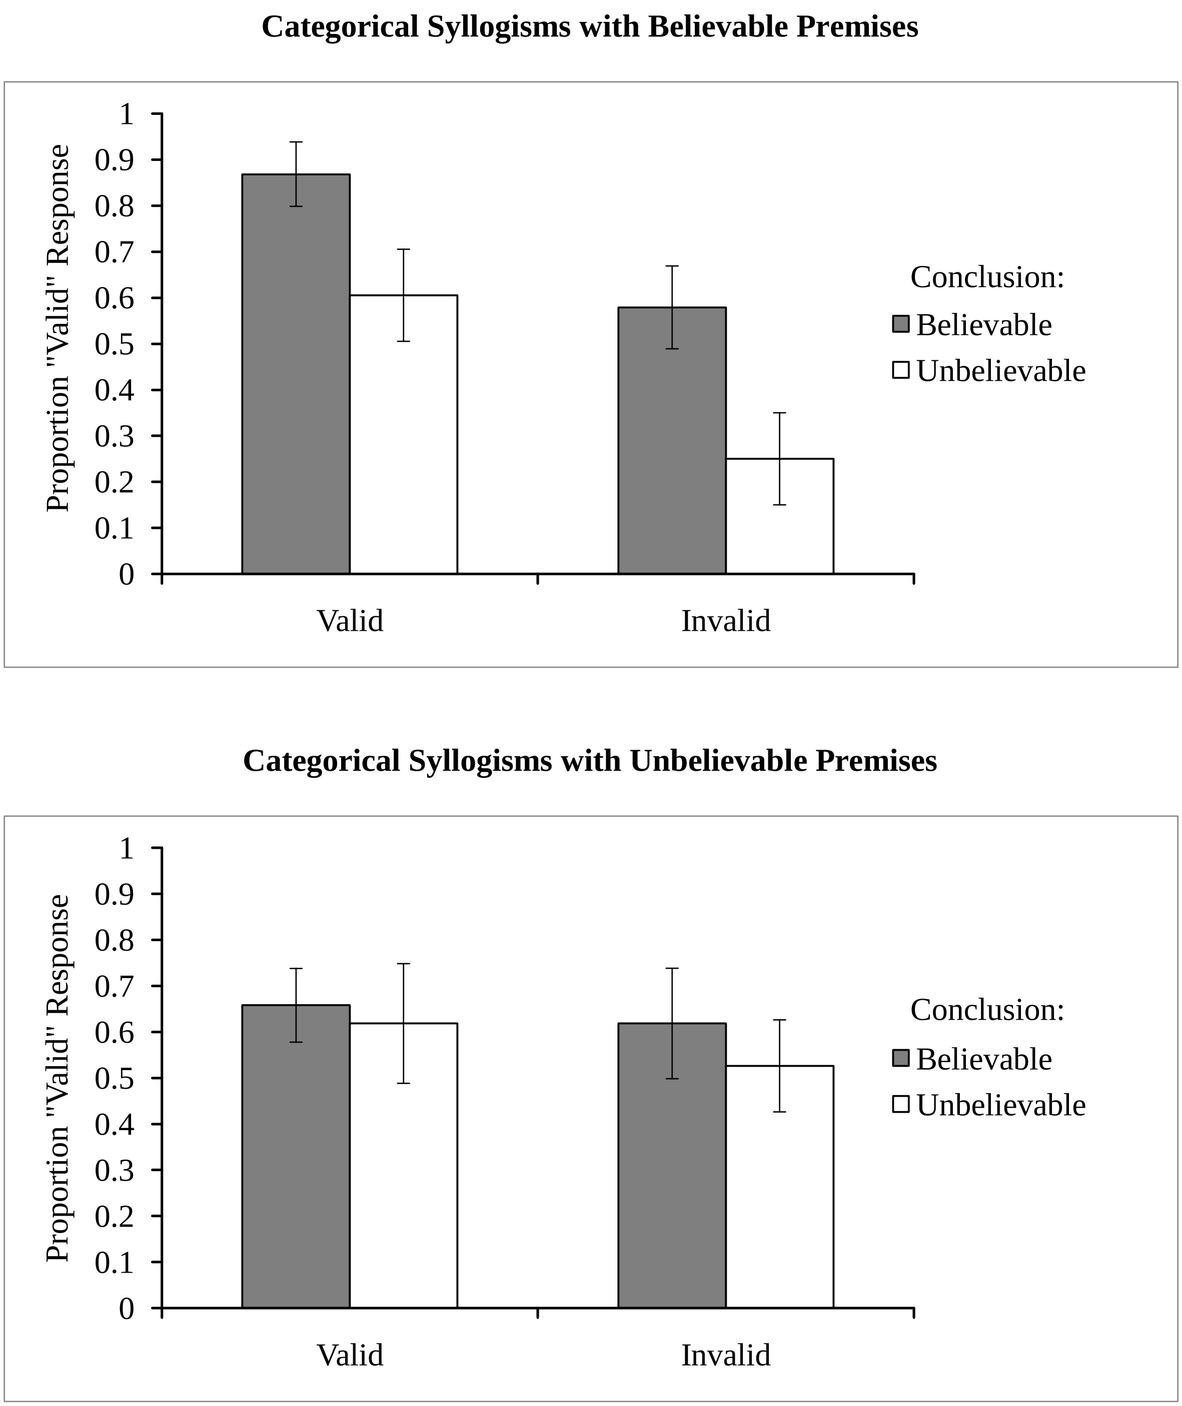

Figure 26. Proportion of conclusions endorsed as valid as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises.

Conclusion Endorsement for Categorical and Conditional Syllogisms as a Function of Premise Believability. There was a significant Type x Premise Believability interaction, $F(1,38) = 8.20$, $MSE = .361$, $p < .01$, $\eta^2 = .178$. Participants endorsed more conclusions following from Believable than from Unbelievable Premises for Conditional Syllogisms, $F(1,37) = 11.08$, $MSE = .343$, $p < .01$, $\eta^2 = .230$; whereas there was no difference in conclusion endorsement based on Premise Believability for Categorical Syllogisms, $F(1,37) = 0.68$, $MSE = .395$, $p = .417$, $\eta^2 = .018$.

Eye-Tracking Data

Of the 39 participants who participated in this Experiment, 25 yielded usable eye-tracking data. Thirteen participants were excluded because the experimenter was unable to accurately calibrate the eye-tracking apparatus to the participants' eyes.

For each syllogism, gaze durations and transitions were summed for two areas of interest (AOI): Premises and Conclusion. Looking time to each AOI is depicted as a function of Premise Believability for Conditional Syllogisms and Categorical Syllogisms in Figure 27. To explore how Premise Believability affected how participants viewed Conditional and Categorical syllogisms, the total of gaze durations was submitted to a 2 (Syllogism Type: Conditional, Categorical) x 2 (Premise Believability: Believable, Unbelievable) x 2 (AOI: Premises, Conclusion) repeated measures ANOVA. There was a significant interaction between Type and AOI, $F(1,24) = 26.86$, $MSE = 1.285$, $p < .001$, $\eta^2 = .528$. This interaction was explored with paired-samples *t*-tests, which revealed that participants viewed premises significantly longer for conditional than for categorical syllogisms, $t(24) = 5.07$, $p < .001$, whereas there was no difference in conclusion gaze duration across Syllogism Type, $t(24) = 1.49$, $p = .150$. There was also a main effect of Premise Believability, $F(1,24) = 21.32$, $MSE = 1.568$, $p < .001$, $\eta^2 = .470$,

Conditional Syllogisms

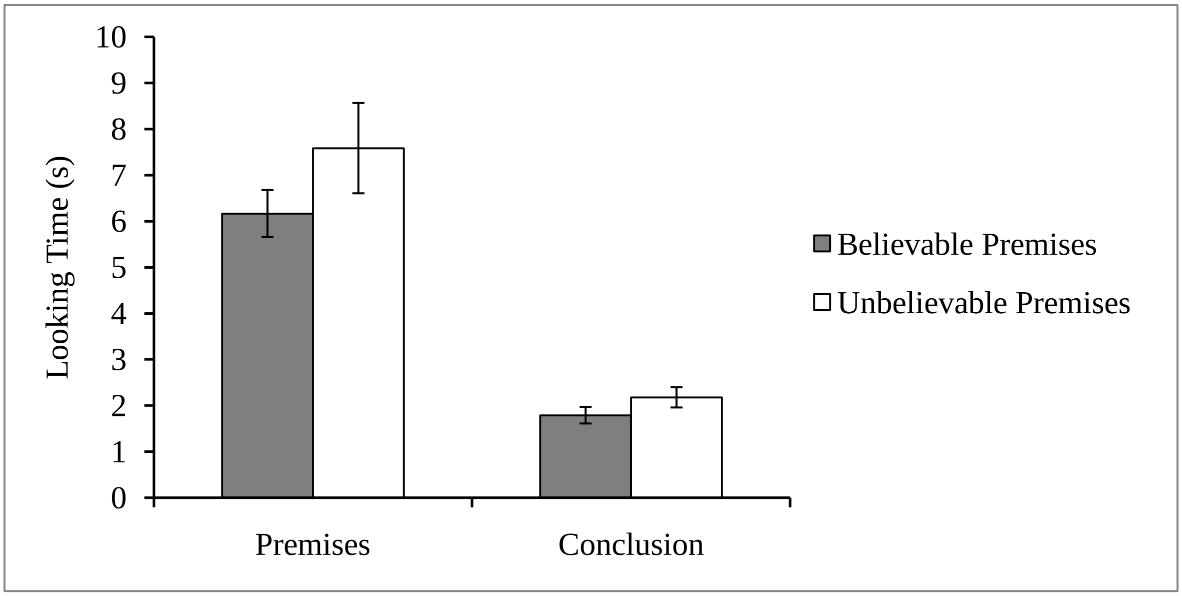

Categorical Syllogisms

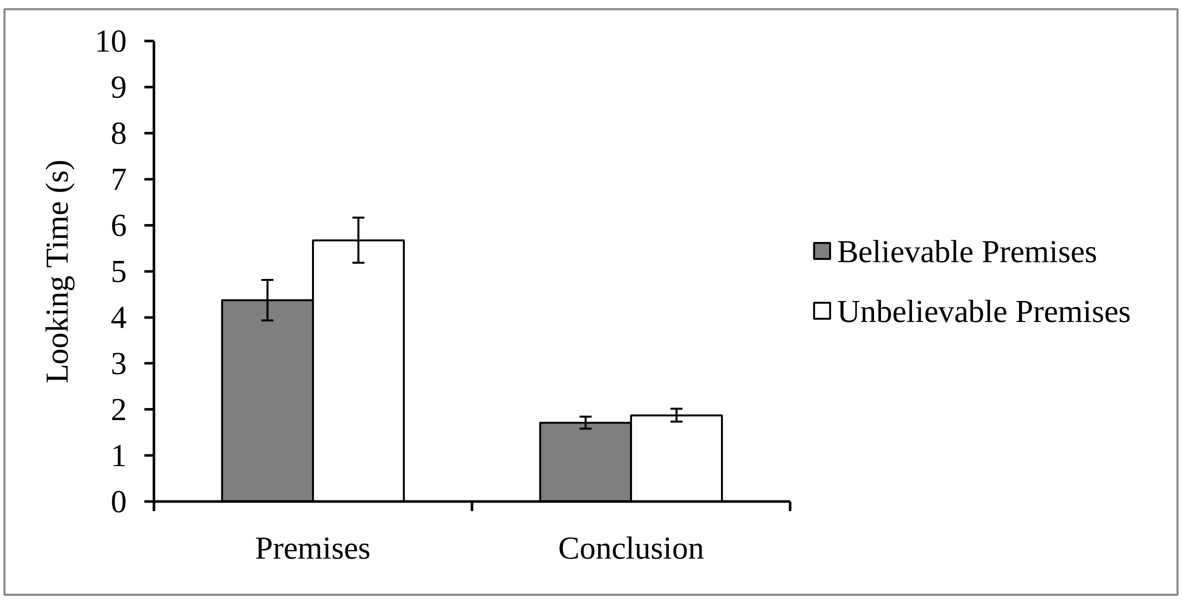

Figure 27. Total looking time to Premises and Conclusion a function of Premise Believability for Conditional and Categorical Syllogisms.

such that participants spent longer looking at syllogisms with unbelievable premises than those with believable premises. This main effect was qualified by an interaction between Premise Believability and AOI, $F(1,24) = 11.52$, $MSE = 1.281$, $p < .01$, $\eta^2 = .324$, which indicates that the longer looking time to syllogisms with unbelievable premises over believable premises is more evident on premise looking time, $t(24) = 4.23$, $p < .001$, than on conclusion looking time, $t(24) = 2.68$, $p < .05$.

Finally, there were main effects of Type, $F(1,24) = 21.16$, $MSE = 2.467$, $p < .001$, $\eta^2 = .469$, and AOI, $F(1,24) = 68.90$, $MSE = 11.967$, $p < .001$, $\eta^2 = .742$, on gaze duration. Not surprisingly, participants spent longer looking at the pair of premises than at the single conclusion, and spent longer looking at Conditional Syllogisms than Categorical Syllogisms.

Gaze Durations as a function of Pre- versus Post- Conclusion Fixation. Next, data was explored in a way that could examine differences in the temporal progression of reasoning about Conditional and Categorical Syllogisms. Gaze duration for premises was examined as a function of when it occurred: before or after the first fixation on the conclusion. Gaze duration on premises was divided into two values; the amount of time that participants spent fixated on the premises before and after conclusion fixation. Figure 28 depicts total premise looking time pre- and post- conclusion fixation as a function of Premise Believability for Conditional and Categorical Syllogisms. These values were submitted to a 2 (Type: Conditional, Categorical) x 2 (Premise Believability: Believable, Unbelievable) x 2 (Order: Pre-Conclusion, Post-Conclusion) repeated measures ANOVA. There was a significant Type x Premise Believability x Order interaction, $F(1,24) = 12.48$, $MSE = .714$, $p < .01$, $\eta^2 = .362$. Here, Premise Gaze Duration for Conditional and Categorical Syllogisms was submitted to two separate 2 (Premise Believability: Believable, Unbelievable) x 2 (Order: Pre-Conclusion, Post-Conclusion) repeated measures

Conditional Syllogisms

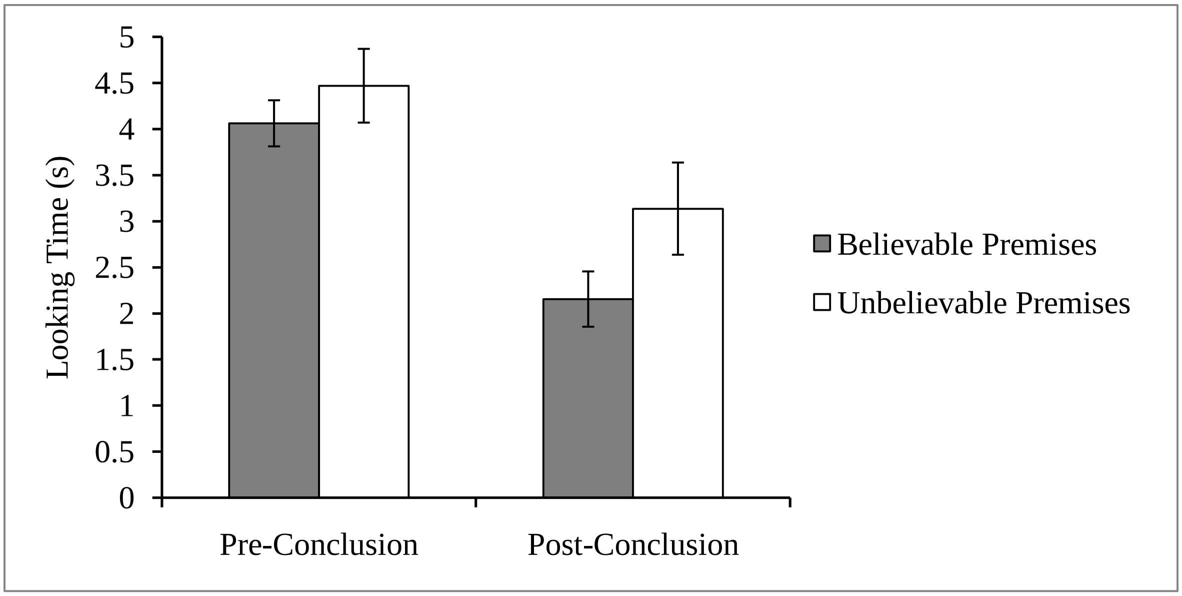

Categorical Syllogisms

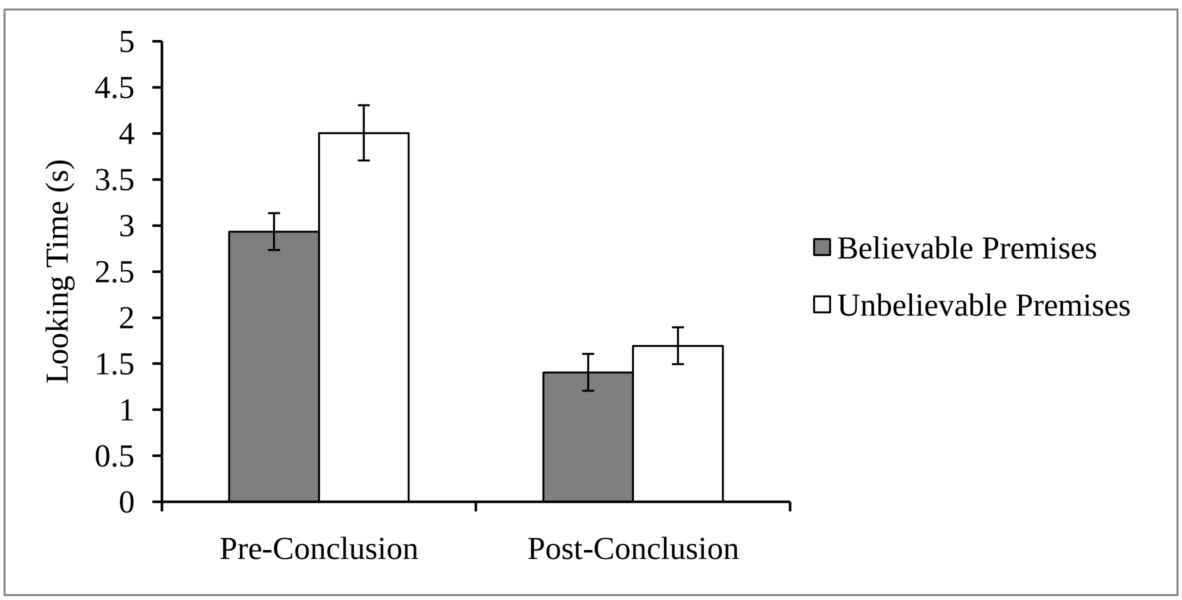

Figure 28. Total looking time to Premises pre- and post-conclusion fixation as a function of Premise Believability for Conditional and Categorical Syllogisms.

ANOVAs. The interaction between Premise Believability and Order was marginally significant for Conditional Syllogisms, $F(1,24) = 3.35$, $MSE = .556$, $p = .081$, $\eta^2 = .132$, and was significant for Categorical Syllogisms, $F(1,24) = 7.87$, $MSE = 1.028$, $p = .010$, $\eta^2 = .264$. To explore these relations, paired-samples *t*-tests were carried out. For Conditional Syllogisms, Premise Believability affected Post-Conclusion looking time, $t(24) = 2.64$, $p <.05$, but not Pre-Conclusion looking time, $t(24) = 1.25$, $p = .226$. For Categorical Syllogisms, the converse was true: Premise Believability affected Pre-Conclusion looking time, $t(24) = 6.32$ $p < .001$, but not Post-Conclusion looking time, $t(24) = 0.26$, $p = .796$.

Gaze Transitions from Conclusion to Premises. To further explore premise processing after at least some reasoning had occurred, transitions from the conclusion to the premises were analyzed. Figures 29 and 30 depict average number of transitions from the conclusion to premises as a function of Conclusion Validity and Believability for Conditional Syllogisms with Believable and Unbelievable Premises, and Categorical Syllogisms with Believable and Unbelievable Premises, respectively. Frequencies of transitions from conclusions to premises were submitted to a 2 (Type: Valid, Invalid) x 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVA. There was a main effect of Type, $F(1,24) = 41.99$, $MSE = .979$, $p <.001$, $\eta^2 = .618$, indicating that, overall, participants made more transitions from the conclusion to premises for Conditional than for Categorical Syllogisms.

There was a significant Premise Believability x Conclusion Validity x Conclusion Believability interaction, $F(1,24) = 4.65$, $MSE = .553$, $p <.05$, $\eta^2 = .152$. Two (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVAs revealed that there was a significant interaction between Conclusion

Conditional Syllogisms with Believable Premises

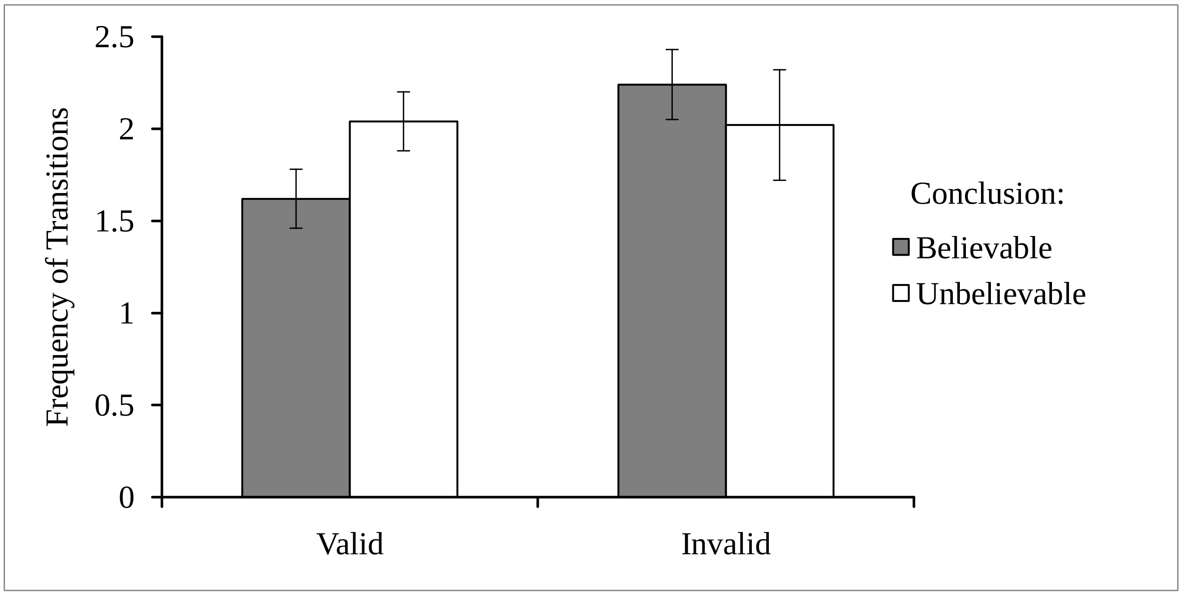

Conditional Syllogisms with Unbelievable Premises

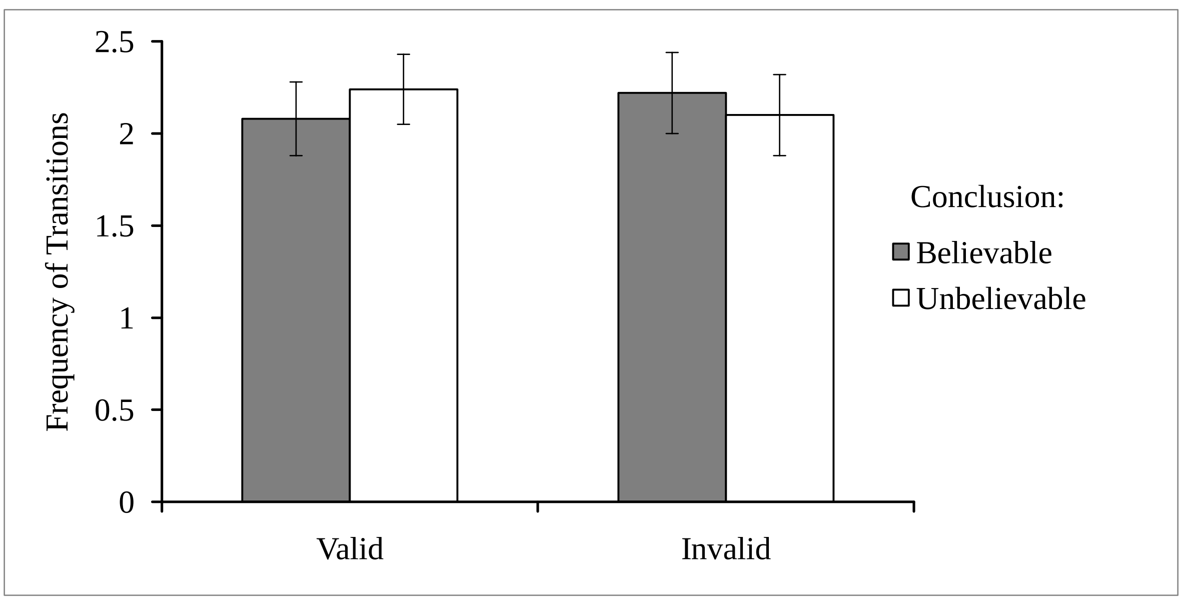

Figure 29. Frequency of transitions from conclusion to premises as a function of Conclusion Validity and Believability for Conditional Syllogisms *with* Believable and Unbelievable Premises.

Categorical Syllogisms with Believable Premises

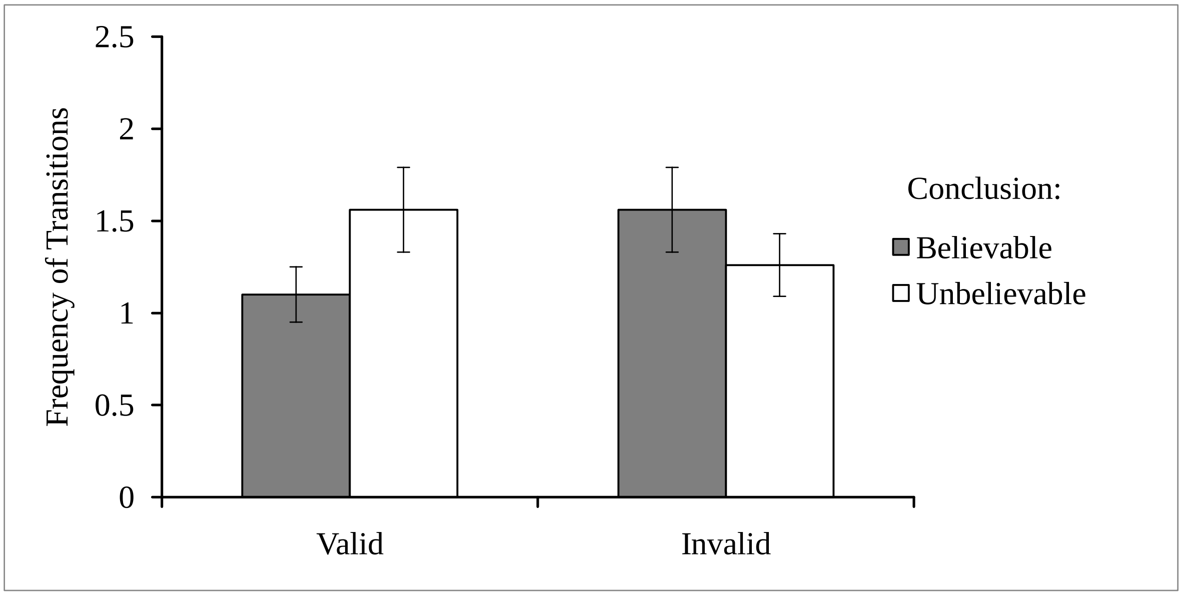

Categorical Syllogisms with Unbelievable Premises

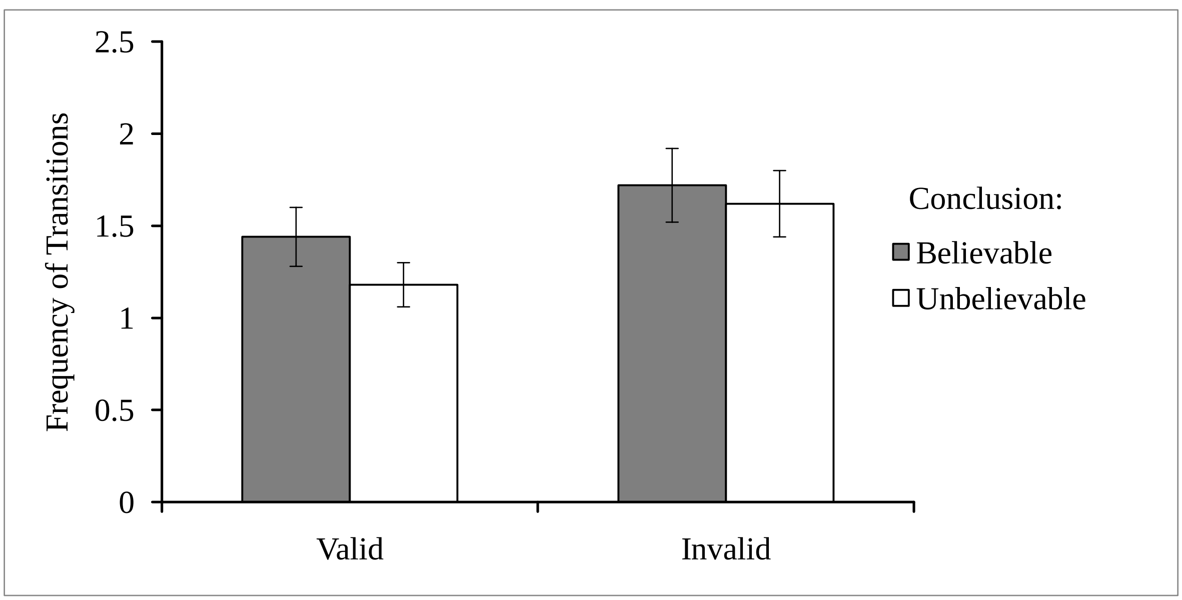

Figure 30. Frequency of transitions from conclusion to premises as a function of Conclusion Validity and Believability for Categorical Syllogisms *with* Believable and Unbelievable Premises.

Validity and Believability, $F(1,24) = 9.09$, $MSE = 1.350$, $p <.01$, $\eta^2 = .259$, when premises were Believable, indicating that participants referred back to the premises more often for conflict problems than for non-conflict problems. This interaction was not significant, however, $F(1,24) =0.07$, $MSE = 1.242$, $p = .792$, $\eta^2 = .003$ when premises were Unbelievable.

Given the hypothesis that there should be no effect of conclusion believability on gaze transition frequency for Categorical Syllogisms with Unbelievable Premises, transition data for Categorical Syllogisms were submitted to a 2 (Premise Believability: Believable, Unbelievable) x 2 (Conclusion Validity: Valid, Invalid) x 2 (Conclusion Believability: Believable, Unbelievable) repeated measures ANOVA. There was a marginally significant Premise Believability x Conclusion Validity x Conclusion Believability interaction, $F(1,24) = 3.76$, $MSE = .716$, $p =.063$, $\eta^2 = .126$. For Believable Premises, there was a significant interaction between Conclusion Validity and Believability, $F(1,24) = 4.55$, $MSE = .807$, $p <.05$, $\eta^2 = .149$. This interaction was not significant for Unbelievable Premises, $F(1,24) = 0.35$, $MSE = .460$, $p =.558$, $\eta^2 = .013$. There was also no main effect of Conclusion Believability, $F(1,24) = 2.06$, $MSE = .398$, $p =.162$, $\eta^2 = .074$.

Memory Data

For each participant, the number of words recalled from each syllogism was tallied. The average number of words recalled for Categorical and Conditional Syllogisms is depicted in Figure 31 as a function of Premise Believability. Number of words recalled was submitted to a 2 (Type: Conditional, Categorical) x 2 (Premise: Believable, Unbelievable) repeated measures ANOVA. Critically, there was a significant interaction between Syllogism Type and Premise Believability, $F(1,36) = 9.68$, $MSE = .4.247$, $p <.01$, $\eta^2 = .212$. Paired-samples *t*-tests indicated that, for Conditional Syllogisms, there was no difference in participants' memory performance as

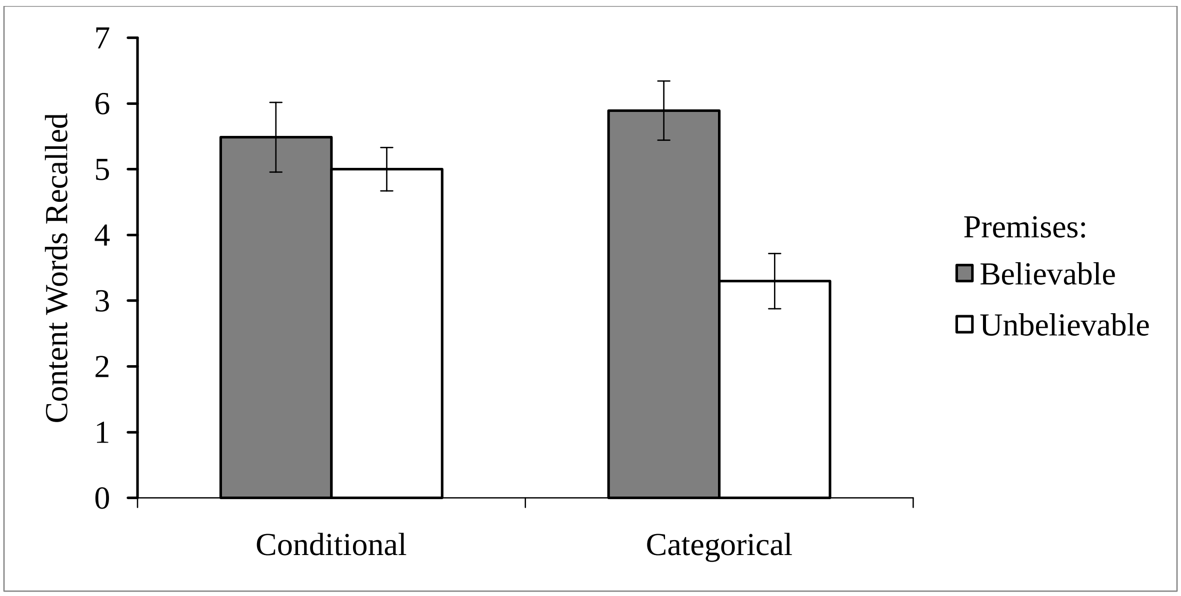

Figure 31. Average number of words recalled from Conditional and Categorical Syllogisms as a function of Premise Believability.

a function of premise believability, $t(37) = 0.88$, $p = .384$. However, for Categorical Syllogisms, participants had significantly better memory for content from syllogisms with Believable than Unbelievable Premises, $t(37) = 6.58$, $p < .001$.[2] There was also a main effect of Premise Believability, $F(1,36) = 20.64$, $MSE = 4.255$, $p < .001$, $\eta^2 = .364$, such that participants recalled more words from syllogisms with Believable than Unbelievable premises.

Discussion

Eye-tracking results will be discussed first, followed by memory results.

Eye-tracking

To test hypotheses regarding the time course of the reasoning process, eye movements were monitored while participants engaged in syllogistic reasoning. Although Experiments 1 through 4 provided evidence that reasoning for categorical and conditional syllogisms differed, discrepancies in the time course of the reasoning process were yet unknown. The predictions offered by my proposed theory of premise evaluation were supported. First, premise believability affected premise looking duration *before* the conclusion was attended to for categorical syllogisms and *after* the conclusion was attended to for conditional syllogisms. This is strong evidence that the believability of categorical premises is evaluated prior to reasoning about the conclusion and that the believability of conditional premises is evaluated after conclusion reasoning. Note that even for categorical syllogisms, participants looked longer at unbelievable premises. This is in line with the theory of Spinoza and Gilbert (1991) that statements are

[2] There was concern that items in Categorical and Conditional syllogisms with Believable and Unbelievable Premises may be differentially memorable for reasons other than how they are treated during the reasoning process. To address this concern, Twenty-eight participants who did not participate in the reasoning portion of the experiment were recruited to serve as a control group. These participants were simply presented with the premises and conclusions from the syllogisms used in the reasoning task and asked to recall as many content words as possible. Critically, in the control group, the Type x Premise Believability interaction was not significant, $F(1,27) = 1.82$, $MSE = 5.013$, $p = .188$, $\eta^2 = .063$, indicating that recall of believable and unbelievable statements was comparable across Categorical and Conditional Syllogisms. Thus, the differential memory performance following the reasoning task was not merely an artefact of the stimuli.

automatically accepted and then discredited in a time-consuming process. This also corroborates experiments that suggest that unbelievable information takes longer to comprehend than believable information (Carpenter & Just, 1975). Thus, it may not be that reasoners do not automatically accept categorical statements, but that they simply “disbelieve” unbelievable premises immediately, prior to reasoning about the conclusion.

Second, participants were more likely to transition back to the premises after looking at the conclusion for conditional than for categorical syllogisms. Again, this suggests that more premise processing happens after reasoning for conditional than for categorical premises. Further, there was evidence that for conditional syllogisms and categorical syllogisms with believable premises, participants referred back to the premises more often when conclusion validity conflicted with conclusion believability. However, this was not the case for categorical syllogisms with unbelievable premises. Yet again, there is evidence that conclusion believability does not play a role in reasoning about categorical syllogisms with unbelievable premises.

Memory for Syllogisms

It was theorized that perhaps reasoners disregard content of categorical syllogisms with unbelievable premises. If so, then participants should have reduced memory for words from these syllogisms than from categorical syllogisms with believable premises and conditional syllogisms. This was indeed the case: Participants had poorer memory for words from categorical syllogisms with unbelievable premises than from the other three types of problems. This suggests that unbelievable categorical premises trigger participants to treat the syllogism as an abstract syllogism, thereby disregarding the content of the syllogism. Examining memory for syllogisms to elucidate mechanisms underlying reasoning—in particular, the depth with which certain components of syllogisms are processed—is a novel approach. This experiment shows that this approach has merit that should be exploited in other future research.

General Discussion

The study of deductive reasoning has a rich history in both philosophy and psychology. Although phenomena associated with syllogistic reasoning, such as belief bias, have been studied for decades, there remain many unanswered questions concerning the mechanisms underlying the reasoning process. For example, the major theories of deductive reasoning (which account for the belief bias effect), including the Mental Models account (Johnson-Laird & Byrne, 1991) and Dual Process accounts (Evans, 2003, 2008), have not explicitly distinguished between mechanisms underlying categorical and conditional syllogistic reasoning. In fact, experiments designed to inform these theories use categorical syllogisms (e.g., Newstead et al., 1992; Klauer et al., 2000) and conditional syllogisms (Thompson, 1996; Torrens et al., 1999; Santamaria et al., 1998) somewhat interchangeably. Given that similar belief bias effects are found for conditional and categorical syllogisms, it was not evident that theorizing different mechanisms was necessary.

Recently, however, theorists are beginning to acknowledge that there may, in fact, be differences in how reasoners treat conditional and categorical syllogisms. For example, evidence now exists that different brain areas are activated during reasoning about categorical and conditional syllogisms (Reverberi et al., 2010), and that task demands, such as introducing time pressure, differentially influence conditional and categorical reasoning (Evans & Bacon, 2009). This dissertation adds to this growing body of evidence by assessing the degree to which manipulations of premise believably differentially influence reasoning with categorical and conditional syllogisms. Here, differential effects of premise believability are used as a tool to make inferences about mechanisms underlying reasoning with these two types of syllogisms. In doing so, a theory was formulated that addresses how categorical and conditional premises are treated differently.

Theory of Premise Evaluation

This theory postulates that categorical and conditional syllogisms differ in how and when premises are evaluated. Drawing on the work of Gilbert (1991; Gilbert et al., 1993; Gilbert et al., 1990) which suggests that statements are automatically assumed to be true, it was theorized that conditional premises are assumed to be true regardless of their actual veracity. As such, the conclusion of conditional syllogisms is reasoned about under this assumption and, after reasoning has occurred, people may return to the premises and disbelieve them in a procedure that requires time.

This theory is supported by the finding that the traditional belief bias effect occurs for conditional syllogisms with both believable and unbelievable premises; that is, reasoners do not seem to treat conclusions following these premises differently (Experiments 1-5). Further, reasoners take longer to reason about conditional syllogisms with unbelievable premises than with believable premises (Experiments 2-4), which is consistent with the postulation that it takes time to "disbelieve" what was initially accepted as true. Eye-tracking data (Experiment 5) support the proposed time course of the reasoning process: Reasoners seem to accept conditional premises as true and then evaluate them *after* reasoning about the conclusion, as evidenced by effects of premise believability influencing post-conclusion looking time rather than pre-conclusion looking time. Participants were also more likely to refer back to the premises after looking at the conclusion for conditional than for categorical syllogisms. With this post-reasoning evaluation process comes a bias for participants to endorse more conclusions following from believable than unbelievable premises.

Although the exact mechanism through which this post-reasoning evaluation process occurs is not fully known, there is other evidence in the literature that has shown that reasoners are less likely to draw or endorse conclusions following from false or unsound premises than from true or sound premises (Simoneau & Markovits, 2003; George, 1995, 1997, 1999). For example, George (1995) found that reasoners were less certain that a conclusion was valid when they were also uncertain that premises were true. Thus, realizing that premises are unbelievable may reduce reasoners' confidence in their valid response, leading them to judge that a conclusion is invalid. Asking participants to rate their confidence in their validity judgements (cf. Shynkaruk & Thompson, 2006) would likely reveal lower confidence for conclusions following from unbelievable premises.

Finally, Experiment 4 indicates that it is the word 'if' in conditional premises that causes reasoners to assume that premises are true. This small word has a rich meaning: It invites reasoners to assume, or imagine, that what follows is true or possible (Evans & Over, 2004). Thus, instructions for conditional syllogisms may not need to specify that participants need to assume that premises are true, as is the standard; these instructions are likely redundant because the truth assumption appears to be the default.

The data for categorical syllogisms support a different theory of premise evaluation. It is theorized that reasoners evaluate the believability of categorical premises immediately upon reading them, prior to reasoning about the conclusion. If premises are true, then all information in the syllogism is taken into account. However, if premises are false, then the content in the problem is ignored or not processed as deeply. Experiments 1 through 5 indicate that belief bias is suppressed when premises are unbelievable; that is, conclusion believability does not affect conclusion endorsement. The absence of the belief bias effect is striking, given that this phenomenon has repeatedly been shown to be extremely robust. Response latency evidence (Experiment 2-4) also indicates that conclusion believability does not affect reasoning time and does not interfere with responding when conclusion believability conflicts with validity. Perhaps the most compelling evidence for this theory comes from Experiment 5, which demonstrated that participants had significantly poorer memory for words from categorical syllogisms with unbelievable premises relative to other syllogisms. This finding supports the theory proposed here and the notion alluded to by Cherubini et al. (1998): Reasoners disregard the content of categorical syllogisms with unbelievable premises.

Mechanisms Underlying the Differential Role of Premise Belief

The specific mechanisms through which this disregarding occurs is not fully known. One possibility is that reasoners simply do not fully encode the content when they are reasoning with syllogisms because they abstract the logical form from the problem; another possibility is that reasoners may actively inhibit the content in the syllogism. There is some evidence for both of these possibilities. Markovits and Lortie-Forgues (in press) report that children who are more accurate at conditional reasoning with unbelievable premises are also more accurate at abstract conditional reasoning, suggesting that the ability to abstract the form from syllogisms may benefit reasoning. Markovits and colleagues (Markovits et al., 2009) also present evidence that inhibition is at work. They presented participants with problems with a major conditional premise and minor categorical premise and asked them to generate conclusions that followed from these premises. They report a curious finding: When the major premise was undeniably false, reasoners were more likely to draw valid, yet unbelievable, conclusions than when the major premise was plausible. They interpret these findings to suggest that, under instructions to assume that premises are true, reasoners inhibit information presented in false premises. Because these concepts are inhibited, reasoners find it easier to draw unbelievable conclusions.

Although inhibition may be necessary under some circumstances, the results of Markovits et al. (2009) do not necessarily transfer to categorical syllogisms with unbelievable premises. First, they used problems with a one major conditional premise and a minor premise (e.g., *If p, then q; p; therefore q*). Although the same mechanisms may underlie reasoning of this form and syllogistic reasoning, there is not yet evidence that this is the case. Second, they claim that it is the instruction to assume that premises are true which elicits inhibition of false information; the experiments presented here demonstrate that reasoners do not take the believability of conclusions following from categorical, unbelievable premises into account even in the absence of this instruction. It seems, then, that there is something about the unbelievable categorical statement itself that leads reasoners to disregard information in the syllogisms in which they are embedded. Regardless of the mechanism, it is clear that reasoners identify unbelievable premises early on in the reasoning process, which then trigger reasoners not to utilize the content of the problem when reasoning.

It is worth noting here that although reasoners do not use the believability of the conclusion to inform their judgment about the validity of the conclusion, they are not necessarily more accurate. That is, they do not compensate for this by using only the validity of the conclusion in their judgment. This is to be expected if reasoners abstract the logical form from the syllogism and then attempt to reason abstractly: Previous work has shown that abstract reasoning is quite difficult (Markovits & Vachon, 1990). Moreover, it is also possible that, in attempting to abstract the logical form, errors are sometimes made that influence subsequent judgments.

Implications for Reasoning Theories

Mental Models theory. The experiments presented here have implications for both reasoning theories and theories of information comprehension. First, the data can speak to the two most influential theories of deductive reasoning: the Mental Models account (Johnson-Laird & Byrne, 1991) and Dual Process accounts. The Mental Models account specifies that the reasoning process starts with a model of the premises, and then the conclusion is tested against this initial model. If the conclusion is believable, reasoners are less likely to search for alternate models that invalidate the conclusion than if the conclusion is unbelievable. This theory makes specific predictions about response latencies that have been fully explored elsewhere (Thompson et al., 2003; Ball et al., 2006). Briefly, because reasoners are thought to search for counter models only when conclusions are unbelievable, unbelievable conclusion should elicit longer response latencies than believable conclusions. In general, the response latency data reported here (along with those of Thompson et al. and Ball et al.), in which participants generally spend longer on problems where the validity and believability of the conclusion conflict, do not support a Mental Models account of reasoning.

More relevant to the current thesis, however, is what can be said about how premises are modelled. Critical assumptions of the Mental Models account of reasoning are the *principle of truth*, which states that, reasoners "construct mental models that represent explicitly only what is true, and not what is false" (Johnson-Laird, 1999, p. 8), and the related *principle of semantic modulation* of conditionals, which asserts that, "the meanings of the antecedent and consequent, and coreferential links between these two clauses, can add information to models, prevent the construction of otherwise feasible models of the core meaning, and aid the process of constructing fully explicit models" (Johnson-Laird & Byrne, 2002, p. 13). Thus, Johnson-Laird and colleagues claim that meaning modulates the representation of premises, and that models elicited by false premises are blocked (Johnson-Laird, 2001).

Contrary to these assumptions is Evans and Overs' (2004) *suppositional theory,* of which they state, "our view of 'if' is that of a linguistic device, the purpose of which is to trigger a process of hypothetical thinking and reasoning" (p. 153). Evans and Over assert that 'if*'* invites reasoners to assume that the content of the conditional premise is true. More specifically, the suppositional theory maintains that reasoners assume that the antecedent of the conditional statement is true, and then judge the probability of the consequent given the antecedent. Individuals do this by taking into account context and the meaning of the conditional. Thus, Evans and Over ascribe unique value to the meaning of the word 'if*'*, but do not assert that it elicits total acceptance of following statement.

The current experiments provide some evidence that can speak to these principles. In typical experiments of deductive reasoning, participants are instructed to assume that premises are true, although four of the five experiments presented here did not include this instruction. In these experiments, belief bias was not found for categorical syllogisms with unbelievable premises, and it is postulated that for these syllogisms, participants extracted the abstract logical form (i.e., stripped the content from the problem). Belief bias resulted for all conditional syllogisms, and it is posited here that the believability of premises acts *post-reasoning,* with participants having a bias to accept conclusions following from believable premises.

It is conceivable, then, that the proposals of Evans and Over (2004) and of Johnson-Laird (2001; Johnson-Laird & Byrne, 2002) are both correct, albeit in different circumstances. In partial accordance with Evans and Over, conditional syllogisms have a built-in instruction of sorts to assume that premises are true: the word 'if,' although it is proposed here that 'if' is even more powerful than Evans and Over assert and causes the reasoner to temporarily assume that both the antecedent *and* consequent are true. Indeed, these experiments yield evidence that both believable and unbelievable premises were successfully modelled when reasoning with conditional syllogisms, as similar belief bias effects resulted regardless of premise believability. For categorical syllogisms on the other hand, which lack this invitation to think hypothetically, reasoners behaved as if they represented unbelievable premises abstractly. Thus, it seems as though participants were only able to reason about meaningful premises when they were true or were assumed to be true (as encouraged through instruction or through the word 'if'). The data from these experiments also uphold this *principle of truth* – it seems as though we do not mentally model information that we know to be false.

The Mental Models account places great emphasis on conclusion believability, asserting that conclusion believability either encourages or inhibits modelling of the premises. However, the data presented here suggest that the reverse relation may also be true: Premise believability may determine how one reasons about the conclusion. Specifically, when premises are false, a model is created with abstract information, which may in turn lead reasoners to abstract the conclusion. In a sense, reasoners are less biased about the conclusion when reasoning about false, categorical, information. Given that reasoning about the premises themselves is the first step in the Mental Models account, it appears that greater attention should be paid to how the content of the premises influences later reasoning about the conclusion.

Dual Process theories. Dual Process accounts of reasoning have also been extremely influential and controversial. As reviewed previously, Dual Process accounts differ on specifics, but most agree that there are two systems that influence reasoning: System 1 is thought to be automatic, quick, resource-independent, and to respond according to heuristics or surface features of stimuli; whereas System 2 is thought to be resource-dependent, slow, and to respond according to logical analysis of stimuli. Dual Process theories differ in how they delineate the relations between these two systems. Default-interventionist models (Kahneman & Frederick, 2002, 2005; Evans, 2006) propose that the response provided by System 1 is always the default response and that, given enough time and resources, System 2 can subsequently override this default response. Parallel-Process theories (e.g., Sloman, 1996) posit that both systems are activated together and operate in parallel, and that the response output will depend on motivation, external factors, and resources. All Dual Process theories posit that System 1 provides a response based on the believability of the conclusion, whereas System 2 provides a response based on the validity of the conclusion.

Overall, the response latency data from the present experiments support general Dual Process accounts, which predict that when the validity and believability of the conclusion conflicts, reasoners will need to spend some time resolving this conflict. For most syllogisms, participants spent longer reasoning about conclusions in which the validity and believability of the conclusion conflicted, suggesting that there was some need to resolve conflicting responses. Interestingly, however, categorical syllogisms with unbelievable premises did not demonstrate this pattern. For these syllogisms, there is no evidence in response latencies, conclusion endorsement patterns, or eye-tracking data of a conflict between two different systems. In fact, there is only evidence that conclusion validity factored into responding, which is thought to be the responsibility of System 2.

This poses a problem for both Default-Interventionist models and Parallel Process models. Both of these accounts state that System 1 provides an automatic response; as Evans and Over (2004) state, “people will sometimes rely on System 1 processes and sometimes much more on System 2 processes (although we hold that the latter processes must always depend on the former to some extent)" (p.24). However, there is no evidence of System 1 involvement for categorical syllogisms with unbelievable premises. It is possible that System 2 always overrode the response provided by System 1, however, this is unlikely given that this is thought to require time, and participants did not spend longer overall on these syllogisms. Rather, it seems as though the unbelievable premises (when participants were not encouraged to assume their truth via instruction or 'if') pre-empted activation of System 1, perhaps because the unbelievable information in categorical syllogisms triggered reasoners to assume that responses provided by System 1, that is, responses based on conclusion believability, would likely be inaccurate. This compelling finding suggests that System 1 may not be automatic at all, and that factors such as context and the task at hand may influence whether System 1 gets activated. Therefore, whether this activation is under volitional control is a question that should be addressed in future research.

Implications for Theories of Information Comprehension

The theory of Spinoza, championed by Gilbert (1991), that to comprehend a statement we must automatically assume that it is true, was used to inform my proposed theory of premise evaluation. Given the data here, what can be said about this theory? The way information is framed seems to influence how people proceed with it, as has been shown in other situations. Specifically, people seem to readily assume that conditional statements are true, and to suspend judgment about the veracity of conditional statements until after the information has been used to reason about a conclusion. On the other hand, people either do not automatically assume that categorical premises are true, or assume that they are true and then immediately disbelieve false statements unless instructed to do otherwise. Support for the latter hypothesis can be drawn from the eye-tracking results of Experiment 5, which show that, prior to looking at the conclusion, participants spent longer looking at unbelievable categorical premises than at believable categorical premises, consistent with the hypothesis that it takes time to disbelieve false information (Gilbert, 1991; Just & Carpenter, 1976). Regardless, the data clearly reveal that conditional and categorical statements are treated differently. Crucially, conditional statements are assumed to be true until at least some reasoning about the conclusion has taken place whereas categorical statements are evaluated prior to reasoning about the conclusion.

Aligning the above discussion of the Mental Model account's *principle of truth* with Gilbert and Spinoza's ideas, there is even more evidence that one must believe something to be true in order to understand it. Johnson-Laird (1999, 2001) postulates that we do not model what is not true, and Gilbert and colleagues (1990) go even further, stating that "the mental representation of a proposition or idea always has a truth value associated with it, and by default this value is *true"* (p. 2). Thus, Gilbert would maintain that it is impossible to represent what is not true. The data in this thesis support this idea: When participants were not encouraged to assume that premises were true (either via instruction or with the word 'if'), they could not represent premises without presumably abstracting the form, as evidenced by their poor memory for the content of categorical syllogisms with unbelievable premises.

Conclusion

The study of deductive reasoning has a rich history which, curiously, has for the most part overlooked the role of premises in a syllogism in favour of the conclusion. This thesis demonstrates that varying the believability of the premises not only yields valuable insights about mechanisms underlying reasoning in general, but also brings to light a key difference between categorical and conditional syllogisms. Specifically, premises in these two types of syllogisms are treated differently; in turn, this influences how the conclusion is judged. For conditional syllogisms, premises, preceded by ‘if,’ are presumed to be true until at least some reasoning about the conclusion has been carried out. On the other hand, categorical premises encourage reasoners to evaluate their believability immediately. If the premises are judged false, then reasoners are in a sense less biased, as they disregard believability information in the rest of the syllogism.

Taken together, the data contained in this thesis add to a small but growing body of literature that is delineating key differences underlying categorical and conditional reasoning. Any comprehensive reasoning theory must account for independent mechanisms underlying the reasoning process for conditional and categorical syllogisms, rather than treating them similarly, as is the current trend. More generally, this thesis corroborates theories of belief bias and indicates that the knowledge and beliefs held by an individual can greatly influence how they reason about information. The theory that I propose suggests that this knowledge not only influences how individuals draw conclusions from information, but how their beliefs about the information itself interact with the format in which the information is presented (i.e., conditionally or categorically, to influence reasoning). Specifically, the data presented here suggests that information presented in a hypothetical manner will be accepted for the duration of the reasoning process, whereas information presented in a factual manner will be critically evaluated immediately. Further, individuals are less likely to be biased when reasoning about false information presented in a factual manner than when reasoning about information that can be interpreted as hypothetical.

Returning to the example from the introduction, a doctor who wants her patient to be less influenced by his beliefs about the uses of aspirin for headaches would be wise to instruct him, “people on this medication cannot take aspirin,” as opposed to “if you take this medication, then do not take aspirin.” Here, even if the patient does not believe the doctor’s orders, he would be less likely to take his beliefs about what is appropriate given a headache into account when drawing a conclusion about whether he is allowed to take aspirin or not. Conversely, advertisers promoting flimsy claims should tell consumers, “If you use our product, you will be happy,” rather than “people who use our product are happy.” Consumers are likely to accept the former claim as they reason about whether they should buy the product or not, whereas consumers are likely to immediately evaluate—and reject—the latter claim. Thus, both the content *and* the format of statements influence how individuals reason about them. Further study of variables that may independently influence reasoning with conditional and categorical syllogisms will surely elucidate new and unique ways that individuals reason about information presented in different ways.

References

Ball, L. J., Lucas, E. J., Miles, J. N. V., & Gale, A. G. (2003). Inspection times and the selection task: What do eye-movements reveal about relevance effects? *Quarterly Journal of Experimental Psychology*, *56*(6), 1053-1077.

Ball, L. J., Phillips, P., Wade, C. N., & Quayle, J. D. (2006). Effects of belief and logic on syllogistic reasoning: Eye-movement evidence for selective processing models. *Experimental Psychology*, *53*, 77-86.

Barnett, D. (2006). Zif is if. *Mind: A Quarterly Review of Philosophy*, *115*(459), 519-565.

Barnett, D. (2010). Zif would have been if: A suppositional view of counterfactuals. *Noûs, 44*(2), 269-304.

Bethell-Fox, C. E., Lohman, D. F., & Snow, R. E. (1984). Adaptive reasoning: Componential and eye movement analysis of geometric analogy performance. *Intelligence*, *8*(3), 205-238.

Byrne, R. M. (1989). Suppressing valid inferences with conditionals. *Cognition, 31*(1), 61-83.

Byrne, R. M. J., & Johnson-Laird, P. N. (2009). 'If' and the problems of conditional reasoning. *Trends in Cognitive Sciences, 13*(7), 282-287.

Carpenter, P. A., & Just, M. A. (1975). Sentence comprehension: A psycholinguistic processing model of verification. *Psychological Review, 82*(1), 45-73.

Cherubini, P., Garnham, A., Oakhill, J., & Morley, E. (1998). Can any ostrich fly? Some new data on belief bias in syllogistic reasoning. *Cognition*, *69*(2), 179-218.

Descartes, R. (1984). Principles of philosophy. In J. Cottingham, R. Stoothoff, & D. Murdoch (Eds. And Trans.), *The philosophical writings of Descartes* (Vol. 1). Cambridge, UK: Cambridge University Press. (Original work published 1644).

De Neys, W. (2006). Automatic-heuristic and executive-analytic processing during reasoning: Chronometric and dual-task considerations. *The Quarterly Journal of Experimental Psychology, 59*(6), 1070 – 1100.

Espino, O., Santamaría, C., Meseguer, E., & Carreiras, M. (2005). Early and late processes in syllogistic reasoning: Evidence from eye-movements. *Cognition, 98*(1), B1-B9.

Evans, J. St. B. T. 1989. *Bias in human reasoning: Causes and consequences*. Brighton: Erlbaum.

Evans, J. St. B. T. (2003). In two minds: Dual-process accounts of reasoning. *Trends in Cognitive Sciences, 7*(10), 454-459.

Evans, J. St. B. T. (2005). *Deductive reasoning*. New York, NY, US: Cambridge University Press.

Evans, J. St. B. T. (2006). The heuristic-analytic theory of reasoning: Extension and evaluation. *Psychonomic Bulletin & Review, 13*(3), 378-395.

Evans, J. St. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, *59*, 255-278.

Evans, J. St. B. T., & Curtis-Holmes, J. (2005). Rapid responding increases belief bias: Evidence for the dual-process theory of reasoning. *Thinking & Reasoning, 11*(4), 382-389.

Evans, J. St. B. T. & Over, D.E. (2004). *If*. Oxford: Oxford University Press.

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*, 77-83.

Evans, J. St. B. T., Over, D. E., & Handley, S. J. (2005). Suppositions, extensionality, and conditionals: A critique of the mental model theory of Johnson-Laird and Byrne (2002). *Psychological Review, 112*(4), 1040-1052.

George, C. (1995). The endorsement of the premises: Assumption-based or belief-based reasoning. *British Journal of Psychology, 86*(1), 93-111.

George, C. (1997). Reasoning from uncertain premises. *Thinking & Reasoning, 3*, 161-189.

George, C. (1999). Evaluation of the plausibility of a conclusion derivable from several arguments with uncertain premises. *Thinking & Reasoning, 5*(3), 245-281.

Gilbert, D. T. (1991). How mental systems believe. *American Psychologist, 46*(2), 107-119.

Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology, 59*(4), 601-613.

Gilbert, D. T., Tafarodi, R. W., & Malone, P. S. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology, 65*(2), 221-233.

Goodman, N. (1947). The problem of counterfactual conditions. *Journal of Philosophy, 44*, 113-128.

Grant, E. R., & Spivey, M. J. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science, 14*(5), 462-466.

Hadjichristidis, C., Handley, S. J., Sloman, S. A., Evans, J. St. B. T., Over, D. E., & Stevenson, R. J. (2007). Iffy beliefs: Conditional thinking and belief change. *Memory & Cognition, 35*(8), 2052-2059.

Hasson, U., Simmons, J. P., & Todorov, A. (2005). Believe it or not: On the possibility of suspending belief. *Psychological Science, 16*(7), 566-571.

Johnson-Laird, P. N. (1988). How is meaning mentally represented? In U. Eco, M. Santambrogio, & P. Viloli (Eds.), *Meaning and mental representations* (pp. 99-118). Bloomington: Indiana University Press.

Johnson-Laird, P. N. (1999). Deductive reasoning. *Annual Review of Psychology, 50*, 109-135.

Johnson-Laird, P. N. (2011) The truth about conditionals. In K. Manktelow, D. Over, & S. Elqayam (Eds.), *The science of reason* (pp. 119-144). New York, NY: Psychology Press.

Johnson-Laird, P. N., & Byrne, R. M. J. (1991). *Deduction*. East Sussex, UK: Laurence Erlbaum Associates Ltd.

Johnson-Laird, P. N., & Byrne, R. M. J. (2002). Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review, 109*(4), 646-678.

Just, M. A., & Carpenter, P. A. (1976). The relation between comprehending and remembering some complex sentences. *Memory & Cognition, 4*(3), 318-322.

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). New York, NY, US: Cambridge University Press.

Kahneman, D., & Frederick, S. (2005). *A model of heuristic judgment*. In K. J. Holyoak & R. G. Morrison (Eds.), *The Cambridge handbook of thinking and reasoning* (pp. 267-294). New York, NY: Cambridge University Press.

Körner, C., & Gilchrist, I. D. (2004). Eye movements in a simple spatial reasoning task. *Perception, 33*, 485-494.

Lewis, D. (1973). Counterfactuals and comparative possibility. *Journal of Philosophical Logic, 2*, 418-446.

Mackie, J. L. (1973). *Truth, probability and paradox: Studies in philosophical logic*. Oxford, UK: Clarendon Press.

Markovits, H., & Lortie-Forgues, H. (in Press). Conditional reasoning with false premises facilitates the transition between familiar and abstract reasoning. *Child Development.*

Markovits, H., Saelen, C., & Forgues, H. L. (2009). An inverse belief–bias effect: More evidence for the role of inhibitory processes in logical reasoning. *Experimental Psychology*, *56*, 112-120.

Markovits, H., & Schroyens, W. (2007). A curious belief-bias effect: Reasoning with false premises and inhibition of real-life information. *Experimental Psychology, 54*, 38-43.

Markovits, H., & Vachon, R. (1990). Conditional reasoning, representation, and level of abstraction. *Developmental Psychology*, *26*(6), 942-951.

Oakhill, J., Johnson-Laird, P. N., & Garnham, A. (1989). Believability and syllogistic reasoning. *Cognition, 31*(2), 117-140.

Politzer, G., & Bourmaud, G. (2002). Deductive reasoning from uncertain conditionals. *British Journal of Psychology*, *93*(3), 345-381.

Quelhas, A. C., Johnson-Laird, P. N., & Juhos, C. (2010). The modulation of conditional assertions and its effects on reasoning. *Quarterly Journal of Experimental Psychology, 63*(9), 1716-1739.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, *124*(3), 372-422.

Reverberi, C., Cherubini, P., Frackowiak, R. S. J., Caltagirone, C., Paulesu, E., & Macaluso, E. (2010). Conditional and syllogistic deductive tasks dissociate functionally during premise integration. *Human Brain Mapping, 31*(9), 1430-1445.

Santamaría, C., García-Madruga, J. A., & Johnson-Laird, P. N. (1998). Reasoning from double conditionals: The effects of logical structure and believability. *Thinking & Reasoning, 4*(2), 97-122.

Shynkaruk, J. M., & Thompson, V. A. (2006). Confidence and accuracy in deductive reasoning. *Memory & Cognition*, *34*(3), 619-632.

Simoneau, M., & Markovits, H. (2003). Reasoning with premises that are not empirically true: Evidence for the role of inhibition and retrieval. *Developmental Psychology*, *39*(6), 964-975.

Spinoza, B. (1982). *The Ethics and selected letters*. (S. Feldman, Ed. and S. Shirley, Trans.). Indianapolis, IN: Hackett. (Original work published 1677).

Stanovich K. E. (1999). *Who is rational? Studies of individual differences in reasoning*. Mahwah, NJ: Lawrence Elrbaum Associates.

Stupple, E. J. N., & Ball, L. J. (2008). Belief-logic conflict resolution in syllogistic reasoning: Inspection-time evidence for a parallel-process model. *Thinking & Reasoning, 14*(2), 168-181.

Thompson, V. A. (1996). Reasoning from false premises: The role of soundness in making logical deductions. *Canadian Journal of Experimental Psychology, 50*(3), 315-319.

Thompson, V. A., Striemer, C. L., Reikoff, R., Gunter, R. W., & Campbell, J. I. D. (2003). Syllogistic reasoning time: Disconfirmation disconfirmed. *Psychonomic Bulletin & Review, 10*(1), 184-189.

Torrens, D., Thompson, V. A., & Cramer, K. M. (1999). Individual differences and the belief bias effect: Mental models, logical necessity, and abstract reasoning. *Thinking & Reasoning, 5*(1), 1-28.

Appendix A

Experiment 1 Stimuli and Pilot Ratings

Note: For the pilot test of the stimuli, 40 participants completed an online measure in which they rated the believability of each statement on a scale from 1 (Very Unbelievable) to 5 (Very Believable).

A) Categorical Syllogisms

a. Believable Premises

	Believable Conclusion	Belief Rating	Unbelievable Conclusion	Belief Rating
Valid	No dogs are felines	4.00	No puppies are cobras	4.31
	Some furry animals are felines	3.78	Some snakes are cobras	4.86
	Therefore, some furry animals are not dogs	4.50	Therefore, some snakes are not puppies	2.87
	No actresses are men	4.42	No birds are salmon	4.56
	Some singers are men	4.67	Some fish are salmon	4.55
	Therefore, some singers are not actresses	4.44	Therefore, some fish are not birds	2.67
Invalid	No teenagers are elderly	4.02	No cod are angelfish	4.20
	Some males are elderly	4.90	Some fish are angelfish	4.48
	Therefore, some teenagers are not males	4.52	Therefore, some cod are not fish	1.43
	No grandparents are children	4.21	No roses are weeds	3.80
	Some females are children	4.48	Some plants are weeds	4.31
	Therefore, some grandparents are not females	4.54	Therefore, some roses are not plants	1.94

b. Unbelievable Premises

	Believable Conclusion	Belief Rating	Unbelievable Conclusion	Belief Rating
Valid	No black animals are rodents	1.77	No animals are dogs	1.15
	Some cats are rodents	1.61	Some lions are dogs	1.31
	Therefore, some cats are not black animals	4.02	Therefore, some lions are not animals	1.16
	No poisonous animals are insects	1.57	No flowers are daisies	1.73
	Some snakes are insects	1.38	Some daffodils are daisies	1.88
	Therefore, some snakes are not poisonous	4.76	Therefore, some daffodils are not flowers	1.81
Invalid	No books are novels	1.69	No butterflies are insects	1.31
	Some magazines are novels	2.09	Some rabbits are insects	1.52
	Therefore, some books are not magazines	3.98	Therefore, some butterflies are not rabbits	2.69
	No vehicles are buses	1.44	No monkeys are primates	1.35
	Some convertibles are buses	1.93	Some alligators are primates	1.46
	Therefore, some vehicles are not convertibles	4.86	Therefore, some alligators are not monkeys	2.31

B) Conditional Syllogisms

a. Believable Premises

	Believable Conclusion	Belief Rating	Unbelievable Conclusion	Belief Rating
Valid	If an animal is a bird, then it has a beak	4.28	If a vehicle is a motorcycle, then it has wheels	4.58
	If an animal has a beak, then it has feathers	3.90	If a vehicle has wheels, then it has a steering wheel	3.38
	Therefore, if an animal is a bird, then it has feathers	3.85	Therefore, if a vehicle is a motorcycle, then it has a steering wheel	2.13
	If an animal is a tiger, then it eats meat	3.92	If an animal is a whale, then it is a mammal	4.52
	If an animal eats meat, then it is a carnivore	3.98	If an animal is a mammal, then it is a land-dweller	3.31
	Therefore, if an animal is a tiger, then it is a carnivore	4.23	Therefore, if an animal is a whale, then it is a land-dweller	1.34
Invalid	If an animal is a feline, then it purrs	3.97	If an object is a grape, then it is a fruit	4.25
	If an animal purrs, then it is a cat	3.74	If an object is a fruit, then it is a food	4.12
	Therefore, if an animal is a cat, then it is a feline	4.67	Therefore, if an object is a food, then it is a grape	2.06
	If an animal is an amphibian, then it croaks	3.81	If an object is an iPod, then it is a music player	4.00
	If an animal croaks, then it is a frog	3.50	If an object is a music player, then it is electronic	4.90
	Therefore, if an animal is a frog, then it is an amphibian	4.38	Therefore, if an object is electronic, then it is an iPod	1.38

b. Unbelievable Premises

	Believable Conclusion	Belief Rating	Unbelievable Conclusion	Belief Rating
Valid	If a person is a pianist, then he is a professor	1.81	If a person is a father, then they are a doctor	4.58
	If a person is a professor, then he is a musician	2.06	If a person is a doctor, then they are a woman	1.62
	Therefore, if a person is a pianist, then he is a musician	4.21	Therefore, if a person is a father, then they are a woman	1.34
	If an animal is a puppy, then it is cold-blooded	1.23	If an animal is a flamingo, then it is a cat	1.57
	If an animal is cold-blooded, then it is a canine	1.50	If an animal is a cat, then it is a reptile	1.73
	Therefore, if an animal is a puppy, then it is a canine	3.92	Therefore, if an animal is a flamingo, then it is a reptile	1.48
Invalid	If a food is candy, then it is meat	1.47	If a plant is a weed, then it is an oak tree	1.50
	If a food is meat, then it is a lollipop	1.14	If a plant is an oak tree, then it is a flower	1.78
	Therefore, if a food is a lollipop, then it is candy	4.31	Therefore, if a plant is a flower, then it is a weed	1.78
	If an animal is cold-blooded, then it is a squirrel	1.05	If a food is a cookie, then it is a vegetable	1.30
	If an animal is a squirrel, then it is a reptile	1.55	If a food is a vegetable, then it is a bean	1.29
	Therefore, if an animal is a reptile, then it is cold-blooded	4.38	Therefore, if a food is a bean, then it is a cookie	1.38

Appendix B

Experiments 2-5 Stimuli and Pilot Ratings

Note: For the pilot test of the stimuli, 20 participants completed an online measure in which they rated the believability of each statement on a scale from 1 (Very Unbelievable) to 5 (Very Believable).

A) Categorical Syllogisms

a. Believable Premises

	Believable Conclusion	Belief Rating	Unbelievable Conclusion	Belief Rating
Valid	All slacks are pants	3.85	All rats are pests	2.12
	All pants are clothing	4.38	All pests are insects	2.07
	Therefore, all slacks are clothing	3.92	Therefore, All rats are insects	1.00
	No dogs are felines	4.00	No puppies are cobras	4.31
	Some furry animals are felines	3.78	Some snakes are cobras	4.86
	Therefore, some furry animals are not dogs	4.50	Therefore, some snakes are not puppies	2.87
Invalid	All infants are newborns	3.38	All toddlers are children	4.03
	All newborns are babies	4.38	All children are humans	4.38
	Therefore, all babies are infants	3.79	Therefore, all humans are toddlers	1.79
	No grandparents are children	4.21	No cod are angelfish	4.20
	Some females are children	4.48	Some fish are angelfish	4.48
	Therefore, some grandparents are not females	4.54	Therefore, some cod are not fish	1.43

b. Unbelievable Premises

	Believable Conclusion	Belief Rating	Unbelievable Conclusion	Belief Rating
Valid Single Model	All reptiles are rabbits	1.50	All cherries are tools	1.06
	All rabbits are cold-blooded	1.62	All tools are wooden things	2.12
	Therefore, all reptiles are cold-blooded	4.19	Therefore, all cherries are wooden things	1.21
Multi-Model	No black animals are rodents	1.77	No flowers are daisies	1.73
	Some cats are rodents	1.61	Some daffodils are daisies	1.88
	Therefore, some cats are not black animals	4.02	Therefore, some daffodils are not flowers	1.81
Invalid Single Model	All copper things are toys	1.57	All diamonds are inexpensive	1.43
	All toys are pennies	1.55	All inexpensive things are gemstones	2.00
	Therefore, all pennies are copper things	3.99	Therefore, all gemstones are diamonds	1.80
Multi-Model	No vehicles are buses	1.44	No monkeys are primates	1.35
	Some convertibles are buses	1.93	Some alligators are primates	1.46
	Therefore, some vehicles are not convertibles	4.86	Therefore, some alligators are not monkeys	2.31

B) Conditional Syllogisms

a. Believable Premises

	Believable Conclusion	Belief Rating	Unbelievable Conclusion	Belief Rating
Valid Single Model	If an animal is a bird, then it has a beak	4.28	If a vehicle is a motorcycle, then it has wheels	4.58
	If an animal has a beak, then it has feathers	3.90	If a vehicle has wheels, then it has a steering wheel	3.38
	Therefore, if an animal is a bird, then it has feathers	3.85	Therefore, if a vehicle is a motorcycle, then it has a steering wheel	2.13
Difficult	If a room is dirty, then it is not clean	4.48	If a person is an employee, then she is not retired	4.08
	If a room is not clean, then it is messy	3.63	If a person is not retired, then she is young	3.48
	Therefore, if a room is dirty, then it is messy	3.89	Therefore, if a person is an employee, then she is young	1.50
Invalid Single Model	If an animal is an amphibian, then it croaks	3.91	If an object is a grape, then it is a fruit	4.25
	If an animal croaks, then it is a frog	3.79	If an object is a fruit, then it is a food	4.12
	Therefore, if an animal is a frog, then it is an amphibian	4.46	Therefore, if an object is a food, then it is a grape	2.06
Difficult	If a person is a girl, then they are not a son	4.37	If it is sunny, then it is not night	3,92
	If a person is not a son, then they are a daughter	4.28	If it is not night, then it is day	3.80
	Therefore, if a person is a daughter, then they are a girl	4.57	Therefore, if it is day, then it is sunny	2.15

b. Unbelievable Premises

	Believable Conclusion	Belief Rating	Unbelievable Conclusion	Belief Rating
Valid Single Model	If a person is a pianist, then he is a professor	1.81	If a person is a father, then they are a doctor	4.58
	If a person is a professor, then he is a musician	2.06	If a person is a doctor, then they are a woman	1.62
	Therefore, if a person is a pianist, then he is a musician	4.21	Therefore, if a person is a father, then they are a woman	1.34
Difficult	If a food is a chilli pepper, then it is not a vegetable	1.61	If an object is a musical instrument, then it is not a piano	1.76
	If a food is not a vegetable, then it is spicy	2.19	If an object is not a piano, then it is not an accordion	1.59
	Therefore, if a food is a chilli pepper, then it is spicy	4.19	Therefore, if an object is a musical instrument, then it is an accordion	2.00
Invalid Single Model	If a food is candy, then it is meat	1.47	If a food is a cookie, then it is a vegetable	1.30
	If a food is meat, then it is a lollipop	1.14	If a food is a vegetable, then it is a bean	1.29
	Therefore, if a food is a lollipop, then it is candy	4.31	Therefore, if a food is a bean, then it is a cookie	1.38
Difficult	If an animal swims, then it is not a fish	1.28	If a person is not a hockey player, then he is rich	1.85
	If an animal is not a fish, then it is a dolphin	1.89	If a person is not rich, then he is an athlete	1.52
	Therefore, if an animal is a dolphin, then it swims	4.51	Therefore, if a person is an athlete, then he is a hockey player	1.29

Appendix C

Complete Overall ANOVA Results

Note: Results in bold were presented in the main text.

A. Experiment 1

a. Conclusion Endorsement Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	1, 68	0.02	.787	.892	<.001
Premise Believability	1, 68	0.11	.521	.740	.002
Conclusion Validity	**1, 68**	**79.36**	**.862**	**<.001**	**.575**
Conclusion Believability	**1, 68**	**48.44**	**.515**	**<.001**	**.416**
Type x Premise Believability	**1, 68**	**2.13**	**.435**	**.149**	**.030**
Type x Conclusion Validity	1, 68	0.45	.511	.503	.007
Type x Conclusion Believability	1, 68	6.23	.364	<.05	.084
Premise Believability x Conclusion Validity	1, 68	2.37	.391	.128	.034
Premise Believability x Conclusion Believability	1, 68	6.49	.295	<.05	.087
Conclusion Validity x Conclusion Believability	**1, 68**	**9.95**	**.368**	**<.01**	**.136**
Type x Premise Believability x Conclusion Validity	1, 68	0.47	.497	.497	.007
Type x Premise Believability x Conclusion Believability	**1, 68**	**3.84**	**.341**	**.054**	**.053**
Type x Conclusion Validity x Conclusion Believability	1, 68	0.63	.284	.432	.009
Premise Believability x Conclusion Validity x Conclusion Believability	1, 68	0.01	.323	.916	<.001
Type x Premise Believability x Conclusion Validity x Conclusion Believability	1, 68	0.10	.323	.752	.001

B. Experiment 2

a. Conclusion Endorsement Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	1, 39	2.81	.160	.101	.067
Models	1, 39	0.06	.200	.804	.002
Premise Believability	1, 39	4.80	.407	<.05	.110
Conclusion Validity	**1, 39**	**37.23**	**.498**	**<.001**	**.488**
Conclusion Believability	**1, 39**	**54.39**	**.341**	**<.001**	**.582**
Type x Models	1, 39	1.88	.167	.179	.046
Type x Premise Believability	**1, 39**	**19.02**	**.066**	**<.001**	**.328**
Type x Conclusion Validity	1, 39	2.83	.217	.101	.068
Type x Conclusion Believability	1, 39	0.23	.192	.134	.057
Models x Premises Believability	1, 39	0.36	.138	.551	.009
Models x Conclusion Validity	**1, 39**	**5.21**	**.117**	**<.05**	**.118**
Models x Conclusion Believability	1, 39	8.15	.203	<.01	.173
Premise Believability x Conclusion Validity	1, 39	2.17	.117	.149	.053
Premise Believability x Conclusion Believability	1, 39	48.30	.131	<.001	.553
Conclusion Validity x Conclusion Believability	**1, 39**	**1.57**	**.162**	**.218**	**.039**
Type x Models x Premise Believability	**1, 39**	**0.70**	**.112**	**.408**	**.018**
Type x Models x Conclusion Validity	**1, 39**	**0.63**	**.125**	**.433**	**.016**
Type x Models x Conclusion Believability	1, 39	1.78	.143	.190	.044
Type x Premise Believability x Conclusion Validity	1, 39	1.99	.157	.166	.049
Type x Premise Believability x Conclusion Believability	**1, 39**	**24.25**	**.167**	**<.001**	**.383**
Type x Conclusion Validity x Conclusion Believability	**1, 39**	**6.55**	**.122**	**<.05**	**.144**
Models x Premise Believability x Conclusion Validity	1, 39	1.06	.189	.310	.026
Models x Premise Believability x Conclusion Believability	1, 39	1.49	.134	.230	.037
Models x Conclusion Validity x Conclusion Believability	1, 39	3.52	.128	.068	.083
Premise Believability x Conclusion Validity x Conclusion Believability	1, 39	0.60	.131	.445	.015
Type x Models x Premise Believability x Conclusion Validity	1, 39	10.68	.155	<.05	.215
Type x Models x Premise Believability x Conclusion Believability	**1, 39**	**1.35**	**.113**	**.252**	**.034**
Type x Models x Conclusion Validity x	1, 39	0.02	.165	.891	<.001

Conclusion Believability					
Type x Premise Believability x Conclusion Validity x Conclusion Believability	1, 39	2.71	.166	.108	.065
Models x Premise Believability x Conclusion Validity x Conclusion Believability	1, 39	2.03	.154	.162	.050
Type x Models x Premise Believability x Conclusion Validity x Conclusion Believability	1, 39	0.03	.091	.854	.001

b. Response Latency Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	1, 39	31.87	1.103×10^8	<.001	.450
Models	1, 39	8.99	7.109×10^7	.005	.187
Premise Believability	1, 39	27.88	8.254×10^7	<.001	.417
Conclusion Validity	1, 39	6.01	1.806×10^8	<.05	.134
Conclusion Believability	1, 39	0.17	8.398×10^7	.684	.004
Type x Models	1, 39	23.84	6.777×10^7	<.001	.379
Type x Premise Believability	**1, 39**	**5.66**	$\mathbf{1.334 \times 10^8}$	**<.05**	**.127**
Type x Conclusion Validity	1, 39	15.06	5.244×10^7	<.001	.279
Type x Conclusion Believability	1, 39	0.64	8.436×10^7	.430	.016
Models x Premises Believability	1, 39	0.06	1.703×10^7	.812	.001
Models x Conclusion Validity	1, 39	1.45	5.307×10^7	.235	.036
Models x Conclusion Believability	1, 39	2.19	8.207×10^7	.147	.053
Premise Believability x Conclusion Validity	1, 39	7.25	1.030×10^8	.01	.157
Premise Believability x Conclusion Believability	1, 39	4.72	4.873×10^7	<.05	.108
Conclusion Validity x Conclusion Believability	1, 39	10.10	9.438×10^7	<.01	.206
Type x Models x Premise Believability	**1, 39**	**0.27**	$\mathbf{6.160 \times 10^7}$	**.602**	**.007**
Type x Models x Conclusion Validity	1, 39	0.12	8.262×10^7	.734	.003
Type x Models x Conclusion Believability	1, 39	3.11	1.242×10^8	.086	.074
Type x Premise Believability x Conclusion Validity	1, 39	0.51	4.371×10^7	.479	.013
Type x Premise Believability x Conclusion Believability	1, 39	0.77	6.318×10^7	.385	.019
Type x Conclusion Validity x Conclusion Believability	**1, 39**	**4.66**	$\mathbf{8.753 \times 10^7}$	**<.05**	**.107**
Models x Premise Believability x Conclusion Validity	1, 39	0.31	4.599×10^7	.579	.008
Models x Premise Believability x Conclusion Believability	1, 39	4.57	1.061×10^8	<.05	.105
Models x Conclusion Validity x Conclusion Believability	1, 39	1.87	6.853×10^7	.179	.046
Premise Believability x Conclusion Validity	1, 39	7.96	5.201×10^8	<.01	.169

x Conclusion Believability					
Type x Models x Premise Believability x Conclusion Validity	1, 39	5.98	7.852×10^7	<.05	.133
Type x Models x Premise Believability x Conclusion Believability	1, 39	6.34	8.629×10^7	<.05	.140
Type x Models x Conclusion Validity x Conclusion Believability	1, 39	1.83	8.590×10^7	.162	.049
Type x Premise Believability x Conclusion Validity x Conclusion Believability	**1, 39**	**12.09**	$\mathbf{1.103 \times 10^8}$	**.001**	**.237**
Models x Premise Believability x Conclusion Validity x Conclusion Believability	1, 39	1.35	1.098×10^8	.252	.033
Type x Models x Premise Believability x Conclusion Validity x Conclusion Believability	**1, 39**	**0.07**	$\mathbf{1.019 \times 10^8}$	**.798**	**.002**

C. Experiment 3

a. Conclusion Endorsement Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	1, 83	21.80	.444	<.001	.208
Type x Condition	1, 83	1.93	.444	.168	.023
Premise Believability	1, 83	15.75	.264	<.001	.160
Premise Believability x Condition	1, 83	0.08	.264	.775	.001
Conclusion Validity	**1, 83**	**102.15**	**1.11**	**<.001**	**.552**
Conclusion Validity x Condition	1, 83	0.13	1.11	.717	.002
Conclusion Believability	**1, 83**	**74.84**	**.814**	**<.001**	**.474**
Conclusion Believability x Condition	1, 83	0.13	.814	.719	.002
Type x Premise Believability	**1, 83**	**19.40**	**.273**	**<.001**	**.189**
Type x Premise Believability x Condition	**1, 83**	**3.10**	**.273**	**.082**	**.036**
Type x Conclusion Validity	1, 83	1.43	.391	.235	.017
Type x Conclusion Validity x Condition	1, 83	0.21	.391	.647	.003
Type x Conclusion Believability	1, 83	0.50	.342	.500	.006
Type x Conclusion Believability x Condition	1, 83	1.66	.342	.201	.020
Premise Believability x Conclusion Validity	1, 83	1.43	.257	.235	.017
Premise Believability x Conclusion Validity x Condition	1, 83	0.27	.257	.607	.003
Premise Believability x Conclusion Believability	1, 83	7.05	.312	<.01	.078
Premise Believability x Conclusion Believability x Condition	1, 83	9.69	.312	<.01	.105
Conclusion Validity x Conclusion Believability	**1, 83**	**11.68**	**.351**	**.001**	**.123**
Conclusion Validity x Conclusion Believability x Condition	1, 83	.014	.351	.906	<.001
Type x Premise Believability x Conclusion Validity	1, 83	0.47	.403	.495	.006
Type x Premise Believability x Conclusion Validity x Condition	1, 83	0.47	.403	.495	.006
Type x Premise Believability x Conclusion Believability	**1, 83**	**21.43**	**.260**	**<.001**	**.205**
Type x Premise Believability x Conclusion Believability x Condition	**1, 83**	**3.68**	**.260**	**.058**	**.043**
Type x Conclusion Validity x Conclusion Believability	1, 83	0.20	.334	.659	.002
Type x Conclusion Validity x Conclusion Believability x Condition	1, 83	0.12	.334	.724	.002
Premise Believability x Conclusion Validity x Conclusion Believability	1, 83	0.64	.215	.426	.008
Premise Believability x Conclusion Validity	1, 83	0.48	.215	.493	.006

x Conclusion Believability x Condition					
Type x Premise Believability x Conclusion Validity x Conclusion Believability	1, 83	0.03	.290	.859	<.001
Type x Premise Believability x Conclusion Validity x Conclusion Believability x Condition	1, 83	1.551	.290	.216	.018

b. Response Latency Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	1, 83	37.03	3.049×10^8	<.001	.309
Type x Condition	1, 83	0.49	3.049×10^8	.486	.006
Premise Believability	1, 83	19.86	1.607×10^8	<.001	.193
Premise Believability x Condition	1, 83	3.49	1.607×10^8	.065	.040
Conclusion Validity	1, 83	16.52	2.454×10^8	<.001	.166
Conclusion Validity x Condition	1, 83	2.93	2.454×10^8	.091	.034
Conclusion Believability	1, 83	1.05	1.857×10^8	.309	.012
Conclusion Believability x Condition	1, 83	0.05	1.857×10^8	.827	.001
Type x Premise Believability	**1, 83**	**8.09**	**1.453×10^8**	**<.01**	**.089**
Type x Premise Believability x Condition	1, 83	0.22	1.453×10^8	.642	.003
Type x Conclusion Validity	1, 83	8.94	1.499×10^8	<.01	.097
Type x Conclusion Validity x Condition	1, 83	1.37	1.499×10^8	.245	.016
Type x Conclusion Believability	1, 83	2.22	1.773×10^8	.140	.026
Type x Conclusion Believability x Condition	1, 83	0.35	1.773×10^8	.558	.004
Premise Believability x Conclusion Validity	1, 83	8.72	1.300×10^8	<.01	.095
Premise Believability x Conclusion Validity x Condition	1, 83	1.26	1.300×10^8	.264	.015
Premise Believability x Conclusion Believability	1, 83	1.33	1.897×10^8	.252	.016
Premise Believability x Conclusion Believability x Condition	1, 83	0.18	1.897×10^8	.673	.002
Conclusion Validity x Conclusion Believability	1, 83	3.94	1.408×10^8	.05	.045
Conclusion Validity x Conclusion Believability x Condition	1, 83	0.24	1.408×10^8	.626	.003
Type x Premise Believability x Conclusion Validity	1, 83	0.27	1.362×10^8	.602	.003
Type x Premise Believability x Conclusion Validity x Condition	1, 83	3.24	1.362×10^8	.075	.038
Type x Premise Believability x Conclusion Believability	1, 83	3.10	1.184×10^8	.082	.036

Type x Premise Believability x Conclusion Believability x Condition	1, 83	0.72	1.184×10^8	.398	.009
Type x Conclusion Validity x Conclusion Believability	1, 83	2.61	1.147×10^8	.110	.031
Type x Conclusion Validity x Conclusion Believability x Condition	1, 83	0.22	1.147×10^8	.643	.003
Premise Believability x Conclusion Validity x Conclusion Believability	1, 83	1.97	1.712×10^8	.165	.023
Premise Believability x Conclusion Validity x Conclusion Believability x Condition	1, 83	0.01	1.712×10^8	.945	<.001
Type x Premise Believability x Conclusion Validity x Conclusion Believability	**1, 83**	**9.41**	$\mathbf{1.128 \times 10^8}$	**<.01**	**.102**
Type x Premise Believability x Conclusion Validity x Conclusion Believability x Condition	1, 83	0.96	1.128×10^8	.330	.011

D. Experiment 4

a. Conclusion Endorsement Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	1, 88	0.66	.424	.418	.007
Type x Condition	1, 88	0.01	.424	.928	<.001
Premise Believability	1, 88	0.84	.395	.361	.010
Premise Believability x Condition	1, 88	0.54	.395	.467	.006
Conclusion Validity	**1, 88**	**109.87**	**.591**	**<.001**	**.555**
Conclusion Validity x Condition	1, 88	1.66	.591	.201	.018
Conclusion Believability	**1, 88**	**110.02**	**1.104**	**<.001**	**.556**
Conclusion Believability x Condition	1, 88	0.38	1.104	.539	.004
Type x Premise Believability	**1, 88**	**13.94**	**.261**	**<.001**	**.137**
Type x Premise Believability x Condition	**1, 88**	**0.26**	**.261**	**.613**	**.003**
Type x Conclusion Validity	1, 88	2.43	.394	.123	.027
Type x Conclusion Validity x Condition	1, 88	1.58	.394	.212	.018
Type x Conclusion Believability	1, 88	0.44	.487	.439	.007
Type x Conclusion Believability x Condition	1, 88	6.08	.487	<.05	.065
Premise Believability x Conclusion Validity	1, 88	11.08	.202	.001	.112
Premise Believability x Conclusion Validity x Condition	1, 88	2.73	.202	.102	.030
Premise Believability x Conclusion Believability	**1, 88**	**11.42**	**.232**	**.001**	**.115**
Premise Believability x Conclusion Believability x Condition	1, 88	1.06	.232	.307	.012
Conclusion Validity x Conclusion Believability	1, 88	**10.32**	**.253**	**<.01**	**.105**
Conclusion Validity x Conclusion Believability x Condition	1, 88	0.60	.253	.441	.007
Type x Premise Believability x Conclusion Validity	1, 88	2.86	.252	.094	.031
Type x Premise Believability x Conclusion Validity x Condition	1, 88	0.04	.252	.851	<.001
Type x Premise Believability x Conclusion Believability	1, 88	11.85	.274	.001	.119
Type x Premise Believability x Conclusion Believability x Condition	**1, 88**	**1.64**	**.274**	**.203**	**.018**
Type x Conclusion Validity x Conclusion Believability	1, 88	4.45	.253	<.05	.048
Type x Conclusion Validity x Conclusion Believability x Condition	1, 88	6.654	.253	<.05	.070
Premise Believability x Conclusion Validity x Conclusion Believability	1, 88	0.26	.256	.610	.003
Premise Believability x Conclusion Validity	1, 88	0.00	.256	.963	<.001

x Conclusion Believability x Condition					
Type x Premise Believability x Conclusion Validity x Conclusion Believability	1, 88	0.14	.285	.708	.002
Type x Premise Believability x Conclusion Validity x Conclusion Believability x Condition	1, 88	0.02	.285	.878	<.001

b. Response Latency Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	1, 88	6.99	2.521×10^8	.010	.074
Type x Condition	1, 88	0.23	2.521×10^8	.634	.003
Premise Believability	1, 88	39.64	6.146×10^7	<.001	.311
Premise Believability x Condition	1, 88	4.72	6.146×10^7	<.05	.051
Conclusion Validity	1, 88	0.12	4.910×10^7	.733	.001
Conclusion Validity x Condition	1, 88	0.94	4.910×10^7	.335	.011
Conclusion Believability	1, 88	22.36	8.099×10^7	<.001	.203
Conclusion Believability x Condition	1, 88	0.49	8.099×10^7	.485	.006
Type x Premise Believability	**1, 88**	**1.98**	$\mathbf{6.133 \times 10^7}$	**.163**	**.022**
Type x Premise Believability x Condition	**1, 88**	**0.41**	$\mathbf{6.133 \times 10^7}$	**.526**	**.005**
Type x Conclusion Validity	1, 88	0.19	3.722×10^7	.668	.002
Type x Conclusion Validity x Condition	1, 88	0.00	3.722×10^7	.969	<.001
Type x Conclusion Believability	1, 88	9.58	6.808×10^7	<.01	.098
Type x Conclusion Believability x Condition	1, 88	0.21	6.808×10^7	.650	.002
Premise Believability x Conclusion Validity	1, 88	1.40	3.300×10^7	.240	.016
Premise Believability x Conclusion Validity x Condition	1, 88	0.13	3.300×10^7	.718	.001
Premise Believability x Conclusion Believability	1, 88	0.06	5.126×10^7	.800	.001
Premise Believability x Conclusion Believability x Condition	1, 88	1.82	5.126×10^7	.181	.020
Conclusion Validity x Conclusion Believability	1, 88	17.17	5.442×10^7	<.001	.163
Conclusion Validity x Conclusion Believability x Condition	1, 88	0.00	5.442×10^7	.976	<.001
Type x Premise Believability x Conclusion Validity	1, 88	0.31	5.764×10^7	.578	.004
Type x Premise Believability x Conclusion Validity x Condition	1, 88	0.04	5.764×10^7	.834	.001
Type x Premise Believability x Conclusion Believability	1, 88	0.00	5.670×10^7	.965	<.001

Type x Premise Believability x Conclusion Believability x Condition	1, 88	0.00	5.670×10^7	.955	<.001
Type x Conclusion Validity x Conclusion Believability	1, 88	0.43	4.107×10^7	.513	.005
Type x Conclusion Validity x Conclusion Believability x Condition	1, 88	0.01	4.107×10^7	.941	<.001
Premise Believability x Conclusion Validity x Conclusion Believability	1, 88	0.16	4.857×10^7	.692	.002
Premise Believability x Conclusion Validity x Conclusion Believability x Condition	1, 88	0.79	4.857×10^7	.375	.009
Type x Premise Believability x Conclusion Validity x Conclusion Believability	1, 88	0.77	5.678×10^7	.382	.009
Type x Premise Believability x Conclusion Validity x Conclusion Believability x Condition	**1, 88**	**1.86**	$\mathbf{5.678 \times 10^7}$	**.176**	**.021**

E. Experiment 5

a. Conclusion Endorsement Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	1, 38	8.20	.432	<.01	.178
Premise Believability	1, 38	2.78	.360	.104	.068
Conclusion Validity	**1, 38**	**7.74**	**1.715**	**<.01**	**.169**
Conclusion Believability	**1, 38**	**48.76**	**.581**	**<.001**	**.562**
Type x Premise Believability	**1, 38**	**8.20**	**.361**	**<.01**	**.178**
Type x Conclusion Validity	1, 38	0.72	.270	.402	.019
Type x Conclusion Believability	1, 38	2.37	.358	.132	.059
Premise Believability x Conclusion Validity	1, 38	10.03	.295	<.01	.209
Premise Believability x Conclusion Believability	1, 38	8.53	.316	<.01	.183
Conclusion Validity x Conclusion Believability	**1, 38**	**0.60**	**.452**	**.444**	**.016**
Type x Premise Believability x Conclusion Validity	1, 38	5.64	.348	<.05	.129
Type x Premise Believability x Conclusion Believability	**1, 38**	**4.95**	**.272**	**<.05**	**.115**
Type x Conclusion Validity x Conclusion Believability	1, 38	0.16	.247	.690	.004
Premise Believability x Conclusion Validity x Conclusion Believability	1, 38	2.71	.213	.108	.067
Type x Premise Believability x Conclusion Validity x Conclusion Believability	1, 38	3.37	.210	.074	.081

b. General Looking Time Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	**1, 24**	**21.16**	**2.467**	**<.001**	**.469**
Premise Believability	**1, 24**	**21.32**	**1.568**	**<.001**	**.470**
AOI	**1, 24**	**68.90**	**11.967**	**<.001**	**.742**
Type x Premise Believability	**1, 24**	**0.31**	**1.171**	**.581**	**.013**
Type x AOI	**1, 24**	**26.86**	**1.285**	**<.001**	**.528**
Premise Believability x AOI	**1, 24**	**11.51**	**1.281**	**<.01**	**.324**
Type x Premise Believability x AOI	1, 24	0.04	.932	.844	.002

c. Pre- vs. Post- Conclusion Looking Time Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	1, 24	20.22	2.132	<.001	.454
Premise Believability	1, 24	21.89	1.932	<.001	.498
Order	1, 24	120.14	1.063	<.001	.846
Type x Premise Believability	1, 24	0.34	1.571	.565	.014
Premises x Order	1, 24	1.22	.880	.280	.053
Type x Premise Believability x Order	**1, 24**	**12.48**	**.714**	**<.01**	**.362**

d. Gaze Transition from Conclusion to Premises Results

Variable(s)	*DF*	*F*	*MSE*	*p*	η^2
Type	**1, 24**	**41.99**	**.979**	**<.001**	**.618**
Premise Believability	1, 24	6.22	.361	<.05	.193
Conclusion Validity	1, 24	5.13	.678	<.05	.165
Conclusion Believability	1, 24	0.04	.689	.951	<.001
Type x Premise Believability	1, 24	0.16	.582	.687	.006
Type x Conclusion Validity	1, 24	0.19	.627	.664	.007
Type x Conclusion Believability	1, 24	0.52	.593	.475	.020
Premise Believability x Conclusion Validity	1, 24	0.01	.434	.930	<.001
Premise Believability x Conclusion Believability	1, 24	2.10	.347	.159	.075
Conclusion Validity x Conclusion Believability	1, 24	4.86	.743	<.05	.158
Type x Premise Believability x Conclusion Validity	1, 24	3.56	.602	.070	.121
Type x Premise Believability x Conclusion Believability	1, 24	0.22	.936	.644	.008
Type x Conclusion Validity x Conclusion Believability	1, 24	0.27	.546	.606	.010
Premise Believability x Conclusion Validity x Conclusion Believability	**1, 24**	**4.65**	**.553**	**<.05**	**.152**
Type x Premise Believability x Conclusion Validity x Conclusion Believability	1, 24	0.68	.747	.415	.026