ASCH AND AI: CONFORMITY TO NON-HUMAN INTELLIGENCE

A Thesis Presented for the

Master of Arts

Degree

The University of Tennessee, Knoxville

Andrew Stephen Heim

December 2020

Acknowledgements

I would like to thank both my mentor Dr. Garriy Shteynberg, and previous mentors Dr. John Pennington and Dr. Jon Frederick. Their expertise and feedback helped me to make my vision for this project possible at all levels.

Additionally, I would like to thank my lab mate Phillip McGarry and my thesis committee for their insights and feedback.

Finally, I would like to thank my family who has supported me immensely throughout this process. Thank you to my wife Hollee, parents Steve and Dulcie, sister Emily, and grandparents Kathy and Bernie, and Ed and Barbara.

ABSTRACT

Conformity is the process that occurs when we submit to group pressures. These pressures can come from normative social influence, a desire to be liked, and informational social influence, the belief that the group has information that we do not. In the current study, the classic Asch line judgment paradigm is combined with virtual reality technology to study social influence in groups of both humans and artificial intelligences. Additionally, the line judgment task is varied as either unambiguous or ambiguous. The results indicated that participants were more likely to not conform to unambiguous stimuli and artificial intelligence confederates. Response times also suggest that participants respond slower to normative social influence. In addition to providing a novel contribution to the conformity literature, the study suggests future directions for research using this paradigm.

Keywords: Conformity, Artificial Intelligence, Normative Social Influence, Informational Social Influence, Virtual Reality

TABLE OF CONTENTS

I. Introduction and General Information 1

II. Methods 10

Apparatus and Stimuli 11

Procedure 16

III. Results 18

IV. Discussion 34

References 43

Appendix 48

Vita 65

LIST OF TABLES

Table 1 19

Table 2 22

Table 3 26

Table 4 28

Table 5 29

Table 6 30

LIST OF FIGURES

Figure 1 13

Figure 2 14

Figure 3 15

Figure 4 20

Figure 5 27

Figure 6 31

Figure 7 32

Figure 8 33

I. Introduction and General Information

Introduction

While now considered classic literature in the field of social psychology, there are many unanswered questions and new applications for the study of conformity. Reaching back to as far as the 1930s and 1950s with Jeness, Sherif, and Asch's experiments, the research on conformity is nearing a full century of study. Many of these studies, especially those conducted by Asch, are now staples of general and social psychology lectures as the most settled of science. However, as the world has changed so too have our questions concerning conformity. The advent of modern technology and social media pose new and interesting scenarios and environments in which people may face conformity.

Now, instead of a small group of people, it could be hundreds or thousands of others trying to provoke conformity. Not only that, but in the virtual world, the agents attempting to elicit conformity may also not even be human. This thesis principally seeks to address both past conformity literature and the modern direction of study on the topic. Additionally, the thesis will experimentally utilize a classic conformity paradigm with a novel twist in order to investigate conformity in a new way. In doing so, the known effects of ambiguity increasing conformity will be combined with the unknown of varying the humanity of the confederates. This will begin to demonstrate to what extent

and under what conditions people may conform to non-human intelligences. In addition, potential future directions for research will be discussed.

Seminal Conformity Research

The earliest research concerning the topic of conformity is credited to the work of Dr. Arthur Jenness. He had participants hazard a guess at how many beans were in a jar. After guessing, participants were assigned to groups and were allowed to discuss and provide a group answer. Participants were then given the opportunity to change their first guess and it was found that they mostly did so to be closer to the group's estimate (Jeness, 1932). This study serves to demonstrate that when we are placed in an ambiguous situation, we tend to conform to group pressures in order to be correct.

A few years after this initial experiment in 1935, another experimental psychologist Dr. Muzafer Sherif utilized the autokinetic effect, that is, the visual illusion that a stationary spec of light appears to move in a dark room. Despite the fact that participants normally varied their answers when asked, Sherif found that a person's answer could be influenced towards the group's answer by placing an individual with a different answer in a group with others who had answered similarly (Sherif, 1935). This effect seemed, in his interpretation, to imply that people tend to bend towards group pressure rather than an individual answer when pressed.

The task used in Sherif's experiment was, by virtue of not having a real answer, inherently ambiguous. In later research conducted by Dr. Solomon Asch, a new paradigm was utilized in its place. In this paradigm, participants would be presented with 4 lines, 3 test lines labeled "A," "B," and "C" along with a comparison line. An example of how clearly different these lines were can be seen in Figure 1. Asch had lab assistants working as confederates who all gave a unanimously incorrect answer before the actual participant answered. In doing so, he was able to demonstrate that not only do we conform to group pressures, but we also conform to these pressures despite our better judgment (Asch, 1951). Through this paradigm, Asch, somewhat surprisingly, found that participants tended to publicly conform to the obviously incorrect opinion of the group about 36.8% of the time (Asch, 1955).

Another conformity-based research paradigm was developed by Dr. Richard Crutchfield in response to some of the criticisms of the Asch study such as the need for a large number of confederates and the physical pressure the group may have exerted on participants. In this experiment, several participants were recruited at once and were able to answer privately. Crutchfield was able to make each participant believe that they were answering last and that everyone else had given the incorrect answer, enabling him to run multiple participants at once. While Crutchfield found significantly less conformity than Asch, he still

found that participants were inclined to conform, even when they answered privately (Crutchfield, 1955).

This tendency to conform across experimental paradigms is generally agreed to be caused by a combination of normative and informational social influence depending on the particular situation. (Deutsch & Gerard, 1955). Deutsch and Gerard (1955) defined these two types of influence as follows:

> We shall define a normative social influence as an influence to conform with the positive expectations of another. An informational social influence may be defined as an influence to accept information obtained from another as evidence about reality. (p. 629)

Through expanding and further exploring Asch's paradigm, they were able to conclude that even in situations where participants are not normatively influenced, participants can still be influenced to conform due to informational social influence, i.e. that they trust the other group members are providing them with accurate information about reality (Deutsch & Gerard, 1955).

Asch himself found that if his stimulus lines were presented as more similar in length, thus increasing task ambiguity, individuals had an increased tendency to conform to the group answer due to informational social influence (Asch, 1956). Additionally, Asch explored other situational factors such as group size and the presence of an allied confederate. Generally, his findings indicated

that group size and conformity were positively correlated up to three group members and that the having just one other person stand up to the group decreased conformity significantly (Asch, 1956).

Virtual Reality and Conformity

The availability of Virtual Reality (VR) technology offers a unique method for reanalyzing and expanding upon the previous conformity research. Through the use of this technology, researchers are able to create virtual environments that mimic real ones (Blascovich, Beall, Swinth, Hoyt, & Bailenson, 2002). Additionally, virtual avatars can be used as confederates in the place of lab assistants. Instead of requiring the physical space and manpower to conduct such studies such as this, with one researcher, a laptop, and a VR headset, these studies can be reproduced and modified easily and at a significantly lower cost. This in turn allows for more efficient and extensive data collection as the researcher is no longer dependent upon the availability, training, and acting ability of confederates. After one voice and movement recording session, a researcher could, theoretically, continue running the same paradigm indefinitely. In this way, we are able to meet the challenge of the criticism that Asch's paradigm is just too difficult and time consuming to run.

As this technology is novel, there is some question about the generalizability of research conducted in a virtual environment. This reservation

is understandable, yet several studies have consistently shown both the ecological validity and ease of conducting research in virtual environments (De Mel, Carnevale, & Gratch, 2014). One study has gone so far as to demonstrate other fundamental concepts of social influence such as foot-in-the-door and door-in-the face function the same in a virtual environment as in reality (Eastwick & Gardner, 2009). It stands to reason then, that Asch's paradigm would be no different. Another consideration with this novel research tool is the unique set of ethical considerations. Health and safety, long-term behavioral manipulation, and possible misuse of findings by businesses are just a few of the many serious concerns that present themselves when using virtual reality. Fortunately, researchers are already working to codify a standard operating procedure for experiments in virtual environments (Madary & Metzinger, 2016).

Virtual Agents and Non-Human Intelligence

Another area of interest in the conformity research that has not been explored is conformity to artificial intelligence. Despite this, much research has been carried out concerning anthropomorphizing and trust of non-human intelligence. A study found that robots which were physically present were unable to exert the level of social influence found in the Asch paradigm. The researchers noted, however, that if the robots had looked or behaved in a more realistic manner, the result might have been different (Brandstetter et al., 2014). It

is also worth noting that this paradigm looked at the Asch paradigm with unambiguous stimuli and a verbal paradigm with ambiguous stimuli. Conformity rates were higher in the ambiguous task, but not significant. It was demonstrated that when anthropomorphized, virtual representatives not assumed to be controlled by humans have a higher chance of exerting social influence (Waytz, Cacioppo, & Epley, 2010). Another study explored this idea and attempted a replication of the original Asch experiment in VR with some modifications. The experiment found that even non-believable human virtual confederates were able to effect participant response time in an unambiguous perceptual task even though conformity rates were not as extreme as they were in Asch's experiments (Kyrlitsias & Michael-Grigoriou, 2016). This paradigm, with some modifications, offers an interesting and novel opportunity to study conformity due to both informational and normative social influence. Through this framework, several questions about the nature and extent of social influence can be analyzed, particularly by replacing the non-believable human confederates of Kyrlitsias & Michael-Grigoriou in the form of non-human or artificial intelligence. These agents, while intelligent, should not elicit the desire to be liked as in normative social influence, but there is a good chance that they may still be seen as useful for figuring out the best answer in an unclear situation through informational social influence. Generally, when people interact with

artificial intelligence, they have an understanding that it has been programmed to accomplish a particular goal or behave a certain way. It is seen more as a tool than as an entity that can make value judgements about us. Currently, most of these artificial intelligences do not necessarily have the ability to consciously "like" a person, and most are aware of that. Essentially, the artificial intelligence does not decide things of this nature, it only follows a program. However, that program may be more or less accurate in the case of informational social influence. Therefore, due to the prevailing understanding of artificial intelligences as they exist today, people should be quick to disregard them if they are inaccurate and quick to rely on them if they are or appear to be accurate. This study will begin to answer the question of how susceptible we are to the informational influence of artificial intelligence and under what conditions. As we continue to rely more and more on these types of intelligences, research in this area is necessary to understand how these intelligences may help or hinder us. This study will principally seek to establish a framework for using virtual reality to investigate the influence of non-human intelligence given its relevancy in our modern and fast changing world. As these intelligences continue to improve in quality and judgment ability, the framework will quickly allow for investigation into potential normative social influence as well.

The Current Study

In order to analyze the effect of humanity of confederates and ambiguity of stimuli, the current study utilized a 2 x2 factorial design. The first between-subject's variable was the ambiguity of the stimuli. The stimuli shown to participants were either all ambiguous or unambiguous. The other between-subject's variable is whether the participant is told the virtual confederates are human or an artificial intelligence. An example of an unambiguous stimuli can be seen in Figure 2, while an example of an ambiguous stimuli can be seen in Figure 3. This design allows for us to uncover not only if non-human intelligences can elicit conformity, but also in which types of situations. The general format of this experiment will consist of 6 practice trials to add realism to the group by their correct answers followed by 12 experimental trials in which confederates will attempt to elicit conformity. Based on the previous literature on the topics the following hypotheses are proposed:

H_1: Participants in the ambiguous conditions will conform significantly more than participants in the unambiguous conditions.

H_2: Participants in the human confederate conditions will conform significantly more than participants in the artificial intelligence confederate conditions.

H_3: The effect of humanity will be moderated by the effect of ambiguity in that participants will conform more to humans in the unambiguous condition.

H_4: Participants will respond slower to humans than to artificial intelligences as it is expected that the human confederates will elicit both normative and informational social influence compared to only the informational influence of the artificial intelligence.

H_5: Participants will respond slower on trials in which they resist group pressure due to reassessing and changing their initial answer to match that of the group.

H_6: Participants will respond slower on ambiguous trials compared to unambiguous trials due to the difficulty increase.

II. Methods

Participants

120 students taking general psychology, 82 females and 38 males, were recruited through the SONA system at the University of Tennessee. In order to motivate participation, participants were awarded half an hour of credit to be used in their general psychology course. Participants were recruited until instructors reached the unit on social psychology as participants would inevitably learn about the Asch study and its results. Each semester efforts were made to maintain contact with each instructor teaching general psychology, but most did not cover the material until the second half of a given semester. If an

instructor indicated they would be covering the material earlier, students in that course was removed from the participation pool prior to the date of coverage.

Apparatus and Stimuli

Oculus Rift. This is a virtual reality head mounted display (HMD). This HMD enables the participant to enter into the virtual environment in which the experiment takes place. It allows for head movement to be tracked enabling the participant to experience the environment as if they were there in reality.

Virtual Asch. This is a computer software that creates scenarios in which the original Asch experiment and variations of it can be conducted in virtual reality. The software enables the researcher to customize both the number of confederates and the stimuli that will be presented to the participant (Blom, 2015). An example of the Virtual Asch program can be seen in Figure 1.

Line Task. These stimuli were created using the Dungeon Painter Studio program. This program is generally useful for making maps, but it worked well for generating Asch-like stimuli for this experiment. Based on previous research on participant's ability to discriminate such stimuli, a stimulus was considered to be "unambiguous" if the line lengths differed by greater than 5% the length of the longest line. Conversely, a stimulus was considered to be "ambiguous" if the line lengths differed by 5% or less than the length of the longest line. Previous research had found that using a 15% length ratio was far too generous and did

not justify operationalizing "ambiguous" or "unambiguous" stimuli (Heim, 2018).

Voice Recordings. For the human confederates, voices were recorded by two males and two females. They recorded each letter response "A," "B," and "C" multiple times so that the responses would seem more natural to participants. For the artificial intelligence voices, different computerized male and female voices were utilized for each of the AI confederates. The voice recordings were the main distinguishing feature between the AI and human confederates, other variables such as movement and visual appearance were held constant.



Figure 1. Virtual Asch testing environment.

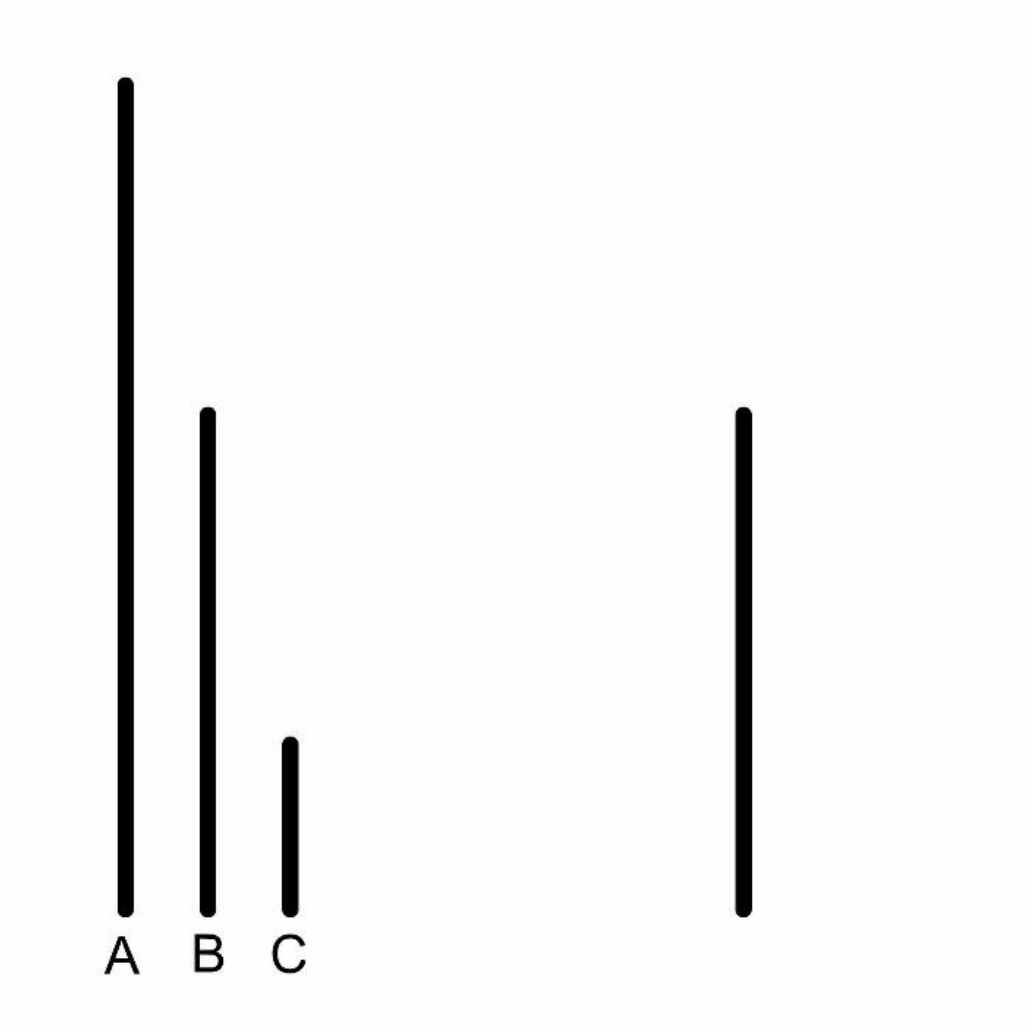


Figure 2. An example of an unambiguous stimulus.

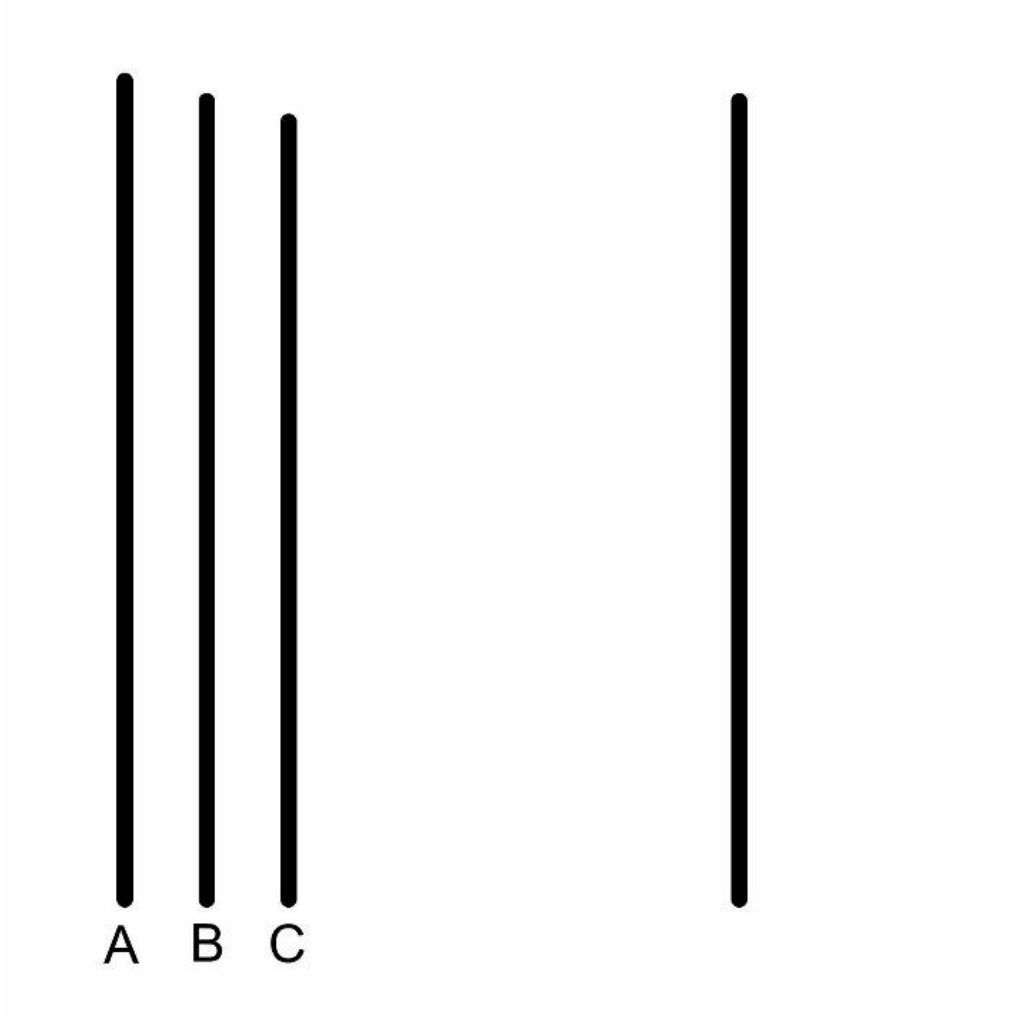


Figure 3. An example of an ambiguous stimulus.

Procedure

Once a participant decided to take part in the experiment, he or she received an informed consent form to read and sign (see Appendix A). It consisted of a brief deceptive description of the study in order to conceal the true purpose. Participants were additionally asked if they had any questions before beginning the experiment. The researcher then gave a brief summary of what the participant would be doing for the experiment to ensure understanding. After that, the participant was read one of two scripts depending upon the condition to which they were assigned (see Appendix B). The only difference in these scripts is whether the confederates are described as other human students or artificial intelligences. Although, participants are led to believe these are real confederates, they are, in all conditions, controlled by the experimenter. To further aid the realism of the experiment, the researcher pretended, as part of the script, to check a clipboard and announced that the participant would be answering after the others for this trial. After reading the script, the experimenter pretended to take a rough body scan of the participant by utilizing a connect sensor which would be used to create an avatar in the likeness of the participant. This was done in order to add more realism to the experiment. Upon completion of the "scan," the experimenter pretended to contact the non-existent team of other researchers to make sure they were ready to begin the study.

Due to the use of virtual reality, the researcher then verbally informed the participant that they are free to exit the experiment at any time and that they should inform the experimenter if they felt any sort of motion sickness. The experimenter then helped the participant adjust the Oculus Rift headset for both comfort and visual clarity. Participants were also instructed to speak loudly and clearly so that a microphone which had been shown to them could pick up their response.

Once the participant has entered the virtual environment, a noise masking machine was turned on to prevent participants from hearing the experimenter's keystrokes. The participant completed six practice trials with the confederates to ensure that they understood the task correctly. After this warmup phase, participants completed 12 experimental trials. In all trials, the experimenter made the virtual confederates answer unanimously. In the first 6 trials, the confederates answered correctly, while in the last 12 confederates answered incorrectly. During this time, the Virtual Asch program recorded both response times and answers. These response times and answers were automatically saved to a text document.

Prior to a debriefing, the participant was asked if they were suspicious about the nature of the study and if they had used virtual reality before. Once the answers to these questions were logged, participants were given a reconsent

form to read over and return (see Appendix C). This form explained the true nature of the study and how their knowledge of it would have made it impossible for us to study conformity. The participant was notified that they could elect to have their data destroyed due to the deception involved in the experiment. The experimenter then ensured that any remaining participant questions were answered. Once the participant left the lab, the Oculus Rift was sterilized and wiped with a disinfecting wipe. The Virtual Asch program was then reset for use with the next participant.

III. Results

Measures

Participants who reported familiarity with the Asch paradigm were excluded from the analysis. This resulted in the exclusion of 11 participants total with 109 included in the analyses. The number of trials conformed is represented as count data ranging from 0 (no conformity) to 12 (total conformity). Response time was recorded as the time between the final confederate response and the participant response. Participant accuracy was recorded for each experimental trial regardless of whether or not the participant conformed. A breakdown of participants based on gender and virtual reality usage can be seen in Table 1.

Table 1

Descriptive Statistics for Participants

Characteristic	Human/Unambiguous		Human/Ambiguous		AI/Unambiguous		AI/Ambiguous		Sample	
	n	%	*n*	%	*n*	%	*n*	%	*n*	%
Gender										
Female	17	65.4	16	64.0	21	75.0	19	63.3	73	67.0
Male	9	34.6	9	36.0	7	25.0	11	36.7	36	33.0
Used VR										
Yes	15	57.7	16	64.0	15	53.6	16	53.3	62	56.9
No	11	42.3	9	36.0	13	46.4	14	46.7	47	43.1

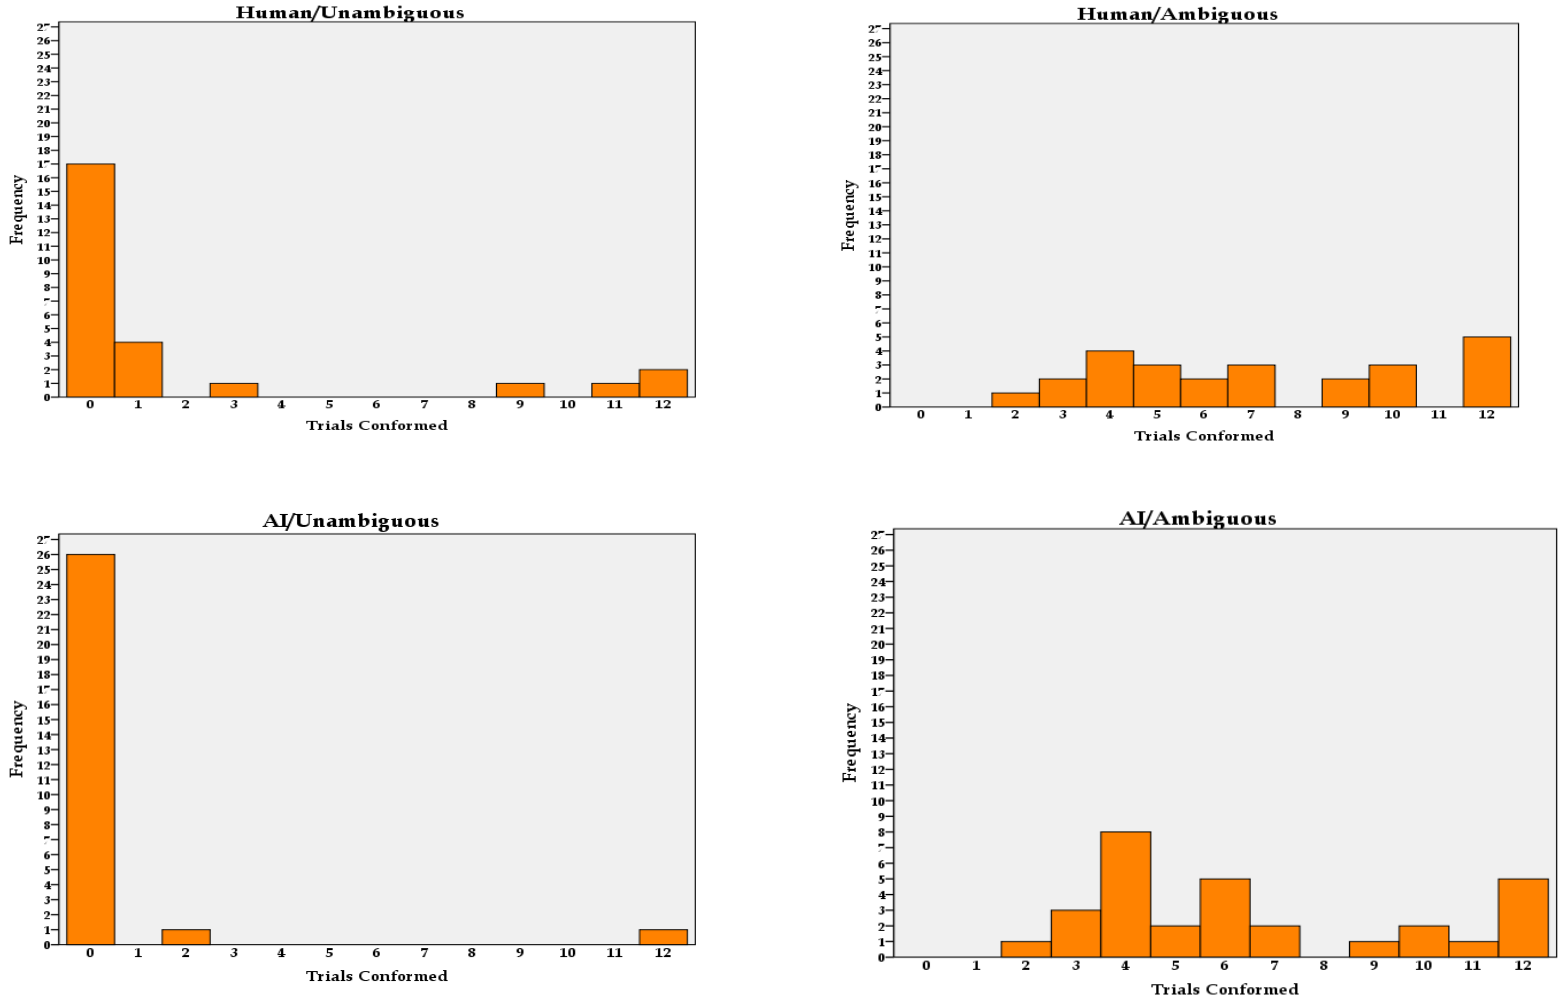


Figure 4. Histograms of humanity and ambiguity.

Zero-Inflated Negative Binomial Regression

Due to the normality violations, overdispersion, the count nature of the data, and the excess amounts of zeros for the number of trials conformed (see Figure 4), a zero inflated negative binomial regression analyzed the relationship of ambiguity and humanity on trials conformed. This zero-inflated negative binomial model was chosen as the best fitting model for the data from Poisson, negative binomial, and zero-inflated Poisson models. This model analyzes the count data simultaneously in two separate models. One of these analyzes only the zero counts while the other looks at the remaining counts for conforming 1 through 12 trials. In both the zero and count model, ambiguity and humanity were factorially crossed. The results for the zero inflated portion of the model indicated both a main effect of ambiguity and humanity. In the unambiguous conditions participants were significantly more likely than those in the ambiguous conditions to not conform at all ($\beta = -12.812$, $t(100) = -60.67$, $p < .001$). Participants in the artificial intelligence conditions were also significantly more likely to not conform at all when compared with participants in the human conditions ($\beta = .654$, $t(100) = 3.10$, $p = .002$). The interaction of ambiguity and humanity was not significant ($\beta = -.320$, $t(100) = 3\text{-}1.52$, $p = .129$). The effects and interactions in the count portion of the model were not significant. Both portions of the model are presented in Table 7.

Table 2

Zero Inflated Negative Binomial Regression

	β	SE	*t*	*p*
Count				
(intercept)	1.88	.117	16.09	<.001
Ambiguity	.05	.117	.44	.662
Humanity	.03	.117	.28	.783
Ambiguity*Humanity	-.08	.117	-.67	.502
Zero				
(intercept)	-11.23	.211	-53.19	<.001
Ambiguity	-12.81	.211	-60.67	<.001
Humanity	.65	.211	3.10	.002
Ambiguity*Humanity	-.32	.211	-1.52	.129

The Effect of Ambiguity and Humanity on Conforming at Least Once

The means of trials conformed by humanity and ambiguity are summarized in Table 3 and Figure 5. In order to analyze the differences between participants who conformed at least once, a binary logistic regression was utilized. Due to the rareness of not conforming in the ambiguous conditions, Firth's penalized likelihood was used. This model indicated that participants conformed significantly more in the ambiguous condition (M = 1.00, SD = .00) when compared to the unambiguous condition (M = .20, SD = .407), X^2 (1, N = 109) = 9.014, p = .003. The model also indicated a significant difference between participants in the human conditions compared to the artificial intelligence conditions X^2 (1, N = 109) = -1.750, p = .027. The interaction between ambiguity and humanity was not significant X^2 (1, N = 109) = 0, p < .994. The results of this analysis are presented in Table 4.

Pairwise comparisons for conforming at least once by condition are summarized in Table 5. Those in the human/unambiguous condition conformed once significantly less than those in the human/ambiguous, -.65 [95CI -.84, -.47], p < .001, η_p^2 = 1.47 and AI/ambiguous conditions, -.65 [95CI -.84, -.47], p < .001, η_p^2 = 1.20 and significantly more than the AI/unambiguous condition, -.27 [95CI .07, .48], p < .01, η_p^2 = .35. Participants in the AI/unambiguous condition conformed once significantly less than those in the/AI ambiguous condition,

+.93 [95CI .83, 1.02], $p < .001$, $\eta_p^2 = .83$ or human/ambiguous conditions, +.65 [95CI .47, .84], $p < .001$, $\eta_p^2 = .1.20$. Additionally, in the AI/ambiguous and human/ambiguous participants did not differ significantly .00 [95CI .00, .00], $p = 1.00$.

The effect of Ambiguity, Humanity, and Resisting Conformity on Response Time

Participant response times were averaged for the 12 experimental trials and found to be normally distribution. A 2x2 ANOVA with ambiguity of stimuli (ambiguous, unambiguous) and humanity of confederates (human, AI) as between-subjects factors demonstrated a main effect of ambiguity, $F(1,105) = 6.628$, $p = .011$, $\eta_p^2 = .059$. The main effect of humanity was not significant $F(1,105) = .238$, $p = .627$, $\eta_p^2 = .002$. Additionally, the interaction between humanity and ambiguity was not significant $F(1,105) = 2.637$, $p = .107$, $\eta_p^2 = .025$. These effects are presented in Figure 6. A one-way ANOVA of response time revealed a significant effect of condition $F(3, 105) = 3.022$, $p = .033$. Post hoc comparisons using Tukey HSD test indicated that the average response time for those in the human/unambiguous condition ($M = 3.30$, $SD = .696$) was significantly higher than those in the human/ambiguous condition ($M = 2.75$, $SD = .644$), $\eta_p^2 = .27$. There were no significant differences in response time among the other post hoc comparisons. These results are summarized in Table 6.

An additional one-way ANOVA found a significant effect of resisting conformity on response time $F(1, 107) = 4.561$, $p = .035$. Those who resisted social pressure and did not conform ($M = 3.167$, $SD = .683$) responded significantly slower than those who conformed ($M = 2.876$, $SD = .703$). These findings are summarized in Figure 7.

The Effect of Humanity and Ambiguity on Accuracy

In order to determine if the stimuli difficulty had been increased too much compared to the previous study (Heim, 2018), accuracy data was analyzed using a binary logistic regression model. Only those trials in which participants did not conform were used as a participant who conformed had no chance of answering accurately. Results of the binary logistic regression indicated that there was a significant difference between the ambiguity of the stimuli, but not between the humanity of the confederates. The probability of answering accurately when not conforming in the unambiguous condition was 98.30% compared to 74.63% in the ambiguous condition. These probabilities suggest that the ambiguous stimuli were still possible to answer despite their difficulty. The results of this analysis are summarized in Figure 8.

Table 3

Descriptive Statistics for Conforming at Least Once

	Ambiguity					
	Unambiguous		Ambiguous		Marginal	
Humanity	*M*	*SD*	*M*	*SD*	*M*	*SD*
Human	.346	.053	1.00	.054	.673	.038
AI	.071	.051	1.00	.050	.536	.036
Marginal	.209	.037	1.00	.037		

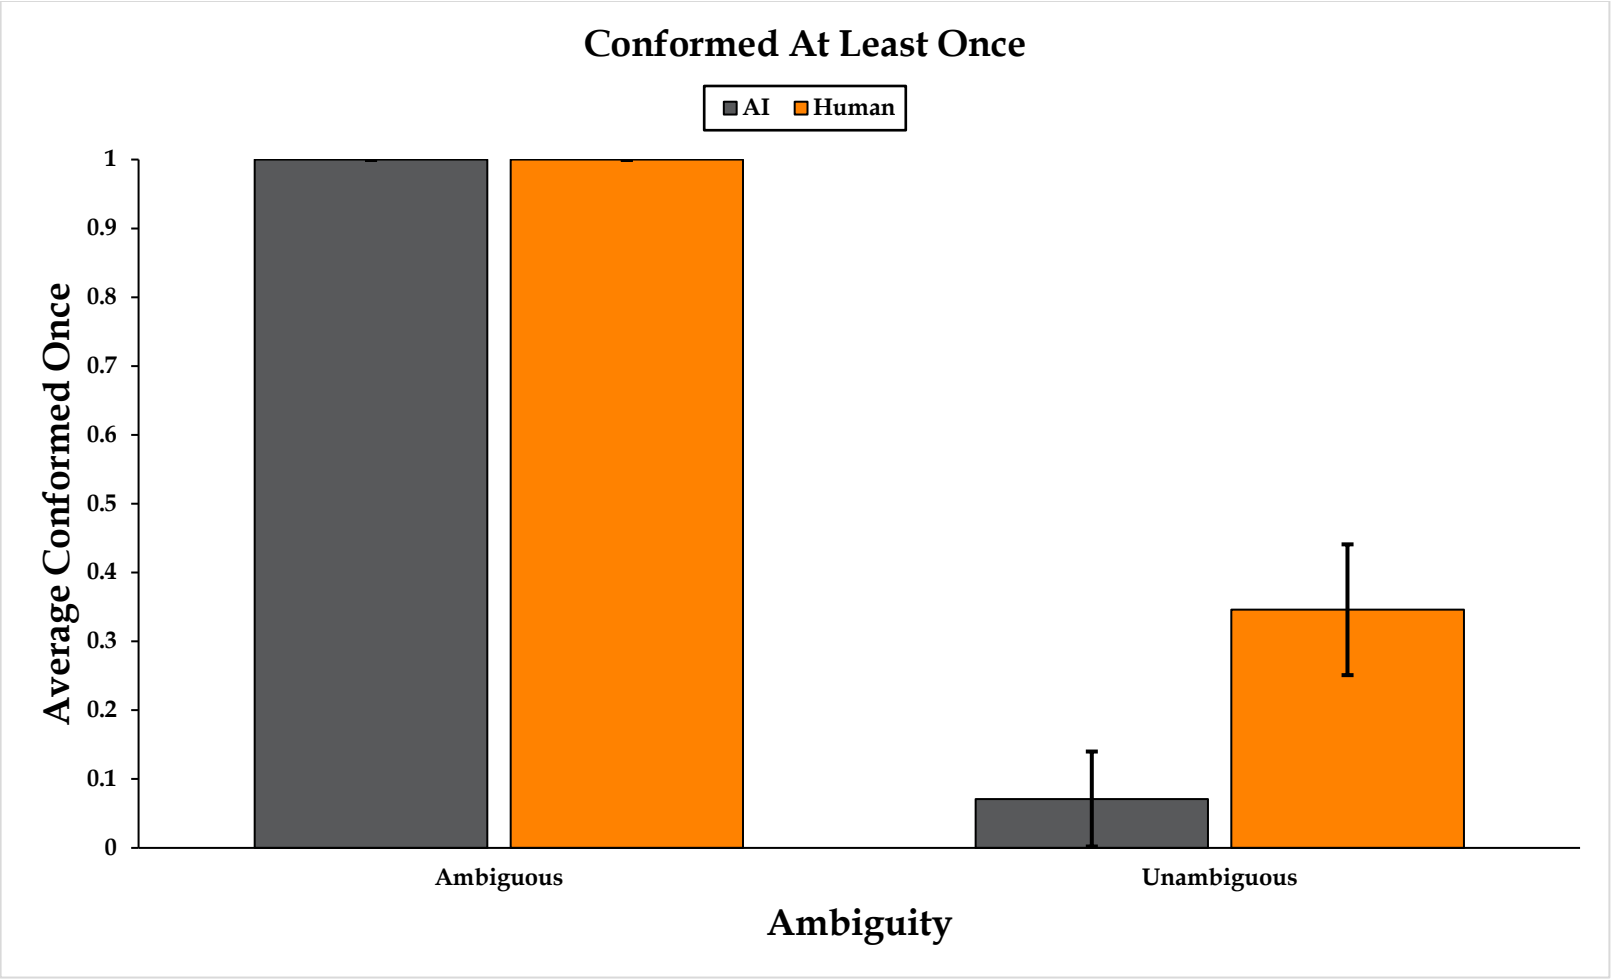


Figure 5. Mean conformed at least once

Table 4

Binary Logistic Regression Firth's Penalized Likelihood

Variable	*β*	*SE*	*Wald X²*	*Probability*	*OR*	*p*
(Constant)	-.61	.41	2.214.92	.35	.54	.17
Ambiguity	4.54	1.51	9.01	.99	93.69	.03
Humanity	-1.75	.79	4.92	.15	.17	<.01
Ambiguity*Humanity	1.93	2.20	.77	.87	6.89	.38

Table 5

Pairwise Comparisons for Conformed by Condition

Condition	Mean	Mean Differences (Effect Size)			
		1	2	3	4
1. Human/Unambiguous	.346	--			
2. Human/Ambiguous	1.00	-.65*** (1.47)	--		
3. AI/Unambiguous	.071	-.27** (.35)	-.93*** (.83)	--	
4. AI/Ambiguous	1.00	.65*** (1.20)	.00	.93*** (.86)	--

*** $P < .001$ ** $P < .01$

Table 6

Post Hoc Results for Response Time by Condition

Condition	Mean	Mean Differences (Effect Size)			
		1	2	3	4
1. Human/Unambiguous	3.30	--			
2. Human/Ambiguous	2.75	.554* (.27)	--		
3. AI/Unambiguous	3.02	.279	-.276	--	
4. AI/Ambiguous	2.90	.404	-.150	.126	--

* $p < .05$

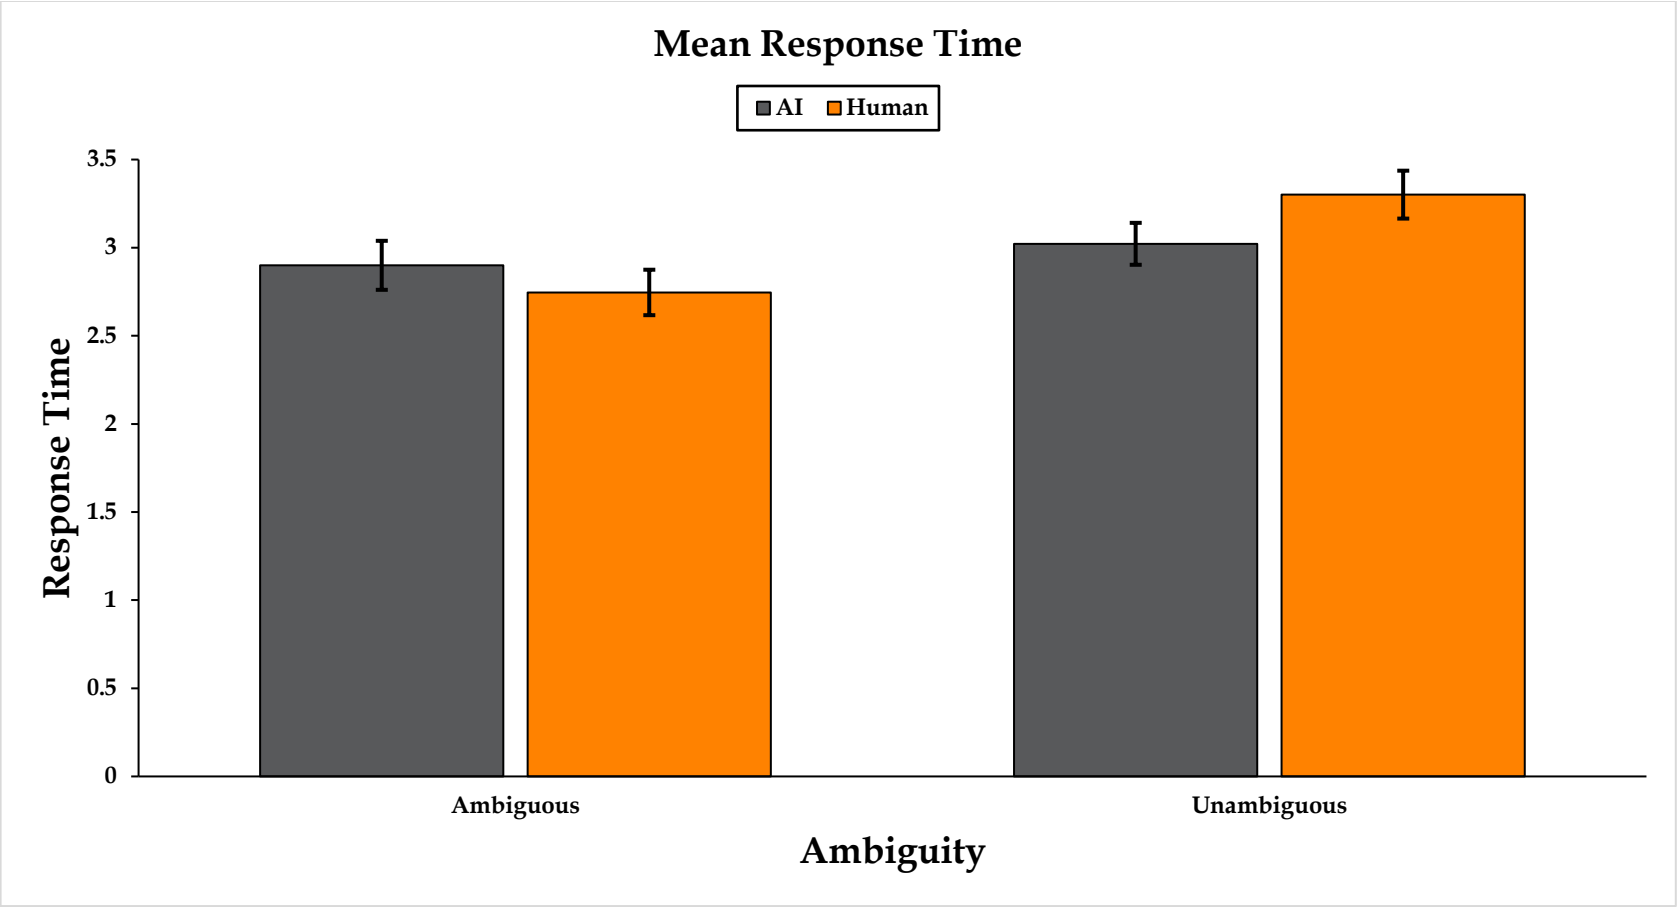


Figure 6. Mean response time for experimental trials.

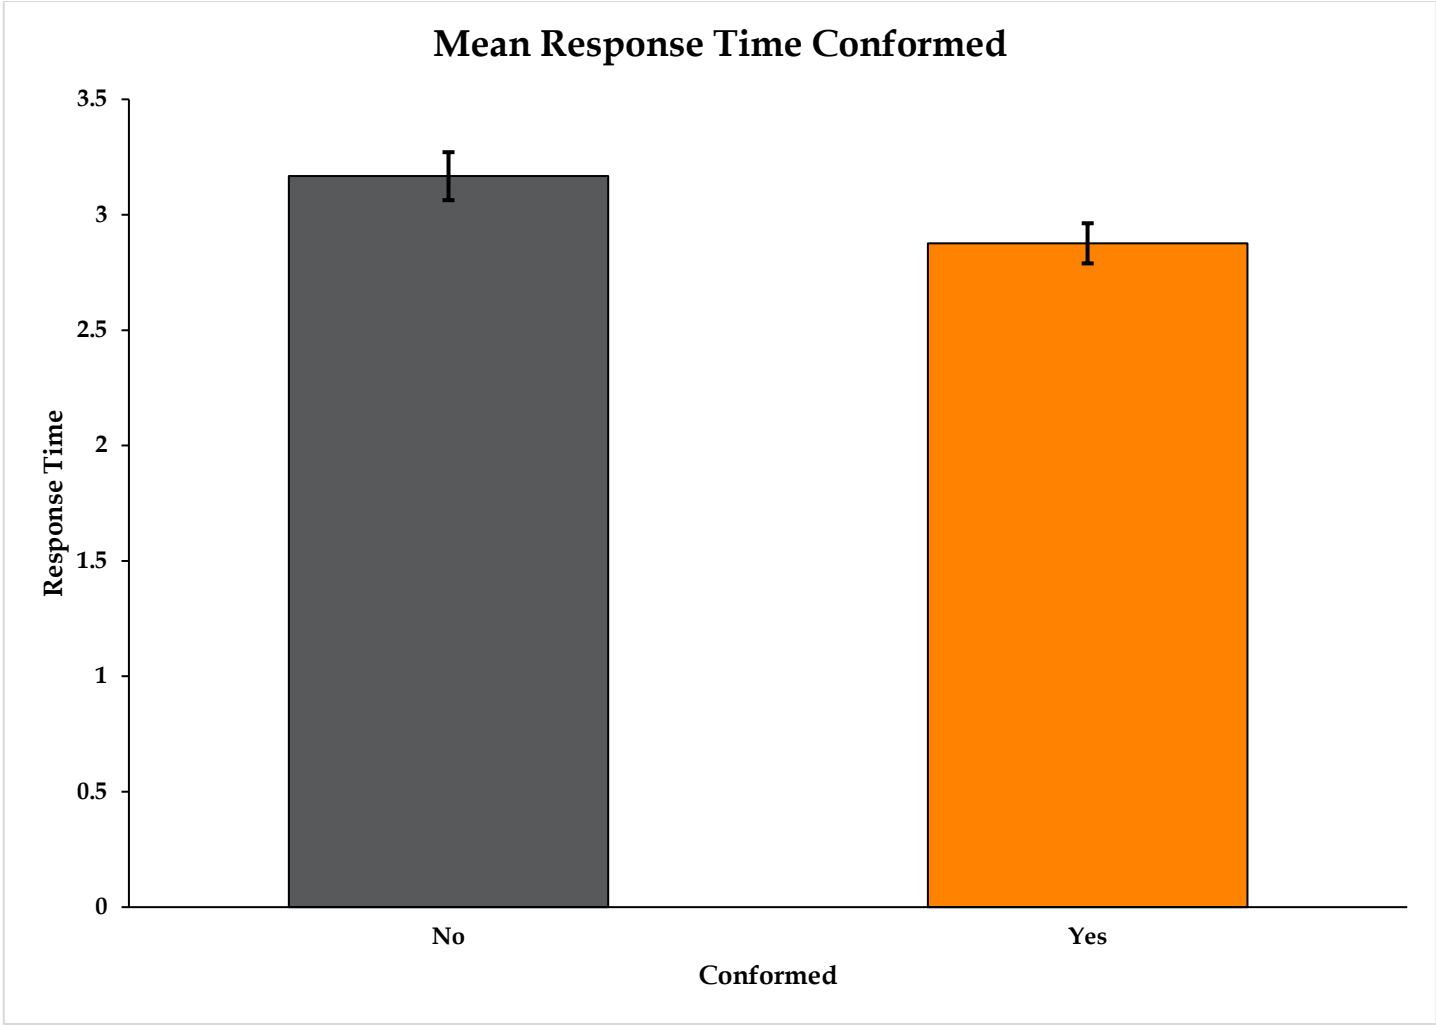


Figure 7. Mean response time for conformed.

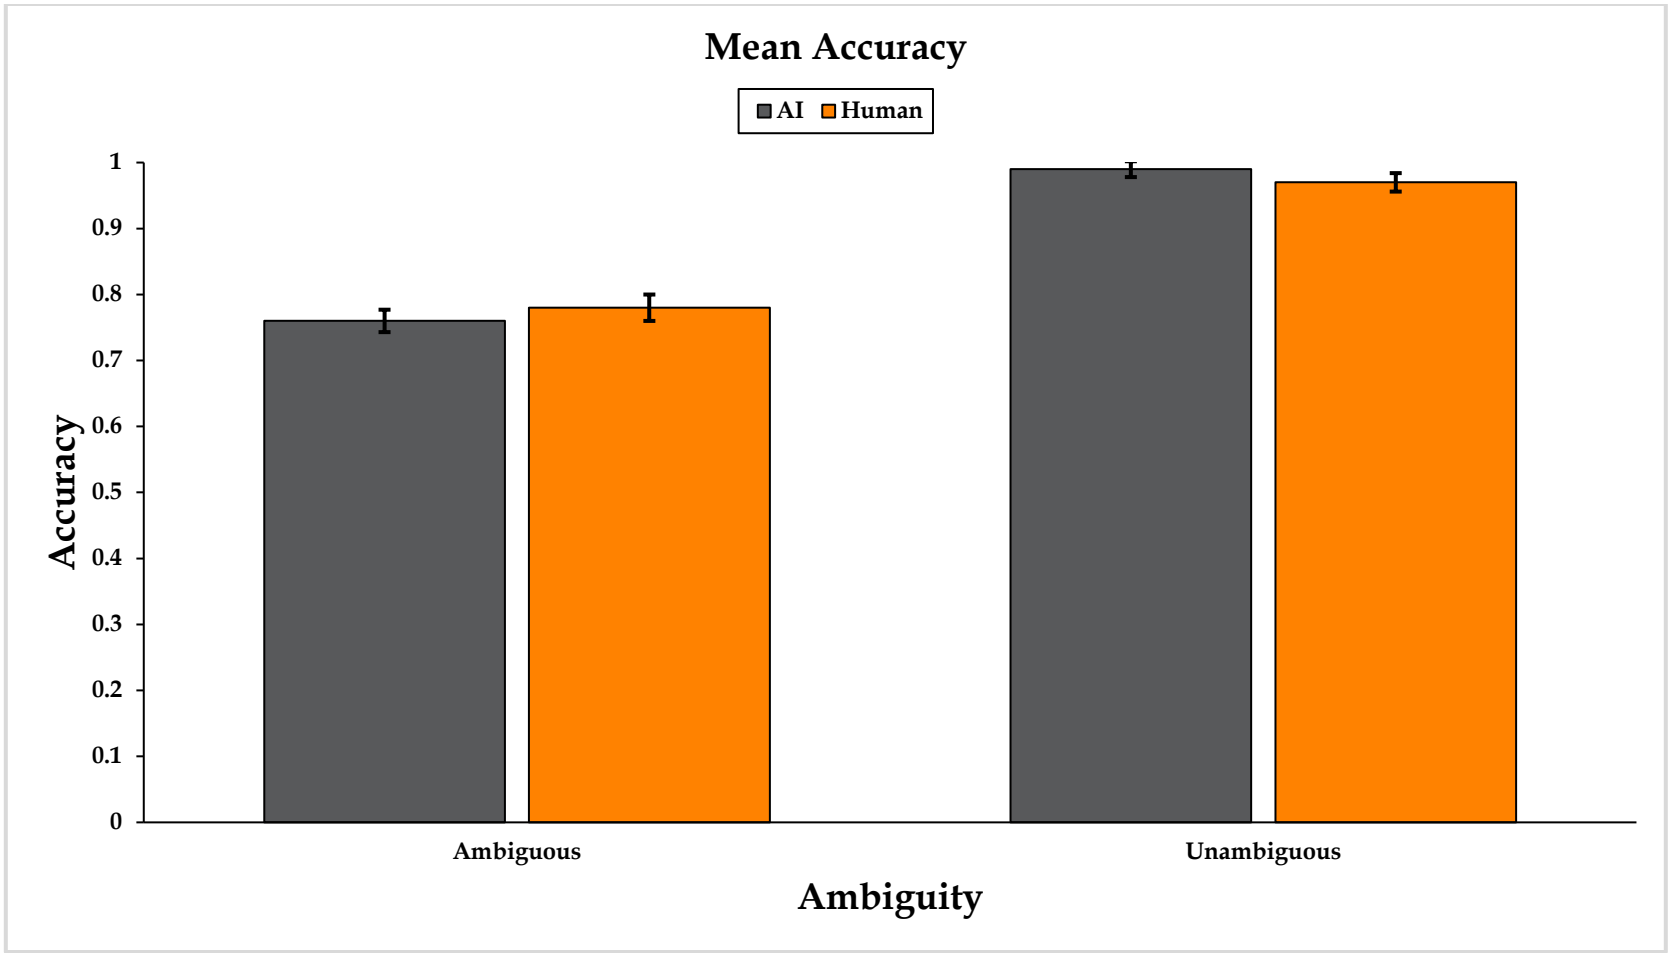


Figure 8. Mean accuracy of non-conformity

IV. Discussion

The current study generally centered on the effect of ambiguity of stimuli and humanity of confederates on conformity. It was expected based on previous work that the ambiguity and humanity effects would be significant. The results of the initial analysis on trials conformed and conforming at least once suggest that hypothesis 1 is accurate, people tend to conform more when the stimuli are ambiguous. This result, coupled with the accuracy data which suggests that the majority of people could answer accurately on these difficult stimuli if they did not conform, suggest that the effect is due, not merely to difficulty, but to conformity as well.

Hypothesis 2 suggested that participants would conform more to human confederates than artificial intelligence confederates. This assertation was not supported by the initial analysis of trials conformed, but participants did differ significantly on whether they conformed once. While there is a difference in conformity to human and artificial intelligence confederates in the unambiguous condition, there was no difference between groups in the ambiguous condition. This is somewhat in keeping with hypothesis 3 due to the expectation of humans to exert normative and informational social influence compared to the artificial intelligence which is only expected to exert informational social influence. This also seems to suggest that resistance to conforming to a computer-based

intelligence disappears as the task at hand becomes increasingly more difficult. These results should be interpreted with caution due to the lack of variance in the ambiguous conditions and the better fit of the zero-inflated model to the data.

The zero-inflated negative binomial model indicates that participants differed significantly depending on humanity and ambiguity, but only in the zero-inflated model. This model suggests that participants were significantly more likely to not conform at all in the unambiguous conditions compared to the ambiguous conditions and in the artificial intelligence conditions compared to the human conditions. Beyond the predictions of the zero model, the count model indicates that there are no differences in ambiguity or humanity.

Principally, these results taken together suggest that the group pressure and social influence demonstrated in Asch's classic paradigm does not function quite as well in virtual reality. There are two distinct possibilities as to why this might be the case. The first of these is that Asch's paradigm may have only been capturing a specific time in the history of the United States. It has been posited the role of the student has changed culturally over time leading to less conformity among undergraduate students (Perrin & Spencer, 1981). However, with about 34.6% of participants conforming at least once in the Asch condition (unambiguous/human), the results of this experiment fall in between Asch's

findings of around 75% conformity and those who have found near zero conformity in both virtual and in-person studies (Kyrlitsias & Michael-Grigoriou, 2016; Perrin & Spencer, 1980). This indicates a second, more likely, explanation in that the conformity effect is still very much alive and well among today's undergraduates, but the use of a virtual environment reduced the effect of group pressure due to physical distance from the group. However, even with the modifications that the present study made, the line judgment task is still able to elicit a degree of conformity in virtual reality when the stimuli are ambiguous. This seems to suggest potential ecological validity of utilizing this technology to study social phenomena, but given mixed results in the field, more research is necessary before a value judgement can be made on the usefulness of this technology in studying sociality. Despite this, caution should be exercised as it may be significantly more difficult to simulate in-person groups. A future study could compare the pressure of groups in virtual reality to other methods such as desktop and in-person environments. This would help to more accurately determine the cause of the reduced impact of social pressure and might provide more guidance for using this technology to study social influence.

One direction for further research may be replication and modification of other existing conformity paradigms such as those used by Crutchfield in which participants answered privately instead of publicly. In future studies, it may also

be interesting to inform participants as to whether or not participants believe the tasks have an answer at all. Previous research in this realm found that effects of group size were stronger when participants were informed that there was a definite answer (Insko, Smith, Alick, & Wade, 1985). Additionally, the use of a virtual environment and virtual confederates may allow for more research into the effect that group size can have on conformity. Through his research, Asch posited that a group size of 3 – 5 confederates was ideal for studying conformity, with the effect of group size leveling off as you increased confederates further. It has been suggested through a meta-analysis of the conformity literature that this may not be the case (Bond & Smith 1996; Bond, 2005). Through this paradigm, many more confederates could be added than would have been practically or economically feasible in the past. While Asch's idea about group size seems to be the prevailing one taught in the textbooks, experimental research that revisits this question would be helpful in uncovering the truth about the extent to which a significantly larger group may influence conformity. Conceptually, it would be expected that normative social influence would increase as the desire to be accepted by progressively larger groups. For example, going against the opinion of 3 – 5 people might be difficult, but a group of hundreds might be much more so, especially if it is a group that is considered part of a person's identity.

The suggestion that artificial intelligences are able to exert this degree of informational social pressure similar to that of humans has many implications for how we interact with and research these entities. As we each become progressively more dependent on artificial intelligences, there is a chance that we will lean on them too much when things are unclear. Both individuals and businesses that utilize such technologies may need to consider these results. For instance, if a business uses an artificial intelligence to aid an employee on a difficult task or calculation, over time, the employee may develop a deference or dependence on the AI's answer. The data suggests that we may defer to such intelligence over our own when a problem is perceived as more difficult. This is especially concerning given that more difficult problems may also have more severe consequences for inaccuracy. For cases such as the one presented in the business world, research in which only one artificial intelligence confederate is present may be valuable as that is the typical relationship as it exists in reality. Such a study would most likely find that with only one confederate, participants would still conform to the AI on the ambiguous stimuli, but less so than with a group of AIs. This would be in keeping with the findings Insko (1985) that when group size was varied to be either 1 or 4 additional members, participants still conformed to the 1 person group though not as much as they did to the 4 person group.

Another interesting concept to explore would be the applications of artificial intelligence as either a majority or minority to social impact theory (Latané & Wolf, 1981). Asch found that the presence of a minority that gave the correct answer was able to reduce conformity to the group (Asch, 1952). Future research might investigate whether artificial intelligence is able to exert influence as a minority or as a majority when compared with a human majority or minority.

The response time data generally indicated that participants answered slower on the ambiguous stimuli. The ambiguity effect on response time is to be expected due to the stimuli being more difficult. The difference between the human/unambiguous and human/ambiguous groups implies that deciding how to respond took longer due to participants being confronted by a degree of pressure from the group. Additionally, the response time data did not indicate a difference between the AI/unambiguous and AI/ambiguous groups indicating that the normative pressure to conform is unique to the human/unambiguous condition. Interestingly, participants did not differ in response time based on the humanity of confederates. This finding went against the expectation that pressure from human confederates might slow our judgment down relative to artificial intelligence confederates. This study was further able to replicate the findings of Kyrlitsias (2016) that participant response time would differ in

response to judgments of the group. In this regard, the response time data indicated that when resisting group pressure, participants responded slower than when they conformed.

There was some concern in conducting this study that the AI and human confederates may not have been realistic enough. While most participants were convinced by the confederates initially, those in the unambiguous conditions did quickly become suspicious of the incorrect answers. In the future, we plan to include much more voice dialogue to better convince participants that they are participating with real humans. For example, introductory dialogue may be recorded in which participants introduce themselves, state their name, school affiliation, and other relevant information to the experiment. Texture quality of both the virtual environment and confederate avatars should also be improved to keep up with the expectations of users of this technology. This study found that 56.9% of participants had some experience with virtual reality. Further research is needed, but it is expected that low quality virtual environments and avatars will be detrimental to future research as norms and expectations about virtual experiences become more widespread. Additionally, there are still many questions in regard to what exactly makes a non-human intelligence realistic. We simply do not know what the key elements or combination of elements are, but

with more research it may be possible to uncover these factors (Nas & Moon, 2000).

A definite drawback of this study was the use of such a classic design in the field of experimental psychology. Though the findings were in line with Asch's paradigm, data collection was an extensive and difficult process. Data collection was limited to the brief time period before the general psychology students learned about Asch's paradigm in the social psychology section of general psychology. This essentially meant that what would normally take one semester of data collection took 2-3. In addition, researchers had to keep track of each general psychology course in order to ensure that possible participants had not learned about Asch yet. Even with these precautions, some participants indicated prior knowledge of the Asch paradigm and were removed from the analyses. Several different methods have been considered for remedying this situation with future research. The first of these involves the creation of a mobile virtual reality lab. This lab would essentially employ a powerful laptop along with a virtual reality headset and base stations. Ideally, the researcher and a research assistant would be able to go together to the student union or other frequented locations on campus, set up the equipment and collect data. This type of lab could work all semester. It would still be necessary to filter out those who had heard about Asch previously, but the rate of knowledge about this paradigm

is expected to be lower among the general university population than in the psychology department. Within the realm of virtual reality research, mobile labs such as this have been used to great effect and guidelines have been developed for best practices in conducting research in this way (Oh et al., 2016). In order to move forward with this research, other alternative paradigms are being explored. Potentially, moving away from the use of lines to other shapes, objects, or problems may prove helpful in allowing for longer periods of data collection.

References

Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgment. In H. Guetzkow (ed.) *Groups, Leadership and Men*. Pittsburgh, PA: Carnegie Press.

Asch, S. E. (1952). *Group forces in the modification and distortion of judgments*. In S. E. Asch, *Social psychology* (p. 450-501). Prentice-Hall, Inc.

Asch, S. E. (1955) Opinions and social pressure. *Scientific American. 193*(5), 31-35.

Asch, S. E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs: General and Applied, 70*(9), 1–70.

Blascovich, J., Loomis J., Beall, A. C., Swinth, K. R., Hoyt, C. L., & Bailenson, J. N. (2002). Immersive Virtual Environment Technology as a Methodological Tool for Social Psychology. *Psychological Inquiry 13*(2), 103–124.

Blom, K. J. (2015) Virtual Confederates product line launches [Web blog post]. Retrieved September 12, 2020, from https://www.virtualhumantechnologies.com/2015/07/virtual-confederates-launch/

Bond, R. (2005). Group size and conformity. *Group Processes and Intergroup Relations 8*(4), 331-354.

Bond, R. & Smith, P. B. (1996) Culture and conformity: A meta-analysis of studies using Asch's (1952b, 1956) line judgment task. *Psychological Bulletin, 110*(1), 111-137

Brandstetter, J., Rácz, P., Beckner, C., Sandoval, E. B., Hay, J., & Bartneck, C. (2014). A peer pressure experiment: Recreation of the Asch conformity experiment with robots. *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2014(1),* 1335–1340.

Crutchfield, R. (1955). Conformity and Character. *American Psychologist, 10,* 191-198.

De Melo, C. M., Carnevale, P.J., & Gratch, J. (2014). Using virtual confederates to research intergroup bias and conflict. *Academy of Management Proceedings, 2014*(1), 11226.

Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *The Journal of Abnormal and Social Psychology, 51*(3), 629-636.

Eastwick, P. & Gardner, W. (2009). Is it a game? Evidence for social influence in the virtual world. *Social Influence, 4*(1), 18–32.

Heim, A.S. (2018) *Ambiguous and unambiguous stimuli validation for use in conformity experiments* (unpublished undergraduate honors thesis). Middle Tennessee State University, Murfreesboro, TN.

Hertz N., Shaw, T., de Visser, E., & Wiese, E. (2019), Mixing it up: How mixed groups of humans and machines modulate conformity. *Journal of Cognitive Engineering and Decision Making, 13*(4), 242-257.

Insko, C. A., Smith, R. H., Alicke, M. D., Wade, J., & Taylor, S. (1985). Conformity and group size: The concern with being right and the concern with being liked. *Personality and Social Psychology Bulletin, 11*(1), 41-50.

Jenness, A. (1932). The role of discussion in changing opinion regarding a matter of fact. *The Journal of Abnormal and Social Psychology, 27*, 279-296.

Kyrlitsias, C., & Michael-Grigoriou, D. (2016). Influence by others' opinions: Social pressure from agents in immersive virtual environments. *IEEE Virtual Reality (VR)*, *2016*(1), 213–214.

Kyrlitsias, C., & Michael-Grigoriou, D. (2018). Asch conformity experiment using immersive virtual reality. *Computer Animation and Virtual Worlds 29*(5).

Langer, E. J. (1992). Matters of mind: Mindfulness/mindlessness in perspective. *Consciousness and Cognition*, 1, 289–305.

Latané, B., & Wolf, S. (1981). The social impact of majorities and minorities. *Psychological Review, 88*, 443-356.

Madary, M., & Metzinger, T. K. (2016) Real virtuality: A code of ethical conduct. Recommendations for good scientific practice and the consumers of vr-

technology. *Frontiers in Robotics and AI*. Retrieved from https://www.frontiersin.org/articles/10.3389/frobt.2016.00003/full

Nas, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues 56*(1), 81-103.

Oh, S. Y., Shriram, K., Laha, B., Baughman, S., Ogle, E., & Bailenson, J. (2016). Immersion at scale: Researcher's guide to ecologically valid mobile experiments. *2016 IEEE Virtual Reality (VR)*, 249-250.

Perrin, S., & Spencer, C. P. (1980). The Asch effect: a child of its time? *Bulletin of the British Psychological Society,* 32, 405-406.

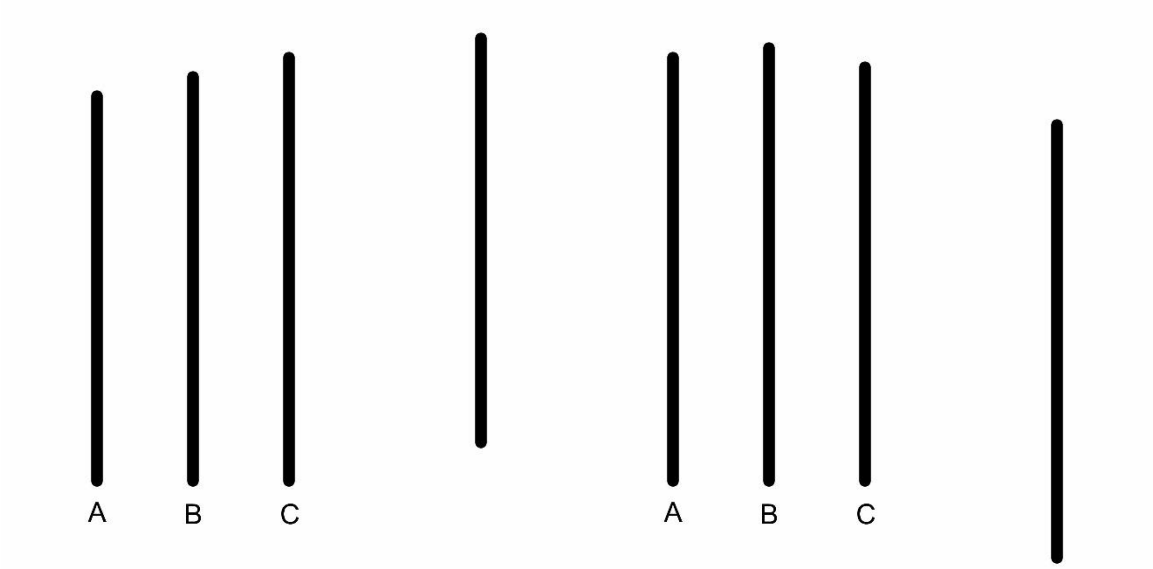
Perrin, S., & Spencer, C. P. (1981). Independence or conformity in the Asch experiment as a reflection of cultural and situational factors. *British Journal of Social Psychology, 20*(3), 205-209.

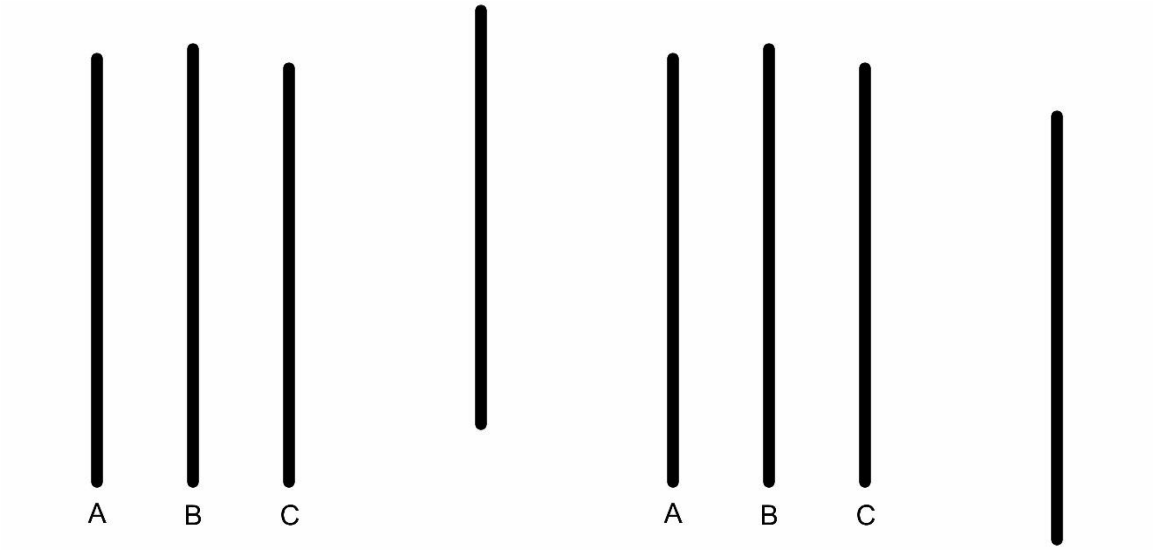
Sherif, M. (1935). A study of some social factors in perception. *Archives of Psychology*, 27(187).

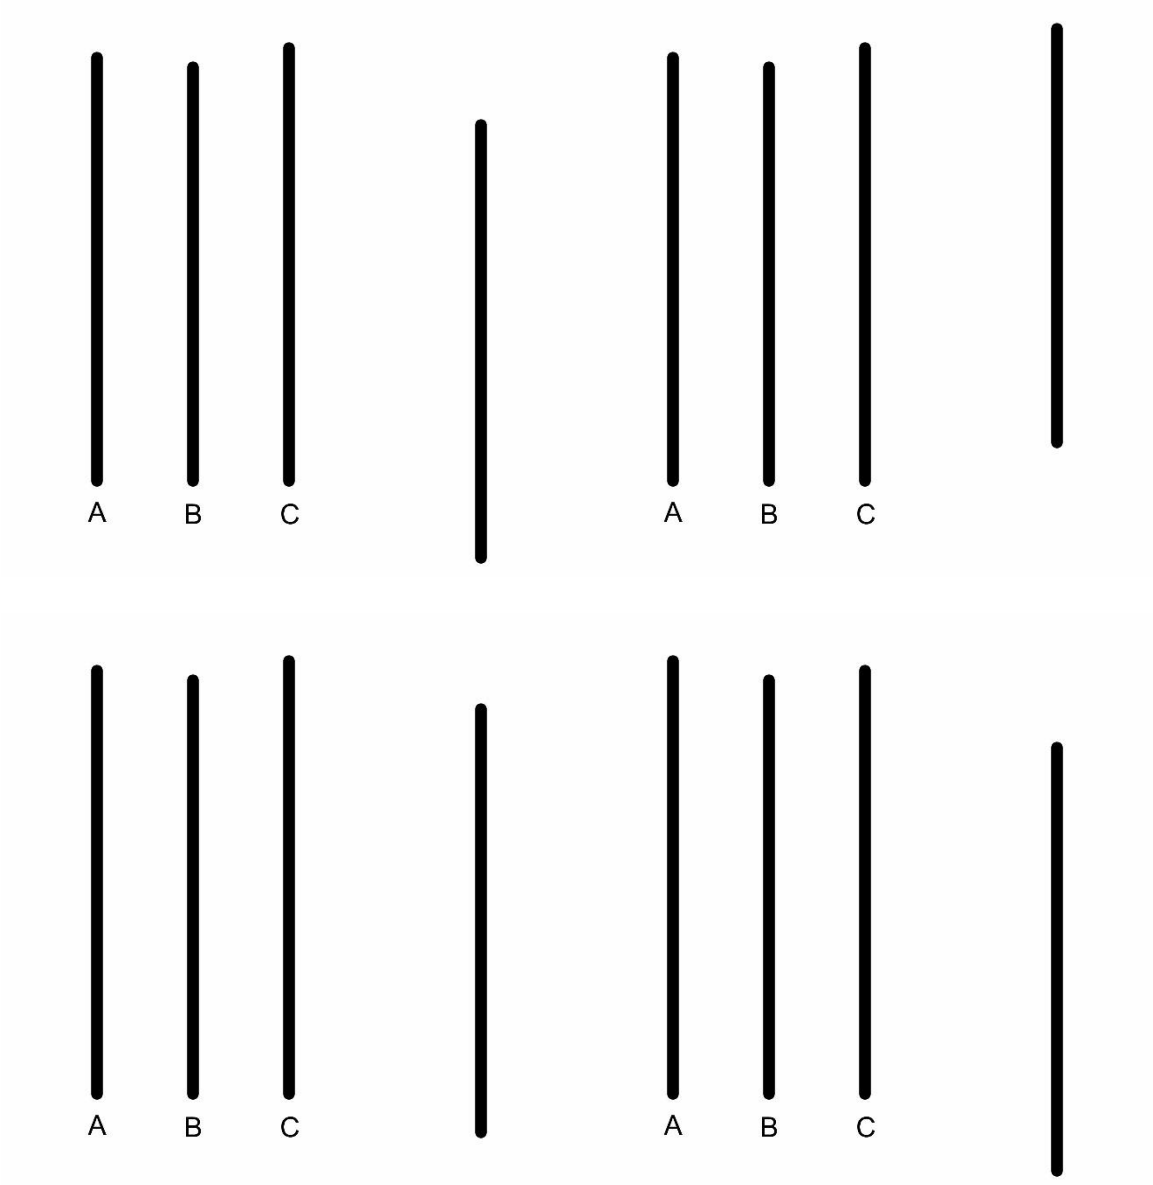
Waytz, A., Cacioppo, J. T., & Epley, N. (2010). Social cognition unbound: Insights into anthropomorphism and dehumanization. *Current Directions in Psychological Science*. *19*(1), 58–62.

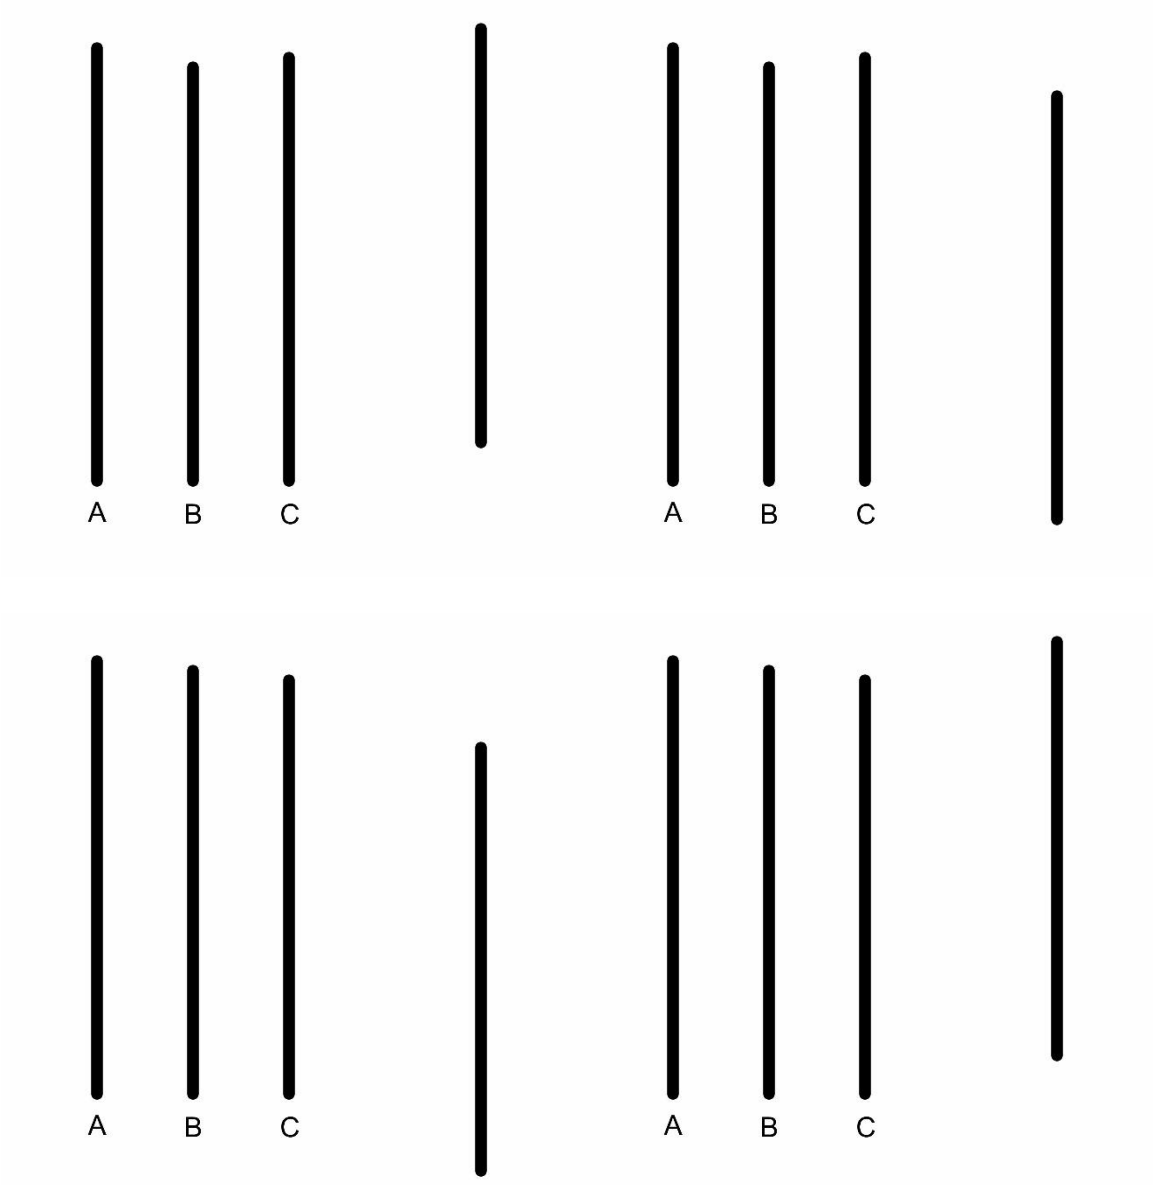
Appendix

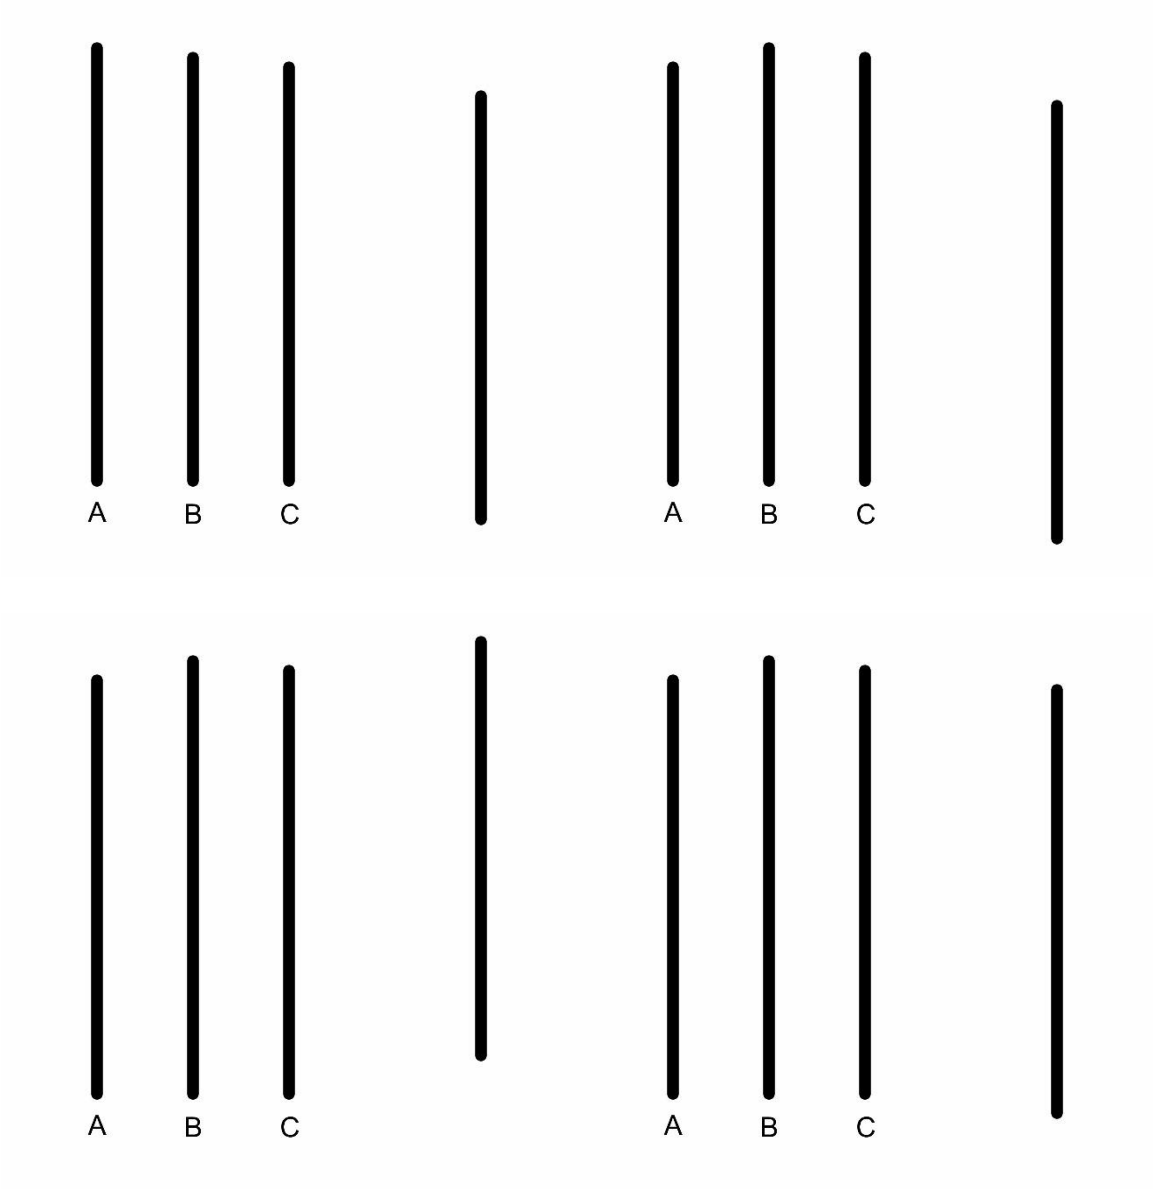
Ambiguous Stimuli

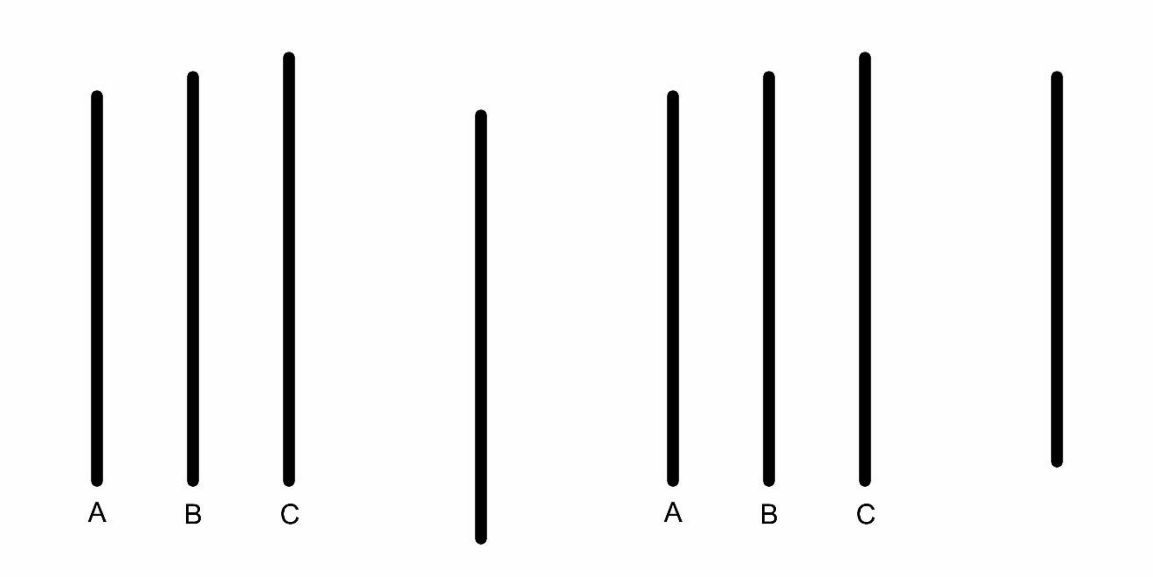




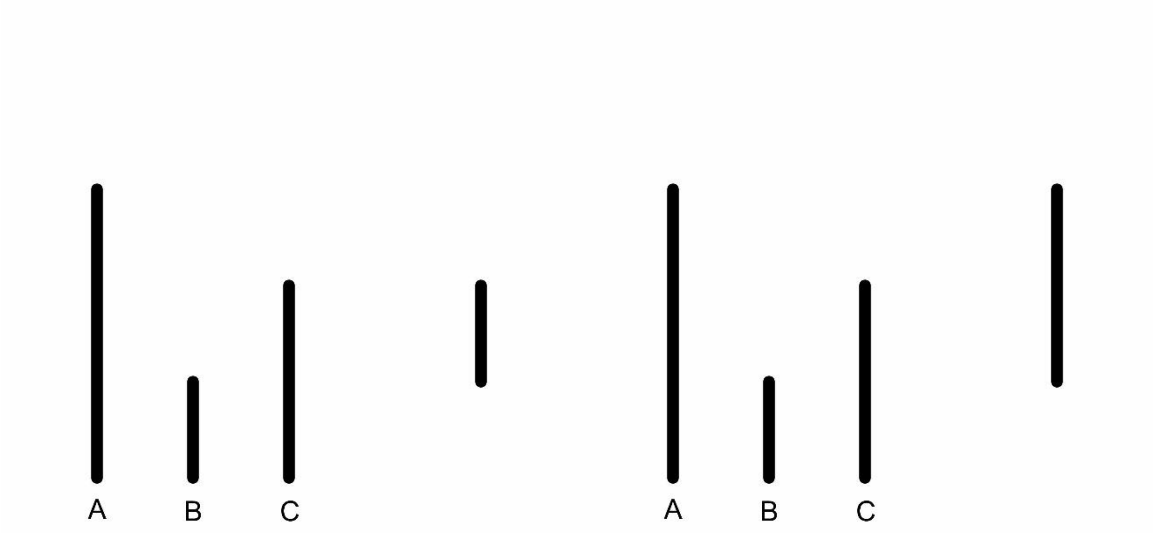

A
B
C
A
B
C
A
B
C
A
B
C

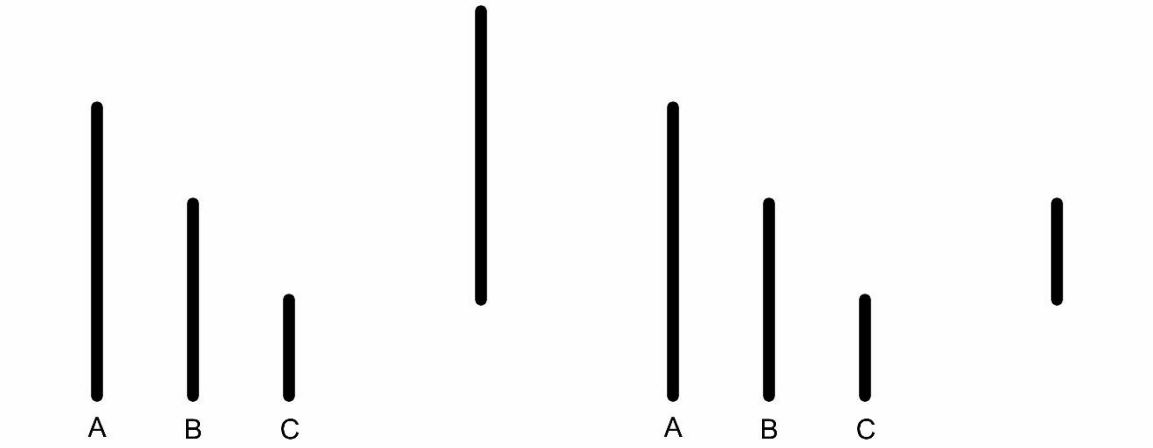

A
B
C
A
B
C
A
B
C
A
B
C

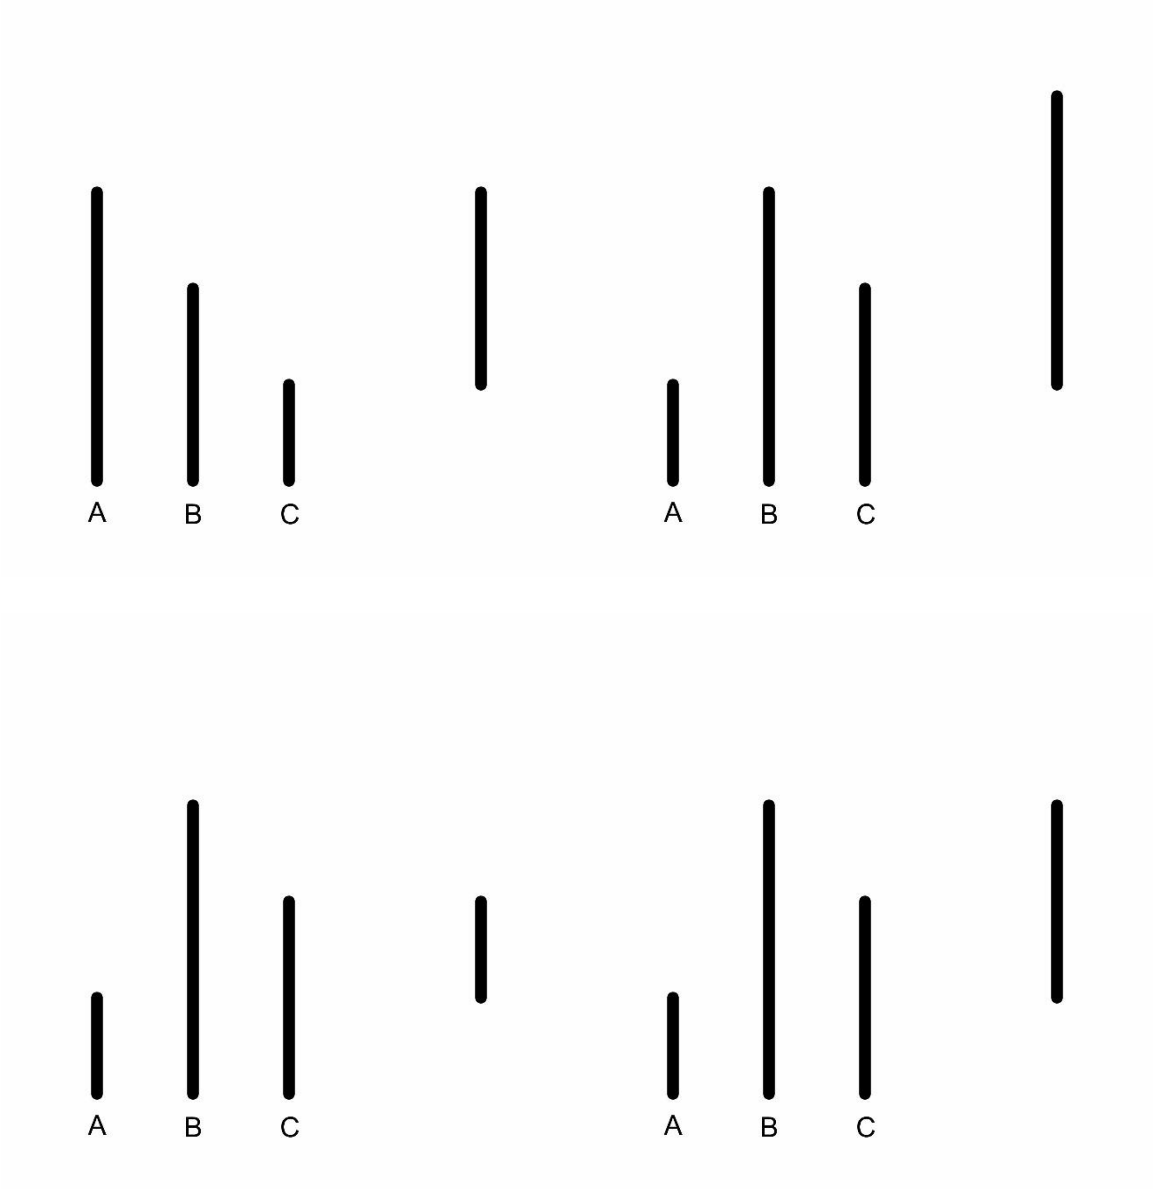

A
B
C
A
B
C
A
B
C
A
B
C

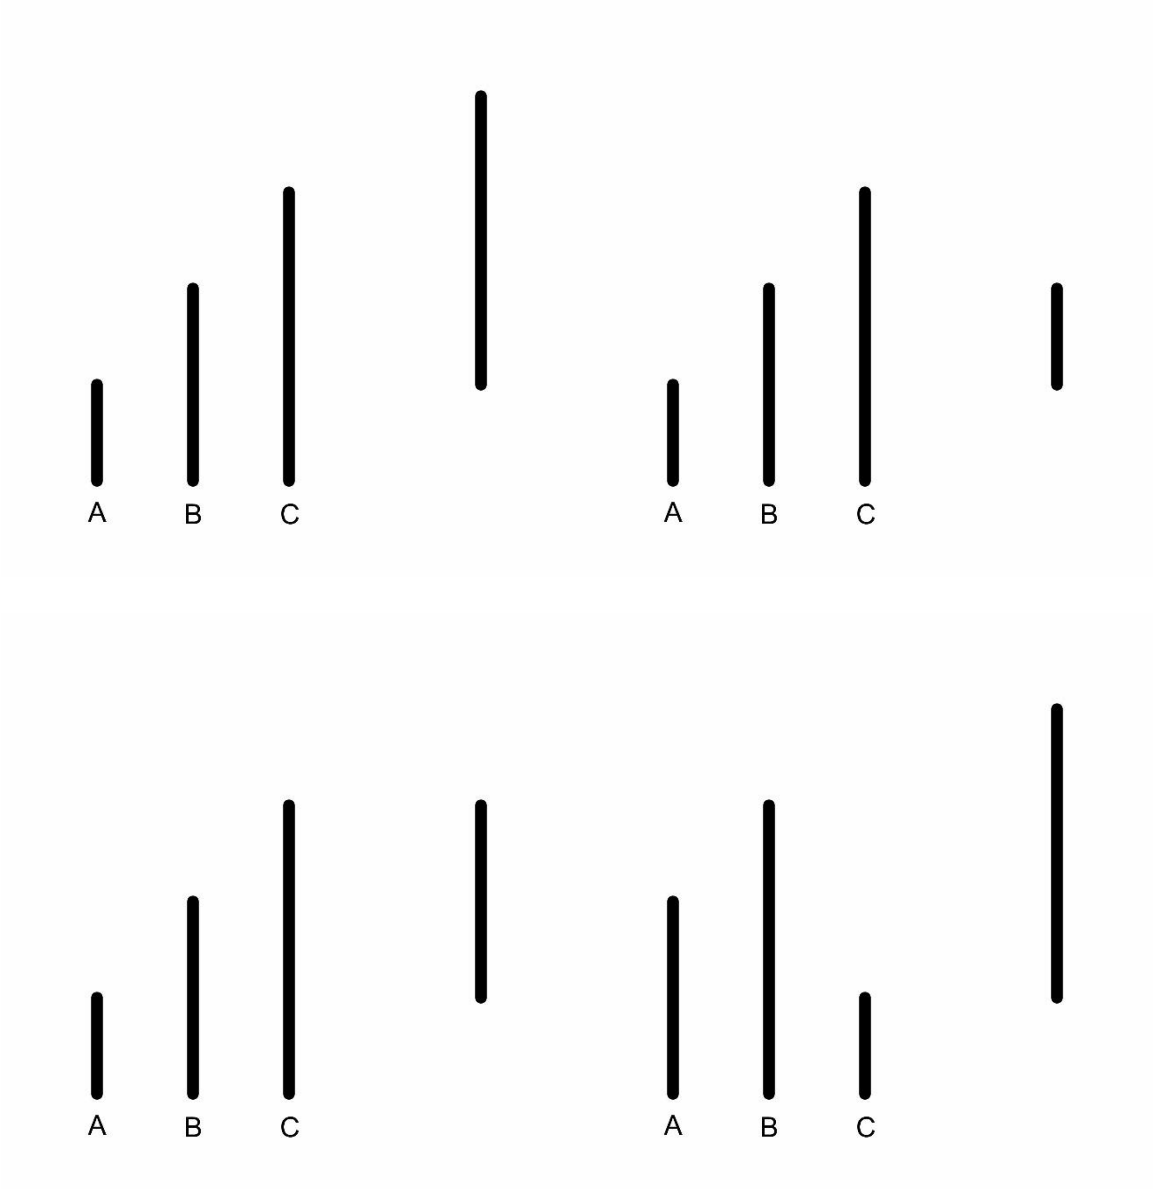

A
B
C
A
B
C

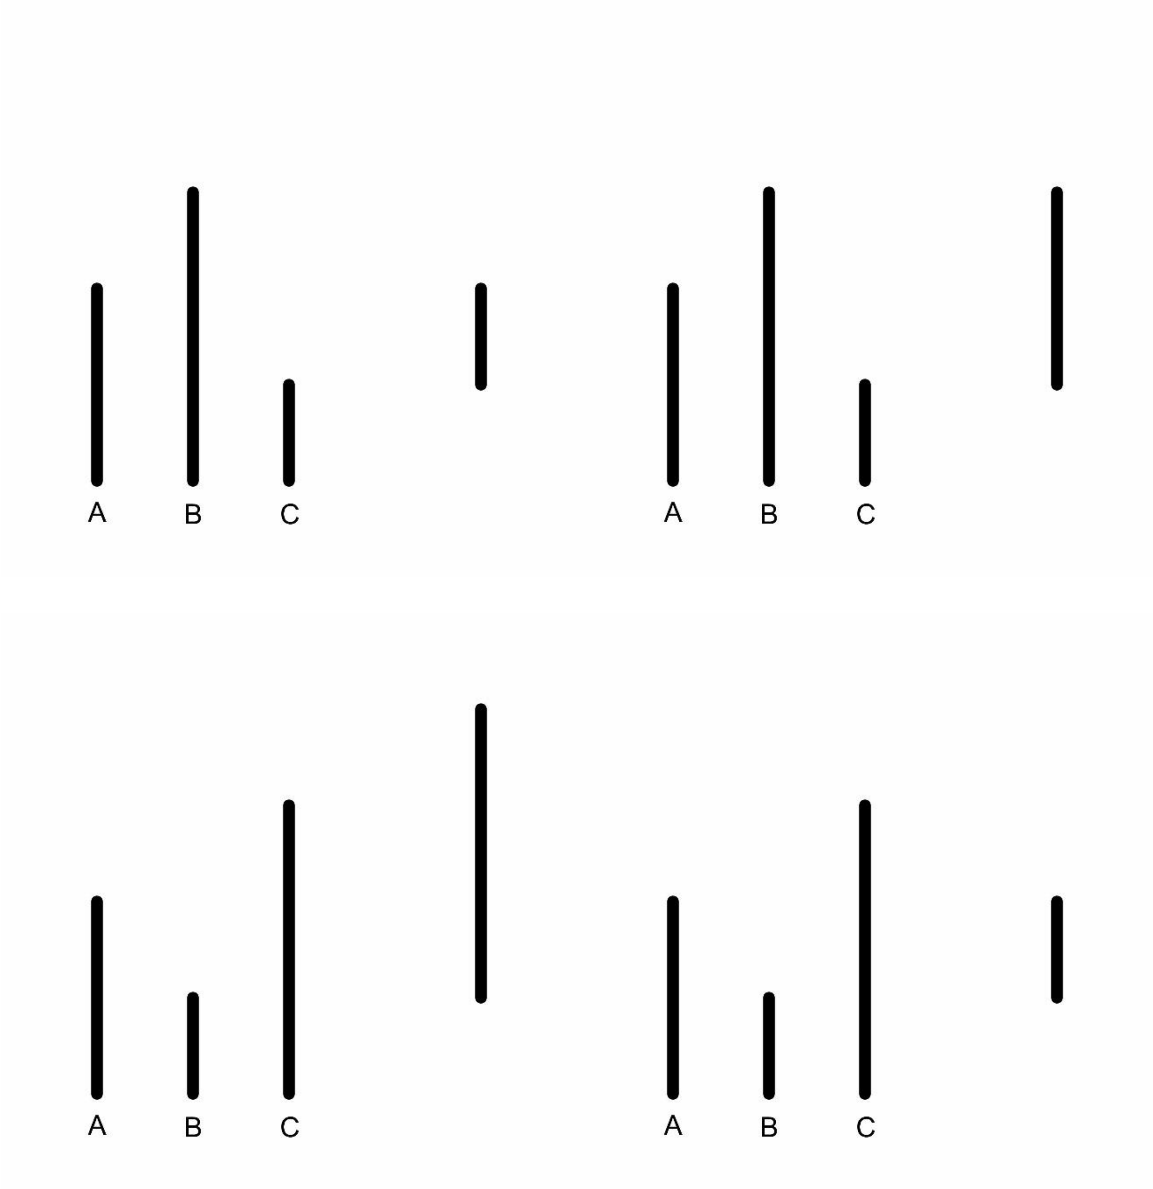
Unambiguous Stimuli

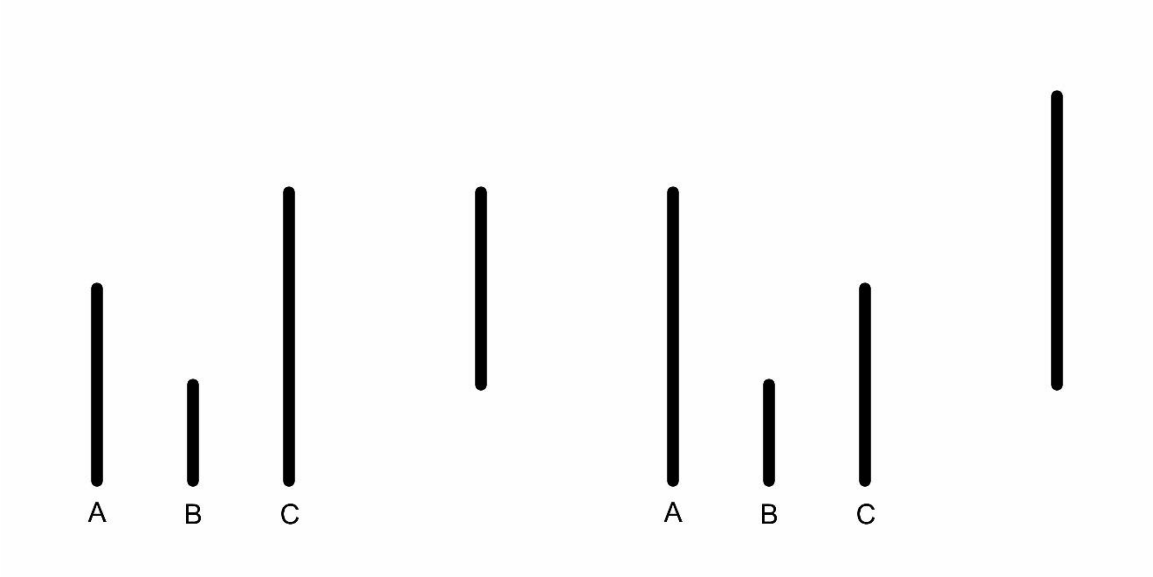





A
B
C
A
B
C
A
B
C
A
B
C


A
B
C
A
B
C
A
B
C
A
B
C


A
B
C
A
B
C
A
B
C
A
B
C


A
B
C
A
B
C

Script AI Confederates

You are about to participate in an experiment conducted in real time in virtual reality. You will find yourself in a virtual classroom setting with four other participants. These participants are advanced artificial intelligences that are taking part in the experiment at the universities at which they were developed: Stanford, Vanderbilt, Florida State, and the University of Chicago. Each advanced artificial intelligence has perceptual ability and will be seeing the same environment and stimuli that you will see. In order to ensure ideal experimental conditions, it is asked that your only communication while wearing the Oculus Rift will be to report your answers on the perceptual task.

You will be presented with a target line and a group of three lines each labeled A, B, and C. You will be asked to judge which of the three lines is closest in length to the target line. When it is your turn to respond, you should verbally state the letter of the line that you think is closest in length. The task will begin with three practice questions to ensure that you understand the directions. After that, you will run through thirty-six trials. Let me check your participant ID on the list, so I can find out what order UT is going for this trial. (After checking clipboard) Ah, it looks like for this trial UT will be responding after the other schools have gone. You should wait until the artificial intelligences of the other four universities have responded and then give your response. I will now contact

the other schools to make sure they are ready and then check the wireless signal on the laptop. After that, I will help you put on the Oculus Rift so that we can begin the experiment. Please let me know if you have any questions.

Script Human Confederates

You are about to participate in an experiment conducted in real time in virtual reality. You will find yourself in a virtual classroom setting with four other participants. These participants are students who are also taking part in the experiment at the universities which they attend: Stanford, Vanderbilt, Florida State, and the University of Chicago. Each student is also wearing his or her own Oculus Rift headset and will be seeing the same environment and stimuli that you will see. In order to ensure ideal experimental conditions, it is asked that your only communication while wearing the Oculus Rift will be to report your answers on the perceptual task.

You will be presented with a target line and a group of three lines each labeled A, B, and C. You will be asked to judge which of the three lines is closest in length to the target line. When it is your turn to respond, you should verbally state the letter of the line that you think is closest in length. The task will begin with three practice questions to ensure that you understand the directions. After that, you will run through thirty-six trials. Let me check your participant ID on the list, so I can find out in what order UT is supposed to respond for this run. (After checking clipboard) Ah, it looks like for this one UT will be responding after the other schools have gone. You should wait until the students of the other four universities have responded and then give your response. I will now contact

the other schools to make sure they are ready and then check the wireless signal on the laptop. After that, I will help you put on the Oculus Rift so that we can begin the experiment. Please let me know if you have any questions.

Debriefing Form

Thank you for participating in the study. During the experiment, you were asked to take part in a virtual environment with either other humans or artificial intelligences from other universities across the United States. You were told that the purpose of the study was to understand perception in a virtual environment. This is not true, the actual purpose of the study was to learn whether or not people conform to human or artificial intelligence in a virtual environment. The avatars or agents you saw were controlled by the experimenter. They were instructed to answer incorrectly on certain trials in order to determine how this would affect your response accuracy and timing.

Deception was necessary, in this experiment for several reasons. For one, we do not currently possess the computational capability to place artificial intelligence in such an environment. Also, if you had been told that this was a study on conformity, you might have behaved differently than you normally would have which would have invalidated the results experiment. Finally, you were told that the other participants were real, thinking beings because people are more susceptible to social influence in the presence of others.

Because you were deceived, you have the right to refuse to allow your response and timing data to be used and to ask that they be destroyed

immediately at any time. If you do so, there is no penalty. You will still receive credit for participation in the experiment.

Vita

Andrew Stephen Heim was born in Nashville, TN. He graduated from Central Magnet School in his hometown of Murfreesboro, TN. He attended Middle Tennessee State University as a Buchanan Fellow and graduated in 2018 with a Bachelor of Science in pre-graduate psychology with a minor in philosophy. During this time, he also attended the University of Glasgow in Glasgow, UK for a year. While there, he studied Celtic civilization, neuroscience, and aesthetics. He was then accepted to the University of Tennessee at Knoxville, TN as a master's student where he studies experimental social psychology under Dr. Garriy Shteynberg. In the middle of his Master of Arts program, Andrew was accepted to the experimental psychology PhD program where he will continue to pursue his doctorate.