# Negotiating task allocation in a team with the help of a negotiator agent

## How the introduction of the agent affects the trust

**Kevin van der Werff**[1]

**Supervisor(s): Myrthe Tielman**[1]**, Carolina Centeio Jorge**[1]

[1]EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
June 25, 2023

Name of the student: Kevin van der Werff
Final project course: CSE3000 Research Project
Thesis committee: Myrthe Tielman, Carolina Centeio Jorge, Rafa Bidarra

## Abstract

With the prevalence of artificial intelligence recently, more attention is being drawn towards the collaboration between humans and agents. Across the many fields where such an agent can be employed, we are going to specifically examine the domain of negotiation. A critical part to ensure success in this human-agent collaboration is to establish trust between the user and the agent. When considering negotiations, different strategies can be employed by the negotiation agents, the specific negotiation agent we shall use in our study employs the so called "hardliner negotiation style", which can be thought of as taking an uncompromising stance. When introducing the agent, we will either show a neutral introduction which does not mention the choice of strategy or show the hardliner introduction which does cover the chosen strategy. We are going to investigate how the choice of introduction affects the trust of humans in the negotiation agent. Leading us to the research question: *What's the effect of an agent using a truthful explanation of their hardliner negotiation style versus a neutral explanation on how much humans trust and rely on the agent to negotiate for them?* A between-subjects study took place to answer the question which has indicated that there is no significant difference (U = 111, $p$ = 0.97 two-tailed, $\alpha$ = 0.05) between the two different introductions.


## 1 Introduction

A big part of a human life consists of negotiation. Negotiation happens in the subtlest of ways but the most obvious situations to think of are sales [Geiger, 2017], interviewing for a job [Marks and Harold, 2011] and allocating team work [Peterson and Thompson, 1997]. Over the course of time, we have employed various negotiation strategies, whether among humans or bots [Ganesan, 1993; Perdue and Summers, 1991]. Negotiation is a human skill that some excel at, while others might be worse at it [Bazerman et al., 2000]. At the core of it, negotiation is a back and forth interaction with the goal of improving the situations for all parties involved [Boothby et al., 2023]. But humans might miss certain allocations that could have achieved better satisfaction for one of, or all, the parties involved, this is only natural. To address this, we would like to explore the usage of negotiating assistant agents. To ensure a successful collaboration between the humans and the negotiating agent, it's important to establish trust [Stanton and Jensen, 2020]. There are several aspects that can influence the trust, such as transparency and task characteristics [Glikson and Woolley, 2020].

One of these aspects is the explainability, which touches upon the field of explainable AI (XAI), within this field there are several ways we can determine the explainability of our agent. One of the ways is through how the agent is introduced to the user.

The negotiation agent can employ different negotiating strategies. These strategies can be thought of as analogous to real life, where people might choose to not stray from their initial propositions, or where others might give in a bit more to try to reach a consensus.

Our participants will be assisted by a negotiation assistant agent named Pocket Negotiator [Jonker et al., 2017]. One of the purposes is to suggest bids to the participant during the back and forth bidding process. Pocket Negotiator can employ different strategies, which will impact the bids that the Pocket Negotiator will suggest. We are going to consider only one of the strategies the Pocket Negotiator can use, the so called "hardliner strategy". This strategy is analogue to not straying from your initial demands.

Along with the chosen strategy, the agent could be introduced in a way that makes it clear to the user which strategy has been chosen, or the introduction could not disclose anything about the chosen strategy. We want to explore the impact of these different introductions on the trust of the humans in the negotiation agent.

This leads us to the research question of the paper; *What's the effect of an agent using a truthful explanation of their hardliner negotiation style versus a neutral explanation on how much humans trust and rely on the agent to negotiate for them?* along with the following hypothesis; *If the introduction of the agent includes the strategy chosen, it will invoke more trust as opposed to a neutral one that does not include the chosen strategy, as the former would be more transparent*.

To this end, we will conduct a between subjects experiment that will make each participant go through a simulated negotiation process in which they shall be assisted by Pocket Negotiator. Before each participant starts on the simulated negotiation process, they shall be introduced to Pocket Negotiator, this introduction shall form the independent variable to study the effect on trust.

In the following section we provide more of the background, including a deeper look at the Pocket Negotiator tool. In section 3 we outline our methodology, we discuss the participants, materials and the procedure of the experiment. In section 4 we cover the results. After that we will have a discussion in section 5, as well as reflecting upon our research in section 6. We finish the paper in section 7 with the conclusion.

## 2 Background

### 2.1 Trust

When we talk about trust, we need to define what trust means in the context of our study, we have adopted the following definition: "Trust (and distrust) are defined as a sentiment resulting from knowledge, beliefs, emotions and other aspects of experience, generating positive or negative expectations concerning the reactions of a system and the interaction with it" by Hoffman et al., 2018.

### 2.2 Related work

Extensive research has been conducted covering the topic of trust and artificial intelligence. Commonly placing emphasis on specific contexts, for example healthcare [Asan et al., 2020; Nundy et al., 2019] or education [Qin et al., 2020; Khosravi et al., 2022]. They commonly touch upon the field

of XAI, with the shared belief that it influences the trust of the user [Ferrario and Loi, 2022]. Models that analyze the various dimensions of an AI, which can influence trust are often represented, such as transparency (how well can the user understand the underlying AI process?) or tangibility (awareness of the presence of the AI) [Glikson and Woolley, 2020]. A general theme that occurs is that the goal is not to maximise the trust the human has in the agent, the goal is to have an *appropriate* amount of trust the human has in the agent. After all, blindly trusting an agent could result in undesired outcomes where the agent might have made a mistake that the human could have discovered through validating the decision [Lee and See, 2004; Stanton and Jensen, 2020; Glikson and Woolley, 2020; Asan et al., 2020].

A small subset of the existing literature covers artificial intelligence in the context of negotiations. This literature leans into modelling the negotiation process as game theory and then investigates how artificial intelligence can contribute to negotiation [Ferreira et al., 2015; Schulze-Horn et al., 2020].

There appears to be limited literature on how specifically the introduction of a negotiation agent might influence the trust of the user, we wish to address that gap.

## 2.3 Pocket Negotiator

We shall explain how the Pocket Negotiator tool functions and how users interact with the tool to simulate a negotiation process. Negotiations in Pocket Negotiator are modelled to suit an one-on-one negotiation process for two parties.

### Domain

The first thing that would need to be set up is the domain, this is where the users define the context, as well as define the specific ”issues”, which represents anything that they will negotiate about. For example, if one wants to set up the negotiation for a new job, they could set up the issues to be ”salary” and/or ”hours per week”. As well as defining the specific outcomes that the users will reach agreements on, for example, the choices of ”24 hours per week”, ”32 hours per week” or ”40 hours per week” for the issue ”hours per week”. The employee might then indicate that they prefer 32 hours per week, but the employer might prefer 40 hours per week, we explain in the following section how these preferences can be indicated in Pocket Negotiator.

### Preferences

Before an user can start the bidding, they need to indicate their preferences so that their negotiation assistant knows what bids would fit the user. In Pocket Negotiator, for each issue the participant can choose their desired outcome, this is done through sliders where the more a value is to the left, the less the participant would be happy with the outcome, and the more it is to the right, the more the participant would be happy with the outcome, see figure 1. So for example, for the issue ”assemble pizza topping” the participant would not be happy if they have to do it, and still not very happy, albeit better, if both of them were to do it and very happy if their friend would do it.

They can also indicate how important each issue is to them. A higher percentage number means they consider that issue to be more important to them (and would prefer if they have their desirable outcome for those issues over others that they consider less important), see figure 2.

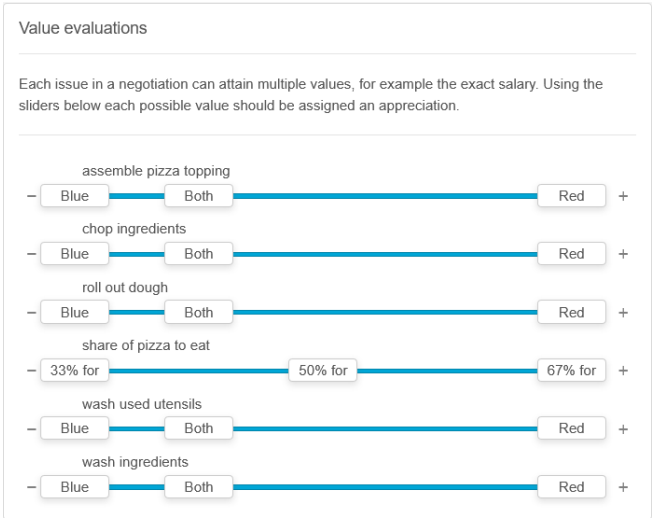


Figure 1: Value evaluations, this is where the user can indicate their preferred outcomes. Blue refers to the participant, red refers to the opposing party. The more an outcome is to the left, the less happy the user is with it. The more to the right, the more happy the user is.

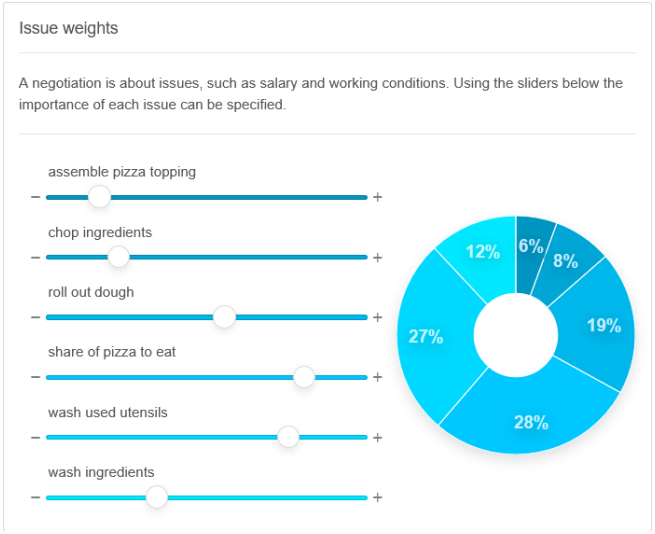


Figure 2: Issue weights, this is where the participant can indicate how important each issue is to them

### Bidding

Now that the preferences have been filled out, the bidding phase can start. This happens in a back-and-forth manner, an user can start with an initial bid, for example, they offer that they will assemble the pizza topping as well as chopping the ingredients, want the pizza to be divided equally, and the opponent can perform the remaining tasks. The opponent can then respond with a counter bid where they might agree on some issues, but want a change in other issues. The negotiation agent assists in this process by suggesting bids to the

user, the user is then able to use this bid, or override it with their own preferences. At any point, each user can accept the last offered bid from the other party or end the negotiations. It is also possible to give control to the negotiating agent which means that they take over the bidding phase from the user, the user does not need to fill in the bids manually anymore, and the bot shall reach an agreement, or not and the bidding will be terminated.

# 3 Methodology

The main goal of the experiment is to observe if two different introductions, which vary in regard to how much information they give out about the strategy of the negotiation assistant, influence the trust of the user in significant ways.

## 3.1 Participants

We had a total of 31 participants, of whom 25 identified as men and 6 as women. The majority originated from Europe, approximately half of the participants fell within the age range of 18-24 and the other half in the age range of 25-34. Most of them having completed either high school or having a bachelors degree. Almost all of them had a background in (software) engineering, whether that be through studying it, having studied it or being employed in a field related to (software) engineering. None of them had used the Pocket Negotiator tool before. The researchers recruited the participants through the researchers' network and fitting the criteria of having a background in (software) engineering. The results of one participant have been excluded from the sample set, as they chose to omit filling in the trust questionnaire.

## 3.2 Materials

The Pocket Negotiator v4 was used to simulate the negotiation process, Microsoft forms that's hosted on the Technological University of Delft's servers were used to conduct the questionnaires. Two questionnaires have been used, one background questionnaire which gathers the background of the participant, e.g. age group and region they grew up in, as well as one question to measure the propensity towards trust in AI of the participant: "I generally trust artificial intelligence (AI) to make accurate and reliable decisions" where participants could answer with a 5-point scale ranging from "strongly disagree" to "strongly agree". The second one being a trust questionnaire which was given after the negotiation process had been completed. This trust questionnaire originates from the paper by Hoffman et al., 2021. Since the questions essentially overlap with questions from the scale by Jian et al., 2000 scale or the Madsen-Gregor Scale [Madsen and Gregor, 2000] where both scales have shown high reliability coefficients we can safely assume content validity.

In this trust questionnaire, participants used a 5-point scale ranging from "strongly disagree" to "strongly agree" to indicate how much they trust various aspects of the negotiation assistant. An example is "I feel safe that when I rely on the negotiation assistant I will get the right answers." A total of eight questions were asked in this questionnaire, the full list can be found in appendix A.

Furthermore, for conducting the calls Discord (https://discord.com/) has been used, to present the introduction of the agent to the participant we have used Google Sheets. To generate figures as well as analyze the data Microsoft Excel and the website Statskingdom (https://www.statskingdom.com/) have been used.

Since the experiments were conducted with different examiners, we established various guidelines to prevent the experiments from deviating too much. For example, the settings to configure and the steps the participant would go through. A total of eight pilot studies have been conducted, after which we have tuned the configuration as well as adjusting some terminology in our explanations.

## 3.3 Procedure

A between-subjects experiment has been conducted. The independent variable is the introduction of the agent, and the dependent variable the trust the participant have in the negotiation assistant. Allocation of the two groups has been done through randomization. The experiments are primarily conducted through voice calls and screen sharing with Discord, an online communication platform. In some cases, the experiments were conducted in person where the examiner and participant met at a mutually agreed upon location. To prepare an experiment, an examiner would set up the negotiation tool on their computer. This means there was little preparation or setup required from the participant and it made sure that the setup was reproducible no matter the hardware of the participant. When the experiment was conducted through voice calls, the examiner would share the screen with the participant while also allowing them to control the users computer, including cursor control as well as keyboard input. An experiment took up to 45 minutes, with the shortest experiments being recorded at around 20 minutes. Each participant had been informed of the length, expectations and what type of data is collected, prior the experiment.

**Consent form and background questionnaire**

Upon starting the experiment, the first step involved providing a detailed explanation of the consent form, making it clear to the participant what is going to happen and having them explicitly give consent. After this, the participant fills in the background questionnaire.

**Tutorial**

Before we start the experiment, we first hold a tutorial to teach the participant how to use the Pocket Negotiator tool before we conduct the actual experiment. During the tutorial, the simulated negotiating situation is simplified to needing to heat up a frozen pizza with a friend (the opponent), there are only two issues the participant is concerned with: (1) decide who is going to put the frozen pizza in the microwave (the participant, their opponent, or both) and (2) the share of the pizza the participant will eat (a minority, an equal share or a majority of the pizza, the remaining portion goes to the opponent). During this time, we explain the user interface, the actions the participant might take and how to fill in their preferences. This is the time where the participant can ask questions about the tool and other related matters, without giving too much away about the tool itself, we will answer

the questions. We considered the tutorial completed once the participant felt they were comfortable with the tool and had no questions remaining.

**Introduction of the negotiation assistant**

At this point we are ready to move on to the negotiating scenario in which the participants shall be introduced to the bot. This is where our independent variable appears. The hardliner group has been shown the following **hardliner introduction**: "You shall be assisted by a negotiation agent in the following negotiation session. In the process of evaluating, the agent shall take your preferences in consideration and based on that it shall evaluate bids to and from the opposing party. This agent will not settle for anything less than the values, preferences and objectives indicated by the user. The goal is to hold a dominant position in the negotiation in order to obtain maximum benefits from the negotiation.". Whereas the neutral group has been shown the following **neutral introduction**: "You shall be assisted by a negotiation agent in the following negotiation session. In the process of evaluating, the agent shall take your preferences in consideration and based on that it shall evaluate bids to and from the opposing party. The goal of the agent is to maximise the overall result attainable for both parties." This introduction is shown in a Google Sheets document while being read aloud by the examiner.

**Simulated negotiation process**

Once the agent has been introduced, we explain the simulated negotiating process they are going to go through. Which is quite similar to the scenario the participants are going through in the tutorial, however, instead of preparing a pizza from the freezer, the participant and their opponent prepare a freshly made pizza. This means that there are more issues to negotiate about, such as washing ingredients, rolling out the dough and assembling the pizza toppings. Once the participant has (not) reached an agreement with their opponent or let the bot finish the negotiations, they fill in the trust questionnaire which measures their trust in the negotiation agent.

## 4 Results

All participants (aside from the one aforementioned participant who had been excluded) have completed the experiment in its entirety.

To measure the trust each participant has in the negotiating assistant, we can quantify the answers they gave to the questions of the trust questionnaire on a scale from 1 (strongly disagree) to 5 (strongly agree). With exception for one question; "I am wary of the negotiation assistant", for that one we inverse the scaling from 5 (strongly disagree) to 1 (strongly agree) since this question indicates a negative sense of trust, whereas the other questions indicate a positive sense of trust. Since the trust questionnaire is content validated, we can sum up these values to get a number indicating the trust they have in the negotiation assistant, with a higher number meaning more trust.

We have two groups, the hardliner group who received the hardliner introduction with sample size of 15 (M = 26.47, SD = 4.41), and the neutral group which had a sample size of 15 (M = 26.47, SD = 4.70). The results can be seen in the box plot in figure 3.

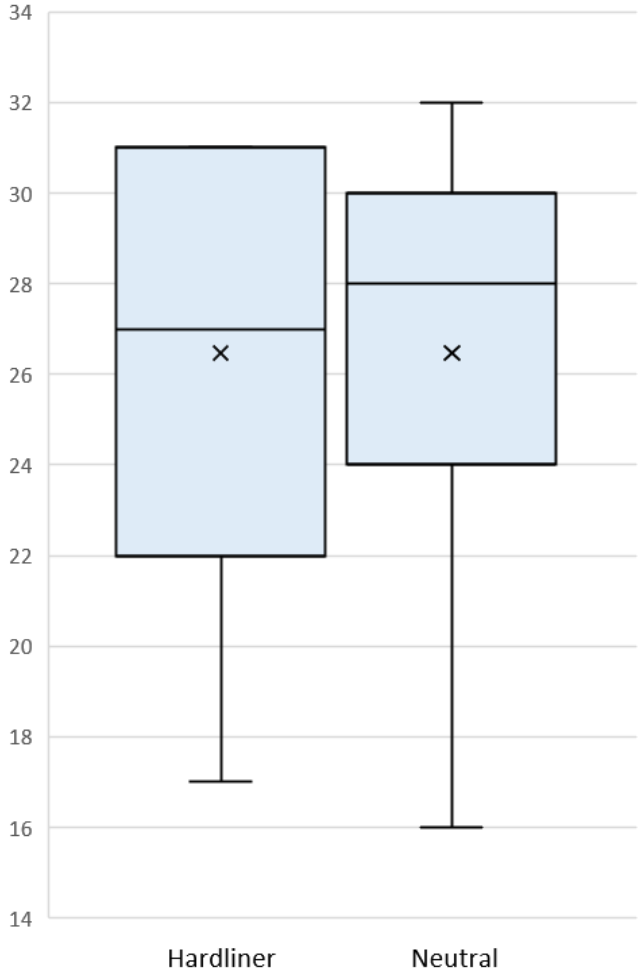


Figure 3: Box plot of the measured results for each sample group

Given that our sample size is small and the choice of our test depends on whether the samples are normally distributed, we use the Shapiro-Wilk test to determine whether our samples are normally distributed. Based on the outcome (hardliner group; W = 0.89, P-value = 0.065, neutral group; W = 0.88, P-value of 0.047), and additionally performing a visual confirmation of the histogram, we assume that the hardliner group is normally distributed, but the neutral group is not. Because of this, we proceed with the Mann-Whitney U Test.

The medians in the groups hardliner and neutral were 27 and 28, respectively. The results indicate that the difference between the groups are not significant (U = 111 , $p$ = 0.97 two-tailed, $\alpha$ = 0.05).

## 5 Discussion and future work

Contrary to our expectations, the results of the study do not affirm the hypothesis we had. It seems that in this case, the introduction of Pocket Negotiator did not make a significant

difference whether humans trust the negotiation assistant or not. A reason for this might be that since the introduction was given through an external document, rather than the negotiation agent presenting themselves with the introduction, the introduction might have had made less of an impact and rather be forgettable for the participants.

Alternatively, the introduction themselves could also take on many shapes and forms, perhaps introductions that do go more into detail and give away more of the strategy might impact the trust of the participant more, as they might notice earlier that the introduction does (not) deviate from the actual strategy of the bot. Inspiration could be taken from Miller, 2019 which goes in depth about explanations and draws inspiration from social sciences. Furthermore, perhaps instead of comparing two specific introductions, one might try to break down essential aspects of the introduction and run studies where these specific aspects vary, to find out if specific aspects of an introduction might impact the trust rather than two specific introductions.

A common question we got when the participants filled in the truth questionnaire was "what did the agent do again?". We speculate that this might because that when the Pocket Negotiator uses the hardliner strategy, the bids that the agent continuously suggests are nearly identical. Which might cause the participant to overlook that there's an underlying process that calculates the bids and instead think that the suggested bids are fixed.

## 6 Responsible Research

Since we have human participants in our experiments, it is a top priority that the participants who partake feel comfortable, know what exactly is expected of them and that they consent to the experiment. The Technology University of Delft has the Human Research Ethics (HREC) to set up guidelines to ensure this is the case. One of these is the usage of a consent form, we have adopted the consent form, and adjusted where needed, to use in our experiment. This consent form is a reference so that (1) as a researcher, we clearly know what we should go over with the participant prior the experiment and (2) for the participant to read, agree and give consent to all the aspects of the experiment.

For the sake of scientific integrity, ensuring reproducibility of the research findings is important. The setup of the experiment should be open for near, if not entirely, identical reproduction, such as the domain, introduction chosen and the questionnaires. However, there are variables that might differentiate between experiments, such as the examiner who conducts the experiments, as even with guidelines established, different examiners might expose different degrees of information to the participant which might influence the results.

## 7 Conclusion

To ensure success in the collaboration between humans and negotiation agents, we need to establish an appropriate amount of trust. To this goal we have explored how two different introductions, each varying in how much they expose the chosen strategy of the negotiation agent, to such a negotiation agent might influence the trust of the human in the negotiation agent. After having conducted an in-between experiment it was found that there was no significant difference between the amount of trust that the participants of the neutral group and the hardliner group had (U = 111, $p$ = 0.97 two-tailed, $\alpha$ = 0.05). We have speculated that while the difference might not be significant for these two explanations, we might still find a significant difference when other explanations, or other models to shape the explanations are used.

## Acknowledgments

This paper is the result of a group effort, for which we would like to thank our supervisors Carolina Jorge and Myrthe Tielman for their valuable guidance and feedback throughout the process. I would also like to personally thank my fellow students with whom I have conducted the study; Deniz Tan Hasdemir, Pavel Germanov, Justin Luu and Kostadin Penchev.

## A Trust Questionnaire

For each question, the participant could pick an answer on a 5-point Likert scale, "strongly disagree", "disagree," "neutral," "agree," and "strongly agree."

1. I am confident in the negotiation assistant. I feel that it works well.
2. The outputs of the negotiation assistant are very predictable.
3. The negotiation assistant is very reliable. I can count on it to be correct all the time.
4. I feel safe that when I rely on the negotiation assistant I will get the right answers.
5. The negotiation assistant is efficient in that it works very quickly.
6. I am wary of the negotiation assistant.
7. The negotiation assistant can perform the task better than a novice human user.
8. I like using the system for decision making.

## References

[Asan et al., 2020] Asan, O., Bayrak, A. E., and Choudhury, A. (2020). Artificial intelligence and human trust in healthcare: Focus on clinicians.

[Bazerman et al., 2000] Bazerman, M. H., Curhan, J. R., Moore, D. A., and Valley, K. L. (2000). Negotiation. *Annual Review of Psychology*, 51.

[Boothby et al., 2023] Boothby, E. J., Cooney, G., and Schweitzer, M. E. (2023). Embracing complexity: A review of negotiation research.

[Ferrario and Loi, 2022] Ferrario, A. and Loi, M. (2022). How explainability contributes to trust in ai.

[Ferreira et al., 2015] Ferreira, A., Cardoso, H. L., and Reis, L. P. (2015). Dipblue: A diplomacy agent with strategic and trust reasoning. volume 1.

[Ganesan, 1993] Ganesan, S. (1993). Negotiation strategies and the nature of channel relationships. *Journal of Marketing Research*, 30.

[Geiger, 2017] Geiger, I. (2017). A model of negotiation issue–based tactics in business-to-business sales negotiations. *Industrial Marketing Management*, 64.

[Glikson and Woolley, 2020] Glikson, E. and Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14.

[Hoffman et al., 2021] Hoffman, R., Mueller, S., Klein, G., and Litman, J. (2021). Measuring trust in the xai context. *Michigan Tech Publications*, PsyArXiv P.

[Hoffman et al., 2018] Hoffman, R. R., Mueller, S. T., Klein, G., and Litman, J. (2018). Metrics for explainable ai: Challenges and prospects institute for human and machine cognition [rhoffman@ihmc.us]. *arXiv:1812.04608*.

[Jian et al., 2000] Jian, J.-Y., Bisantz, A. M., and Drury, C. G. (2000). Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics*, 4.

[Jonker et al., 2017] Jonker, C. M., Aydoğan, R., Baarslag, T., Broekens, J., Detweiler, C. A., Hindriks, K. V., Huldtgren, A., and Pasman, W. (2017). An introduction to the pocket negotiator: A general purpose negotiation support system. volume 10207 LNAI.

[Khosravi et al., 2022] Khosravi, H., Shum, S. B., Chen, G., Conati, C., Tsai, Y. S., Kay, J., Knight, S., Martinez-Maldonado, R., Sadiq, S., and Gašević, D. (2022). Explainable artificial intelligence in education. *Computers and Education: Artificial Intelligence*, 3.

[Lee and See, 2004] Lee, J. D. and See, K. A. (2004). Trust in automation: Designing for appropriate reliance.

[Madsen and Gregor, 2000] Madsen, M. and Gregor, S. (2000). Measuring human-computer trust. *Proceedings of Eleventh Australasian Conference on Information Systems*.

[Marks and Harold, 2011] Marks, M. and Harold, C. (2011). Who asks and who receives in salary negotiation. *Journal of Organizational Behavior*, 32.

[Miller, 2019] Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences.

[Nundy et al., 2019] Nundy, S., Montgomery, T., and Wachter, R. M. (2019). Promoting trust between patients and physicians in the era of artificial intelligence.

[Perdue and Summers, 1991] Perdue, B. C. and Summers, J. O. (1991). Purchasing agents' use of negotiation strategies. *Journal of Marketing Research*, 28.

[Peterson and Thompson, 1997] Peterson, E. and Thompson, L. (1997). Negotiation teamwork: The impact of information distribution and accountability on performance depends on the relationship among team members. *Organizational Behavior and Human Decision Processes*, 72.

[Qin et al., 2020] Qin, F., Li, K., and Yan, J. (2020). Understanding user trust in artificial intelligence-based educational systems: Evidence from china. *British Journal of Educational Technology*, 51.

[Schulze-Horn et al., 2020] Schulze-Horn, I., Hueren, S., Scheffler, P., and Schiele, H. (2020). Artificial intelligence in purchasing: Facilitating mechanism design-based negotiations. *Applied Artificial Intelligence*, 34.

[Stanton and Jensen, 2020] Stanton, B. and Jensen, T. (2020). Trust and artificial intelligence. *National Institute of Standards and Technology*.