

Stockholm School of Economics
Department of Economics
5350 Thesis in Economics
Spring 2025

AI and Altruism

– An Experiment on Joint Decision-Making

Hanna Riese (42414)

Abstract: Humans increasingly interact with artificial intelligence (AI) in various contexts. Through a pre-registered experimental survey with university students (N=115), I investigate how charitable giving varies between people who make a joint decision with an imagined AI agent compared to those who make the decision with an imagined human. I also investigate whether people who decide with an imagined AI tend to revise their share donated to a different degree in an individual decision, compared to those who make the same decision with an imagined human. I do not find any statistically significant support for the hypothesis that people who make the decision with an imagined AI, on average, give a different amount to charity or change their share of endowment given compared to those who make the decision with an imagined human. The ability to draw conclusions from this finding is, however, limited by the aspect that the co-player is hypothetical and that I have a relatively small sample.

Keywords: Joint decision-making, altruism, charitable giving, artificial intelligence

JEL: C90, D64, D91

Supervisor: Anna Dreber Almenberg
Date submitted: 2025-08-25
Date examined: 2025-08-26
Discussant: John Gorham
Examiner: Jaakko Meriläinen

Acknowledgements

I want to extend my most sincere thank you to my supervisor, Anna Dreber Almenberg. Anna has enthusiastically and patiently provided me with invaluable guidance, encouragement, and support throughout this process.

I also want to express my gratitude towards my family and fellow classmates for their kind support and encouragement.

Table of Contents

1	INTRODUCTION	4
2	PREVIOUS RESEARCH AND THEORETICAL FRAMEWORK	6
2.1	<i>Joint decision-making and human-computer interactions</i>	6
2.2	<i>The dictator game and social preferences</i>	7
2.3	<i>Hypotheses</i>	8
3	THE EXPERIMENT	9
3.1	<i>A modified joint dictator game for charity</i>	9
3.2	<i>Experimental procedure</i>	10
3.3	<i>Empirical strategy</i>	16
4	RESULTS	17
4.1	<i>Primary analysis</i>	18
4.2	<i>Secondary analysis</i>	19
5	DISCUSSION	20
6	SUMMARY AND CONCLUSION.....	23
	REFERENCES.....	24
	APPENDIX A – PRE-REGISTRATION	29
	APPENDIX B – SURVEY	34
	APPENDIX C - E-MAIL.....	38

List of Tables and Figures

Table 1	Descriptive statistics.....	18
Table 2	Attention question.....	18
Table 3	T-tests and robustness tests primary analysis.....	19
Table 4	T-test secondary analysis	20

1. Introduction

Economists have traditionally assumed that decisions are made by individual economic agents (Cason & Mui, 1997) but have increasingly become interested in investigating group decision-making (Luhan et al., 2009). Parallel to this development, people more often make decisions together with non-human agents (Kirchkamp & Strobel, 2019). Artificial intelligence (AI) has spread in society and entered everyday life of many (Sheikh et al., 2023). In recent years, so called-generative AI and large language models, such as ChatGPT from OpenAI, have become influential in both professional and private settings (Gross, 2023). While the concept of AI can be traced to the 1950s, there is no single definition of what it constitutes (Sheikh et al., 2023). In a communication on AI for Europe in 2018, the European Commission used the following definition (p.1): “Artificial intelligence (AI) refers to systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals.” In the same communication, the European Commission proposes that AI can change the world in a way that can be paralleled to electricity or the steam engine (2018). While AI is expected to create considerable changes in our economy and society, its effect on various economic parameters will also depend on the behavior it elicits from humans who interacts with it (March, 2021).

Research speak towards that behaviors when humans interact with other humans do not necessarily readily translate into contexts where humans interact with computers or automated agents (Chugunova & Sele, 2022; March, 2021). Research has covered numerous aspects of interacting with various forms of automated agents (Chugunova & Sele, 2022) but the increased spread of AI in various contexts of society (Sheikh et al., 2023) motivates me to study further how the interaction with an AI can affect the outcome of a decision. Economists have studied social preferences and investigated actions that are unselfish (Drouvelis, 2021), such as the act of donating money to charity. This type of decision, addressed for example in the non-strategic dictator game, appears to be one decision-making context that is largely unexplored when it comes to addressing the effect of deciding together with a non-human agent.

The purpose of this thesis is to investigate how a joint decision with an AI agent affects the outcome in a situation where people can choose to be altruistic or act in a way that is beneficial to themselves, compared to when such a decision is made with another human. The thesis aims at shining light on if and to what extent an AI agent can affect altruistic decisions made by humans.

To address this aim, I design and pre-register an experimental survey based on the dictator game. I ask participants to allocate 100 Swedish crowns (SEK) between an option that is more beneficial to themselves (more “selfish”) and a charity, together with either an imagined co-player in the form of an AI agent or a human. Participants are told to imagine that the co-player has already allocated their part of the endowment (25 SEK) to the more selfish option and are then asked to allocate the remaining 75 SEK. The participants are then asked to allocate 100 SEK independently in an individual decision.

Within the context of the experiment conducted, the primary research question is: Do people who make a joint money allocation decision together with an imagined human choose to donate a different amount to a charity compared to people who make the same decision together with an imagined AI agent? In a secondary analysis, I also address whether there is a difference in regard to which extent people alter their choice of donation in an individual setting compared to when making a joint decision with an imagined AI agent or human.

Prior research has found that humans tend to act differently in strategic interactions with humans compared to interactions with computers (March, 2021). I argue that whether such interaction would render a more or less altruistic outcome in a non-strategic environment remains unclear. I hypothesize that there is a difference in the mean amount given to charity between participants who make a decision with an imagined human compared to those who make the same decision with an imagined AI agent. As an extension, I also hypothesize that participants who decide with an imagined AI on average change their share given to charity in an individual setting to a different extent than those who decide with an imagined human. These hypotheses are non-directional.

I do not find any statistically significant support for the hypothesis that people who imagine making a decision together with an AI on average donate a different amount to charity than those who imagine making the decision with a human. Neither do I find that the average change in share of the endowment donated between those who make the decision with an imagined AI statistically significantly differs from those who make the decision with an imagined human. While the null hypotheses cannot be rejected at the five percent significance level, I am careful to infer that this study has excluded the existence of an effect on altruistic choices from the type of co-player that participants have made their decision with. Among other things, the treatment might not have been considered sufficiently strong, given that participants are solely asked to imagine making the

decision with a co-player and do not interact in person or reality with somebody else. I also have a relatively small sample.

The remainder of this paper is structured as follows. Section 2 covers the previous research and theoretical framework that I draw on. Section 3 describes the experiment and the empirical strategy. Section 4 displays the results and section 5 constitutes the discussion. Finally, section 6 concludes.

2. Previous research and theoretical framework

2.1 Joint decision-making and human-computer interactions

Conventionally, economists have focused on individual decision-making theories and contexts, but a growing literature studies the effects on outcomes when deciding in teams instead of individually (Luhan et al, 2009). For example, in a meta-analysis of dictator games, Engel (2011) found that group decisions have a statistically significant negative effect on dictator game giving. This notion is also strengthened by Luhan et al. (2009), who claim that most of the literature across various bargaining games found that people in teams tend to behave more selfishly than people who decide individually. Furthermore, Luhan et al. (2009) also observe that teams exhibit more selfish behavior compared to individuals in their study. However, one study contradicts these findings. Cason and Mui (1997) found that people rather exhibited more other-regarding behavior in groups.

A related research topic covers the effect of strategic behavior from humans from computer interactions, which has gained interest among economists (March, 2021). In a review of strategic interactions between humans and computers consisting of 162 studies that employ computers as players in various experiments, March (2021) concludes that the literature spans several disciplines, such as economics, computer science and psychology. Further, earlier use of computer players in experiments can partly be attributed to various research objectives, such as reducing the risk of confounding factors from human behavior (March, 2021). In a recent review on 138 experiments investigating the interactions between humans and automated agents, Chugunova and Sele (2022) claim that the research has led to inconsistent results across numerous fields. In parallel, March (2021) argues that the research on strategic interactions between humans and computers has lacked common procedures that render the experimental studies incomparable (March, 2021).

When it comes to the effect that computers as players in various strategic games have on the outcomes and behaviors of human participants, March (2021) finds that in numerous social dilemmas, such as the prisoner's dilemma or the public goods game, the effect of social motives in the games appears reduced, making participants act more selfishly with computer players. March (2021) similarly concludes that in bargaining and contracting games that employ computer players, it is common to find that players exhibit more rational and selfish behavior that is less driven by emotions.

On the other hand, Kirchkamp and Strobel (2019) propose that humans behave less selfishly with a computer as a partner in a dictator game compared to players who are sharing the decision with a human. Based on arguments of to what extent a computer could be seen as responsible for an outcome, they argue that "the responsibility for a selfish outcome cannot be shared with a computer to the same extent as it can with a human. Thus, upholding a positive self- and social image while deciding selfishly together with a computer should not be as easy as when deciding with another human" (Kirchkamp & Strobel, 2019, p. 30). Kirchkamp and Strobel (2019) hypothesize that people who are in a joint decision with another human choose the selfish option in the dictator game to a larger extent than those who make decisions alone or with the computer. Kirchkamp and Strobel (2019) find that the number of selfish choices are higher when people make decisions together with a computer, compared to those who make an individual decision. Further, those who make the decision together with another human, exhibit an even higher frequency of choosing the selfish option. These effects are, however, not deemed as large and nor are they statistically significant.

In summary, findings from experiments on the interaction between humans and computers appear to provide reasons to believe that people behave differently when interacting with non-human agents compared to humans. However, evidence from some contexts point towards that humans seem to act more rational and selfish in some interactions, while others argue that humans would become less selfish with computers in other contexts.

2.2 The dictator game and social preferences

Several studies use the dictator game to study social preferences, non-selfish behavior and charitable giving (see e.g. Eckel & Grossman, 1996; Livingston & Rasulkhamedov, 2023). The dictator game does not contain any strategic aspects (Levitt & List, 2007) and has been widely used

in various forms (Engel, 2011). In its standard form, the game consists of a participant (typically student) being asked how to allocate an endowment of 10 dollars between themselves and another anonymous person (typically student) from the same pool of participants (Engel, 2011). While standard economic theory predicts that participants would keep their endowment to themselves, research has observed that a substantial share of people tend to share some amount of their endowment with the recipient (Engel, 2011). In dictator games where the recipient is a charity, the amount given tends to be higher compared to games in which a student is the recipient. For example, Eckel and Grossman (1996) study altruism in a dictator game, comparing the outcomes between a student and a charity recipient. Eckel and Grossman (1996) find that the charity received an amount that was three times larger than the student recipients and concluded that altruism can motivate human behavior.

In terms of theories suggested to motivate giving, Andreoni (1989;1990) propose a model of impure altruism where people can give to charity both driven by purely altruistic motives and be driven by the more self-interested motive of achieving a “warm glow”, from which the person would benefit from the act of giving itself. In addition, Andreoni and Bernheim (2009) suggest that giving can be motivated by social image concerns and that people achieve utility from being perceived as fair. Further, social norms prescribing how people should behave have also been proposed as an explanatory factor for giving in dictator games (Eckel et al., 2023). In contrast to theories on social preferences, factors considering the design of dictator games have also been suggested as explanatory factors. For example, Hoffman et al. (1996) argue that the degree of anonymity can be decisive when studying the motives of giving in dictator games and that observability from the experimenter can affect the outcome.

2.3 Hypotheses

Based on the previously outlined empirical observations in prior research and the theoretical framework, there are reasons to believe that humans behave differently when interacting with humans compared to non-human agents. However, I do not deduct a clear prediction from prior literature regarding in which direction such difference would occur in a non-strategic context of a dictator game where the charity is the recipient. Therefore, I derive the following non-directional hypotheses for the null and alternative hypotheses in the primary analysis.

H0: There is no statistically significant difference in the mean amount donated to the charity between the treated and the control group.

H1: There is a statistically significant difference in the mean amount donated to the charity between the treated and the control group.

As an extension of the primary analysis, I also formulate the following secondary null and alternative hypotheses to compare the joint and individual decisions between the two groups.

H0: There is no statistically significant difference in the mean change in shares donated to the charity between the treated and the control group.

H1: There is a statistically significant difference in the mean change in shares donated to the charity between the treated and the control group.

In the experiment, the treatment group consists of participants who are asked to imagine making the allocation decision together with an AI and the control group consists of participants who imagine making a decision together with another human.

3 The Experiment

3.1 A modified joint dictator game for charity

To address the research question, an experimental survey is constructed as it appears to be a suitable method to study the topic in question. An experiment generates data on the outcome of interest in a controlled environment, which yields the ability to infer causal effects (Drouvelis, 2021). Further, treatment can be randomly assigned to participants, which should hold constant other factors that affect the outcome across the two groups (Drouvelis, 2021).

My experimental design draws upon some features of the experiment conducted by Kirchkamp and Strobel (2019)¹, in which participants (students) either alone, with another human or a computer, decide in a binary dictator game with an equal or unequal option on how to divide money between themselves and a recipient. While the general design and aim of the study bears resemblance through its employment of the dictator game and the investigation of decisions

¹ Kirchkamp and Strobel (2019) also test for outcomes other than the choice in the dictator game, such as the perception of responsibility and guilt.

involving both humans and automated agents, this experimental survey has several distinct features. First and most importantly, Kirchkamp and Strobel's (2019) design uses a binary dictator game in which the team outcome depends on uncertainty about the other player's decision. To better isolate the effect on the outcome of which co-player a participant has, my experiment is designed without any component of strategy or uncertainty. Second, the recipient in this experiment is a charity, to serve as a proxy for an altruistic decision. This can be compared to the standard set-up used by Kirchkamp and Strobel (2019), where the recipient was a student. Third, Kirchkamp and Strobel (2019) use computers as their co-player and the computer acted in accordance with the choices that had been made by human players earlier in the experiment. This contrasts with the aim of my study, which specifically investigates how participants are affected by making a decision together with a co-player described as an AI agent. Fourth, in contrast to Kirchkamp and Strobel (2019), the individual dictator game decision is not included as its own treatment but in a sequence following the joint decision. This resembles the design employed by, for example, Luhan et al. (2009) and is also practically motivated, as an additional treatment requires a larger sample size to achieve sufficiently many observations in each treatment.

I write a detailed pre-analysis plan² outlining the experimental procedure and the analyses to be conducted. The plan was registered on the 3rd of May 2024, prior to the data collection, using the Open Science Framework (OSF)³. By specifying the analyses and questions to be addressed prior to observing the outcome, the pre-registration increases the credibility of the results (Nosek et al., 2018). I follow the pre-analysis plan consistently throughout this thesis. The pre-analysis plan can be found in Appendix A.

3.2 Experimental procedure

Monetary stakes and incentives

In the experimental survey (see Appendix B), participants allocate 100 SEK between two alternatives in two rounds. I use the national currency, in accordance with many other dictator game experiments (Engel, 2011). The national currency is used based on the assumption that SEK is the currency closest at hand when considering economic decisions in a Swedish context, which should make it easier for participants to allocate the money in accordance with their true preferences. The sum of 100 SEK approximates the standard dictator game endowment of 10

² The questions in the plan derive from the AsPredicted registration tool (Wharton Credibility Lab, n.d).

³ https://osf.io/j39eb/?view_only=1a73cf3878ca4aff8b01a472919e3a89

dollars (Engel, 2011). As has been common procedure in various other dictator games (Engel, 2011), the allocation choice is restricted to be stated in increments of one SEK and did not allow for cents (<1 SEK, “ören”) to be distributed.

All participants' decisions in the dictator game are not implemented (in line with, e.g., Carpenter et al., 2008), but the participants are nonetheless incentivized to act according to their preferences. I draw four random allocation choices from the pool of joint and individual decisions and donate actual money in accordance with these decisions. This procedure allows the provision of monetary incentives without collecting any form of identifying information from the participants. The participants are informed that their choices can have real monetary consequences both in the survey e-mail and on all occasions where they are being asked to make an allocation decision in the survey. The purpose is to enhance the probability that participants will answer in accordance with true preferences. As is a common practice in surveys to reduce potential hypothetical bias, the participants are explicitly encouraged to behave as if the allocation was made in a real situation (Freeman et al., 2014). In the joint decision, it is emphasized that the donation would be made in accordance with both the participant's choice as well as the decision made by the imagined co-player, to mirror a joint decision-making situation where it is ultimately the team's total outcome that matters.

Recipients and action set

The dictator game design in this experimental survey differs from the standard procedure, where participants are endowed with money which they can either choose to give to a recipient or keep for themselves (Engel, 2011). While practical limitations do not allow me to let participants choose to keep any money for themselves, I design the experimental procedure in this study using a proxy for a more selfish option. For this purpose, participants have the option to allocate money to the student association at the Stockholm School of Economics (SSE) (named SASSE). Given that SASSE arranges various activities for SSE students (SASSE, n.d.), I propose that it can serve as a substitute for the standard selfish option in the dictator game and that participants should be able to (more or less) directly take advantage of a donation to SASSE.

I select Barncancerfonden (the Swedish Childhood Cancer Fund) as the charity recipient. Barncancerfonden works against cancer in children by, for example, funding research (Barncancerfonden, n.d.). In the context of this thesis, I interpret giving to the charity as an act of

altruism. However, I do not test for the motives directly and therefore, I cannot exclude other drivers of charitable giving in the experiment.

First round – “joint” allocation decision

In the first round, the participants are asked to imagine that they make a joint decision together with an AI agent (treatment) or another human in the form of a student (control) in which they should allocate 100 SEK between SASSE and Barncancerfonden. I design the survey to mirror a joint decision-making context within the scope of what is practically implementable. Even though the experiment does not allow for students to interact with another player in the game, I frame the decision in ways that emphasize its ambition to mimic a joint decision. First, participants are asked to answer an attention question that emphasizes the key aspect that they are asked to imagine taking a decision together with another player. Second, participants are informed that if the allocation chosen in their joint allocation decision is chosen for an actual donation, both the imaginary co-player's and the participant's own decisions will be implemented.

As I could not arrange for an actual interaction between real players, it necessitates that the decision of the imaginary co-player is artificially constructed⁴. It is also necessary to hold the decision from the co-player constant across both the treatment and the control group, in order to facilitate a valid comparison between them. It was not an option to use deception and state that a human or an AI had taken a decision that they had not, in fact, taken as I follow the non-deceptive standard in economics (Niederle, 2015; Alós-Ferrer & Yechiam, 2020). A more rigorous experimental setting would have allowed for the pre-determined actions of the co-players to act in accordance with human behavior in prior studies. This is a procedure that has been adopted in studies using computers as players (March, 2021). However, due to the limited scope of this experiment, it was deemed necessary for simplicity to determine only one action of the imaginary co-player. I choose to let the imaginary co-player act in accordance with the more selfish option. This mirrors the finding in Luhan et al. (2009) that team decisions tend to invoke more selfish behaviors in a dictator game.

⁴ In the case where an actual non-human agent would have been used in the experiment, such practice would also have necessitated an action (programming) on my part deciding which strategy/behavior the agent should pursue (see e.g. Chen et al., 2018).

For simplicity, I keep the total pay-offs constant across the two decisions. This means that 100 SEK is the total money allocated in both the individual and the joint decision, but the participant only allocates 75 SEK in the first (joint) decision and 100 SEK in the second (individual) decision. The imaginary co-player is endowed with 25 SEK out of the 100 SEK. The reason for not providing the imaginary co-player and the participant with an equal endowment is to enable the participant's choice to be decisive of the outcome of the joint allocation. This is important to be able to assess to what extent the participants wished to either strengthen or reduce the effect of the prior (selfish) decision made by the co-player. Simultaneously, it is important that the imaginary co-player allocated a non-trivial sum (arbitrarily chosen as 25 percent of the full endowment of 100 SEK), in order to create a situation where the participant considered that the co-player's decision had somewhat of an impact on the total allocation.

To avoid that the treatment and control group receive a different amount of information, I do not formally define AI in the treatment condition. This should reduce the risk of participants leaving the survey due to the treatment, which could cause selection bias (Arechar et al., 2018). Attrition tends to occur more often in online experiments compared to traditional lab experiments (Arechar et al, 2018) and if the treatment condition would be lengthier and more technical in comparison to the control condition, there is a risk that participants randomized to the treatment group drop out to a larger extent. Instead of providing a formal definition in the survey, I inform the participants that they can consider the AI agent to be a chatbot, exemplified by ChatGPT from OpenAI. This specific example is chosen as it has become familiar to many given its recent spread, in which it was estimated to have gained approximately 100 million users in a short time after its introduction (Hu, 2023).

Second round – individual allocation decision

In the second round, participant are asked to allocate 100 SEK between SASSE and Barncancerfonden fully on their own. The second round could in isolation be described as a standard, individual dictator game for charity, with the main difference from the typical design being that there is no option for the participant to keep the money themselves. Instead, the SASSE-alternative serves as a proxy for the selfish option, as described earlier.

The participant is told that if the allocation in this second decision is randomly chosen for implementation, it will fully reflect the participant's own choice of allocation. This is stated to contrast the outcome of the joint decision, in which the allocation implemented consists of both

the respondents' own choice as well as the choice that was imagined having been made by another co-player.

In line with Luhan et al. (2009), I employ both an individual and a joint dictator game. However, it is important to distinguish that the main goal of this study does not lie in the difference between the joint and the individual donation per se, but comparing whether this difference varies between the treatment and control group.

Treatment and randomization

Participants are randomly assigned to either the control or the treatment group, using a between-participants randomizer provided in the Qualtrics survey tool. The treatment consists of the random assignment to one of two types of an imaginary co-player in the first round. Participants in the control group are assigned to imagine making a decision with another human (a student), which served as the baseline for comparison, while participants in the treatment group are told to imagine making a decision with an AI agent.

Pool of participants

I employ students as my pool of participants, which is the typical population in economic experiments (Fréchette, 2015). There are three main motives for this.

First, as proposed by Drouvelis (2021), one favorable aspect of using the standard pool of students is the ability to compare and replicate the experiment. Experiments in economics tend to include bachelor students and, sometimes, graduate students in business and economics (Fréchette, 2015). By using this standard participant pool, I enhance the comparability of my results.

Second, given that my objective is to investigate AI-agents as a co-player without emphasis on its technical aspects, I aim for a sample which could be assumed to have some familiarity with the concept. While I do not have information about the state of AI knowledge in my pool of participants, I draw on findings from a report of a representative sample in Sweden from 2022 on the knowledge of AI in Sweden (Insight Intelligence et al., 2022) in which they found that people who are younger, higher educated and reside in the capital area (Stockholm) tend to have larger knowledge about AI than other groups. These characteristics are expected to correspond well to the participant pool addressed in this study.

Third, using university students is convenient as I can easily contact them as a student myself at SSE. While I do not have the ability to recruit respondents in the traditional way and thus cannot directly control the sample size in this study, I aim to maximize the number of participants by contacting all current students at SSE apart from PhD-students. PhD-students are not included in the potential sample as I consider them different from other students as they are employed by the university and compensated monetarily. While I only contact students at SSE, there are no limitations for recipients of the study to further distribute the survey link to non-students or students outside SSE. In practice, however, I do not have any indications of such action having occurred, nor can I consider any obvious motives for why people would be interested in doing so.

The participants are fully anonymous, and I do not collect any demographic information. Participants are required to explicitly consent to participate in the study, confirming that they are 18 years or above. As pre-registered, I omit observations that are incomplete by the end of data collection.

As pre-registered, the analyses are performed regardless of the final number of respondents, but cautiously interpreted if the final sample contain less than 30 respondents in either the treatment and/or control group, following what the statistics book by Dahmström considers a “large” sample (2011, p. 272).

Neutrality and transparency

I aim to follow the standard procedure in economic experiments of refraining from so-called loaded framing, which could potentially lead to biased outcomes (Drouvelis, 2021). I strive for neutral language and avoid the use of concepts that can be perceived as loaded, such as the “dictator game”.

I take several measures to be transparent in the experimental design, to avoid any form of deception, as is common practice in economics (Niederle, 2015; Alós-Ferrer & Yechiam, 2020). I inform the students who receive the survey e-mail that the survey is experimental and that with whom they make the joint decision can vary and is hypothetical. I also explicitly state that I ask the attention question to assess the comprehension of the task.

Data collection

I use Qualtrics to construct an experimental survey. I send out the survey to students at SSE on Friday the 3rd of May 2024. The survey and the e-mail is provided in Appendix B and C, respectively. As stated in the pre-registration, the data collection ceased on Thursday the 9th of May 2024 at 7:00 pm UTC+02:00 time. The Qualtrics survey is also uploaded to my OSF project page.

3.3 Empirical strategy

To test the main hypothesis, I perform a two-tailed t-test to assess whether there is a statistically significant difference in the mean amount donated to charity between the treated and the control group. The main outcome variable is the amount of money in SEK donated to the charity (Barncancerfonden) in the joint decision.

To test the robustness of the findings in the primary analysis, I also perform two additional tests. First, I carry out a Mann-Whitney U-test, which is a non-parametric test and relaxes some of the assumptions of the t-test as it does not require normally distributed observations or identical variance in the underlying populations of the two samples (Körner & Wahlgren, 2006). Second, I perform a two-tailed t-test for difference in means in the outcome variable when excluding respondents that have not answered the attention question correctly. As the attention question serves as a measure on how well participants paid attention and understood the context of the task, this robustness test can indicate whether the main analysis contains noise from respondents who were answering the survey without actually considering the task at hand.

To test the secondary hypothesis, I carry out a t-test for the difference in means between the treatment and the control group regarding the change in the share donated between the joint and individual decision. The secondary outcome variable is the change in share donated to the charity, calculated as the difference between the share donated in the joint and the share donated in the individual decision and measured in percentage points.

Drawing upon the recommendations by Benjamin et al. (2018) aimed at increasing reproducibility in research and credibility of results, I characterize statistical significance in the analyses as results with a p-value of less than or equal to 0.005, and regard results with a p-value of 0.05 or less (often seen as the standard p-value threshold, see, e.g., Simmons et al., 2011) as suggestive evidence.

4 Results

Descriptive statistics for the sample are shown in Table 1. The total sample consists of 115 observations, out of which 56 are randomly assigned to the treatment group and 59 to the control group. The share of participants in the control and treatment group is 51 and 49 percent, respectively.

The mean amount donated to charity in the joint decision for the full sample is 69.02 SEK out of 75 SEK possible. The corresponding number for the individual decision is 88.21 SEK out of 100 SEK possible.

Participants on average donate 92.02 percent of their endowment in the joint decision and 88.21 percent in the individual allocation decision. The mean change in the share donated between the individual and joint decision is -3.79 percentage points, in which the mean in the treatment group is -5.14 percentage points and -2.55 percentage points in the control group.

As depicted in Table 1, the treatment group, in which people were asked to imagine making a decision together with an AI, on average donate a lower amount in point estimates and share of endowment to charity, compared to the control group. The treatment group also lowers the mean share given to charity in the individual decision to a larger extent than the control group, closely approximated as the double amount of change in the control group.

Table 1 Descriptive statistics

	Total sample	Treatment	Control	Difference (Treatment-Control)
Mean donation to charity joint (SEK)	69.017	67.607	70.356	-2.749
Mean share donated to charity joint (%)	92.023	90.143	93.808	-3.665
Mean donation to charity individ. (SEK)	88.209	85.000	91.254	-6.254
Mean share donated to charity individ. (%)	88.209	85.000	91.254	-6.254
Mean change in share donated to charity (p.p.)	-3.814	-5.143	-2.554	-2.589
N	115	56	59	

Table 2 below reports the share of participants who provides a correct answer to the attention question. 97 percent in total reports a correct answer, 96 percent in the treatment group and 98 percent in the control group.

Table 2 Attention question

	Total sample	Treatment	Control
Correct answers (%)	97.39	96.43	98.31

4.1 Primary analysis

The primary analysis tests whether there is a difference in the mean amount donated to charity between the treatment and the control group. The results from the main analyses and their corresponding robustness tests are provided in Table 3.

Table 3 T-tests and robustness tests primary analysis

	Two-tailed t-test	Mann Whitney	Excl. incorrect
	Difference in means	U-test	attention answers
P-value	0.285	0.384	0.187
Standard errors	2.560		2.553
N	115	115	112

* $p \leq 0.05$, ** $p \leq 0.005$

The two-tailed t-test for two independent samples testing the difference in means do not show a statistically significant difference between the treatment and control group ($p=0.29$). This finding is robust when excluding respondents who do not answer the attention question correctly ($p=0.19$). The Mann-Whitney U-test, employed as a robustness test, do not show any statistically significant difference between the treatment and control group ($p=0.38$), in line with the primary t-test. The main null hypothesis cannot be rejected at the five percent or 0.5 percent significance level. In other words, I do not find statistically significant or suggestive evidence that there is a difference between the treatment and control group in regard to the mean amount donated to charity in the joint decision.

Following the formula provided in Djimeu and Houndolo (2016) and their accompanying online calculator, I approximate the minimum detectable effect sizes (MDE:s) using the observed standard errors and the critical t-value for a two-tailed t-test with 80 percent statistical power and a significance level of 0.05 and 0.005 respectively. This yielded estimated MDE:s of 7.168 ($2.8 * 2.560$), for evidence on the five percent significance level and 9.344 ($3.65 * 2.560$) for evidence on the 0.5 percent significance level.

4.2 Secondary analysis

The secondary analysis tests whether there is a difference in the change in share donated to the charity in the joint and individual decision between the treatment and control group. The results from the secondary analysis are provided in Table 4. The two-tailed t-test testing the difference in means between the groups do not show a statistically significant difference between the treatment and the control group ($p=0.11$). The secondary null hypothesis cannot be rejected at the five percent or 0.05 percent significance level. There is no statistically significant evidence or suggestive

evidence that there is a difference between the treatment and control group regarding the change in share donated to the charity in the joint and individual decision.

Table 4 T-test secondary analysis

	Two-tailed t-test
	Difference in means
P-value	0.113
Standard errors	.0162
N	115

* $p \leq 0.05$, ** $p \leq 0.005$

5 Discussion

In summary, I did not find any statistically significant or suggestive evidence that there is a difference in the mean amount donated to charity or the mean change in shares donated to charity between people who imagine making a decision with an AI compared to people who imagine making a decision with a human co-player.

While I am not aware of any study using the specific experimental procedure of this thesis, the closest study that lends itself to comparison is arguably the paper by Kirchkamp and Strobel (2019), from which I have used features of their general design. In line with the findings in this thesis, Kirchkamp and Strobel did not find an effect of the type of partner (human or computer) on the choice made in the binary dictator game employed in their paper. In contrast, however, Kirchkamp and Strobel (2019) found a lower frequency of selfish choices in the computer-partner treatment compared to the human-partner treatment, while this finding was not statistically significant.

While I did not find any statistically significant differences in means between the treatment and control group, I am careful to interpret this finding as support for the notion that there is no effect on altruistic decisions and charitable giving decisions made with humans compared to an AI agent. This is due to several reasons.

First, the final sample in this study was relatively small. The number of final observations was 115 in total, out of which 56 were in the treatment and 59 were in the control group. Based on the

minimum detectable effect sizes calculated, the sample size seems to allow for only sizable differences between the treatment and control group to be detected.

Second, the results could have been affected by the fact that I could not let participants interact with an AI agent or with another person in reality, which limited this study to investigate the case in which participants were solely asked to imagine that they are making joint decisions. The answers on the attention question showed that almost all participants paid attention to the important detail describing that they were asked to imagine making a decision together with another player. However, since it was solely an instruction to imagine the scenario, there is a risk that the treatment was not sufficiently strong and that the experiment did not provide an impactful test of how decisions with others are made in an actual context.

Third, as in many other economic experiments, one can debate to what extent the student population is representative of the larger population. For example, Carpenter et al. (2008) found that non-students (community members) tended to donate more to charity in a dictator game than students. Carpenter et al. (2008) argued that findings from student dictator games where the recipient is a charity do not necessarily generalize to other populations. Given the variations in AI knowledge along different demographic characteristics found to be prevalent in a recent representative study in Sweden (Insight Intelligence et al., 2022), one might want to be especially cautious about generalizing the conclusions to other populations when it comes to AI specifically.

While the purpose of this paper was not to investigate the amount of the money donated to charity per se, the results in terms of the amount donated by participants can be indicative of potential threats to the internal validity, which should be acknowledged. As seen in the results, the mean share of endowment donated to charity was 92 percent in the joint decision and 88 percent in the individual decision. This can be compared to the number found in the meta-analysis of 616 dictator game treatments by Engel (2011), in which people on average gave 28.35 percent of their endowment, or in the dictator game by Eckel and Grossman (1996), where people who had a charity as the recipient in the dictator game gave on average 31.0 percent of their endowment.

The high amount of money donated to charity in the experiment could possibly be attributed to the way in which the dictator game was designed. While the dictator game was not hypothetical as it implemented four random allocation choices, participants never experienced being endowed with money which they could keep for themselves. It could also be that the student association did not

serve as a functional substitute for this aspect, as it was less directly beneficial to the participant compared to the ability to keep money for themselves. According to Kettner and Weichman (2016), in dictator games, “hypothetical bias occurs when self-reported measures overstate money transfers to peers compared to situations where participants actually bear the monetary consequences of their action” (p. 120). However, research on the effect of this bias in dictator games has been mixed (Kettner & Waichman, 2016), and it is therefore not easy to conclude to what extent this could have had an effect on the outcome.

Other aspects should also be kept in mind when addressing the results. I chose to interpret giving to the charity in this experiment as an act of altruism. This interpretation could be debated. I did not test for the participants’ motives for giving and can for example, not exclude the motive of warm-glow (Andreoni, 1989;1990) or social norms (Eckel et al., 2023). However, due to participants’ full anonymity both in regard to me as an experimenter and to other participants, I argue that it was less likely that social image (Andreoni & Bernheim, 2009) or observability (Hoffman et al., 1996) would serve as an important factor explaining the amount and share of charitable giving observed in this experiment.

The practical limitations required that I constructed the choice of the imaginary co-player and I chose to let it act in accordance with traditional economic theory, which predicts that the player should keep all their endowment (Engel, 2011). Participants were only faced with a single decision from their imagined co-player and the experiment did not investigate whether and how this aspect affected the amount of donations chosen by participants. One possible avenue for future research is to test whether the result found in this study would be robust if the co-player was being altruistic or acted in accordance with a distribution from human decisions in a specific dictator game (which, e.g., is the procedure that Kirchkamp and Strobel (2019) uses for their computer players). As done in some dictator games (Engel, 2011), one could also consider using the strategy method and let participants make their choice of allocation in the joint decision for various choices made by the imaginary co-player.

Further, it is also important to consider that the monetary incentives for participants were different in the individual and in the joint decision. Participants allocated 75 SEK in the joint decision and 100 SEK in the individual. For example, a participant that gave 75 SEK in the individual decision might have considered this to be identical to giving 75 SEK in the prior joint decision, while not taking into consideration that 75 SEK constituted only 75 percent of the endowment in the

individual decision while it constituted their full endowment in the first joint decision. Therefore, one should be careful to make comparisons between the two decisions within subjects in this experiment.

6 Summary and conclusion

The aim of this paper was to investigate whether the outcomes differ when people make joint decisions with an AI agent compared to with a human. I conducted an experimental survey following a pre-registered plan published before data collection. I did not find any statistically significant evidence that people who imagine making a decision together with an AI on average donated a different amount to charity than those who imagined making the decision with another human. I also did not find a statistically significant difference in the mean change in share of endowment donated between the two groups, when comparing the joint and an individual allocation decision. While the null hypotheses cannot be rejected at the five percent significance level, I am very careful to interpret this finding as support for the notion that there is no difference in the outcomes of decisions taken with humans compared to AI agents in this context. Rather, I acknowledge that, among other things, the hypothetical aspects and relatively small sample size in this experiment could have contributed to the result.

References

- Arechar, A. A., Gächter, S., & Molleman, L. (2018). Conducting interactive experiments online. *Experimental Economics : A Journal of the Economic Science Association*, 21(1), 99–131. <https://doi.org/10.1007/s10683-017-9527-2>
- Alós-Ferrer, C., & Yechiam, E. (2020). At the eve of the 40th anniversary of the Journal of Economic Psychology: Standards, practices, and challenges. *Journal of Economic Psychology*, 80, 102309. <https://doi.org/10.1016/j.joep.2020.102309>
- Andreoni, J. (1989). Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence. *The Journal of Political Economy*, 97(6), 1447–1458. <https://doi.org/10.1086/261662>
- Andreoni, J. (1990). Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving. *The Economic Journal (London)*, 100(401), 464–477. <https://doi.org/10.2307/2234133>
- Andreoni, J., & Bernheim, B. D. (2009). Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects. *Econometrica*, 77(5), 1607–1636. <https://doi.org/10.3982/ECTA7384>
- Barncancerfonden (n.d.). *Det här är Barncancerfonden*, Retrieved May 11, 2024 from <https://www.barncancerfonden.se/om-barncancerfonden/>
- Benjamin, D. J., Berger, J. O., Johannesson, M., Nosek, B. A., Wagenmakers, E.-J., Berk, R., Bollen, K. A., Brembs, B., Brown, L., Camerer, C., Cesarini, D., Chambers, C. D., Clyde, M., Cook, T. D., De Boeck, P., Dienes, Z., Dreber, A., Easwaran, K., Efferson, C., ... Johnson, V. E. (2018). Redefine statistical significance. *Nature Human Behaviour*, 2(1), 6–10. <https://doi.org/10.1038/s41562-017-0189-z>
- Carpenter, J., Connolly, C., & Myers, C. K. (2008). Altruistic behavior in a representative dictator experiment. *Experimental Economics : A Journal of the Economic Science Association*, 11(3), 282–298. <https://doi.org/10.1007/s10683-007-9193-x>

Cason, T. N., & Mui, V.-L. (1997). A Laboratory Study in Group Polarisation in the Team Dictator Game. *The Economic Journal (London)*, 107(444), 1465–1483.

<https://doi.org/10.1111/1468-0297.00233>

Chen, Y., Jiang, M., Kesten, O., Robin, S., & Zhu, M. (2018). Matching in the large: An experimental study. *Games and Economic Behavior*, 110, 295–317.

<https://doi.org/10.1016/j.geb.2018.04.004>

Chugunova, M., & Sele, D. (2022). We and It: An interdisciplinary review of the experimental evidence on how humans interact with machines. *Journal of Behavioral and Experimental Economics*, 99, 101897. <https://doi.org/10.1016/j.jsocec.2022.101897>

Dahmström, K. (2011). *Från datainsamling till rapport – att göra en statistisk undersökning* (5th ed.). Studentlitteratur.

Djimeu, E. W., & Houndolo, D.-G. (2016). Power calculation for causal inference in social science: sample size and minimum detectable effect determination. *Journal of Development Effectiveness*, 8(4), 508–527. <https://doi.org/10.1080/19439342.2016.1244555>

Drouvelis, M. (2021). *Social preferences : an introduction to behavioural economics and experimental research*. Agenda Publishing.

Eckel, C. C., & Grossman, P. J. (1996). Altruism in Anonymous Dictator Games. *Games and Economic Behavior*, 16(2), 181–191. <https://doi.org/10.1006/game.1996.0081>

Eckel, C. C., Hoover, H. G., Krupka, E. L., Sinha, N., & Wilson, R. K. (2023). Using social norms to explain giving behavior. *Experimental Economics : A Journal of the Economic Science Association*, 26(5), 1115–1141. <https://doi.org/10.1007/s10683-023-09811-z>

Engel, C. (2011). Dictator games: a meta study. *Experimental Economics : A Journal of the Economic Science Association*, 14(4), 583–610. <https://doi.org/10.1007/s10683-011-9283-7>

European Commission (2018). COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE EUROPEAN COUNCIL, THE COUNCIL, THE

EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS on Artificial Intelligence for Europe. Brussels 25.4.2018 COM(2018) 237 final.
<https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52018DC0237>

Fréchette, G. R. (2015). Laboratory Experiments: Professionals Versus Students. In *Handbook of Experimental Economic Methodology*. Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780195328325.003.0019>

Freeman, A. M. Herriges, J.A. & Kling, C.L. (2014). *The Measurement of Environmental and Resource Values : THEORY AND METHODS* (3rd ed.). Resources for the Future.

Gross, N. (2023). What ChatGPT Tells Us about Gender: A Cautionary Tale about Performativity and Gender Biases in AI. *Social Sciences (Basel)*, 12(8), 435.
<https://doi.org/10.3390/socsci12080435>

Hoffman, E., McCabe, K., & Smith, V. L. (1996). Social Distance and Other-Regarding Behavior in Dictator Games. *The American Economic Review*, 86(3), 653–660.

Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base - analyst note. *Reuters*. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>

Insight Intelligence, Bolagsverket, Postnord, Lunds universitet, Göteborgs universitet (2022). *SVENSKA FOLKET OCH AI : Svenska folkets attityder till Artificiell Intelligens 2022*.
<https://bolagsverket.se/download/18.6e6c29d18481cff0f2214/1670481583044/svenska-folket-och-ai-2022.pdf>

Kettner, S. E., & Waichman, I. (2016). Old age and prosocial behavior: Social preferences or experimental confounds?. *Journal of Economic Psychology*, 53, 118–130.
<https://doi.org/10.1016/j.joep.2016.01.003>

Kirchkamp, O., & Strobel, C. (2019). Sharing responsibility with a machine. *Journal of Behavioral and Experimental Economics*, 80, 25–33. <https://doi.org/10.1016/j.socec.2019.02.010>

Körner, S., & Wahlgren, L. (2006). *Statistisk dataanalys* (4th ed.). Studentlitteratur.

Levitt, S. D., & List, J. A. (2007). What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?. *The Journal of Economic Perspectives*, 21(2), 153–174. <https://doi.org/10.1257/jep.21.2.153>

Livingston, J. A., & Rasulkhamedov, R. (2023). On the Interpretation of Giving in Dictator Games When the Recipient is a Charity. *Journal of Economic Behavior & Organization*, 208, 275–285. <https://doi.org/10.1016/j.jebo.2023.01.025>

Luhan, W. J., Kocher, M. G., & Sutter, M. (2009). Group polarization in the team dictator game reconsidered. *Experimental Economics : A Journal of the Economic Science Association*, 12(1), 26–41. <https://doi.org/10.1007/s10683-007-9188-7>

March, C. (2021). Strategic interactions between humans and artificial intelligence: Lessons from experiments with computer players. *Journal of Economic Psychology*, 87, 102426. <https://doi.org/10.1016/j.joep.2021.102426>

Niederle, M. (2015). Intelligent Design: The Relationship Between Economic Theory and Experiments: Treatment-driven Experiments. In *Handbook of Experimental Economic Methodology*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195328325.003.0006>

Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115(11), 2600–2606. <https://doi.org/10.1073/pnas.1708274114>

Sheikh, Haroon. (2023). *Mission AI The New System Technology* (1st ed). Springer International Publishing. <https://doi.org/10.1007/978-3-031-21448-6>

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science*, 22(11), 1359–1366. <https://doi.org/10.1177/0956797611417632>

Student Association at Stockholm School of Economics (n.d). *Our Organization*. Retrieved May 6, 2024 from <https://www.sasse.se/about-us/organization/committees/>

Wharton Credibility Lab (n.d). *AsPredicted*. Retrieved March 11, 2024 from <https://aspredicted.org>

Appendix A – Pre-registration

https://osf.io/j39eb/?view_only=1a73ef3878ca4aff8b01a472919e3a89

1) Have any data been collected for this study already?

No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?

Primary analysis: The main research question is: Do people who make a joint money allocation decision together with an imagined human choose to donate a different amount to a charity compared to people who make the same decision together with an imagined artificial intelligence (AI) agent?

The main hypothesis being tested is that the average amount donated to a charity ("Barncancerfonden" - the Swedish Childhood Cancer Fund) is different between the group of participants who are asked to imagine making a joint donation decision with an imagined human (control group) and the group of participants who are asked to imagine making the same joint decision with an imagined AI agent (treatment group). The corresponding main null and alternative hypotheses are:

H0: There is no statistically significant difference in the mean amount donated to the charity between the treated and the control group.

H1: There is a statistically significant difference in the mean amount donated to the charity between the treated and the control group.

I do not have a clear prior of the direction of the hypothesis.

Secondary analysis: I will investigate whether there is a difference between the treatment and control group in regard to what extent they change their choice of donation to the charity in an individual allocation decision compared to the "joint" decision. After participants have made a decision where they were asked to imagine that they were making a decision with a human or AI agent, they are asked to make the decision again but this time individually and without any imagined input from someone else. I will construct a variable measuring the difference between each participant's individual and "joint" allocation decision, calculated as the share donated to

charity of the maximum amount that they are able to allocate. In the “joint” decision, participants can allocate a maximum amount of 75 SEK to the charity, which means that the share of money donated will be calculated as the participant’s choice of donation to the charity divided by 75. In the individual decision, participants can allocate a maximum amount of 100 SEK to the charity, which means that the share of money donated will be calculated as the participant’s choice of donation to the charity divided by 100. A donation of 0 will be coded as 0 percent. I will then calculate the difference between the shares in the individual and “joint” decision. Then I will test whether there is a statistically significant difference in the means between the treated and control group for this variable measuring the change in shares donated. The corresponding secondary null and alternative hypotheses are:

H0: There is no statistically significant difference in the mean change in shares donated to the charity between the treated and the control group.

H1: There is a statistically significant difference in the mean change in shares donated to the charity between the treated and the control group.

This hypothesis is also non-directional.

3) Describe the key dependent variable(s) specifying how they will be measured.

The primary dependent variable is the amount of SEK, in incremental of 1 SEK, donated to the charity (Barncancerfonden) in the “joint” decision.

The secondary dependent variable is the change in shares of money donated to the charity (Barncancerfonden), measured in percentage points as the difference between the individual and “joint” decision as a share of the maximum amount the participant can allocate.

4) How many and which conditions will participants be assigned to?

Participants will be told that they have 100 SEK to allocate between two alternatives, a charity (Barncancerfonden) and the student association at Stockholm School of Economics. In the first round, participants will make a “joint” decision and allocate the 100 SEK together with an imagined co-player. In the second round, the participants will allocate the 100 SEK fully on their own. I interpret donating to the charity as an altruistic act. Participants will also be asked an attention question.

In the “joint” decision, participants are asked to allocate 100 SEK between two options (a charity and the school’s student’s association), and they are told that 25 SEK of these have already been allocated to the student association. Participants are asked to imagine that this part of the decision was made by an imaginary co-player. Participants are then asked to allocate the remaining 75 SEK between a charity and the student association in any way they would like.

Participants will randomly be assigned to one out of two conditions. In the treatment group, participants are told to imagine that the 25 SEK has been allocated by an AI agent, whereas in the control group, participants are instead told to imagine that the 25 SEK has been allocated by another student.

After making the “joint” decision with an imagined co-player, all participants will be asked to allocate 100 SEK between a charity and the school’s student association fully on their own. Participants are told that 4 randomly drawn allocations (from the pooled sample of decisions), will be implemented with real money.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.

I will conduct a two-tailed t-test for difference in means between the treated and control group for both the primary and secondary analysis. I will conduct a Mann-Whitney U-test as a robustness test for the primary analysis.

In all analyses, I define a p-value of less than or equal to 0.05 as suggestive evidence and a p-value of less than or equal to 0.005 as statistically significant evidence, as suggested by Benjamin et al. (2018).

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.

To proceed with answering the survey, participants must consent. Survey responses that have not been finished (i.e. where the respondent has not reached the end of the survey at the time when the survey ends) will be excluded. The survey tool is designed so that respondents must provide answers to all questions to complete the survey. Further, the allocation decisions are designed so that respondents must provide an answer that sums to exactly 75 SEK and 100 SEK respectively.

I will include responses that have an incorrect answer on the attention question in both the primary and secondary analyses. For the primary analysis, I will however conduct a robustness test where I will exclude those with incorrect answers while conducting the two- tailed t-test for a difference in means between the treatment and control group.

7) How many observations will be collected or what will determine the sample size? No need to justify decision, but be precise about exactly how the number will be determined.

The survey will be sent out to students at Stockholm School of Economics, with the aim to maximize the number of respondents.

The data collection will stop at 7.00 pm UTC+02:00 time Thursday the 9th of May 2024.

The analyses will be conducted irrespective of the final sample size. However, if the final sample size for the primary and/or secondary analysis is less than 30 observations in either the treatment and/or the control group, the results and its implications will be interpreted very carefully.

I will estimate and report the minimum detectable effect size (MDE) that I had 80 percent statistical power to detect in the primary analysis. Based on the formulae in Djimeu and Houndolo (2016), the MDE will be estimated based on the observed standard errors as follows. For suggestive evidence (p-value of less than or equal to 0.05), the MDE will be estimated as 2.8 multiplied by the standard error. For statistically significant evidence (p- value of less than or equal to 0.005), the MDE will be estimated as 3.65 multiplied by the standard error.

8) Anything else you would like to pre-register? (e.g., data exclusions, variables collected for exploratory purposes, unusual analyses planned?)

The treatment and control group will be asked an attention question on their understanding of the task. The question also serves as a way to emphasize the crucial aspect that participants are asked to imagine making a joint decision:

Control group

I was asked to imagine making a decision together with...

- A professor
- Another student

- A family member

Treatment group

I was asked to imagine making a decision together with...

- Another student
- An AI agent
- A family member

I will report the share of participants that answer the attention question correctly as an indication of the respondents' attention in the survey. I will not exclude participants with an incorrect answer in the primary or secondary analysis but exclude them in a robustness test for the primary analysis.

9) Give a title for this pre-registration

Pre-registration for Master Thesis Study (Hanna Riese, 2024)

10) Type of Study

Experimental survey for master thesis

11) Data Source

Survey with university students

References

Benjamin, D. J., Berger, J. O., Johannesson, M., Nosek, B. A., Wagenmakers, E.-J., Berk, R., Bollen, K. A., Brembs, B., Brown, L., Camerer, C., Cesarini, D., Chambers, C. D., Clyde, M., Cook, T. D., De Boeck, P., Dienes, Z., Dreber, A., Easwaran, K., Efferson, C., ... Johnson, V. E. (2018). Redefine statistical significance. *Nature Human Behaviour*, 2(1), 6–10.

<https://doi.org/10.1038/s41562-017-0189-z>

Djimeu, E. W., & Houndolo, D.-G. (2016). Power calculation for causal inference in social science: Sample size and minimum detectable effect determination. *Journal of Development Effectiveness*, 8(4), 508–527. <https://doi.org/10.1080/19439342.2016.1244555>

Appendix B – Survey⁵

Information and consent

Welcome!

This is an experimental survey for a master thesis in Economics at SSE. The survey is conducted by Hanna Riese, master student at SSE. If you have any questions, please contact me at 42414@student.hhs.se.

The survey consists of 3 questions and takes around 3 minutes. You will be asked to allocate 100 SEK between two alternatives. First, you will make the decision together with an imagined co-player. Then, you will make the decision on your own. I will also ask a question on your understanding of the task. There are no personal or demographic questions.

While you do not allocate your own money in this survey, I will randomly draw 4 anonymous allocations and make real donations. This means that your choice can have actual monetary consequences in reality. Therefore, please try to act as if the choice you make is real. Your response is fully anonymous and your participation is completely voluntary.

To start, please state below that you agree to participate in the survey, confirming that you are at least 18 years old.

- ☐ I agree
- ☐ I do not agree

⁵ Note: Participants in the survey could not observe the question titles, nor whether they were in the treated or the control group. The emphasis in italics solely serves the purpose to illustrate how the information and parts of the survey varied between the treatment and the control group. It was not displayed in red to participants.

Treatment Task

You are asked to allocate 100 SEK between two options. Imagine that this allocation choice is made together with an artificial intelligence (AI) agent. You can think of this agent as for example a chatbot, like ChatGPT from OpenAI.

You will be asked to allocate money between the two following options:

Option A – SASSE, the Student Association at SSE that, among other things, organizes various activities for students at SSE.

Option B – Barncancerfonden (the Swedish Childhood Cancer Fund), a charity that combats cancer in children by, among other things, funding research.

Before you make the decision, I will ask you a question about your understanding of the task.

Which of the following is true about the information that was shown to you above?

I was asked to imagine making a decision together with...

- *Another student*
- *An AI agent*
- *A family member*

Now, 25 SEK of the 100 SEK have already been allocated to option A – SASSE. *Imagine that it is the AI agent which has made this part of the decision.*

Next, you will be asked how you wish to allocate the remaining 75 SEK.

Control Task

You are asked to allocate 100 SEK between two options. Imagine that this allocation choice is made together with another student.

You will be asked to allocate money between the two following options:

Option A – SASSE, the Student Association at SSE that, among other things, organizes various activities for students at SSE.

Option B – Barncancerfonden (the Swedish Childhood Cancer Fund), a charity that fights cancer in children by, among other things, funding research.

Before you make the decision, I will ask a question on your understanding of the task.

Which of the following is true about the information that was shown to you above?

I was asked to imagine making a decision together with...

- ☐ *A professor*
- ☐ *Another student*
- ☐ *A family member*

Now, 25 SEK of the 100 SEK have already been allocated to option A – SASSE. *Imagine that it is the other student who has made this part of the decision.*

Next, you will be asked how you wish to allocate the remaining 75 SEK.

Joint Allocation

You are now asked to allocate the remaining 75 SEK freely between option A – SASSE and option B – Barncancerfonden.

Please choose your allocation below. Your allocation must sum to exactly 75.

If your choice of allocation in this question is randomly chosen for a real donation, I will donate the total outcome of the 100 SEK. This means that I will both donate the 25 SEK to SASSE decided earlier and then add your decision of how you allocate the remaining 75 SEK between SASSE and the charity Barncancerfonden.

- *Option A: SASSE [0-75 SEK]*
- *Option B: Barncancerfonden [0-75 SEK]*

Individual Allocation

Now you are asked to allocate 100 SEK between option A – SASSE and option B – Barncancerfonden **fully on your own**.

Please choose your allocation below. Your allocation must sum to exactly 100 SEK.

If your choice of allocation in this question is randomly chosen for a real donation, I will allocate the full 100 SEK in accordance with your choice of amount to donate to SASSE and/or the charity Barncancerfonden.

- *Option A: SASSE [0-100 SEK]*
- *Option B: Barncancerfonden [0-100 SEK]*

Appendix C - E-mail

Subject: 3-minute survey for master thesis

Dear SSE student,

My name is Hanna Riese. I am conducting an experimental survey* as part of my master thesis in Economics here at SSE and I would be very happy if you could participate. This is a short survey with only **3 questions** that takes about **3 minutes** to complete. I will not ask for any personal or demographic information. You will be asked to imagine making an allocation of money between two options.

You do not allocate your own money, but your choice has potentially real monetary consequences as I will make actual donations in accordance with four random and anonymous allocation decisions.

Your response is very valuable, as a sufficiently large sample size is needed to be able to draw valid conclusions.

You will find the survey here: *[Qualtrics Survey Link]*

Data protection: Your response will be fully anonymous and your participation in this experimental survey is completely voluntary.

Thank you so much for contributing. If you have any questions, do not hesitate to contact me at 42414@student.hhs.se

Kind regards,
Hanna

*Experimental survey, in the context of this master thesis, means that you will make a decision that can have real monetary consequences, but with whom you make it with will be hypothetical and can vary between participants.