Lying under time pressure and how to promote honesty

Master Thesis – Behavioral Economics

Laura Weyergans

Student Number: 483107

Supervisor: Dr. Jan Stoop

Second assessor: Cem Peker

Date final version: 26/06/2019

Abstract

Undoubtedly, lying is a significant part of our everyday lives – to a larger extent than we might grasp. When it comes to lying under time pressure, scientific literature suggests that people fail to contemplate and exert self-control. However, recent literature cannot distinguish whether this lack of deliberation yields increased truth-telling or lying. To understand those two opposing lines of research, I conducted an experiment in which people had the opportunity to cheat in order to enhance self-benefit, both under time pressure and with ample time. Moreover, as tools to promote honesty under time pressure are yet unexplored, I bridged the concept of linguistic cues to decision-making under limited time. Scholars found that exposing people to the self-relevant noun ‘cheater’ decreases lying behavior. Thus, I further implemented this insight in my experiment. The results show that there is no difference in lying or truth-telling behavior under time pressure compared to when individuals have ample time. Furthermore, this experiment fails to find a significantly lower level of lying behavior when people are exposed to the word ‘cheater', in both the time pressure and no time pressure treatment. Nevertheless, this research may serve as a stepping stone for future research in the field of dishonesty prevention under time pressure, as understanding lying behavior and how to prevent it remain crucial.

Keywords: lying, honesty, morality, time pressure, decision making, linguistic cue, phrasing differences

Table of Contents

List of Tables

List of Figures

1. Introduction

Once in a while, everyone may catch themselves telling a lie. Regardless of whether it is lying out of pure self-interest or just an exaggeration like saying "I went for a run this morning, for 10km", (even though it was only 8.5 km) dishonesty is present in everybody's life. It is undeniable that in some situations, one cannot resist the temptation to lie to get a higher reward. Astonishingly, DePaulo, Kirkendol, Kashy, Wyer and Epstein (1996) found that 20% to 31% of social interactions are coined with lies. Additionally, three out of four lies are told out of pure self-interest (Camden, Motley, & Wilson, 1984). Taking these results together, people engage in dishonesty in 15% to 23% of their personal interactions. Such interactions can be both in a social or business context. Yet, people are unaware of the consequences even a supposedly small lie might entail. The annual costs due to fraud in business and employee theft are estimated to be $600 billion (Ariely, 2008). This figure only depicts the amount measured in business without taking social or private interactions into account. Thus, annual costs of lying potentially even surpass $600 billion.

Who is not familiar with the proverb "Time is money"? Nowadays, people generally live under time pressure and often lack time to deliberately think decisions through. This comes into play in social and business interactions. As people under limited time rely on heuristics, scholars tried to figure out whether individuals per se are honest or dishonest. However, the literature fails to give a definite result whether lying or truth-telling prevails under time pressure. This leaves valuable room for research and calls for more exploration.

Inspired by this discrepancy in the literature, I try to figure out if people lie less or more under time pressure. Furthermore, I aim to find a tool to promote honesty under lack of time. Results given by Bryan, Adams and Monin (2013) show that making use of subtle linguistic cues can increase honesty. For instance, exposing people to the noun 'cheater' in the introduction text of a task, when afterward provided with an opportunity to cheat, made a notable difference. The mere exposure to the word 'cheater' reduced lying behavior in comparison to no exposure to the word. As this insight seems quite promising,

I make use of it in this study. I conduct an experiment to bridge the concept of linguistic cues with lying under time pressure. Hence, this study formulates the following research question:

Does lying or truth-telling prevail under time pressure, and can this lying behavior be reduced by subtle linguistic cues?

Since this, to the best of my knowledge, has never been tested before, I step into a yet unexplored area of research. Motivated by the importance of promoting honesty, this study tries to fill the void in the scientific literature about lying behavior under time pressure.

The present study is structured as follows. First, I present related scientific literature about lying behavior under time pressure and subtle linguistic cues. After that, the hypotheses for this research are derived. Thirdly, the study elaborates on the methodology and the research design used. Moreover, the results of the conducted experiment are reported. Lastly, the findings of the study are put into context by discussing the findings in light of previous scientific literature, and recommendations for future research are provided.

2. Literature Review

This chapter provides an overview of the existing literature about lying behavior and subtle linguistic cues. At the end of the first subsection, I combine these two threads of literature. The second and last part of this chapter derives the hypotheses of the research at hand.

2.1 Lying

2.1.1 Definition of lying

Literature fails to state a unified definition of lying. Papers by Carson (2006) and Fallis (2009) that solely focus on the meaning and concept of lying underpin the prevalent discrepancy between scholars. On the one hand, the definitions vary around the intention of the lie-teller, whether the purpose of telling a lie is to deceive someone or not (Fallis, 2009). On the other hand, the concept of lying varies regarding the lie-teller's beliefs about

the statement's truthfulness. The Oxford English Dictionary (2019), for instance, defines lying as the deliberate action of telling falsehoods to mislead. According to Carson (2006), this definition, however, does not capture the fact that the lie-teller can never believe his own statement to be true.

This research uses a fairly broad definition by DePaulo et al. (1996), stating that lying is "intentionally [trying] to mislead someone" (p.981).

2.1.2 Different kinds of lies

Before diving into the motivations of people to lie, it is crucial to establish a more profound understanding of the aspects of a lie. A common taxonomy for a lie is divided into four main components, namely the content, the reason, the type and the referent of the lie (DePaulo et al., 1996). As the name content already implies, that component describes what the lie is about, such as feelings, achievements, actions, facts or possessions (DePaulo et al., 1996). The reason for the lie differentiates between self-oriented or other-oriented lies. Other-oriented lies can be told to guard someone from being harmed, so-called white lies (Erat & Gneezy, 2012). Self-oriented lies on the other side can be black-lies if the lie is told to benefit at the expense of others (Capraro, 2017; Erat & Gneezy, 2012). The third component, the type of lie, refers to the degree of falsehoods. This implies, for instance, whether the content is contradictory to the truth or an exaggeration. Lastly, the referent of the lie describes what the lie refers to, such as another person or an object or event (Cappelen, Sørensen, & Tungodden, 2013; DePaulo et al., 1996). Variations in all the above-stated aspects of a lie influence whether someone lies and under which circumstances she does so.

2.1.3 When do people lie?

The abundance of studies about lying emphasizes the vast interest among scholars in trying to explain why people engage in unethical behavior. Several reviews on the literature of lying attempt to guide researchers to a more profound understanding of dishonesty and the motivations behind it (Gino, 2015; Jacobsen, Fosgaard, & Pascual-Ezama, 2018; Rosenbaum, Billinger, & Stieglitz, 2014). The subsequent paragraphs present a glimpse at these results.

A primary explanation of why people lie is based on rational analysis (Becker, 1968). Thus, whether to lie merely depends on an intrinsic cost and benefit calculus. In case the benefits of lying outweigh the costs of lying, humans switch from honesty to dishonesty. The benefit of lying is usually described as a higher monetary reward, whereas costs include chances of getting caught and the possible fine that comes with it (Becker, 1968). Today, this economic rational approach is still used in explaining why humans tell lies (Kajackaite & Gneezy, 2017). However, this cost-benefit analysis is a rigid concept and does not leave room for any situation in which lying might be justified. Even though Immanuel Kant in 1785 famously formed a clear stance against lying, lying is not seen as a dichotomy of either being morally right or wrong anymore (Kant, 2012). It is argued that white lies, for instance, might be justified, as such lies prevent others from being harmed (Ariely, 2012; Erat & Gneezy, 2012). Hence, lying is not always regarded as entirely immoral anymore, even though telling a lie is still mostly referred to as being unethical (Perkins & Turiel, 2007).

Today's explanations of dishonesty provide a more sophisticated notion of lying. As explained by Ariely (2012), dishonesty stems from an economic motivation as well as from a psychological motivation. This means that there is more to dishonesty than just rational intrinsic calculations. On the one hand, there is an economic reward, such as money, when engaging in dishonesty (DePaulo et al., 1996; Kajackaite & Gneezy, 2017). On the other hand, the lie-teller wants to maintain a certain self-image of being an honest person (Mazar, Amir, & Ariely, 2008). It seems to be a win-lose situation in which one side will be sacrificed to a certain extent. Thus, the decision of being dishonest is influenced by more factors than just some costs and benefits (Mazar et al., 2008).

Other forces further shape dishonesty outside the economic or psychological scope. Research has shown that the more creative people are, the better their abilities in storytelling and the higher the chances to lie (Ariely, 2012; van der Zee, Anderson, & Poppe, 2016). Both studies also state that depletion influences the likelihood of lying. If people are cognitively depleted, they lose the ability of deliberation and hence act more impulsively, which sometimes includes telling lies. Besides, pursuing a clear goal might induce people to put morality aside and focus on achieving the agenda. Therefore, the

chances of lying increase (DePaulo et al., 1996; Schweitzer, Ordonez, & Douma, 2004). Furthermore, Kouchaki and Desai (2015) and Kouchaki and Wareham (2015) found that experiencing anxiety or feeling socially rejected elevates the probability of behaving dishonestly as people mainly focus on getting a higher reward. Moreover, the more time people spend contemplating, the better they are in developing justifications for engaging in lying. Utilizing those justifications, people substantiate the decisions they make as they often feel compelled to explain themselves (Gunia, Wang, Huang, Wang, & Keith, 2012). The explanations enable humans to act more dishonestly but still uphold a certain moral self-image of being an honest person (Mazar et al., 2008; Shalvi, Dana, Handgraaf, & De Dreu, 2011). Lastly, the saliency of lies in social interactions can elevate the number of lies being told. The more people can observe other people lying, the lower the barrier of being dishonest themselves (Gino, Ayal, & Ariely, 2009; van der Zee et al., 2016).

In sum, the literature suggests that a cost-benefit analysis of lying is a primary driver of lie telling; however, other factors such as justifications, maintaining an honest self-image, or the content of the lie may also influence lie telling behavior.

2.1.4 Lying under time pressure

After acknowledging that the process of lying is not as straight-forward as one might expect, it is important to understand how this internal lying mechanism works under limited time. Generally, as people have limited time to make a decision, they lack time for deliberation and self-control (Bereby-Meyer & Shalvi, 2015). Such a shortage of contemplation induces people to rely on heuristics and to make decisions with little effort based on the automatic and intuitive system 1 (Capraro, 2017; Kahneman, 2011).

There is no clear evidence on how time pressure affects lying behavior. On the one hand, research has shown that honesty requires time (Bereby-Meyer & Shalvi, 2015; Gunia et al., 2012; Shalvi, Eldar, & Bereby-Meyer, 2013). In the experiment conducted by Shalvi et al. (2013), participants had to anonymously roll a die and report the number accordingly. This reported number then determined their payoff. Participants had either ample time to roll the die and report the number or just 20 seconds. Furthermore, Gunia et al. (2012) tested lying behavior with a controlled laboratory experiment by anonymously

matching two players. Player 2 had to decide on a payoff distribution that is paid out to players 1 and 2, based on information received by player 1. Only player 1 knew the payoff structure that there was no equal payoff. Player 1 then had to decide whether to send a truthful or untruthful message to player 2 explaining the payoff structure. The results of these studies show that the yield to temptation of a higher reward elevates lying behavior. People under time pressure initially serve self-interested desires before having time to contemplate and time to understand the full extent of possible costs or consequences. Further, it has been found that in time pressure situations subjects are more receptive to give in to the personal benefit of the lie and hence lie more often. Lying then can be regarded as the default option (Bereby-Meyer & Shalvi, 2015). Hence, especially when it comes to self-oriented kinds of lies, dishonesty seems to prevail.

On the other hand, some scholars also found the opposite. Capraro (2017), for instance, tested lying behavior with the deception game. Participants got assigned to one out of two groups and then anonymously matched with a partner. Answering truthfully to which group the subject got assigned to yielded equal payoff for both, lying yielded a lower payoff for the matched partner. The results of this study indicate that the truth comes naturally and faster than a lie. Spence et al. (2001) used another way of testing response time and lying behavior. They let participants answer the same set of questions twice. A color cue signaled whether they should lie or tell the truth. Such a tool to test lying behavior is coined as the Sheffield lie test. The results also showed that response time when answering truthfully was significantly lower than when answering with a lie. Verschuere, Spruyt, Meijer and Otgaar (2011) confirmed this finding by also making use of a variation of the Sheffield lie test in their experiment. Additionally, Foerster, Pfister, Schmidts, Dignath and Kunde (2013) tested lying by making use of the die roll experiment explained in the preceding paragraph. However, they manipulated the time participants had between rolling the die and reporting the number. They further affirm that honesty appears to be the immediate response. Therefore, these studies indicate that honesty seems to be the default option. Regardless of the kind of lie or the economic motivation, the tendency to be honest is dominant.

Verschuere and Shalvi (2014) acknowledged this divergence and tried to explain the conflicting results based on two theories, the Truth-default theory (TDT) and the Information manipulation theory 2 (IMP2). The TDT states that being honest is the automatic tendency in social conversations (Levine, 2014). This theory focuses mainly on the evaluation of incoming messages by the receiver, saying that people tend to believe what is said. Levine (2014) only implicitly states that humans are inclined to opt for the truth. Hence, the theory does not focus enough on the sender's perspective of lying or truth-telling under time pressure (Verschuere & Shalvi, 2014). The IMP2 theory was introduced by McCornack, Morrison, Paik, Wisner and Zhu in 2014. They specify that dishonesty might be just as natural or even more natural than telling the truth. In general, humans always strive to maximize efficiency in all aspects of behavior. Lying needs to be efficient regarding potential costs and benefits (McCornack et al., 2014). Therefore, this theory focuses on the lie teller's intention rather than the assessment of incoming messages. Trying to explain this disparity even further, Verschuere and Shalvi (2014) propose motivation to be one of the main drivers of lying behavior. A motivation might be lying to obtain a higher self-profit, for instance, in the form of money or praise. In order to evaluate the researchers' findings of whether truth-telling or lying prevails under time pressure, one should carefully map the circumstances under which a lie or the truth is told. Truth-telling might, in fact, be the default option without a strong motivation to behave dishonestly. If indeed, the motivation of lying is clear and dominant, dishonesty might be the automatic tendency (Verschuere & Shalvi, 2014).

In sum, whether lying or truth-telling dominate under time pressure is debatable. However, scholars have agreed upon lying behavior under time pressure to be a question of what the default option is since there is no time for deliberation. One proposed bridge between the divergent findings on lying under time pressure is the motivation behind the lie.

2.2 Subtle linguistic differences

The second part of this literature review addresses subtle linguistic differences. Such linguistic differences can be described as alterations in question or introduction phrasing (Minson, VanEpps, Yip, & Schweitzer, 2018). They influence the responder's behavior

and may affect the content of her response (Gelman & Heyman, 1999; Minson et al., 2018; Walton & Banaji, 2004). Minson et al. (2018) revealed, for instance, that carefully phrased questions induced people to be more honest and disclose more sensitive private information.

Studies on linguistic changes suggest that word manipulations in the instructions or question phrasing using a noun are more powerful than using a verb (Bryan, Walton, Rogers, & Dweck, 2011; Bryan et al. 2013; Gelman & Heyman, 1999; Walton & Banaji, 2004). This particularly applies to nouns which allow self-identification, such as 'carrot-eater' or 'voter'. Exposing people to those self-relevant nouns creates the opportunity of self-identification as people can easily be a carrot-eater or a voter. Experimenters can highlight the desirable identity by including that noun in the task instructions or survey questions (e.g. placing the word 'voter' in a questionnaire about elections) (Bryan et al., 2011). This effect is different for verbs, as verbs do not have the ability to evoke self-perception. Reading a verb instead of a noun does not facilitate the assessment of possible self-preferences (Walton & Banaji, 2004).

Furthermore, especially regarding dishonest behavior, Bryan et al. (2013) found that using the power of such linguistic differences evoked honesty. In their study, they exposed participants to the sentences "Please do not be a cheater" then afterward provided an opportunity to easily cheat and claim more money. The results indicate that the presentation to such a subtle but pervasive linguistic change reduced the lying behavior (Bryan et al., 2013). They explain this shift in behavior by people's desire to maintain a positive and honest self-image. Therefore, invoking identity concerns with self-relevant nouns by linking the unethical behavior with the associated undesired identities is a powerful tool to influence lying behavior.

2.3 Hypothesis derivation

This chapter combines the aforementioned research results and derives the hypotheses for this research. Even though the literature indicates that no clear pattern emerges concerning lying under time pressure, I suspect that time pressure inflates dishonesty, which is in line with the findings by Bereby-Meyer and Shalvi (2015), Gunia et al. (2012)

and Shalvi et al. (2012). Firstly, if people lack time to contemplate, they yield to temptation and lie to increase their benefits. Secondly, as there usually is a clear motivation when it comes to tempting situations, people serve self-interest without thinking about possible consequences. As explained in chapter 3, I provide a tempting situation evoking a clear motivation to lie in order to have a higher chance to win a reward. Thus, I postulate the first hypothesis as follows:

Hypothesis 1: People under time pressure lie more compared to people having ample time.

In accordance with the insights given by Bryan et al. (2013), subtle linguistics changes are a powerful tool to alter people's behavior. Reminding people of moral standards by exposing them to self-relevant nouns such as 'cheater', reduces dishonesty. Hence, I hypothesize the following:

Hypothesis 2: Adding the linguistic change in phrasing with 'Please do not be a cheater' decreases lying behavior compared to no exposure to the phrase.

A combination of both influences has not been tested, yet. Thus, it is unknown how people react under time pressure with moral reminders. However, as time pressure is supposed to increase the lying behavior and subtle linguistic changes decrease lying behavior, a conjunction may decrease the proportion of lying. Consequently, the effect of adding the word 'cheater' is expected to mitigate the effect of time pressure. This leads me to derive the following:

Hypothesis 3: Adding the phrase 'Please do not be a cheater' under time pressure decreases lying behavior compared to no exposure to the phrase under time pressure.

3. Methodology

This section elaborates on the methodology and the research design of the experiment at hand.

3.1 Experimental design

For the study, an anonymous one-shot experiment was conducted. Each subject participated only once so that observations are independent on the subject level. This implies that the choices of other participants could not influence the participants. Independent observations are crucial for the statistical analysis afterwards so that there is no interaction of the data points. Furthermore, a between-subject design was chosen. Such a design is more applicable in this particular experiment. Firstly, even though more powerful statistical tests can be performed and a smaller sample size is sufficient for a within-subject design, learning confounds between the tasks can bias the results in a within-subject design. Secondly, other possible confounds, such as the order effect (the order in which the participants do the tasks) need to be taken into consideration when using within-subject designs (Charness, Gneezy, & Kuhn, 2012). In this study, it is crucial that the participants were exposed to the treatments in isolation and hence did not realize a change in the environment, which might yield changing tactics.

Additionally, the experiment was of a single-blind character. Thus, subjects were anonymous towards other subjects. As anonymity strongly influences lying behavior it is crucial to provide the highest degree of anonymity as possible.

In addition, the subjects were randomly assigned to the treatment conditions. Randomization is important for multiple reasons: Firstly, it overcomes confounding effects. Imagine wealth effects: Participants having a higher income than others might influence the decision of lying, as wealthier participants do not need some extra money besides their income. Another example is the subject's mood. Being in a particularly good or bad mood might influence the lying behavior. Therefore, making use of randomization minimizes the risk of correlation of unobservable characteristics with individual characteristics. Secondly, randomization avoids selection-bias. If subjects self-select into treatments, the results can readily be biased, as the potential outcome is not independent of receiving the treatment anymore (Heckman & Smith, 2005). Thus, making use of randomization facilitates inferring causalities.

3.1.1 Experimental description

The experiment was conducted by means of an online survey tool[1]. Features such as randomization could easily be set-up, to ensure the experimental design explained in the preceding paragraph. After deciding to participate in the experiment and after being randomly assigned to one treatment, each participant had to go through two stages during the experiment.

In the first stage, participants had to roll a die and report the number of the first die roll. The reported number determined the possible payoff. In the introduction of the experiment, I explained the payoff structure as follows: Rolling a one yielded one lottery ticket, rolling a two yielded two tickets, rolling a three yielded three tickets and so forth. This payoff structure was inspired by Shalvi et al. (2012).

As the experiment was conducted online, I simulated the die roll online. Participants were led to an external website[2] to roll the die. Making use of an external website eliminated potential concerns that the die result could be traced. Furthermore, as the die results could not be tracked, the degree of anonymity was strengthened by reducing possible fears of getting caught.

If lying or truth-telling under time pressure is indeed driven by motivation (see Chapter 2.1.4), I tried to lower the barrier of lying as much as possible so that there is a clear motivation for lying: Firstly, the die result was anonymous, as described in the preceding paragraph. Secondly, the higher the number reported, the more lottery tickets the participants could earn. Lastly, to ensure legitimacy, subjects were allowed to roll the die as many times as they wanted. This further gave room for justifications when not reporting what was rolled the first time (see Chapter 2.1.3).

During the second stage of the experiment, subjects answered standard demographical questions on age, gender, whether they are a student or not, income besides the study loan and the country of origin. On the one hand, this enables me to obtain descriptive

[1] The tool used is Qualtrics.
[2] Website used: https://www.google.com/search?q=roll+a+die

statistics of the sample. On the other hand, such demographics add further control for the later analysis. Yet, this was not expected, due to randomization into treatments. I chose to control for these specific explanatory variables, as those are the most common ones affecting lying behavior (Rosenbaum et al., 2014). Please consult Appendix A1 and A2 for the survey itself.

3.1.2 Incentives and precepts

In order to conduct a controlled economic experiment, five precepts introduced by Smith (1982) need to be satisfied. A controlled experiment implies, that subjects do not react randomly to the incentives provided. In the following five paragraphs, I describe each precept separately.

The presented experiment made use of a reward structure using monetary incentives. By means of money as a reward the first out of five precepts 'Non-satiation' was satisfied. The utility function is a monotone increasing function of monetary reward and for each participant, it holds that more is better. However, as the budget was limited, there were no resources to pay out all participants. Therefore, a binary lottery incentive structure was implemented so that the subjects could earn lottery tickets for the decisions taken in the experiment. Three participants who were willing to leave their email address were randomly drawn and monetary prizes of 25€, 15€ and 10€ were paid out[3].

Moreover, as all subjects received the same information about the payoff structure and the money was actually paid out, there was no deception regarding the reward. As it was impossible for me to trace the die result of any roll, as it was stated, deceiving in any other way was not used. However, it could not easily be defined what the bad and the good outcome were, and that the reward was increasing in the good outcomes. Hence, the second precept, 'saliency', was only partially complied.

Even though a researcher can never know what is an adequate payment to incentivize participants to think hard enough about the task, the chosen reward structure should have

[3] For the selection of the winners, the following random generator was used: https://www.randomlists.com/list-randomizer

been sufficient to dominate all subjective costs incurred by the participants and impede decreasing performance with too high incentives, as other researchers have used similar incentive structures (Ariely, Gneezy, Loewenstein, & Mazar, 2009). Therefore I conclude that the ‘dominance’ precept by Smith (1982) was satisfied.

Each participant was only informed about his or her payoff structure without giving any information about payoff alternatives to other subjects. Ensuring such privacy is essential to rule out any possibility that subjects can attach weights to other subjects’ payoff. Therefore, the ‘privacy’ precept was fulfilled.

The last precept, ‘parallelism’, is not mentioned as a sufficient condition in setting up a controlled experiment. Nevertheless, this precept was also satisfied during the experiment. It can be assumed that general laws of behavior apply everywhere. Underpinned by using a common method eliciting lying behavior during experiments, it can be stressed that the condition is met (Fischbacher & Föllmi-Heusi, 2008; Foerster et al., 2013; Shalvi et al., 2011, 2012).

As the sufficient conditions were only partially complied, the study at hand can only partially be considered a valid controlled economic experiment.

3.2 The treatments

The randomized controlled experiment itself divided subjects into four different treatments. The treatments differed in two disparate aspects, so that a two by two design evolved, depicted in table 1.

In the baseline treatment, participants were exposed to a neutrally framed introduction explaining the task and the payoff structure in detail. They had ample time to roll the die and to report the figure of the first die roll. Such a design is proven to be effective, as it has been used multiple times by other researchers (Fischbacher & Föllmi-Heusi, 2008; Foerster et al., 2013; Shalvi et al., 2011, 2012).

In the first treatment, participants received the same instructions but had limited time to fulfill the task. They only had 20 seconds to roll the die and report the die result. This

experimental set-up was inspired by the design chosen by Shalvi et al. (2012). The authors also enforced time pressure by limiting the time to 20 seconds. Highlighting the information about limited time to fulfill the task in the introduction ensured common knowledge among participants assigned to this treatment. After the time passed by, answering was not possible anymore, since the site automatically advanced to the next page.

In the second treatment, the participants were exposed to a different instruction. In addition to the explanation of the task, the instruction contained a linguistic cue: the phrase: "Please do not be a cheater" was added at the end of the introduction text. The phrase was highlighted so that it quickly caught the subject's attention. Exposing participants to self-relevant nouns was successfully implemented by Bryan et al. (2013).

The third and last treatment combined the elements linguistic cue and time pressure. Therefore, subjects only had 20 seconds to report the die roll and were exposed to the same phrase as in treatment 2. Both the time pressure instruction and the linguistic cue were made salient in the instructions, whereas the linguistic cue was highlighted even stronger, to make sure that all participants in that treatment group saw the moral reminder.

	No time pressure	**Time pressure**
No linguistic cue	Baseline treatment	Treatment 1
Linguistic cue with "Please do not be a cheater"	Treatment 2	Treatment 3

Table 1: *2 x 2 design of treatments*

3.3 Measurements of interest variables

3.3.1 Lying measure

The main variable of interest in this experiment is the participants' lying behavior. As the external website used to simulate the die roll does not show any results of the participants' rolls, it is impossible to detect lying on the individual level. Thus, lying can only be measured on an aggregate level. Measuring lying behavior by letting participants

anonymously roll a die and report the number rolled is, as previously mentioned, a common method to measure lying behavior (Fischbacher & Föllmi-Heusi, 2008; Foerster et al., 2013; Shalvi et al., 2011, 2012). Due to probability, the expected distribution of die results should be equal to 16.67%. Hence, if people reported honestly, the distribution should not statistically differ from 16.67% per die result. Accordingly, if for instance, the participants over-reported a five, the percentage of participants reporting a five should be significantly higher than 16.67%.

3.3.2 Time pressure measure

Another variable of interest in the presented research is time pressure. In order to detect if the time pressure manipulation is successful, I make use of the average response time of participants. If on average people take more time to report the die result when having ample time, it indicates that the enforcement of time pressure works. Measuring time manipulation by comparing the average response time was also used by Shalvi et al. (2012).

3.3.3 Linguistic cue measure

The next variable of interest is whether the exposure to the self-relevant noun ‘cheater’ is successful in the way that it decreases lying behavior. To measure the effectiveness of the subtle linguistic cue, I examine the general distribution of reported die results. Since it is expected that implementing linguistic cues would increase honesty, the probability of each die result should be closer to 16.67%, so that overall fewer numbers should be over- or under-reported.

3.4 Data collection

To collect the data, I made use of the online survey platform Qualtrics. Before going live with the final survey, I conducted a small external pilot study on the 16th and the 17th of April 2019 with 15 respondents. Only some of the respondents were familiar with the current topic in order to get the best insights on the survey's quality. Running a pilot study is important for multiple reasons. First, it enables the assessment of the survey’s feasibility. It allowed me to check the retrieved data and if the measures described in chapter 3.3 are quantifiable as mentioned. Secondly, I was able to test whether the survey

is correctly and commonly understood by the participants. Lastly, the retrieved data facilitated a check of the randomization procedure (Lancaster, Dodd, & Williamson, 2004; Thabane et al., 2010). After running the pilot study, I made some small adjustments to the theme of the survey, as it was not mobile friendly, and I adjusted the wording of the introduction and explanation of the task.

After implementing the aforementioned adjustments, I started gathering data for the final study on the 25th of April 2019. The survey subjects were recruited online without revealing the objective of the survey. I deliberately chose not revealing it upfront, to ensure that the participants can truly answer without being aware of the aim of the study. The survey was shared in Facebook groups such as “Survey Sharing 2019” or “Dissertation Survey Exchange”, on LinkedIn, as well as on websites such as “SurveyTandem”, “SurveySwap” and “SurveyCicle”. The last data point was collected on the 12th of May 2019.

4. Data and Results

The following chapter presents the results of the experiment at hand. First, it discusses the descriptive statistics of the sample before it dives into the analysis in order to test the hypotheses.

4.1 The dataset

In total, 313 respondents started the survey. This is the number of respondents who were reasonable to gather within the scope of this research. 17.3% of all subjects dropped out before completing the survey. There is no clear pattern with regard to the point where participants dropped out. Subjects who finished the survey but did not report a number are excluded from the dataset, as they cannot be used for analyzing lying behavior on the aggregate level. Furthermore, some participants have to be precluded from the analysis, as they are outliers regarding response time. Such outliers emerged because the software measured response time by counting the time until the next page was fully loaded. This measure is flawed for subjects that had a bad internet connection and resulted in response times beyond the 20 seconds enforced in the time pressure

condition. Thus, observations above 20 seconds response time are omitted from the data, as those response times would bias the results. After cleaning the data, 244 independent observations on the subject level are left for the analysis.

4.2 Descriptive Statistics

4.2.1 Demographics

Table 2 provides an overview of the demographics of the total sample as well as for the individual subsamples. As mentioned in the preceding paragraph, the total sample consists of 244 subjects. 68 subjects were assigned to the baseline treatment, in which participants were not exposed to linguistic cues and had ample time to report the die roll. 59 participants were assigned to the time pressure treatment, 61 subjects to the linguistic cue treatment and 56 to the combination of both, time pressure and the linguistic cue. As shown in the table, no variation can be found on the age dimension. The average age is similar, with 24 to 26 years among treatments. However, it can be seen that the data is skewed towards women and students. In the total sample as well as in each subsample, the share of women and students is above 50%. However, these shares vary across treatments. Notably, the skewness towards women varies from 51.82% to 71.43% among the treatments. Using the 2x6 Fisher Exact test, I checked whether this variation is balanced or if the distribution of gender is different across treatments. The result showed that gender is unbalanced among treatments at a 10% level (FE 0.087).

	Total	**No linguistic cue & ample time**	**Time pressure**	**Linguistic cue**	**Time pressure with linguistic cue**
N	244	68	59	61	56
Gender	Female 63.52% Male 36.48%	Female 69.12% Male 30.88%	Female 62.71% Male 37.29%	Female 50.82% Male 49.18%	Female 71.43% Male 28.57%
Age (in years)	Min 17 Max 78 Ø 24.97	Min 18 Max 78 Ø 26.11	Min 18 Max 52 Ø 24.47	Min 18 Max 56 Ø 24.64	Min 17 Max 35 Ø 24.46
Student share	87.70%	86.76%	88.14%	83.61%	92.86%

Table 2: *Demographics of total sample and subsamples*

4.2.2 Outcome variable lying

The variable of interest for this research is whether the subjects on the aggregate level lied. Before diving into the analysis, however, I checked the distribution of the overall sample to see whether subjects acted dishonestly in general. Assuming total honesty, the distribution for each number should be equal to 16.67%, as explained in chapter 3.3.1. Figure 1 depicts the distribution accordingly. It shows that dishonesty is present in the data. A pattern of over-reporting can be observed for the numbers above three and numbers below four are under-reported. Particularly, the number five is notably over-reported and the number two under-reported.

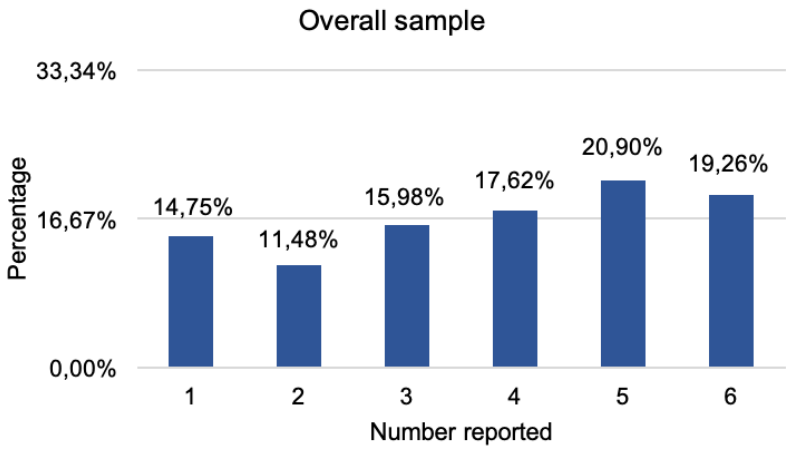

Figure 1: *Distribution of reported numbers in the overall sample*

However, several binomial tests showed that the differences are not statistically different from 16.67%[4].The higher tendency to report a five instead of a six might be explained by individuals' tendency to lie only partially. Subjects wanted to have a higher chance of winning the prize but still had the desire to uphold the image of being an honest person. Therefore, subjects did lie but did not maximize the own monetary payoff either by reporting a five rather than a six[5]. Despite the deviations from the hypothesized distribution, the observed distribution is not statistically different, as already indicated by

[4] When running multiple tests after another, just by chance, one can be significant. Hence, I applied the Bonferroni correction as I was running six binomial tests. Using the conventional p-value of 0.05, a corrected p-value of 0.05/6=0.0083 was left for the analysis.

[5] Taking into account that there were also participants rolling a five and reporting truthfully. This argumentation is in line with Fischbacher & Föllmi-Heusi (2008).

the binomial tests. Applying the Chi-squared test shows no significant difference in expected and observed frequencies (χ^2= 3.922, p=0.561).

An additional notion to check is whether the non-students add noise to the data so that they would bias the results. As almost 90% of the sample are students, the non-students could be left out for analysis in case of differences in results. Blocking for students would decrease the error term's variance as it increases the subject pool's homogeneity and therefore would elevate the power of the tests. However, the 2x6 Fisher Exact test shows no statistically significant differences in the distribution of reported numbers between the student sample and the whole sample (FE 0.901). Thus, the full sample is used for the in-depth analyses.

Generally, to analyze experimental data researchers make use of non-parametric tests, as the assumption of normality usually is violated. Even though non-parametric tests do not harness the full richness of the data, normality is not required. The results of the Kernel density function, as well as the result of the Shapiro-Wilk-test (see Appendix B1), confirm that the data violates the normality assumption and advocates the utilization of non-parametric tests[6].

4.3 Analyses

4.3.1 Time pressure manipulation

As described in chapter 3.3.2, the time pressure manipulation is tested by comparing the average response times in the different treatments. As expected, the average response time is lower in the time pressure treatment (M=11.026s, SD=4.353s) as compared to the baseline treatment (M=14.726s, SD=10.377s). Similarly, the average response time in the linguistic cue treatment (M=13.963s, SD=10.096s) is higher than in the time pressure with linguistic cue treatment (M=11.530s, SD=3.598s). Notably, both treatments with no time pressure have a huge standard deviation implicating a high variability of the response time. To now distinguish whether a parametric or non-parametric test can be applied for testing the differences, I make use of the Kernel density function as well as

[6] Most important assumption of independent observations necessary for performing non-parametric tests is satisfied as mentioned in section 4.1

the Shapiro-Wilk-test. The results clearly show a violation of normality (see Appendix B1), so that the non-parametric Mann-Whitney-U test is applied. The disparity in average response time for the baseline treatment and the time pressure treatment is statistically different at a 10% level (z=1.682, p=0.093), indicating that the time pressure manipulation was successful. However, for the treatments with linguistic cues, the average response times with and without time pressure constraints are not significantly different (z=0.900, p=0.368).

4.3.2 Lying under time pressure

To investigate if people under time pressure lied more in comparison to people having ample time, I compare the baseline treatment to the time pressure treatment. The distributions of the reported numbers are depicted in figure 2. As shown in this figure, the distributions differ slightly. Numbers five and six are over-reported in both treatments, and subjects reported the numbers one and two less than expected by chance. Again, as in the overall sample, the over-reporting of a five over a six is seen in both treatments, signifying the presence of partial liars in both treatments. However, subjects in the time pressure treatment had an even higher tendency to report higher numbers such as five and six and higher tendency to report fewer numbers such as two and three. Furthermore, subjects in the time pressure treatment also reported number one more often than people in the baseline treatment. This might be indicative that participants in the time pressure treatment tended to lie more for their own benefit than in the baseline treatment, albeit there is only a slim difference in means. On average, the subjects in the baseline treatment reported a 3.838[7] (SD=1.626) and subjects in the time pressure treatment reported a 3.932 (SD=1.711). In general, it can be said that the higher the mean, the more people were inclined to lie for their own benefit by reporting a higher number and vice versa. Full honesty would result in a mean of 3.5.

[7] Interval scale of the outcome variable was assumed. The difference between all numbers was the same enabling the computation of the mean.

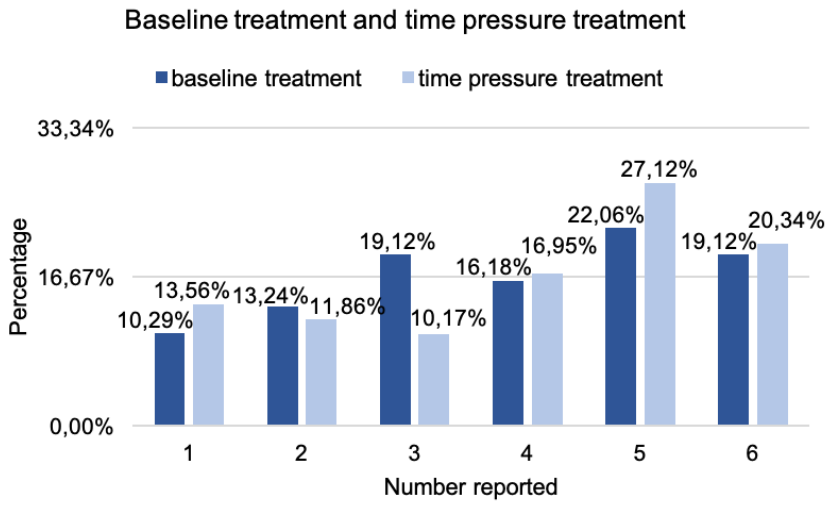

***Figure 2**: Distribution of reported numbers in the baseline treatment and time pressure treatment*

To investigate the first hypothesis of this research, I compare the difference in the distributions of the reported numbers among the treatments. To do so, a Mann-Whitney-U test is used. The results show that there is no significant difference in lying behavior between the two treatments (z=-0.396, p=0.692) (see Appendix B2). Yet, as one of the main drawbacks of using that test is adjusting for ties, and the result clearly shows adjustments for ties, I further used the 2x6 Fisher Exact test. In line with the result of the Mann-Whitney-U test, the Fisher Exact test confirms that there exists no difference in the frequency distribution between the treatments (FE 0.798). This leads me to say that there is no evidence in support of hypothesis 1.

4.3.3 Lying with linguistic cues

To answer the second hypothesis that adding the linguistic cue 'Please do not be a cheater' in the introduction decreases lying behavior, I compare lying behavior in the baseline treatment with lying behavior in the linguistic cue treatment. Thus, the linguistic cue is the only difference between the treatments. Before diving into the non-parametric tests applied in this case, I looked at the distribution of reported numbers in the two treatments. As seen in figure 3, the distributions differ slightly. Subjects exposed to the linguistic cue reported five and six less frequently and one and four more frequently than

subjects in the baseline treatment, signaling a potential direction regarding lying as hypothesized.

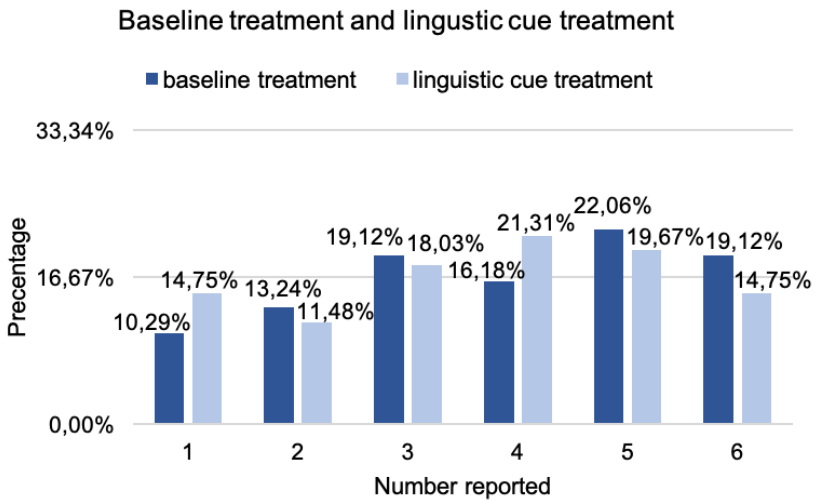

Figure 3: *Distribution of reported numbers in the baseline treatment and linguistic cue treatment*

Taking a further look at the means per treatment, the hypothesized suggested lying behavior seems to be satisfied: The mean of the baseline treatment is 3.838 (SD= 1.626), the mean of the linguistic cue treatment is 3.639 (SD= 1.633). To compare the differences in means, I again use the Mann-Whitney-U test. A resulting p-value of 0.495 (z=0.683) does not lend support for the second hypothesis (see Appendix B2). Participants in the linguistic cue treatment did not lie significantly less than in the baseline treatment. To further strengthen this statement and to overcome the adjustments for ties, the 2x6 Fisher Exact test is applied. Nevertheless, the result of the test affirms that there is no evidence to support hypothesis 2.

4.3.4 Lying under time pressure with linguistic cues

Unlike in the investigation of the first two hypotheses, for the examination of the third hypothesis, the time pressure treatment serves as the baseline. To draw sufficient conclusions whether the linguistic cue influenced lying under time pressure, it is important that only one factor is changed – the addition of the phrase. Thus, the time pressure treatment and the time pressure with linguistic cue treatment are used in order to infer

results for the third hypothesis. Figure 4 depicts the distributions of numbers for the two treatments.

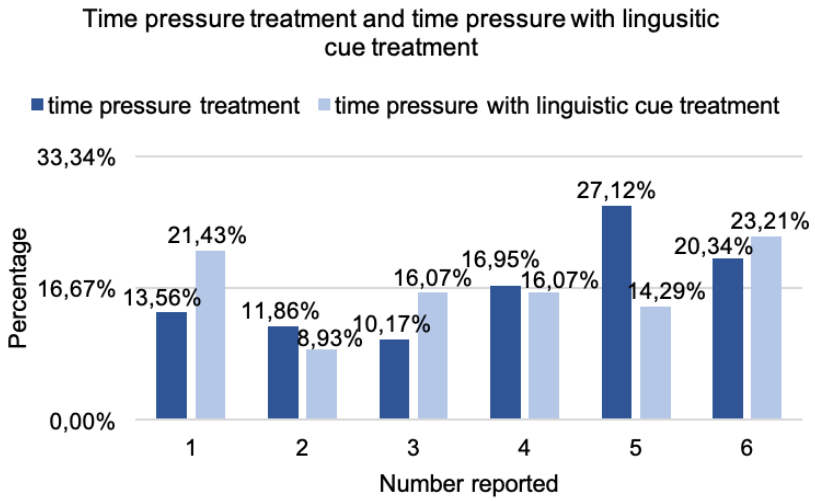

Figure 4: *Distribution of reported numbers in the time pressure treatment and time pressure with linguistic cue treatment*

Interestingly, there exists a remarkable difference in the frequency the number five was reported. There are numerous partial liars in the time pressure treatment, whereas in the time pressure with the linguistic cue treatment there are not. An opposing discrepancy is seen for the number one: 21.43% of the subjects in the time pressure with the linguistic cue treatment reported a one compared to only 13.56% in the time pressure treatment. Additionally, the number three is fairly over-reported in the time pressure treatment with the linguistic cue compared to the time pressure treatment. In all, the distribution of reported numbers seems to be somewhat different, indicating that the linguistic cues influenced lying behavior under time pressure.

To test whether the addition of the phrase altered behavior in the postulated way, I make use of the Mann-Whitney-U test to compare the means of the two treatments. As mentioned in chapter 4.3.2 the mean of the time pressure treatment is 3.932 (SD=1.711). The mean for the time pressure with the linguistic cue treatment is 3.625 (SD=1.854). The test displays that there is no difference in distribution (z=-0.850, p=0.395) (see Appendix B2). Including the adjustments for ties, the 2x6 Fisher Exact test confirms these findings (FE 0.499). Accordingly, there is no statistical support for hypothesis 3.

4.3.5 Robustness of results

Even though I do not find evidence in support of any hypothesis in the preceding paragraphs, I use linear regression to check for the robustness of the results. The results of the regression analysis are reported in table 3. In order to infer any conclusions from the non-parametric tests, it is vital that solely the treatments influenced the decision of reporting a number and not any other confounding variable. When looking at the restricted and unrestricted model, one can observe a slight difference. The coefficients for the different treatments change in magnitude when control variables are introduced. Particularly, the coefficient of being in the time pressure treatment compared to being in the baseline treatment changes the sign and hence, the direction of the effect. For the other two remaining coefficients, the sign stays the same, only the magnitude changes slightly. This might be caused by the omitted controls or multicollinearity. However, after computing the variance inflation factors to assess multicollinearity, one can rule out multicollinearity as the standard errors are not inflated upwards (see Appendix B3). Omitted variable bias cannot be precluded. Yet, omitted variable bias is not caused by the controls, as adding those variables do not yield a great change in magnitude or significance. This further supports the utilization of non-parametric tests throughout the preceding paragraphs as those tests do not account for controls. Overall, as all coefficients in both the unrestricted and restricted model are highly insignificant and all explanatory variables are jointly insignificant, a change in magnitude and sign can also be arbitrary.

It needs to be acknowledged that the results of the linear regression are just approximations. First, the dependent variable with 6 different outcomes may not be considered as a continuous variable. A continuous scale is a necessary assumption for OLS regressions. Secondly, a histogram of the residuals shows (see Appendix B4) that the residuals are not normally distributed. As the normal distribution of the error terms as well as the continuous scale are necessary assumptions of OLS, the inferences made are treated with caution.

***Table 3**: Results of the linear regression predicting the number reported*

	Number reported	
	(I)	(II)
Time pressure treatment	0.094	-0.020
	(0.298)	(0.304)
Linguistic cue treatment	-0.199	-0.278
	(0.288)	(0.293)
Time pressure & linguistic cue	-0.213	-0.294
	(0.317)	(0.337)
Age		-0.005
		(0.021)
Female		-0.368
		(0.242)
Student		-0.089
		(0.422)
Income		-0.383
		(0.244)
East-Europe		-0.245
		(0.350)
Northern America		0.292
		(0.413)
Asia		-0.488
		(0.511)
Middle East		-0.506
		(0.778)
Australia		0.296
		(0.710)
Africa		0.225
		(1.481)
Central & South America		-0.407
		(0.890)
Constant	3.838***	4.610***
	(0.197)	(0.870)
Observations	244	244
R2	0.006	0.037

Note: For each variable, coefficients and the heteroskedasticity-robust standard error in parentheses are reported. Column (I) shows the restricted model, with only the treatment variables as independent variables and the baseline treatment as the reference category. Column (II) shows the unrestricted model, the effect after controlling for age (measured in years), gender, student, income (assuming that all non-students have income, dichotomous variable) and country of origin, with West-Europe as the reference category. (* : $p < 0.1$, ** : $p < 0.05$, *** : $p < 0.01$).

4.4 Exploratory analysis

As seen in the preceding chapter, the results of the linear regression showed some sensitivity when the control variables were added. This induced me to go beyond the formal theory and hypothesis testing to take a closer look at the data. Parametric tests might only give approximate results as the assumption of normality is violated. Still, as such tests better harness the richness of the data, they are used for this ex-post exploratory analysis to get a better understanding of the data.

Firstly, I use the unpaired t-test to test for possible differences in reporting behavior between females and males. Research has shown that men have a higher tendency to lie than women (DePaulo et al., 1996; Dreber & Johannesson, 2008). Interestingly, this result is also shown in the current data. Women on average reported significantly lower numbers than men on a 10% level (μ_{Female}=3.632, μ_{Male}=3.989, SD_{Female}=1.671, SD_{Male}=1.729, Pr(T<t)=0.057). Thus, women lied less than men. These results prompted me to analyze if women react differently to the treatments than men. However, after running multiple unpaired t-tests and applying the Bonferroni-correction accordingly, the results show that women did not react significantly differently than men to any manipulation.

Next, I analyze if there is a difference in reporting behavior and the country of origin. Different cultural virtues or different religious beliefs might cause a potential discrepancy in lying behavior. As Lytle and Rivers (2007) found, culture virtues and beliefs can influence ethical behavior. Using the one-way ANOVA enables me to test whether there is a difference in average reporting behavior between different country categories. The results indicate that there is no systematic pattern in lying behavior among country categories (F=0.35, p= 0.932). Therefore, cultural values do not seem to influence lying behavior.

Thirdly, it is interesting to understand whether students who are financially better off reacted differently to the treatments. It might be, for instance, that students with some income besides their study loan do not feel the same necessity to earn extra money than students without income, and therefore lie less. Hence, I make use of the unpaired t-test.

Students with no income lied significantly more on a 10% level than students with income ($\mu_{Student\text{-}income}$= 3.602, $\mu_{Student\text{-}noincome}$= 3.941, $SD_{Student\text{-}income}$ = 1.709, $SD_{Student\text{-}noincome}$= 1.660, Pr(T<t)=0.072). Remarkably, in the current data set there is a difference in lying behavior and income, even though there is no clear evidence in the literature saying that higher income leads to a lower level of lying when being exposed to a financial opportunity to lie (Abeler, Becker, & Falk, 2014; Rosenbaum et al., 2014). When testing if these students with income reacted differently to the treatment manipulations, it becomes clear that there is no distinction in reporting behavior. After using the Bonferroni-correction, all unpaired t-tests are insignificant. Hence, students without income did not react differently to the treatments but lied more than students with income.

Lastly, after categorizing age in three distinct age groups, the one-way ANOVA is applied. This test indicates whether there is a difference in lying behavior between the different age categories. As age is one main factor that can influence lying behavior, it is interesting to see if this also applies to the current dataset. The literature has shown that older people tend to lie less than younger ones (Conrads, Irlenbusch, Rilke, & Walkowitz, 2013; Ross & Robertson, 2010). Yet, the result of the ANOVA displays that lying behavior among groups is the same (F=1.40, p= 0.249). Thus, age did not drive lying behavior in this study.

Remarkably, this exploratory analysis shows that women lied less than men and that students without income lied more than students with income. However, these findings remain proxies of the real effect as parametric tests were applied.

4.5 Power

A post-hoc power analysis is conducted in order to obtain the power of the performed analysis[8]. Additionally, as all null hypotheses are accepted, the risk of a type 2 error is present. The type 2 error can be defined as the false acceptance of the null hypothesis. To approximate the power calculations, I make use of a rule of thumb and assume that all analyses are based on a student t-test. As the tool allows differentiating between two treatments only, I perform three power calculations: Once I distinguish between time

[8] The power analysis is conducted by using the G*power tool.

pressure and no time pressure, once between linguistic cue and no linguistic cue and once between being exposed to linguistic cues with and without time pressure. This approach is chosen since such comparisons of treatments are used in order to answer the hypotheses of the research at hand. For all post-hoc analyses, the power is remarkably lower than the conventional level of power of 0.8. The maximum power achieved is 0.15, comparing the time pressure treatment with and without the linguistic cue. (All results of the calculations can be found in Appendix B5). Acknowledging that these results only give proxies of the true power, it is indicative that the study is highly underpowered and hence the risk of the type 2 error is present.

5. Discussion

5.1 General discussion

This study fails to find support for increased lying behavior under time pressure. The null hypothesis that there is no difference in reporting under time pressure compared to having ample time could not be rejected. In this study, participants under time pressure did not lie more compared to subjects with ample time. Furthermore, subjects under time pressure also did not lie less than those with sufficient time. The current study does not find any difference in lying behavior regarding time pressure, as the means were not significantly different. This result might be in line with the discrepancy of research trying to understand lying under time pressure. As there are divergent opinions on whether people tend to lie more or less under time pressure, this study indicates that there may be no clear answer to the question of whether people lie more or less under lack of time. Rather, it may be influenced by other factors than just time pressure. As pointed out in chapter 2.1, the process of lying is complex and shaped by multiple factors. This, in fact, also applies to lying under time pressure. Thus, the assessment of lying under time pressure is rather sophisticated and may not be answered on the aggregate level.

Likewise, this study does not find any evidence in support of the hypothesis that people exposed to the linguistic cue lied less than people not exposed to it. Nonetheless, the lower mean of the linguistic cue treatment might be indicative that the linguistic cue can work as hypothesized. In line with the findings by Mazar et al. (2008) and Pruckner and

Sausgruber (2013), an explicit moral reminder can mitigate the tendency to lie. Even an implicit reminder of morality, such as looking in the mirror when making a decision or signing a sheet before filling it in, proved to decrease levels of dishonesty (Gino & Mogilner, 2014; Shu, Mazar, Gino, Ariely, & Bazerman, 2012; Vincent, Emich, & Goncalo, 2013). Maybe one might think that confronting participants with the word 'cheater' might induce the opposite (an increased level of lying) by reminding them that lying is possible. Then this might explain the non-significance of the results. Yet, and to the best of my knowledge, there is no evidence in the literature showing that a moral reminder, regardless whether it is implicit or explicit, evoked the opposite effect. Therefore, I note that the lack of evidence is due to the underpowered study and might well work when replicating the study with more resources.

A combination of both linguistic cue and time pressure was tested with hypothesis 3. In line with the results of the previously discussed hypotheses, there was no evidence that the linguistic cue under time pressure alleviates lying behavior. Such non-significance might be attributed to the failure of the time pressure manipulation. As assessed in chapter 4.3.1, the time pressure manipulation did not work for the linguistic cue with time pressure treatment as compared to the treatment with only the exposure to the linguistic cue. As it did work in the case of time pressure alone, it might be that the lack of significance was only due to too few participants. This point of view is further underpinned by looking at the means since the treatment with linguistic cue under time pressure is about two seconds lower than in the linguistic cue treatment alone. Besides the flawed time pressure manipulation, the lack of evidence for hypothesis 3 might also stem from the fact that the linguistic cue is not strong enough to mitigate the effect of time pressure. Time pressure may have been a stronger factor and may have dominated the linguistic cue. As this was the first time a combination of both was tested, I can only speculate why I did not find evidence in support of the hypothesis. However, the direction of means clearly indicate that lying under time pressure with exposure to the linguistic cue seems to be less than without the linguistic cue. Therefore, I assume that the lack of statistical difference is rather due to the little power of the study and not primarily because the two concepts do not work in combination.

Furthermore, I made use of linear regression to check for the robustness of the results. As randomization suggests, such controls did not affect the reported number significantly. Although the coefficients were all not significant, the coefficients in the restricted model changes slightly when the controls were added. This result, in line with the findings of the exploratory analysis, emphasizes the importance of further research on the effect of specific characteristics on lying behavior. Gender might be of interest, as lying behavior seems to be different between men and women (DePaulo et al., 1996; Dreber & Johannesson, 2008; Erat & Gneezy, 2012). Additionally, in light of the exploratory analysis, income might be interesting to investigate in more detail. Yet, as this study was the first to combine linguistic cues and time pressure, it was of importance to investigate a general potential tool to prevent lying under time pressure. Further research should also focus on potential differences among characteristics such as gender, age and income.

5.2 Limitations

One main limitation of this research is the only partial compliance of the five precepts of Smith (1982). As elaborated in chapter 3.1.2, the precepts were not entirely fulfilled. Particularly, the saliency precept poses a problem. This precept is needed to achieve a valid and controlled experiment. If a study fails to satisfy the sufficient conditions, the participants might react randomly to the incentives. Furthermore, the dominance precept might have been violated due to the usage of a binary lottery incentive. Hence, future research should use more resources to pay out all participants according to their reported number so that the saliency precept is satisfied. However, in light of the given resources, I note that applying the lottery was enough to satisfy the dominance precept. Still, future research should carefully take all the precepts into account.

Additionally, another limitation of the current study is that it is notably under powered. Due to the lack of statistical power, the risk of the type 2 error is high. This indeed poses a problem, as all null hypotheses could not be rejected. The study might have well worked and significant differences could have been found if the power was higher. Yet, this remains pure speculation. For further research in this field of study, this should clearly be considered. One possible way to increase power is to reduce the error variance, for instance with blocking on characteristics such as gender or using a within-subject design.

Still, as explained in the preceding paragraph, this research was focused on finding a general way to promote honesty under time pressure, so that blocking was no option. Besides the increased power, there are more perks of using a within-subject design (see chapter 3.1). However, possible learning confounds dominated all benefits of using within-subject design. Furthermore, an elevated level of power can be achieved by increasing the treatment level variance. If a linear relationship is assumed, this can be accomplished by only using two treatment levels with the greatest spread possible. Lastly, power can be increased by using more observations. As the a-priori power calculation shows, approximately 4954 participants would have been needed in each treatment to increase the power to the conventional level of 0.8 (Please find the calculation in Appendix B6). However, within the scope of this research, this number of observations was not feasible to achieve.

Another drawback of this study is that lying could only be measured on the aggregate level. As already mentioned in chapter 5.1, this study does not find evidence of lying or truth-telling being the dominant strategy under time pressure. As the literature is also still debating on this particular topic, it might be of interest to measure lying with a lack of time on the individual level. Lying as a noteworthy complex process might be more insightful to see per individual, which to the best of my knowledge has never been tested. Potentially, a natural field experiment might constitute an option, as participants then do not know that they are being watched. This might yield valuable insights as the subjects then behave naturally within the scope of such a sensitive topic. However, lying measured on the individual level might not be suitable to draw any externally valid conclusions and hence questions the concept of conducting research on the individual level. Furthermore, as high anonymity remains a crucial topic so that people are not compromised in the way they behave during the experiment, dishonesty measured on the aggregate level stays a valuable and useful tool to apply within the scope of this research (Batson, Kobrynowicz, Dinnerstein, Kampf, & Wilson, 1997; Shalvi et al., 2012).

Moreover, the sample used for this research is not representative of the population as almost 90% are students and the share of females is above 60%. Trying to spread the

survey beyond my personal network did not work sufficiently to increase heterogeneity among participants.

Acknowledging these above-mentioned limitations, it needs to be said that they clearly are a drawback and should be taken into consideration when conducting further research within this field of study. Still, given the timeframe and limited resources, this study provides some notable indications about dishonesty under time pressure and how to potentially promote honesty.

6. Conclusion

Sometimes, lying might just pave the way for a higher benefit, particularly when there is just a short time window to make a decision. And who does not like having some extra money or recognition? Exactly this inspired me to conduct the present research as people are not aware of the consequences even a small lie might entail. Accordingly, the present paper tested lying under time pressure and how to possibly prevent it.

Scientific papers cannot agree on whether lying or truth-telling prevails under decision-making within a short time frame. Therefore, I tested this by having participants self-report their results with and without time constraints. The results indicate that there is no difference in the level of dishonesty under time pressure as compared to having unlimited time. In line with the prevalent discrepancy within the scientific literature, the results of this study might display that it is not a simple question of what the default option is. Rather, it might be of importance to look at specific decisions individually and not at the aggregate level.

Furthermore, as tools to prevent lying with lack of time have never been tested before, this study fills this void in the literature about lying. In doing so, it bridges two concepts to find a successful tool to prevent dishonesty under lack of time. It exposes people to a moral reminder when facing a decision under time pressure. The results only indicate that exposing people to the phrase ‘Please do not be a cheater’ might promote honesty under time pressure. Because it is unclear of whether the non-significance is due to the lack of statistical power or simply because such linguistic cues do not work under time pressure,

I do not find substantial evidence that moral reminders help to promote honesty under time pressure.

This study may serve as the foundation for future research trying to find a way to decrease lying behavior under time pressure. Further studies may exceed the methodology of this paper by considering all the aforementioned drawbacks and limitations. Still, it remains highly interesting and important to understand how to make people stick with the truth under time pressure.

References

Abeler, J., Becker, A., & Falk, A. (2014). Representative evidence on lying costs. *Journal of Public Economics*, *113*, 96–104.

Ariely, D. (2008). *Predictably irrational*. New York: HarperCollins Publishers.

Ariely, D. (2012). *The (honest) truth about dishonesty: How do we lie to everyone - especially ourselves*. HarperCollins Publishers.

Ariely, D., Gneezy, U., Loewenstein, G., & Mazar, N. (2009). Large Stakes and Big Mistakes. *The Review of Economic Studies*, *76*(2), 451–469.

Batson, C. D., Kobrynowicz, D., Dinnerstein, J. L., Kampf, H. C., & Wilson, A. D. (1997). In a very different voice: Unmasking moral hypocrisy. *Journal of Personality and Social Psychology*, *72*(6), 1335–1348.

Becker, G. S. (1968). Crime and Punishment: An Economic Approach. *The Economic Dimensions of Crime*, 13–68.

Bereby-Meyer, Y., & Shalvi, S. (2015). Deliberate honesty. *Current Opinion in Psychology*, *6*, 195–198.

Bryan, C. J., Walton, G. M., Rogers, T., & Dweck, C. S. (2011). Motivating voter turnout by invoking the self. *Proceedings of the National Academy of Sciences*, *108*(31), 12653–12656.

Bryan, Christopher J., Adams, G. S., & Monin, B. (2013). When cheating would make you a cheater: Implicating the self prevents unethical behavior. *Journal of Experimental Psychology: General*, *142*(4), 1001–1005.

Camden, C., Motley, M. T., & Wilson, A. (1984). White lies in interpersonal communication: A taxonomy and preliminary investigation of social motivations. *Western Journal of Speech Communication*, *48*(4), 309–325.

Cappelen, A. W., Sørensen, E. T., & Tungodden, B. (2013). When do we lie? *Journal of Economic Behavior and Organization*, *93*, 258–265.

Capraro, V. (2017). Does the truth come naturally? Time pressure increases honesty in one-shot deception games. *Economics Letters*, *158*, 54–57.

Carson, T. L. (2006). The Definition of Lying*. *Nous*, *40*(2), 284–306.

Charness, G., Gneezy, U., & Kuhn, M. A. (2012). Experimental methods: Between-subject and within-subject design. *Journal of Economic Behavior and Organization*, *81*(1), 1–8.

Conrads, J., Irlenbusch, B., Rilke, R. M., & Walkowitz, G. (2013). Lying and team incentives. *Journal of Economic Psychology*, *34*, 1–7.

DePaulo, B. M., Kirkendol, S. E., Kashy, D. A., Wyer, M. M., & Epstein, J. A. (1996). Lying in Everyday Life. *Journal of Personality and Social Psychology*, *70*(5), 979–995.

Dreber, A., & Johannesson, M. (2008). Gender differences in deception. *Economics Letters*, *99*(1), 197–199.

Erat, S., & Gneezy, U. (2012). White Lies. *Management Science*, *58*(4), 723–733.

Fallis, D. (2009). What Is Lying? *The Journal of Philosophy*, *106*(1), 29–56.

Fischbacher, U., & Föllmi-Heusi, F. (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association*, *11*(3), 525-547.

Foerster, A., Pfister, R., Schmidts, C., Dignath, D., & Kunde, W. (2013). Honesty saves time (and justifications). *Frontiers in Psychology*, *4*, 1–2.

Gelman, S. A., & Heyman, G. D. (1999). The Effects of Lexicalization on Children ' s Inferences About. *Psychological Science*, *10*(6), 489–493.

Gino, F. (2015). Understanding ordinary unethical behavior: Why people who value morality act immorally. *Current Opinion in Behavioral Sciences*, *3*, 107–111.

Gino, F., Ayal, S., & Ariely, D. (2009). Contagion and Differentiation in Unethical Behavior. *Psychological Science*, *20*(3), 393–398.

Gino, F., & Mogilner, C. (2014). Time, Money, and Morality. *Psychological Science*, *25*(2), 414–421.

Gunia, B. C., Wang, L., Huang, L. I., Wang, J., & Murnighan, J. K. (2012). Contemplation and conversation: Subtle influences on moral decision making. *Academy of Management Journal*, *55*(1), 13-33.

Heckman, J. J., & Smith, J. A. (2005). Assessing the Case for Social Experiments. *Journal of Economic Perspectives*, *9*(2), 85–110.

Jacobsen, C., Fosgaard, T. R., & Pascual-Ezama, D. (2018). Why Do We Lie? a Practical Guide To the Dishonesty Literature. *Journal of Economic Surveys*, *32*(2), 357–387.

Kahneman. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Kajackaite, A., & Gneezy, U. (2017). Incentives and cheating. *Games and Economic Behavior*, *102*, 433–444.

Kant, I. (2012). *Fundamental principles of the metaphysics of morals*. (T. K. Abbott, Ed.). Courier Corporation.

Kouchaki, M., & Desai, S. D. (2015). Anxious, threatened, and also unethical: How anxiety makes individuals feel threatened and commit unethical acts. *Journal of Applied Psychology*, *100*(2), 360–375.

Kouchaki, M., & Wareham, J. (2015). Excluded and behaving unethically: Social exclusion, physiological responses, and unethical behavior. *Journal of Applied Psychology*, *100*(2), 547–556.

Lancaster, G., Dodd, S., & Williamson, P. (2004). Design and analysis of pilot studies: recommendations for good practice. *Journal of Evaluation in Clinical Practice*, *10*(2), 307–312.

Levine, T. R. (2014). Truth-Default Theory (TDT). *Journal of Language and Social Psychology*, *33*(4), 378–392.

Lytle, A. L., & Rivers, C. (2007). Lying, cheating foreigners!! Negotiation ethics across cultures. *International Negotiation*, *12*(1), 1-28.

Mazar, N., Amir, O., & Ariely, D. (2008). The Dishonesty of Honest People: A Theory of Self-Concept Maintenance. *Journal of Marketing Research*, *45*(6), 633–644.

McCornack, S. A., Morrison, K., Paik, J. E., Wisner, A. M., & Zhu, X. (2014). Information Manipulation Theory 2: A Propositional Theory of Deceptive Discourse Production. *Journal of Language and Social Psychology*, *33*(4), 348–377.

Minson, J. A., VanEpps, E. M., Yip, J. A., & Schweitzer, M. E. (2018). Eliciting the truth, the whole truth, and nothing but the truth: The effect of question phrasing on deception. *Organizational Behavior and Human Decision Processes*, *147*, 76–93.

Oxford English Dictionary. (2019). Oxford English Dictionary. Retrieved March 21, 2019, from http://www.oed.com/view/Entry/108036?rskey=RI5FgH&result=1&isAdvanced=false#eid

Perkins, S. A., & Turiel, E. (2007). To Lie or Not to Lie: To Whom and Under What Circumstances. *Child Development*, *78*(2), 609–621.

Pruckner, G. J., & Sausgruber, R. (2013). Honesty on the streets: A field study on newspaper purchasing. *Journal of the European Economic Association*, *11*(3), 661–679.

Rosenbaum, S. M., Billinger, S., & Stieglitz, N. (2014). Let's be honest: A review of experimental evidence of honesty and truth-telling. *Journal of Economic Psychology*, *45*, 181–196.

Ross, W., & Robertson, D. (2010). Lying : The Impact of Decision Robertson, *10*(2), 409–440.

Schweitzer, M. E., Ordóñez, L., & Douma, B. (2004). Goal setting as a motivator of unethical behavior. *Academy of Management Journal*, *47*(3), 422-432.

Shalvi, S., Dana, J., Handgraaf, M. J. J., & De Dreu, C. K. W. (2011). Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior. *Organizational Behavior and Human Decision Processes*, *115*(2), 181–190.

Shalvi, S., Eldar, O., & Bereby-Meyer, Y. (2012). Honesty Requires Time (and Lack of Justifications). *Psychological Science*, *23*(10), 1264–1270.

Shalvi, S., Eldar, O., & Bereby-Meyer, Y. (2013). Honesty requires time—a reply to Foerster et al. (2013). *Frontiers in Psychology*, *4*, 1–2.

Shu, L. L., Mazar, N., Gino, F., Ariely, D., & Bazerman, M. H. (2012). Signing at the beginning makes ethics salient and decreases dishonest self-reports in comparison to signing at the end. *Proceedings of the National Academy of Sciences*, *109*(38), 15197–15200.

Smith, V. L. (1982). Microeconomic systems as an experimental science. *The American Economic Review*, *72*(5), 923–955.

Spence, S. A., Farrow, T. F., Herford, A. E., Wilkinson, I. D., Zheng, Y., & Woodruff, P. W. (2001). Behavioural and functional anatomical correlates of deception in humans. *Neuroreport*, *12*(13), 2849–2853.

Thabane, L., Ma, J., Chu, R., Cheng, J., Ismaila, A., Rios, L., … Goldsmith, C. (2010). A tutorial on pilot studies: the what, why and how. *BMC Medical Research Methodology*, *10*(1), 1–10.

Van der Zee, S., Anderson, R., & Poppe, R. (2016). When lying feels the right thing to do. *Frontiers in Psychology*, *7*, 734.

Verschuere, B., & Shalvi, S. (2014). The Truth Comes Naturally! Does is? *Journal of Language and Social Psychology*, *33*(4), 417–423.

Verschuere, B., Spruyt, A., Meijer, E. H., & Otgaar, H. (2011). The ease of lying. *Consciousness and Cognition*, *20*(3), 908–911.

Vincent, L. C., Emich, K. J., & Goncalo, J. A. (2013). Stretching the Moral Gray Zone: Positive Affect, Moral Disengagement, and Dishonesty. *Psychological Science*, *24*(4), 595–599.

Walton, G. M., & Banaji, M. R. (2004). Being What You Say: The Effect of Essentialist Linguistic Labels on Preferences. *Social Cognition*, *22*(2), 193–213.

Appendices

Appendix A. The Survey

Appendix A1. The survey as text

The survey was only available in English. Introduction text part 1 and 2 were shown to every participant.

Introduction text – part 1

Welcome!
Thank you for deciding to participate in this small experiment as part of my master thesis at Erasmus University Rotterdam. Filling in this survey will approximately take 2 minutes. I appreciate your participation. All information will be treated confidentially.

By completing this survey, you can enter a lottery to win money. I will randomly select 3 participants who can win a prize of 25€, 15€ or 10€.

Whoever I draw first, wins 25€, whoever I draw second, wins 15€ and whoever I draw third wins 10€.

This survey consists of two parts:

1. A die roll (The rolling of the die will be done through an external website, thus, I am not able to trace the result of your die roll.)

2. Some general questions

If you have any questions, email to 483107lw@student.eur.nl

Thank you for participating! Feel free to forward the survey to your friends.

Best,
Laura

Introduction text – part 2

The result of the die roll influences your probability of winning the prize as it determines how many tickets you can get for the lottery.

The ticket structure is as follows:

If you rolled 1, you get 1 lottery ticket

If you rolled 2, you get 2 lottery tickets

If you rolled 3, you get 3 lottery tickets

If you rolled 4, you get 4 lottery tickets

If you rolled 5, you get 5 lottery tickets

If you rolled 6, you get 6 lottery tickets

Baseline treatment

On the next page, you are asked to roll a die and report the number you rolled. You can roll the die as often as you like, to see that the die is legitimate. Please remember the figure of your first roll and fill in this number.

Remember that the number you report determines how many lottery tickets you get.

If you do not have a die close by, you can click on a link which leads you to an external website. On this website, a die roll is simulated. The result of your die roll will not be traced in any way.

Q Click here to roll the die

Q What number did you roll? _____

Treatment 1 – Time pressure

On the next page, you are asked to roll a die and report the number you rolled. You can roll the die as often as you like, to see that the die is legitimate. Please remember the figure of your first roll and fill in this number.

Remember that the number you report determines how many lottery tickets you get.

Attention! You have 20 seconds to do so. A countdown displays how much time you have left. After 20 seconds the page will automatically advance to the next page and no further entry will be possible.

If you do not have a die close by, you can click on a link which leads you to an external website. On this website, a die roll is simulated. The result of your die roll will not be traced in any way.

The time continues to count down while you roll the die on the external website this website, a die roll is simulated. The result of your die roll will not be traced in any way.

Q Click here to roll the die

Q What number did you roll? _____

Treatment 2 – Linguistic cue

On the next page, you are asked to roll a die and report the number you rolled. You can roll the die as often as you like, to see that the die is legitimate. Please remember the figure of your first roll and fill in this number.

Remember that the number you report determines how many lottery tickets you get.

If you do not have a die close by, you can click on a link which leads you to an external website. On this website, a die roll is simulated. The result of your die roll will not be traced in any way.

Please do not be a cheater.

Q Click here to roll the die

Q What number did you roll? _____

Treatment 3 – Time pressure with linguistic cue

On the next page, you are asked to roll a die and report the number you rolled. You can roll the die as often as you like, to see that the die is legitimate. Please remember the figure of your first roll and fill in this number.

Remember that the number you report determines how many lottery tickets you get.

Attention! You have 20 seconds to do so. A countdown displays how much time you have left. After 20 seconds the page will automatically advance to the next page and no further entry will be possible.

If you do not have a die close by, you can click on a link which leads you to an external website. On this website, a die roll is simulated. The result of your die roll will not be traced in any way.

The time continues to count down while you roll the die on the external website this website, a die roll is simulated. The result of your die roll will not be traced in any way.

Please do not be a cheater.

Q Click here to roll the die

Q What number did you roll? _____

This part was again shown to everyone.

Q Please specify your gender

- ○ Female (1)
- ○ Male (2)
- ○ Other (3)

Q How old are you? _____

Q Are you a student?

- ○ Yes (1)
- ○ No (2)

Q Do you have income besides your study loan?

- ○ Yes (1)
- ○ No (2)

Q What is your country of origin?

▼ Afghanistan (1) ... Zimbabwe (1357)

Q If you wish to participate in the lottery to win the money, you can leave your email address here. This information will not be shared. Please click on next to finish the experiment. Thank you.

_____

Appendix A2. The survey in screenshots

This part of the appendix shows an excerpt of how the survey looked like online. The treatment shown is the treatment of time pressure with linguistic cues.

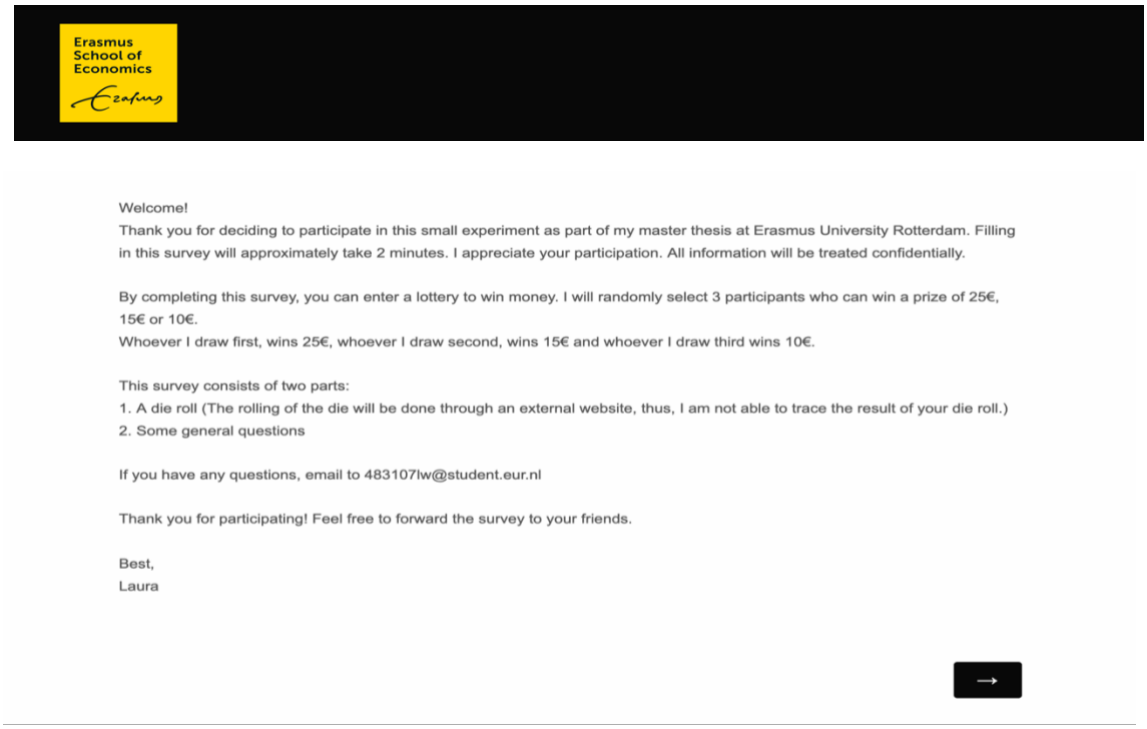

Figure 5: *Survey screenshot of introduction 1*

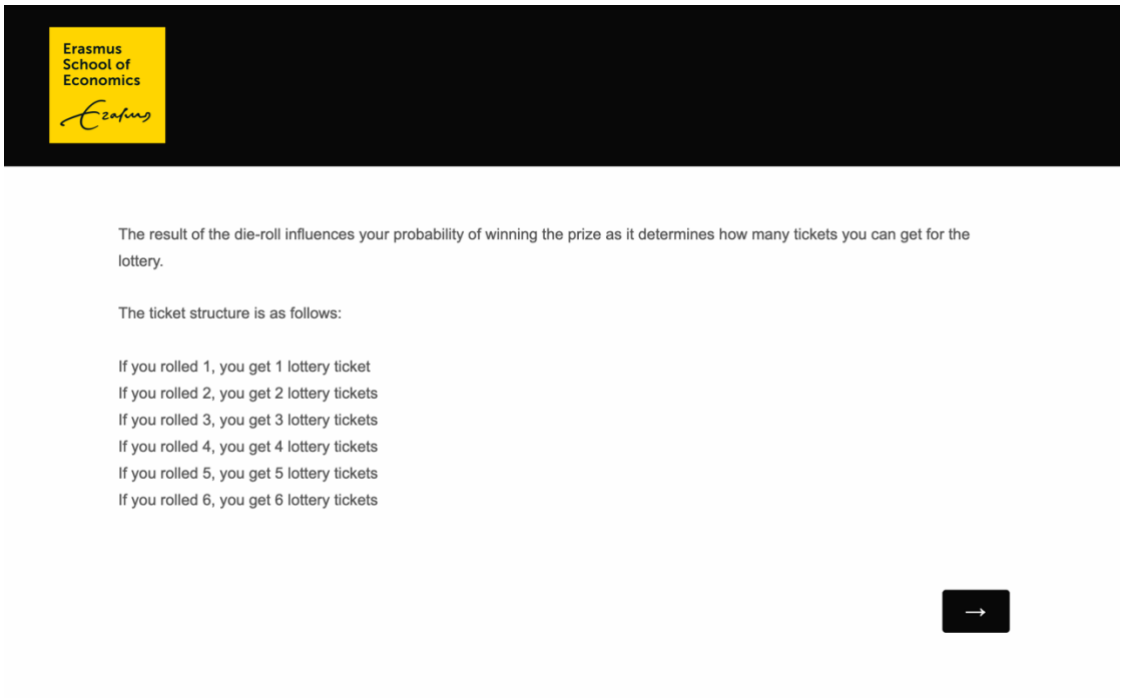

Figure 6: *Survey screenshot of introduction 2*

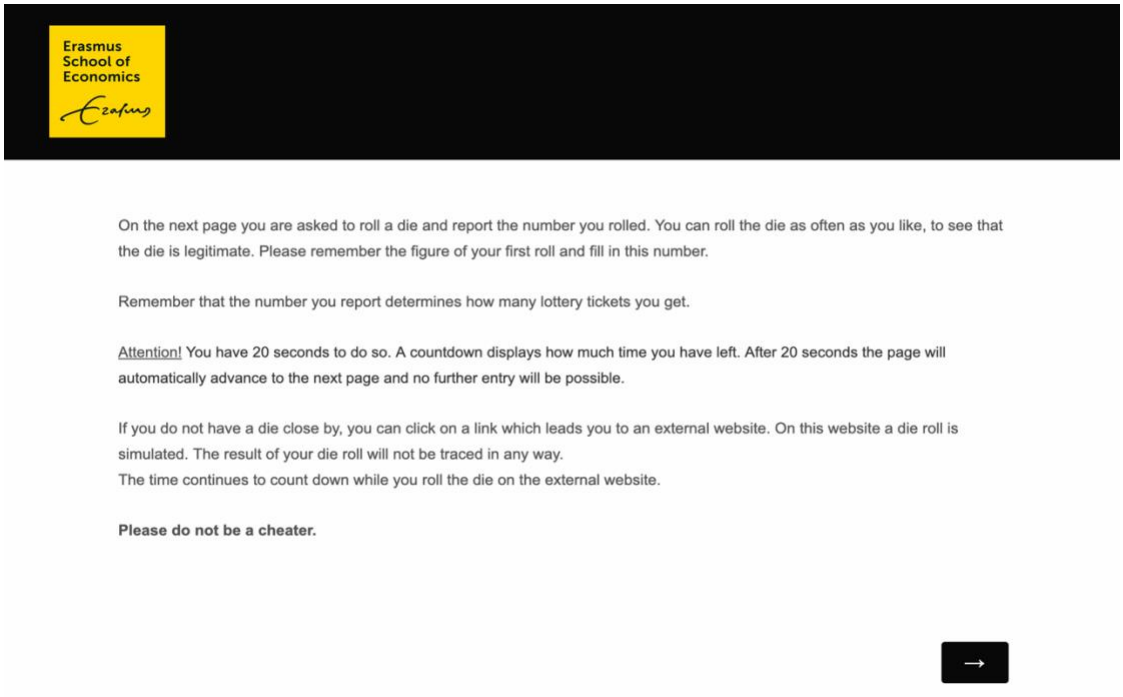

Figure 7: *Survey screenshot of introduction text treatment 3*

Appendix B. Data and Results

Appendix B1. Test for normality

Normality is tested making use of the Kernel density estimate as well as the Shapiro-Wilk test for the die result as well as for the timing variable.

Variable “dice_result”

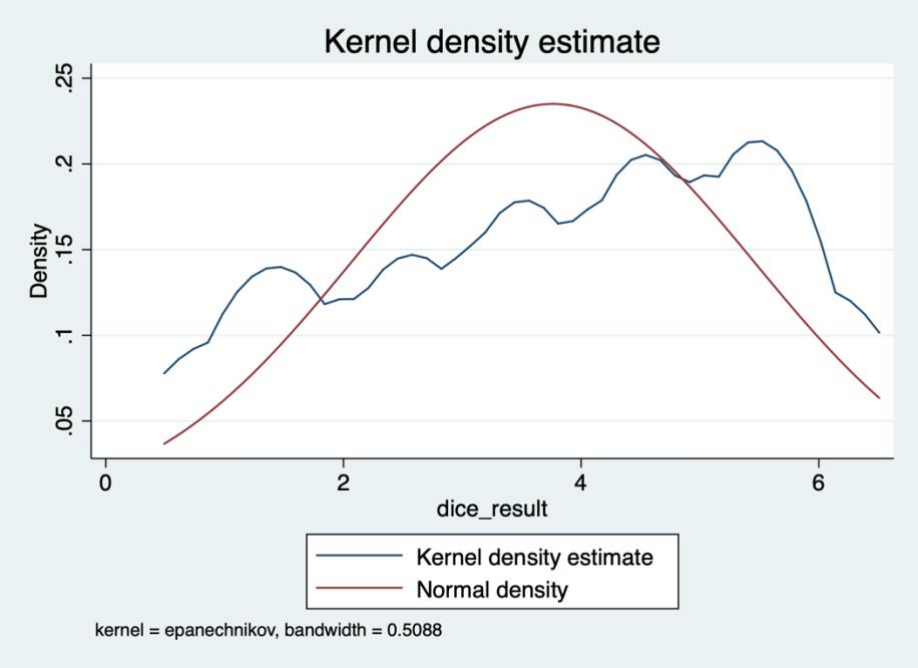

Figure 8: *Kernel density graph variable "dice_result"*

N	SW test statistic	Z-statistic	p-value
244	0.98838	1.682	0.04625

Table 4: *Shapiro-Wilk normality test for variable "dice_result"*

Variable “timing”

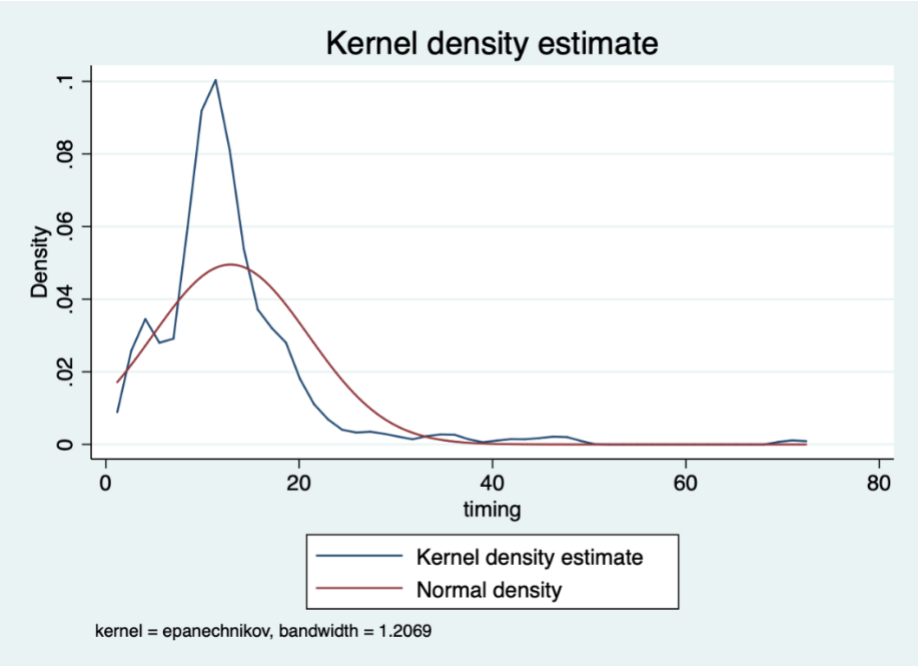

Figure 9: *Kernel density graph variable "timing"*

N	SW test statistic	Z-statistic	p-value
244	0.76349	8.685	0.00000

Table 5: *Shapiro-Wilk normality test for variable "timing"*

Appendix B2. Mann-Whitney-U tests in Chapter 4.2

The table shows the results of the Mann-Whitney-U tests applied to test for the hypotheses of this research in chapter 4.2.

	Z-statistic	p-value	Adjustment for ties
Baseline treatment vs. time pressure treatment	-0.396	0.6922	1439.13
Baseline treatment vs. linguistic cue treatment	0.683	0.4947	1382.54
Time pressure treatment vs. time pressure treatment with linguistic cues	0.850	0.3951	1035.30

Table 6: *Mann-Whitney-U tests to compare lying behavior among treatments*

Appendix B3. Variance inflation factors

The variance inflation factors can be used to test for multicollinearity in OLS regressions. The higher the variance inflation factor, the more severe the problem of multicollinearity. The factors for the unrestricted model in this study however does not show evidence for severe multicollinearity.

Unrestricted model

	VIF
Time pressure treatment	1.49
Linguistic cue treatment	1.50
Time pressure & linguistic cue	1.48
Age	1.37
Female	1.15
Student	1.46
Income	1.21
East-Europe	1.12
Northern America	1.09
Asia	1.06
Middle East	1.07
Australia	1.06
Africa	1.06
Central & South America	1.02
Mean VIF	1.22

Table 7: *Variance inflation factors for the unrestricted model*

Appendix B4. Residuals of the linear regression

The histogram, as well as the Kernel density function, clearly shows that the residuals are not normally distributed.

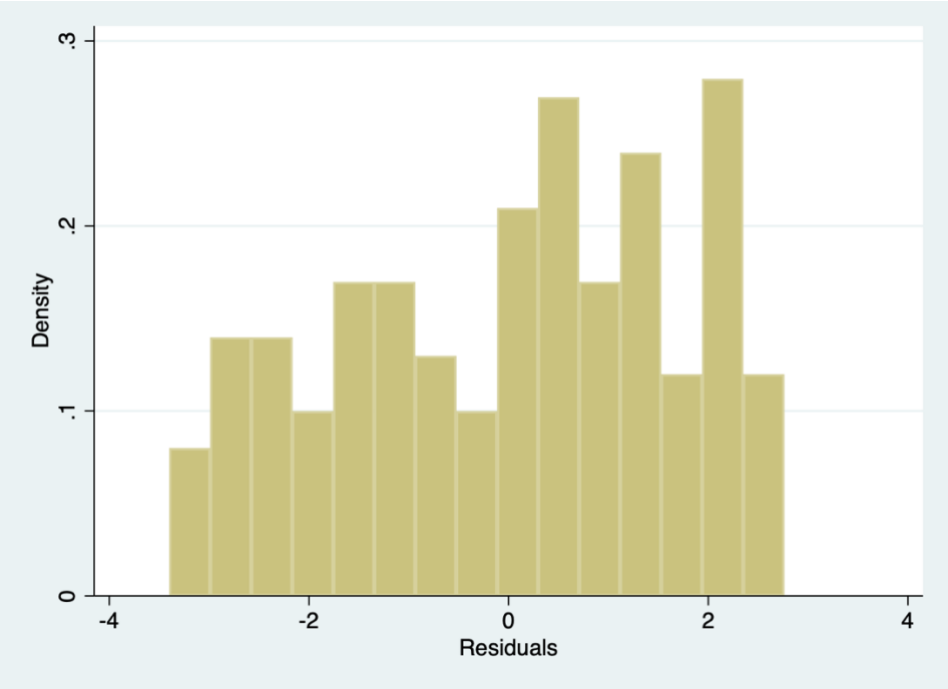

Figure 10: *Histogram of regression residuals*

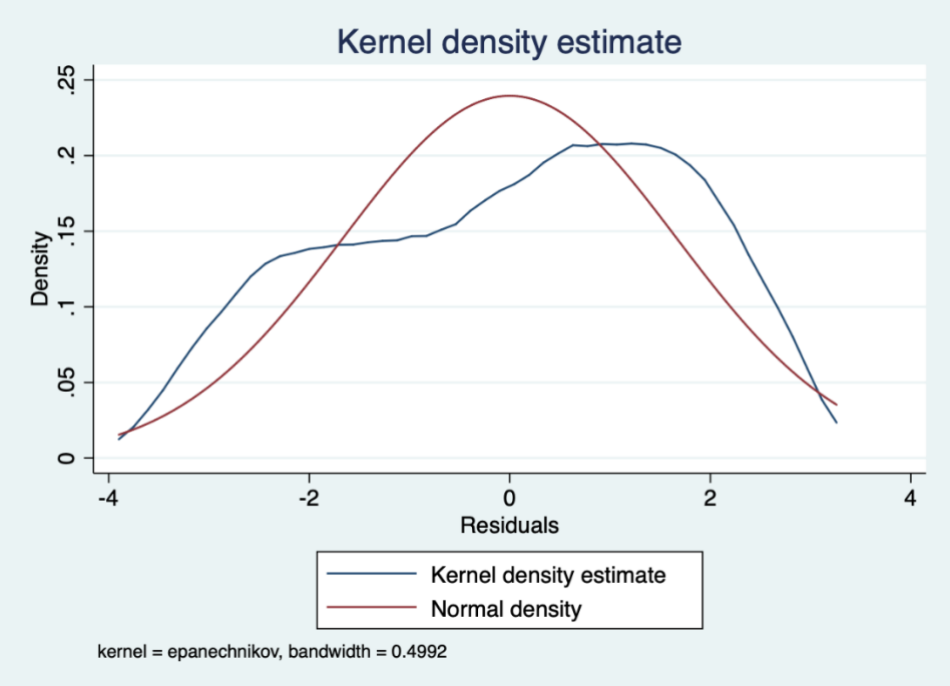

Figure 11: *Kernel density graph regression residuals*

Appendix B5. Post-hoc power analysis

For all subsequently shown post-hoc analyses the conventional level of α=0.05 was used. The effect size d was calculated using the effect size drawer by plugging in the values for the mean and standard deviation respectively.

Baseline treatment and time pressure treatment

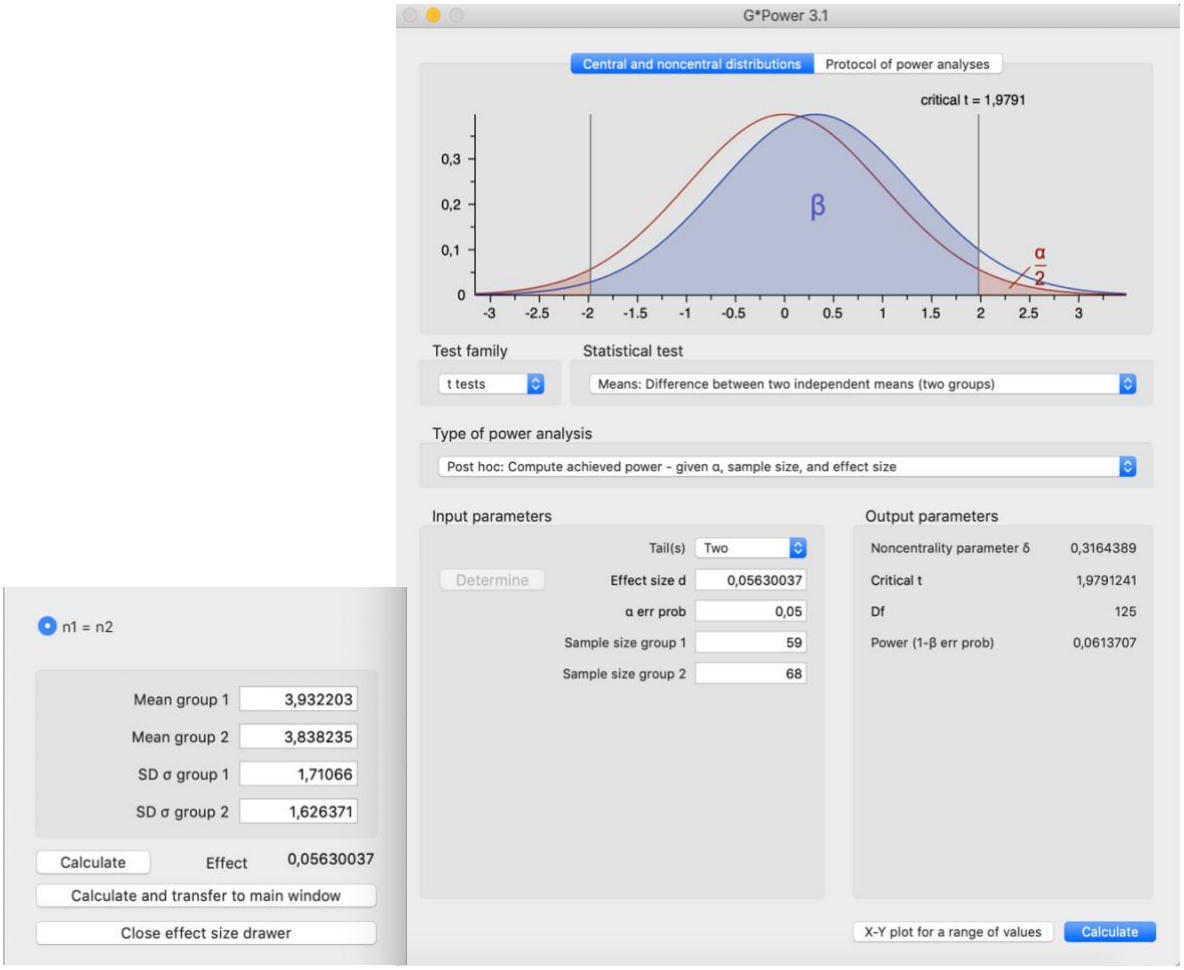

Figure 12: *Post-hoc power analysis: baseline treatment and time pressure treatment*

Baseline treatment and linguistic cue treatment

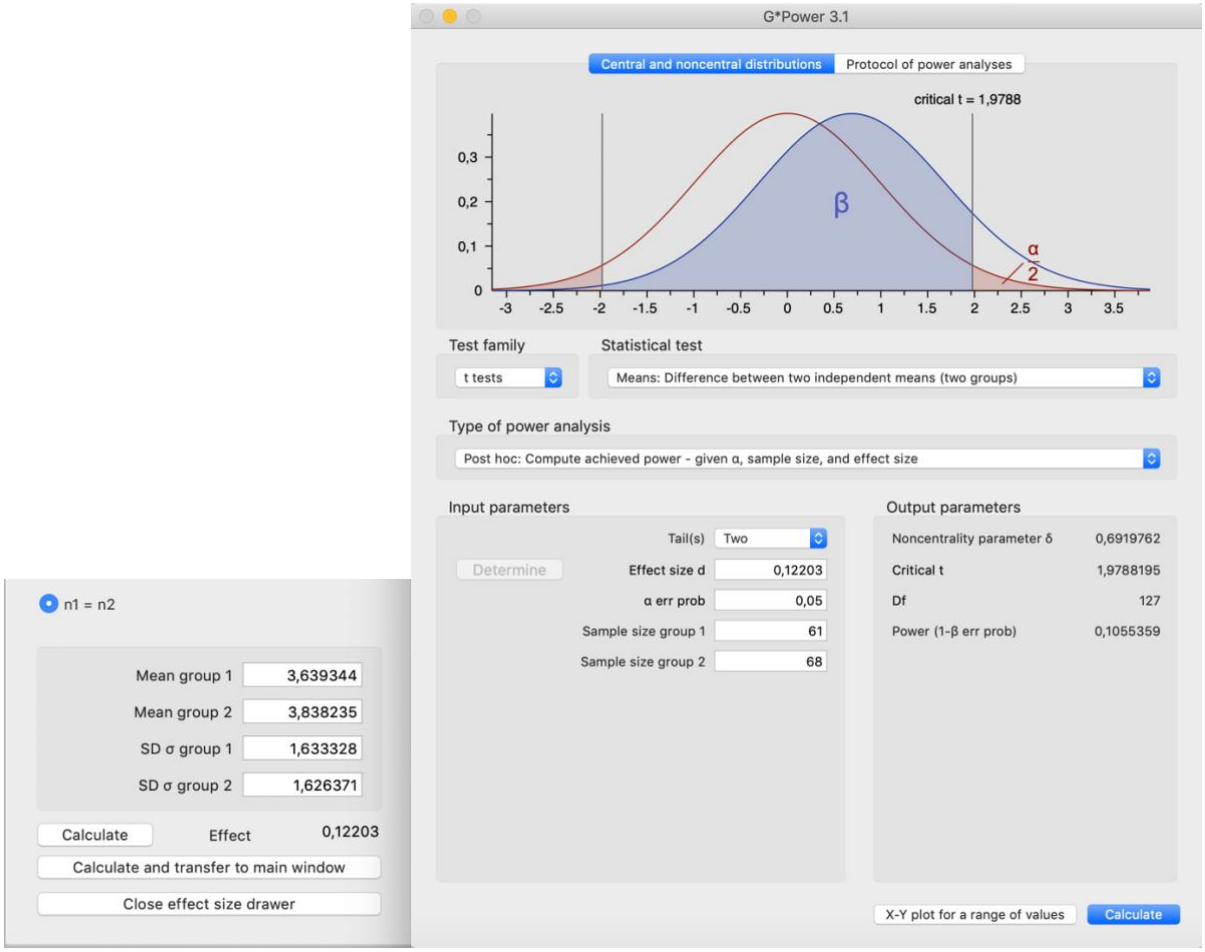

Figure 13: *Post-hoc power analysis: baseline treatment and linguistic cue treatment*

Time pressure treatment and time pressure with linguistic cue treatment

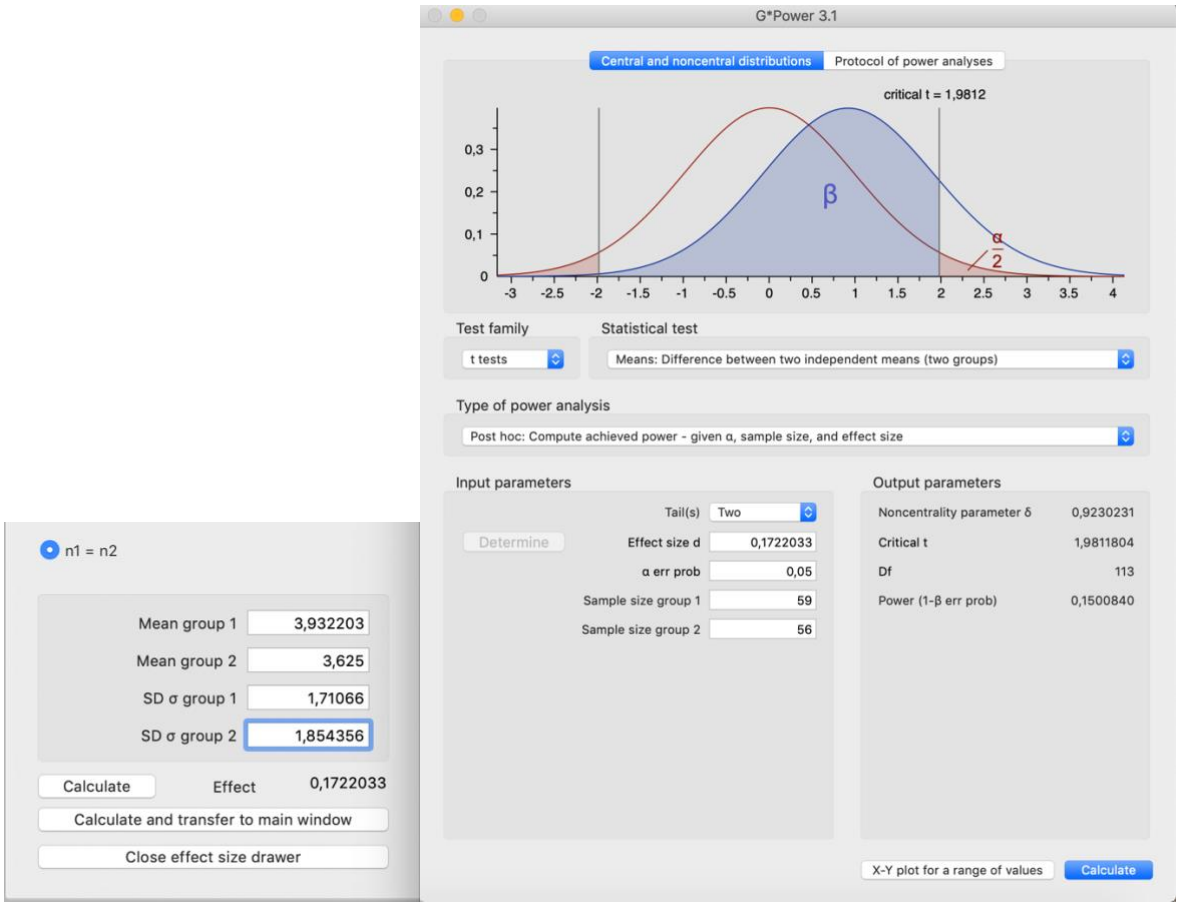

Figure 14: *Post-hoc power analysis: time pressure treatment and time pressure with linguistic cue treatment*

Appendix B6. A-priori power analysis

As seen in Appendix B4, three different post-hoc power calculations were used to get the power of the analyses made in this research. In order to estimate the minimum sample size required to achieve a power of 0.8 with the significance level of α=0.05 for all calculations made in this research, I used the smallest achieved effect size of 0.056. This effect size was calculated using the baseline and time pressure treatment (see Appendix B4). Assuming again the usage of the student t-tests for the analyses, the power calculation shows that approximately 4954 participants are needed in each treatment to achieve a power of 0.8.

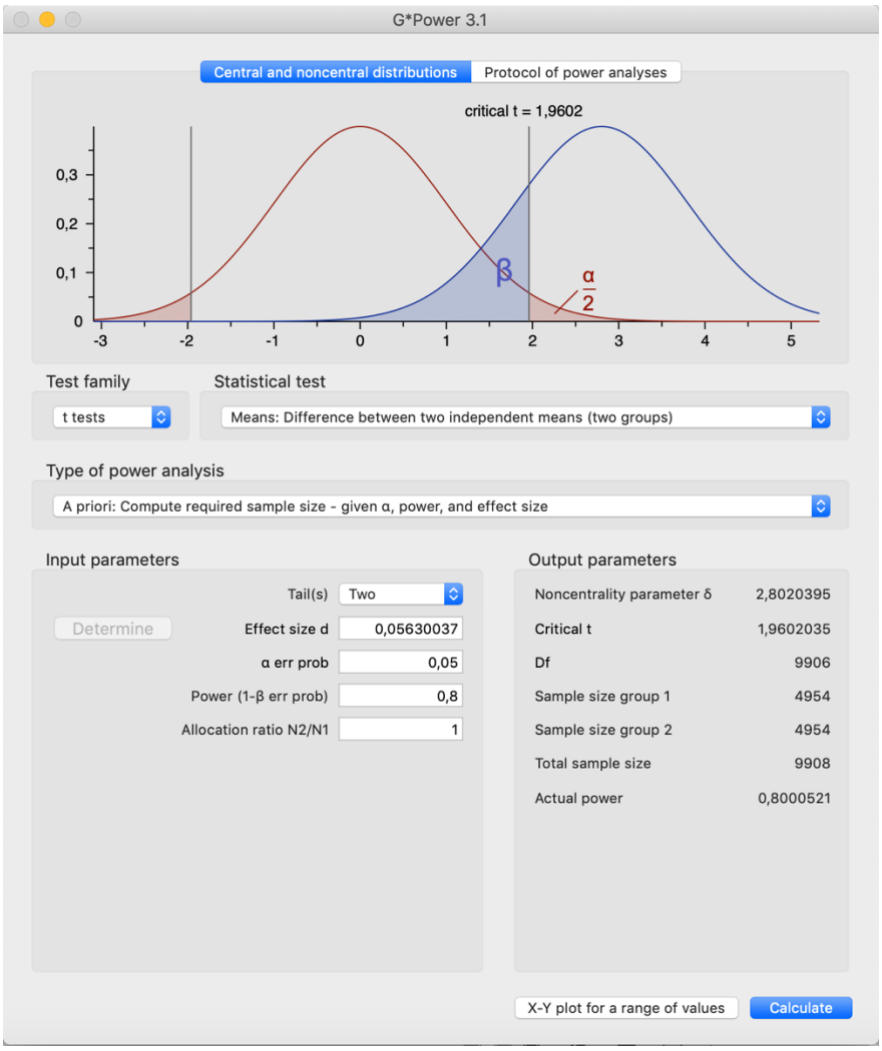

Figure 15: *A-priori power analysis: baseline treatment and time pressure treatment*