

Essays on Information Preferences and Morality

F. Ceren Ay

A thesis presented for the degree of

Philosophiae Doctor (Ph.D.)



NORGES HANDELSHØYSKOLE

Department of Economics

NHH Norwegian School of Economics

Bergen

2021

Acknowledgements

First and foremost, I would like to thank my supervisor Erik Ø. Sørensen for his consistent guidance and support over the course of my Ph.D.. The papers in this thesis greatly benefited from his knowledge, advices and most importantly patience. Under his supervision, Ph.D. was a constant journey of learning not only about research, but also about life, software and typography. Without his great efforts, I would not be able to grow as a researcher and my projects would not be as carefully thought as they are now. I frankly feel privileged for having the chance to have hours of meetings, detailed revisions and always insightful comments from Erik.

I would like to thank my co-authors in the second chapter of this thesis, Joel W. Berge and Katrine B. Nødvedt, for their role in this paper. I would also like to thank Stefan Meißner, who is a great colleague and a very good friend, and my co-author in the third chapter of this thesis. I would like to thank Hallgeir Sjøstad, for inviting me to join the research collaboration to investigate social and moral psychology during Covid-19, which resulted as the fourth chapter of this thesis. I am grateful to Hallgeir also for interesting and always fruitful discussions and for always being positive and encouraging, working with him has been inspiring.

Being a part of Department of Economics at NHH and FAIR was a very valuable opportunity for me. I am grateful to all my colleagues at NHH for creating a very supportive research environment. I am particularly thankful to Bertil Tungodden for being always supportive and open for discussions, and creating an amazing research group, FAIR, which feels more like a family. I cannot thank enough to Catalina Franco, for her help on my research and overall being a Ph.D.

student, and for being very motivating and supportive. She has been a great support for me in the last 2 years and her suggestions improved my papers extensively. I am also grateful to the Ph.D. coordinators Aline Bütikofer and Katrine V. Løken for being very helpful about any problem that I faced. I would like to express my gratitude particularly to Adriana, Anna, Janina and Karen for being the greatest force when struggling to handle the practical part of being a Ph.D. student and doing research.

I was very lucky to be in the same cohort with Stefan and Sara, I am grateful to both for being inspiring friends and researchers. I would also like to thank to fellow Ph.D. students, particularly to Alessandro, Charlotte, Daniel, Felix, Ingvild, Kjetil, Mads, Mirjam, Ranveig, Shrey, and Xu for being great colleagues and great friends. This thank you should be extended to my fellow Ph.D. student Oda and Nina Serdaravic from UiB, whose research minds and friendship have added joy, especially to the last stages of my Ph.D..

During my Ph.D., I had the chance to visit Brown University for a semester. I am very grateful to Prof. Louis Putterman for firstly hosting me, and for enjoyable discussions. I am also grateful to my secondary supervisor Pedro Dal Bó, for his comments on my papers which contributed a lot in the final versions. I would also like to thank Steven Sloman, for the insightful discussions at the Sloman Lab.

Lastly, I want to thank my family and friends for their support through my Ph.D.. I am particularly grateful to Mert, for simply everything, it would be much more difficult without him.

Contents

Introduction	1
1 Information and Strategic Avoidance in Reciprocal Decisions	15
1.1 Introduction	17
1.2 Experimental Design	23
1.3 Sample and Procedure	31
1.4 Results	31
1.4.1 Information and Ignorance	32
1.4.2 Information Acquisition and Avoidance	34
1.4.2.1 Willingness to Pay for Information and Ignorance	34
1.4.3 Motivations for Information Preferences	36
1.4.4 Effects of Unwanted Information	41
1.5 Concluding Remarks	43
Appendices	
1.A Supplementary Analysis	50
1.A.1 Information Avoidance and Reciprocity	51
1.A.2 Ex-Post Information and Revisions	51
1.A.3 Trustors' Sending Decision	52
1.B Instructions	57

1.B.1	Information about the roles and endowments	57
1.B.2	Information for trustors	57
1.B.3	Information for trustees	57
1.B.4	Information decision of trustees in Hidden Information and Ex-Ante Avoid- ance	58
1.B.5	Information decision of trustees in Ex-Post Avoidance	58
1.B.6	WTP elicitation with BDM	58
1.B.6.1	Hidden Information	58
1.B.6.2	Ex-Ante Avoidance	59
1.B.6.3	Ex-Post Avoidance	59
1.B.7	Conditional returning decision	60
1.B.7.1	In Hidden Information and Ex-Ante Avoidance	60
1.B.7.2	In Ex-Post Avoidance	60
1.C	Post-Experiment Questionnaire	61
2	Strategic Curiosity:	
	An Experimental Study of Curiosity and Dishonesty	63
2.1	Introduction	65
2.2	Experimental Design	69
2.2.1	Treatments	71
2.2.1.1	Baseline	71
2.2.1.2	Exogenous availability of information	71
2.2.1.3	Endogenous availability of information	72
2.3	Sample and Procedure	76
2.4	Results	78
2.4.1	Dishonesty	79
2.4.2	Curiosity	81

2.4.3	Motivations for Curiosity	84
2.4.4	Information, Curiosity and Dishonesty	88
2.5	Concluding Remarks	91
Appendices		
2.A	Supplementary Analysis	96
2.A.1	Details of Dishonesty	96
2.A.2	Dishonesty and Ex-Post Information Collection	96
2.B	Instructions	99
3	Investigating Motivations for Information Avoidance—The Role of Certainty, Rewards and Overconfidence	117
3.1	Introduction	119
3.2	Ego Utility and Information	123
3.2.1	Experimental Design	123
3.2.2	Sample and Procedures	127
3.2.3	Results	128
3.3	Neutral Information and Information Preferences	134
3.3.1	Experimental Design	134
3.3.2	Sample and Procedure	137
3.3.3	Results	137
3.4	Conclusion	143
Appendices		
3.A	Supplementary Analysis	148
3.B	Ego Utility Instructions	149
3.B.1	Examples of Progressive Matrices	156

4 National Identity Predicts Public Health Support During a Global Pandemic: Results

from 67 Nations	157
4.1 Introduction	163
4.2 Overview	165
4.3 Results	166
4.4 Discussion	171
4.5 Implications	172
4.6 Limits of Generalizability	173
4.7 Methods and Analysis	175
4.7.1 Participants	175
4.7.2 Survey	175
4.8 Analysis	177
4.9 References	181
4.10 Supplement	190

Introduction

“Facts do not cease to exist because they are ignored.”

Aldous Huxley

In the light of the information age, immense amounts of information on almost every topic and from numerous sources are easily accessible and provided even when people do not ask for it. People are often exposed to information about their environment, abilities, health, past, present and future. With this overload of information, to filter and choose what they really want to know, people must put time and energy on seeking, avoiding and processing information. This means, being ignorant or informed about something is a choice, with time and processing power being the important constraints. One important question to answer to be able to understand changes in the social and economic interactions is: how do people decide what they want to know and to not know?

Research in economics and psychology show a clear pattern of behavior for information choice: people want to know things that make them feel good, and to not know things that make them feel bad. People may feel good or bad about many kinds of information and particularly “*news*” about one’s own abilities, inner traits, dispositions and intelligence can fall in both categories. People often enjoy learning that they are seen as a good person and liked by others, have higher abilities and are more successful than they believed. As it might be pleasant to learn such positive things, when the information contains a judgment that contradicts one’s own existing beliefs, it may bring disutility to the acquirer. If the disutility is anticipated in advance, people may choose not to know certain things by avoiding information. Andries and Haddad (2020) describe

this behavior as “closing eyes to avoid the stress of the roller coaster of life”.

In many real life situations people tend to shape their preferences for information in a self-serving way to preserve their beliefs. For example, Karlsson et al. (2009) show that, investors monitor their portfolios frequently in rising markets, whereas they do not when the markets are flat or falling. In bad times they “put their heads in the sand” to avoid bad news. Another example shows that such behavior can be observed even in professional investors as a group who can be “a color-blind in a sea of red flags” and can cause immense financial loss (Norris, y 16).¹ Not only in financial investments but also in the decisions that consumers make a similar behavior is observed. More information in the label increases demand for eco-labeled products by compelling people to choose these products over others (Loureiro et al., 2002). Reczek et al. (2018); Onwezen and van der Weele (2016) show that when making consumption decisions, people ignore information about ethically relevant products.

Even when there is no observer, people may worry about learning unpleasant things about themselves. This concern may stem from the fact that people strive for being successful and good people, and learning that they could not manage may be painful. Bénabou and Tirole (2003) and Bénabou and Tirole (2006) explain this behavior with the concerns on how people think of and see themselves—*self-image*. Similarly, Bodner and Prelec (2003) argue that people learn things about themselves—*get diagnosis*—from their actions and decisions. They can learn how altruistic and fair, and how skillful and intelligent they are from the decisions they make and activities they engage in. When this diagnosis is not congruent with the existing beliefs, it may be hurtful to learn and people may try to avoid this information by not learning their scores in tests or the consequences of their decisions. This behavior might be caused by the uncertainty that people feel about their own morality and abilities. Köszegi (2006) explains this with the confidence that people have about their own beliefs: if people are confident about their inner traits and abilities, information that contradicts those beliefs cannot change the opinions much. However, if people

¹Norris’s (y 16) NYT article covers the mortgage securitization of Merrill Lynch in 2007. It was reported in 2019 that there were more than 100 “red flags”.

are not confident about their own beliefs, then information can change the existing beliefs which could be detrimental for ego. For example, Eil and Rao (2011) show that people incur monetary costs to avoid learning how they performed in an intelligence test and how attractive they are seen by other people.

As it is important for people to appear skillful to others and to themselves, to pursue a positive moral image has great importance as well. In many decision, people need to forego private benefit to make moral decisions or to increase other's welfare. In the previous studies, it has been shown that in such situations people are eager to appear moral, either to others (Ariely et al., 2009; Lacetera and Macis, 2010) or to themselves (Dana et al., 2007; Mazar et al., 2008; Falk, 2020). In such decisions, people may feel compelled to act according to moral constraints for their desire to be regarded as good people by others and by themselves. Absence of information about the consequences of their decisions may help to circumvent these moral constraints in such cases. Not knowing what a decision can cause or did cause can make the selfish decisions easier without having concerns about morality. A lot of research show that when they are in situations in which increasing private monetary benefit is at odds with causing a loss for someone else or making unfair decisions, people choose to opt out from information about the consequences for others and selfishness increases substantially (Dana et al., 2007; Grossman and Van Der Weele, 2017; Broberg et al., 2007). People deliberately choose ignorance self-servingly to make selfish decisions when faced with a moral conflict. Golman et al. (2017) provide a detailed review of situations in which information avoidance can be chosen strategically, such as consumption and political decisions that may potentially hurt others.

As people can actively shape their information environment, avoidance may not be the only tool to circumvent moral constraints. People sometimes, may collect "additional" information to excuse the selfish decisions. The demand for unnecessary information is shown in the previous literature for a desire receive anticipated positive news about one's own intelligence (Eliaz and Schotter, 2010), and appetite for knowledge or a desire to close the gap between what is known and

what information is available (Golman and Loewenstein, 2018; Golman et al., 2019; Loewenstein, 1994). People may demand additional information not only to receive the anticipated utility or pure curiosity, but also to justify morally questionable decisions, such as dishonesty.

Information is central in almost all economic and social decisions that people face since what is known matters. For example, people need to know about their abilities when deciding what to study or what career to choose, about the specification of products they consume and about the consequences of their decisions for themselves and for others, and this is particularly important when there is a potential impact on other people. This means absence of information can contribute to the inefficiencies that are caused by the decisions that people make. People may end up choosing the wrong education and career path, may consume products that are harmful for themselves and for others, and may cause a loss for others as a result of their decisions. This means investigating motivations for information preferences and the effects on individual decisions is crucial to understand these inefficiencies. The main focus of the chapters of this thesis is the association between information preferences and moral decisions.

In the first three chapters of this thesis, I explore how people actively shape their information environment when they face a trade-off between their self-image and monetary gain. In the first two chapters, the self-image is about moral relevance of decisions that can affect others badly or might involve deception to increase private benefit. In the third chapter, self-image is in a context in which people face a trade-off between learning how well they performed in an intelligence test and monetary gains. In the fourth chapter, I report the findings of a study on social and moral psychology during Covid-19 from an international collaboration. All chapters include pre-registered laboratory, online or a survey experiments to investigate behavioral motivations of individual decisions.

In the first chapter of this thesis which is titled as “Information and Strategic Avoidance in Reciprocal Decisions”, I investigate motivations for information preferences by focusing on decisions in which concerns about how others and oneself can see it—*image concerns*—may have

an impact. In a reciprocal setting, I document that people avoid information even when it cannot affect their decision. In such settings, people make decisions in response to another person's decision in a positive or a negative way, such as kindness to kindness and rudeness to anger. Not reciprocating means betraying another person's intentions. To investigate how information and preferences over information can change reciprocity, I conducted a laboratory experiment with a manipulated trust game which is played in pairs (trustors and trustees). Trustors decide whether to send money to the trustee, and then trustees decide whether to return any to the trustor. In the standard game, people make decisions with complete information about each other's endowment, whereas in the present paper, both players are not informed about the consequences of their decisions in terms of final distributions. Having such a design enables me to investigate the effects of exogenously given information and ignorance. However, trustees have the chance to *acquire* or to *avoid* information by incurring a monetary cost. Acquiring information can help to reciprocate more and lead to more equal distribution, whereas avoiding information can help not to feel bad about making selfish decisions and not reciprocating.²

I aim to answer mainly two questions in this paper. The first question is: *Can exogenously provided information change reciprocal decisions?* To investigate this question, trustees in one treatment are exogenously given the information about their trustor's initial endowment whereas in another treatment they made their decision under ignorance. I document that exogenously given information increases reciprocity substantially compared to uninformed decisions. The second question is: *Do individuals avoid information strategically to make selfish decisions and to preserve a good image even when the decision has already been made?* To answer this question, trustees are given the chance to deliberately choose whether they want to know trustor's endowment level. I document that there is a link between information preferences and reciprocity; low levels of reciprocity are positively correlated with information avoidance and higher levels of reciprocity are correlated with information acquisition. This pattern is observed both before and after

²I use *selfish* to describe the decision that is not prosocial in the experiment.

the reciprocal decision has been made, which suggests that even after the decision has been made, learning that it caused selfish outcomes may make one feel bad.

This paper contributes to understanding the insights of reciprocal decisions with three main results on the effects of exogenous information settings: the association between reciprocity and information, and ex-ante and ex-post information preferences. The first main result is, when there is exogenous information, average returned amounts are approximately 2 times higher than when there is no chance to receive information ex-ante. Second, I document that information avoidance is negatively related with the returned amounts and, thus, is motivated by the monetary rewards. People who choose to know the consequences of their decisions return significantly higher amounts on average than people who choose to avoid information. And the third result of this paper is, I report that even after they have already made the decision, people who choose to return lower amounts are more likely to avoid information. 60% of the people who wanted to avoid information returned 0 to their partners. Information about the consequences, particularly in reciprocal decisions, does not only have an instrumental value in terms of reaching better distribution decisions, but also a psychological value by signaling how prosocial the decision was even when it had already been made.

The second chapter is a joint paper with two fellow PhD students at NHH, Joel W. Berge and Katrine B. Nødvedt, and titled as “Strategic Curiosity: An Experimental Study of Curiosity and Dishonesty”. In this study, we provide experimental evidence on a novel phenomenon concerning information preferences: people strategically collect additional unnecessary information to justify morally questionable decisions. We conduct a virtual dice-rolling experiment where participants roll a dice and self-report the outcome of the first roll for monetary rewards. In this setting, we vary the extent to which participants can continue rolling the dice before reporting as well as the displayed content of those additional roll-outcomes. We document that people systematically roll the dice more—*are more curious*—when tempted to misreport. We find that curiosity is positively correlated with the size of the lie. However, contrary to previous studies, we observe no variation

in dishonesty across treatments regardless of the possibility to collect additional non-instrumental information. This study provides new insights into how individuals actively shape their information environment in pursuit of self-interest.

In a pre-registered experiment, we recruited 1580 US participants on Amazon Mechanical Turk (mTurk). In the base treatment (*Single Roll*), participants roll the dice once and then report the outcome. In this condition, people can continue to roll the dice after reporting while knowing that they cannot change their report. To study how people search for additional information when tempted to misreport, we introduce three variations to identify causal effects of different information environments. In the *Three Rolls* treatment, the number of rolls is limited to three, and people do not have a chance to roll less or more than three times before reporting the outcome of the first roll. In the other two treatments, participants can roll as many times as they want to before reporting. In the *Multiple Numbers* treatment, the dice's sides always display numbers, whereas in the *Multiple Figures* treatment, the dice displays random figures after the first roll. Since only the outcome of the first roll should be reported, additional rolls have no instrumental value in the case of honest reporting.

Similar to avoiding information that makes it difficult to excuse selfish decisions (Grossman, 2014), we find evidence that curiosity is driven by a desire to justify selfish behavior. Our main finding is that non-instrumental information is collected strategically, implying that people are *strategically curios*. We call the behavioral phenomenon that people collect related and unrelated additional information to relax moral constraints that are at odds with their self-interest *strategic curiosity*. Contrary to the previous literature, we find no variation in dishonesty across treatments despite the possibility to collect additional information and the content of this information. With our explicit focus on deliberate decisions to acquire additional non-instrumental information when tempted to misreport, we contribute to the research on information and moral decisions by showing that curiosity can also be driven by a desire to justify selfish behavior.

Our results provide evidence that people are systematically more curious when collecting

additional information can help justify dishonesty. As this is an observed game, we can disentangle whether collecting additional information that is related or unrelated to the outcome leads to higher misreporting. We find that people who observe lower outcomes in the first roll are more likely to roll more times before reporting the outcome of the first roll. This result indicates that people acquire additional information—not only to search for justifications—but also to distract themselves from moral conflict.

This study contributes to research on preferences toward information in moral dilemmas, which can arise both before and after a decision has been made, and both when the decision-maker possesses incomplete and full information. Various motivations that shape information preferences have been documented previously, and curiosity is one of the prominent drivers of information acquisition. Loewenstein (1994) provides a review on curiosity and posits that demand for information is intrinsic; it is “appetite for knowledge.” This kind of curiosity directed toward all kinds of information is defined as *epistemic curiosity* (Litman et al., 2005). We find that curiosity need not only arise from an intrinsic desire for information; it might also be driven by a strategic desire to justify selfish behavior or distract oneself from moral conflict. More recently, Golman and Loewenstein (2018) introduce the concept of “information gaps”, which refers to people’s desire to collect information to close the gaps between what is already known and what information is available. Eliaz and Schotter (2010) provide experimental evidence that individuals are willing to pay to receive information regarding the results of an intelligence test even when this information has no value in terms of achieving higher outcomes. We find that people acquire unnecessary information particularly when tempted to make morally questionable decisions. This extends the literature on curiosity by showing that curiosity may be motivated by strategic reasons rather than only an innocuous desire to collect information.

The third chapter is in collaboration with a fellow PhD student, Stefan Meißner and titled as “Investigating Motivations for Information Avoidance—The Role of Certainty, Rewards and Overconfidence”. In this paper we investigate an objective for active information avoidance: self-

image protection. In two experiments, we study how information preferences can change when the information is ego related and hurtful for self-image, and when it is neutral. To vary the neutrality of information, in the first experiment we use information about one's relative rank in an intelligence test and in the second experiment we use "neutral" information that is about a randomly chosen number. Our results show that people avoid ego related information and are willing to pay to do so, however, this behavior is not related with confidence and certainty about beliefs in one's own skills. We observe that a high share of people avoid "neutral" information even though it is not ego-related, and the willingness to pay to do so is positive. Besides avoidance, we provide findings that show people's willingness to pay to receive neutral information respond to the value of information. In the first experiment, we only use ego utility and the information is about relative rank in an intelligence test. We find high willingness to pay for both information and avoidance, which did not clarify the motivations for information avoidance. To test whether ego related information is the reason for high avoidance, we ran the second experiment a year after the first one, with using only neutral information. In the second experiment, the neutral information is about a randomly chosen number which is not related to the intelligence level or any personal attribution. We observe that a high share of participants avoid information even though the information cannot show anything about their abilities. When we compare the amounts people are willing to pay in these two experiments, we find that demand for avoidance from ego related information is significantly higher compared to avoidance from neutral information. In this paper, we show that when making decisions on information, people may have behavioral noise which affects the measures of our treatments. However, by comparing two kinds of information we show that despite the behavioral noise they may have, people are willing to pay higher amounts to avoid ego related information. Our findings contributes to the literature on information and ego utility by showing that these preferences might not be driven solely by self-image concerns.

The last chapter of the thesis is titled as "National identity predicts public health support during a global pandemic: Results from 67 nations". This paper is a result of an international

collaboration of more than 150 researchers from 67 countries, to investigate moral and social psychology during Covid-19. In the last year of my PhD, the Covid-19 pandemic has caused devastating disruptions in economic and social systems all over the world. Millions of people lost their health and even lives, and many others fall into extreme poverty. In the mean time, public health policies differed in every country, and thus, the influence of the pandemic was varied. Understanding how the pandemic will affect individual behavior and the aggregate systems is crucial to be able to evaluate policies and regulations and their future effects. I decided to join this research collaboration in the late stages of my PhD to contribute the behavioral research on the Covid-19 pandemic. The paper in this chapter is the first paper of this project investigating support for public health and its psychological determinants. As of January 2021, more than 90 million people worldwide have been infected by the new coronavirus and more than 1.9 million have died, making the COVID-19 pandemic one of the greatest health crises of the last 100 years. Until a vaccine or effective medical treatment is developed, the public response to the current pandemic is largely limited to policy-making and collective behavior change (Lewnard and Lo, 2020). To reduce virus transmission, it is crucial that people engage in public health behavior (e.g., maintain spatial distance and improve physical hygiene) and support COVID-19 protective policies (e.g., limiting travel and closing bars and restaurants). Indeed, the Director of the World Health Organization recently argued: “That is why behavioural science is so important—it helps us to understand how people make decisions, so we can support them to make the best decisions for their health” (WHO, 2020). In the current investigation, we respond to this call for behavioral science and present the results from a global study across 67 countries, testing key predictors of public health support.

Bibliography

Andries, M. and Haddad, V. (2020). Information aversion. *Journal of Political Economy*, 128(5):1901–1939.

- Ariely, D., Bracha, A., and Meier, S. (2009). Doing Good Or Doing Well? Image Motivation And Monetary Incentives In Behaving Prosocially. *American Economic Review*, 99(1):544–55.
- Bénabou, R. and Tirole, J. (2003). Self-Knowledge And Self-Regulation: An Economic Approach. *Psychology of Economic Decisions*, 1:137–167.
- Bénabou, R. and Tirole, J. (2006). Incentives and Prosocial Behavior. *American Economic Review*, 96(5):1652–1678.
- Bodner, R. and Prelec, D. (2003). Self-Signaling And Diagnostic Utility In Everyday Decision Making. *Psychology of Economic Decisions*, 1:105–126.
- Broberg, T., Ellingsen, T., and Johannesson, M. (2007). Is Generosity Involuntary? *Economics Letters*, 94(1):32–37.
- Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting Moral Wiggle Room: Experiments Demonstrating An Illusory Preference For Fairness. *Economic Theory*, 33(1):67–80.
- Eil, D. and Rao, J. M. (2011). The Good News-Bad News Effect: Asymmetric Processing Of Objective Information About Yourself. *American Economic Journal: Microeconomics*, 3(2):114–138.
- Eliasz, K. and Schotter, A. (2010). Paying For Confidence: An Experimental Study Of The Demand For Non-Instrumental Information. *Games and Economic Behavior*, 70(2):304–324.
- Falk, A. (2020). Facing yourself – a note on self-image. *Journal of Economic Behavior Organization*.
- Golman, R., Hagmann, D., and Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1):96–135.
- Golman, R. and Loewenstein, G. (2018). Information Gaps: A Theory of Preferences Regarding the Presence and Absence of Information. *Decision*, 5(3):143.

- Golman, R., Loewenstein, G., Molnar, A., and Saccardo, S. (2019). The demand for, and avoidance of, information. *Information (May 31, 2019)*.
- Grossman, Z. (2014). Strategic Ignorance And The Robustness Of Social Preferences. *Management Science*, 60(11):2659–2665.
- Grossman, Z. and Van Der Weele, J. J. (2017). Self-Image and Willful Ignorance in Social Decisions. *Journal of the European Economic Association*, 15(1):173–217.
- Karlsson, N., Loewenstein, G., and Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty*, 38(2):95–115.
- Köszegi, B. (2006). Ego utility, Overconfidence, and Task Choice. *Journal of the European Economic Association*, 4(4):673–707.
- Lacetera, N. and Macis, M. (2010). Do All Material Incentives For Pro-Social Activities Backfire? The Response To Cash And Non-Cash Incentives For Blood Donations. *Journal of Economic Psychology*, 31(4):738–748.
- Lewnard, J. A. and Lo, N. C. (2020). Scientific and ethical basis for social-distancing interventions against covid-19. *The Lancet. Infectious diseases*, 20(6):631.
- Litman, J., Hutchins, T., and Russon, R. (2005). Epistemic curiosity, feeling-of-knowing, and exploratory behaviour. *Cognition & Emotion*, 19(4):559–582.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, 116(1):75.
- Loureiro, M. L., McCluskey, J. J., and Mittelhammer, R. C. (2002). Will Consumers Pay A Premium For Eco-Labeled Apples? *Journal of Consumer Affairs*, 36(2):203–219.
- Mazar, N., Amir, O., and Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of marketing research*, 45(6):633–644.

Norris, F. (2008, May 16). Color-Blind Merrill in a Sea of Red Flags. *N.Y. Times*.

Onwezen, M. C. and van der Weele, C. N. (2016). When Indifference is Ambivalence: Strategic Ignorance about Meat Consumption. *Food Quality and Preference*, 52:96–105.

Reczek, R. W., Irwin, J. R., Zane, D. M., and Ehrich, K. R. (2018). That's Not How I Remember It: Willfully Ignorant Memory for Ethical Product Attribute Information. *Journal of Consumer Research*, 45(1):185–207.

Chapter 1

Information and Strategic Avoidance in Reciprocal Decisions

Abstract

People sometimes avoid information about the consequences of a decision for other people, when it can signal an undesired characteristic about themselves such as being selfish or untrustworthy. Even when the only observer is themselves, people can have *self-image concerns*. This paper contributes to the literature on information preferences and moral decisions by showing the relation between self-image concerns and information preferences in a reciprocal setting. In a laboratory experiment with 520 participants, I show that when information is exogenously given, reciprocity is higher compared to the setting in which information is hidden. People who reciprocate less are more likely to avoid information about the consequences of their decisions for their partner. This behavior is present even when the reciprocal decision has already been made and the information

⁰This work is partially supported by the Research Council of Norway through its Centres of Excellence Scheme, FAIR project No 262675. I am grateful to Alexander Cappelen, Pedro Dal Bó, Catalina Franco, Botond Közsegi, George Loewenstein, Stefan Meissner, Marco Piovesan, Louis Putterman, Hallgeir Sjåstad, Eirik Strømmland, Erik Ø. Sørensen, Sigve Tjøtta, Bertil Tungodden and the participants of PhD Workshop at Norwegian School of Economics and University of Bergen, Brown Bag Seminar at Department of Economics - FAIR at the Norwegian School of Economics, Brown Bag Seminar at the Department of Economics at Brown University, BBE Workhsop at Humboldt University, ESA Job Market Candidate Seminars and Nordic Exchange seminar at Lund University for their helpful comments and contributions. I thank to M. Yiğit Gürdal at Boğaziçi University for making it possible to run the experiment at the Boğaziçi Econ Lab. I would like to thank Centre of Ethics at the Norwegian School of Economics for financial support.

cannot affect the allocation. Overall, the findings show reciprocal decisions are not only motivated by the desire for fair outcomes or being seen as a good person by others, but also the desire to self identify as a good person.

Keywords: Information avoidance; self-image; reciprocity

JEL-Classifications: C91, D82, D83, D91

1.1 Introduction

People often make decisions in which increasing one's own benefit means causing a loss for others. In such situations, morality is often considered as the prominent obstacle to make the decision that serves one's own benefit. Evidence from experiments shows that *avoiding* information about the consequences of a decision for others can reduce these constraints from morality (Dana et al., 2006, 2007; Grossman and Van Der Weele, 2017). There is also evidence from the field: more information in the label increases demand for eco-labeled products by compelling people to choose these products over others (Loureiro et al., 2002). Reczek et al. (2018); Onwezen and van der Weele (2016) show that when making consumption decisions, people ignore information about ethically relevant products. People can choose what information to collect not to feel bad about the decision or to be regarded as a prosocial person by both oneself and others. Economic models often assume that avoidance is a strategy only in the ex-ante perspective—prior to decision making—when the information about consequences can influence the decision and make one feel compelled to act in a certain way. That means, there is no need to avoid information to preserve a desired image if the decision has already been made and changing it upon receiving the information is not an option, or in other words when information is not *instrumental* in decision making. This argument contradicts with people having concerns about past decisions for causing a bad outcome for others such as guilt and regret, and does not explain how preferences over information could change when these concerns are present.

This paper experimentally investigates motivations for information preferences by focusing on decisions in which concerns about how others and oneself can see it—*image concerns*—may have an impact. In a reciprocal setting, I document that people avoid information even when it is not instrumental. In such settings, people make decisions in response to another person's decision in a positive or a negative way, such as kindness to kindness and rudeness to anger. Not reciprocating means betraying another person's intentions. To investigate how information and preferences over information can change reciprocity, I conducted a laboratory experiment with a manipulated

trust game which is played in pairs (trustors and trustees). Trustors decide whether to send money to the trustee, and then trustees decide whether to return any to the trustor. In the standard game, people make decisions with complete information about each other's endowment, whereas in the present paper, both players are not informed about the consequences of their decisions in terms of final distributions. Having such a design enables me to investigate the effects of exogenously given information and ignorance. However, trustees have the chance to *acquire* or to *avoid* information by incurring a monetary cost. Acquiring information can help to reciprocate more and lead to more equal distribution, whereas avoiding information can help not to feel bad about making selfish decisions and not reciprocating.¹

The relation between information and image concerns has been a topic in both theoretical and experimental literature. Bodner and Prelec (2003) defines *diagnostic utility* of a decision as the diagnosis a decision can signal about one's own inner traits and characteristics. Bénabou and Tirole's (2006) theory focuses on the role of diagnostic utility by proposing a framework with a *decision maker* and an *observer*. Bénabou and Tirole (2006) argue that people care how they are seen by the observer which could be others (*social image*) and by themselves (*self image*) when making a decision. Grossman and Van Der Weele (2017) extend these two models to explain the relation between self-image concerns and information preferences in a self signalling model in which the observer is oneself and the decision introduces a conflict between a self serving choice and welfare. Their model concludes that information avoidance is an opt out mechanism only when information can be used to make better decisions in terms of welfare and when some people are tempted to choose the self serving option. Evidence from experiments has shown similar behavior when people are tempted to make selfish decisions or when information is potentially bad news about one's health or intelligence (Eil and Rao, 2011; Ganguly and Tasoff, 2017). Related literature established various motivations for information avoidance such as procrastination, anticipatory utility, guilt aversion and political concerns (Grossman and Van Der Weele, 2017; Attanasi et al.,

¹I use *selfish* to describe the decision that is not prosocial in the experiment.

2015; Golman et al., 2017; Köszegi, 2006). In addition to the experimental evidence, information avoidance is documented in various field settings. Ehrich and Irwin (2005) and Reczek et al. (2018) show that people are reluctant to get ethical information on cheap products and are more likely to remember good information about well working products. It is yet unclear how exogenously given (or not given) information can affect how much people reciprocate and whether people avoid information to prevent feelings of letting someone down and not reciprocating positive intentions even after they have already made the decision and information has no instrumental value.

To fill this gap in the literature I aim to answer mainly two questions in this paper. The first question is: *Can exogenously provided information change reciprocal decisions?* To investigate this question, trustees in one treatment are exogenously given the information about their trustor's initial endowment whereas in another treatment they made their decision under ignorance. I document that exogenously given information increases reciprocity substantially compared to uninformed decisions.

The second question is: *Do individuals avoid information strategically to make selfish decisions and to preserve a good image even when the decision has already been made?* To answer this question, trustees are given the chance to deliberately choose whether they want to know trustor's endowment level. I document that there is a link between information preferences and reciprocity; low levels of reciprocity are positively correlated with information avoidance and higher levels of reciprocity are correlated with information acquisition. This pattern is observed both before and after the reciprocal decision has been made, which suggests that even after the decision has been made, learning that it caused selfish outcomes may make one feel bad.

The present paper contributes to the recently growing literature on information preferences and moral decisions by providing evidence that what motivates moral decisions may not only be the desire for fair outcomes but also what these decisions signal about one's own characteristics to oneself and to others. Results show that lack of information can create a *moral wiggle room* which circumvents reciprocity and makes selfish decisions easier. Similar findings have been documented

in non-reciprocal settings (Dana et al., 2007; Grossman and Van Der Weele, 2017; Golman et al., 2017).

A total of 520 undergraduate students were recruited to participate in the experiment. To have uncertainty about the final distribution after the trustee choose how much to return, the endowment levels are heterogenous—low (15 ECU) or high (35 ECU). Each player is assigned to one of the two endowment levels randomly and are not informed about each other's endowments. The trustees can choose to receive or to avoid this information by incurring a monetary cost, revealing their willingness to pay to receive, and to avoid information (WTP, henceforth). I vary the timing of this decision to investigate information preferences before (*ex-ante*) and after (*ex-post*) deciding how much to return.

This paper contributes to understanding the insights of reciprocal decisions with three main results on the effects of exogenous information settings: the association between reciprocity and information, and *ex-ante* and *ex-post* information preferences. When there is exogenous information, average returned amounts are approximately 2 times higher than when there is no chance to receive information *ex-ante*. I document that information avoidance is negatively related with the returned amounts and, thus, is motivated by the monetary rewards. People, who choose to know the consequences of their decisions return significantly higher amounts on average than people who choose to avoid information. One of the most important contributions I make with this paper is to show that even after they have already made the decision, people who choose to return lower amounts are more likely to avoid information. 60% of the people who wanted to avoid information returned 0 to their partners. Information about the consequences, particularly in reciprocal decisions, does not only have an instrumental value in terms of reaching better distribution decisions, but also a psychological value by signaling how prosocial the decision was even when it had already been made.

Most of the experimental evidence on information preferences and moral decisions are documented from dictator games. In dictator games, a sender is asked to decide how much to share

over a given endowment with a passive recipient. Even though dictator games provide a very simple and abstract setup to investigate prosociality, a decision maker in the game is less socially entangled than in many real life decisions. In addition, in dictator games, on the one hand, players only have information from the experimenter and the resources could be considered as “*mana from heaven*”. Trust games, on the other hand, provide a setup in which the decision maker receives a signal from the partner about their intentions and the decision is entangled with reciprocity. Reciprocal decisions can provide a more realistic setup by introducing the behavioral signal from the trustor’s sending decision which creates a social pressure. Although dictator games are commonly used in the behavioral research and are abstract forms of many social decisions, when the decisions are under greater social pressure, people can avoid information to preserve a good self-image even when the decision has already been made and cannot be changed.

Reciprocal decisions give the chance to respond to someone else’s behavior in a positive or a negative way and are often considered as measures of trust and trustworthiness. When making decisions that affect others, intentions and beliefs play an important role. Arrow (1974) describes trust as “an important lubricant of a social system”. Even though it cannot be easily valued in monetary terms, trust and trustworthiness have a pragmatic value in relations with other individuals and institutions by decreasing inefficiencies. With Arrow’s (1974) words, “*It is extremely efficient; it saves a lot of trouble to have a fair degree of reliance on the other people’s word*”. With advancing technology and numerous devices, information is easily accessible and individuals are able to choose what to know and what not to. It is important to investigate the role of information in decision making to better understand dynamics of morality.

In classical economic theory, almost all models assume people pursue their material benefit when making decisions. However, evidence shows that individuals act according to their concerns for fairness and inequality when faced with a conflict between moral values and self-serving opportunities. This inconsistency between classical theory and evidence shows that classical economic theories lack psychological realism. Behavioral models incorporate psychological motiva-

tions such as fairness motives and concerns about inequality into individual utility, which incorporates psychological realism in economic theories. Outcome based models, propose that concerns about inequality motivate altruistic behavior (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). Cappelen et al. (2007) show that fairness ideals could vary and lead people to make different distribution decisions. These models suggest that individuals suffer from unfair outcomes and this affects how they decide to distribute an amount of money between themselves and others. Outcome based models provide explanations on motivations for distributive decisions, however, they do not explain how concerns about one's image and information can change these decisions. Questions on how information about the consequences of a decision could affect "perception of oneself", *self image*, remain unanswered in outcome-based models. By providing evidence on the relation between selfish decisions and information avoidance, I show that self-image concerns may play an important role for reciprocal decisions and choices toward information may be shaped by these concerns. This finding does not support the classical theory of information economics which considers information as a desirable asset only in the ex-ante perspective when it can help to make superior decisions both for individual utility and social welfare (Stigler, 1961). However, it supports more recent theories that introduce information as a direct source of utility—or disutility—by signalling one's own characteristics and dispositions to oneself and to others, rather than being only a tool (Bénabou and Tirole, 2003; Bodner and Prelec, 2003; Bénabou and Tirole, 2004, 2011). Following this, findings of the present paper also support the argument that when there is a potentially bad signal, it can be painful for people to learn that they are not in the desired good profile and can avoid information even when there is no one else observing (Grossman and Van Der Weele, 2017). Besides this support, I extend Grossman and Van Der Weele's (2017) predictions and findings on information avoidance by showing that in a more complex moral setting, people avoid information ex-post and are willing to incur a monetary cost not to learn the consequences of their decisions.

Grossman and Van Der Weele's (2017) model introduces a trade-off between self image and

self interested choices. They investigate situations in which the information is about the welfare effects of the decision. They argue that knowing the welfare effects of a decision may make one feel compelled to act less selfishly. Avoiding ex-ante information about the consequences of the decision helps to protect the self-image and prevent feeling bad about causing a loss for someone else. In their model, people should avoid information only in the ex-ante perspective when information can potentially change the decision. Because in that case, making the selfish decision while knowing the consequences can cause the decision maker to act in a certain way. They argue that higher willingness to be perceived in a desired way increases willingness to avoid information for individuals who are tempted to act selfishly.² I present results that support and extend Grossman and Van Der Weele's (2017) argument by showing that in reciprocal decisions information avoidance can be chosen to make more selfish decisions. I also show that in trust games, image concerns are more important than in dictator games and people avoid information even when there is no instrumental value of information.

In the next section, the design of the experiment is explained in detail. The rest of the paper is as follows: The third section documents the procedure and the details of the sample, the fourth section discusses the results and the fifth section concludes the paper by discussing the implications of the findings. In the appendices, details of the post-experiment questionnaire and supplementary analysis can be found.

1.2 Experimental Design

The general setup of this experiment includes variations of the standard trust game (Berg et al., 1995) to investigate information preferences in reciprocal decisions when participants face a conflict between monetary gain and moral constraints. The standard trust game (investment game, lost wallet) is a sequential game played in pairs. Players are assigned to roles of *trustors* and *trustees*

²In Grossman and Van Der Weele (2017) *homo economicus* types in the model are the extreme selfish types and are assumed to make preferences to only increase monetary payoffs. These people are assumed to have no concerns about their image. *High social types*, on the other hand, always make the welfare increasing decisions and always acquire information to reach superior decisions in terms of final distributions.

and the trustor decides whether to send a part of her initial endowment (if any) to the trustee. If the trustor decides to send, this amount is multiplied by three and transferred to the trustee. The trustee then decides how much to return out of the total amount he received (see Figure 1.1). In the standard trust game, both trustor and trustee make decisions with complete information about the rules and the endowments of both parties. This means when deciding whether to send and how much to return, players are able to calculate final distributions and make their decisions accordingly. By sending money to the trustee, trustors signal intentions and trust and reciprocity expectations to trustees. Trustees can choose whether to reciprocate or not, which is the main moral dilemma in the game. Although returned amounts vary under different settings, Johnson and Mislin (2011) report that the average returned amount is 37% of the total received amount. Variations in the returned amount when the endowment levels are different are attributed to inequality aversion. This means, people respond to unequal endowments in the standard trust game. By adding uncertainty on the initial levels of endowments, this paper investigates how information preferences can serve to circumvent the moral constraints in the trust game.

The design is mainly built on two manipulations of the standard trust game: participants' initial endowment levels are not unified and they do not have complete information on each other's endowment. Information is disclosed if the trustee decides to do so and two timings are introduced for this decision: before the returning decision is made *ex-ante* and after the decision is made *ex-post*.

At the beginning of the experiment participants are randomly paired and assigned to the roles of trustors and trustees and one of the two endowment levels, high (35 ECU) or low (15 ECU). Assignment to endowment levels are randomized on the participant level and participants are informed that each of them is assigned to one of the endowment levels with equal chance. When the game starts, participants learn their role and endowment level but not their partner's endowment level. For trustors, instructions and rules are the same in all treatments. They are informed that they do not have the chance to learn their trustee's endowment and information

preferences in any stage of the experiment. After learning their role and endowment, the trustors are asked to decide whether to send 10 ECU or not. The trustor's choice is a binary choice; they can send 10 ECU or 0 ECU. If the trustor decides to send, 10 ECU is multiplied by 3 and 30 ECU is received by the trustee. Trustors who chose not to send and their trustees are directed to another screen in which they answer unincentivized questions and received a participation fee at the end. The complete ignorance in trustors' decisions, provides a chance to investigate the motivations of trust. Since 10 ECU is lower than both of the potential endowment levels, it is not possible for trustees to make any assumptions about the trustor's endowment level. Thus, the trustors' sending decision signals nothing but *trust* to the trustees. After the trustors' decision, the trustees are assigned one of the three treatments and asked to decide how much to return. Because the endowments are heterogeneous in the game, making a decision to reach equal final distributions requires the trustees to know the trustor's initial endowment (for an overview of endowment levels and equalizing returns see Figure 1.2).

In the Hidden Information treatment, trustees are given the chance to decide whether to acquire information about their partner's endowment level or not ex-ante, and they are informed that if they decide to acquire such information, they have to incur a monetary cost. In this treatment,




Figure 1.1: Trust Game

Note: In the ordinary trust game, trustor starts the game with e_1 endowment level and trustee starts with e_2 . First, trustor decides how much to send to trustee, 0 or x . If trustor decides not to send anything both players end up with their initial endowments: e_1 and e_2 . If trustor decides to send x , it is multiplied by 3 and sent to trustee as $3x$. Then in the second stage, trustee decides how much to return to trustor between 0 and $3x$. If trustee decides to send y , trustor will receive y and have a final payoff of $e_1 - x + y$ and trustee will have a final payoff as $e_2 + 3x - y$.

the default option is *not* to acquire the information and it is costless.

In the Ex-Ante Avoidance treatment, similar to the Hidden Information treatment, trustees decide whether to acquire the information ex-ante. In this treatment, trustees are informed that they incur a monetary cost if they decide to *avoid* information. In this treatment, the default is to acquire the information and it is costless. In both the Hidden Information and the Ex-Ante Avoidance treatments, trustees decide whether to get the information about the trustors' endowment ($I = 1$) or not ($I = 0$) before deciding how much to return.

In the Ex-Post Avoidance treatment, trustees decide how much to return under ignorance, and then they are asked to decide whether to acquire information on the final distribution or not (ex-post). In this treatment, the default is to acquire the information and it is costless. In Ex-Post Avoidance, participants make the information decision after they decide how much to return without knowing the trustor's endowment. After the decision has been made and the information preferences are stated, a revision screen, which participants were not informed about in advance, pops up for the ones who receive information (for those both who choose to acquire information and who want to avoid but lost the auction). The possibility of revising the returned amount is introduced to investigate the effects of unwanted information.

In all treatments, the trustees' conditional returning amounts for the possible endowment levels of the trustors are elicited by using the strategy method. If they receive the information (if they win the auction in Hidden Information and lose in other treatments), the decision for the realized condition is implemented and they are informed about it (see Figure 1.3 for the timing of the game). If they do not receive the information, they are asked to decide another unconditional amount to return which is to be implemented for any endowment level of the trustor. Formulating the returning decision in this way gives the chance to collect informed and uninformed decisions in all treatments. Ex-Post Avoidance treatment enables me to collect revision decisions after they submit amounts that they want to return unconditionally. This gives the chance to investigate whether receiving information which they wanted to avoid has any effect on return (revised) decisions.

Introducing two different time settings for information preferences helps to investigate whether the motivation for the preference is use of information for the decision or potential psychological loss of learning one's own type. On the one hand, in both Hidden Information and Ex-Ante Avoidance treatments, information has instrumental value in the classical terms since it can be used for deciding how much to return. It can help to make superior decisions in terms of final distribution for both players and social welfare. On the other hand, without information the feeling of being compelled to act prosocially decreases and selfish actions can be justified by *not being aware* of the consequences. This impact of information is predicted by Grossman and Van Der Weele (2017). In *Ex-Post Avoidance*, trustees make the returning decision under ignorance and then state their preference on information. Thus, information has no instrumental value since the return decision had already made. However, it can still have a psychological impact on image utility and ignorance can be preferred strategically to protect self-image. Avoiding information about the consequences for the other party helps trustees to maintain a desired self-image. In the ex-post setting, information has a diagnostic value by showing one's own prosociality rather than an instrumental value as argued by Bodner and Prelec (2003) and Bénabou and Tirole (2011).

In all three treatments, trustees' WTP for information preferences (avoidance or acquisition) to be implemented is elicited. Trustees are given 1 ECU bonus to use for this purpose. To reveal




Figure 1.2: Returning Decision in the game

Note: After Player 1's decision, Player 2 can cooperate and return a part of the investment or defect and take all the investment for herself. Depending on $E1$ and $E2$, the equalizing amount of return differs as shown in parentheses.

true WTPs, a Becker-DeGroot-Marschak (BDM) auction is used (Becker et al., 1964; Broberg et al., 2007). After stating their preferences on acquiring the information or avoiding the information, in a separate screen trustees are asked to state the price that they are willing to pay for that decision to be implemented over 1 ECU bonus. A randomly chosen price is compared with the player's willingness to pay. When the auction resolves, the randomly chosen price is announced to the trustees. If the participant's WTP is greater than the randomly chosen price, the decision is implemented and he pays the game price and the rest of 1 ECU bonus is added to the final payoff. If a participant's WTP is lower, the decision is not implemented and he does not pay anything so the 1 ECU bonus is added to the final payoff.³ For those who do not want to change the default setting, I impute WTP of zero. Bids in the BDM auction gives a continuous measure of WTP to acquire / avoid information for each person. At the end of the experiment, final payoffs are realized and the post experiment questionnaire takes place. In this experiment, each person makes a single decision (sending or returning) to prevent priming effects which could emerge when multiple decisions are made by the same person.

The design that is explained above helps to answer the main questions of the study. By assigning participants to default information and ignorance, this experiment yields results on the effects of exogenous information in reciprocal decisions. Introducing two different time settings in which WTP is elicited for the information avoidance to be implemented provides documentation on information preferences and motivations for these preferences. Strategic ignorance is motivated by image concerns; learning the consequences makes the selfish option less desirable since it has negative signals about ones own characteristics. Bénabou and Tirole (2006, 2011) show that image concerns as a component of the individual utility function effects social decisions when there is an observer. This observer of the action is not necessarily another person, but oneself can observe the consequences and learn about own inner dispositions and traits. In the previous literature it

³At the time of the experiment, the monetary equivalent of 1 ECU is 1 TRY and a soft drink at the school cafeteria like tea, coffee or 2 bottles of water. To find an optimal bonus for BDM, participants' opinions in a discussion after the pilot session are taken into account.

has been well documented that ex-ante excuses might increase selfish behavior and decreases prosocial behavior substantially (Grossman and Van Der Weele, 2017; Dana et al., 2006, 2007; Regner, 2016; Broberg et al., 2007). Yet it is still not very clear whether the motivation is to prevent being compelled to act in a certain way, or to prevent receiving the negative signals that may potentially hurt self image. If the main motivation is to prevent being compelled in a certain way, in the Ex-Post Avoidance treatment information avoidance should not be observed since it has no instrumental value, but if it is to avoid the negative signals, then learning about consequences might still yield diagnostic utility and signal one's own characteristics. Comparison of ex-ante and ex-post ignorance enables me to pin down whether it is the instrumental value or the ex post diagnostic utility that is important.

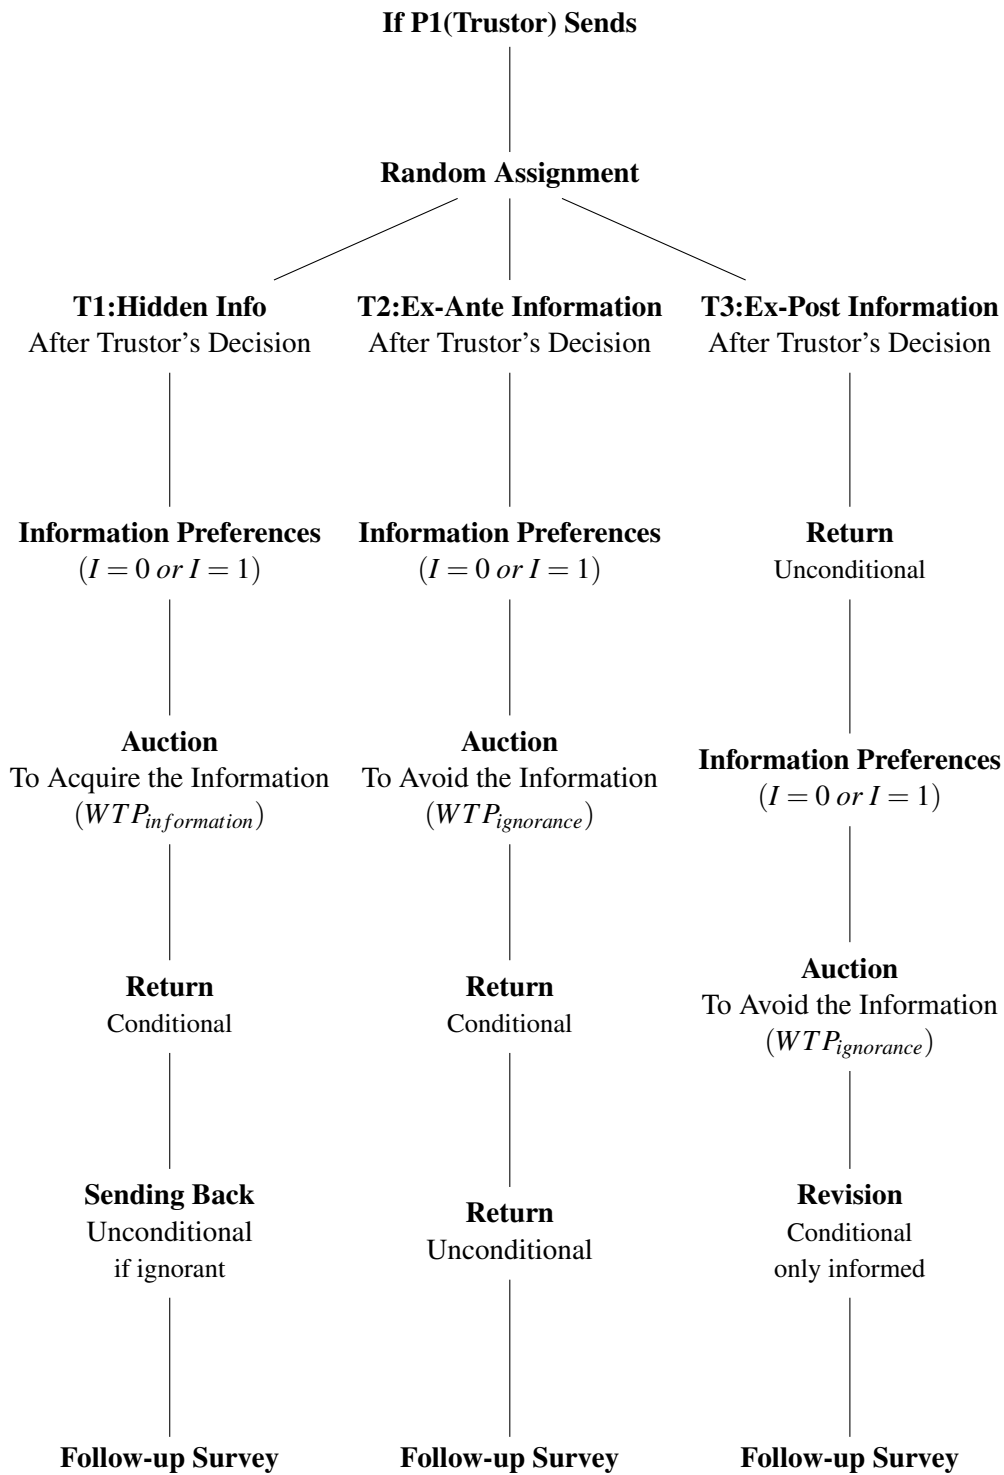


Figure 1.3: Experimental Design

Note: In all groups the endowments are assigned randomly. First, the trustor makes a decision on how much to send(x). This amount is tripled by the experimenter($3 \cdot x$). The trustee decides how much money to return (and so keep) out of $3x$. In the Hidden Info and Ex-ante Information treatments trustees decide whether to receive the information before deciding how much to return. In the Ex-Post Avoidance treatment, trustees make the decision on receiving information after they decide how much to return. A BDM auction is conducted to receive/avoid the information in all treatments. If the auction is lost, players' decision on revealing information won't be implemented. Note that if the trustor chooses not to send, both parties end up with initial endowments without proceeding the game.

1.3 Sample and Procedure

The experiment was designed by using Z-tree (Fischbacher, 2007) and conducted in Boğaziçi University—İstanbul. 520 undergraduate students were recruited via ORSEE (Greiner, 2015). Average payments were 40 Turkish Liras which is equal to 8.6 Euros at the time of the experiment. A pre-analysis plan was submitted before the data collection with the number "AEARCTR-000279". The pre-analysis plan for sample size was 450 participants. The data are collected between April 2018 and October 2018. Assignment to treatments is randomized within each session. Details about the sample can be seen in Table 1.1. 60% of the trustors decided to send 10 ECU which supports the sample size estimation for 0.8 power in the pre-analysis plan.⁴ For both endowment levels of trustors this fraction was almost the same. 61% of the participants were male and the average age was 21. Assignment to treatments was well balanced in terms of observable characteristics. Besides the behavioral measures, unincentivized decisions of trustors and trustees (who did not receive any amount from their partner) on how much to return and send were collected by using the strategy method. Questions on demographic information, behavioral covariates and social preferences such as reciprocity and risk preferences were included in the post-experiment questionnaire. Social and risk preferences were measured using questions from Falk et al.'s (2016) preference survey module.⁵

1.4 Results

To investigate motivations for the choices over information, I focus on trustees' decisions in this section. Only a subset of the pre-registered hypotheses are analyzed and discussed in this section.⁶ I present results in four steps: first, the effects of exogenously given information on reciprocity, second, information preferences and WTP for these, third, motivations for information preferences,

⁴Sample size estimations were made by using the findings from Johnson and Mislin's (2011) meta-study on trust games. The main aim was to reach a sample size to have significant variation in returned amounts between treatments.

⁵Details about the questionnaire can be found in the Appendix.

⁶I report the results for the rest of the pre-registered hypotheses and some exploratory analyses for the trustors' sending decision in Appendix A.

and fourth the effects of ex-post unwanted information on revised final distributions. I discuss each result and the relation to previous literature at the end.

1.4.1 Information and Ignorance

In the Hidden Information and Ex-Ante Ignorance treatments, participants are asked to chose whether they want to be informed prior to deciding how much to return. In the Ex-Post treatment, participants decide how much to return without any information. In the Hidden Information and Ex-Ante Avoidance treatments, trustees first decide how much they would return conditioned on the levels of initial endowment and are told that in case of receiving information, this conditional decision will be implemented, whereas if they do not receive the information, they are asked to state how much they want to return regardless of the trustor’s initial endowment. This structure of the design provides the opportunity to compare the exogenously given ignorance in the Ex-Post condition with the conditionally decided return amounts in the Hidden Information and Ex-Ante Avoidance.

Results show significant impact of ignorance on returned amounts across treatments. To test

Table 1.1: Summary descriptives by treatment

	Ex-Ante N=178	Ex-Post N=164	Hidden Info N=178
Age	20.8 (4.45)	21.1 (4.00)	20.5 (5.25)
Gender:			
F	65 (36.5%)	59 (36.0%)	78 (43.8%)
Endowment	25.4 (10.0)	24.1 (9.99)	25.3 (10.0)
Trustors who sent(%)	0.64	0.63	0.57
Share of trustors who send by endowment level:			
	Endowment		
	35 ECU	15 ECU	
Share of Senders:	73 (58.9%)	80 (58.8%)	

Note: Standars errors are in parentheses. Participants are assigned randomly assigned to treatments within each session. Almost 60% of the participants were male. “Trustors who sent” implies that on average 60% of the trustors decided to send 10 ECU. This proportion is almost the same for both endowment levels as can be seen in the bottom part of the table.

whether ignorance has any impact on returned amounts, I compare the ignorant (unconditional) returned amounts in the Ex-Post Avoidance with conditional returned amounts in the Hidden Information and Ex-Ante Avoidance. Returned amounts are significantly lower under exogenous ignorance than exogenous information. In the Ex-Ante Avoidance treatment, the average returned amount is 10.97(sd. = 6.69) and in the Hidden Information treatment 11.30 (sd. = 7.23, $p = 0.83$ for Welch's t-test for paired differences between Ex-Ante Avoidance and Hidden Information). The average returned amount is 6.02 (sd. = 5.4) in the Ex-Post Avoidance treatment and significantly lower than the average returns both in Hidden Information ($p < 0.001$ for paired differences between Ex-Post Avoidance and Hidden Information) and in Ex-Ante Avoidance ($p < 0.001$ for paired differences between Ex-Post Avoidance and Ex-Ante Avoidance).

Conditional decisions in Ex-Ante and Hidden Information are not different in terms of timing and structure. Returned amounts in these two treatments are not significantly different. Under exogenous ignorance, participants send lower amounts than the ones under exogenous information.⁷ This result supports the proposed effects of ignorance in the previous literature, when there is an uncertainty about the consequences of the decision, psychological cost decreases and selfish decisions increase substantially.

Result 1 *Exogenously given information increases reciprocity substantially whereas ignorance lowers it.*

Being ignorant about the consequences, can lower the psychological cost of decisions and alleviate the feeling of being compelled to act in a certain way as estimated and shown in the previous literature (Bénabou and Tirole, 2011; Grossman, 2014; Golman and Loewenstein, 2018; Grossman and Van Der Weele, 2017; Dana et al., 2006, 2007). In the present experiment, participants who do not know the consequences of their decisions return lower amounts than those who are informed ex-ante. This result supports the previous literature by showing that when there is no

⁷In the Ex-Ante Avoidance, informed (conditional) choices were weighted by probabilities: $pY_h + (1 - p)Y_l$ where participants were informed that both endowment levels are equally likely to be assigned. In the Ex-Post Avoidance, the default setting is ignorance and participants make an unconditional decision on how much to return.

exogenously given information or the chance to acquire information about the consequences of a decision, people can act selfishly without feeling bad about it. This first result shows that exogenously given information increases reciprocity compared to exogenously given ignorance which suggests that not knowing about what the decision can cause can decrease the psychological disutility of the selfish decision and leads to lower returned amounts.

1.4.2 Information Acquisition and Avoidance

In this experiment, both information and ignorance requires participants to incur a monetary cost. Besides its monetary cost, receiving information may have cognitive and psychological costs to players since they have to process the information. Ignorance in the Hidden Information treatment—or information in the Ex-Ante and Ex-Post Avoidance treatments—does not require subjects to incur any monetary costs. Thus, if information has no other effect than its instrumental use, people are expected to choose what is free in the treatments.

Before continuing with the details, the proportion of trustees who chose to acquire information and avoid information in all treatments is around 75% and 25% respectively. In no treatments does this share deviate significantly, which shows that exogenously given conditions of information (exogenous information or ignorance) does not have a significant effect on the fraction of information seekers and avoiders.

1.4.2.1 Willingness to Pay for Information and Ignorance

In the Ex-Ante Avoidance treatment, information can affect the decision and so it has instrumental value in the view of classical theory. In the Ex-Post Avoidance, the returning decision has already been made and information cannot affect the decision. However, if information has any psychological effect, it may persist even after the decision has already been made. In this case, avoidance can be observed ex-post as well.

Although the proportion of evaders does not differ significantly, to investigate whether in-

strumental value of information differs WTP to avoid information, I compare the WTPs in the Ex-Ante Avoidance and the Ex-Post Avoidance treatments. Figure 1.4 shows that average WTP to avoid information is 0.08 (sd. = 0.2), and is significantly lower than WTP to acquire information which is 0.23 (sd. = 0.27, $p < 0.001$). This shows that people who intend to receive information have a stronger motivation to do so. When calculating these averages, I added the WTPs for participants who did not change the default option as 0. Average WTP to avoid information among only those who chose to avoid in the Ex-Ante Avoidance and Ex-Post Avoidance treatments are 0.48 (sd. = 0.28) and 0.48 (sd. = 0.29) respectively. WTP for information among only those who wanted to receive information in the Hidden Information treatment is 0.39 (sd. = 0.26) Interestingly, WTP to avoid information does not vary significantly between Ex-ante and Ex-Post Avoidance treatments ($p = 0.29$).

Result 2

- a) *People avoid ex-ante information even when it can help to make better distribution decisions in terms of reciprocity*
- b) *People avoid ex-post information even when the distribution decision has already been made and it is not possible to change the returned amount*

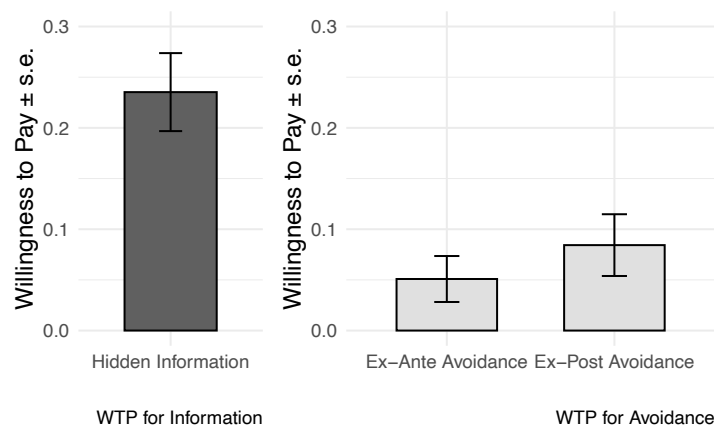


Figure 1.4: WTP for Information and Ignorance

Note: This figure shows that willingness to pay for information is significantly higher than willingness to pay for both ex-ante and Ex-Post Avoidance. Willingness to pay for avoidance is almost the same when it is decided ex-ante and ex-post.

Result 2a and 2b support the image utility argument by showing that the proportion of evaders and willingness to pay for avoidance do not differ even after a decision has been made. Not only when it can effect the decision but even after the decision has already been made, learning about the consequences of one's own actions causes information avoidance to maintain beliefs on desired self image. Information can be used to make better decisions in terms of social welfare but at the same time it can be strategically chosen/avoided to manage self image (Murnighan et al., 2001).

Result 2b contradicts Grossman and Van Der Weele (2017) who predict lower WTP for avoidance when instrumental value is lower or zero and provide evidence from a dictator game. Bénabou and Tirole (2006) on the other hand, emphasize the image utility of signaling to the observer—others or one's self—even when the actions are already done. Result 2b differs from the dictator game results that are documented in the literature (Dana et al., 2007; Grossman and Van Der Weele, 2017) by showing that in reciprocal games, not reciprocating is more hurtful than selfishness in dictator games since it means betraying another person's intentions. The main difference between these two is: if information about the consequences has more than instrumental value in terms of image utility then ignorance can be observed even when the decision has already been made. Results 2a and 2b show that in reciprocal decisions, the psychological effect of the information is greater since the social setting is more complex than dictator game decisions. In the next section, I present the mechanisms that lead to information acquisition and information avoidance.

1.4.3 Motivations for Information Preferences

To investigate the motivations for information preferences, the relation between information preferences and returned amounts are analyzed in Hidden Information and Ex-Ante Avoidance treatments. In the previous literature, this behavior is predicted by the models and supported by the experiments that individuals can choose not to learn about potential consequences and effects of

their decision not to look selfish to themselves and to others (Bénabou and Tirole, 2003; Bodner and Prelec, 2003; Grossman and Van Der Weele, 2017). Following the previous literature, in the present experiment, participants' information preferences are expected to be in line with their returned amounts; if they chose not to learn about the consequences ex-ante, returned amounts are expected to be lower and if they returned lower amounts they are expected to say "No" to information ex-post. Figure 1.5 provides supporting evidence to the models of image concerns as reported by Grossman and Van Der Weele (2017); Bénabou and Tirole (2011); Bodner and Prelec (2003): in all treatments individuals who return lower amounts are more likely to say 'No' to information (Avg. Returned Amounts for "Yes"=10.3 (sd. = 7.34), for "No"= 6.87 (sd. = 7.37), $p = 0.014$). In the Ex-Post treatment, 60% of participants who said "No" to information returned 0. This pattern is in line with the idea that information avoidance is a strategic tool to cover up the psychological cost of selfish decisions.

Even though WTP for ignorance does not differ significantly across treatments, to see whether it is related with returned amounts, I report the difference in returning decisions of participants who are willing to pay any positive amount to avoid information and those who are not willing to pay at all. In Figure 1.6, the vertical axis exhibits 2 groups of WTP: equal to 0 and greater than 0. It can be seen that participants who were willing to pay for ignorance are the ones who returns relatively lower amounts. Particularly in the Ex-Post Avoidance treatment, the proportion of 0 returns is much higher for participants who stated positive WTP for ignorance. This suggests a pattern of behavior: people who returned or intend to return lower amounts are more likely to avoid learning the consequences and are more willing to pay for it.

To investigate the association between information preferences and final returned amounts, I report the results from an OLS regression in which final returned amounts are explained with information decision ($I = 0$ or $I = 1$), willingness to pay (WTP) for information and avoidance, if their decision was implemented (Lost Auction), initial endowment level and other control variables,

$$Y_i = \alpha + \beta_1 I_i + \beta_2 WTP_i + \beta_3 IfLost_i + X_i + \varepsilon_i. \tag{1.1}$$

Table 1.2 explains the returned amounts by information preferences and WTP for them. In the table I report a significant relationship between final returned amounts and information preferences, particularly in Hidden Information and Ex-Post Avoidance treatments. As can be seen, WTP to acquire information is associated with higher reciprocity. In the Ex-Post Avoidance treat-




Figure 1.5: Return by Information Decision

Note: This figure exhibits the returned amounts by information choices, as can be seen people who do not want to receive information about the consequences of their decisions return lower amounts in all treatments. The fraction of people who did not return anything back is much higher for those who say ‘No’ to information.

ment, wanting to acquire the information ($I = 1$) has an increasing effect on returned amounts. Participants who chose to acquire information reciprocate more. WTP has a positive effect both in the Hidden Information and the Ex-Post Avoidance treatments. In the Ex-Post Avoidance treatment, those who avoid information have image concerns to some extent and they are they are the low social types (in the terminology of Grossman and Van Der Weele (2017)), they are not completely selfish. As can be seen in the table, participants who have higher WTP to avoid information are more likely to return higher amounts compared to ones who do not avoid information. The coefficient of the “Lost Auction” variable in the table shows if a person’s WTP was positive but lower than the game price, so the decision was not implemented. In the Hidden Information




Figure 1.6: WTP for Ignorance by Returned Amounts

Note: This figure shows that players who returned lower amounts are more likely to have positive WTP for ignorance. In Ex-Post treatment, a high proportion of participants who submitted positive WTP returned 0 before they made information decision. In Ex-Ante treatment, most of the participants with positive WTP returned relatively lower amounts considering the equalizing amount is minimum 10 for all combinations of endowment levels (see Figure 1.2).

treatment it refers to the case that participants who were willing to pay positive amounts to acquire the information, but they lost the auction so their decision was not implemented and they made the decision under ignorance. As it can be seen from regression results in Table 1.2, this is positively related with returned amounts in the Hidden Information treatment. The estimated coefficients in Table 1.2 show that returned amounts are affected by information preferences and initial endowments. Information preferences play an important role in the returned amounts, particularly in the Hidden Information and Ex-Post Avoidance treatments. In the Ex-Ante Avoidance treatment, this relation is not as strongly observed.

Table 1.2: Final Returns

	<i>Dependent variable:</i>		
	Final Return		
	(Hidden Info)	(Ex-Ante)	(Ex-Post)
Receive Info. (1 or 0)	-0.332 (3.402)	2.105 (2.252)	8.692*** (3.000)
WTP	10.954** (5.187)	-3.370 (6.022)	15.442** (5.751)
Lost Auction	4.693* (2.703)	0.924 (2.284)	-2.840 (1.857)
Endowment (=35)	4.042 (2.612)	4.708*** (1.734)	-1.817 (1.754)
Gender (Male)	-6.175** (2.962)	-2.328 (1.953)	2.796 (1.944)
Controls	Yes	Yes	Yes
Observations	51	57	51

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Note: Standard errors are in parentheses. This table shows OLS regression for the final returned amounts. Each column represents a treatment and explanatory variables are the behavioral variables in the experiment and controls are unlisted background characteristics which were collected in the post experiment survey. WTP is willingness to pay for information decision to be implemented and it is between 0 and 1. Lost Auction is a binary variable which equals to 1 if participant's willingness to pay is positive and lower than randomly chosen game price, so the auction was lost. Control variables include demographics and post-experiment questionnaire which can be seen in Appendix B.

Taken together, these results support the argument for social types that are defined by Grossman and Van Der Weele (2017): extremely selfish—*homo economicus*—types return 0 and do not avoid learning the consequences, whereas low social types return slightly higher amounts and are more willing to avoid learning the final distribution they caused. These results show a clear motivation for the choice of information: people who are willing to pay to acquire information are high-social types and even when they cannot acquire the information, they return significantly higher amounts compared to those who want to avoid information. These results on the relation between choice of information and returned amounts show that information avoidance is motivated by monetary rewards; people avoid ex-ante information to make selfish decisions.

1.4.4 Effects of Unwanted Information

In the Ex-Post Avoidance treatment, after the decision has been made and information preferences are elicited, participants who wanted to avoid information but lost the auction or did not want to avoid in the first place, receive the information eventually and are then given the opportunity to revise the initially entered amounts to return. If the motivation for avoidance is not to feel compelled to act in a certain way, seeing the consequences can lead to revision. However, once the decision has been made, participants might not want to change the returned amounts to not lose from their final payoff.

Results in Figure 1.7 show that participants revise and increase returned amounts if they learn that their partner started with the low endowment level (15 ECU). Even though revisions change the distribution, mostly they do not lead to equal final payoffs. After revisions, the gap between the initial returned amounts and the amounts that equalize the final payoffs decreased substantially; the decrease is 10.3 for high and 16.5 for low endowed trustors ($p = 0.004$). The average revision to the low endowed trustors is 8.18 (sd. = 6.92), that is significantly higher than the average revision to the high endowed ones which is 4.37(sd. = 5.54) ($p = 0.003$).

Figure 1.7 is split in panels which shows the endowment levels of trustor (that is shown

when the information is revealed) in columns and whether the trustee lost the auction (wanted to avoid information but could not). As can be seen in the figure, people who wanted to avoid information but lost the auction and received the information unwillingly (the upper two panels) make lower revisions after the information is received. On the other hand, people who wanted to receive the information (the lower two panels) revise to a higher extent for both income levels of trustors. The first column of this figure shows the revisions to low-endowed trustors. Revised amounts to the low-endowed trustors are higher compared to those to the trustors with high endowment levels.

***Result 4** If the auction is lost and information is unwillingly received ex-post, people tend to revise their decision to reduce inequality. However, those who avoid ex-post information are less likely to revise their initial returned amounts compared to those who wanted to receive it.*

In the revision decisions, information gains an instrumental value when people can use it in revisions. The revisions show that information preferences are shaped by the initial intention of

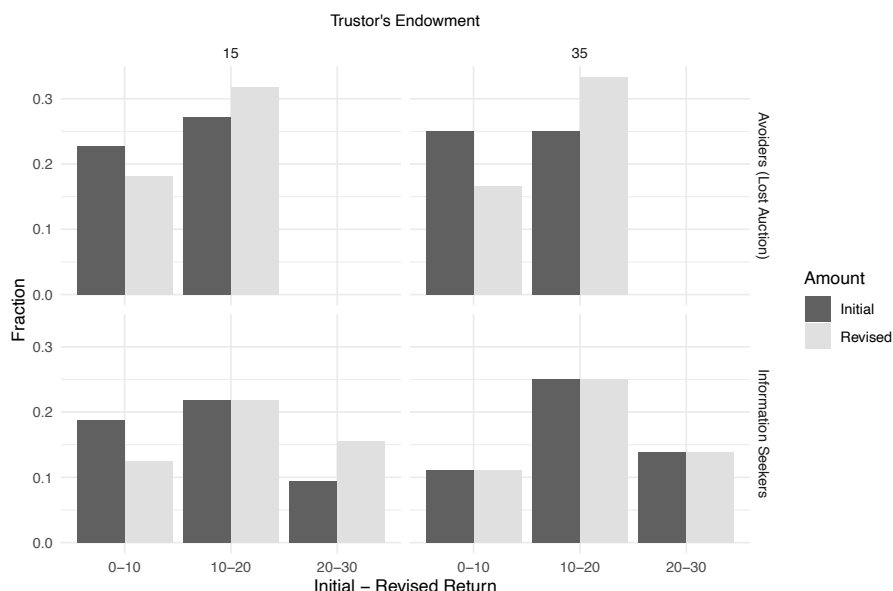


Figure 1.7: Conditional Revisions

Note: This figure shows revisions of participants who lost and did not lose but preferred to get the information. Participants who submitted positive amounts to avoid but could not win the auction revise lower amounts compare to those who wanted to get information. However, revised amounts are higher when the trustor had low endowment (15) compared to revision to trustors with high endowment (35).

people; if they intend to make selfish decisions, ex-post information has a smaller effect compared to those who are willing to acquire ex-post information. Although on average revisions are made to reduce inequality, those that avoided information are less likely to make revisions, and instead they keep their initial decisions (which on average are less reciprocal and quite selfish). Non-revising behavior of those who want to avoid information may be driven by the objective that they wanted to receive the amount they stated in the first place and this desire is stronger than guilt or regret induced by the information. This result contradicts Grossman and Van Der Weele (2017) by showing that people avoid information even when it is not instrumental and this behavior is related to lower reciprocity. This introduces a new mechanism for information avoidance in reciprocal decisions: people can avoid information about past decisions even when the decision was made under ignorance. All results together, provide evidence for the effects of information on one's self image; knowing that the past decision caused less reciprocal distributions and let someone down by not reciprocating their trust is something people would like to avoid, can motivate them to incur a monetary cost to avoid this information.

1.5 Concluding Remarks

In today's world, information is very easy to reach and is sometimes imposed on people even when it is not asked for. This gives a new direction to behavioral studies: how do people attempt to control the flow of information that they are exposed to? For decades, effects of different levels and asymmetry of information have been discussed and many economic and social inefficiencies are considered to be related to problems regarding information. For example, misreporting is assumed to occur when one has private information and can misuse it, and in the principle-agent frameworks, principals may receive false signals about agents which can cause suboptimal outcomes. Many economic models incorporate morality as a factor preventing people from engaging in these inefficiencies. In the present paper I show that by avoiding information about consequences, people can circumvent moral constraints. I provide evidence for the relation between preferences about

information and the change in moral decisions by documenting results from a reciprocal game.

The reciprocal setting in this experiment is a rich moral environment, as it introduces the chance to respond to someone else's behavior. In many economic theories, reciprocity is considered mainly in forms of distribution and intentions (Sobel, 2005). When making decisions in such settings, having equal payoffs often requires an individual to forego self utility which creates a trade-off between moral values and opportunities. Since this mechanism is relevant to many social interactions, reciprocity is assumed as a key concept in forming social capital (Glaeser et al., 2000), and has fundamental implications in many economic settings (Bartling et al., 2018; Bohnet and Huck, 2004; Berg et al., 1995). I use a reciprocal setting to investigate the effects of information and choices over information on reciprocity. I show that when information about the consequences of a decision for another person is exogenously given, people make less selfish choices. This result shows the importance of information in reciprocal decisions that people face in many different contexts.

One of the underlying mechanisms that may influence the information choice and distribution decision is the concerns about inequality. Some research is devoted to investigating inequality aversion in reciprocal decisions by using trust games. Glaeser et al. (2000) show that trustors with lower endowments are less likely to send a part of their endowments and interpret this as an indication for reciprocity expectation from trustees. Xiao and Bicchieri (2010) provide supporting evidence from a trust game in which endowments are heterogeneous, participants with high endowment level return higher amounts. Rodriguez-Lara's (2018) results, in contrast to Xiao and Bicchieri (2010), show that inequality aversion is not very effective in a trust game with endowment heterogeneity. Different than Xiao and Bicchieri (2010), Rodriguez-Lara (2018) uses a within subject design to allow seeing the difference between two endowment levels, which may create experimenter demand effects on the decisions. Lower effects of inequality aversion are observed particularly when inequality aversion is introduced in a trade off between efficiency (Kritikos and Bolle, 2001; Engelmann and Strobel, 2004) or altruism (Chowdhury and Jeon, 2014). I use an

experiment in which opting-out from information about the consequences may cause unequal distribution of final outcomes. I show that when people face information unwillingly, they revise their initial decisions on how much to return for achieving a more equal distribution at the end. This behavior is particularly strong when they learn that their partner started with the low endowment level. This result supports the argument that people may have concerns about inequality in terms of reciprocating to a previous decision in reciprocal decisions (Glaeser et al., 2000; Xiao and Bicchieri, 2010).

In this paper, I show that when people are given the chance to avoid information about the consequences, reciprocity decreases substantially whereas when the information is exogenously given or self-selected, reciprocity is higher and leads to less inequality in the final distribution of outcomes. These results help understanding the dynamics of distribution decisions in trust games by introducing the role of information and information choice for inequality aversion. This paper documents a clear behavioral pattern between information avoidance and selfish decisions. An important contribution of this paper is the finding that in reciprocal settings, ex-ante and ex-post avoidance do not differ significantly. This shows that, in the reciprocal decisions, psychological effect of information about consequences is greater, it brings a disutility to the acquirer even after a decision has already been made and information has no instrumental value. These results extend the Grossman and Van Der Weele's (2017) predictions on information avoidance and self image concerns, by showing that when the decisions introduce a more morally complex situation than dictator games, image concerns may be higher and can lead to avoidance even when the information is not instrumental and the decision cannot be changed. The results of the present paper support the argument that decisions and actions signal one's own characteristics and inner traits to oneself and to others, which motivates distributive decisions (Bodner and Prelec, 2003; Bénabou and Tirole, 2011; Akerlof and Kranton, 2000).

With this study, I provide answers to questions that are relevant to many economic and social decisions that depend on what and how much people know and how such knowledge influences

interactions between individuals. People reciprocate more when exogenously given the information about consequences. Many decisions in daily life depend on reciprocal relations such as social interactions with other people, taxpayers' relation with government, charities collecting donations and local organizations which depend mostly on private contributions. Ariely et al.'s (2009) and Lacetera and Macis's (2010) experiments show that prosocial actions, such as being a blood donor and participating in prosocial activities for good causes e.g. donations to charities could be driven by *social image motivations*—to be liked or seen as a good person by others—. In this paper, I document evidence for *self image motivations*. Even when there is no observer other than one-self, knowing the consequences of a decision can decrease selfishness substantially. People choose to avoid information in a self-serving way to circumvent moral constraints when they are given the chance. This helps them to make self serving decisions, even the morally questionable ones, without feeling bad about the consequences.

Bibliography

- Akerlof, G. A. and Kranton, R. E. (2000). Economics and Identity. *Quarterly Journal of Economics*, 115(3):715–753.
- Ariely, D., Bracha, A., and Meier, S. (2009). Doing Good Or Doing Well? Image Motivation And Monetary Incentives In Behaving Prosocially. *American Economic Review*, 99(1):544–55.
- Arrow, K. J. (1974). *The Limits Of Organization*. WW Norton & Company.
- Attanasi, G., Battigalli, P., and Manzoni, E. (2015). Incomplete-Information Models Of Guilt Aversion In The Trust Game. *Management Science*, 62(3):648–667.
- Bartling, B., Fehr, E., Huffman, D., Netzer, N., et al. (2018). The Causal Effect Of Trust. Technical report, Department of Economics-University of Zurich.

- Becker, G. M., DeGroot, M. H., and Marschak, J. (1964). Measuring Utility By A Single-Response Sequential Method. *Systems Research and Behavioral Science*, 9(3):226–232.
- Bénabou, R. and Tirole, J. (2003). Self-Knowledge And Self-Regulation: An Economic Approach. *Psychology of Economic Decisions*, 1:137–167.
- Bénabou, R. and Tirole, J. (2004). Willpower And Personal Rules. *Journal of Political Economy*, 112(4):848–886.
- Bénabou, R. and Tirole, J. (2006). Incentives and Prosocial Behavior. *American Economic Review*, 96(5):1652–1678.
- Bénabou, R. and Tirole, J. (2011). Identity, Morals, And Taboos: Beliefs As Assets. *Quarterly Journal of Economics*, 126(2):805–855.
- Berg, J., Dickhaut, J., and McCabe, K. (1995). Trust, Reciprocity, And Social History. *Games and Economic Behavior*, 10(1):122–142.
- Bodner, R. and Prelec, D. (2003). Self-Signaling And Diagnostic Utility In Everyday Decision Making. *Psychology of Economic Decisions*, 1:105–126.
- Bohnet, I. and Huck, S. (2004). Repetition and Reputation: Implications for Trust and Trustworthiness When Institutions Change. *American Economic Review*, 94(2):362–366.
- Bolton, G. E. and Ockenfels, A. (2000). ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review*, 90(1):166–193.
- Broberg, T., Ellingsen, T., and Johannesson, M. (2007). Is Generosity Involuntary? *Economics Letters*, 94(1):32–37.
- Cappelen, A. W., Hole, A. D., Sørensen, E. Ø., and Tungodden, B. (2007). The Pluralism Of Fairness Ideals: An Experimental Approach. *American Economic Review*, 97(3):818–827.

- Chowdhury, S. M. and Jeon, J. Y. (2014). Impure Altruism or Inequality Aversion?: An Experimental investigation Based on Income Effects. *Journal of Public Economics*, 118:143–150.
- Dana, J., Cain, D. M., and Dawes, R. M. (2006). What You Don't Know Won't Hurt Me: Costly (But Quiet) Exit In Dictator Games. *Organizational Behavior and Human Decision Processes*, 100(2):193–201.
- Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting Moral Wiggle Room: Experiments Demonstrating An Illusory Preference For Fairness. *Economic Theory*, 33(1):67–80.
- Ehrich, K. R. and Irwin, J. R. (2005). Willful Ignorance in the Request for Product Attribute Information. *Journal of Marketing Research*, 42(3):266–277.
- Eil, D. and Rao, J. M. (2011). The Good News-Bad News Effect: Asymmetric Processing Of Objective Information About Yourself. *American Economic Journal: Microeconomics*, 3(2):114–138.
- Engelmann, D. and Strobel, M. (2004). Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments. *American Economic Review*, 94(4):857–869.
- Falk, A., Becker, A., Dohmen, T. J., Huffman, D., and Sunde, U. (2016). The preference survey module: A Validated Instrument for Measuring Risk, Time, and Social Preferences. *Netspar discussion paper*.
- Fehr, E. and Schmidt, K. M. (1999). A Theory Of Fairness, Competition, And Cooperation. *Quarterly Journal of Economics*, 114(3):817–868.
- Fischbacher, U. (2007). Z-Tree: Zurich Toolbox For Ready-Made Economic Experiments. *Experimental Economics*, 10(2):171–178.
- Ganguly, A. and Tasoff, J. (2017). Fantasy And Dread: The Demand For Information And The Consumption Utility Of The Future. *Management Science*, 63(12):4037–4060.

- Glaeser, E. L., Laibson, D. I., Scheinkman, J. A., and Soutter, C. L. (2000). Measuring Trust. *Quarterly Journal of Economics*, 115(3):811–846.
- Golman, R., Hagmann, D., and Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1):96–135.
- Golman, R. and Loewenstein, G. (2018). Information Gaps: A Theory of Preferences Regarding the Presence and Absence of Information. *Decision*, 5(3):143.
- Greiner, B. (2015). Subject Pool Recruitment Procedures: Organizing Experiments With Orsee. *Journal of the Economic Science Association*, 1(1):114–125.
- Grossman, Z. (2014). Strategic Ignorance And The Robustness Of Social Preferences. *Management Science*, 60(11):2659–2665.
- Grossman, Z. and Van Der Weele, J. J. (2017). Self-Image and Willful Ignorance in Social Decisions. *Journal of the European Economic Association*, 15(1):173–217.
- Johnson, N. D. and Mislin, A. A. (2011). Trust Games: A Meta-Analysis. *Journal of Economic Psychology*, 32(5):865–889.
- Köszegi, B. (2006). Ego utility, Overconfidence, and Task Choice. *Journal of the European Economic Association*, 4(4):673–707.
- Kritikos, A. and Bolle, F. (2001). Distributional Concerns: Equity-or Efficiency-Oriented? *Economics Letters*, 73(3):333–338.
- Lacetera, N. and Macis, M. (2010). Do All Material Incentives For Pro-Social Activities Backfire? The Response To Cash And Non-Cash Incentives For Blood Donations. *Journal of Economic Psychology*, 31(4):738–748.
- Loureiro, M. L., McCluskey, J. J., and Mittelhammer, R. C. (2002). Will Consumers Pay A Premium For Eco-Labeled Apples? *Journal of Consumer Affairs*, 36(2):203–219.

- Murnighan, J. K., Oesch, J. M., and Pillutla, M. (2001). Player Types And Self-Impression Management In Dictatorship Games: Two Experiments. *Games and Economic Behavior*, 37(2):388–414.
- Onwezen, M. C. and van der Weele, C. N. (2016). When Indifference is Ambivalence: Strategic Ignorance about Meat Consumption. *Food Quality and Preference*, 52:96–105.
- Reczek, R. W., Irwin, J. R., Zane, D. M., and Ehrich, K. R. (2018). That’s Not How I Remember It: Willfully Ignorant Memory for Ethical Product Attribute Information. *Journal of Consumer Research*, 45(1):185–207.
- Regner, T. (2016). Reciprocity Under Moral Wiggle Room: Is It A Preference Or A Constraint? *Experimental Economics*, (4):1–14.
- Rodriguez-Lara, I. (2018). No Evidence Of Inequality Aversion In The Investment Game. *PLoS One*, 13(10):e0204392.
- Sobel, J. (2005). Interdependent Preferences And Reciprocity. *Journal of Economic Literature*, 43(2):392–436.
- Stigler, G. J. (1961). The Economics Of Information. *Journal of Political Economy*, 69(3):213–225.
- Xiao, E. and Bicchieri, C. (2010). When Equality Trumps Reciprocity. *Journal of Economic Psychology*, 31(3):456–470.

1.A Supplementary Analysis

In this section, I document findings that are pre-registered but not listed as main results. In addition to pre-registered findings, I provide an explanation to sending decisions on the trustors. Although these results are mostly quite similar to the main results, they help to understand the underlying mechanisms for information preferences better.

1.A.1 Information Avoidance and Reciprocity

Although the relation between information preferences and reciprocity is listed among the main results, Figure 1.A.1 shows this result in detail for each treatment. As can be seen in the figure, in all treatments people who reciprocate less are more likely to say “No” to receiving information ($p = 0.014$). This relation is present even when the decision has already been made. In the Ex-Post Avoidance treatment people who reciprocated less are significantly more likely to avoid information. Results provide evidence for the motivation to choose ignorance. In all treatments, people who avoid information return lower amounts compared to those who want to receive information. Ignorant decisions are closer to informed returns to the high-endowed trustors and are significantly lower than the informed decisions to low-endowed trustors ($p < 0.001$).

Figure 1.A.3 shows how much time people spend on making the information choice. They make a binary choice (“Yes” or “No”) in all three treatments before stating their WTP for this choice to be implemented. As can be seen in the figure, people spend more time to make this choice in the ex-ante conditions (Hidden Information and Ex-Ante Avoidance) compared to ex-post choice (Ex-Post Avoidance) ($p = 0.003$). This shows that deciding whether they want to receive information or not is a more difficult decision in the ex-ante perspective compared to the ex-post perspective.

1.A.2 Ex-Post Information and Revisions

In this section, I document the findings on revisions which are made in the Ex-Post Avoidance treatment after the initial decisions on how much to return is made. Figure 1.A.2 shows the revised returns with how the gap between the final amount and the equalizing amount.

If people avoid information to not feel guilty about it, they are expected to revise this decision after learning the consequences willingly or unwillingly. However, those who willingly want to receive that information may have two motivations: they may be homo economicus types as defined by Grossman and Van Der Weele (2017) and do not have image concerns or bad feel-

ings when doing something selfish, or they may be social types who always want to make social decisions and want to learn if they manage to do so.

Table 1.A.1 supports this finding by showing the significantly increasing effect of losing the auction on the gap between the equalizing and final returned amounts. This shows that, people who wanted to avoid information ex-post however did not have sufficiently high WTP to win the auction and caused a 4-unit higher gap between the equalizing return and the final return compared to those who willingly received the information.

1.A.3 Trustors' Sending Decision

In this game, trustors make the sending decision without having the chance to receive information about their trustees' initial endowment levels. In this case, what is the motivation to *trust* could be due to risk preferences. As mentioned earlier, no matter what their endowment level is, 60% of the trustors decide to send. Table 1.A.2 explains the likelihood of sending money to the trustee. As can be seen in the table, Risk preferences significantly explains this behavior and the higher scores of self-reported risk loving increases the likelihood of sending money significantly. Another unobservable characteristic that estimates this relation significantly is the WTP for good causes (charities etc.) which is self reported in the post-experiment questionnaire.

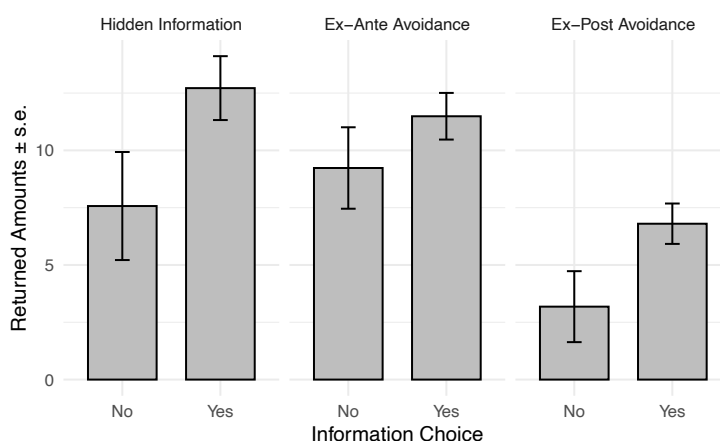


Figure 1.A.1: Returned Amounts and Information Preferences

Note: This figure shows average returned amounts by treatments and information preferences. On the y axis, Returned Amounts +/- standard error is reported.




Figure 1.A.2: Equalizing Amounts - Revised Returns

Note: This figure shows the difference between the amount which can equalize final payoffs for each player and the actual returned amount after revision by the player's endowment and the conditional endowment of the partner. Light gray bars show the conditional decisions (low or high endowed partner) of low endowed trustors and dark gray bars show for high endowed trustors. *Bars are se.s.*




Figure 1.A.3: Time spent on information choice.

Note: This figure shows how much time (in seconds) was spent when deciding whether to receive the information or not by treatments. As can be seen, when the information choice is made ex-ante, people spend significantly more time compared to ex-post choice.

Table 1.A.1: Average Gap Between Equalizing Amounts and Final Returns after Revision

<i>Dependent variable:</i>	
Gap After Revision	
Endowment (=35)	0.538*** (0.094)
WTP	-4.159 (4.637)
If Lost	4.127** (1.975)
Controls	Yes
Observations	51
F Statistic	3.660*** (df = 13; 37)
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

Table 1.A.2: Trustors' Decision on Sending 10ECU or not

<i>Dependent variable:</i>	
	Send
Endowment (=35)	−0.032 (0.287)
Gender (=Male)	−0.316 (0.312)
Risk Pref.	0.355*** (0.072)
WTP for Good	0.164** (0.064)
Other Controls	Yes
Observations	260
$dydx_{risk}$	0.070***
$dydx_{wtp}$	0.033**
Log Likelihood	−150.831
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

1.B Instructions

In this section, instructions in the experiment are included. The experiment was conducted in Turkish, the instructions here are the translated text to English.

1.B.1 Information about the roles and endowments

In the experiment there are two possible endowment levels: 15 and 35. Also there are two roles: X and Y. Player X always has the same options: Send 10 ECU or Send 0. After this decision, Player Y decides how much to return to Player X (if received any). At the beginning of the experiment you are assigned to these roles and endowments randomly. Now, in that screen you are given information on your role and endowment. Your role: X (Y) and Your Initial Endowment: 35 (15). Please use the buttons below to confirm that you understood your role and endowment level.

1.B.2 Information for trustors

You're assigned to role "Player X". As explained in the beginning of the game you have two options as "Send 10ECU" and "Send 0". If you send 10 ECU to your partner it will be multiplied by 3 by the experimenter, and received by Player Y as 30. Your final payoff depends on the Player Y's return decision after that. You are not informed about Player Y's initial endowment. Player Y can make a choice that affects whether he or she learns your endowment level. This decision might not be realized and you will not be informed about your partner's decision on learning your endowment level at any stage of this experiment.

1.B.3 Information for trustees

You are assigned to the role "Player Y". You start the game without being informed about Player X's endowment. Player X started the game without knowing your endowment as well. As explained at the beginning of the game in case Player X decides to send 10 ECU this amount will be multiplied by 3 by the experimenter, and you will receive 30ECU. You will be asked how much

to return over that amount to Player X. You will be given the chance to learn your partner's initial endowment level. Player X does not have the chance to learn yours. Player X will not be informed whether you choose to get information or not. At the end of the game Player X will only learn her/his own final payoff.

1.B.4 Information decision of trustees in Hidden Information and Ex-Ante Avoidance

Your partner sent 10 ECU to you. That amount is multiplied by 3 by the experimenter and you received 30 ECU. Before you decide how much to return, you can decide whether you want to learn Player X's endowment or not by using the buttons below. Your partner will not be informed about that decision. Please submit your decision: "Reveal information", "Do not reveal the information"

1.B.5 Information decision of trustees in Ex-Post Avoidance

You choose and confirmed the amount you want to return :[Returned Amount]. You will learn your partner's initial and final endowment levels on the next screen. But you have the chance to stop that information. You can decide whether you want to learn Player X's endowment or not by using the buttons below. Your partner will not be informed about your decision. Please submit your decision: "Reveal information", "Do not reveal the information"

1.B.6 WTP elicitation with BDM

1.B.6.1 Hidden Information

You decided that you want the information. There is a price for learning player X's endowment. You get a bonus of 1 ECU that you can use to pay this price (or keep for yourself). How much are you willing to pay for the information about player X's endowment?

Please enter a number between 0 and 1. The price will be a random number between 0 and 1, and if your willingness to pay is above that realized price, your decision be implemented and

you pay the price.

If your willingness to pay is below the realized price, you keep all of your bonus but won't learn player X's endowment level.

1.B.6.2 Ex-Ante Avoidance

You decided that you do not want the information. There is a price for not learning player X's endowment. You get a bonus of 1 ECU that you can use to pay this price (or keep for yourself). How much are you willing to pay for not learning the information about player X's endowment?

Please enter a number between 0 and 1. The price will be a random number between 0 and 1, and if your willingness to pay is above that realized price, your decision be implemented and you pay the price.

If your willingness to pay is below the realized price, you keep all of your bonus but will learn player X's endowment level.

1.B.6.3 Ex-Post Avoidance

You decided that you do not want the information. There is a price for not learning player X's endowment. You get a bonus of 1 ECU that you can use to pay this price (or keep for yourself). How much are you willing to pay for not learning the information about player X's endowment?

Please enter a number between 0 and 1. The price will be a random number between 0 and 1, and if your willingness to pay is above that realized price, your decision be implemented and you pay the price.

If your willingness to pay is below the realized price, you keep all of your bonus but will learn player X's endowment level.

1.B.7 Conditional returning decision

1.B.7.1 In Hidden Information and Ex-Ante Avoidance

For your partner Player X there are two possible endowment levels with equal chances. Please enter the amount you would return for each of the two possible levels of the endowment:

15:

35:

If your price to learn [to not learn] Player X's endowment level is [not] higher than the randomly chosen price, you will [will not] learn the true endowment level of Player X, and only the amount you entered for this endowment level in this screen will be implemented.

If your price is lower [higher] than the randomly chosen price, you will be asked to enter another amount to return which will be realized regardless of the initial endowment level of Player X.

1.B.7.2 In Ex-Post Avoidance

For your partner Player X there are two possible endowment levels with equal chances. Please enter the amount you would return for each of the two possible levels in case you knew it:

15:

35:

If your price to not learn Player X's endowment level is higher than the randomly chosen price, you will not learn the true endowment level of player X, and the first amount you wanted to return will be realized. (THE INITIAL RETURNED AMOUNT)

If your price is not higher than the randomly chosen price, you will learn the true endowment level of player X, and the amount you entered for this endowment level in this screen will be implemented.

1.C Post-Experiment Questionnaire

1. Gender
2. Age
3. Imagine that you are living in a society with income levels from 0 to 10, where would you place your income in that scale?
4. How do you see yourself: are you a person who is generally willing to take risks, or do you try to avoid taking risks? Please use a scale from 0 to 10, where a 0 means you are "completely unwilling to take risks" and a 10 means you are very willing to take risks".
5. Please again indicate your answer on a scale from 0 to 10, where 0 means you are "completely unwilling to do so" and a 10 means you are "very willing to do so". You can also use any numbers between 0 and 10 to indicate where you fall on the scale, like 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10.
 - (a) How willing are you to punish a person who treats you unfairly even if there may be costs for you?
 - (b) How willing are you to punish a person who treats others unfairly even if there may be costs for you?
 - (c) How willing are you to give to good causes without expecting anything in return?
6. Please indicate your answer on a scale from 0 to 10. A 0 means "does not describe me at all" and a 10 means "describes me perfectly". You can also use any numbers between 0 and 10 to indicate where you fall on the scale, like 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10.
 - (a) When someone does me a favor I am willing to return it.
 - (b) If I am treated very unjustly, I will take revenge at the first occasion, even if there is a cost to do so.

(c) I assume that people have only the best intentions.

Chapter 2

Strategic Curiosity:

An Experimental Study of Curiosity and Dishonesty

Abstract

In this study, we provide experimental evidence on a novel phenomenon concerning information preferences: people strategically collect additional non-instrumental information to justify morally questionable decisions. We conduct a virtual dice-rolling experiment where participants roll a dice and self-report the outcome of the first roll for monetary rewards. In this setting, we vary the extent to which participants can continue rolling the dice before reporting as well as the displayed content of those additional roll-outcomes. We document that people systematically roll the dice more—*are more curious*—when tempted to misreport. We find that curiosity is positively correlated with the size of the lie. However, contrary to previous studies, we observe no variation

⁰This work is partially supported by the Research Council of Norway through its Centres of Excellence Scheme, FAIR project No 262675. We are grateful to Catalina Franco, Thomas De Haan, Stefan Meissner, Lars Ivar Oppedal Berge, Nina Serdaravic, Hallgeir Sjøstad, Oda Sund, Erik Ø. Sørensen, Helge Thorbjørnsen, Bertil Tungodden and the seminar participants at Norwegian School of Economics during the Behavioral Economics Course with Lise Vesterlund and FIBE 2019 Conference for their helpful comments and contributions. We would like to thank Centre of Ethics at Norwegian School of Economics for financial support.

in dishonesty across treatments regardless of the possibility to collect additional non-instrumental information. This study provides new insights into how individuals actively shape their information environment in pursuit of self-interest.

Keywords: Information seeking; curiosity; dishonesty; lying cost

JEL-Classifications: C99, D82, D83, D91

2.1 Introduction

Classical theories of information economics define information as a tool that individuals use to reach superior decisions (Stigler, 1961). Under this theory, information is valuable only from an ex-ante perspective and if it can be used to make better decisions. However, evidence from behavioral research shows that people can avoid relevant information to create a *moral wiggle room* in which morally questionable decisions can be excused (Dana et al., 2007; Golman et al., 2017). Avoidance may not be the only tool to achieve moral wiggle room. Collecting additional information can be used strategically to interpret facts toward one's own preferences. Even though theoretical and experimental research shows that people tend to strategically avoid information to excuse self-interested choices (Dana et al., 2007; Grossman and Van der Weele, 2017; Golman et al., 2017), little research has examined to what extent individuals *collect* additional information to excuse the pursuit of self-interest.¹ We fill this gap in the literature by providing an experimental analysis of the tendency to collect more information when confronted with the temptation to misreport. We conducted a one-shot dice-rolling game (Shalvi et al., 2011; Fischbacher and Föllmi-Heusi, 2013) in which participants roll a fair virtual dice and report the outcome of the first roll for monetary rewards, with higher reported numbers resulting in higher payments. We vary whether people can collect non-instrumental information and the content of the additional information. We implement these treatment variations by restricting how many times people can roll the dice before reporting, and whether the dice displays numbers or random figures after the first roll.

Similar to avoiding information that makes it difficult to excuse selfish decisions (Grossman, 2014), we find evidence that curiosity is driven by a desire to justify selfish behavior. Our main finding is that non-instrumental information is collected strategically, implying that people

¹We use the term additional information to refer information that strictly speaking, is superfluous to the reporting decision but that can be useful to individuals when trying to justify dishonest behavior. We use 'additional' instead of 'superfluous' because we introduce a new goal for the information collection: justifying dishonesty. Hence, the information can be useful in justifying dishonesty even though the information is not instrumental according to classical theories.

are *strategically curious*.² Contrary to the previous literature, we find no variation in dishonesty across treatments despite the possibility to collect additional information and the content of this information. With our explicit focus on deliberate decisions to acquire additional non-instrumental information when tempted to misreport, we contribute to the research on information and moral decisions by showing that curiosity can also be driven by a desire to justify selfish behavior.

In a pre-registered experiment, we recruited 1580 US participants on Amazon Mechanical Turk (mTurk). In the base treatment (*Single Roll*), participants roll the dice once and then report the outcome. In this condition, people can continue to roll the dice after reporting while knowing that they cannot change their report. To study how people search for additional information when tempted to misreport, we introduce three variations to identify causal effects of different information environments. In the *Three Rolls* treatment, the number of rolls is limited to three, and people do not have a chance to roll less or more than three times before reporting the outcome of the first roll. In the other two treatments, participants can roll as many times as they want to before reporting. In the *Multiple Numbers* treatment, the dice's sides always display numbers, whereas in the *Multiple Figures* treatment, the dice displays random figures after the first roll. Since only the outcome of the first roll should be reported, additional rolls have no instrumental value in the case of honest reporting.

Our results provide evidence that people are systematically more curious when collecting additional information can help justify dishonesty. As this is an observed game, we can disentangle whether collecting additional information that is related or unrelated to the outcome leads to higher misreporting. We find that people who observe lower outcomes in the first roll are more likely to roll more times before reporting in the *Multiple Numbers* and the *Multiple Figures* treatments. People in the *Multiple Numbers* treatment observe numeric outcomes whereas those in the *Multiple Figures* see only non-order symbols in the additional rolls. The average number of rolls is not significantly different in these two treatments. In the *Multiple Numbers* treatment partici-

²We call the behavioral phenomenon that people collect related and unrelated additional information to relax moral constraints that are at odds with their self-interests *strategic curiosity*.

pants roll 4.8 times on average, whereas in the *Multiple Figures* treatment, they roll 4.9 times on average. This result indicates that people acquire additional information—not only to search for justifications—but also to distract themselves from moral conflict.

We find that dishonesty does not respond to the availability of additional information. The average *size of the lie*—the distance between the reported number and the actual outcome—is 0.50 units for all participants and does not significantly vary across treatments. However, we document a positive relationship between information collection and dishonesty. Rolling the dice additional times is correlated with higher size of the lie and this relationship is particularly strong among dishonest participants.

Our supplementary results show that even after the report is submitted, people continue collecting information ex-post in the *Single Roll* treatment. We find a significant difference in rolling behavior between ex-ante and ex-post rolling, when the outcome of the first roll is low—but no difference when the outcome of the first roll is high. This suggests that people are more curious when the additional information can be used to justify reporting dishonestly. Further evidence shows that people use curiosity to justify morally questionable decisions even after the decision has already been made. In this ex-post rolling, we observe that participants who misreported roll significantly more times than those who reported honestly. Dishonest reporters roll the dice 5.5 additional times after the first roll, whereas honest reporters roll it only 3.1 more times. Although not pre-registered, these findings support our main hypotheses and findings on the strategic use of curiosity to justify dishonesty.

This study contributes to research on preferences toward information in moral dilemmas, which can arise both before and after a decision has been made, and both when the decision-maker possesses incomplete and full information. Various motivations that shape information preferences have been documented previously, and curiosity is one of the prominent drivers of information acquisition. Loewenstein (1994) provides a review on curiosity and posits that demand for information is intrinsic; it is “appetite for knowledge.” This kind of curiosity directed toward all kinds

of information is defined as *epistemic curiosity* (Litman et al., 2005). We find that curiosity need not only arise from an intrinsic desire for information; it might also be driven by a strategic desire to justify selfish behavior or distract oneself from moral conflict. More recently, Golman and Loewenstein (2015) introduced the concept of “information gaps”, which refers to people’s desire to collect information to close the gaps between what is already known and what information is available. Eliaz and Schotter (2010) provide experimental evidence that individuals are willing to pay to receive information regarding the results of an intelligence test even when this information has no value in terms of achieving higher outcomes. We find that people acquire unnecessary information particularly when tempted to make morally questionable decisions. This extends the literature on curiosity by showing that curiosity may be motivated by strategic reasons rather than only an innocuous desire to collect information.

Information can play an important role in people’s self-image management when it has *diagnostic utility*, revealing people’s own moral type or disposition (Bodner and Prelec, 2003; Rabin, 1995). In this case, both avoiding and collecting information can help preserve a desired self-image, helping people to feel good about their abilities and traits (Golman et al., 2019). Our study expands this literature by showing that information acquisition can serve a similar function as information avoidance. When the temptation to misreport is present, people tend to acquire non-instrumental information. In support of this, we find that people tend to acquire non-instrumental information that is also unrelated to the moral decision. This shows that the curiosity might not only stem from the search for justifications, but also from the search for distractions from the moral conflict. Prior research documents that rational inattention can cause various behavioral biases like present bias and correlation neglect (Sims, 2006; Gabaix, 2019). We extend this research by showing that, even when information is not related to the decision or the outcome, it can serve as a tool for self-distraction and inattention, which can be used strategically to stick with certain decisions.

The rest of the paper proceeds as follows. Section 2 presents the details of the experimental

design. Section 3 explains the procedure and the details about the sample. Section 4 presents the results from our experiment. Section 5 concludes the paper.

2.2 Experimental Design

To investigate the relation between curiosity and dishonesty, we use modified dice-rolling game (Shalvi et al., 2011; Fischbacher and Föllmi-Heusi, 2013). The dice-rolling game has been widely used to study dishonesty in the previous literature (Abeler et al., 2019). The standard structure of the game is that participants roll a fair six-sided dice and report the outcome of the first roll for monetary rewards. In the instructions, participants are informed that higher reported numbers result in higher payments. Because participants roll the dice privately, the experimenter cannot infer whether an individual misreports the roll outcome but can only infer dishonesty on the aggregate level. Together with the payment structure, the privacy of the game provides those who roll low numbers with a monetary incentive to misreport by reporting a higher number than the one they rolled. In the case of honest reporting, the expected average reported number is 3.5, and outcomes are uniformly distributed on the integers 1 to 6. Using this game has several advantages for investigating dishonesty since the game is of a simple nature that is easy to understand for participants. Since the theoretical distribution is known, experimenters can detect overall dishonesty.

Because our research question required us to observe both the number of times participants roll the dice and the outcomes of the dice rolls, we used an online dice-rolling game. In our experiment, participants roll a virtual dice on a computer screen while a software records how many times a participant rolls the dice and the outcomes of each dice-roll. Participants are informed about this procedure in advance and are told that their payoff will only be dependent on the reported number and not the factual outcome. Participants then report the outcome of their first throw and receive payment according to their report (reporting 1 yields a payment of USD 0.5 and reporting higher numbers increase the payment with increments of USD 0.5).

Though observed games are becoming more common in the dishonesty literature (Abeler

et al., 2019; Gneezy et al., 2018), the observability of the reporting situation could potentially create an experimenter demand effect and affect participants' perceptions of the game. This could lower participants' level of dishonesty. However, the focus of our paper is to detect information preferences and strategic use of information collection. Using an online experiment requires using a virtual dice in the browser, because with a private dice we cannot collect information on how many times the dice has been rolled and the outcome of the first roll. Aware of the potential concerns about observability, participants are ensured that their choices remain anonymous. Before knowing about the dice-rolling game, participants are informed that they would not be rejected based on the submission of an incorrect answer. Before reading further instructions, all the participants had to answer correctly on questions related to the study's terms and conditions. In the following instructions, participants read about the dice-rolling game and the rules of their treatment condition. Another concern when using a virtual dice is that participants might suspect that the dice is not fair. To combat this concern, we allowed participants to freely roll the dice before knowing about the reporting task. We explicitly informed all participants that the dice was programmed to be fair. To avoid priming participants on numeric values while practicing, the sides of the dice displayed non-ordered and random symbols.

Using an observed version of the basic dice-rolling game enables us to investigate whether the outcome of the first roll affects the likelihood that participants roll more than once, and whether the distance between the observed first-roll outcome predicts information acquisition. This enables us to measure the exact size of the lie and what drives information acquisition, which is crucial for our research questions. By reaching a sample of 1,580 participants, we aimed to provide valid findings for strategic curiosity. In addition, the virtual set-up enables us to scale up the experiment and post the experiment on online platforms where participants can choose when and where to complete the experiment. The sample size enables us to run a well-powered study after making the estimations for at least 80% power (see Ay et al. (2019)).³

³Sample size is estimated with the mean values for reported die outcome from Shalvi et al. (2011). In the control group, participants are allowed to roll the die only once, whereas in treatment participants roll multiple times. We

2.2.1 Treatments

To study the relation between curiosity and dishonesty, we manipulate whether participants are able to choose how many times they can roll the virtual dice. Restricting participants' ability to collect additional information provides exogenous variation in the decision to collect additional information, which facilitates causal analyses between treatment groups. In total, we introduce four treatment variations to our dice-rolling setting.

2.2.1.1 Baseline

To establish a baseline, we implemented a *Single Roll* treatment where the participants only roll the virtual dice once before they submit their report. In the *Single Roll* treatment, the availability of additional information is (exogenously) restricted along with their ability to justify misreporting using additional outcome-related information. To obtain a proxy measure for pure (epistemic) curiosity in our setting, we allowed participants in the *Single Roll* treatment to continue to roll the dice as many times as they would like after they reported. Because participants could not change their report after submission, collecting additional information has no instrumental use for the decision.

Including our baseline, we provide a design in which we manipulate exogenous and endogenous information availability.

2.2.1.2 Exogenous availability of information

To investigate whether the amount of counterfactual information—without self-selecting to collect it—affects dishonesty, we limit the number of rolls prior to reporting in two of our treatments. In addition to the *Single Roll* baseline, we implemented a *Three Rolls* treatment where participants are forced to roll the dice three times before they submit their report. In this treatment, additional

used the reported numbers (so the earnings) in the control and treatment groups for our estimations. $\mu_{control}$ shows the average reported outcome in the control group whereas μ_{treat} shows in the treatment group. In condition *Single Roll* (control group) where only one roll is possible $\mu_{control} = 3.97$ and $\sigma_{control} = 1.56$. In the treatment where multiple rolls are allowed $\mu_{treat} = 4.45$ with $\sigma_{treat} = 1.59$.

outcome-related information is exogenously given to participants as they cannot proceed to the reporting page before the dice has been rolled exactly three times. This enables us to investigate whether observing additional outcomes-related information in itself increases misreporting and whether this effect is driven by the counterfactual outcomes they observe.

In *Single Roll* and *Three Rolls* treatments, the number of rolls is exogenously limited prior to reporting. This restriction helps us study how participants respond to this limitation and whether it affects dishonesty.

2.2.1.3 Endogenous availability of information

To study whether participants are strategically curious, we implemented two treatments that allow for endogenous information collection. In these two treatments, participants can roll the dice as many times as they want to, however the content of the dice is different after the first roll.

Multiple Numbers

In the *Multiple Numbers* treatment, participants choose how many times to roll the dice before reporting the first outcome. To investigate whether the content of the subsequent roll-outcomes matters for the decision to collect additional information, we add a description underneath the “roll” button that indicates the content of the next dice-roll. In the *Multiple Numbers* treatment, the description stated, “*Potential outcomes: Numbers from 1 to 6*”. Because participants must actively choose to roll the dice additional times to obtain additional outcome-related information, we can compare whether having access to this information affects misreporting differently by being endowed with such information. This allows us to also investigate whether those who have access to outcome-related (counterfactual) information systematically roll more when they observe low roll-outcomes in their first (actual) roll.

Multiple Figures

In the *Multiple Figures* treatment, participants choose how many times they roll the dice before reporting but, after the first roll, subsequent roll-outcomes display non-ordered symbols instead of numbers. Changing the displayed content of the sides of the dice restricts participants' access to additional outcome-related information that can (directly) help justify misreporting by reporting the best throw. Rolling the figures-dice only generates unrelated information. Yet, acquiring unrelated information can serve as a distraction from their moral standard to report honestly and therefore make misreporting less threatening to their moral self-view (Mazar et al., 2008). To be rationally inattentive as Gabaix (2019) proposes, participants can search for distraction by collecting unrelated information. The description underneath the "roll" button states "*Potential outcomes: Random symbols only*" to remind participants.⁴ In contrast to the *Multiple Numbers* treatment where participants know that additional roll-outcomes display (counterfactual) numeric information, participants in the *Multiple Figures* treatment know that rolling the dice additional times only generates unrelated symbolic information. This allows us to examine whether participants are more curious about the additional information when the information has more potential to justify misreporting and whether observing related information is more effective in justifying misreporting than observing unrelated information.

To eliminate the effects of self-selection to rolling multiple times, our design enables us to compare exogenous (*Three Rolls*) and endogenous choice of multiple rolls (*Multiple Numbers*). See Figure 2.1 for an overview of the design.

⁴The symbols displayed on the sides of the dice are identical to the ones on the practice dice.













	Single Roll	Three Rolls	Multiple Numbers	Multiple Figures
1st				
2nd	-			
3rd	-			
4th	-	-		

Figure 2.1: Potential Outcome of the Dice by Treatments

Note: This figure shows the potential outcomes of the dice in each treatment. Just to represent potential outcomes and the limitations of the treatments, we visualize four random outcomes. In *Single Roll* and *Three Rolls* treatments number of rolls is limited by design whereas in the *Multiple Numbers* and *Multiple Figures* treatments participants can roll the dice as many times as they want (unlimited). In the *Multiple Figures* treatment participants only see figures instead of numbers.

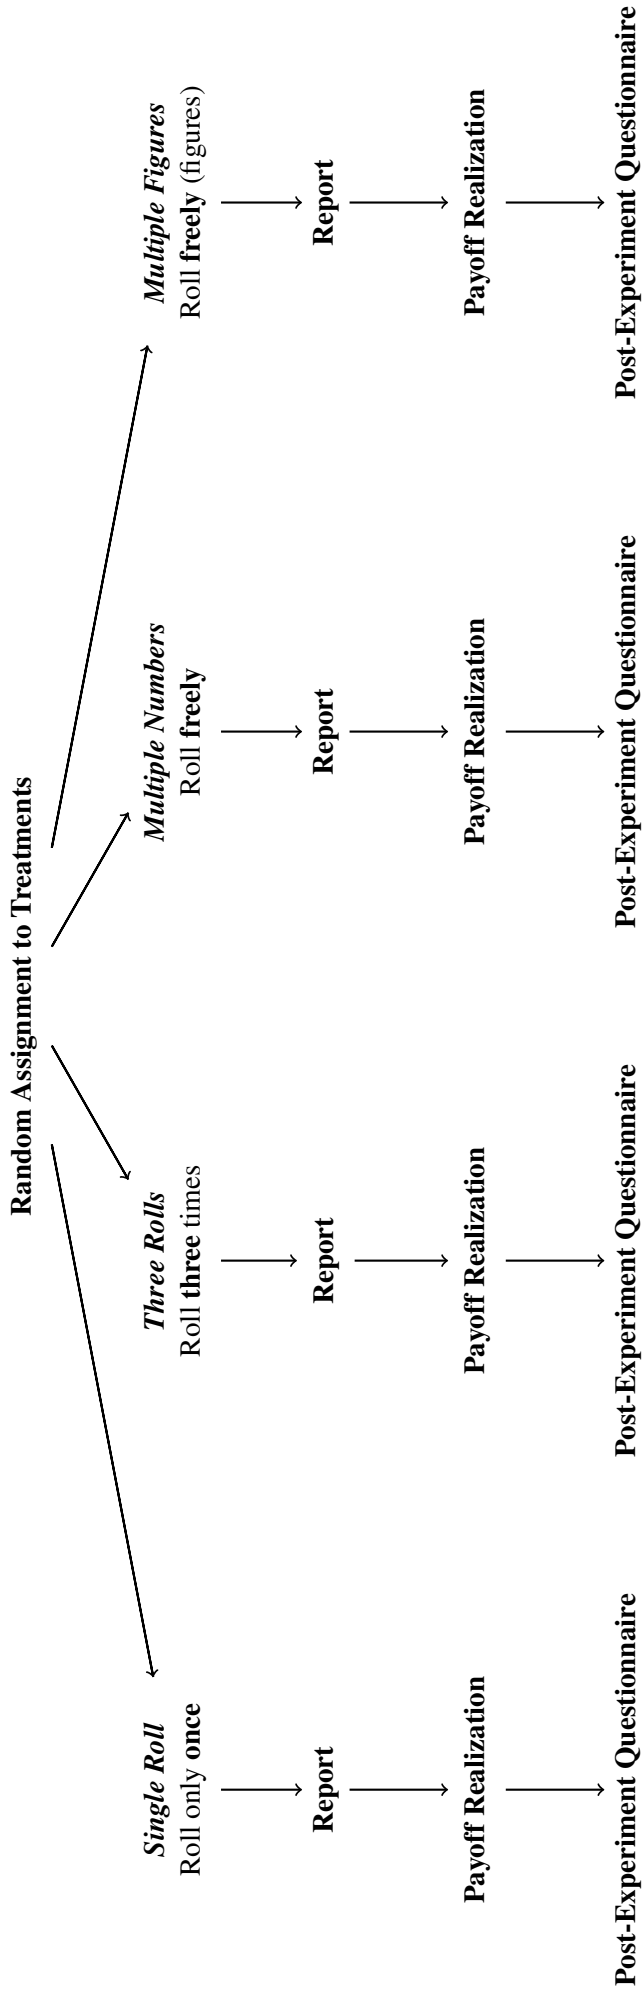


Figure 2.2: Experiment procedure

Note: Treatments are built on the variation of rolling rules. After participants are randomly assigned to one of the treatments, they roll the dice as many times as allowed (or they chose in the *Multiple Rolls* and *Multiple Figures*). After rolling they are asked to report the outcome of the first roll with being informed that the final payoff is calculated over the reported number.

2.3 Sample and Procedure

The experiment was posted as a Human Intelligence Task (HIT) on the Amazon Mechanical Turk (mTurk) crowdsourcing platform in June-July 2019. The interface of the experiment was programmed using oTree (Chen et al., 2016). Before data was collected for this experiment, our design was approved by Institutional Review Board at the NHH Norwegian School of Economics.⁵ Our hypotheses are pre-registered with AEA-RCT Registry (Ay et al., 2019).⁶ Each of the 1,580 participants participated in only one treatment and was not aware of the other experimental treatments.

All participants received the same instructions about the task, the payoff structure, and the overall procedure of the game. Participants were informed that their answers are recorded but would be kept anonymous and that the researchers cannot trace their choices back to their personal identities or their MTurk profiles. The instructions inform participants that all submitted work would be accepted regardless of accuracy and that payments would be transferred without any further questions shortly after the completion of the task. We reassured participants about their anonymity and the exclusion rules because participants could refrain from acting dishonestly due to reputation concerns on the online crowdsourcing platform.⁷ To avoid potential experimenter-demand effects, we provided this information to participants before the instructions about the dice-rolling task.⁸

Initially, participants started by practicing with a trial dice for as long as they wanted before reading about the main experiment. The sides of the trial dice displayed only non-ordered and random symbols and all participants had to roll the practice dice at least once before proceeding. This

⁵IRB Application number: NHH-IRB 07/19.

⁶Before running the main experiment, we conducted a pre-test on the same platform with 125 participants.

⁷Comments that participants give during the experiment also suggest that many were experiencing a moral dilemma in the reporting situation. For example, one participant wrote, "I did report the correct first roll of '2'. (I was curious and did roll the dice other times, but my report was truthful and accurate.)" Another wrote, "I wanted to report a higher dice roll for more money. I did not."

⁸Payments to participants is automated in the experiment, which allowed us to pay bonuses without storing worker IDs. We did not store worker IDs and IP-addresses to ensure anonymity. We deliberately abstained from using contentious words such as dishonesty, lying, or misreporting to avoid experimenter-demand effects. Participants were given a participation code once they accepted the HIT, which prevented them from retaking the HIT. No duplication of participation code was found.

method was chosen to reassure participants about the fairness of the dice without priming them with numeric outcomes. After practicing with the trial dice, participants read the instructions for the dice-rolling task along with the payment structure. After reporting their roll-outcomes, participants answered questions related to the experiment (e.g., perceived descriptive norms of similar dice-rolling games, self-reported feelings of being observed, and perceived legitimacy of rolling more than once), along with demographic questions. Overall, the experiment took participants approximately 7 minutes to complete, and the average payment was USD 2.5, which included a participation fee of USD 0.5.

Table 2.1: Summary descriptives table by groups of treatment

	<i>Multiple Figures</i>	<i>Multiple Numbers</i>	<i>Single Roll</i>	<i>Three Rolls</i>
	N = 386	N = 397	N = 409	N = 388
<i>Panel A</i>				
Age	37.8 (11.5)	38.7 (12.1)	39.0 (12.5)	38.3 (10.5)
Gender (= <i>F</i>)	0.47 (0.50)	0.48 (0.50)	0.45 (0.50)	0.44 (0.50)
Education:				
High School or Less	38 (11.5%)	52 (14.7%)	45 (12.5%)	41 (12.2%)
Higher Than High School	292 (88.5%)	301 (85.3%)	316 (87.5%)	295 (87.8%)
<i>Panel B</i>				
Norms	4.93 (2.80)	4.88 (2.60)	4.81 (2.78)	5.15 (2.76)
Feeling Observed	7.02 (3.25)	6.98 (3.21)	6.77 (3.13)	6.82 (3.36)
Political Views	3.39 (2.39)	3.48 (2.49)	3.18 (2.43)	3.45 (2.37)

Note: Standars errors are in parentheses. Values in the upper part of the table are self-reported demographics and beliefs elicited with a survey. It is clear from this table that sample was well balanced across treatments in terms of observable characteristics and beliefs. Political views are scaled from 0 (very liberal) to 8 (very conservative).

Table 2.1 contains descriptive statistics across the different treatment treatments. In *Panel A* we report the demographic measures, and in *Panel B* we report some of the self reported beliefs

that were collected the post-experiment questionnaire. Assignment to treatments is balanced in terms of observable characteristics and beliefs. The average age of participants is 38, and 46% of our sample is female. The majority of participants have at least a high school degree. *Panel B* in Table 2.1 provides an overview of covariates included in our analyses and estimations. *Norm* shows the beliefs on the levels of dishonesty in general, whereas *Feeling Observed* refers to how observed they felt during the experiment on a scale of 0 to 10. Political Views are elicited on a scale from 0 (Very Liberal) to 10 (Very Conservative). Our results show that participants are more liberal overall. We observe that participants' reported beliefs about the prevailing dishonesty norm or feelings of being observed do not differ significantly across treatments (Kruskal Wallis test results for norms $\chi^2_3 = 2.87$, $p = 0.41$ and for feeling observed $\chi^2_3 = 2.77$ and $p = 0.43$).

2.4 Results

In this section, we report the findings of our experiment. Reported results include both pre-registered and supplementary analyses to clarify our findings. Our results based on two sources of randomization: treatment variations and the outcome of the dice in the first roll. The first result we provide on dishonesty stems from the first source of randomization, whereas the results on curiosity are based on the second source of randomization. All the hypotheses and analyses are pre-registered except for the analyses on the number of rolls in the *Single Roll* treatment in Result 2b and Result 3b.

In the first subsection, we provide findings on dishonesty across treatments. Although our design has the similar features to Shalvi et al.'s (2011), the main goal of our paper is to investigate curiosity in a moral context. For this reason, in the following sections we document our findings on information collection and whether it could be strategically chosen. Using an observed game enables us to provide analyses on information collection based on the outcome of the first roll whereas in studies using unobserved games, it is not possible to investigate such behavior and motivations for it (Shalvi et al., 2011; Fischbacher and Föllmi-Heusi, 2013). Finally, we document

the correlation between curiosity and dishonesty.

2.4.1 Dishonesty

Our design provides two channels to investigate the impact of additional information on dishonesty: exogenous (*Single Roll* and *Three Rolls*) and endogenous variation of availability (*Multiple Numbers* and *Multiple Figures*). Collecting additional information on the outcome by rolling the dice multiple times enables participants to observe outcome-related (numeric) or unrelated (symbolic) information before reporting. In the case of honest reporting, the average expected reported number from the dice-rolling is 3.5.

Looking at the average reported numbers, we observe deviation from the theoretical expectation under honesty ($p < 0.001$). In our baseline *Single Roll*, average reported number is 3.94 (sd. = 1.79). We find no significant statistical difference on dishonesty across treatments (*Three Rolls*, 3.92 (sd. = 1.76) in *Multiple Numbers*, 4.12 (sd. = 1.66) in *Multiple Figures*, $p = 0.125$). Figure 2.1 shows the distribution of reported numbers in each treatment, which reveal that distributions of reported numbers are similar across treatments. The dashed line across Figure 2.1 indicates each number's theoretical frequency, i.e., $1/6$. It can be seen in the figure that the share of reports below "4" are lower than the theoretical fraction, whereas reports of "6" are higher and also the highest of all reports. This figure shows that participants tend to report higher numbers than "3" and mostly "6" if they decide to misreport.

In the experiment, we could observe both the actual and the reported number for each participant and are therefore able to develop a precise measure of dishonesty. These precise measures reveal that 19.9% of the participants in the *Single Roll* treatment, 24.6% in the *Three Rolls* treatment, 20.4% in the *Multiple Numbers* treatment, and 20.5% in *Multiple Figures* treatment misreported the outcome of the first roll. The share of dishonest reports is in line with the values shown in Abeler et al.'s (2019) meta-study on dishonesty experiments. Findings from the previous literature suggest that observing higher numeric outcomes than the actual outcome makes lying easier

by enabling participants to report the best outcome that they observe instead of making a fictitious report (Shalvi et al., 2011, 2015).

Our result shows that participants in all treatments are dishonest, and the level of dishonesty is not significantly different across treatments. This result contrasts with previous experimental findings in unobserved settings and our hypothesis on dishonesty in the pre-analysis plan. For example, Shalvi et al.'s (2011) study shows an increase in dishonesty when the participants are instructed to roll multiple times compared to only rolling once. In our experiment, the *Three Rolls* treatment is similar to the "Multiple Rolls" treatment in (Shalvi et al., 2011) regarding how many times participants are instructed to roll. In the *Three Rolls* treatment in which participants are instructed to roll three times and were not given a chance to change it, we do not find any significant difference in the level of misreporting (reported number (4.15) and distance (0.64)). The *Three Rolls* treatment provides both a comparison to Shalvi et al. (2011) and serves to control for




Figure 2.1: Distribution of Reported Numbers

Note: This figure shows the distribution of reported numbers for each treatment separately. The dashed line represents theoretical expectation on the share of each outcome from rolling a dice which is equal to 0.16. As can be seen fraction of 5 and 6 is higher than the theoretical expectation whereas fraction of reporting lower numbers fell below the theoretical expectation.

self-selection in our experiment. However, we do not observe an increase in misreporting in the *Multiple Numbers* treatment compared to *Three Rolls*.

Our results show that there is no difference in dishonesty between the settings where access to additional information is exogenously limited and those where it is endogenously decided. Although the level differences are not significant between treatments, mean reported values are higher than 3.5 in all treatments. This shows dishonesty in all treatments, no matter the availability and the content of the information.

Since we observe the actual outcome of the first roll, we can analyze the “size of the lie” or the distance between the reported and actual outcome. The “distance” shows the deviation from the actual outcome.⁹ We find no significant variation in distance between treatments ($p = 0.383$). In our pre-analysis plan, our hypothesis was that availability of additional information increases dishonesty. Our findings instead suggest that the availability of additional information does not affect misreporting.

Result 1: Availability of additional information - either exogenously or endogenously given - does not increase misreporting compared to having no additional information available before reporting.

2.4.2 Curiosity

In this section, we provide our findings on curiosity by first analyzing participants’ rolling behavior before reporting (ex-ante) and then the same behavior after the report is submitted (ex-post). For ex-ante information collection, our two treatments enable us to investigate curiosity based on the content of information: outcome-related information in *Multiple Numbers* and unrelated information in *Multiple Figures*. When additional roll-outcomes display numbers (*Multiple Numbers*), rolling more times generates counterfactual information, which has been previously shown to make misreporting more frequent (Shalvi et al., 2011). Observing desired counterfactuals helps

⁹Distance is 0.58 (sd. = 1.39) in *Single Roll*, 0.64 (sd. = 1.49) in *Three Rolls*, 0.47 (sd. = 1.40) in *Multiple Numbers* and 0.53 (sd. = 1.30) *Multiple Figures*.

to decrease the psychological distance between the actual outcome and the desired one.¹⁰ By contrast, when the additional roll-outcomes display random symbols (*Multiple Figures*), participants cannot use the additional information to reduce the psychological distance between the actual roll outcome and the desired one. In *Multiple Numbers* treatment, the information is always outcome related since the outcomes of the dice are always numeric, whereas in the *Multiple Figures* treatment, outcomes are unrelated since the dice shows random figures after the first roll.

In the pre-analysis plan, we hypothesized a higher number of rolls in the *Multiple Numbers* treatment compared to the *Multiple Figures* treatment. We find however, no difference in rolling behavior between observing outcome-related information (numeric) and information that is random and unrelated (symbols). Participants who can roll freely before they report throw the dice 4.98 (15.5) times after the first roll when the additional rolls have numeric outcomes, whereas this number is 4.82 (17.6) when the rolls have symbolic outcomes ($p = 0.96$). Figure 2.2 shows the average number of rolls prior to reporting in the *Multiple Numbers* and *Multiple Figures* treatments.




Figure 2.2: Number of Rolls Before Reporting

Note: Bars are standard errors. This figure shows average number of rolls for each treatment in which participants can roll the dice as many times as they want. For each treatment, the number of rolls is shown by the outcome of the first roll as below or equal 3 (≤ 3) and above 3 (> 3). As can be seen in the figure, participants who saw lower first outcome roll significantly more times than those who saw higher numbers. This holds for both Multiple-rolls treatments (numbers and figures) in which participants can roll as many times as they want before reporting. ($p < 0,001$)

¹⁰By desired counterfactuals, we refer to the higher numbers than the outcome of the first roll.

As shown in the figure, when faced with a lower outcome in the first roll, the number of rolls is significantly higher in both treatments. This result is in line with our pre-registered hypothesis.

Result 2a: People are curious and collect additional information even when the additional information is irrelevant to the task. Curiosity is higher when the outcome of the first roll is low.

Although not pre-specified, we want to further investigate whether people who rolled more when they observed a low outcome do so out of “pure curiosity” or whether they are searching to justify misreporting. This former type of curiosity is defined as “epistemic curiosity,” which is a desire to collect information even when it is not targeted to a certain end (Loewenstein, 1994; Litman et al., 2005). To assess this, we let those in the *Single Roll* treatment have the opportunity to keep rolling the dice after they submit the report. These participants are explicitly told that they would not be able to change their report but could keep on rolling if they wanted to do so. We use their ex-post rolling (after reporting) behavior in *Single Roll* treatment as a benchmark for seeking non-instrumental information out of pure curiosity. Figure 2.3 shows number of rolls in both *Single Roll* and *Multiple Numbers* treatments based on the outcome of the first roll (greater than 3 and lower or equal than 3). In both treatments, participants who roll lower outcomes in the first roll rolled significantly more times than those with a high outcome in the first roll ($p < 0.01$ for both treatments). Distribution of the number of rolls can be seen in Figure 2.4. Number of ex-post rolls is significantly lower than number of ex-ante rolls (4.98 in *Multiple Numbers*, 2.75 in *Single Roll*, $p = 0.005$).

Result 2b: Curiosity is observed even after the decision has been made—when additional information has no instrumental value. A low outcome of the first roll increases the number of ex-post rolls.

¹⁰We use instrumentality in terms of affecting the final result. Even though it cannot change the reported value, ex-post information may help people validate their dishonest reports. We provide further analysis of this behavior in the next section.

Our design enables us to provide findings on the rolling behavior based on the outcome of the dice. Conditioning on the outcome of the first roll has not been possible in previous research that has relied on hidden rolls. Our results show that the level of information collection is not significantly affected by the content of information. Our results document that participants continue rolling the dice even after reporting. Although additional information has no use, this behavior could be motivated by a desire to justify dishonesty. In the next section, we continue documenting our results on motivations for curiosity.

2.4.3 Motivations for Curiosity

A tension between reporting honestly and reporting self-servingly arises when there is a large distance between the desired outcome (rolling a high number) and the actual outcome. We conjecture that this tension gives rise to a demand for information that reduces the gap between the desired




Figure 2.3: Ex- Ante and Ex-Post Number of Rolls

Note: This figure shows average number of rolls for each treatment in which participants can roll the dice as many times as they want and the outcomes of rolls are numbers. The only difference is that in the *Multiple Numbers* treatment they can roll the dice before reporting (ex-ante) whereas in the *Single Roll* treatment they can only roll additional rolls after reporting (ex-post). Number of rolls is shown by the outcome of the first roll as below or equal 3 (≤ 3) and above 3 (> 3). As can be seen in the figure, participants who saw lower first roll-outcomes roll significantly more times than those who saw higher numbers. This holds for both *Multiple Numbers* and *Single Roll* treatments in which participants can roll as many times as they want before reporting. Although this result is not pre-registered we find it valuable to show the information seeking even when there is no instrumental value that can affect the final reporting behavior and the outcome.

and the actual outcome. Demand for information is driven by the need to find justifications that can reduce the perceived distance or gap (Shalvi et al., 2011; Fischbacher and Föllmi-Heusi, 2013). The demand will be greater the larger the distance is between the actual outcome (e.g. rolling a ‘one’) and the number that one would prefer to report (e.g. ‘six’). We pre-specified in the pre-analysis plan that people are more likely to acquire information that could reduce the perceived distance between the factual outcome and the wealth-maximizing outcome when this distance is large (e.g. rolling a one) compared to the when there is less or no distance (e.g. rolling a five). That is, when honesty concerns are in conflict with self-interest, individuals actively try to reduce the intrinsic cost of lying by acquiring information that may reduce the perceived size of the potential lie.¹¹ To investigate whether curiosity is used to reach higher monetary outcomes, we document estimations for the information collection in relation to the outcome of the first roll and other behavioral parameters in this section.




Figure 2.4: Distribution of Number of Rolls

Note: this figure shows the distribution of how many times participants chose to roll by treatments in the rows and the outcome of the first roll on the columns.

¹¹Our estimation is based on Gneezy et al. (2018), however, the method we use deviates to correctly specify impacts of our observed design.

Our analysis on motivations to collect additional information prior to reporting focuses on how many times participants choose to roll the dice in the *Multiple Numbers* and *Multiple Figures* treatments. In these treatments, participants have the opportunity to roll as many times as they would like to before they report the outcome of the first roll. We hypothesized that observing a low outcome on the first roll produces a demand for justification to misreport. Among participants who could freely roll the dice with numbers, we find that the lower their first roll, the more likely they are to roll again. Table 2.1 shows that for both treatments, higher outcome of the first roll significantly decreases the likelihood that the subject would roll more than once ($p < 0.01$).

Result 3a: Outcome of the first roll affects the likelihood of rolling more than once, even when the additional outcome is not related. The lower the outcome of the first roll, the higher the likelihood of continued rolling ex-ante.

This result supports our hypothesis on curiosity; people can be strategically curious to justify morally questionable decisions when the actual outcome is low. For lower outcomes of the first roll (≤ 3), the number of rolls is significantly higher before reporting (ex-ante) in the *Mul-*

Table 2.1: Logistic Regression for Rolling More than Once

	<i>Dependent variable:</i>	
	<i>(Multiple Numbers)</i>	<i>(Multiple Figures)</i>
First Roll	-0.349*** (0.072)	-0.421*** (0.078)
Controls	Yes	Yes
Observations	397	386
<i>dydx</i>	-0.065*** (0.021)	-0.081*** (0.024)
Log Likelihood	-225.443	-218.645

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$ *se.s* are in parentheses

Note: Standars errors are in parentheses. Results show that, in both treatments higher observed numbers in the first roll decreases the likelihood of “Keep Rolling” significantly. *dydx* shows the marginal effect of “First Roll”. Unlisted controls: Age, Gender, Norms, Feeling Observed, Political views, Income level, Education Level, Belief on Lying in the Experiment.

Multiple Numbers treatment than after reporting (ex-ante) in the *Single Roll* (4.61 in *Multiple Rolls*, 3.23 in *Single Roll*, $p = 0.03$). We do not observe a significant difference between ex-ante and ex-post number of rolls for participants who see higher outcomes in the first roll (2.88 in *Multiple Numbers*, 2.16 in *Single Roll*, $p = 0.49$).

This finding is in line with Gneezy et al.'s (2018) finding which states that the distance between reported and actual outcomes drives the intrinsic lying costs. We find that when the distance between the desired and actual outcomes increases, participants seem to desire information that could reduce this psychological distance. This finding suggests that people actively attempt to shape their information-set according to what serves their self-interest. Participants search for more information when they face lower outcomes in the first roll, and they search for such information significantly more when it can be acquired before rather than after the reporting decision.

Figure 2.5 shows number of rolls after reporting by honest versus dishonest participants (2.56 for honest reporting, 4.35 for misreporting, $p = 0.043$). As shown in the figure, participants who reported honestly roll significantly fewer times than participants who reported dishonestly. Showing that participants roll more after they misreport compared to if they reported honestly is important as it provides additional evidence that participants seem to use additional information as a way to assess the credibility of both potential and past lies—even when there is no monetary gain from rolling more and continued rolling has an opportunity cost.

Participants are more likely to roll the dice to a greater extent when the outcome of the first roll is low. However, in the ex-ante treatment, participants acquire more information than the ex-post situation, supporting our hypothesis on “strategic” curiosity. By design, the additional rolls before the reporting can be considered when reporting, but additional rolls after the reporting cannot affect the submitted report.

Result 3b: Participants who reported dishonestly are more likely to collect additional information after reporting.

2.4.4 Information, Curiosity and Dishonesty

Our results show a clear behavioral pattern on low values of the first roll and curiosity. When we define strategic curiosity, we posit that it is a tool to make self-serving decisions easier. To investigate whether participants use the additional information for this goal, we provide evidence for the relationship between the distance of the lie and the number of rolls.

We observe no significant difference in the level of misreporting between *Multiple Numbers* and *Multiple Figures* treatments. Although we do not observe treatment differences in reported numbers, Table 2.2 shows that distance (size of the lie) is significantly higher among those who chose to roll the dice more, regardless of the content of the dice—numbers or figures. As mentioned earlier, *Multiple Numbers* and *Multiple Figures* treatments are the ones in which participants “endogenously” decide how many times to roll the dice. In the *Single Roll* and *Three Rolls* treatments, the number of rolls is exogenously decided.

In terms of relevancy of the content, *Multiple Numbers* and *Multiple Figures* treatments provide variation of the content of the endogenously collected information.¹² We find a very similar




Figure 2.5: Number of Ex-Post Rolls and Dishonesty

Note: This figure shows the mean number of rolls for the Single Roll treatment by groups of dishonest (misreport) and honest reporting. After reporting, participants who misreported roll significantly more times than participants who reported honestly ($p = 0.042$).

¹²Relevancy is used in terms of additional information being relevant to the first roll and the reporting decision.

effect of number of rolls on distance in both treatments; it is positively and significantly associated with the distance of the lie. This shows that even though the effect size is different, curiosity toward both related and unrelated information is positively related, with greater distance between the actual outcome and the reported number.

These results can provide insights on motivations of dishonesty for two different types of information: with related information, dishonesty is driven by the fact that there is a potential higher outcome, and with unrelated information, the mechanism might be similar to distraction. Since it was not one of the main concerns of our research, the latter mechanism is not documented in detail. Effects of inattention and distraction have been previously documented in different behavioral concepts than dishonesty previously (Gabaix, 2019; Falk and Zimmermann, 2016). We show the same dishonesty level in these treatments and the same effect of rolling more on this level, suggesting that different mechanisms cause similar effects on dishonesty.

Result 4: Curiosity is associated with misreporting regardless of the content of additional information.

Table 2.2 also shows that, in the *Single Roll* treatment in which participants could roll the dice only once prior to reporting the outcome, the time participants spent on the rolling page (Time Rolling) is positively associated with misreporting.¹³ This result suggests that, even without rolling the dice more times, people who report dishonestly spend more time on the page before doing so. This finding supports the argument that morality is the intuitive choice in social decisions. For example, Cappelen et al. (2016) show a strong association between short response time and fair behavior, which means fairness is the intuitive choice in social decisions. In a meta study on dishonesty experiments, Köbis et al. (2019) show that in situations where dishonesty affects others, honesty is the intuitive choice.

¹³Note that in the rolling page of *Single Roll* treatment, participants are not allowed to roll after the first one, but they can deliberately stay on the page deliberately until they click the "Next" button. This variable is not added for the analysis in other treatments because people stay on the page while rolling the dice. In the *Single Roll* treatment, they cannot do anything on this page.

In none of the treatments is the feeling of being observed is related to individual misreporting. This shows that feeling observed during the game did not contribute to the treatment effects that are documented in this paper. In all treatments, beliefs on norms about dishonesty are associated with larger lies (i.e., larger distance from the first outcome). Beliefs on norms are elicited by asking participants how likely others are to report dishonestly in similar experimental settings. This is a self-reported measure on participants' beliefs about others' dishonest behavior in similar settings. Our findings show that beliefs about others' dishonesty is correlated with larger-distance lies. Although these mentioned effects are low, they support related research suggesting that beliefs about the prevailing descriptive norms (i.e., others' dishonesty in similar games) influence people's tendency to be dishonest (Bicchieri et al., 2019).

Table 2.2: Predictors of "distance" (Reported - Real Outcome) across treatments

<i>Dependent variable: distance</i>				
	<i>(Single Roll)</i>	<i>(Three Rolls)</i>	<i>(Multiple Numbers)</i>	<i>(Multiple Figures)</i>
# of Rolls			0.041*** (0.014)	0.041*** (0.011)
Time Rolling	0.010** (0.005)			
Time Report	0.002 (0.007)	0.013 (0.008)	0.012 (0.011)	0.005 (0.005)
Age	-0.011* (0.006)	-0.018** (0.009)	-0.027*** (0.007)	-0.007 (0.007)
Sex (F)	-0.400** (0.157)	-0.136 (0.180)	-0.270* (0.154)	-0.469*** (0.149)
Controls	Yes	Yes	Yes	Yes
Observations	347	317	324	331

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$ *se.s* are in parentheses

Size of the lie: Standard errors are in parentheses. To explain distance of the lie, we use behavioral measures that are collected in the game and other covariates like self-reported unobservable characteristics. *Time Report* variables show the seconds participants spent on reporting and rolling the dice. A subset of data used to make this estimation to exclude participants with first outcome of 6 (since including these participants masks the effects of other covariates with 0 lying). Unlisted controls: Political view, Income level, Education Level, Belief on Lying in the Experiment. When we make the same estimation for only dishonest people, we observe that 1 additional roll increases the distance by 0.23 units, which means 4 additional rolls increases the lie by 1 unit for dishonest reporters.

2.5 Concluding Remarks

Morality is often considered to restrict people from making self-serving decisions that are morally questionable. As opposed to treating morality as a goal in itself, it is sometimes treated as a set of internalized constraints on people's real goal of pursuing self-interest. In that case, people strategically acquire information to create moral wiggle room that makes otherwise morally unacceptable decisions appear morally acceptable. As people increasingly have access to an abundance of information, understanding how people use their discretion over various kinds of information in moral dilemmas is increasingly important.

In this study, we investigate how people strategically collect additional non-instrumental information to justify morally questionable decisions. We use a modified dice-rolling experiment where the availability and the content of additional information vary between subjects. This design enables us to disentangle what motivates collecting additional information, and to investigate how this affects dishonest reporting. We provide novel evidence that curiosity about related and unrelated information is heightened when being curious can help circumvent the moral obligation to report honestly. Even though people's curiosity is associated with more dishonest behavior, we find no treatment effect on dishonest reporting, suggesting a more complex underlying mechanism.

We obtain evidence that people acquire additional information—not only to search for justifications—but also to distract themselves from the moral conflict. Even after the reporting decision has been made in the *Single Roll* treatment, we find that people who misreport are more likely to acquire additional information to evaluate the credibility of their past lies. Further analyses provide additional evidence that people tend to use additional information to assess the credibility of potential lies rather than merely searching for justifications for selfish behavior.

Our study fills an important gap in the behavioral research literature by showing that information acquisition can be a strategic behavior. Previous research on information preferences in moral context has devoted considerable attention on information avoidance (Golman et al., 2017; Dana et al., 2007; Grossman and Van der Weele, 2017). We provide experimental evidence that

people actively collect non-instrumental information when tempted by the benefits of being dishonest. Our findings support the notion that people attempt to circumvent moral constraints rather than having a preference for morality (Rabin, 1995). As people often have discretion over how much information to collect and consider before making moral decisions, understanding endogenous information collection and processing choices is essential to improve our understanding of behavior and to design better policies.

We provide evidence that information collection is linked with higher levels of dishonesty even when the information is not related to the task. Although our design limits us to providing a more detailed investigation of the motivations, previous literature suggests that people might rationally seek inattention to distract themselves when making decisions (Gabaix, 2019; Sims, 2003, 2006). In the domain of moral cognition, becoming inattentive to moral standards can make it easier for people to excuse diverging from their standards (Mazar et al., 2008). Our findings on the systematic collection of unrelated information provide another interesting trait: people seem to collect information to distract themselves from moral dilemmas.

Even though our experimental investigation is concerned with endogenous information collection, we also contribute to research on how counterfactual information affects moral cognition (Shalvi et al., 2011; Bassarak et al., 2017; Effron, 2018). Previous literature finds that people tend to process new pieces of information to confirm their own beliefs. Though we find evidence of this, we find that people seem to not change their decisions based on the content of the additional information. Instead, our results suggest that dishonest people tend to acquire additional information to justify their dishonest decisions.

Our results provide insights that increase our understanding of the link between information and moral decisions by suggesting that people could use unnecessary information to justify morally questionable decisions. This insight is important as people have access to extensive information about many different topics in their daily life. Thus, we believe our results can increase our comprehension of social and economic decisions as our research serves as a step towards

understanding the psychological mechanisms that drive the strategic use of curiosity.

Bibliography

Abeler, J., Nosenzo, D., and Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4):1115–1153.

Ay, F. C., Nødtvedt, K. B., and Berge, J. (2019). Strategic curiosity. *AEA RCT Registry-AEARCTR-000440*.

Bassarak, C., Leib, M., Mischkowski, D., Strang, S., Glöckner, A., and Shalvi, S. (2017). What provides justification for cheating—producing or observing counterfactuals? *Journal of Behavioral Decision Making*, 30(4):964–975.

Bicchieri, C., Dimant, E., and Sonderegger, S. (2019). It’s not a lie if you believe it: Lying and belief distortion under norm-uncertainty. *Available at SSRN*.

Bodner, R. and Prelec, D. (2003). Self-signaling and diagnostic utility in everyday decision making. *The Psychology of Economic Decisions*, 1:105–26.

Cappelen, A. W., Nielsen, U. H., Tungodden, B., Tyran, J.-R., and Wengström, E. (2016). Fairness is intuitive. *Experimental Economics*, 19(4):727–740.

Chen, D. L., Schonger, M., and Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1):67–80.

Effron, D. A. (2018). It could have been true: How counterfactual thoughts reduce condemnation of falsehoods and increase political polarization. *Personality and Social Psychology Bulletin*, 44(5):729–745.

- Eliasz, K. and Schotter, A. (2010). Paying for confidence: An experimental study of the demand for non-instrumental information. *Games and Economic Behavior*, 70(2):304–324.
- Falk, A. and Zimmermann, F. (2016). Beliefs and utility: Experimental evidence on preferences for information. *CESifo Working Paper*.
- Fischbacher, U. and Föllmi-Heusi, F. (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association*, 11(3):525–547.
- Gabaix, X. (2019). Behavioral inattention. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 2, pages 261–343. Elsevier.
- Gneezy, U., Kajackaite, A., and Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review*, 108(2):419–53.
- Golman, R., Hagmann, D., and Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1):96–135.
- Golman, R. and Loewenstein, G. (2015). Curiosity, information gaps, and the utility of knowledge. *Information Gaps, and the Utility of Knowledge (April 16, 2015)*.
- Golman, R., Loewenstein, G., Molnar, A., and Saccardo, S. (2019). The demand for, and avoidance of, information. *Information (May 31, 2019)*.
- Grossman, Z. (2014). Strategic ignorance and the robustness of social preferences. *Management Science*, 60(11):2659–2665.
- Grossman, Z. and Van der Weele, J. J. (2017). Self-image and willful ignorance in social decisions. *Journal of the European Economic Association*, 15(1):173–217.
- Köbis, N. C., Verschuere, B., Bereby-Meyer, Y., Rand, D., and Shalvi, S. (2019). Intuitive honesty versus dishonesty: Meta-analytic evidence. *Perspectives on Psychological Science*, 14(5):778–796.

- Litman, J., Hutchins, T., and Russon, R. (2005). Epistemic curiosity, feeling-of-knowing, and exploratory behaviour. *Cognition & Emotion*, 19(4):559–582.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1):75.
- Mazar, N., Amir, O., and Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research*, 45(6):633–644.
- Rabin, M. (1995). Moral preferences, moral constraints, and self-serving biases. *mimeo-UC Berkeley: Department of Economics*.
- Shalvi, S., Dana, J., Handgraaf, M. J., and De Dreu, C. K. (2011). Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior. *Organizational Behavior and Human Decision Processes*, 115(2):181–190.
- Shalvi, S., Gino, F., Barkan, R., and Ayal, S. (2015). Self-serving justifications: Doing wrong and feeling moral. *Current Directions in Psychological Science*, 24(2):125–130.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.
- Sims, C. A. (2006). Rational inattention: Beyond the linear-quadratic case. *American Economic Review*, 96(2):158–163.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3):213–225.

2.A Supplementary Analysis

2.A.1 Details of Dishonesty

Shalvi et al. (2011) argue that observing higher counterfactuals causes misreporting to reach higher potential payoff. On the other hand, Gneezy et al. (2018) predicts that participants would refrain from fully using the strategic advantages when lying as the size of the lie increases. In this section we report results about lying behavior that are not pre-registered. Figure 2.A.1 shows the average reported number by the first roll. As can be seen, for every outcome we observe lying which is higher for values lower than 4 compare to higher outcomes. Our results show that participants who misreported mostly reported 6, while some share report lower numbers as can be seen in Figure 2.A.2. Distribution of overall (pooled sample of honest and dishonest reporters) reports also show high share of 6 reports in Figure 2.1. The dashed line in Figure 2.1 shows the theoretical share of each number which is $1/6$ (0.16).

2.A.2 Dishonesty and Ex-Post Information Collection

As shown in the main results, we observe that people who see a lower outcome in the first roll and who misreport are more likely to roll the dice ex-post, after the report has been submitted. To complete the analysis on that here, we provide an an extra analysis on the relation between dishonesty and ex-post rolls. As can be seen in the Table 2.A.1, there is a positive relation between the distance (the size of the lie) and the ex-post number of rolls. As mentioned earlier, although this result is not pre-registered we find it crucial to understand the relation between moral decisions and curiosity. In a digital dice rolling game, people who misreport collect information by rolling the dice not only before reporting but also even after the report has been submitted.

Figure 2.A.1: Dishonesty and the outcome of the first roll

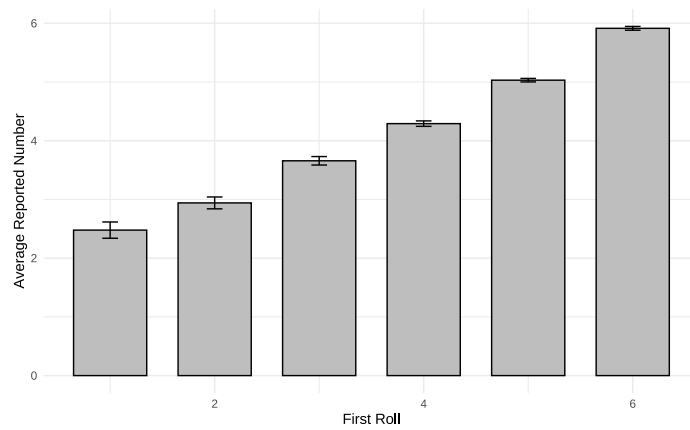



Figure 2.A.2: Distribution of dishonesty



Note: This figure shows the distribution of reported numbers for the participants who reported dishonestly. The figure is separated by the potential outcomes of the first roll and each column corresponds to a level as indicated in the upper titles. Most of the participants reported 6 while for lower outcomes of the first roll slightly higher variation can be seen. We do not observe any “negative lying”.

Table 2.A.1: Dishonesty in the Single Roll treatment and Number of Ex-post Rolls

	<i>Dependent variable:</i>
	<i>distance</i>
# of Ex-Post Rolls	0.058*** (0.022)
Time Rolling (Ex-Ante)	0.009* (0.005)
Time Report	0.003 (0.007)
Age	-0.010 (0.006)
Sex (F)	-0.392** (0.156)
Controls	Yes
Observations	347

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Ex-post rolls and dishonesty: Standard errors are in parentheses. To explain the relation between ex-post rolls and dishonesty, we use behavioral measures that are collected in the game and other covariates such as self-reported unobservable characteristics. A subset of data was used to make this estimation to exclude participants with first outcome of 6 (since including these participants masks the effects of other covariates with 0 lying). Unlisted controls: Political view, Income level, Education Level, Belief on Lying in the Experiment.

2.B Instructions

Welcome!

This is a study about decision-making conducted by researchers at the Norwegian School of Economics (NHH). You earn at least \$1.00 from participating and you will have chance to earn up to \$3.50 based on your decisions.

The data collected for the study will be completely anonymous, and the researchers will not be handling any personal data as a part of this project. In potential publications resulting from the study, it will not be possible to recognize any individual participant.

Important information about this HIT:

- Estimated time: 7.5 minutes
- Your choices are **completely anonymous** and will not be traced back to you by the researchers
- Generally, all HITs are accepted as long as workers read instructions and answers all questions
- This study does not involve deception

Do not refresh pages or open multiple tabs with the task URL. This HIT is not compatible with mobile devices. You will not be able to navigate back once you have proceeded to the next page

Participation is voluntary. You can discontinue the study at any time **but** you have to complete all three stages to be eligible to receive your payment.

I agree to participate in this study

Start the experiment

Test your understanding

In this study, your choices are partially anonymous:

- True
- False

According to the study's HIT policy, you get rejected from submitting incorrect answers:

- True
- False

In this study, the answers you submit will not be traced back to you personally:

- True
- False

Next

Payment information

From now on, you will roll a regular six-sided digital die with numbers from 1 to 6.

Your task is to roll the die, and anonymously report the outcome. The number on the die (i.e. 1, 2, 3, 4, 5, 6) indicates the payment you receive from this part of the experiment. For example, if the number is "1" you earn \$0.5, and if the number is "6" you earn \$3.0.

Here is a full overview:

Number	Payment
1	\$0.5
2	\$1.0
3	\$1.5
4	\$2.0
5	\$2.5
6	\$3.0

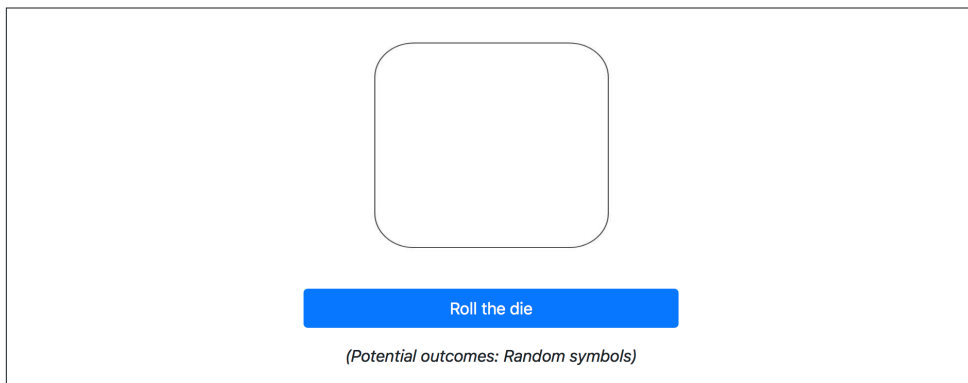
I have read and understood the instructions

Next page

Practice rolling a digital die

Below is a six-sided digital die. Each side of this die contain a random and non-ordered symbol instead of a number. When you press the 'roll-the-die' button, you roll the die and see the outcome of that throw. Every time you press the button, the die will roll again and show the outcome of the next throw.

Try rolling the die and press the 'Next-page' button once you understand how the digital die works.



Instructions

Before you roll the digital die, read the following instructions carefully:

- **First:** Roll the digital die
- **Second:** Check the outcome of your roll
- **Third:** Proceed to the next page to anonymously report the number you got

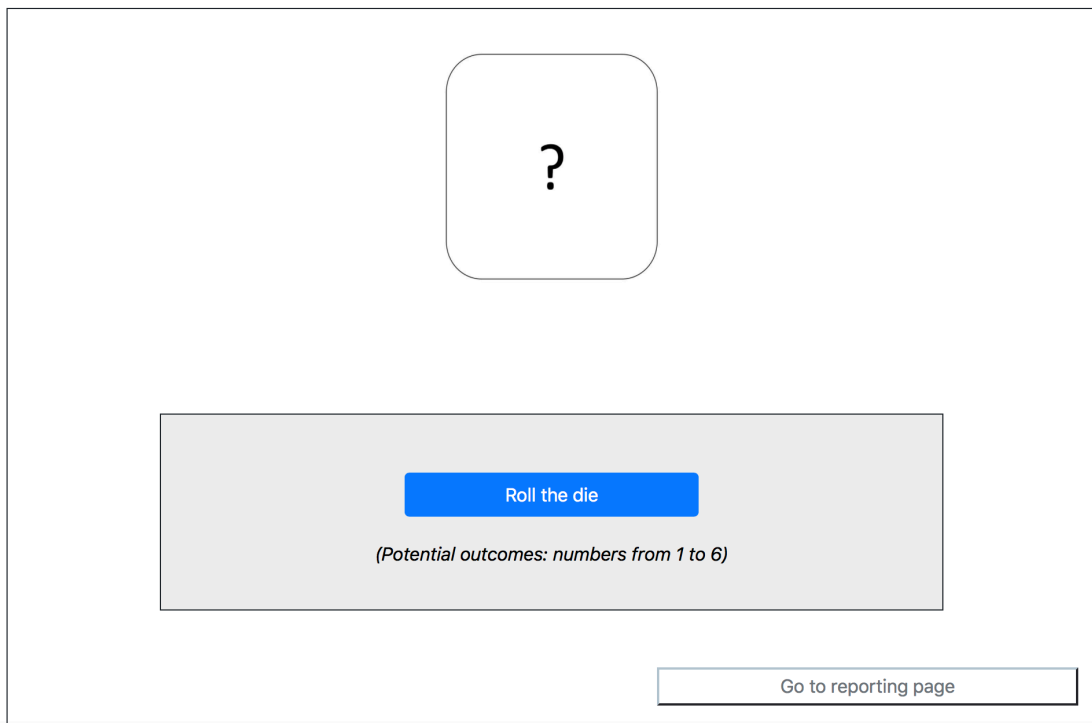
By pressing the "next" button below, you will be able to roll the digital die.

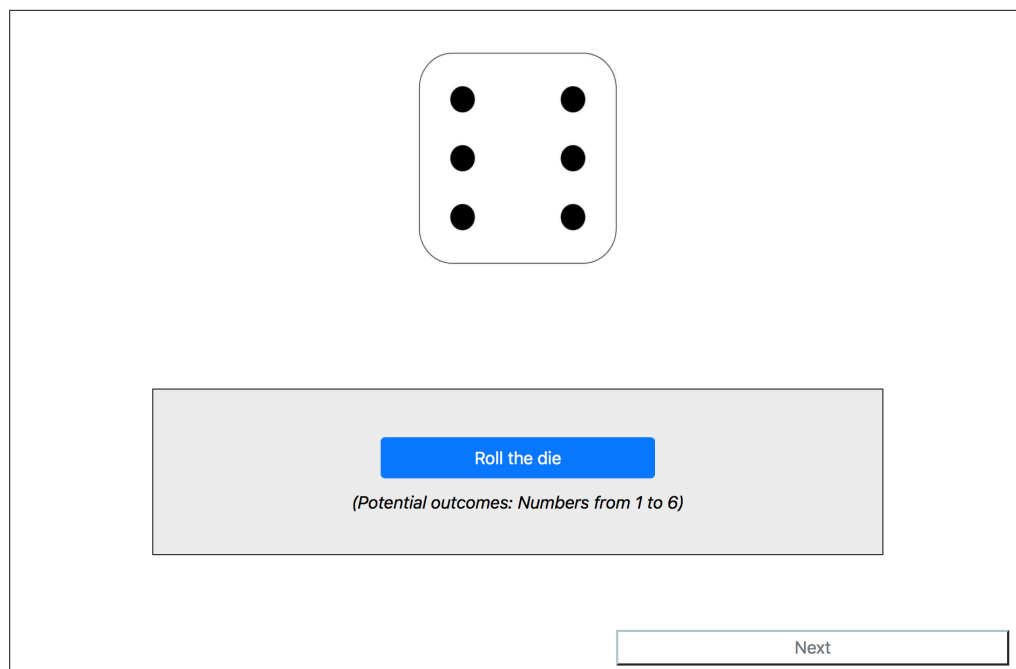
I have read and understood the instructions

Next

You will now be able to roll the die

I am ready to roll the die





Keep rolling?

You may now continue rolling the digital die as many times as you would like to verify to yourself that the die is fair. However, you cannot change the report you submitted. If you do not want to continue rolling, just press the "next" button on the following page.

[Go to the digital die](#)

Your report

Report the number you got:

1 2 3 4 5 6

In this last part of the study, you answer some questions related to the experiment and some about yourself. Once you are done with the survey, you will receive your MTurk code and be informed about your payment.

[Next: Survey](#)

Questions related to the study

In this study, you were asked to roll a digital die and report the outcome of the roll. The following are questions related to this task.

After you rolled the die once, how curious did you feel about what the die would have shown if you rolled it one more time (or a few more times)? On a scale from 0 (not at all curious) to 10 (very curious):

Some research studies use a similar die-rolling task where participants can report dishonestly to earn more money. To what extent do you think people are honest in these studies?

Did you feel that your reporting choice was anonymous?

Next page

Questions related to the study

Rate your political views on a scale from 0 (Very liberal) to 10 (Very conservative):

0 1 2 3 4 5 6 7 8 9 10

Do you have any other comments regarding this experiment?

Next page

Background information

How old are you?

What is your gender?

- Male
- Female
- non binary
- do not want to report my gender

What is your highest academic achievement?

- Less than high school
- High school graduate
- College graduate
- Professional degree
- Doctorate

How many years of work experience do you have? :

How would you describe the annual income of your household?

Thank you for participating in this experiment.

- Your bonus payment is **\$2.50**, which corresponds to the outcome you reported (that is, 5).
- You have also earned an additional **\$0.50** for participating.
- In total, your pay is: **\$3.00**

Instructions

Before you roll the digital die, read the following instructions carefully:

- **First:** Roll the digital die once
- **Second:** Check the outcome of your roll
- **Third:** After the first roll, roll the die *two more times*.

When you are done rolling, you may proceed to the next page to anonymously report the number you got on your first roll.

By pressing the "next" button below, you will be able to roll the digital die.

I have read and understood the instructions

Next

Figure 2.B.1: Instructions for the Three Rolls treatment

Instructions

Before you roll the digital die, read the following instructions carefully:

- **First:** Roll the digital die once
- **Second:** Check the outcome of your roll
- **Third:** After the first roll, you may continue to roll the die as *many times as you would like*.

When you are done rolling, you may proceed to the next page to anonymously report the number you got on your first roll.

By pressing the "next" button below, you will be able to roll the digital die.

I have read and understood the instructions

Next

Figure 2.B.2: Instructions for the Multiple Numbers treatment

Instructions

Before you roll the digital die, read the following instructions carefully:

- **First:** Roll the digital die once
- **Second:** Check the outcome of your roll
- **Third:** After the first roll, you may continue to roll die as *many times as you would like*
- **NB!** After rolling the die once, the die reverts back to only containing random symbols. That is, the sides of the die change from numbers to random symbols.

When you are done rolling, you may proceed to the next page to anonymously report the number you got on your first roll.

By pressing the "next" button below, you will be able to roll the digital die.

I have read and understood the instructions

Next

Figure 2.B.3: Instructions for Multiple Figures treatment

Chapter 3

Investigating Motivations for Information Avoidance—The Role of Certainty, Rewards and Overconfidence

Abstract

In this paper we investigate an objective for active information avoidance: self-image protection. In two experiments, we study how information preferences can change when the information is ego related and hurtful to self-image, and when it is neutral. To vary the neutrality of information, in the first experiment we use information about one's relative rank in an intelligence test and in the second experiment we use "neutral" information about a randomly drawn number. Our results show that people avoid ego related information and are willing to pay to do so. However, this behavior is not related to confidence and certainty about beliefs in one's own skills. We observe that a high share of people avoid "neutral" information even though it is not ego-related, and the willingness to pay to do so is positive. Besides avoidance, we provide findings that show people's

⁰This work was partially supported by the Research Council of Norway through its Centres of Excellence Scheme, FAIR project No 262675. We are grateful to Alexander Cappelen, Erik Ø. Sørensen, Bertil Tungodden and the seminar participants at FAIR, Norwegian School of Economics and Brown University, Department of Economics. We would like to thank Björn Bartling and Botond Közsegi for helpful and fruitful comments on our paper.

willingness to pay to receive neutral information responds to the value of information. Our findings contributes to the literature on information and ego utility by showing that these preferences might not be driven solely by self-image concerns.

Keywords: Information avoidance; ego utility; self-image; random choice

JEL-Classifications: C91, D82, D83, D91

3.1 Introduction

Economic theories and experiments show that people tend to process information about their own abilities in a biased manner. Previous studies show that this behavior stems from the fact that people strive for a positive image of being skillful (Bénabou and Tirole, 2002). This creates a motive to avoid and to demand information other than the usefulness of it to maintain a positive self-image. People tend to avoid information if they anticipate bad news (Ganguly and Tasoff, 2017; Eil and Rao, 2011), particularly when faced with a judgement about attributes that are important to people, such as one's own intelligence. If the anticipated content is positive about their skills, intelligence and/or monetary gains, people demand information (Eliaz and Schotter, 2010; Karlsson et al., 2009). Kőszegi (2006) relates this behavior to confidence about one's own skills and abilities: people with low confidence may avoid information if they are not certain of their beliefs about their abilities and thus, have fragile self esteem, whereas those who are overconfident and certain are not expected to avoid information, because their beliefs are less likely to change and their self esteem is not fragile.

These studies suggest that people want to protect themselves and to preserve a positive *self-image* by not learning potentially unpleasant things about themselves such as being less intelligent or less successful than expected. In this paper we experimentally investigate preferences over ego related information when receiving or avoiding this information directly influences monetary outcomes. We give people the chance to earn a bonus payment by learning their rank relative to their peers in an intelligence test. We introduce two treatments in which we vary whether receiving or avoiding information is free. With this variation we elicit the effects of the default settings on the information choice. We show that approximately half of all participants do not want to learn their relative rank even when learning can increase final monetary payoffs. However, we also document that a substantial share of people demand ego related information. To establish the effects of ego utility on information choice, we run a second study with Neutral Information that has no ego relevant content. Introducing this second study enables us to compare information acquisition

behavior when the information is ego related and when it is not.

Our findings contribute to the literature on information preferences and ego utility by showing that people are willing to avoid learning potentially hurtful news even when that means that they forego a certain monetary payoff. These results provide evidence on the importance of ego utility when making choices over information and extends the experimental evidence on avoidance from potentially bad news about one's own intelligence by documenting the trade-off between ego utility and monetary payoffs. Contrary to the previous arguments on ego utility, we find no significant relation between confidence, certainty about one's own intelligence and information avoidance behavior.

To investigate information preferences in relation to ego utility we ran the experiment (*Ego Utility*) in October 2018. We recruited 405 undergraduate students to participate in a laboratory experiment at the Norwegian School of Economics (NHH), one of the most selective business schools in Norway. Participants were recruited from first-year students who have not finished any obligatory courses and thus, have not taken any exams at the time of the experiment. This means, at the time of the experiment, participants had not yet received prior information about their relative success in their cohort at NHH yet. In the experiment, first an intelligence test is administered, and then each participant is asked to guess the share of participants that performed better than them. Participants are told that if their guess is correct, they receive a bonus of 80 NOK (approx. \$10) in addition to their show-up fee. Following this, they are informed that they have the chance to learn their true rank and revise their previous guess accordingly. Revising the guess based on true information about their rank guarantees that they receive the 80 NOK bonus at the end. At this stage, participants face the choice whether they want to receive the information or not. In one treatment, *Costly Avoidance*, receiving information is costless and while participants would have to pay to avoid information. In the second treatment, *Costly Information*, participants must pay to receive information while avoiding is costless. In both treatments, participants are asked to state their Willingness to Pay (WTP, henceforth) for their choice to be implemented. By using this, we

can measure WTP for avoidance in the *Costly Avoidance* treatment, and WTP for Information in the *Costly Information* treatment.

Our results from the first experiment show that half of the participants avoid learning their relative rank even when it is costly to avoid information. We observe high demand for avoidance of ego related information. However, neither overconfidence nor certainty about the guess are significantly correlated with information choice. High shares of positive payers for both avoidance and information indicates that in such decisions people may have substantial behavioral noise. This makes it difficult to reach a clear conclusion about Ego Utility and information choice. To clarify the results of Ego Utility, based on the results of the first study, we conduct a second study, *Neutral Information*.

The *Neutral Information* experiment is identical to the *Ego Utility* experiment, with one main difference: information is about a randomly drawn number instead of the rank in an intelligence test. Information is neutral because it is not related to any individual ability or success. The *Neutral Information* experiment functions as a control for the WTP measures in the *Ego Utility* experiment. If the demand for information and avoidance stem from the behavioral noise that people have, we expect to see similar results in the second study. By comparing these two studies, we are trying to disentangle the effects of ego Utility from the behavioral noise that can be seen in both studies. In total of 399 undergraduate students participated in the second experiment in October 2019. The second experiment was conducted a year after the first experiment with first year students at the time.

Results from the second experiment show lower avoidance from Neutral Information compared to ego-related information. However, a substantial share of people still avoid information even when it is about a randomly drawn number. A higher proportion of participants chose to receive neutral information. The demand to receive information responds to the value of information in this experiment; demand for high value information is higher than demand for low value information. Comparing these two experiments provides evidence that demand for information reacts

to the value of neutral information whereas with ego related information, self-image concerns play an important role and lead to higher avoidance. Overall, findings from these two experiments show that when calculating WTP for information choice, people may have behavioral noise. However, despite the noise in behavioral measures, we document higher demand for information avoidance when information is ego related.

In this paper, the design of our first experiment enables us to investigate the relation between overconfidence and information choice. Overconfidence is extensively investigated in the previous literature. It has been shown that overconfidence motivates decisions to compete or not and causes significant variations through gender (Niederle and Vesterlund, 2007). We focus on the relation between overconfidence and information preferences. This relation is investigated in different aspects in the previous literature. Particular questions that have attracted researchers' attention focus on how people process information regarding their abilities and what kind of biases could emerge. When the information challenges existing beliefs, people may process the information in a biased manner or they may choose to avoid information if they are given the chance to do so. Evidence from economic experiments show that people process information in a self-serving way to preserve self-confidence (Möbius et al., 2011), they update their beliefs when faced with both bad and good news to strengthen confidence (Coutts, 2019), and they are less responsive to information about their own personal traits (Buser et al., 2018). Different from these studies, we investigate whether confidence can shape preferences over information by motivating people to choose what to learn and what information to avoid. Our results show that confidence does not significantly affect information preferences. However, certainty about one's own guess significantly decreases the likelihood of positive WTP to avoid information. These findings extend the previous literature by showing that even though how successful people think they are is not highly effective, the more certain they are about their guess the lower is the likelihood to pay any positive amount to avoid information.

We use a Becker-DeGroot-Marschak (BDM, henceforth) auction in the experiment to elicit

true WTP for information and avoidance (Becker et al., 1964). In this mechanism, it is the dominant strategy to truthfully reveal the willingness to pay for a good or information. If the willingness to pay is higher than a randomly drawn price, the player wins the auction and vice-versa. BDM is one of the commonly used methods to reveal true WTP for a good or a decision. Comparing the effects of ego related information to neutral information enables us to show how WTP is affected by ego utility. Our results from these two experiments contribute to the usage of the BDM mechanism in information choices when information is ego related or neutral. We provide evidence that there is behavioral noise in such scales. However, WTP measured in a BDM-like mechanism still responds to the type of utility which shows that the measurement errors may stem from behavioral noise that is observed in individual decisions.

In section 2 we document the details of the design and results of first experiment, Ego Utility, and in section 3 documents the second experiment, Neutral Information.

3.2 Ego Utility and Information

In this section, we describe the first of the two experiments that we cover in this paper: Ego Utility. Our main goal is to investigate whether preferences over ego related information are influenced by factors such as overconfidence, certainty and monetary rewards. For this reason, we varied the costly information choice—receiving or avoiding information—in one treatment variation.

In subsection 2.1, we document the details of the design, in subsection 2.2 we cover the information about the procedure and sample, and in 2.3, we document the findings from the *Ego Utility* experiment.

3.2.1 Experimental Design

In this experiment, we aim to measure willingness to pay to receive or to avoid information on one's own relative intelligence. Learning one's own relative intelligence could be detrimental to self-image in case the results do not meet the expectations. Yet, the information can also increase

participants' earnings in the experiment. Therefore, we introduce a trade-off between monetary utility and potentially undesirable information.

First, participants take an IQ test that is taken from a larger test and contains 25 progressive Raven-style Matrices in total.¹ The original time for completing the test is 10 minutes however participants in our experiment are given 8 minutes to complete all the questions. Participants are informed that the test is different from the original test and the accuracy of the IQ measurement might not be very high. During the test, it is possible to go back and forth and change answers. Participants are instructed to solve as many matrice-tasks as possible in 8 minutes. After taking the test, they are asked to guess their rank compared to their peers in the same session. If their guess is correct (± 5) they win 80 NOK.² After making their guesses, participants are randomly assigned to one of two treatments: Costly Information and Costly Avoidance. In these treatment variations, the default information setting is either receiving (Costly Avoidance) or not receiving (Costly Information) information. After being informed about whether they must pay to receive or to avoid information, participants are asked to state whether they want to learn the information on their true relative rank or not. If their preference is different from the default setting of the treatment, they must pay an amount for that decision to be implemented. In both treatments we use a Becker-DeGroot-Marschak (BDM) (Becker et al., 1964) auction to measure willingness to pay (WTP) for implementation of the participants' information choice. We use the BDM auction to elicit participants' true information preferences and WTP for that to be implemented. All participants receive 50 NOK that they can use in the BDM auction. Then they are asked to submit how much they would be willing to pay for their decision (receiving/avoiding information) to be implemented. Each participant can submit a price between 0 and 50 NOK. The submitted price is compared to a randomly drawn game price in the next stage and if the submitted WTP is higher, the participant pays the price, and learns the true rank. If the auction is lost, the participant do not

¹The intelligence test is taken from a longer test (<https://openpsychometrics.org/printable/>) and the original time is 10 minutes instead of 8 minutes.

²They are asked to guess 'What percentage of the participants in the same room at that moment performed better than them?' by entering a number between 0 and 100. 80 NOK was approximately 9 Euros at the time of the experiment.

learn anything about their rank, and they can keep the 50 NOK bonus payment. If the participant gets the information (with or without choosing it) they have the chance to revise their guess based on the correct rank information. At the end of the game they will receive an extra payment from their guess (0 NOK or 80 NOK) and the rest of the bonus payment after BDM results (see Figure 3.1).

The BDM mechanism allows us to elicit preferences toward information since it enables us to capture variations in the WTP for their decisions to be implemented. Participants are given the chance to enter “0” if they want to implement the default setting. By entering the maximum amount, “50” they ensure that their decision is implemented. The best strategy in such a mechanism is to enter their “true” WTP (Becker et al., 1964; Karni, 2009). Note that participants who received information are shown the percentage who performed better than them as the fraction of total number of participants in the same room at that session. By doing so we aim to prevent information being more hurtful than aimed. Since in some cases participants show willingness to avoid information but not sufficiently high to make it be implemented, we aimed for less hurtful information by showing only relative performance instead of intelligence scores.

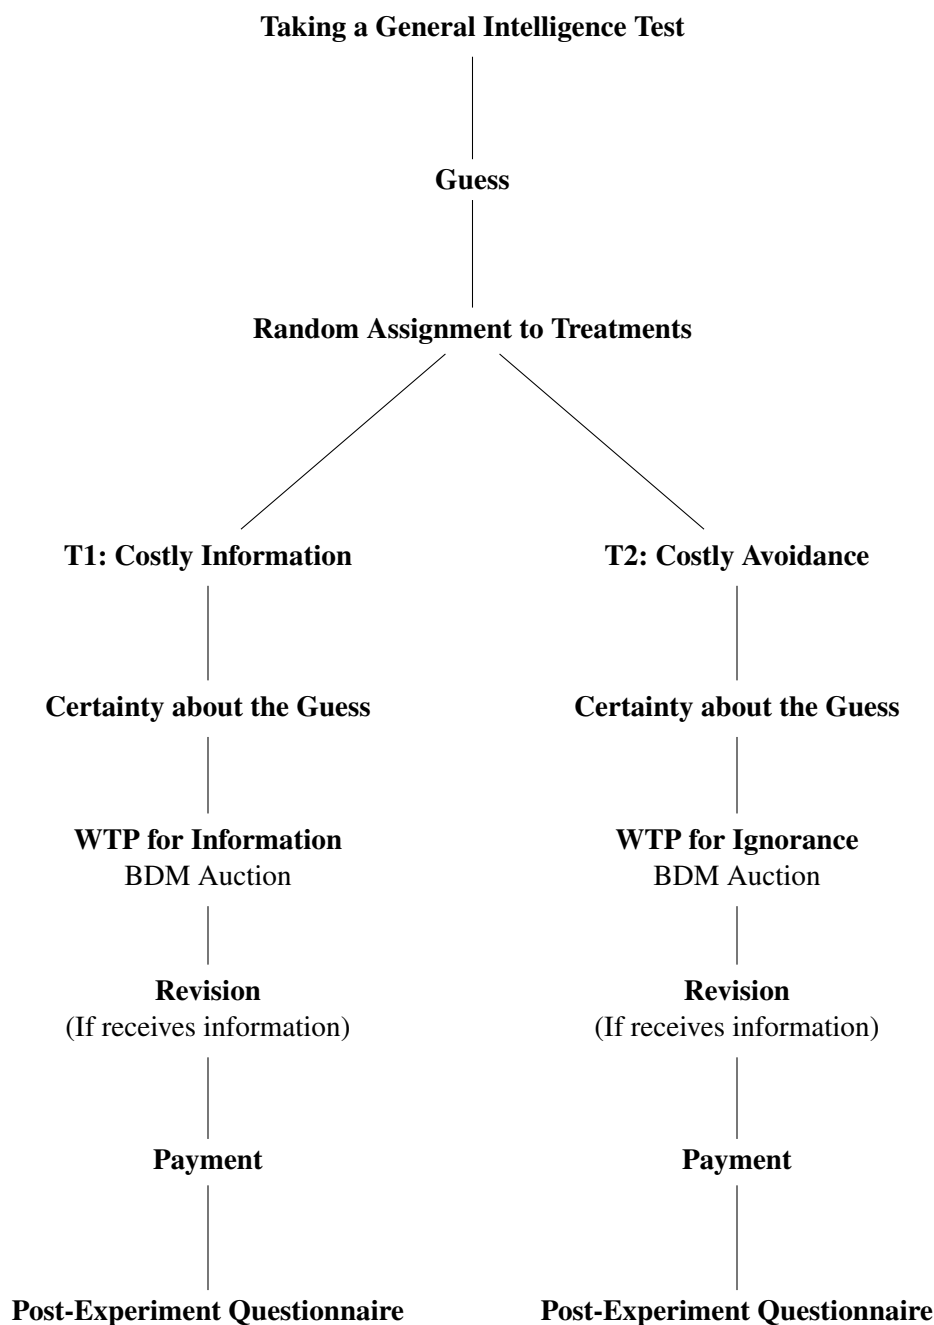


Figure 3.1: Ego Utility Experiment

Note: In this experiment, people decide whether to receive or avoid information about their relative performance in an intelligence test. In T1: Costly Information, they are told that they have to pay and they have to incur a monetary cost if they choose to receive the information whereas avoidance is free. In T2: Costly Avoidance, they have to incur a cost if they choose to avoid information whereas information is free. Receiving the information guarantees earning the 80 NOK bonus. To elicit people's willingness to receive and avoid information we use a BDM auction. After final payoffs are realized, a post-experiment took place to elicit demographics and other covariates.

3.2.2 Sample and Procedures

405 first year undergraduate students at the Norwegian School of Economics (NHH) participated in a total of 7 sessions. The experiment is programmed in Zurich Toolbox for Ready-made Economic Experiments (Z-tree) (Fischbacher, 2007) and participants are recruited via Expmotor (Sørensen, 2011) (participant interface can be seen in Appendix B). Prior to data collection, we registered a pre-analysis plan with a detailed description of our hypotheses with AEA RCT Registry (Ay and Meißner, 2018).

After entering the room, participants are randomly placed in the booths that prevent them seeing each other's monitors. Instructions and consent are given both verbally and on the monitor. All the participants are randomly assigned to one of the treatments within each session.

At the time of the experiment, first year undergraduate students have taken only one obligatory course and they have not received any grades yet. This means, they did not have course grades to compare themselves with their peers. This enables us to eliminate some factors that might indicate people's relative success or intelligence compare to their peers. We recruited first year students who have been admitted to NHH, one of the most selective business schools in Norway. This might be self-evident about their success and intelligence while not knowing their exact relative situation compared to their cohort. In the experiment, a total of 8 minutes is given to complete 25 questions. In the original version, the test is 10 minutes long, however we reduce the time to create uncertainty with time pressure which reduces the risk that beliefs are either on the upper or lower ends of the distribution.³ By doing so, we aim to have variation in the guess on how well they did and the certainty about the guess. Our sample is not comparable to the general population since they are successful and competitive undergraduate students in a selective business school. Moreover, we manipulated the time given in the intelligence test, which makes the results of this test not comparable to the other samples that took the test.

Table 3.1 summarizes the descriptives of the sample. The majority of participants are male

³We report the progress in the test in terms of number of people seeing the questions and answering correctly in Appendix A.1.

(65%) and the average age is 20.5. As can be seen, our sample is well balanced in terms of observable characteristics. In the post-experiment questionnaire we ask people how likely they are to take risk in general on a scale from 0 to 10. The average score for willingness to take risk is 5.4 and the sample is well-balanced in terms of self-reported risk preferences. Participants in this experiment answered 8 questions correctly on average. This does not vary across treatments.

3.2.3 Results

In this section we document the results from the *Ego Utility* experiment. We start by reporting the findings on the WTP for information and for avoidance. We find a positive WTP for both avoiding (23.7, sd.=14.8) and receiving (29.7, sd.=14.5) information (see Figure 3.2). However, the average WTP to receive information is significantly higher than WTP to avoid ($p < 0.001$). 90% of the participants in the *Costly Avoidance* treatment entered a positive WTP to avoid information whereas this share is 95% to receive information in the *Costly Information* treatment.

Result 1: The willingness to pay for both, information and avoidance, is positive.

Both, information and avoidance, are valued by a substantial share of participants. For this valuation to be reasonable, we hypothesized that information and avoidance should be preferred

Table 3.1: Descriptive Statistics for Ego Utility

	Costly Avoid N=198	Costly Info N=207
Gender:		
Female	63 (31.8%)	80 (39.0%)
Male	135 (68.2%)	125 (61.0%)
Age	20.3 (1.41)	20.6 (1.99)
Risk	5.33 (1.93)	5.51 (1.83)
# of Correct Answers	8.10 (3.48)	8.23 (3.15)

Note: Standars errors are in parentheses. This table provides an overview of our sample in the Ego Utility experiment. As can be seen, between two treatments we observe a statistically significant difference for levels of Willingness to Pay and the share of participants who entered positive amounts.

more when they are free. However we observe that this does not hold for either information or avoidance.

One direct motivation for wanting to receive information is clear; it increases monetary payoff. Information avoidance, however, is driven by more psychological motives; the information could be detrimental to their ego by showing a low level of intelligence. Having both, a fraction of positive payers and a positive average values of WTP in both treatments indicates that, participants consider both the monetary and ego utility of information. The monetary value of information is significantly valued at an higher level than the ego—self-image—effect. The distribution of WTP for information and for avoidance can be seen in Figure 3.3. The value of information is determined by participant’s subjective beliefs about their own rank; for a rational risk-neutral agent without ego-utility, it is bounded by completely non-informed (uniform) priors, for which the value of information is 47 NOK (the optimal WTP is 50). The treatments were intended to measure WTP for information and ignorance separately by changing the default setting of the game. This feature of the game makes these results not comparable between receiving and avoiding information.

Based on the previous literature, we hypothesized that the main motivation for information avoidance might be image (ego utility) protection. According to previous research (Eil and Rao, 2011; Kőszegi, 2006; Kőszegi et al., 2019) ego utility is one of the prominent reasons to avoid




Figure 3.2: Average WTP for Information and Avoidance

Note: This figure shows the average WTPs for both information and avoidance. For avoidance the average WTP is 23.7 (sd.=14.8) and for information it is 29.7 (sd.=14.5). This difference is significant ($p < 0.001$).

information. If self-image is vulnerable, avoiding potentially hurtful information is preferable to prevent one’s own ego being hurt. To measure how vulnerable their guess is, we asked participants how certain they are about the correctness of their guess on the relative performance. We pre-specified that if the uncertainty about the guess is high, avoidance is preferable since the information is potentially more hurtful. To explain WTP for information and avoidance, we document results from an OLS regression in Table 3.2. The first variable we use in this table is “# of Questions Seen”. This variable is the number of questions that a participant saw in 8 minutes during the intelligence test. the number of questions that a participant saw in 8 minutes during the intelligence test. We introduce this variable as an indicator of the progress in the test, which is the only signal for participants as well since they can only see how many questions they have answered and not how many correct answers they had. In Table 3.2, the other variables are the guess on what share of all participants performed better than them (between 0 and 100), how certain they are about this guess (between 0 and 100) and the self-reported risk preferences. As seen in the table,

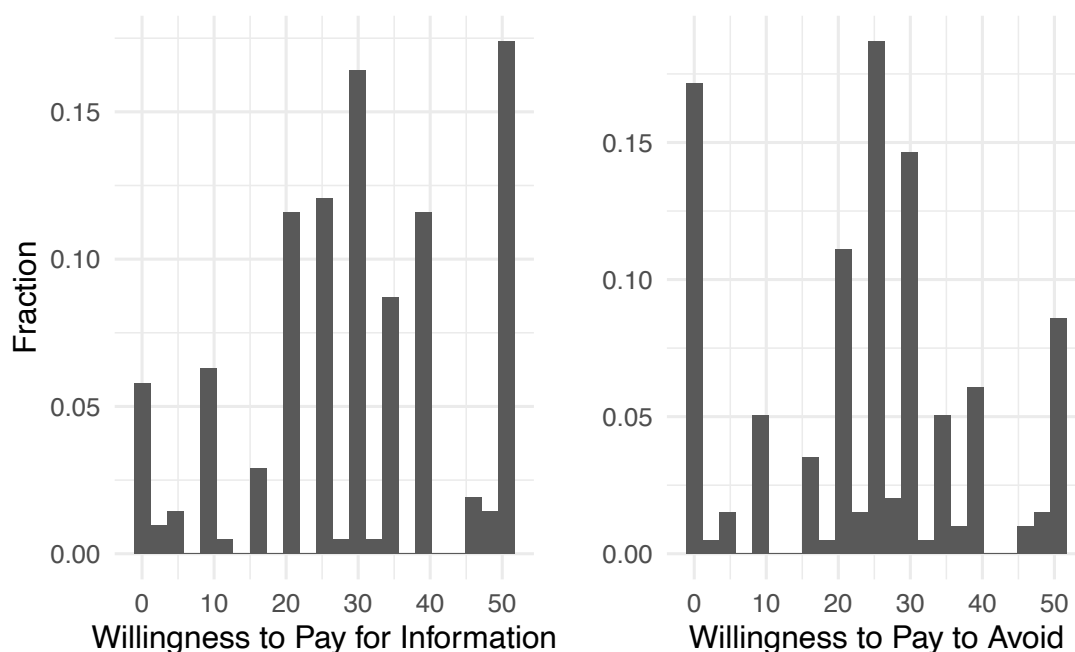


Figure 3.3: WTP for Information and Avoidance

Note: This figure shows the distribution of WTPs for both information and avoidance. As can be seen, a higher share of participants state high WTP for information compared to WTPs for avoidance.

the covariates have little explanatory power for the WTP for information and avoidance.

Result 2: The willingness to pay to receive and avoid information is not explained by confidence and / or certainty about one's own relative performance.

This result is in contradiction with the previous literature and our pre-registered hypotheses. For example, Eil and Rao (2011) and Kőszegi (2006) show that information avoidance in case of ego related tasks is driven by self-confidence. In our analysis we use guess about one's own

Table 3.2: Willingness to pay

	<i>Dependent variable:</i>	
	Willingness to Pay	
	Costly Info.	Costly Avoid.
	(1)	(2)
# of Quest. Seen	-0.0004 (0.186)	0.070 (0.177)
Guess	3.534 (5.303)	0.912 (5.484)
Certainty	-0.002 (0.045)	-0.005 (0.046)
Risk	-0.486 (0.596)	0.484 (0.584)
Gender (M)	-0.861 (2.517)	-1.518 (2.195)
Age	-0.606 (0.779)	0.123 (0.523)
Observations	197	204

Note: *p<0.1; **p<0.05; ***p<0.01

Note: Standars errors are in parentheses. In this table, we show the WTP in relation with a group of covariates. # of Quest. Seen shows the number of questions that are seen during the test. Guess shows the guess on what share of all participants performed better than them. Certainty shows the answer for the question "What are the chances that your guess is within 5% range of actual share of participants that performed better than you?" on a scale between 0 and 100. Risk is the self-reported risk preferences.

relative performance and the certainty—how certain they are about their guess—as indicators of confidence. We do not observe any significant relation between these variables and WTP for information/avoidance. Even though our results document positive WTP for avoidance, behavioral measures that are collected in our experiment do not explain the demand for avoidance.

Figure 3.4 shows the relation between the number of questions seen and a participant’s guess on their relative performance. As can be seen in the figure, the number of questions they see during the test is negatively related with the guess. This means that, as they see more questions, they guess a lower share of participants performed better than them. We find this for both male and female participants and does not vary by gender. This result can also be seen in Table 3.3. Although the magnitude is small, the number of questions seen during the task is positively related to a participant’s guess on how well they did in the test.

Result 3: Progress in the test is positively related with the guess on one’s own relative performance. Participants, who see a higher number of questions tend to guess lower




Figure 3.4: Guess and Number of Questions Seen

Note: This figure shows how the guess on own’s rank in the intelligence test is related to the number of questions they see in the experiment. We use the number of question seen as an indicator of performance in the test, because participants do not see whether their answers were correct or not but only how many questions they went through. Number of questions is the only indicator to have a perception of how well they did in the test.

shares for people performed better than them.

Our findings show that people avoid information about their relative rank in an intelligence test, even when they can earn higher monetary payoffs by receiving it and avoidance requires them to incur a monetary cost. However, WTP for avoidance does not seem to be affected by the cost of avoidance in terms of both WTP and share of avoiders. Our design closely follows the previous literature on information avoidance and Ego Utility by introducing *sensitive* information about one's own relative intelligence. This information can potentially hurt ego by showing lower than expected rankings if the self-esteem is fragile (Kőszegi, 2006; Kőszegi et al., 2019). Even though we document that avoidance is preferred and individuals are willing to pay for it, we are not able to explain it with their self-reported risk preferences, their guess on how well they did and how certain they think the guess is correct. In the previous literature, female participants are shown to have lower confidence. However, our results do not confirm these previous findings and we do not observe any gender difference in the guess that they make about their own performance. Although we do not have any behavioral measure to investigate this result further, this might be

Table 3.3: Guess on “What share of all participants performed better than you?”

	<i>Dependent variable:</i>	
	Guess	
	(Costly Info.)	(Costly Avoid.)
# of Quest. Seen	−0.005** (0.002)	0.003 (0.002)
Certainty	−0.0002 (0.001)	−0.001* (0.001)
Gender (M)	−0.033 (0.031)	−0.019 (0.028)
Observations	198	205

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Note: Standars errors are in parentheses. This table shows the estimation for the guess on how well they performed in the task. As can be seen in the table, number of questions they see during the test has significant effect on the guess only in one treatment.

due to females being equally confident since they are in a selective school. This partially supports the findings from Almås et al. (2016) who show there are very few differences between adolescent boys and girls in Norway.

Overall, our results from the *Ego Utility* experiment show a high demand for information and avoidance. This indicates that a behavioral noise exists when people make decisions over information collection. Therefore, it is necessary to investigate whether our finding that people avoid potentially ego-relevant information is entirely driven by that noise. To that purpose, we conducted a second experiment in which ego utility is not present and the information is not related to any individual characteristics or ability. We provide findings from the *Neutral Information* experiment in the following section. This second experiment provides us with measures for participants' demand for or avoidance of non-ego related, neutral information which allows us to see the effects of ego utility on information choice.

3.3 Neutral Information and Information Preferences

After conducting the *Ego Utility* experiment, we decided to run an identical experiment with one significant difference: information is neutral, that is does not concern any important personal or social characteristics of people. If information avoidance is motivated by self-image concerns, then people should be unwilling to avoid information that does not concern their self-image.

In the first subsection we present the details of our design, in the second subsection we give information about procedures of the experiment and sample, and lastly in the third subsection we report the findings from the *Neutral Information* experiment and compare *Ego Utility* and *Neutral Information*.

3.3.1 Experimental Design

For neutral information, we use a random number as the content in this second experiment. People choose whether they want to receive information on the randomly drawn number. As in the first

experiment, this information can be used to increase monetary payouts from the experiment. We introduce another variation in this experiment: The value of information. We vary the value by having two different probabilities that the information will help the participants to increase their earnings. Varying the earning from information enables us to see whether choice over information and WTP for it to be implemented react to the monetary value of information. The mechanism we are after is very simple: if the potential earnings from information is higher, demand and WTP for it should be higher. By providing results on information preferences for Neutral Information we aim to explain results from the Ego Utility experiment better.

Different from guessing their own rank in an intelligence test, we ask subjects to guess a randomly generated number which is neither related to their performance nor their rank in the intelligence test. If the motivation to avoid information is to protect self-image, such information should not lead to a demand for avoidance. Our design aims to pin down two sources of information preferences when the information has no effect on ego utility. The first one is the value of information. We investigate it by introducing two different levels of value of receiving the information. The second is the willingness to receive and avoid neutral information. To investigate these two channels we used a 2×2 design in which the default settings on information (costly avoidance and costly information) are combined with 2 levels of value of information; high and low. With this design we aim to investigate avoidance and acquisition when neutral (not ego related) information has a clear monetary value.

The experiment proceeds very similarly to the *Ego Utility* experiment which can be seen in Figure 3.1. The design first differs on the “Guess” stage, where participants enter their guess on their relative rank in the *Ego Utility* experiment. In the *Neutral Information* experiment participants are informed that a random number between 1-9 is chosen by the computer and if they guess it correctly they will earn an additional monetary payoff. As in the *Ego Utility* experiment, in the “Certainty” stage people stated how certain they are about their guess on the random number. We wanted to have this measure first to make the design identical to the Ego Utility experiment,

and second to see whether there is a difference on how accurate people think their guess is when the content of information changes. In the *Neutral Information* experiment, in case of receiving information about the randomly drawn number, participants are given the chance to update their guess and earn the additional payoff, 80 NOK as in the *Ego Utility* experiment.

The information set in this experiment is different from the *Ego Utility* experiment, in case of receiving information they see a set of numbers (2 or 4). All the participants are informed that one of the numbers in this information set is the correct number. We manipulate the value of information by changing how many numbers are shown in case of receiving the information. In the high value treatment, participants are given the chance to see a set of 2 numbers and informed that one of them is the randomly drawn number. In the low value treatment they can choose to see a set of 4 numbers and are informed that one of them is the randomly drawn number. If they chose to receive information they are shown the set of 2 (4) numbers and given chance to update their guess. The manipulation on the value of information changes the likelihood of the revised guess being correct.

The rest of the experiment follows Figure 3.1. Participants state their WTP out of 50 NOK in a BDM auction for their information choice to be implemented. Although the interval and the bonus are the same as in the *Ego Utility* experiment, here the value of information is different both in terms of ego utility and monetary utility. In *Neutral Information*, the information does not have any impact on ego utility because it is irrelevant to any individual information. In terms of monetary utility, after receiving the information their updated guess may be correct with 50% or 25% chance whereas in the *Ego Utility* experiment the information is only about the exact rank and the updated guess is correct with 100% chance. As in the *Ego Utility*, participants are randomly assigned first to one of the two conditions; *Costly Information* and *Costly Avoidance*. After this first variation they are assigned one of the two variations: high or low value information.

The design of the control experiment with *Neutral Information* closely follows our *Ego Utility* experiment with the difference of information being neutral. This feature enables us to

estimate the real size of the Ego Utility that we measured in the first experiment.

3.3.2 Sample and Procedure

We followed an almost identical procedure to the first experiment and registered a pre-analysis plan with AEA RCT Registry prior to data collection (Ay and Meißner, 2019). As in the *Ego Utility* experiment, first year students were invited to participate in the experiment before any assignments and exams. Table 3.1 shows the descriptives of the data collected in the Neutral Information experiment. As can be seen, our sample is well randomized across treatments. In terms of observable characteristics and performance measures, our two samples (*Ego Utility* and *Neutral Information*) are quite similar.

3.3.3 Results

The *Neutral Information* experiment focuses mainly on two dimensions of information preferences: value of information and avoidance vs. receiving information. We aim to investigate information preferences when the information does not have any ego-relevant component. This design enables us to document how preferences toward Neutral Information differ from preferences toward ego related information.

We observe that the majority of participants entered positive WTP for information. For both

Table 3.1: Descriptive Statistics for Neutral Information

	Costly Avoid _{high} N=102	Costly Avoid _{low} N=104	Costly Info _{high} N=97	Costly Info _{low} N=96
Gender:				
Female	39 (38.2%)	40 (39.6%)	40 (41.7%)	32 (33.3%)
Male	63 (61.8%)	61 (60.4%)	56 (58.3%)	64 (66.7%)
Age	20.6 (1.41)	20.4 (1.46)	20.6 (3.64)	20.7 (1.41)
Risk	6.64 (2.11)	6.56 (2.04)	6.19 (1.78)	6.76 (2.15)
Correct Answer	8.97 (2.36)	9.32 (2.07)	8.61 (1.98)	8.83 (2.24)

Note: Standars errors are in parentheses. This table provides an overview of our sample in the *Neutral Information* experiment. As can be seen, our sample is quite similar to the one in the *Ego Utility* experiment in terms of demographics, score and risk behavior. Between treatments we observe a statistically significant difference for levels of Willingness to Pay and the share of participants who entered positive amounts.

high and low value information, 93.8% of the participants entered a positive amount for WTP to receive information (see Figure 3.1). This result is partly expected since we hypothesized that when the information has no effect on ego and has monetary value it should be demanded on a higher level. For the same reason, we expected almost all the participants to enter a positive WTP to receive information. However, our findings show that approximately 6% of participants did not enter a positive amount. Although we do not observe any difference between information value treatments on the share of participants who are willing to pay a positive amount for the information, we observe a significant change in the average WTP. The WTP for information reacts to the value of information in the hypothesized way: when information has high value (seeing 2 numbers) average WTP (28.2, sd.=13.4) is significantly higher than when the information has low value (seeing 4 numbers, 24.5, sd.=12.3, $p = 0.013$).

For avoidance, our findings deviate from our pre-registered hypotheses. High shares of people are willing to pay a positive amount to avoid information in both treatments. When the information has high value (2 numbers are shown), 83.3% and when the information has low value

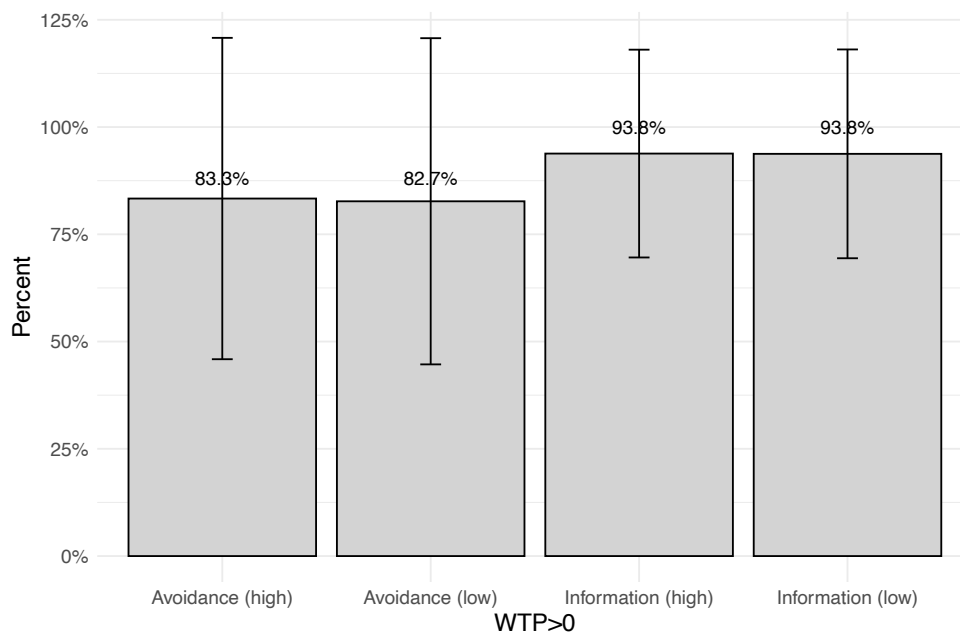


Figure 3.1: Share of Positive WTP

Note: Bars are standard errors. In this figure, we report the share of participants who avoided information by treatments. We observe a high share of participants for both choosing avoidance and information.

(4 numbers are shown) 82.7% of the participants entered positive WTP to avoid information. Although we expected to capture the behavioral noise that people have when making decisions over information, we did not hypothesize avoidance from Neutral Information to be as high as in our results. In Figure 3.2, we report the average WTPs for information and avoidance. We do not observe any effect of value of information on WTP to avoid information. WTP to avoid high value information (13.9, $sd.=13$) is not significantly different from WTP for low value information (14.8, $sd.=11.8$, $p = 0.29$). This might suggest that the willingness to pay a positive amount for avoidance is mostly behavioral noise and is not driven by any feature of the information.

Figure 3.3 shows the distribution of WTP by treatments for the *Neutral Information* experiment. As can be seen, WTP for information in both high and low values condition is much higher compared to WTP for avoidance. The difference in WTP to receive high and low value information is clear in the figure. This difference is not observed for avoidance.

Result 1: WTP to receive information increases when the value of information is higher. WTP to avoid information does not respond to value of information.




Figure 3.2: Average WTPs for Information and Avoidance

Note: In this figure, we report the average WTPs for both avoidance and information. We compare WTPs for high and low value information. For avoidance, on the one hand, average WTP does not vary by the value of information ($p = 0.29$). On the other hand, for information we observe significantly higher WTP when the value is high ($p = 0.013$).

This result shows that demand for neutral information responds to the value of information whereas demand for avoidance does not vary. In this experiment motivation for avoidance is not clear since the information is about a random number rather than any personal attributes. This finding suggests that people have a substantial behavioral noise when making decisions over information in a BDM auction.

When these findings are compared to ego related information although the level of WTP to avoid is lower, we find that the majority of participants enter positive amounts in both cases. If the measured levels of WTP in the *Ego Utility* experiment reflects only ego utility, we would not see high shares of positive payers avoiding neutral information. Based on the findings from the *Neutral Information* experiment, our results suggest that the scales that are used to measure Ego Utility might contain substantial noise. BDM auctions are commonly used for such measures and theoretically lead subjects to enter their true preferences since the dominant strategy is entering one's own valuation. WTPs that are measured with BDM, enable us to see the difference between neutral information and ego utility and the findings suggest that if this mechanism is used only for one of these, measures might capture more channels than Ego Utility.

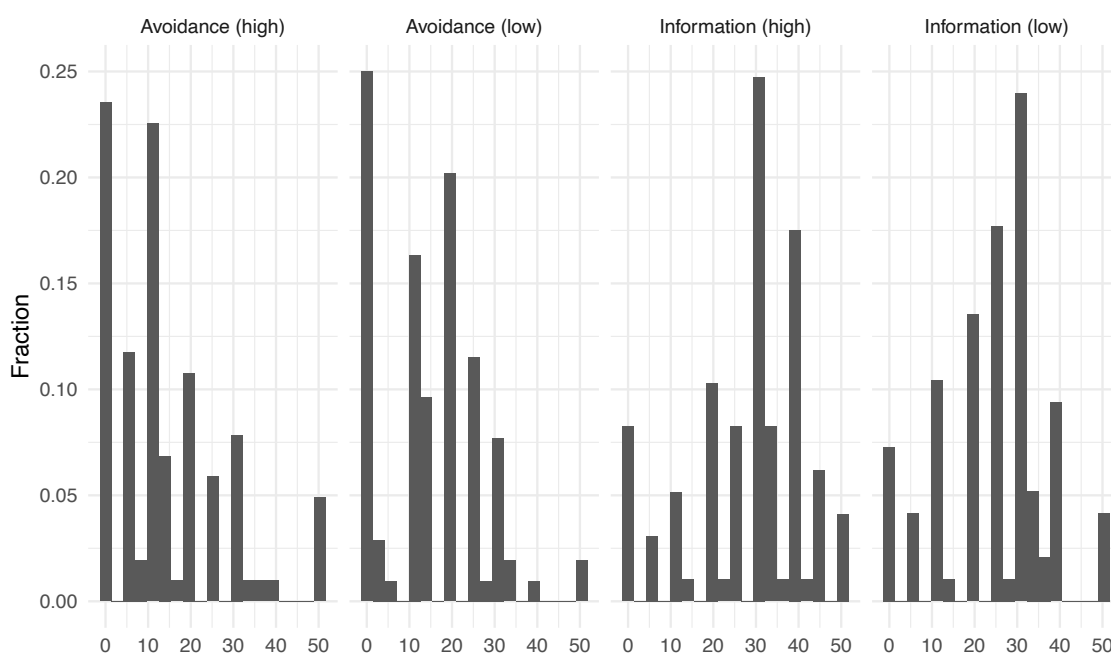


Figure 3.3: Distribution of WTP by treatments

Table 3.2 aims to find determinants of WTP to receive and to avoid neutral information. Each column of the table refers to one treatment in the experiment. In the fourth column, when the information has low value males enter significantly lower WTPs. However neither the magnitude nor the sign of the gender variable is constant across treatments. Other than that we do not observe a significant relation with any of the covariates.

In the previous literature it is shown that individuals demand “useless” information which has no instrumental value to feel good about themselves (Eliaz and Spiegler, 2006).⁴ In some cases receiving information could be preferred if the information signals positive attributes of oneself, such as high intelligence. Avoiding information that does not contain anything potentially hurtful has not been a topic of interest in behavioral economics. This is partly due to the problem of ratio-

Table 3.2: Regressions of WTA and WTP on situation variables in the Control Experiment

	<i>Dependent variable:</i>			
	Willingness to Pay			
	Costly Avoid _{high}	Costly Avoid _{low}	Costly Info _{high}	Costly Info _{low}
	(1)	(2)	(3)	(4)
# of Quest. Seen	0.390 (0.290)	0.388 (0.284)	0.539* (0.313)	-0.150 (0.271)
Risk	-0.561 (0.791)	0.418 (0.734)	0.941 (0.853)	0.053 (0.670)
Gender (M)	4.972 (2.994)	-2.626 (2.983)	1.839 (3.251)	-7.200** (3.307)
Controls	Yes	Yes	Yes	Yes
Observations	102	100	96	96

Note:

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Note: Standars errors are in parentheses. Control variables are the unobserved self-reported characteristics that are collected in the post-experiment survey. It includes self-reported general success, intelligence and importance of intelligence with demographics.

⁴In this paper, individuals demand to learn how they performed in an intelligence test even when there is no monetary gain. The authors argue that this is to satisfy Ego Utility.

nalizing such behavior. In psychology, however, some research documented that uncertainty can give pleasure. Wilson et al. (2005) provide experimental evidence that after good news, avoiding information could extend the positive feelings. Wilson et al. (2013) argue that uncertainty could make “thrill” persist longer which can lead to avoidance. Sunstein (2019) describes this situation as “ruining” the joy by revealing information. Our experiment does not document any of these channels, yet we find it valuable to note as potential motivations to avoid neutral information.

Our findings provide a weak evidence for the motivations for information avoidance. We are unable to explain avoidance of neutral information with any mechanism that we have in the experiment. However, we define this as a behavior that stems from the noise in the individual choices. People may fail to choose the best option in terms of monetary outcome or ego utility, and they might have some random utility from making other, even dominated, alternatives (Luce, 1959). This means, when people are given the alternatives, some may always choose the options that are not the dominant ones. In both *Ego Utility* and *Neutral Information* experiments, we gave people the chance to use the BDM scale to state their preference only on one direction—receiving or avoiding information. At this point, we continue our analysis by considering the WTP measures that we collected in the *Ego Utility* experiment may contain a substantial amount of behavioral noise due to use of the scale which is constrained to only positive numbers. Incentives are similar in both experiments, which means similar behavioral noise might be the reason for the heterogeneity in WTP for avoidance in the *Neutral Information* experiment as well. By comparing the effects of the ego utility and image concerns with the effects of random utility, we show that despite the behavioral noise in both experiments, being exposed to ego related information and self-image concerns increases the demand for information avoidance. Even if some part of the WTP for avoidance in the *Ego Utility* experiment stems from the behavioral noise and not the concerns about self-image, we aim to capture the part that is caused by self-image concerns with this comparison. We use a regression analysis for such comparison in which we show the effects of being in the *Ego Utility* experiment on WTP and then the interaction between being in the *Ego Utility*

experiment and in the *Costly Avoidance* treatment. By using such specification, we expect to have an increasing effect of being in the Ego Utility experiment and in the Costly Avoidance treatment on WTP if the concerns about self-image are effective in the *Ego Utility* experiments. Our analysis captures the diff-in-diff effects of two experiments and avoidance treatments. Another benefit of using such model is to capture any differences between two samples in our two experiments.

We report this analysis in table 3.3. The first column of Table 3.3 shows the analysis without control variables. In this column the positive effect of being in the *Ego Utility* experiment is highly significant and can be seen in the first row. “Ego×Avoidance” variable captures the effect of Ego Utility on WTP for avoidance. As can be seen, facing a trade off between Ego Utility and monetary gains is related to a significantly higher WTP to avoid information compared to the information that is neutral—not ego related—and about a randomly drawn number. This shows a higher demand for avoidance in the *Ego Utility* experiment.

Result 2: When the information is about the relative rank in the intelligence test, WTP to avoid information is higher compared to WTP to avoid Neutral Information.

This result clarifies our findings by showing that even though people avoid both, ego relevant and random information, the willingness to pay to avoid ego relevant information is higher. Our results for the *Neutral Information* experiment shows that when faced with similar scales people may have behavioral noise, and the measured WTP for information choice might be driven by various factors. However, even with the noise, ego utility is stronger than “random utility” in the *Neutral Information* experiment and leads to higher WTP to avoid information.

3.4 Conclusion

In many social and economic decisions beliefs about one’s own abilities play an important role. For example decisions on what to study, which career path to follow, which tasks to sign up for are all based on one’s perception of own abilities and skills. Evidence from economic theory and

experiments show that these beliefs are important to people and when they receive information about their own abilities, they tend to process it in a biased manner (Karlsson et al., 2009; Kőszegi, 2006; Kőszegi et al., 2019; Eil and Rao, 2011; Buser et al., 2018; Möbius et al., 2011). If the information shows lower abilities than expected, it might be detrimental to ego. In this paper, we show that when there is a chance of receiving “bad news”, people choose to avoid information about their abilities even when that means that they do not receive a monetary reward. Our results show that people avoid information about their relative intelligence and they are willing to forego

Table 3.3: Overall effect of Ego Utility on WTP

	<i>Dependent variable:</i>	
	Willingness to Pay	
	(Without Controls)	(With Controls)
Ego Utility	3.382** (1.371)	3.172** (1.401)
Avoidance	-11.983*** (1.373)	-12.177*** (1.381)
Ego×Avoidance	5.946*** (1.934)	6.410*** (1.948)
Risk		0.068 (0.263)
Gender (=M)		-1.192 (1.069)
Constant	26.347*** (0.986)	26.739*** (1.893)
Observations	804	797

Note: *p<0.1; **p<0.05; ***p<0.01

Note: Standars errors are in parentheses. This table shows the relation between Ego Utility and WTP for the decision on information to be implemented. As can be seen, being in the Ego Utility experiment has an increasing effect on the WTP by itself. Effects of being in the Ego Utility experiment and in the Avoidance treatment are captured by the interaction term Ego×Avoidance. As can be seen, WTP to avoid information is significantly higher in the Ego Utility experiment compared to Neutral Information. The first column of the table shows this analysis without controlling for gender and self-reported risk whereas the second column shows the results when controlling for those two covariates. In the second column 7 observation is lost due to absence of responses to control questions.

monetary gains to do so.

Contrary to the previous experiments, we do not find an association between confidence and information preferences. Instead, we provide evidence in favor of no relationship between information avoidance, guess about own performance and certainty about this guess. In the second experiment, we investigate whether avoidance is solely motivated by Ego Utility. Our results show that a high fraction of participants avoid information even when it is not ego related and even when avoidance means foregoing higher monetary payoffs. We find that WTP to receive information is positively associated to the value of information, whereas this relationship is not present for avoidance. Overall, the findings from the second experiment show that people may avoid information even when it has no effect on ego utility—self-image. This indicates that behavioral noise has some impact on people’s decisions on whether they want to pay to avoid information.

Taken together, the two experiments show that people are willing to pay to avoid information in settings in which this protects their self-image and in settings in which it does not. This finding suggests that it is not easy to interpret participants’ demand for information avoidance in the first experiment as purely driven by self-image protection, as it indicates that part of the demand is due to behavioral noise. However, the comparison of the WTP to avoid in the first experiment to the second experiment shows that participants in the *Ego Utility* experiment are willing to pay significantly more to avoid information. This points to the conclusion that ego utility does in fact at least partially explain participants’ willingness to pay in the first experiment.

In this paper, we document findings on avoidance of ego related information. Our findings are relevant to many economic and social decisions that require processing information about abilities and skills such as education and career choices. We show that no matter how confident the people are, they avoid receiving information about their intelligence. Our findings contribute to the understanding of information preferences which in real life can cause inefficiencies.

Bibliography

- Almås, I., Cappelen, A. W., Salvanes, K. G., Sørensen, E. Ø., and Tungodden, B. (2016). Willingness to compete: Family matters. *Management Science*, 62(8):2149–2162.
- Ay, F. C. and Meißner, S. (2018). Investigating motivations for information avoidance - the role of value of information. *AEA RCT Registry, AEARCTR-0003473*.
- Ay, F. C. and Meißner, S. (2019). Investigating motivations for information avoidance - the role of value of information. *AEA RCT Registry, AEARCTR-0004800*.
- Becker, G. M., DeGroot, M. H., and Marschak, J. (1964). Measuring utility by a single-response sequential method. *Systems Research and Behavioral Science*, 9(3):226–232.
- Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly journal of economics*, 117(3):871–915.
- Buser, T., Gerhards, L., and Van Der Weele, J. (2018). Responsiveness to feedback as a personal trait. *Journal of Risk and Uncertainty*, 56(2):165–192.
- Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, 22(2):369–395.
- Eil, D. and Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):114–138.
- Eliaz, K. and Schotter, A. (2010). Paying for confidence: An experimental study of the demand for non-instrumental information. *Games and Economic Behavior*, 70(2):304–324.
- Eliaz, K. and Spiegel, R. (2006). Can anticipatory feelings explain anomalous choices of information sources? *Games and Economic Behavior*, 56(1):87–104.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

- Ganguly, A. and Tasoff, J. (2017). Fantasy and dread: The demand for information and the consumption utility of the future. *Management Science*, 63(12):4037–4060.
- Karlsson, N., Loewenstein, G., and Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty*, 38(2):95–115.
- Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, 77(2):603–606.
- Kőszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, 4(4):673–707.
- Kőszegi, B., Loewenstein, G., and Murooka, T. (2019). Fragile self-esteem. *Available at SSRN 3412006*.
- Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. John Wiley, New York.
- Möbius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2011). Managing self-confidence: Theory and experimental evidence. *NBER working paper*, 17014.
- Niederle, M. and Vesterlund, L. (2007). Do women shy away from competition? do men compete too much? *Quarterly Journal of Economics*, 122(3):1067–1101.
- Sunstein, C. R. (2019). Ruining popcorn? the welfare effects of information. *Journal of Risk and Uncertainty*, 58(2-3):121–142.
- Sørensen, E. (2011). Expmotor: Innføring for brukere. *mimeo*.
- Wilson, T. D., Centerbar, D. B., Kermer, D. A., and Gilbert, D. T. (2005). The pleasures of uncertainty: prolonging positive moods in ways people do not anticipate. *Journal of Personality and Social Psychology*, 88(1):5.
- Wilson, T. D., Ndiaye, D. G., Hahn, C., and Gilbert, D. T. (2013). *Still a thrill: Meaning making and the pleasures of uncertainty*. American Psychological Association.

3.A Supplementary Analysis




Figure 3.A.1: Progress of the IQ Test in Ego Utility




Figure 3.A.2: Progress of the IQ Test in Neutral Information

3.B Ego Utility Instructions

Please look at your screen now. I am reading from the instructions displayed on your screen.

I will now describe the test which makes up the second part of the session. The test is an intelligence test made up of 25 questions, taken from a longer test.

For every question, there is a pattern with a piece missing and a number of pieces below the pattern. You have to choose which of the pieces below is the right one to complete the pattern. You will see 8 pieces that might complete the pattern for each question. In every case, one and only one of these pieces is the right one to complete the pattern.


For each question, please enter your answer in the column to the right of the pattern. You will score 1 point for every right answer. You will not be penalized for wrong answers.

You will have 8 minutes to complete all questions. During the test, you can move back and forth between the 25 questions in that part and you can change your previous answers. The top right-hand corner of the screen will display the time remaining (in seconds).

Before we start the test, please raise your hand if you have any questions. During the test, please raise your hand if you have a problem with your computer or a question regarding the experiment.

Remaining time: 592

Please enter your answer to this question in the column to the right of the pattern.
You can move back and forth between the 25 questions in this part using the red buttons and you can change your previous answers.
Once the 10 minutes for this part are over, you will not be able to change your answers to this part of the test.



Please enter your answer to this question in this column.

Which one is the missing part of the pattern?

- A
- B
- C
- D
- E
- F
- G
- H

Next Question

Period 1 of 1 Remaining time (sec): 41

You are now asked to estimate what percentage of the participants in this room performed better than you. The number of correct answers in the test will be compared to all other students in the room. Your payment will be determined by how well you guess.

The precision of your estimate determines your payment for the experiment. If your guess is within 0.05 range of actual share of people performing better than you, you receive 80 kr. Otherwise, you earn nothing.

Please enter your guess about your relative rank. For this purpose, please indicate your belief over what percentage of participants did better than you. If you for example believe that 50% of participants did better, please enter 0.5. If you believe everyone did worse than you, please enter 1.0. If you believe everyone did better than you, please enter 0.0. Your answer has to be between 0 and 1.

Your estimate:

OK

Period 1 of 1 Remaining time [sec]: 39

What are the chances that your guess is within 5% range of actual share of the participants that performed better than you?

Please indicate a probability in percentage, your answer has to be between 0 and 100. Eg. If you enter 20 that means you are 20% certain that your guess is correct. Your answer should be between 0 and 100.

OK

Period 1 of 1 Remaining time [sec]: 40

You are now given the chance to find out the exact percentage of participants who performed better than you. If you learn that, you will be given the chance to change your guess in the next screen. There is a price to get that information. For this purpose, you are given 50kr. You can decide how much are you willing to pay to learn that information.

The price you enter will be compared to a price which is randomly drawn by a computer. This price can be between 0 and 50 and all numbers within this range are equally likely. Consider following example scenarios:

- 1) If the computer chooses a price that is below your maximum willingness to pay for the information, you will receive the information and pay the price that the computer chose. If you for example stated 35kr and the computer draws a 30, you will receive the information, pay 30kr and keep 20kroners from the bonus.
- 2) If the computer chooses a price that is above your maximum willingness to pay for the information, you will not receive the information and keep the 50kr. If for example you stated 35kr and the computer draws a 40, you will not receive the information and keep the 50kr.

Once you have decided how much you want to pay, the best strategy in terms of monetary gain is always to enter that amount. Remember that a correct guess payout 80kr. You can state a price between 0kr and 50kr as your maximum willingness to pay for that information.

Enter the amount you are willing to pay for information:

Submit My Answer

Period 1 of 1 Remaining time [sec]: 25

The price you entered was higher than the randomly drawn price by the computer, you are given the information below.

Randomly drawn price	30
Amount you earned from bonus	20
Your rank:	1.00

As mentioned at the beginning, you can earn 80 kroners from your guess. You will learn if you earned it or not at the final payoff screen.

OK

Periode

1 von 1

Verbleibende Zeit [sec]: 84

Please answer the questions below considering general situations.

How well do you think other participants guessed their share of participants who performed better than them?

- Higher than their actual rank
- Lower than their actual rank
- Accurate

Would the share of participants who performed better than you change if there were students from 2nd and 3rd years taking part?

- Would be lower
- Would be higher
- Wouldn't change

Would your guess over the share of participants who performed better than you change if there were students from 2nd and 3rd years taking part?

- Would be lower
- Would be higher
- Wouldn't change

Think about the other participants in this room, would the share of participants who performed better than them change if there were students from 2nd and 3rd years taking part?

- They would report lower guesses
- They would report higher guesses
- Their guess wouldn't change

Think about the other participants in this room, would their guess over the share of participants who performed better than them change if there were students from 2nd and 3rd years taking part?

- They would report lower guesses
- They would report higher guesses
- Their guess wouldn't change

How do you think about your academic success compared to other students?

- More successful than average
- Average
- Less successful than average

How do you think your own intelligence compared to other students?

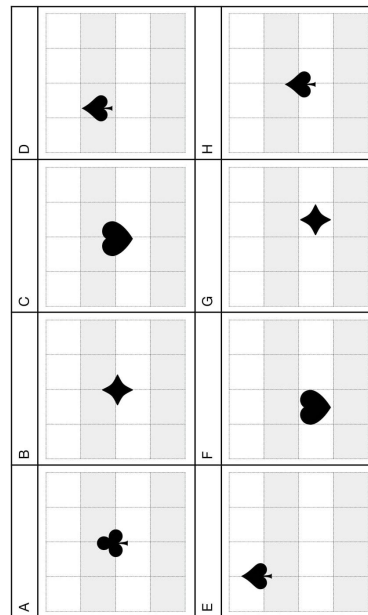
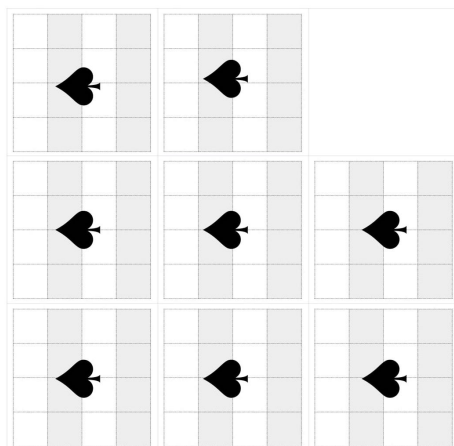

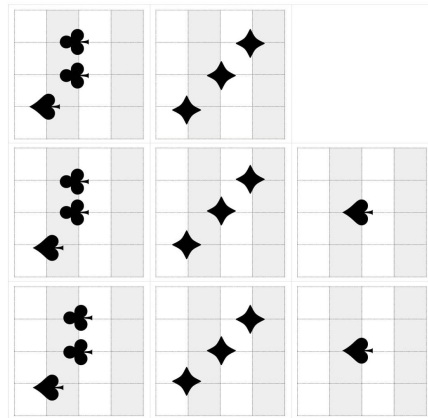
- More intelligent than average
- Average
- Less intelligent than average

How important is it for you to think yourself as an intelligent person?

Not Important At All Very Important

Continue

3.B.1 Examples of Progressive Matrices



Chapter 4

National Identity Predicts Public Health

Support During a Global Pandemic:

Results from 67 Nations

⁰This work is from an international research collaboration with more than 150 researchers from 67 countries. This work was partially supported by the Research Council of Norway through its Centres of Excellence Scheme, FAIR project No 262675. For further details see [The project website](#). For the most recent version of the publicly available pre-print, visit <https://psyarxiv.com/ydt95>.

Title: *National identity predicts public health support during a global pandemic: Results from 67 nations*

Authors & Acknowledgments:

Jay J. Van Bavel¹, Aleksandra Cichocka^{2,3}, Valerio Capraro⁴, Hallgeir Sjåstad^{5,6}, John B. Nezlek^{7,8}, Tomislav Pavlović⁹, Mark Alfano¹⁰, Michele J. Gelfand¹¹, Flavio Azevedo¹², Michèle D. Birtel¹³, Aleksandra Cislak⁷, Patricia L. Lockwood^{14,15}, Robert M. Ross¹⁰, Koen Abts¹⁶, Elena Agadullina¹⁷, David M. Amodio^{1,18}, Matthew A. J. Apps¹⁹, John Jamir Benzon R. Aruta²⁰, Sahba Besharati²¹, Alexander Bor²², Becky L. Choma²³, Charles Crabtree²⁴, William A. Cunningham²⁵, Koustav De²⁶, Waqas Ejaz²⁷, Christian T. Elbaek²⁸, Harry Farmer¹³, Andrej Findor²⁹, Daniel Flichtentrei³⁰, Renata Franc⁹, Biljana Gjoneska³¹, Estrella Gualda³², Yusaku Horiuchi²⁴, Toan L. D. Huynh³³, Mostak A. Imran³⁴, Agustin Ibanez^{35,36,37,38,39}, Jacob Israelashvili⁴⁰, Katarzyna Jasko⁴¹, Jaroslaw Kantorowicz⁴², Elena Kantorowicz-Reznichenko⁴³, Andre Krouwel⁴⁴, Yordan Kutiyski⁴⁵, Michael Laakasuo⁴⁶, Claus Lamm⁴⁷, Jonathan Levy^{48,49}, Caroline Leygue⁵⁰, Ming-Jen Lin⁵¹, Mohammad S. Mansoor⁵², Antoine Marie^{53,54}, Lewend Mayiwar⁵⁵, Honorata Mazepus⁴², Cillian McHugh⁵⁶, Panagiotis Mitkidis^{28,57}, Andreas Olsson⁵⁸, Tobias Otterbring⁵⁹, Dominic J. Packer⁶⁰, Jussi Palomäki⁴⁶, Anat Perry⁴⁰, Michael Bang Petersen²², Arathy Puthillam⁶¹, Julian Riano-Moreno^{62,63,64,65}, Tobias Rothmund¹², Hernando Santamaría-García^{65,66}, Petra C. Schmid⁶⁷, David Stadelmann⁶⁸, Catalin Augustin Stoica⁶⁹, Drozdstoy St. Stoyanov⁷⁰, Kristina K. Stoyanova⁷⁰, Shruti Tewari⁷¹, Bojan Todosijević⁹, Benno Torgler^{72,73}, Manos Tsakiris^{74,75,76}, Hans H. Tung^{77,78}, Radu Umbres⁶⁹, Edmunds Vanags⁸⁰, Madalina Vlasceanu⁸¹, Andrew J. Vonasch⁸², Meltem Yucel⁸³, Yucheng Zhang⁸⁴, Mohcine Abad⁵⁴, Eli Adler⁴⁰, Narin Akrawi⁸⁵, Hamza Alaoui Mdarhri⁵⁴, Benedict G. Antazo²⁰, Ceren F. Ay⁵, Mouhamadou El Hady Ba⁸⁶, Sergio Barbosa⁸⁷, Brock Bastian⁸⁸, Anton Berg⁴⁶, Maria Paula Bernal Zárata⁶², Michał Białek⁸⁹, Ennio Bilancini⁹⁰, Natalia Bogatyreva¹⁷, Leonardo Boncinelli⁹¹, Jonathan E. Booth⁹², Sylvie Borau⁹³, Ondrej Buchel⁹⁴, C. Daryl Cameron⁹⁵, Chrissie F. Carvalho⁹⁶, Tatiana Celadin⁹⁰, Chiara Cerami^{97,98}, Hom Nath Chalise⁹⁹, Xiaojun Cheng¹⁰⁰, Luca Cian¹⁰¹, Kate Cockcroft²¹, Jane R. Conway¹⁰², Mateo A. Córdoba-Delgado⁶⁵, Chiara Crespi¹⁰³, Marie Crouzevialle⁶⁷, Jo Cutler¹⁹, Marzena Cypryńska⁷, Justyna Dabrowska¹⁰⁴, Michael A. Daniels¹⁰⁵, Victoria H. Davis²⁵, Pamala N. Dayley¹⁰⁶, Sylvain Delouvé¹⁰⁷, Ognjan Denkovski¹⁸, Guillaume Dezechache¹⁰⁸, Nathan A. Dhaliwal¹⁰⁵, Alelie B. Diato¹⁰⁹, Roberto Di Paolo^{90,110}, Marianna Drosinou^{46,111}, Uwe Dulleck^{73,112,113}, Jānis Ekmanis⁸⁰, Tom W. Etienne⁴⁵, Hapsa H. Farhana³⁴, Fahima Farkhari¹², Kristijan Fidanovski¹¹⁴, Terry Flew¹¹⁵, Shona L. Fraser²¹, Raymond B. Frempong⁶⁸, Jonathan A. Fugelsang¹¹⁶, Jessica Gale⁸², E. Begoña García-Navarro³², Prasad Garlandinne⁷¹, Oussama Ghajjou¹¹⁷, Theofilos Gkinopoulos¹³, Kurt Gray¹¹⁸, Siobhán M. Griffin⁵⁶, Bjarki Gronfeldt², June Gruber¹¹⁹, Eran Halperin⁴⁰, Elizabeth Harris¹, Volo Herzon⁴⁶, Matej Hruška²⁹, Guanxiong Huang¹²⁰, Matthias F. C. Hudecek^{121,122}, Ozan Isler⁷², Simon Jangard⁵⁸, Frederik J. Jørgensen²², Frank J. Kachanoff¹¹⁸, John Kahn²⁴, Apsara Katuwal Dango⁵², Oleksandra Keudel¹²³, Lina Koppel¹²⁴, Mika Koverola⁴⁶, Anton Kunnari⁴⁶, Josh Leota¹²⁵, Eva Lerner^{113,122}, Neil Levy¹⁰, Chunyun Li⁹², Juliana Linares Puerta⁶⁵, Elizabeth U. Long²⁵, Chiara Longoni¹²⁶, Marina Maglič⁹, Darragh McCashin¹²⁷, Alexander L. Metcalf¹²⁸, Igor Mikloušić⁹, Asako Miura¹²⁹, Juliana Molina-Paredes⁶⁵, César Monroy-Fonseca¹³⁰, Elena Morales-Marente³², David Moreau¹³¹, Rafał Muda¹³², Annalisa Myer¹³³, Kyle

Nash¹²⁵, Jonas P. Nitschke⁴⁷, Matthew S. Nurse¹¹², Yohsuke Ohtsubo¹³⁴, Victoria Oldemburgo de Mello^{25,96}, M. Soledad Palacios-Galvez³², Jussi Palomäki⁴⁶, Yafeng Pan⁵⁸, Zsófia Papp¹³⁵, Philip Pärnamets^{1,58}, Mariola Paruzel-Czachura¹³⁶, Zoran Pavlović¹³⁷, César Payán-Gómez⁸⁷, Silva L. L. Perander⁴⁶, Michael M. Pitman²¹, Joanna Pyrkosz-Pacyna¹³⁸, Steve Rathje¹³⁹, Ali Raza¹¹⁹, Kasey Rhee²⁴, Gabriel G. Rêgo¹⁴⁰, Claire E. Robertson¹, Iván Rodríguez-Pascual³², Teemu J. Saikkonen¹⁴¹, Octavio Salvador-Ginez⁵⁰, Waldir M. Sampaio¹⁴⁰, Gaia C. Santi⁹⁷, Natalia Santiago Tovar⁶², David A. Savage^{73,142}, Julian A. Scheffer⁹⁵, Philipp Schönegger¹⁴³, David T. Schultner¹⁸, Enid M. Schutte²¹, Andy Scott¹²⁵, Madhavi Sharma⁵², Pujan Sharma¹⁴⁴, Ahmed Skali¹⁴⁵, Clara A. Stafford¹⁴⁶, Dragan T. Stanojevic¹³⁷, Anna Stefaniak¹⁴⁷, Anni Sternisko¹, Brent Strickland⁵³, Jukka Sundvall⁴⁶, Jeffrey P. Thomas⁹², Gustav Tinghög¹²⁴, Iris J. Traast¹⁸, Raffaele Tucciarelli⁷⁵, Michael Tyralla¹⁴⁸, Nick D. Ungson⁶⁰, Mete S. Uysal¹⁴⁹, Jan-Willem van Prooijen⁴⁴, Dirk Van Rooy¹¹², Daniel Västfjäll¹²⁴, Peter P. J. L. Verkoeijen⁴³, Joana B. Vieira⁵⁸, Christian von Sikorski¹⁵⁰, Alexander C. Walker¹¹⁶, Jennifer Watermeyer²¹, Erik Wetter¹⁵¹, Robin Willardt⁶⁷, Michael J. A. Wohl¹⁴⁷, Adrian D. Wojcik³, Kaidi Wu¹⁵², Yuki Yamada¹⁵³, Onurcan Yilmaz¹⁵⁴, Kumar Yogeeswaran⁸², Caroline-Therese Ziemer¹², Rolf A. Zwaan⁴³, & Paulo S. Boggio¹⁴⁰.

¹Department of Psychology & Neural Science, New York University, New York, NY, USA. ²School of Psychology, University of Kent, Kent, UK. ³Department of Psychology, Nicolaus Copernicus University, Toruń, Poland. ⁴Department of Economics, Middlesex University London, London, UK. ⁵Norwegian School of Economics, Bergen, Norway. ⁶Center for Applied Research, Norwegian School of Economics, Bergen, Norway. ⁷Institute of Psychology, SWPS University of Social Sciences and Humanities, Warsaw, Poland. ⁸Department of Psychological Sciences, College of William and Mary, Williamsburg, VA, USA. ⁹Institute of Social Sciences, Belgrade, Serbia. ¹⁰Department of Philosophy, Macquarie University, Sydney, Australia. ¹¹Department of Psychology, University of Maryland, College Park, MD, USA. ¹²Department of Communication and Media Use, Friedrich Schiller University Jena, Jena, Germany. ¹³University of Greenwich, London, UK. ¹⁴Department of Experimental Psychology, University of Oxford, Oxford, UK. ¹⁵Center for Human Brain Health, University of Birmingham, Birmingham, UK. ¹⁶Department of Sociology, Tilburg University, Tilburg, The Netherlands. ¹⁷National Research University Higher School of Economics, Moscow, Russia. ¹⁸Faculty of Social and Behavioral Sciences, University of Amsterdam, Amsterdam, The Netherlands. ¹⁹School of Psychology, University of Birmingham, Birmingham, UK. ²⁰Counseling and Educational Psychology Department, De La Salle University, Manila, Philippines. ²¹Wits University, Johannesburg, South Africa. ²²Department of Political Science, Aarhus University, Aarhus, Denmark. ²³Department of Psychology, Ryerson University, Toronto, Ontario, Canada. ²⁴Dartmouth College, Hanover, NH, USA. ²⁵Department of Psychology, University of Toronto, Toronto, Ontario, Canada. ²⁶Gatton College of Business & Economics, University of Kentucky, Lexington, KY, USA. ²⁷School of Social Sciences and Humanities, National University Science and Technology, Islamabad, Pakistan. ²⁸Department of Management, Aarhus University, Denmark. ²⁹Faculty of Social and Economic Sciences, Comenius University, Bratislava, Slovakia. ³⁰IntraMed, Buenos Aires, Argentina. ³¹Macedonian Academy of Sciences and Arts, Skopje, North Macedonia. ³²University of Huelva, Huelva, Spain. ³³WHU -

Otto Beisheim School of Management, Vallendar, Germany. ³⁴Department of Educational and Counselling Psychology, University of Dhaka, Dhaka, Bangladesh. ³⁵Global Brain Health Institute, University of California San Francisco, San Francisco, CA, USA. ³⁶University of San Andrés, Victoria, Buenos Aires, Argentina. ³⁷National Scientific and Technical Research Council, Buenos Aires, Argentina. ³⁸Autonomous University of the Caribbean, Barranquilla, Barranquilla, Colombia. ³⁹Center for Social and Cognitive Neuroscience, Adolfo Ibañez University, Santiago, Chile. ⁴⁰Department of Psychology, Hebrew University of Jerusalem, Jerusalem, Israel. ⁴¹Institute of Psychology, Jagiellonian University, Kraków, Poland. ⁴²Leiden University, Leiden, The Netherlands. ⁴³Erasmus University Rotterdam, Rotterdam, The Netherlands. ⁴⁴Vrije Universiteit Amsterdam, Amsterdam, The Netherlands. ⁴⁵Kieskompas (Election Compass), The Netherlands. ⁴⁶Department of Digital Humanities, University of Helsinki, Helsinki, Finland. ⁴⁷Social, Cognitive and Affective Neuroscience Unit, Department of Cognition, Emotion, and Methods in Psychology, Faculty of Psychology, University of Vienna, Vienna, Austria. ⁴⁸Baruch Ivcher School of Psychology, Interdisciplinary Center Herzliya, Herzliya, Israel. ⁴⁹Department of Neuroscience and Biomedical Engineering, Aalto University, Espoo, Finland. ⁵⁰School of Psychology, National Autonomous University of Mexico, Mexico City, Mexico. ⁵¹Department of Economics, National Taiwan University, Taipei, Taiwan. ⁵²Central Department of Psychology, Tribhuvan University, Kathmandu, Nepal. ⁵³Institut Jean Nicod, Ecole Normale Supérieure, Paris, France. ⁵⁴School of Collective Intelligence, Ben Guérir, Morocco. ⁵⁵Department of Leadership and Organizational Behavior, BI Norwegian Business School, Oslo, Norway. ⁵⁶Department of Psychology, University of Limerick, Limerick, Ireland. ⁵⁷Social Science Research Institute, Duke University, Durham, NC, USA. ⁵⁸Department of Clinical Neuroscience, Karolinska Institute, Solna, Sweden. ⁵⁹Department of Management, University of Agder, Kristiansand, Norway. ⁶⁰Department of Psychology, Lehigh University, Bethlehem, PA, USA. ⁶¹Department of Psychology, Monk Prayogshala, Mumbai, India. ⁶²Faculty of Medicine, Cooperative University of Colombia, Bogotá, Colombia. ⁶³Department of Bioethics, El Bosque University, Bogotá, Colombia. ⁶⁴Intellectus Memory and Cognition Center, Hospital Universitario San Ignacio, Bogotá, Colombia. ⁶⁵Pontifical Xavierian University, Bogotá, Colombia. ⁶⁶Intellectus Memory and Cognition Center, Hospital Universitario San Ignacio, Bogotá, Colombia. ⁶⁷Swiss Federal Institute of Technology (ETH Zurich), Zürich, Switzerland. ⁶⁸University of Bayreuth, Bayreuth, Germany. ⁶⁹National University of Political Studies and Public Administration, Bucharest, Romania. ⁷⁰Department of Psychiatry and Medical Psychology, Medical University of Plovdiv, Plovdiv, Bulgaria. ⁷¹Indian Institute of Management Indore, Indore, India. ⁷²School of Economics and Finance, Queensland University of Technology, Brisbane, Australia. ⁷³Center for Behavioral Economics, Society and Technology, Queensland University of Technology, Brisbane, Australia. ⁷⁴Department of Psychology, Royal Holloway, University of London, London, UK. ⁷⁵Warburg Institute, School of Advanced Study, University of London, London, UK. ⁷⁶Department of Behavioral and Cognitive Sciences, Faculty of Humanities, Education and Social Sciences, University of Luxembourg, Luxembourg. ⁷⁷Department of Political Science, National Taiwan University, Taipei, Taiwan. ⁷⁸Center for Research in Econometric Theory and Applications, National Taiwan University, Taipei, Taiwan. ⁸⁰University of Latvia, Riga, Latvia. ⁸¹Department of Psychology, Princeton University, Princeton, NJ, USA. ⁸²University of Canterbury, Christchurch, New Zealand. ⁸³Department of Psychology, University of Virginia, Charlottesville, VA, USA. ⁸⁴School of Economics and

Management, Hebei University of Technology, Tianjin, China. ⁸⁵Institute for Research and Development-Kurdistan. ⁸⁶University of Dakar, Dakar, Senegal. ⁸⁷School of Medicine and Health Sciences, Del Rosario University, Bogotá, Colombia. ⁸⁸Melbourne School of Psychological Sciences, University of Melbourne, Parkville, Australia. ⁸⁹Institute of Psychology, University of Wrocław, Wrocław, Poland. ⁹⁰IMT School for Advanced Studies Lucca, Lucca, Italy. ⁹¹University of Florence, Florence, Italy. ⁹²Department of Management, London School of Economics and Political Science, London, UK. ⁹³TBS Business School, Toulouse, France. ⁹⁴Social Policy Institute of the Ministry of Labor, Social Affairs and Family of the Slovak Republic. ⁹⁵Department of Psychology, Pennsylvania State University, Centre County, PA, USA. ⁹⁶Federal University of Santa Catarina, Florianopolis, Brazil. ⁹⁷University Institute of Advanced Studies, Pavia, Italy. ⁹⁸Neurological Institute Foundation Casimiro Mondino, Pavia, Italy. ⁹⁹Geriatric Society of Nepal, Kathmandu, Nepal. ¹⁰⁰School of Psychology, Shenzhen University, Shenzhen, China. ¹⁰¹Darden School of Business, University of Virginia, Charlottesville, VA, USA. ¹⁰²Institute for Advanced Study in Toulouse, Toulouse, France. ¹⁰³Department of Brain and Behavioral Sciences, University of Pavia, Pavia, Italy. ¹⁰⁴Cavalry Captain Witold Pilecki State University of Malopolska, Oświęcim, Poland. ¹⁰⁵Sauder School of Business, University of British Columbia, Vancouver, British Columbia. ¹⁰⁶Department of Psychology, University of California, Los Angeles, CA, USA. ¹⁰⁷Department of Psychology, University of Rennes, Rennes, France. ¹⁰⁸University of Clermont Auvergne, CNRS, LAPSCO, Clermont-Ferrand, France. ¹⁰⁹Cavite State University, Indang, Philippines. ¹¹⁰University of Alicante, Alicante, Spain. ¹¹¹Faculty of Medicine, University of Helsinki, Helsinki, Finland. ¹¹²Australian National University, Canberra, Australia. ¹¹³Ludwig Maximilian University of Munich, Munich, Germany. ¹¹⁴Department of Social Policy and Intervention, Oxford University, Oxford, UK. ¹¹⁵Creative Industries Faculty, Queensland University of Technology, Brisbane, Australia. ¹¹⁶Department of Psychology, University of Waterloo, Waterloo, Ontario, Canada. ¹¹⁷University of Bradford, Bradford, England, UK. ¹¹⁸University of North Carolina, Chapel Hill, NC, USA. ¹¹⁹University of Colorado Boulder, Boulder, CO, USA. ¹²⁰Department of Media and Communication, City University of Hong Kong, Hong Kong. ¹²¹University of Regensburg, Regensburg, Germany. ¹²²Munich University of Applied Sciences, Munich, Germany. ¹²³Free University of Berlin, Berlin, Germany. ¹²⁴Linköping University, Linköping, Sweden. ¹²⁵Department of Psychology, University of Alberta, Edmonton, Canada. ¹²⁶Questrom School of Business, Boston University, Boston, MA, USA. ¹²⁷University College Dublin, Dublin, Ireland. ¹²⁸University of Montana, Missoula, MT, USA. ¹²⁹Osaka University, Suita, Japan. ¹³⁰Seele Neuroscience, Mexico City, Mexico. ¹³¹University of Auckland, Auckland, New Zealand. ¹³²Maria Curie-Skłodowska University, Lubin, Poland. ¹³³Frank Batten School of Leadership & Public Policy, University of Virginia, Charlottesville, VA, USA. ¹³⁴Kobe University, Kobe, Japan. ¹³⁵Center for Social Sciences Institute for Legal Studies, Hungarian Academy of Sciences, Budapest, Hungary. ¹³⁶University of Silesia, Katowice, Poland. ¹³⁷University of Belgrade, Belgrade, Serbia. ¹³⁸AGH University of Science and Technology, Kraków, Poland. ¹³⁹Department of Psychology, University of Cambridge, Cambridge, England, UK. ¹⁴⁰Mackenzie Presbyterian University, São Paulo, Brazil. ¹⁴¹University of Turku, Turku, Finland. ¹⁴²University of Newcastle, Callaghan, Australia. ¹⁴³University of St. Andrews, St. Andrews, Scotland. ¹⁴⁴Manmohan Memorial Medical College and Teaching Hospital, Kathmandu, Nepal. ¹⁴⁵Deakin University, Geelong, Australia. ¹⁴⁶Department of Psychology, University of Western Ontario, London, Ontario,

5

Canada. ¹⁴⁷Carleton University, Carleton University, Northfield, MN, USA. ¹⁴⁸Institute for Emerging Market Studies of the Hong Kong University of Science and Technology, Hong Kong. ¹⁴⁹Dokuz Eylül University, İzmir, Turkey. ¹⁵⁰Department of Psychology, University of Koblenz-Landau, Landau, Germany. ¹⁵¹Stockholm School of Economics, Stockholm, Sweden. ¹⁵²University of Michigan, Ann Arbor, MI, USA. ¹⁵³Kyushu University, Fukuoka, Japan. ¹⁵⁴Kadir Has University, Istanbul, Turkey.

Contact: Jay Van Bavel jay.vanbavel@nyu.edu

Abstract: The ongoing COVID-19 pandemic is a devastating global health crisis. Without a vaccine or effective medication, the best hope for mitigating virus transmission is collective behavior change and support for public health interventions (e.g., spatial distancing, physical hygiene, and endorsement of health policies). In a large-scale international collaboration (N = 46,769 across 67 countries), we investigated why people adopted public health behaviors and endorsed public policy interventions (e.g., closing bars and restaurants) during the pandemic (April-May, 2020). Results revealed that respondents who identified more strongly with their nation consistently reported engagement in public health behaviors and greater support for public health policies. We also found a small effect of political orientation, suggesting that left-wing individuals were more likely to endorse public health behaviors and support for public health measures than right-wing individuals. We discuss the implications of links between national identity, leadership, and public health for managing the COVID-19 and future pandemics.

Keywords: COVID-19, national identity, public health, pandemic, human development

Introduction

As of December 2020, more than 64 million people worldwide have been infected by the new coronavirus and more than 1.4 million have died,¹ making the COVID-19 pandemic one of the greatest health crises of the past century. Until a vaccine or effective medical treatment is administered, the public response to the current pandemic is largely limited to policy-making and collective behavior change (Lewnard & Lo, 2020). To reduce virus transmission it is crucial that people engage in public health behavior (e.g., maintain spatial distance and improve physical hygiene) and support COVID-19 protective policies (e.g., limiting travel and closing bars and restaurants). And even after effective vaccines are administered, it will be critical to convince people to take them. This is why the Director of the World Health Organization recently argued: *“That’s why behavioural science is so important – it helps us to understand how people make decisions, so we can support them to make the best decisions for their health”* (WHO, 2020). In the current investigation, we respond to this call for behavioral science and present the results from a global study across 67 countries, testing key predictors of public health support.

What drives people to support public health initiatives during a pandemic? Answering this question may help countries, health organizations, and political leaders develop effective interventions to increase compliance with actions such as maintaining spatial distance and restricting travel (Van Bavel et al., 2020). Governments can coordinate individuals and help them avoid behaviors no longer considered socially responsible. However, recent evidence suggests this type of leadership requires cultivating a shared sense of solidarity to increase compliance with recommended health behaviors (Biddlestone, Greene, & Douglas, 2020; Haslam et al., 2011; Martinez-Brawley & Gualda, 2020). Solidarity with other members of one’s group is a component of ingroup identification (Leach et al., 2008), that is the personal significance that being part of a nation holds for an individual (Cameron, 2004; Leach et al., 2008; Postmes et al., 2012; Tajfel, 1978). Identifying with a group is associated with mutual cooperation and adherence to norms (e.g., Brewer, 1999; Buchan et al., 2011; De Cremer & Van Vugt, 1999), motivation to help other members of their group (Ellemers, Spears, & Doosje, 1999; Levine, Prosser, Evans, & Reicher, 2005), and willingness to engage in collectively-oriented actions aimed at improving the group’s welfare (Bilewicz & Wójcik, 2010; Klandermans, 2002; Tajfel, 1978; van Zomeren et al., 2008). We reasoned that national identities, in particular, might play a key role in promoting public health in the current pandemic (see Van Bavel et al., 2020).

¹ These figures were reported by the World Health Organization: www.who.int/emergencies/diseases/novel-coronavirus-2019

National identity plays an important role in motivating people to engage in greater civic involvement (Huddy & Khatib, 2007) and costly behaviors that benefit other members of their national community (Kalin & Sambanis, 2018). Accordingly, a strong sense of shared national identity might help promote collective efforts to combat the pandemic within a country (e.g., Dovidio et al., 2020). Moreover, border closures, travel bans, and national task forces have likely made national identities even more salient during the pandemic (Bieber, 2020). The existence and activation of strong collective identities can allow political leaders to mobilize large populations to adhere to emergency public health measures. For instance, political leaders and public health officials often foster a sense that “we are in this together” and that we can manage the crisis through collective action (Gkinopoulos & Hegarty, 2018; van Zomeren et al., 2008). This might be particularly important since partisanship and polarization within countries can reduce health behavior and increase the risk for infections and mortality (Gadarian, Goodman, Pepinsky, 2020; Gollwitzer, Martel, Brady, Pärnamets, Freedman, Knowles, & Van Bavel, 2020; Van Bavel et al., 2020). However, recent work in the United States has found that threats to national identity can lead to less support for public health initiatives (Kachanoff et al., 2020).

The goal of the current paper is to examine whether national identification explains global adherence to the collective health measures during a pandemic (Haslam, 2020; Haslam & Reicher, 2006; Jetten et al., 2020). We examined the associations between the strength of identification with one’s nation and whether people adopted public health behaviors (e.g., spatial distancing, hand washing) and endorsed public policy interventions (e.g., closing bars and restaurants). Prior work suggests these actions could substantially reduce the number of COVID-19 infections (Block et al., 2020; Ferguson, 2006; Koo et al., 2020; Lewnard & Lo, 2020). Our primary hypothesis is that people who identify strongly with their nation will, on average, express stronger support for public health measures than those who identify less with their nation.

National identity is distinct from beliefs about national superiority (e.g., Huddy & Del Ponte, 2019; Kosterman & Feshbach 1989; Roccas, Klar, & Liviatan, 2006). The latter has been conceptualized as national narcissism: a belief in the greatness of one’s nation that is unappreciated by others (Golec de Zavala, Cichocka, Eidelson, & Jayawickreme, 2009). National identity tends to correlate positively with national narcissism because they both assume a positive evaluation of one’s nation. However, they predict different outcomes. For example, outgroup prejudice is *negatively* associated with national identification but *positively* with national narcissism (Golec de Zavala, Cichocka, & Bilewicz, 2013). Research also finds that collective narcissism is associated with frustrated individual needs, such that people high in collective narcissism are especially concerned with how their group reflects on them (Cichocka, 2016). National narcissism then predicts greater preoccupation with maintaining a positive image of the nation than with the well-being of fellow citizens (Cislak et al., 2018; Marchlewska et al., 2020). Thus, national narcissists may be less inclined to engage in behaviors to prevent the spread of COVID-19—or even acknowledge the risks associated with the pandemic in their home country (Nowak et al., 2020). Therefore, in predicting compliance with public health measures, we sought to distinguish national identification from national narcissism.

In addition, there is some evidence right-wing political ideology is associated with both national identity (e.g., van der Toorn, Nail, Liviatan & Jost, 2014) and national narcissism (e.g., Cichocka, Marchlewska, Golec de Zavala, & Olechowski, 2016; Cislak et al., 2018; Marchlewska, Cichocka, Panayiotou, Castellanos, & Batayneh, 2018). Moreover, supporters of right-wing political parties have tended to downplay risks associated with COVID-19 (e.g., Calvillo et al., 2020; Capraro & Barcelo, 2020; Sjästad & Van Bavel, 2020) and were less likely to comply with preventative measures compared to left-leaning or liberal individuals (van Holm et al., 2020; Capraro & Barcelo, 2020). For that reason, we differentiated the effects of national identification and narcissism from political preferences.

Overview

The COVID-19 pandemic is a truly *global* crisis with over 200 countries reporting infections. To understand the variables that account for public health support around the globe, we launched a collaborative, international project in April 2020 collecting large-scale data from as many nations as possible. We collected a large sample consisting of citizens from 67 countries ($n = 46,769$; see **Figure 1**). See **Methods** for details about the sample.

We analysed the data using multi-level models, which allows us to account for country-level variation in economic, political and health measures while examining whether national identification predicted public health support. We modeled country-level factors, such as the Human Development Index (HDI)² to examine whether the impact of person-level factors, like national identification, would predict public health support, including after adjusting for the general health and standard of living in a country (see **Methods** for details and sample items). We also adjusted for the COVID-19 infection and mortality rates within each country to ensure that public health support was not merely a function of local risks. Due to the large sample size, we focused our interpretations on findings that were highly statistically significant ($p < .001$).

Our results provide robust evidence that national identification is a reliable predictor of all three forms of public health support measured in our survey: self-reported spatial distancing, physical hygiene, and policy support. Citizens who identified more strongly with their nation reported greater support for critical public health measures, even after adjusting for national narcissism and political ideology (as well as the country-level Human Development Index and local rates of COVID-19 infections and mortality). Right-wing participants reported lower levels of support for all three protective measures than left-wing participants. National narcissism was only weakly, albeit positively, related to support for two measures: physical hygiene and policy support. In sum, the results provide clear evidence that national identification is related to public health support across national and cultural contexts.

² The Human Development Index (HDI) is a measure of achievement in key dimensions of human development: a long and healthy life, being knowledgeable, and having a decent standard of living. The HDI is the mean of normalized indices for each of the three dimensions (see <http://hdr.undp.org/en/content/human-development-index-hdi>).

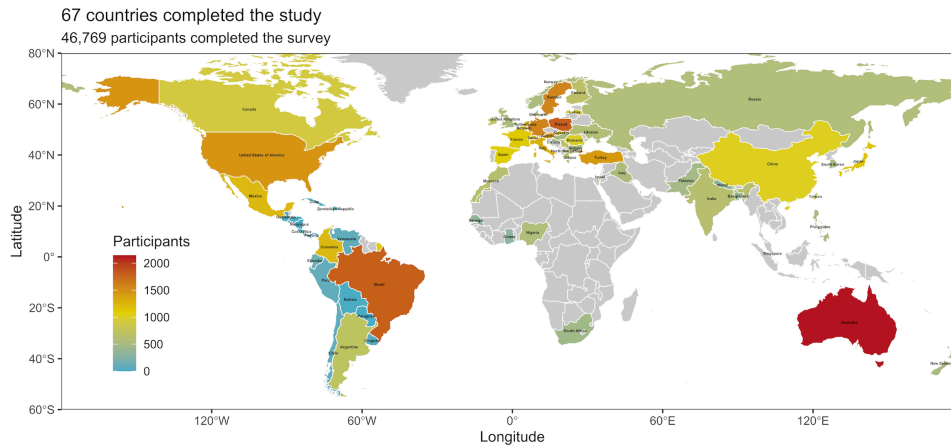


Figure 1. Map of the 67 participating countries and territories with sample size scaled to color (we did not obtain samples from countries in grey). The total sample included 46,769 participants.

Results

Participants generally reported following the guidelines for contact and hygiene and they supported policies that were intended to reduce the impact of COVID-19 (i.e., means for all three measures were above 8, on scales ranging from 0 to 10; see **Table 1**). The public health measures were all highly correlated with one another (Pearson’s $r_s > .38$). Consistent with prior work, national identification was positively correlated with both national narcissism ($r = .38$) and right-wing political ideology ($r = .18$).

Table 1. Summary statistics and multi-level correlations for person-level measures. The mean score for each scale is presented along with the variance explained within and between participants and the scale reliability (alpha). This is no alpha for ideology since it is a one-item measure. Higher scores reflect greater support for each measure (and stronger right-wing political beliefs in the case of ideology).

	Mean	Variance		Alpha	Correlations				
		Between	Within		2	3	4	5	6
1.Spatial distancing	8.61	.21	2.13	.74	.43	.44	.02	.15	-.02
2. Physical hygiene	8.21	.47	2.31	.72		.38	.12	.17	.02

3. Policy support	8.29	.94	3.45	.81			.06	.13	-.03
4. National narcissism	5.36	2.18	4.89	.82				.38	.27
5. National identification	8.04	.83	4.00	.72					.18
6. Political ideology	4.99	.37	5.03	NA					

For all measures, except political ideology, there were negative relationships between HDI scores and country-level means (see **Table 2**). In other words, citizens in countries with higher scores on the global Human Development Index also reported less support for two of our COVID-19 public health measures. However, we should note that our dataset includes data from very few African countries, many of which have relatively low HDI scores but seem thus far to have fared better in the pandemic than higher-HDI countries (see Ghosh, Bernstein, & Mersha, 2020).³

Table 2: Relationships between HDI (Human Development Index) scores and means of person-level variables. *T*-ratios and variance provide the test statistic and the percentage of variance explained for each variable.

Measure	HDI	<i>t</i>-ratio	Variance
<i>Spatial distancing</i>	-.13	2.59	8%
<i>Physical hygiene</i>	-.40*	4.98	34%
<i>Policy support</i>	-.59*	5.81	37%
<i>National narcissism</i>	-.94*	6.39	11%

³ Although it is beyond the scope of the current paper, future research should examine whether people in less developed countries indeed placed a greater faith in their political and medical institutions during the early stages of the pandemic, or whether they were motivated to simply report higher compliance with governmental guidelines.

11

<i>National identity</i>	<i>-.52*</i>	<i>7.33</i>	<i>31%</i>
<i>Political ideology</i>	<i>-.12</i>	<i>1.44</i>	<i>2%</i>

Note: * $p < .001$

As seen in **Table 3**, national identification was the strongest predictor of all three COVID-19 public health support measures (see **Figure 2** and **Figure 3** for the coefficients in each country). It was significantly and positively related to all three measures (even after adjusting for national narcissism and political ideology). Individuals with stronger national identification (relative to other people within their own nation) reported stronger support for increasing spatial distance and improving physical hygiene than individuals with weaker identification, and they also endorsed COVID-19 public health policies to a greater extent.

Table 3. Relations between outcomes and predictors (including the slope and *t*-ratio of each relationship). National identification was the strongest predictor of all three COVID-19 public health support measures.

Outcome	Predictor	Slope	<i>t</i>-ratio
Spatial distancing	National narcissism	-.006	< 1
	National identification	.128*	8.49
	Political ideology	-.027*	4.31
Physical hygiene	National narcissism	.060*	6.40
	National identification	.124*	11.24
	Political ideology	-.015	1.89
Policy support	National narcissism	.029*	2.82

	National identification	.127*	10.30
	Political ideology	-.050*	4.72

Note: * $p < .001$

National narcissism was significantly positively related to two of the three protective measures (albeit weakly). Individuals scoring higher in national narcissism supported physical hygiene and COVID-19 policies more strongly compared to individuals with lower levels of national narcissism. The relations between political ideology and public health support were negative (albeit weakly) for two of the outcome measures, indicating that individuals with more left-leaning or liberal political orientation tended to support spatial distancing and COVID-19 policies more strongly than those with more right-leaning or conservative political orientation. Taken together, the three predictors accounted for 8% of the person-level variance of the contact measure, for 7% of the person-level variance of the hygiene measure, and 5% of the person-level variance of the policy support measure.

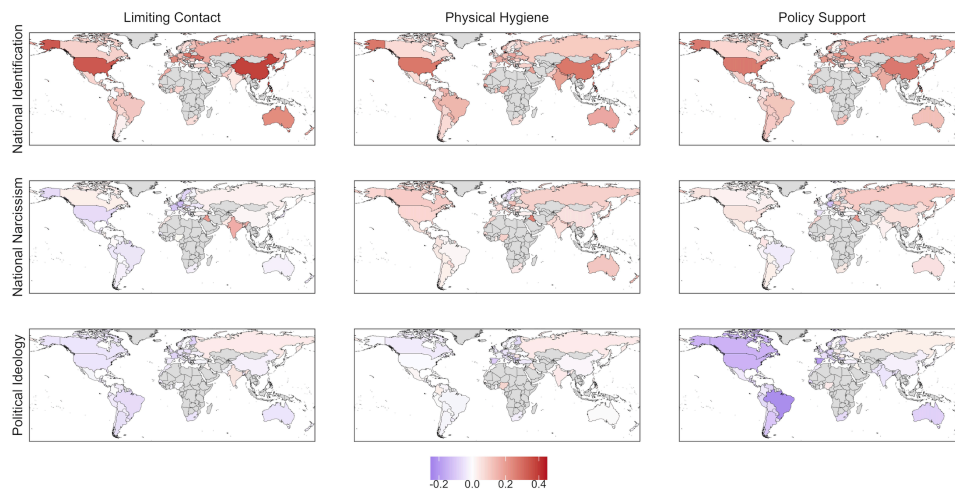


Figure 2. Relation between collective concerns and public health measures in 67 countries and territories. Heat index depicts the slope coefficients in each country. Blueish colors mean negative associations between our predictors and our outcomes while reddish colors mean positive associations (higher scores reflect stronger relationships between national identification, greater national narcissism and greater conservatism, and limiting contact, physical hygiene, and policy support measures).




Figure 3. Relation between collective concerns and public health measures in 67 countries and territories. The coefficients reflecting the relationship between national identity and each of the health measures are presented for each country from strongest (top) to weakest (bottom). The relationship with physical contact (red), policy support (green), and hygiene (blue) are color coded.

Finally, we explored whether the relations between predictors and outcomes varied across countries as a function of HDI. The relations between national identification and each of the three public health measures were not moderated by HDI. On the other hand, we found that HDI moderated the relation between national narcissism and spatial distancing, such that this relation was negative in countries that had higher HDI scores but positive in countries that had lower HDI scores. We also found that HDI moderated the relation between right-wing political ideology and hygiene, such that this negative relation was stronger in countries that had higher HDI scores than in countries that had lower HDI scores. See **Methods and Analysis** for details.

Discussion

Our research suggests that national identities might play an important role in the fight against a global pandemic. Following World War II, early work in social psychology had a tendency to focus on the negative side of nationalism and leadership persuasion, such as destructive obedience to authority (Milgram, 1963) and group conformity to incorrect beliefs held by others (Asch, 1956). In the decades since then, research on social identity (Tajfel, 1978) and a “social cure” approach to mental health (Jetten, Haslam & Haslam, 2011) has revealed that there is also a pro-social side to group identity. Based on this latter perspective we predicted, and found, that national identification was *positively* associated with public health support around the globe.

Individual differences in the strength of identification with one’s nation robustly predicted public health support, operationalized as behavioral health intentions (physical distance and physical hygiene) and support for COVID-19 policy interventions. In short, people who identified more strongly with their nation reported greater engagement with critical public health measures around the globe.⁴ These results are consistent with the social psychological literature on the benefits of identifying with one’s social groups. It also underscores a potential benefit of national identification which might be particularly salient during a national or global health crisis (Beider, 2020). Our research provides evidence that this form of identification might help to understand public health behaviors and deliver effective campaigns to promote those behaviors during a global pandemic that confined citizens to their homes and countries. Future work should examine the potential benefits of identity appeals during a pandemic and aim to isolate the causal influence of national identity on real behavior.

⁴ It is unclear how immigrants and refugees, who may not be living in their country of citizenship, respond to issues of identity during a public health crisis. We flag this as an important area for future research.

There is reason to believe that other forms of group identification can undercut public health. For instance, recent research has found that partisanship within countries (i.e., when people strongly identify with a specific political party) is associated with risky behavior (Gadarian et al., 2020), including spatial distancing (Gollwitzer et al., 2020; see also Alcott et al., 2020). One paper that used geo-tracking data from 15 million smartphones in the US found that counties that voted for a Republican (Donald Trump) over a Democrat (Hillary Clinton) exhibited 14% less spatial distancing during the early stages of the pandemic. These partisan gaps in distancing predicted subsequent increases in infections and mortality in counties that voted for Donald Trump. Moreover, partisanship was a stronger predictor than many other economic or social factors (e.g., county-level income, population density, religion, age, and state policy). This may be due to leadership, social norms, and media consumed by people from different identity groups. As such, stronger group identification is not always associated with healthy behavior.

It is tempting to conclude that political ideology might account for these relationships. However, we found that right-wing political ideology had a positive, moderate correlation with both national identity and narcissism, but very weak correlations with support with public health measures in our multi-country sample. Specifically, right-wing political beliefs were associated with less support for COVID-19 protection, compared to left-wing political beliefs. This relationship between political beliefs and compliance has been observed in several countries (e.g., Capraro & Barcelo, 2020; Gollwitzer et al., 2020; Ponce, 2020). Similarly, while national identification and national narcissism were associated positively with support for public health measures, right-wing political ideology was negatively associated. This suggests that a collective identity might lead people to value the protection of the entire group during a pandemic, even after adjusting for their ideological differences. This can be seen as the positive side of national identity and how it might be leveraged in a global crisis to promote activities that benefit a society.

It is also important to note that the relationship between national identity and public health support was distinct from national narcissism. In past research, national narcissism has predominantly been linked to problematic attitudes towards both out-group and in-group members (Cichocka, 2016; Cichocka & Cislak, 2020; Marchlewska et al., 2020). However, we found that national narcissism was positively associated with self-reported physical hygiene and support for COVID-19 preventative policies. Still, these effects were much smaller than those for national identity and depended on the context (i.e., national narcissism was associated with a lower likelihood of reporting spatial distancing in relatively more economically developed countries). Future work should thus carefully consider cross-national differences in human development as well as local norms associated with national identity. It is also noteworthy that there was an unexpected *negative* link between HDI and self-reported health behaviors. It is unclear why this link exists but future work might examine if people in less developed countries place greater faith in their political and medical institutions during the early stages of a pandemic or whether they are motivated to report higher support for public health guidelines.

Implications

Our evidence suggests that national identity may have modest predictive value for people's endorsement of and adherence to public health measures in the context of COVID-19. This information may be leveraged to create a sense of inclusive nation-based in-groups, potentially increasing engagement with recommended policies. Political and public health leaders might be able to similarly adapt effective communication strategies to appeal to a sense of national identification. Indeed, this might be particularly helpful in highly polarized countries where adherence to public health recommendations has become a partisan issue (see Gollwitzer et al., 2020). For instance, Canadian leaders across the political spectrum adopted similar messaging about the serious risks of the current pandemic which resulted in a rare moment of cross-partisan consensus among the public (Merkley, Bridgman, Loewen, Owen, Ruths, & Zhilin, 2020). Such recategorizations to overarching inclusive national groups (e.g., Gaertner et al., 2016) may be effective for preventing unhealthy behaviors. As such, leaders who wish to inspire public health behavior might benefit from connecting the issue to feelings of national identity. Framing these messages at the level of the nation rather than, for instance, a partisan group, region, or municipality also makes sense in the context of a pandemic, which requires national (and international) coordination in the public health sector (Dovidio et al., 2020; Harari, 2020).

However, we note the effective application of these appeals requires future research as national identity is also implicated in intergroup conflict. This is more likely in the case of national narcissism (Cichocka & Cislak, 2020; Golec de Zavala et al., 2009), which tends to be associated with lower solidarity with other nations in the crisis (e.g., Gorska et al., 2019). In the absence of such narcissism, national identity could reflect not only concerns about protecting one's own country, but also into concern for other nations. Indeed, prior research has found that national identification is associated with more positive attitudes towards other nations--especially when adjusting for national narcissism (Cichocka et al., 2016; Golec de Zavala, Cichocka, & Bilewicz, 2013). Thus, the nature of national identity might be an important determinant of the effectiveness of identity and the potential for international cooperation. In addition, it could turn out that a commitment to cosmopolitanism or other supranational identities and ideologies may play a role that bolsters what we have seen in the case of national identity (Liu, Zhang, Leung, Zuniga, Gastardo-Conaco, Vasiutynskyi, & Kus-Harbord, 2020).

Limits on generalizability

The vast majority of published research in psychology and social sciences has been conducted in so-called WEIRD cultures (Henrich, Heine, & Norenzayan, 2010), typically restricted to the narrow western and educational setting of American or European university students, and non-representative participants from industrialized, rich and democratic countries. The COVID-19 pandemic, however, is a truly *global* issue, currently affecting over 200 countries and territories all over the world. We aimed to collect representative samples to help make generalizations to the wider population within each country as well as the broader sample of countries around the globe. Although we managed to collect data from 67 countries and territories, we were nevertheless unable to ensure representative samples from many countries or even conduct our survey in other countries (especially in many African countries as well as the Middle East). Therefore, our inferences apply to nations where we managed to complete this research and the specific samples we obtained.

The current research was correlational and conducted during the early phase of the pandemic. While a causal relationship between national identification and public health behavior makes sense from a theoretical perspective, we cannot rule out the possibility that public health behavior actually causes national identification, or that both are caused by a third variable (e.g., Bieber, 2020). Moreover, we have no evidence whether this pattern would apply during later stages of the current or future pandemics. Indeed, it is possible that national identity increases during times of crisis as people recognize their duty as citizens to help respond to this issue. We encourage future work to examine ways to experimentally manipulate the salience of national identification or frame health messages in a way that highlights the link between identification and the public health measures. A related issue is the critical need to measure actual behavior rather than behavioral intentions or policy support (see Isler, Isler, Kopsacheilis, & Ferguson, 2020) as well as consider access to the necessary resources to engage in health behavior.

Another limitation, hinted at above, is the exclusive focus on national groups rather than, for instance, participants' tendency to identify with their city, region, religion, or ethnic group — or, for that matter, all of humanity. Some research suggests that local leaders might be ineffective if their advice contradicts a national leader (see Gollwitzer et al., 2020). In the current pandemic, nation-states have been among the most important actors for implementing policy or promoting national health guidelines, but sub-national units and international organizations such as the World Health Organization (WHO) have also played an important role. Our data do not examine whether sub- or supra-national identities play a similar role in driving public health attitudes and behaviors but we encourage future work on this topic.

Conclusion

The COVID-19 pandemic spreading across the world is one of the most devastating global health crises of the century. Until a verifiably safe and effective vaccine or therapeutic treatment is widely available, efforts to inspire collective action for greater compliance with public health measures become the central challenge when mitigating the transmission of the SARS-CoV-2 virus (e.g., spatial distancing, physical hygiene and support for health policies). Our large-scale international collaboration across 67 countries suggests that identification with one's nation is positively associated with support for critical public health measures. Understanding the role of social identity appears to be an important issue when addressing public health crises.

Methods and Analysis

In April 2020, we launched a call using social media to collect data all over the world on psychological factors that might be related to COVID-19 pandemic response, with public health support as the primary outcome. Each team was asked to collect data from at least 500 participants, representative with respect to gender and age, in their own country or territory. We created a survey in English (see below) that we sent to each team. The survey was approved by the ethics board at the University of Kent.⁵ Where necessary, each team translated the survey into the local language, using the standard forward-backward translation method, and then collected the data. The datasets were then collated and analyzed using multi-level models. We report how we determined our sample size, all data exclusions (if any), all manipulations, and all measures in the study (see **supplement**).

Participants

We collected data from a total of 67 countries ($n = 46,769$ participants; Mean age = 43.09; Gender = 51.8% females)⁶. **Figure 1** shows the geographical distribution of countries included in the project (For a full list and sample characteristics from each country, please see **supplement**). The sample includes countries from all continents (except for Antarctica).⁷

We encouraged teams to collect nationally representative samples. Of the 67 countries in which data were collected, representative samples were achieved in 30 and convenience samples were collected in 34, and both types of sampling were used in three countries. All the analyses reported in our paper were repeated controlling for differences in sampling method. Although some coefficients differed as a function of sampling procedure, none of these differences compromised or altered the main effects we report.⁸

Survey

Questionnaires were administered online. Each participant completed a series of psychological measures and self-reported public health behaviors (see complete survey with all items in **supplement**). Participants completed the scales in random order.

⁵ Each team was allowed to include additional items after the main survey under their own institutional protocol.

⁶ 50,944 participants answered the survey. We first excluded those participants that did not answer more than 75% of the survey. This first cut resulted in a sample of 48,895 participants. Then, we cut those younger than 18 y/o or older than 100 y/o which resulted in a sample of 48,764. Finally, we cut those that failed on an attentional check task; with this, we achieved the final sample of 46,769. 20 countries were able to collect representative samples with respect to sex and age. 44 countries were able to collect more than 500 subjects.

⁷ Some are overrepresented (e.g., Europe, Americas) while others are underrepresented (e.g., Africa, Middle East).

⁸ For example, type of sampling moderated the slope between physical hygiene and national identification. The mean slope was .124, and the estimated slope for countries that collected representative samples was .146, whereas it was .102 for countries that collected convenience samples. In other words, the main finding appeared to be slightly stronger in the more representative samples.

For the current paper, we focused on three potential predictors of public health support. Our primary predictor was a two-item *national identification* measure (which included one item from Postmes, Haslam, & Jans, 2012 and an additional item measuring identity centrality from Cameron, 2004): “I identify as (nationality)” and “Being a (nationality) is an important reflection of who I am”. One secondary predictor was a three-item *national narcissism* scale (Golec de Zavala, Cichocka, Eidelson, & Jayawickreme, 2009), which include the following sample item: “My (national group) deserves special treatment.” The nationalities were provided by the survey researchers. These measures used an 11-point slider scale with three labels items: 0 = “strongly disagree”, 5 = “neither agree nor disagree”, 10 = “strongly agree”.

As a third predictor, we included a one-item measure of *political ideology*: “Overall, how would you best describe yourself in terms of political ideology?”. This measure used a scale from 0 = *extremely liberal/left-leaning* to 10 = *extremely conservative/right-leaning*.⁹ This single-item measure of ideology has been found to account for a significant proportion of the variance in presidential voting intentions in American National Election studies between 1972 and 2004 (Jost, 2006). We included the terms left-leaning and right-leaning to make the item generalizable to numerous political systems.

As the primary outcome variable, we included three measures of public health support. A *spatial distancing* scale, consisting of five items, as, for example, “During the days of the coronavirus (COVID-19) pandemic, I have been staying at home as much as practically possible”.¹⁰ A *physical hygiene* scale, consisting of five items, as, for example, “During the days of the coronavirus (COVID-19) pandemic, I have been washing my hands longer than usual”. A *policy support* scale, consisting of five items, as, for example, “During the days of the coronavirus (COVID-19) pandemic, I have been in favor of closing all schools and universities”. We used an 11-point “slider scale with three labels: 0 = “strongly disagree”, 50 = “neither agree nor disagree”, 100 = “strongly agree”, which was re-coded to a scale from 0 to 10.

To see if these relationships varied as a function of socio-economic factors and the state of the pandemic in each country, we examined several country-level factors. Specifically, we included the 2018 (most recent available) Human Development Index (ranging from 0 to 1), which represents a combined index of life expectancy at birth, level of education (mean years of schooling for adults over 25 and expected years of schooling for children), and national wealth (gross national income per capita; Human Development Report, 2019). To ensure our results were not confounded with the pandemic rate across countries, we measured the total COVID-19 infection and mortality cases (as well as the infection and mortality rate per capita) in each country at the start of data collection for this project.¹¹ These variables had no impact on the results and are not discussed further. All measures will be made publicly available upon publication at the *Open Science Framework* website.

⁹ Note that the meanings of 'left' and 'right' vary by cultural and political context (see Aspelund, Lindeman, & Verkasalo, 2013).

¹⁰ Prior to conducting our analyses, we learned that the five-item scale had low reliability ($\alpha = .002$). However, after dropping one bad item the scale had acceptable reliability ($\alpha = .72$). As such, all analyses reported in the paper use this four-item version of the scale.

¹¹ Our main findings were robust even after adjusting for total infections and deaths as well as infections and deaths per capita at the start of data collection for this project (April 17, 2020; Dong, Du, & Gardner, 2020).

Analysis

We conceptualized the data as a multi-level data structure in which persons were nested within countries, and we analyzed the data with a series of multi-level models (MLM) using the program HLM (Raudenbush et al., 2011; see Nezlek, 2010, for a description of using MLM to analyze data from multinational studies). The analyses examined within-country (person-level) relationships between behavioral health protective responses to COVID-19 (i.e., spatial distancing, physical hygiene, and policy support) and individual differences in collective concerns (i.e., national identification, national narcissism, and political ideology). We also examined the moderating effects of country-level differences on these person-level relations. For instance, we examined if these person-level relations between collective concerns and health protective measures varied as a function of between-country differences in overall human development as measured by the Human Development Index (HDI) or national rates of COVID-19 infections and mortality.

Descriptive statistics and reliability of person-level measures

Before examining relations between COVID-19 protection and socio-political attitudes, we examined the reliability of our measures (with the exception of political ideology, which was measured with only one item). These analyses consisted of models in which the i items in a scale were nested within j persons, which were nested within k countries. Such analyses provide the multi-level equivalent of a Cronbach's alpha (Nezlek, 2017; Raudenbush & Bryk, 2002). The model is below.

$$\begin{array}{ll} \text{Level 1 (item-level):} & y_{ijk} = \pi_{0jk} + e_{ijk} \\ \text{Level 2 (person-level):} & \pi_{0jk} = b_{00k} + r_{0jk} \\ \text{Level 3 (country-level):} & b_{00k} = g_{000} + u_{00k} \end{array}$$

In the level 1 model, y_{ijk} is response i , for person j , in country k , π_{0jk} is a random coefficient representing the mean response for person j in country k , b_{0j} is a random coefficient representing the mean of y for country k (across the j persons in each country), e_{ijk} represents the error associated with each measure, and the variance of r_{ijk} constitutes the within-country variance. In multi-level modeling, the coefficients from one level of analysis are passed up to the next. In the level 3 model, g_{000} represents the grand mean of the country level means (b_{00ks}) from the person-level model, u_{00k} represents the error of b_{00k} , and the variance of u_{00k} constitutes the level 3, country level variance.

These analyses suggested that, with the exception of spatial distancing, our scales were at least "moderately" reliable ($\alpha > .60$; Shrout, 1998). The reliability estimates and descriptive statistics are presented in **Table 1**. For spatial distancing, follow-up analyses indicated that a reliable scale could be created from items 1, 3, 4, and 5. Item 2 asking about visiting friends, family or colleagues was therefore dropped from the final analyses.

The estimated means suggest that people generally reported following the guidelines for contact and hygiene and they supported policies that were intended to reduce the impact of COVID-19 (i.e., means for all three measures were above 8, on scales ranging from 0 to 10). Moreover, although the majority of the variance in national identification, national narcissism, and political ideology was within-country, there was also notable between-country variance. This justified further analyses of relations between country-level means of these measures and HDI. We calculated scale means and used Mplus (Muthén & Muthén, 2017) to estimate multi-level correlations for person-level measures, controlling for the nested structure of the data (see **Table 1**).

Country-level relations between HDI and the means of person-level measures

The next set of analyses examined relations between scores on the HDI and the means of the person-level measures. This model was a variant of the unconditional model. HDI scores were entered as a predictor in the country level model presented above (level 3). MLM analyses do not estimate standardized coefficients, and to simplify the interpretation of the results, HDI scores were standardized prior to analysis (and, therefore, were entered uncentered).

Note that these analyses took into account the reliability of scales. By nesting items within persons, we estimated a latent mean for each construct. Although the results of analyses of two-level models that did not estimate latent means but used observed means (persons within countries) were indistinguishable from the results we present here, we present the results of the three models because they are more accurate than the results of the two-level analyses.

The results of these analyses are presented in **Table 2**. For all measures, except political ideology, there were negative relationships between HDI scores and country-level means. Note that the coefficients in the table represent the change in a country-level mean associated with a 1 SD increase in HDI scores. In other words, citizens in countries with higher scores on the global Human Development Index also reported less support for COVID-19 public health measures. Effect sizes are defined as the percent reduction in the country-level variance of a null model (**Table 2**) associated with the inclusion of HDI scores at the country-level. Because political ideology was measured with only one item, the variance estimates and effect size for political ideology are from a two-level model (persons nested within countries). Estimating effect sizes for multilevel analyses such as those used in the present study are discussed in Nezlek (2010).

Person-level relations between collective concerns and public health support.

Next, we examined person-level relations between the three COVID-19 protection measures (modeled as outcomes) and national identification (NI), national narcissism (NN), and political ideology (PI) (modeled as predictors). Predictors were defined as the mean scores for each scale. To account for relations among the predictors, all predictors were entered at the person level of the model. Predictors were entered group-mean centered and were modeled as randomly varying. Again, because this was done using a three-level model, outcomes were measured in terms of latent means.

Entering predictors group-mean centered meant that estimates of coefficients controlled for country-level differences in means (Nezlek, 2010). Entering predictors as randomly varying meant that the model took into account the possibility that slopes varied between countries. In essence, a regression equation, consisting of an intercept and a slope, was estimated for each country, and these estimates were “passed up” to the country level where they were tested for significance. The model is below (item level is not shown).

$$\begin{aligned} \text{Level 2 (person-level):} & \quad \pi_{0jk} = b_{00k} + b_{01k}*(NN) + b_{02k}*(NI) + b_{03k}*(PI) + r_{0jk} \\ \text{Level 3 (intercept):} & \quad b_{00k} = g_{000} + u_{00k} \\ \text{Level 3 (NN slope):} & \quad b_{00k} = g_{010} + u_{01k} \\ \text{Level 3 (NI slope):} & \quad b_{00k} = g_{020} + u_{02k} \\ \text{Level 3 (PI slope):} & \quad b_{00k} = g_{030} + u_{03k} \end{aligned}$$

The hypothesis of interest was tested by assessing the significance of the g_{010} , g_{020} , and g_{030} coefficients in this model. Was the mean slope between an outcome and a predictor significantly different from 0? The results of these analyses are summarized in **Table 3**. These unstandardized coefficients represent the expected change in an outcome for a one-unit increase in a predictor, i.e., an increase of one on a scale (out of 11). Also, the random error terms for all predictors were significant at $p < .001$.

As can be seen from the coefficients presented in **Table 3**, national identification was the most reliable and strongest predictor of our COVID-19 public health support measures (see **Figure 2** for the coefficients in each country). It was significantly and positively related to all three measures (even after adjusting for national narcissism and political ideology). Individuals with stronger national identification (relative to other people within their own nation) reported stronger support for limiting physical distance and improving physical hygiene than individuals with weaker identification, and they also endorsed COVID-19 public health policies to a greater extent.

National narcissism was significantly positively related to two of the three protective measures (albeit weakly). Individuals scoring higher in national narcissism supported recommendations for physical hygiene and endorsed COVID-19 related policies more strongly compared to individuals with lower levels of national narcissism.

The relations between political ideology and public health support were negative (albeit weakly) for all three outcome measures, indicating that individuals with more left-leaning or liberal political orientation tended to adhere to recommendations for contact, hygiene and supported COVID-19-related policies more strongly than those with more right-leaning or conservative political orientation.

Effect sizes were estimated using a similar procedure that was used for estimating effect sizes at the country-level. Effect sizes were defined as the percent reduction in the person-level variance of a null model (**Table 2**) associated with the inclusion of the three predictors (collective narcissism, national identification, and political ideology) at the person-level. The three predictors accounted for 8% of the person-level variance of the contact measure, for 7% of the person-level variance of the hygiene measure, and 5% of the person-level variance of the policy support measure.

Next, we examined if person-level relations (slopes) between collective concerns and COVID-19 public health support varied across countries as a function of HDI by adding HDI scores to the level 3 model that examined slopes. The relations between national identification and each of the three public health measures were not moderated by HDI. Indeed, we observed only two modest moderating effects.

We found that HDI moderated the relations between national narcissism and spatial distancing ($g_{011} = -.031, t = 3.44, p < .05$). The relation between national narcissism and spatial distancing was negative in countries that had higher HDI scores (the estimated slope for a country +1 *SD* on the HDI was .037) but positive in countries that had lower HDI scores (the estimated slope for a country -1 *SD* on the HDI was .025). We also found that HDI moderated the relation between political ideology and hygiene ($g_{031} = -.017, t = 2.23, p < .05$). The overall negative relation between right-wing political ideology and hygiene was stronger in countries that had higher HDI scores (the estimated slope for a country +1 *SD* on the HDI was -.032) than in countries that had lower HDI scores (the estimated slope for a country -1 *SD* on the HDI was .002, functionally 0). However, we note that these effects were not statistically significant at the $p < .001$ threshold so we recommend interpreting them with caution.

References

- Alfano, M., Iurino, K., Stey, P., Robinson, B., Christen, M., Yu, F., & Lapsley, D. (2017). Development and validation of a multi-dimensional measure of intellectual humility. *PLoS ONE*, 12(8): e0182950.
- Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., & Yange, D. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics*, 04254.
- Aquino, K., & Reed, A. II. (2002). The self-importance of moral identity. *Journal of Personality and Social Psychology*, 83, 1423-1440.
- Aspelund, A., Lindeman, M., & Verkasalo, M. (2013). Political Conservatism and Left-Right Orientation in 28 Eastern and Western European Countries. *Political Psychology* 34, 409–417.
- Back, M. D., Küfner, A. C. P., Dufner, M., Gerlach, T. M., Rauthmann, J. F., & Denissen, J. J. A. (2013). Narcissistic admiration and rivalry: Disentangling the bright and dark sides of narcissism. *Journal of Personality and Social Psychology*, 105, 1013-1037.
- Barkan, R., Ayal, S., Gino, F., and Ariely, D. (2012). The pot calling the kettle black: distancing response to ethical dissonance. *Journal of Experimental Psychology: General*, 141, 757–773.
- Biddlestone, M., Green, R., & Douglas, K. (2020). Cultural orientation, powerlessness, belief in conspiracy theories, and intentions to reduce the spread of COVID-19. *British Journal of Social Psychology*, 59: 663-673
- Bieber, F. (2020). Global Nationalism in Times of the COVID-19 Pandemic. *Nationalities Papers*, 1–13.
- Bilewicz, M., Witkowska, M., Pantazi, M., Gkinopoulos, T., and Klein, O. (2019). Traumatic rift: How conspiracy beliefs undermine cohesion after societal trauma. *Europe's Journal of Psychology*, 15(1), 82-93.
- Bjørnskov, C. (2010). How comparable are the Gallup World Poll life satisfaction data? *Journal of Happiness Studies*, 11, 41-60.
- Block et al., (2020). Social network-based distancing strategies to flatten the COVID-19 curve in a post-lockdown world. *Nature Human Behaviour*, 4, 588-596.

25

Buchan, N. R., Brewer, M. B., Grimalda, G., Wilson, R. K., Fatas, E., & Foddy, M. (2011). Global social identity and global cooperation. *Psychological Science, 22*(6), 821-828.

Caluori, N., Jackson, J. C., Gray, K., & Gelfand, M. (2020). Conflict changes how people view god. *Psychological Science, 31*, 280-292.

Cameron, J. E. (2004). A three-factor model of social identity. *Self and Identity, 3*, 239-262.

Capraro, V., & Barcelo, H. (2020). The effect of messaging and gender on intentions to wear a face covering to slow down COVID-19 transmission. *Journal of Behavioral Economics for Policy, 4*(2), 45-55

Cichocka, A. (2016). Understanding defensive and secure in-group positivity: The role of collective narcissism. *European Review of Social Psychology, 27*(1), 283-317.

Cichocka, A., & Cislak, A. (2020). Nationalism as collective narcissism. *Current Opinion in Behavioral Sciences, 34*, 69-74.

Cichocka, A., Marchlewska, M., Golec de Zavala, A., & Olechowski, M. (2016). 'They will not control us': In-group positivity and belief in intergroup conspiracies. *British Journal of Psychology, 107*, 556-576.

Cislak, A., Wojcik, A. D., & Cichocka, A. (2018). Cutting the forest down to save your face: Narcissistic national identification predicts support for anti-conservation policies. *Journal of Environmental Psychology, 59*, 65-73

Curry, O. S., Chesters, M. J., & Van Lissa, C. J. (2019). Mapping morality with a compass: Testing the theory of 'morality-as-cooperation' with a new questionnaire. *Journal of Research in Personality, 78*, 106-124.

De Cremer, D., & Van Vugt, M. (1999). Social identification effects in social dilemmas: A transformation of motives. *European Journal of Social Psychology, 29*, 871-893.

Dong, E., Du, H., & Gardner, L. (2020). An interactive web-based dashboard to track COVID-19 in real time. *The Lancet Infectious Diseases, 20*(5), 533-534.

Douglas, K. M., Sutton, R. M. & Cichocka, A. (2017). The psychology of conspiracy theories. *Current Directions in Psychological Science, 26*, 538-542.

Dovidio, J. F., Ickler, E. G., Kunst, J. R., & Levy, A. (2020). Common identity and humanity. *Together apart the psychology of COVID-19*, 142-146.

- Ellemers, N. (2017). *Morality and the regulation of social behavior: Groups as moral anchors*. Milton Park, UK: Routledge.
- Ellemers, N., Van der Toorn, J. (2015). Groups as moral anchors. *Current Opinion in Behavioral Sciences*, 6, 189-194.
- Ellemers, N., Spears, R., & Doosje, B. (1999). *Social identity*. Blackwell: Oxford.
- Ellis, E.G. (2020). The coronavirus outbreak is a petri dish for conspiracy theories. *Wired*, retrieved from <https://www.wired.com/story/coronavirus-conspiracy-theories/>.
- Ferguson, N. M. (2006). Strategies for mitigating an influenza pandemic. *Nature*, 442, 448-452.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19, 25-42.
- Gadarian, S. K., Goodman, S. W., & Pepinsky, T. B. (2020). Partisanship, health behavior, and policy attitudes in the early stages of the COVID-19 pandemic. *SSRN*.
- Gaertner, S., Guerra, R., Rebelo, M., Dovidio, J., Hehman, E., & Deegan, M. (2016). The common ingroup identity model and the development of a functional perspective: A cross-national collaboration. In J. Vala, S. Waldzus, & M. M. Calheiros (Eds.), *The social developmental construction of violence and intergroup conflict* (p. 105–120). Springer International Publishing.
- Ghosh, D., Bernstein, J. A., & Mersha, T. B. (2020). COVID-19 pandemic: The African paradox. *Journal of Global Health*, 10(2).
- Gkinopoulos, T. and Hegarty, P. (2018). Commemoration in Crisis: A Discursive Analysis of Who 'We' and 'They' have been or become in ceremonial political speeches before and during the Greek Financial Downturn. *British Journal of Social Psychology*, 57(3), 591-609.
- Golec de Zavala, A., Cichocka, A., Eidelson, R., & Jayawickreme, N. (2009). Collective narcissism and its social consequences. *Journal of Personality and Social Psychology*, 97, 1074–1096. doi: 10.1037/a0016904
- Golec de Zavala, A., Cichocka, A., & Iskra-Golec, I. (2013). Collective narcissism moderates the effect of in-group image threat on intergroup hostility. *Journal of Personality and Social Psychology*, 104(6), 1019-1039.

Golec de Zavala, A., Peker, M., Guerra, R., & Baran, T. (2016). Collective narcissism predicts hypersensitivity to in-group insult and direct and indirect retaliatory intergroup hostility. *European Journal of Personality, 30*(6), 532-551.

Gollwitzer, A., Martel, C., Brady, W. J., Pärnamets, P., Freedman, I. G., Knowles, E. D., & Van Bavel, J. J. (2020). Partisan differences in spatial distancing predict infections and mortality. *Nature Human Behaviour, 4*, 1186–1197.

Gollwitzer, M., Platzer, C., Zwarg, C., & Göritz, A. S. (2020). Public acceptance of potential covid-19 lockdown scenarios. *Unpublished manuscript*.

Good, A., Harris, P. R., Jessop, D., and Abraham, C. (2015) Open-mindedness can decrease persuasion amongst adolescents: The role of self-affirmation. *British Journal of Health Psychology, 20*(2), 228-242.

Graham, J. (2013). Mapping the moral maps: From alternate taxonomies to competing predictions. *Personality and Social Psychology Review, 17*, 237-241.

Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology, 96*(5), 1029-1046.

Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review, 108*, 814-834.

Haidt, J., Kesebir, S. (2010). Morality. In Fiske, S., Gilbert, D., Lindzey, G. (Eds.), *Handbook of social psychology* (5th ed., pp. 797-832). Hoboken, NJ: Wiley.

Harari, Y. N. (2020). The world after coronavirus. *Financial Times, 20*. Retrieved from: <https://www.ft.com/content/19d90308-6858-11ea-a3c9-1fe6fedcca75>

Haslam, S. A., & Reicher, S. (2006). Stressing the group: social identity and the unfolding dynamics of responses to stress. *Journal of Applied Psychology, 91*(5), 1037-1052.

Haslam, S.A., Reicher, S.D. & Platow, M.J. (2011). *The New Psychology of Leadership: Identity, Influence, and Power*. Routledge.

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world?. *Behavioral and Brain Sciences, 33*, 61-83.

Horton, R. S., & Sedikides, C. (2009). Narcissistic responding to ego threat: when the status of the evaluator matters. *Journal of Personality, 77*, 1493–1525.

- Isler, O. I., Isler, B., Kopsacheilis, O., & Ferguson, E. (2020). Limits of the social-benefit motive among high-risk patients: A field experiment on influenza vaccination behaviour. *BMC Public Health*, *20*, 240.
- Jetten, J., Haslam, C., & Haslam, S. A. (2011). *The social cure*. Psychology Press.
- Jost, J. T. (2006). The end of the end of ideology. *American Psychologist*, *61*, 651-670.
- Jordan, J., Leliveld, M. C., & Tenbrunsel, A. E. (2015). The moral self-image scale: Measuring and understanding the malleability of the moral self. *Frontiers in Psychology*, *6*, Article 1878.
- Kachanoff, F. J., Bigman, Y. E., Kapsaskis, K., & Gray, K. (2020). Measuring realistic and symbolic threats of COVID-19 and their unique impacts on well-being and adherence to public health behaviors. *Social Psychological and Personality Science*.
- Kalin, M., & Sambanis, N. (2018). How to think about social identity. *Annual Review of Political Science*, *21*, 239-257.
- Khan, S. S., Garnett, N., Hult Khazaie, D., Liu, J. H., & Gil de Zúñiga, H. (2020). Opium of the people? National identification predicts well-being over time. *British Journal of Psychology*, *111*(2), 200–214.
- Koo, J. R., et al. (2020). Interventions to mitigate early spread of SARS-CoV-2. *The Lancet*, *20*, 678-688.
- Kramer, A., Guillory, J., & Hancock, J. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, *111*(24), 8788–8790.
- Leach, C. W., Van Zomeren, M., Zebel, S., Vliek, M. L., Pennekamp, S. F., Doosje, B., ... & Spears, R. (2008). Group-level self-definition and self-investment: a hierarchical (multicomponent) model of in-group identification. *Journal of Personality and Social Psychology*, *95*, 144-165.
- LeDoux, J. (2012). Rethinking the emotional brain. *Neuron*, *73*(4), 653-676.
- Lewnard, J. A., & Lo, N. C., (2020). Scientific and ethical basis for social-distancing interventions against COVID-19. *The Lancet*, *20*, 631-633.

Liu, J. H., Zhang, R. J., K.-Y. Leung, A., de Zúñiga, H. G., Gastardo-Conaco, C., Vasiutynskyi, V., Kus-Harbord, L. (2020). Empirical correlates of Cosmopolitan Orientation: Etiology and functions in a worldwide representative sample. *Political Psychology*.

Lyons, P. A., Coursey, L. E., & Kenworthy, J. B. (2013). National identity and group narcissism as predictors of intergroup attitudes toward undocumented Latino immigrants in the United States. *Hispanic Journal of Behavioral Sciences*, 35, 323-335. doi: 10.1177/0739986313488090

Lyons, P., Kenworthy, J., Popan, J. (2010). Ingroup identification and group-level narcissism as predictors of U.S. citizens' attitudes and behavior toward Arab immigrants. *Personality and Social Psychology Bulletin*, 36, 1267-1280.

Malone, G. P., Pillow, D. R., & Osman, A. (2012). The general belongingness scale (GBS): Assessing achieved belongingness. *Personality and Individual Differences*, 52, 311-316.

Marchlewska, M., Cichocka, A., Jaworska, M., Golec de Zavala, A., & Bilewicz, M. (2020). Superficial ingroup love? Collective narcissism predicts ingroup image defense, outgroup prejudice, and lower ingroup loyalty. *British Journal of Social Psychology*, 59, 857-875.

Marchlewska, M., Cichocka, A., Panayiotou, O., Castellanos, K., & Batayneh, J. (2018). Populism as identity politics: Perceived in-group disadvantage, collective narcissism, and support for populism. *Social Psychological and Personality Science*, 9(2), 151-162.

Martinez-Brawley, E. & Gualda, E. (2020). Transnational Social Implications of the Use of the "War metaphor" Concerning Coronavirus: A Birds' Eye View. *Culture e Studi del Sociale*. 5(1), 259-272.

McCauley, C. & Jacques, S. (1979). The popularity of conspiracy theories of presidential assassination: a Bayesian analysis. *Journal of Personality and Social Psychology*, 37, 637-644.

Merkley, E., Bridgman, A., Loewen, P. J., Owen, T., Ruths, D., & Zhilin, O. (2020). A rare moment of cross-partisan consensus: Elite and public responses to the COVID-19 pandemic in Canada. *Canadian Journal of Political Science*, 53, 311-318.

Michaels, J., and Norton, A. (2015). The complex interaction between self-esteem and self-control. Society for Personality and Social Psychology, Annual Conference. At: Long Beach, CA. Affiliation: Society for Personality and Social Psychology.

Muthén, L. K., & Muthén, B. O. (2017). *Mplus User's Guide* (8th ed.). Muthén & Muthén.

Nezlek, J. B. (2010). Multi-level modeling and cross-cultural research. In D. A. Matsumoto & A. F. R. van de Vijer (Eds.), *Cross-Cultural research methods in psychology* (pp. 299–347). Oxford University Press.

Nezlek, J. B. (2017). A practical guide to understanding reliability in studies of within-person variability. *Journal of Research in Personality*, *69*, 149–155.

Nowak, B., Brzóška, P., Piotrowski, J., Sedikides, C., Żemojtel-Piotrowska, M., & Jonason, P. K. (2020). Adaptive and maladaptive behavior during the COVID-19 pandemic: The roles of Dark Triad traits, collective narcissism, and health beliefs. *Personality and Individual Differences*, *167*, 110232.

Ponce, D. (2020). The impact of coronavirus in Brazil: politics and the pandemic. *Nature Review Nephrology*, *16*, 483.

Postmes, T., Haslam, S. A., & Jans, L. (2012). A single-item measure of social identification: Reliability, validity, and utility. *British Journal of Social Psychology*, *52*, 597–617.

Primi, C., Morsanyi, K., Chiesi, F., Donati, M. A., & Hamilton, J. (2016). The development and testing of a new version of the cognitive reflection test applying item response theory (IRT). *Journal of Behavioral Decision Making*, *29*, 453–469.

Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (2nd ed.). Sage Publications.

Raudenbush, S. W., Bryk, A. S., Cheong, Y. F., Congdon, R. T., & du Toit, M. (2011). *HLM 7: Linear and nonlinear modeling*. Chicago: Scientific Software International.

Rhodewalt, F., & Morf, C. C. (2005). Reflections in troubled waters: Narcissism and interpersonal self-esteem regulation. In A. Tesser, J. Wood, & D. Stapel (Eds.), *On building, defending, and regulating the self* (pp. 127–151). New York: Psychology Press.

Rhodewalt, F., Peterson, B. (2010). The self and intergroup attitudes: Connecting “fragile” personal and collective self-concepts. In Forgas, J. P., Cooper, J. E., Crano, W. D. (Eds.), *The psychology of attitudes and attitude change* (pp. 263–282). New York, NY: Psychology Press.

- Robins, R. W., Hendin, H. M., & Trzesniewski, K. H. (2001). Measuring global self-esteem: Construct validation of a single-item measure and the Rosenberg self-esteem scale. *Personality and Social Psychology Bulletin*, 27, 151–161. <https://doi.org/10.1177/0146167201272002>.
- Routledge, C. (2012). Failure causes fear: The effect of self-esteem threat on death-anxiety. *The Journal of Social Psychology*, 152(6), 665-669.
- Schatz, R. T., Staub, E., & Lavine, H. (1999). On the varieties of national attachment: Blind versus constructive patriotism. *Political Psychology*, 20(1), 151-174.
- Shrout, P. E. (1998). Measurement reliability and agreement in psychiatry. *Statistical Methods in Medical Research*, 7(3), 301–317.
- Sjåstad, H. (2019). Short-sighted greed? Focusing on the future promotes reputation-based generosity. *Judgment and Decision Making*, 14, 199-213.
- Sjåstad, H., & Van Bavel, J. J. (2020). The best-case heuristic in risk prediction: Hopes and fears in a global health pandemic (COVID-19). Pre-print retrieved at: [10.31234/osf.io/pcj4f](https://doi.org/10.31234/osf.io/pcj4f)
- Scheier, M. F., Carver, C. S., & Bridges, M. W. (1994). Distinguishing optimism from neuroticism (and trait anxiety, self-mastery, and self-esteem): A re-evaluation of the Life Orientation Test. *Journal of Personality and Social Psychology*, 67, 1063-1078.
- Tangney, J. P., Boone, A. L., & Baumeister, R. F. (2018). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. In *Self-regulation and self-control* (pp. 181-220). Routledge.
- Tomasello, M., Vaish, A. (2013). Origins of human cooperation and morality. *Annual Review of Psychology*, 64, 231-255.
- Tunçgenç, B., El Zein, M., Sulik, J., Newson, M., Zhao, Y., Dezechache, G., Deroy, O. (2020). We distance most when we think our social circle does. *Unpublished manuscript*.
- Twenge, J. M., & Campbell, W. K. (2003). Isn't it fun to get the respect we're going to deserve? Narcissism, social rejection, and aggression. *Personality and Social Psychology Bulletin*, 29, 261–272.
- Van Bavel, J. J., Baicker, K., Boggio, P. S., Capraro, V., Cichocka, A., Cikara, M., ... & Drury, J. (2020). Using social and behavioural science to support COVID-19 pandemic response. *Nature Human Behaviour*, 1-12.

Van der Toorn, J., Nail, P. R., Liviatan, I., & Jost, J. T. (2014). My country, right or wrong: Does activating system justification motivation eliminate the liberal-conservative gap in patriotism?. *Journal of Experimental Social Psychology, 54*, 50-60.

van Prooijen, J.-W. & Douglas, K. M. (2017). Conspiracy theories as part of history: the role of societal crisis situations. *Memory Studies, 10*, 323–333.

Van Zomeren, M., Postmes, T., & Spears, R. (2008). Toward an integrative social identity model of collective action: a quantitative research synthesis of three socio-psychological perspectives. *Psychological Bulletin, 134*(4), 504-535.

Velotti, P., Garofalo, C., Bottazzi, F., & Caretti, V. (2017a). Faces of shame: Implications for self-esteem, emotion regulation, aggression, and well-being. *The Journal of Psychology, 151*(2), 171–184.

Waytz, A., Iyer, R., Young, L. et al. (2019). Ideological differences in the expanse of the moral circle. *Nature Communications, 10*, 4389.

Supplement**Summary statistics****Table S1**

Summary of Samples. N stands for the sample size. Gender is coded as follows: 1=male, 2=female, 3=other (only 0.3% of the total sample reported gender=3). Standard deviations are reported in brackets. Cases/million and Deaths/million represent the number of cases (deaths) per million inhabitants at the start of data collection. Some countries have smaller sample sizes that fall short of $n = 500$; even so, we decided to include them in the analysis as MLM takes into account different numbers of observations.

Country	N	Mean gender	Mean age	Cases / million	Deaths / million
Argentina	698	1.70 (0.46)	47.32 (15.36)	89.97	4.42
Australia	2141	1.52 (0.51)	46.98 (17.52)	268.95	3.32
Austria	1390	1.54 (0.51)	49.82 (14.13)	1723.93	61.97
Bangladesh	389	1.57 (0.52)	32.31 (11.08)	36.64	0.94
Belgium	1153	1.41 (0.50)	46.24 (18.65)	4073.91	628.88
Bolivia	35	1.62 (0.49)	41.94 (12.85)	89.34	4.67
Brazil	1807	1.61 (0.50)	37.19 (13.95)	321.94	21.97
Bulgaria	500	1.66 (0.47)	29.80 (10.67)	194.71	8.29
Canada	901	1.61 (0.49)	43.57 (17.45)	1319.93	75.58
Chile	93	1.65 (0.48)	47.78 (15.33)	737.48	10.57
China	1030	1.49 (0.50)	43.24 (14.02)	70.16	3.47

34

Colombia	1211	1.65 (0.48)	40.77 (15.58)	112.73	5.10
Costa Rica	25	1.36 (0.49)	44.64 (12.73)	139.4	1.2
Croatia	473	1.51 (0.50)	45.95 (14.64)	502.22	14.53
Cuba	40	1.52 (0.51)	48.82 (12.74)	122.48	4.94
Denmark	553	1.49 (0.50)	48.81 (17.54)	1531.15	73.49
Dominican Republic	34	1.71 (0.46)	40.09 (13.93)	592	26.53
Ecuador	142	1.55 (0.50)	40.37 (11.82)	1360.66	38.82
El Salvador	28	1.54 (0.51)	46.42 (11.51)	50.31	38.82
Finland	679	1.55 (0.58)	38.47 (13.34)	850.54	34.96
France	1119	1.55 (0.50)	43.18 (16.20)	2477.06	348.16
Germany	1571	1.51 (0.51)	49.57 (16.13)	1912.29	73.79
Ghana	455	1.33 (0.47)	31.36 (7.73)	232.99	1.14
Greece	637	1.35 (0.49)	29.77 (11.43)	236.38	12.69
Guatemala	45	1.47 (0.50)	44.64 (13.20)	30.72	0.87
Honduras	24	1.71 (0.46)	39.25 (14.30)	68.93	6.36
Hungary	506	1.52 (0.50)	48.52 (16.54)	264.38	28.66
India	631	1.42 (0.50)	33.02 (11.83)	21.77	0.69
Iraq	544	1.49 (0.53)	29.94 (10.85)	48.06	2.29

35

Ireland	722	1.67 (0.47)	38.18 (14.46)	4009.80	224.90
Israel	1241	1.51 (0.50)	41.23 (15.23)	1751.69	22.97
Italy	1283	1.54 (0.50)	46.61 (16.61)	3303.74	446.93
Japan	1161	1.51 (0.51)	47.10 (15.18)	111.88	3.04
Latvia	935	1.67 (0.49)	45.40 (13.40)	426.04	6.77
Macedonia	696	1.57 (0.51)	37.95 (11.30)	672.6	31.25
Mexico	1236	1.49 (0.50)	47.61 (13.78)	123.05	11.36
Morocco	647	1.54 (0.51)	32.13 (12.43)	114.35	4.5
Nepal	346	1.55 (0.52)	27.99 (7.59)	1.85	0
Netherlands	1297	1.46 (0.50)	49.63 (16.83)	2224.54	262.38
New Zealand	509	1.50 (0.50)	45.78 (17.63)	301.02	3.89
Nicaragua	16	1.62 (0.50)	42.75 (14.84)	2.01	0.46
Nigeria	564	1.49 (0.50)	32.49 (16.30)	6.82	0.20
Norway	526	1.54 (0.50)	46.87 (17.26)	1399.45	37.75
Pakistan	474	1.54 (0.50)	26.72 (7.58)	65.57	1.38
Panama	17	1.65 (0.49)	44.12 (17.85)	1440.43	39.95
Paraguay	26	1.85 (0.37)	39.88 (9.64)	32.76	1.29
Peru	88	1.44 (0.50)	46.37 (14.38)	897.12	24.44

36

Philippines	513	1.51 (0.51)	36.74 (12.27)	72.89	4.79
Poland	1800	1.50 (0.50)	46.37 (17.02)	313.46	14.8
Romania	503	1.50 (0.51)	42.52 (14.46)	584.18	33.02
Russia	556	1.53 (0.50)	45.01 (15.49)	603.09	5.49
Senegal	343	1.37 (0.49)	35.38 (12.97)	46.44	0.57
Serbia	685	1.73 (0.44)	42.62 (11.83)	949.86	17.91
Singapore	542	1.51 (0.50)	43.71 (24.08)	2557.27	2.48
Slovakia	1225	1.50 (0.50)	44.09 (15.88)	252.93	3.3
South Africa	453	1.75 (0.44)	38.58 (13.11)	82.95	1.56
South Korea	495	1.48 (0.51)	41.90 (13.82)	208.21	4.72
Spain	1089	1.33 (0.47)	46.02 (13.68)	4887.56	501.09
Sweden	1563	1.41 (0.49)	52.92 (15.42)	1850.05	222.29
Switzerland	1052	1.51 (0,50)	47.93 (16.60)	3403.03	194.28
Taiwan	833	1.50 (0.50)	43.99 (13.25)	18.04	0.25
Turkey	1464	1.51 (0.50)	37.21 (15.20)	1369.04	35.37
Ukraine	577	1.53 (0.51)	37.45 (8.03)	214.60	5.24
United Kingdom	547	1.52 (0.50)	45.61 (15.57)	2375.81	317.43
United States	1471	1.52 (0.50)	44.32 (16.43)	3010.96	171.42

Uruguay	47	1.68 (0.47)	52.65 (13.94)	179.71	4.34
Venezuela	84	1.58 (0.50)	45.37 (12.67)	11.4	0.35

Details about the survey

Participants took, in random order, the following scales:

- A *limiting physical contact* scale, consisting of five items, as, for example, “During the days of the coronavirus (COVID-19) pandemic, I have been staying at home as much as practically possible”.
- A *physical hygiene* scale, consisting of five items, as, for example, “During the days of the coronavirus (COVID-19) pandemic, I have been washing my hands longer than usual”.
- A *policy support* scale, consisting of five items, as, for example, “During the days of the coronavirus (COVID-19) pandemic, I have been in favor of closing all schools and universities”.
- A *generosity* measure, measuring the proportion of the daily wage in the corresponding country a participant would keep for themselves vs. give to a national charity vs. give to an international charity (Sjåstad, 2019).
- A two-item *psychological well-being* scale (Bjørnskov, 2010).
- A three-item *collective narcissism* scale (Golec de Zavala, Cichocka, Eidelson, & Jayawickreme, 2009).
- A two-item *national identification* scale (one item from Postmes, Haslam, & Jans, 2012 and additional item measuring identity centrality).
- A *COVID-19 conspiracy beliefs* scale, consisting of four items such as “The coronavirus (COVID-19) is a bioweapon engineered by scientists.”
- A six-item *open mindedness* scale (Alfano et al, 2017).
- A seven-item *morality-as-cooperation* scale (Curry, Chesters, & Van Lissa, 2019).
- A two-item *trait optimism* scale (Scheier, Carver, & Bridges, 1994).
- A four-item *social belonging* scale (Malone, Pillow, & Osman, 2012).
- A four-item *trait self-control* scale (Tangney, Boone, & Baumeister, 2018).
- A one-item *self esteem* scale (Robins, Hendin, & Trzesniewski, 2001).
- A six-item *narcissism* scale (Back, Küfner, Dufner, Gerlach, Rauthmann, & Denissen, 2013).
- A ten-item *moral identity* scale (Aquino & Reed, 2002).
- A *COVID-19 risk perception* scale, consisting of two items such as “By April 30, 2021: How likely do you think it is that you will get infected by the Coronavirus (Covid-19)?” Available answers from 0% to 100%, with 10% increments.
- A one-item *political orientation* scale. Participants were asked: “Overall, what would be the best description of your political views?”. Available answers from 0 = very left leaning to 10 = very right leaning.

- A *moral circle* measure (Waytz, Iyer, Young, et al., 2019).
- A subjective physical health measure. Participants were asked: “In general, how would you rate your physical health as it is today?”. Available answers from 0 = “extremely bad” to 10 = “extremely good”.

After these scales, participants took a three-item *cognitive reflection test* (CRT). The test was a reworded version of the test proposed by Frederick (2005); we reworded the items because the classic CRT is very well known. After the CRT, participants answered some demographic questions. Full instructions, in English, are available at:

https://docs.google.com/document/d/1nAPFtxYIMsnYxdZbdCqpFqge3rrwar2f9FvmsYlwm_Y/edit?usp=sharing