Essays on Economic Decision-Making in Times of Crises

Inauguraldissertation
zur
Erlangung des Doktorgrades
der
Wirtschafts- und Sozialwissenschaftlichen Fakultät
der
Universität zu Köln

2024

vorgelegt
von
Sören Harrs, M.Sc. Economics
aus
Kiel

Referent*in: Prof. Dr. Bettina Rockenbach
Korreferent*in: Prof. Dr. Frederik Schwerter
Tag der Promotion: 25.07.2024

Acknowledgment

This dissertation would not have been possible without the help of many people and institutions that have supported me along the way.

First and foremost, I thank my main supervisor Bettina Rockenbach for her guidance and enduring support. Throughout my dissertation, Bettina has given me great freedom to explore my own research interests while always providing me with guidance and advice whenever it was needed. From working on our joint research projects and from teaching our joint course in game theory, I have learned a lot from her over the years. I am grateful that Bettina always encouraged my attendance at summer schools and conferences, and supported my research visits in Bergen and San Diego, which provided me with great opportunities to experience life in academia outside our department.

Second, I thank my supervisor Frederik Schwerter for his support, advice, and for always being willing to offer help, even on short notice. Through our discussions over the years, I have gained many valuable insights about doing research in our field of economics and about navigating academia as a PhD student.

Third, I am very thankful to my co-authors Maj-Britt Sterba, Lukas Wenner, Marvin Gleue, Christoph Feldhaus, Andreas Löschel, and Lara Marie Berger. They have immensely contributed to the research in this dissertation.

I am also especially thankful to everyone who supported me in the past year on the academic job market, including my supervisors, Bertil Tungodden, and the faculty and colleagues at the University of Cologne. I would also like to thank my current and past colleagues at the Chair of Behavioral and Experimental Economics: Heike Kirch, Sebastian Tonke, Susanna Grundmann, Lukas Reinhardt, Sebastian Schneiders, Thomas Lauer, Eugenio Verrina, Viet A. Nguyen, and Lukas Wenner.

I am grateful for the financial support of the Cluster of Excellence ECONtribute, the Cologne Graduate School (CGS), the Center for Social and Economic Behavior (C-SEB), the Max Planck Institute for Research on Collective Goods (MPI), and the Chair of Behavioral and Experimental Economics. By funding my research projects, conferences, and research visits, these institutions have provided me with excellent research conditions during my dissertation.

In particular, I thank Bertil Tungodden and Alexander Cappelen for hosting me during my research visit at the FAIR Center at NHH in Bergen, and Uri Gneezy for hosting me during my research visit at the UC San Diego.

I would also like to thank those who are not listed above but who have contributed to this dissertation in one way or another. Many people have given comments, feedback, or provided inspiration for the research in this dissertation. Looking back, I am particularly grateful for the friends I have made in academia along the way. They have made working on this dissertation a very enjoyable experience.

Last but certainly not least, I would like to thank my family and friends. They have supported me in many visible and invisible ways over the years. In particular, I am deeply thankful to my parents, Birte and Torsten. Without their efforts and support, from my childhood, throughout my education, until today, I would not have arrived at writing this dissertation. I therefore dedicate this dissertation to my parents.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In the past years, people around the world have experienced a "polycrisis" consisting of the unprecedented global coronavirus pandemic, the looming ecological crisis of climate change, and increasing geopolitical conflicts.[1] These societal crises have led to profound changes in people's lives and in their economic circumstances. Throughout history, similar pandemics, natural disasters, and wars have often led to long-term changes in economic and political institutions (e.g. Acemoglu et al., 2021). Thus, times of crises can be turning points in the lives of individuals and in the history of societies.

From a societal point of view, most crises typically cause a loss of income and wealth, an increase in uncertainty about the future, and a general reduction in well-being. From a scientific point of view, however, societal crises can provide unique opportunities. Societal crises, on the one hand, provide the opportunity to study the impacts of rare and unique shocks. During the coronavirus pandemic, for example, economists could suddenly study how governmental lockdowns impact the macroeconomy, or how large-scale cash transfers impact poverty rates. Studying these unique shocks can provide novel insights into the functioning of the economic system. On the other hand, societal crises can provide ideal settings to learn something fundamental about human decision-making. Societal crises, for example, provide ideal settings to study how people react to economic shocks and economic inequalities, how people cooperate to mitigate a crisis, or how people make economic decisions under uncertainty.

This dissertation makes use of these opportunities and studies people's economic decision-making in times of crises. The topic of this dissertation is to a large part the result of the economic, ecological, and societal crises that took place at the time of its writing (from 2019 to 2024). In particular, the dissertation focuses on economic decision-making in the context of the coronavirus crisis and in the context of the climate crisis.

[1] See, for example, historian Adam Tooze explaining the term "polycrises" in an interview with the World Economic Forum https://www.weforum.org/agenda/2023/03/polycrisis-adam-tooze-historian-explains/, accessed on June 1, 2024.

This dissertation consists of four chapters. While each chapter is based on an independent research paper, the chapters are connected through three common themes that extend beyond the context of crises. The first common theme is the use of experimental methods, in particular the use of randomized experiments and the use of abstract economic games. The second common theme is a focus on understanding how preferences and beliefs - two fundamental building blocks of economic theory - interact in shaping economic decisions. The third common theme is that the chapters make contributions to the literature in behavioral economics by advancing our understanding of how "non-standard" motives like fairness and identity influence economic decisions. Subsequently, I will summarize each of the chapters in detail.

In **Chapter 2**, titled *Fairness and Support for Welfare Policies*, (with Maj-Britt Sterba), we study a central economic decision in times of crisis: how much the government should financially support those individuals that are hit by the crisis. Using representative online surveys in the US during the coronavirus pandemic, we provide novel evidence about why some people support welfare policies and others oppose them. Specifically, we study the role of people's fairness views about economic inequality. A novel feature of our study is that it collects detailed data about two determinants of fairness views: fairness preferences - revealed through consequential transfer choices in an experimental game - and beliefs about the causes of inequality in society. Our analysis proceeds in two parts. In Part 1, we establish that there are large and robust differences in support for welfare policies between individuals with egalitarian, libertarian, and meritocratic fairness preferences. The findings in Part 1 reject the assumption of homogeneous fairness preferences in prominent models in political economy (Alesina and Angeletos, 2005; Alesina and Giuliano, 2011), and extend previous findings from laboratory experiments (Cappelen et al., 2007, 2013). Thereby, Part 1 addresses one of the key criticisms faced by experimental economic research: that the findings in laboratory experiments may not generalize to choices outside the lab (Levitt and List, 2007).

In Part 2, we then test how stable people's fairness views about economic inequality are during the coronavirus pandemic by using individual-level panel data collected between Spring 2020 and Fall 2021. We find that the US population does not become more egalitarian or libertarian over time, but instead remains predominantly meritocratic in their fairness preferences. However, over the course of the pandemic, US Americans on average reduce their belief in merit as the main cause of economic inequality in US society, and increase their belief in the importance of good and bad luck. Those individuals who change their beliefs about the causes of inequality also change their support for welfare policies. Thereby, the findings in Part 2 confirm the central mechanism in prominent models in political economy: that belief shocks can explain the evolution of policy support over time (Alesina and Angeletos, 2005; Alesina and Giuliano, 2011; Bénabou and Tirole, 2006; Piketty, 1995). By documenting that beliefs in merit decline over the course of the pandemic, our findings suggest that the pandemic could lead to lasting changes in people's support for welfare policies, and potentially, to long-term changes in institutions of the welfare state.

Chapters 3 and 4 of this dissertation then study a second central economic decision in times of crisis: whether individuals voluntarily contribute to the prevention and mitigation of a crisis. In the context of the coronavirus crisis, individuals could voluntarily contribute to the mitigation of the pandemic by engaging in social distancing to contain the spread of the coronavirus. In the context of the climate crisis, individuals can voluntarily contribute to the mitigation of climate change by buying products that are produced in an eco-friendly way or by directly donating to climate protection. But what motivates individuals to engage in costly voluntary actions to prevent and mitigate a crisis? And under what conditions can high levels of such voluntary actions be reached?

In **Chapter 3**, *Revealing Good Deeds: Disclosure of Social Responsibility in Competitive Markets* (with Bettina Rockenbach and Lukas Wenner, published in **Experimental Economics** in 2022), we use market experiments to study how the purchase of products that are produced in a fair and socially responsible way depends on the disclosure of information by producers. In five treatment conditions, we experimentally vary how information about the social externality of products can be disclosed by producers to consumers. Our data show that voluntary disclosure of information can lead to high levels of socially responsible consumption through market competition, but only if the disclosed information is reliable. The results of Chapter 3 show that even though many consumers have preferences for fair and socially responsible products, these preferences are not sufficient to achieve high levels of socially responsible consumption in markets. Only with standardized and unambiguous information can consumers form accurate beliefs about social externalities, and thus make socially responsible consumption decisions. Therefore, our results justify governmental policies that regulate the information disclosure in markets, for example, through standardized eco-labels.

In **Chapter 4**, *Identity and Voluntary Efforts for Climate Protection* (with Marvin Gleue, Christoph Feldhaus, and Andreas Löschel, published in the **Journal of Economic Behavior & Organization** in 2024), we conduct a field experiment to study how people's identity concerns shape their choice to generate donations for climate protection. As our experimental manipulation, we ask questions designed to change subjects' beliefs about how environmentally friendly they personally are. We implement positive and negative shocks to subjects' environmental identity beliefs compared to a control group. We find that a negative shock to identity beliefs increases voluntary efforts for climate protection, especially among individuals with a strong prior identity belief. The findings in Chapter

4 advance the literature on identity in economics by providing evidence consistent with the theory of Bénabou and Tirole (2011). According to Bénabou and Tirole (2011), an identity is a malleable belief that people infer from their own actions, rather than a stable preference as in Akerlof and Kranton (2000). Our findings in Chapter 4 also have implications for NGOs and policymakers that would like to encourage more voluntary actions for climate protection. Our findings highlight the opportunities and limitations of using people's identity concerns for that purpose.

In **Chapter 5**, *How Narratives Influence Economic Decision-Making: Experimental Evidence* (with Lara Marie Berger and Bettina Rockenbach), we study how narratives about societal crises can change economic decision-making. In an information provision experiment, we provide subjects with naturalistic news articles that either contain an optimistic, a pessimistic, or a balanced narrative about the future course of the coronavirus pandemic. We then document the effects of these narratives on three fundamental determinants of economic decision-making: stock market expectations, risk preferences, and time preferences. We find that after reading a more pessimistic narrative about the pandemic, subjects hold more pessimistic beliefs about the development of the stock market. However, reading these narratives also has strong effects on people's risk and time preferences: the more pessimistic the narrative, the more risk-averse and impatient subjects become in incentivized economic games. These results provide evidence that narratives about societal crises can influence economic decision-making through two persuasion mechanisms: belief-based persuasion and preference-based persuasion. Thereby, Chapter 5 provides novel insights about how communication about societal crises can influence economic decision-making.

Contribution

Below I describe my contributions to each chapter of this dissertation.

Chapter 2 is joint work with Maj-Britt Sterba. Maj-Britt and I were equally involved in the design of the study and the data collection. I conducted the data analysis. We jointly wrote the paper.

Chapter 3 is joint work with Bettina Rockenbach and Lukas Wenner. Bettina and Lukas designed and collected the data for the main experiment. I collected the data for the online follow-up experiment. The first version of the paper was written by Bettina and Lukas. We then jointly worked on the final version of the paper. Lukas conducted most of the data analysis. I prepared the replication package.

Chapter 4 is joint work with Marvin Gleue, Christoph Feldhaus, and Andreas Löschel. All authors were involved in the design of the experiment. Marvin, Christoph, and I collected the field data. Marvin and I conducted the data analysis and mainly wrote the paper. I prepared the replication package.

Chapter 5 is joint work with Lara Marie Berger and Bettina Rockenbach. All authors were involved in the design of the experiments, the data collection, and in writing the current version of the paper. Lara and I conducted the data analysis. I prepared the replication package.

Chapter 2

Fairness and Support for Welfare Policies

This chapter is based on the paper "*Fairness and Support for Welfare Policies*" which is joint work with Maj-Britt Sterba. This paper was my job market paper for the 2023/24 academic job market and has been published in a very similar version on my personal homepage.[1]

Abstract How do fairness views shape people's support for welfare policies? Using large surveys in representative samples of US Americans, we study the roles of two determinants of fairness views: fairness preferences - revealed through transfer choices in a spectator game - and beliefs about the causes of inequality. We establish three novel findings: First, people with egalitarian, libertarian, and meritocratic fairness preferences differ strongly in their support for welfare policies. Second, beliefs about the causes of inequality have a strong effect on the policy support of meritocrats, but a much weaker effect on non-meritocrats. Third, leveraging individual-level panel data collected during the coronavirus pandemic, we show that shifts in support for welfare policies over time are rather caused by shocks to beliefs, than by shocks to fairness preferences. Our findings demonstrate that heterogeneous fairness preferences and beliefs interact in shaping people's support for welfare policies, which has theoretical implications for models in political economy. Our paper also has practical implications because it documents a declining belief in a meritocratic US society, which may have long-term consequences for the US welfare state.

[1] We thank Bertil Tungodden, Alexander Cappelen, Eliana La Ferrara, Bettina Rockenbach, Chris Roth, Frederik Schwerter, Sebastian Tonke, Oliver Kirchkamp, Christoph Engel, Love Christensen, and participants at various seminars and conferences for very helpful comments. We gratefully acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany´s Excellence Strategy – EXC 2126/1– 390838866, from the Center for Social and Economic Behavior (C-SEB) and from the Max Planck Society. IRB approval has been obtained from the University of Cologne (210013SH).

2.1 Introduction

People disagree about the fairness of economic inequalities. Across the social sciences, scholars have argued that people's disagreements about fairness - and not just their economic self-interest - are central to understanding people's polarized views about welfare policies. However, two influential strands of literature in economics have provided competing theories why people disagree about the fairness of economic inequalities. One strand of literature explains disagreements about fairness with opposing beliefs about the causes of inequality (e.g. Alesina and Angeletos, 2005; Alesina and La Ferrara, 2005; Fong, 2001), another strand of literature with egalitarian, libertarian, and meritocratic fairness preferences (e.g. Almås et al., 2020; Cappelen et al., 2007, 2013), that is with opposing normative views about what type of inequalities should be considered fair.

Understanding whether conflicts about welfare policies are based on polarized beliefs about the causes of inequality, or on polarized fairness preferences, seems to be of great societal relevance, in particular in times of rising economic inequality and strong political polarization. Yet, the existing evidence about this question is surprisingly limited. While the literature has repeatedly established a strong correlation between people's beliefs about merit and luck as causes of inequality and their support for redistributive policies (e.g. Alesina and Giuliano, 2011; Fong, 2001; Fong and Poutvaara, 2019; Stantcheva, 2021), no study has yet systematically disentangled the roles of fairness preferences and beliefs as determinants of people's policy preferences.

Fundamental questions about how fairness shapes people's policy preferences are therefore unresolved: Do people with egalitarian, libertarian, and meritocratic fairness preferences differ in their support for welfare policies? Are beliefs about the causes of inequality in society equally important to all people, or does the importance of beliefs vary depending on people's fairness preferences? And can we explain shifts in support for welfare policies over time with changes in people's fairness preferences or rather with changes in their beliefs about the causes of inequality?

In this paper, we study these questions using large surveys conducted with representative samples of US Americans (N=1975) in the context of the coronavirus pandemic. In our surveys, we collect detailed data about people's fairness preferences, their beliefs about the causes of inequality in society, and their support for welfare policies. Our empirical analysis is structured by a simple theoretical framework: People's preferences for taxes and transfers in society are explained by a trade-off between economic self-interest and fairness views. Fairness views, in turn, are determined by an interaction of heterogeneous fairness preferences with heterogeneous beliefs about the causes of inequality. In the first part of this paper, we test predictions from this theoretical framework. This empirical exercise shows how well fairness preferences and beliefs can descriptively explain why people are polarized in their support for welfare policies, that is, why some people strongly support welfare policies, while others strongly oppose welfare policies. In the second part of this paper, we then study if and how the exogenous shock of a major societal crisis changes people's views about the fairness of economic inequalities and their support for welfare policies. For that purpose, we use a subset of our data that consists of individual-level panel data (N=499), collected in two waves during the coronavirus pandemic (in Spring 2020 and Fall 2021), and a survey experiment (N=745), in which subjects write about personal experiences from the first wave of the coronavirus pandemic. This second empirical part allows us to move beyond the cross-sectional analysis of the first part and to study the stability and dynamics of fairness views and support for welfare policies over time.

There are at least three important benefits to understanding how fairness shapes people's support for welfare policies. First, these insights can help predict whether and under what conditions political agreements on central policy issues may be reached. If political conflicts are based on stable and opposing fairness preferences, any agreement seems hard to reach - even if all individuals in society share the same beliefs about reality. However, if political conflicts are based on malleable beliefs, changes in opinions and agreements appear more likely - especially in contexts in which transparent information about the causes of inequality is available. Second, it may allow us to better understand what type of policies the public perceives as fair. Analyzing the fairness views of the public could prove crucial for evaluating the political feasibility of any policy reform that causes or reduces economic inequalities. Welfare policies seem to be of first-order importance for economists due to their far-reaching impacts on labor markets, the macroeconomy, and the livelihoods of individuals affected by illnesses and unemployment. Third, these insights may eventually improve our understanding of why institutions of the welfare state change over time and vary across countries - one of the central puzzles discussed in the related theoretical literature (Alesina and Angeletos, 2005; Alesina et al., 2012; Bénabou and Tirole, 2006; Piketty, 1995).

Answering our research questions comes with one main empirical challenge: it requires a measure of fairness preferences types at the individual level. However, the existing literature has not yet established and validated an individual-level measure of the egalitarian, libertarian and meritocratic fairness preference type (see Cappelen et al., 2022a; Fehr and Charness, 2023, for excellent reviews).[2] To address this challenge, we use a simple within-subjects design in a spectator game to identify individual-level heterogeneity in fairness

[2]Instead, the most recent literature has typically used spectator games with between-subject designs (see Almås et al., 2020; Cohn et al., 2023), which only allow to estimate a distribution of fairness preference types at the population level. The foundational work of Cappelen et al. (2007, 2013) has used within-subject designs, but Cappelen et al. (2007) do not classify subjects at the individual level, and Cappelen et al. (2013) are not concerned with the egalitarian, libertarian and meritocratic fairness preference type.

preferences. Based on their transfer choices in our spectator game, *individuals* can be classified into the distinct fairness preference types (egalitarians, libertarians, meritocrats). This is crucial for our analysis because only this individual-level type classification allows us to test whether fairness preference types differ in their support for welfare policies and whether they interact with beliefs about the causes of inequality as predicted by theory.

In our within-subjects design, each subject takes two transfer decisions between two anonymous workers, one in which an inequality in payments (4$ vs 0$) is caused by a lottery ("luck"), and one in which it is caused by differences in productivity in a real effort task ("merit"). *Egalitarians* equalize inequalities in both choice contexts (11.9% of our sample who transfer \$2 in both contexts). *Libertarians* never equalize any inequality (5.3% of our sample who transfer \$0 in both contexts). *Meritocrats* choose to fully redistribute inequalities caused by luck (t_{Luck} =\$2), but are willing to accept inequalities caused by merit (t_{Merit} <\$2) (49.7% of our sample). While the meritocratic type predominates in our US sample (in line with the previous literature), our within-subject design reveals that there is also a lot of heterogeneity among *Meritocrats*: *Meritocrats* can be classified into four subtypes (with *strong* to *weak meritocratic preferences*) according to the difference they make between transfers on inequalities due to luck and merit. The choice behavior of a third of subjects (33.1%) is not consistent with any of the narrowly defined fairness preference types ("*Other*"). Taken together, our within-subjects design reveals that there is substantial heterogeneity in fairness preferences at the individual level. To speak to the literature, we focus our main analysis on the egalitarian, libertarian, and meritocratic fairness preference types, and extend it to the meritocratic subtypes where it is insightful.

We use established survey questions to measure heterogeneity in beliefs about the causes of economic inequality in US society and policy support for a comprehensive set of welfare policies. The set of welfare policies comprises policies implemented by the US government to support people in economic need during the COVID crisis, such as the temporary expansion of unemployment benefits, as well as long-term policy reforms of the US welfare state, such as universal health care, and subjects' general support for redistribution in society.

As our first main result, we show that fairness preferences strongly and independently predict support for welfare policies in line with theoretical predictions. *Egalitarians*, on average, show the strongest support for welfare policies, *Libertarians* show the strongest opposition against welfare policies, and *Meritocrats* lie in-between these two extreme positions. The mean differences in support for welfare policies between all three fairness preference types are significant and remarkably large. *Egalitarians* on average have a 0.85 standard deviations higher support for welfare policies than *Libertarians*. Across our set of welfare policies, the predictive power of the fairness preference types for policy support proves to be stronger than that of major socio-demographics including income and employment status. Fairness preferences remain strong predictors of policy preferences even when controlling for a variety of variables, such as socio-demographics, beliefs about the causes of inequality, left-right political ideology, altruism/selfishness and trust in government.

As our second main result, we find that fairness preferences and beliefs interact in line with theoretical predictions. The policy preferences of *Meritocrats* strongly depend on their beliefs about the causes of economic inequality in US society. The policy preferences of non-meritocrats, in contrast, depend to a significantly weaker degree on these beliefs. Moreover, beliefs have a significantly stronger influence on the policy preferences of *Meritocrat*s with *strong* meritocratic preferences compared to *Meritocrats* with *weak* meritocratic preferences. Hence, those subjects whose transfer choices in the spectator game are more sensitive to the causes of inequality, also put more weight on beliefs about the causes of inequality in US society when forming their policy preferences.

The sizes of the estimated coefficients indicate that heterogeneity in fairness preferences and heterogeneity in beliefs are of similar importance for explaining people's polarized views about welfare policies. *Meritocrats* who believe that economic inequalities in US society are mainly due to luck are as supportive of welfare policies as *Egalitarians*. *Meritocrats* who, in contrast, believe that economic inequalities are mainly due to merit are at least as opposed to welfare policies as *Libertarians*. In that way, differences in policy positions between individuals with opposing fairness preferences (*Egalitarians* versus *Libertarians*) are comparable in size to differences between *Meritocrats* with opposing beliefs (merit vs luck).

As our third main result, we find that shifts in support for welfare policies in times of societal crises are rather caused by shocks to beliefs about the causes of inequality, than by shocks to fairness preferences. In our panel data and in our experiment, we do not find significant changes in the distribution of fairness preference types. Moreover, changes in transfer choices at the individual level over time are not meaningfully related to changes in support for welfare policies. Beliefs about the causes of inequality in US society, in contrast, change considerably over time. Between the two waves of our panel data, between Spring 2020 and Fall 2021, beliefs in merit as the main cause of economic inequality in US society decreased by about 0.11 SD. This result is consistent with data from the General Social Survey, which reveals that over the pandemic (from 2018 to 2022) the belief in a meritocratic US society has declined substantially, and is now at its lowest value since the 2007/2008 financial crisis. Leveraging our individual-level panel data, we show that those subjects who reduce their belief in merit over time also increase their support for welfare policies.

Our data also provide evidence for one specific mechanism that can explain changes in beliefs about the causes of inequality in times of societal crises: personal experiences in which people lose control over their own lives. These types of experiences were omnipresent in the pandemic and seem to shatter the belief in a meritocratic society. In our panel data, only those subjects who report a lower sense of control over their own lives in the second wave of data collection, on average, reduce their beliefs in a meritocratic society. With our experiment we provide causal evidence for the proposed mechanism: it shows that subjects who are asked to recall experiences of *low control* from the first months of the pandemic reduce their beliefs in merit as a cause of inequality compared to those who are asked to recall an experience of *high control*.

Taken together, our results contribute to a better understanding of how fairness preferences and beliefs about the causes of inequality interact in shaping people's fairness views and policy preferences. In that way, our paper bridges the two separated strands of literature on fairness in economics: the survey-based literature on beliefs (e.g. Alesina and Giuliano, 2011; Alesina and La Ferrara, 2005; Fong, 2001; Fong and Poutvaara, 2019), and the laboratory-based literature on fairness preferences (e.g. Cappelen et al., 2007, 2013; Konow, 2000).[3]

Our paper in particular contributes to the rapidly growing literature on fairness preferences (e.g. Almås et al., 2020; Andre, 2024; Cappelen et al., 2023; Cappelen et al., 2022b; Cohn et al., 2023). While many recent papers advance our understanding of fairness preferences by experimentally introducing new features in the spectator game, our paper advances the literature by (i) providing evidence that heterogeneity in fairness preferences is important for understanding people's policy preferences and (ii) studying the stability of fairness preferences over time.

Closest to our paper regarding the external validity of fairness preferences are Almås et al. (2020) and Cohn et al. (2023). Almås et al. (2020) show that there is a simple correlation between the level of inequality implemented in the spectator game and the view that society should aim to equalize outcomes. Cohn et al. (2023) show that differences in inequality acceptance in the spectator game may explain differences in policy preferences between the top 5% and the bottom 95% of the US income distribution. Our approach differs from the existing literature in multiple ways. Our within-subjects design allows us to show for the first time that there are large and robust differences in policy preferences between individuals with egalitarian, libertarian and meritocratic fairness preferences. We thus base our analysis on the key theoretical distinction between different fairness preference types and not on general inequality acceptance. Further, we provide first evidence

[3] A third strand of literature in economics studies fairness views directly and does not empirically distinguish between fairness preferences and beliefs (Hvidberg et al., 2023; Stantcheva, 2021).

that these types meaningfully interact with beliefs in shaping policy preferences, and we control for a variety of potentially confounding variables in our descriptive analysis.[4]

Our paper also contributes to the literature on the stability of social preferences and fairness views (e.g. Almås et al., 2010; Fehr et al., 2008, 2013; Kosse et al., 2020), especially the strand of literature that is interested in how life-events in adulthood may shape fairness views (Barr et al., 2016; Cappelen et al., 2022a; Hvidberg et al., 2023; Roth and Wohlfart, 2018). In contrast to the findings by Barr et al. (2016) and Cappelen et al. (2021), our study provides evidence that changes in fairness views in times of personal and societal crises are rather driven by changes in beliefs than by changes in fairness preferences. Our data thereby also show that experience effects in belief formation (see Malmendier, 2021) seem to have an important application in political economy, confirming Piketty (1995).

Taken together, our findings provide two important insights for theories in political economy. First, our data provide novel empirical evidence that shocks to beliefs about the causes of inequality can explain changes in policy preferences over time, which is a central mechanism in many prominent models in political economy (Alesina and Angeletos, 2005; Alesina et al., 2012; Bénabou and Tirole, 2006; Piketty, 1995).[5] Second, our findings imply that theories in political economy should not only focus on beliefs about the causes of inequality but should also incorporate heterogeneous fairness preferences. Our findings demonstrate that heterogeneity in fairness preferences matters for explaining people's policy preferences, even in a country like the US where people are predominantly meritocratic. In other countries with much more heterogeneity in fairness preferences, the integration of fairness preferences should prove even more important.

Methodologically, our paper contributes to the literature by showing that a simple within-subjects design - based on just two transfer choices - can recover meaningful heterogeneity in fairness preferences. Given its simplicity, this measure can well be implemented in large-scale surveys. An individual-level measure of fairness preferences seems important for studying where disagreements about fairness come from. It may also prove useful in many other settings, ranging from wage setting in firms to support for affirmative action, in which heterogeneity in fairness preferences may be relevant for understanding people's demand for fair institutions and fair policies.

The remainder of this paper proceeds as follows. Section 2.2 introduces our simple theoretical framework. Section 2.3 describes our survey design, and Section 2.4 our mea-

[4] Our findings are also related to a broader literature which studies how other types of social preferences shape policy preferences and political ideology (Enke et al., 2023; Epper et al., 2020; Fisman et al., 2017; Kerschbamer and Müller, 2020). This literature studies social preferences that are not sensitive to the causes of inequality and are therefore distinct from fairness preferences. Closest to our paper methodologically are Epper et al. (2020), who follow a type-based approach and find that inequality-averse and altruistic individuals in Switzerland are more likely to support plebiscites about redistributive policies than selfish individuals.

[5] A seminal empirical paper in this literature has recently been retracted (Giuliano and Spilimbergo, 2014).

surement of fairness views. Sections 2.5 and 2.6 present our main results. Section 2.7 concludes.

2.2 Theory on Fairness Views: Preferences and Beliefs

In this section, we present a simple theoretical framework that guides our survey design, our measurement of fairness preferences and beliefs, and the subsequent empirical analyses. The theoretical framework incorporates heterogeneous fairness preferences and heterogeneous beliefs about the causes of inequality, building on Cappelen et al. (2007), into a simple political economy framework along the lines of the classical Meltzer and Richard (1981) model.

Individuals are motivated by own income and by fairness considerations. An individual's income before taxes, y, is exogenously determined on the market. The disposable income depends on the size of a tax and transfer system implemented in society. A transfer system specifies the percent of the total income collected as a tax, $\tau \in [0, 1]$, which is then distributed as a lump-sum transfer to all individuals in society, $T = \frac{1}{N}\sum_{i=1}^{N} \tau \cdot y_i = \tau \cdot \bar{y}$. For simplicity, individuals do not pay attention to the efficiency costs of redistribution.[6] Individuals maximize the following utility function when voting for their preferred tax rate τ in society:

$$U_i = (1 - \tau) \cdot y_i + T - \gamma \cdot \Omega \tag{2.1}$$

The parameter $\gamma \geq 0$ determines the weight people attach to fairness considerations and Ω represents the disutility generated by unfair social outcomes, as in Alesina and Angeletos (2005). In contrast to Alesina and Angeletos (2005), but similar to Cappelen et al. (2007), we assume that this disutility takes the following form:

$$\Omega = \frac{(\tau - F_i(b_i))^2}{2} \tag{2.2}$$

An individual's fairness view, $F_i(b_i) \in [0, 1]$, specifies what the individual perceives to be a fair tax rate in society. These fairness views are jointly shaped by fairness preferences, $F_i(\cdot)$, and beliefs about the causes of income inequality, $b_i \in [0, 1]$. Specifically, $b_i = 0$ corresponds to the belief that inequality is entirely caused by merit and $b_i = 1$ to the belief that inequality is entirely caused by luck. Fairness preferences are characterized by a function that maps any level of beliefs into a fair tax rate. For individuals with meritocratic fairness preferences, the fair tax rate is belief-dependent: it increases with

[6] Stantcheva (2021) shows that, when it comes to the formation of people's policy preferences, efficiency concerns seem to play surprisingly little role. In a similar way, Almås et al. (2020) show experimentally that fairness concerns seem to be more important than efficiency concerns.

the belief in luck as the main cause of inequality in society, $F'_{Mer}(b_i) > 0$. Meritocrats do not accept inequalities due to luck, $F_{Mer}(Luck) = 1$, but accept inequalities due to merit (to some extent), $F_{Mer}(Merit) < 1$. For individuals with other fairness preferences, fairness views are independent of beliefs. Egalitarians, for example, always view equality in economic outcomes as fair: $F_{Ega}(b_i) = 1$ for all beliefs. Libertarians, on the other hand, always view market outcomes as fair no matter how they were generated: $F_{Lib}(b_i) = 0$ for all beliefs. Therefore, fairness preferences determine whether and to what extent beliefs about the causes of inequality matter for fairness views.

Given an interior solution, the preferred tax rate, τ^*, then corresponds to:

$$\tau^* = \frac{1}{\gamma}(\bar{y} - y_i) + F_i(b_i) \tag{2.3}$$

As in the classical Meltzer-Richard model, the preferred tax rate depends on the difference between the income of a voter and the mean income in society: a lower relative income increases the preferred tax and transfer level. However, fairness concerns, γ, now reduce the influence of economic self-interest on policy preferences. Individuals also take their fairness views, $F_i(b_i)$, into account. In our paper, we aim to provide novel insights into the fundamental properties of people's fairness views $F_i(b_i)$. In contrast to the existing literature, we therefore directly measure people's fairness preferences, $F_i(\cdot)$, and their beliefs about the causes of inequality in society, b_i, in order to disentangle their effects on people's policy preferences.

2.3 Data Collection and Survey Design

2.3.1 Data Collection

We ran our surveys on Prolific, a large online survey platform focused on scientific research (`https://www.prolific.com/`).[7] Three key advantages of using Prolific for our data collection are, first, that it provides samples that are broadly representative of the US population (see 2.3.2), second, that it is known for high data quality (see 2.3.3), and third, that it allows to re-invite the same participants, which enabled us to collect our individual-level panel data.

We collected our data in two waves. Data for Wave 1 were collected between May 15th and May 17th 2020 (N=745). Data for Wave 2 were collected around 1 1/2 years later, between September 25th 2021, and January 3rd 2022. Wave 2 consisted of re-sampling

[7] In recent years online platforms like Prolific and Amazon MTurk have been increasingly used in economic research (for example in Cappelen et al. (2023), DellaVigna and Pope (2018), and Kuziemko et al. (2015)). Replications of classic experiments on online platforms show that results on online platforms are consistent with results in the laboratory (Snowberg and Yariv, 2021).

around two-thirds of participants from Wave 1 (N=499)[8] and a new sample of participants (N=729). See Appendix Figure A.13 for a timeline of the data collection.

2.3.2 Sample

Through Prolific we obtained representative samples of US Americans, stratified by age, gender and race. Participation was restricted to subjects who were 18 years or older. Table 2.1 shows that our sample matches the general population of the US well on a number of key characteristics such as age, sex, race, household income, employment status and geographic location. At the same time, Hispanics and people with very low education turn out to be underrepresented in our sample. Also, liberals tend to be overrepresented compared to the US population at large: 51.5% of subjects identify as liberal or very liberal on economic policy issues, 24.2% identify as moderate, and 24.3% as conservative or very conservative.[9] These types of imbalances are however typical in representative online samples that are not based on random sampling (Stantcheva, 2023). Given that our sample covers such a diverse set of subgroups of the US population, it should provide sufficient heterogeneity in fairness views and policy preferences to answer the research question at hand. The imbalances compared to the US population should however be taken into account when generalizing from our study to the US population at large.

2.3.3 Data Quality

Prolific is known to have high data quality in comparison to other online platforms and even in comparison to standard student subject pools (Douglas et al., 2023).[10] In addition to the measures implemented by Prolific, we took several measures to ensure a high quality of our submissions: In Wave 1 and in the new sample in Wave 2, access to our survey was restricted to participants using a laptop or desktop computer.[11] To ensure that all subjects live within the United States, we used a screening protocol provided by Winter et al. (2019) that screens out users who try to hide their geographic location using VPNs. For the main socio-demographics, we double-checked the self-reported data with administrative data provided by Prolific. For our panel study (N=499), we validated that

[8] Two observations in Wave 2 could not be matched to an observation in Wave 1 because they entered a wrong Prolific ID. This reduces the final sample from 501 to 499 observations in our panel.

[9] As a comparison, in the American National Election Survey 2020, the share of people that self-identify as liberal is 30.1%, the share of moderates is 22.0% and the share of conservatives is 33.1% using a similar question. The ANES question (V201200) is measured on a 7-point Likert scale (Extremely Liberal/ Liberal/ Rather Liberal/ Moderate/ Rather Conservatives/ Conservative/ Extremely Conservatives) and does have an additional response category "haven't thought much about this" which is chosen by the remaining 14.5% of participants.

[10] Confirming these findings, we ran two small pilot studies on Prolific and Amazon MTurk (via Cloud Research) prior to Wave 1 in which the effort and engagement of participants with our writing task proved to be much higher on Prolific.

[11] In Wave 2 of our panel, we allowed the use of tablets to increase the retention rate. Just 4.0% of subjects (N=20) used a tablet.

TABLE 2.1: Sample Characteristics

	Wave 1	**Wave 2 - Panel**	**Wave 2 - New**	**Full Sample**	**US Population 2020**
	(1)	(2)	(3)	(4)	(5)
Socio-demographics					
Female	0.52	0.54	0.52	0.51	0.51
Median Age	40	45	45	43	38
Race					
White	0.75	0.75	0.75	0.75	0.75
Black	0.13	0.14	0.12	0.13	0.14
Asian	0.07	0.06	0.06	0.06	0.07
Race Other	0.05	0.05	0.06	0.05	0.04
Ethnicity					
Hispanic	0.07	0.07	0.09	0.08	0.18
Education					
High School/GED or lower	0.34	0.33	0.26	0.31	0.58
College Degree	0.47	0.50	0.42	0.46	0.29
Graduate Degree	0.19	0.18	0.32	0.23	0.12
Economic Background					
Income					
Less than $49,999	0.42	0.43	0.32	0.38	0.39
$50,000 to $74,999	0.20	0.20	0.21	0.21	0.17
$75,000 to $99,999	0.14	0.14	0.16	0.15	0.13
$100,000 - $150,000	0.14	0.14	0.20	0.16	0.16
$150,000 or more	0.09	0.09	0.11	0.10	0.15
Employment Status					
Employed	66.3	71.1	65.3	67.1	72.9
Unemployed	12.3	8.4	10.2	10.5	5.0
Not in Labor Force	21.3	20.6	24.6	22.3	22.1
Census Regions					
Northeast	21.3	19.8	20.4	20.6	17.4
Midwest	18.3	18.2	20.4	19.0	20.8
South	45.1	48.3	38.7	43.5	38.1
West	15.3	13.8	20.4	16.8	23.7
Observations	N=745	N=501	N=729	N=1975	-

Notes: US population estimates are provided by the US Census Bureau (https://data.census.gov/) and are based on the 2020 US Census and the American Community Survey. US population estimates for education are based on the highest level of educational attainment of the population 25 years and over; the share of people not in the labor force is based on the population 20 years to 64 years. Income brackets in our sample were combined to match the US Census data.

participants entered consistent main socio-demographics in both waves. Both comparisons show that subjects enter this data with very high consistency.[12]

2.3.4 Survey Design

Procedures The surveys from all data collections are provided in the Online Appendix A.2.2. This paragraph gives an overview of the main procedures. Subjects first enter their socio-demographic characteristics. Then, some subjects are randomly exposed to one of our treatment manipulations (we discuss the experiment in detail in Section 2.6). Next, all subjects report their current psychological state. In the main part of the survey, we elicit the main variables of our analysis in the following order: (1) beliefs about the causes of inequality, (2) fairness preferences in a spectator game, (3) support for welfare policies. We also elicit a transfer choice with self-interest measured in a modified dictator game.[13] The survey concludes with questions about subjects' exposure to the pandemic in the health and financial domain. At the very end, subjects can provide feedback.

The experiment was implemented in the survey software Qualtrics. In Wave 1, participants received a flat payment of $1.40 plus a bonus payment between $0 and $1.20 depending on their choice with self-interest. In Wave 2, participants additionally received a surprise bonus payment of $1.00 and $0.50, respectively.[14] The median time to complete the study was 12.5 minutes in Wave 1 and 10.5 minutes in Wave 2.

Differences between Wave 1 and Wave 2 The surveys in both waves follow the procedures described above and differ only in minor aspects. In Wave 2, we implement the following main changes: We include a set of additional control variables at the end of the survey, including trust in government, liberal vs conservative ideology on economic policy issues, and voting choices in the 2016 and 2020 elections. We also adapt the health and financial exposure questions to measure personal experiences between the two waves. We further drop a third transfer choice in the spectator game that featured ambiguity about the inequality-generating process to keep the survey in Wave 2 as short as possible.[15]

[12] In the entire sample, age matches the data provided by Prolific in 98.2% of cases (+/- 2 years as margin of error), gender in 99.0% of cases and race/ethnicity in 95.9% of cases. In our panel study, age is entered consistently across waves in 99.0% of cases (+/- 2 years as margin of error), gender in 99.8% of cases, and race/ethnicity in 98.0% of cases.

[13] Participants are matched with one other participant. They are told that one of them will receive a bonus of $1.20. Who gets the bonus is determined by a lottery. Subjects are then asked how much they want to give to the other participant in case that they win the bonus.

[14] The bonus was announced after the part that replicated Wave 1. The additional bonus was paid to nudge high attention and honest reporting in the exposure measures based on a "gift exchange" motive. The size of the bonus differs due to different amount of additional questions in the two surveys.

[15] In Wave 1, we elicit the third choice at the end of the spectator game. The ambiguity rule is always displayed last (after the luck and merit rule), so omitting it in Wave 2 should not confound the comparison of fairness preferences between Waves 1 and 2 in our panel. In Wave 1, most subjects simply redistribute the average of their transfers on luck and merit in the ambiguity rule, which is why we dropped the choice to keep the survey in Wave 2 as short as possible.

2.3.5 Measurement of Outcome and Control Variables

Support for Welfare Policies We measure subjects' support for a diverse set of policies. Subjects are first asked how much they approve of four specific policies contained in the pandemic support package of the US government (Economic Impact Payments, the increase and expansion of unemployment benefits, the expansion of Medicaid, and paid sick leave). We then also ask for their general approval of economic redistribution in society, to capture preferences for the overall level of taxes and transfers in society. We also measure approval of universal health care, which has arguably been one of the most controversial policy proposals to reform the US welfare state in the past decade. All policy preferences are elicited on a 5-point Likert scale (1 "strongly disapprove" to 5 "strongly approve").

To reduce the dimensionality of our analysis, we present graphical illustrations for "*Support for Welfare Policies*" which is the first principal component of all policy preferences (see Appendix A.1.1.7 for details). In our main regression analysis, we present results for the separate policy preferences. For that purpose, the four pandemic policies are aggregated to a "*Pandemic Support*" index by taking a simple average of the four subitems.[16]

Socio-Demographics As our standard socio-demographic characteristics we measure age in years, gender, race, Hispanic ethnicity, and level of education. We further measure subjects' income and current employment status. Income is measured as gross household income in the previous calendar year using seven income brackets [<$20,000, $20,000-$34,999, $35,000-$49,999, $50,000-$74,999, $75,000-$99,999, $100,000-$149,999, >$150,000]. We ask for more detailed income and employment data as part of measuring exposure to the pandemic. We also collect data on subjects' place of residence at the US state and US county level.

Political Ideology Political ideology is self-reported on a scale from 1 to 10, where 1 refers to left/liberal and 10 to right/conservative. In Wave 2, we additionally ask for a self-classification of political orientation on economic policy issues on a 5-point scale [Very Liberal/ Liberal/ Moderate/ Conservative/ Very Conservative], and for voting behavior in the 2016 and 2020 presidential elections.

2.4 Measuring Fairness Views: Preferences and Beliefs

In this section, we describe and graphically illustrate how we classify subjects into distinct fairness preference types based on their transfer choices in the spectator game (2.4.1)

[16]Using the first principal component produces almost identical results, as it assigns similar weights to all four policy preferences.

and how we measure beliefs about the causes of inequality (2.4.2). We then discuss the correlation between fairness preferences and beliefs in our data set (2.4.3).

2.4.1 Measuring Fairness Preferences

Spectator Game Fairness preferences are revealed through transfer choices in a spectator game, in which we implement a new within-subjects design.[17] In the spectator game, subjects can transfer earnings between two workers as an anonymous and impartial spectator. The workers were recruited in a different sample via Amazon MTurk prior to the data collection and participated in a tedious real effort task.[18] The workers earn a fixed show-up fee of $0.50 plus a variable compensation for working on the real effort task ($0 to $4), which depends on the transfer choices of the spectator.

Each of the spectators chooses a transfer for two different states of the world. In both states of the world, there is the same level of inequality: one worker is initially endowed with $4 and the other one with $0 as their variable compensation for the real effort task. The causes of inequality, however, differ between the two states. The causes of inequality are either:

- **Luck:** A lottery determines who receives 4$.
- **Merit:** The more productive worker receives 4$.

Spectators can decide to transfer between $0 to $4 (in steps of 10 cents) to the worker with $0.[19] In a classic between-subjects design, one group of subjects would choose a transfer on inequalities due to luck, and another group on inequalities due to merit. In our **within-subjects design**, each subject takes two transfer choices, one for each cause of inequality, in random order. Spectators are not aware of the second cause of inequality when they make their first transfer choice, they only know that a second choice is pending. We informed spectators that for every fourth spectator, one of their choices (chosen randomly) would be implemented and would determine the payoffs for a pair of workers. The decisions of spectators are hence probabilistically incentivized to limit costs of data collection (see Andre (2024) and Bartling et al. (2023) for a similar approach).[20]

[17] The instructions for subjects closely follow the wording used in Almås et al. (2020), except that we make the necessary changes to implement the within-subjects design.

[18] The task consisted of counting the number 1 in a line of symbols. Their productivity was measured as the number of correctly solved lines. Spectators are, however, not aware of these details about the real-effort task.

[19] Almås et al. (2020), instead, offer a choice between seven allocations [($0,$6), ($1,$5), ($2,$4), ($3,$3), ($4,$2), ($5,$1), ($6,$0)]. The motivation for our design is that we did not want to limit the heterogeneity in fairness preferences that could be observed by restricting the choice set, and also to potentially observe more precise variation in fairness preferences in our experiment and panel data.

[20] In total we recruited N=988 workers via MTurk. Shortly after the data was collected, spectators were matched to worker pairs and the choices of spectators were implemented by paying out the earnings after transfers to the MTurk workers.

Link to Theory In the spectator game, subjects can implement a transfer between the two workers as a social planner without any confounding effect of economic self-interest. Implementing a transfer between the workers is not associated with any costs or benefits for the spectator: $y_i'(\tau) = 0$, $T_i'(\tau) = 0$. The chosen "tax rate" in the spectator game (a transfer of $2 would be equivalent to a tax rate of 100%) should therefore only be determined by what a subject thinks is fair given beliefs about the causes of inequality:

$$\tau_i^* = F_i(b_i) \tag{2.4}$$

Subjects make two transfer choices in the spectator game in which we exogenously fix beliefs, b_i, by providing information about the causes of inequality. This arguably allows us to identify the fairness preferences on inequalities due to luck, $F_i(b_i = Luck)$, and on inequalities due to merit, $F_i(b_i = Merit)$[21], for each individual in our sample.

Classification of Fairness Preference Types The within-subjects design allows us to classify each individual in our sample into egalitarian, libertarian or meritocratic fairness preferences based on their two transfer choices in the spectator game ($\mathrm{T_{Luck}}$ and $\mathrm{T_{Merit}}$ hereafter). The classification is as follows: Egalitarians equalize all inequalities in both transfer choices: $\mathrm{T_{Luck}} = \mathrm{T_{Merit}} = \$2$. Libertarians, on the other hand, accept both types of inequality: $\mathrm{T_{Luck}} = \mathrm{T_{Merit}} = \$0$. Meritocrats equalize inequalities due to luck but accept inequality due to merit: $\mathrm{T_{Luck}} = \$2 > \mathrm{T_{Merit}}$. In a between-subjects design, these types can only be estimated at the population level under a number of assumptions about the consistency of choices across treatments.

Heterogeneity in Fairness Preference Types Figure 2.1 shows a jittered scatter plot of subjects' transfers from the rich worker ($4) to the poor worker ($0) when the inequality is caused by merit (x-axis) and when it is caused by luck (y-axis). Each dot represents one observation in our sample. Egalitarian fairness preferences (red dots) are implemented by 11.9% of our subjects (N=235), while Libertarian fairness preferences (yellow dots) are implemented by 5.3% of our subjects (N=104). Meritocratic fairness preferences (blue dots) are implemented by 49.7% of our subjects (N=982). Around 33.1% of our sample cannot be classified into one of these three distinct fairness preference types. They are hence classified as "Other" (grey dots).

[21] We are aware that the concept of merit is not clearly defined in this situation as the abilities needed to be productive in the task might themselves be a matter of luck rather than of effort. However, the debate on whether ability should be seen as a factor within or partly out of the control of the individual is a philosophical question that is beyond the scope of our paper and likely also not resolved in the minds of our respondents. The situation we present to our respondents thus reflects the uncertainties that surround productivity and merit in any real world situation, in line with the existing literature.

FIGURE 2.1: Transfer Choices and Fairness Preference Types

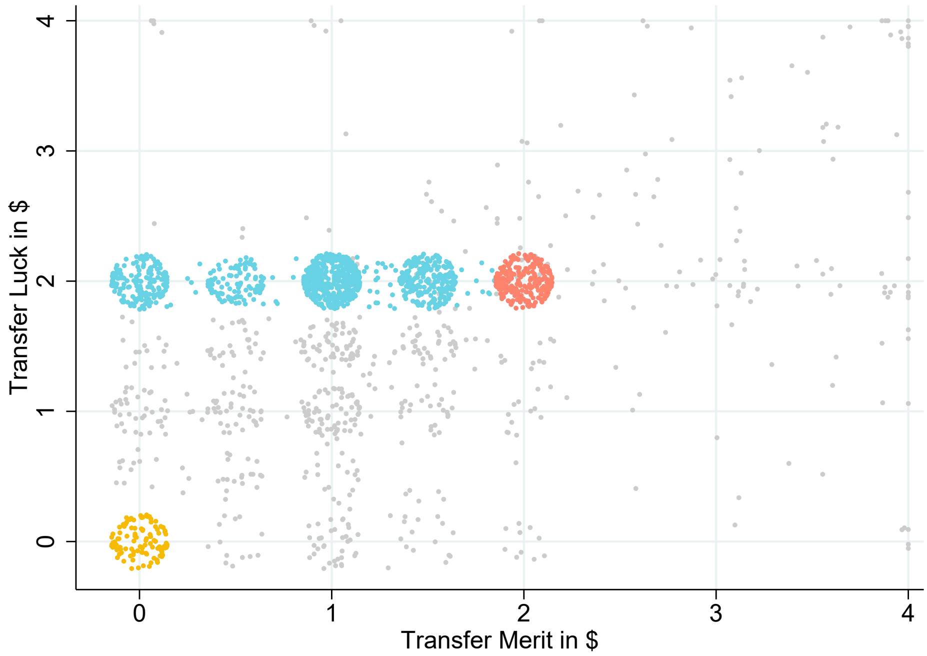

	Fairness Preference Type	Classification	N	Share
●	Egalitarians	$T_{\text{Luck}} = T_{\text{Merit}} = \$2$	235	11.9%
●	Libertarians	$T_{\text{Luck}} = T_{\text{Merit}} = \$0$	104	5.3%
●	Meritocrats	$T_{\text{Luck}} = \$2 > T_{\text{Merit}}$	982	49.7%
●	Other	-	654	33.1%

Notes: The figure shows the distribution of transfer choices on inequalities due to luck and merit in the spectator game. Each dot represents one observation in our sample (N=1975). The colors of the dot indicate to which fairness preference type an individual is assigned where red stands for the egalitarian, yellow for the libertarian, and blue for the meritocratic fairness preference type. Dots are jittered, so that clusters of individuals are visible.

Are there more types? The case of strong and weak meritocrats The most apparent novel heterogeneity in fairness preferences revealed by our within-subjects design is heterogeneity among the meritocrats. While the meritocratic fairness ideal specifies that inequalities caused by luck are unfair, it does not specify how much inequality is fair when inequalities are caused by merit (Roemer and Trannoy, 2015). Heterogeneity along this dimension among the meritocrats is illustrated in Figure 2.1 by the four clusters of meritocrats around ($0,$2), ($0.5,$2), ($1,$2) and ($1.5,$2) that all fully equalize inequalities due to luck but accept inequalities due to merit to varying degrees. Hence, meritocrats differ in the strength of their meritocratic preferences, which we define as the difference between transfers on luck and merit: Meritocratic Preferences = T_{Luck} - T_{Merit}. The four clusters of meritocrats in Figure 2.1 can be classified as having (from left to right) very strong, rather strong, moderate, and weak meritocratic preferences (also see Figure A.3).

What about heterogeneity among subjects classified as having "other" fairness preferences? Among the "other" type, subjects do not seem to strongly cluster around one point in Figure 2.1: the largest cluster is at ($1,$1) with around 3.2% of subjects in our sample. While some of the choices among the "other" type may reflect measurement error, most choices, for example of the clusters around ($1,$1) or ($1.5,$1.5), may very plausibly reflect subject's genuine fairness preferences in the spectator game.

In order to speak to the existing literature, we focus our main analysis on a comparison of the egalitarian, libertarian, and meritocratic fairness preference types. To strike a balance between parsimony and richness of the type classification, we extend the analysis to the meritocratic subtypes where it provides insightful contributions. Our analysis of different fairness preference types therefore focuses on those six clusters of subjects that contain 5% or more of subjects in our sample. All of these empirically relevant fairness preference types are consistent with one fairness ideal derived from influential theories of distributive justice, namely, egalitarianism, libertarianism, or meritocracy.

Validating the Type Classification To validate our new individual-level classification of fairness preference types, we present four pieces of evidence. First, we estimate the type distribution with the between-subject approach of Almås et al. (2020) using just the first transfer choice of each subject (see Appendix Table A.1). The within and between type distributions do not differ by much, except that the shares of egalitarians and libertarians are slightly higher using the between-subjects approach (by 3-4pp). This is a "mechanical" effect if subjects violate the assumptions that are made about the consistency of their choices in the between-subjects approach. Second, while we observe some order effects in our data, they do not cause substantial changes in the type distributions, neither in the between- nor in the within-classification (see Appendix Table A.1). Third, if we classify subjects who are close to ($0,$0) as libertarians and those close to ($2,$2) as egalitarians (diff$<$ $0.25), the shares increase only marginally, by 1 and 15 observations respectively.[22] Last, the shares of egalitarians and meritocrats in our sample are very similar to the estimated shares in representative samples of the US in Almås et al. (2020) and Cohn et al. (2023), while the share of libertarians is lower (see Appendix A.2 for a detailed comparison).

2.4.2 Measuring Beliefs

We measure beliefs with two established questions from the World Values Survey about the relative importance of factors within individual control (e.g. merit) and factors outside

[22]Given their small numbers, classifying these observations as libertarians or egalitarians does not meaningfully change any of our results.

individual control (e.g. luck) in causing economic success and poverty on 10-point Likert Scales (see Figure 2.2).[23] The first question measures beliefs about the causes of success in US society, namely, whether success is rather caused by hard work or by luck and connections. The second question measures beliefs about the causes of poverty in a similar way (which we slightly modified by including "bad luck" as a reason for poverty).

FIGURE 2.2: Beliefs about the Causes of Inequality

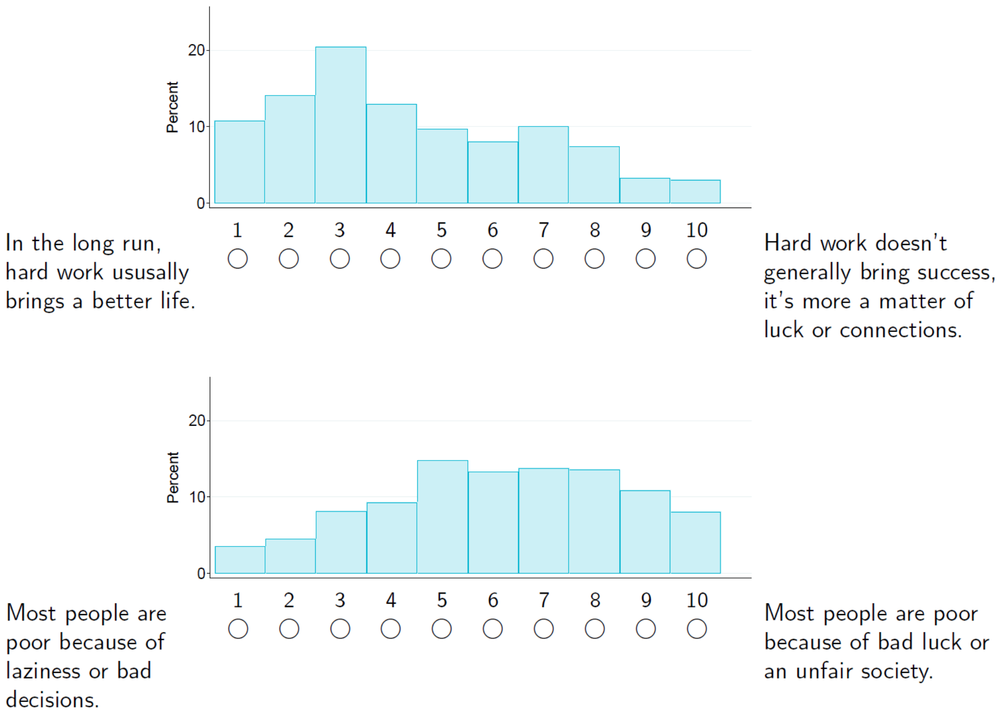

Notes: The figure shows the distributions of answers to our survey questions about the causes of inequality in US society. Subjects are asked whether they rather agree with the statement on the left or on the right and to place their views on the scale from 1 to 10 accordingly.

On average, we find that subjects believe strongly that hard work leads to economic success (58.3% state 4 or lower, mean=4.4, see Figure 2.2), in line with ample evidence from previous studies using the same or similar questions in the US (Alesina and Giuliano, 2011). At the same time, most subjects believe that poverty is often due to factors outside individual control such as bad luck (only 22.0% state 4 or lower, mean=6.2, see Figure 2.2). This shows that people do not seem to believe that economic success and poverty are caused to the same degree by factors within and outside individual control, which highlights the importance of measuring beliefs about both processes, in line with the recent findings by Fong and Poutvaara (2019).

[23]Note that beliefs about the causes of inequality should be conceptually distinguished from beliefs about the degree of social mobility, studied for example in Alesina et al. (2018), which measure how much people move between societal strata. While these two beliefs may be related, beliefs about the causes of inequality measure why people move between societal strata, which arguably makes it a better measure of how meritocratic society is. For example, a society with a very high degree of social mobility may not be meritocratic if positions in society are determined randomly.

To construct the variable "*Beliefs in Merit*", we reverse code the items and take a simple average of the answers to these two questions. Note that higher beliefs correspond to a higher belief in factors within individual control, that is, a belief that US society is more meritocratic.

2.4.3 Correlation between Preferences and Beliefs

Given that few papers have jointly measured fairness preferences and beliefs about the causes of inequality, we know very little about the empirical relationship between these two measures. A priori, it seems likely that these two measures are correlated, given that people may form their beliefs about the causes of inequality based on their normative ideas about the fairness of inequality, or vice versa.

When looking at the correlation in our sample, we find that egalitarians, libertarians, and meritocrats not only hold very different normative views about fairness but also on average have different beliefs about the causes of inequalities in US society: the average belief in a meritocratic society is highest among libertarians (6.4), followed by meritocrats (5.8) and lowest among egalitarians (4.7) (see Appendix Figure A.2 for histograms of beliefs by type). While these mean differences are statistically significant ($p < 0.001$ for all pairwise comparisons), there is a lot of unexplained variation in beliefs among the types: only 3.8% percent of the variation in beliefs can be attributed to subjects holding different fairness preferences. Hence, the two elements determining fairness views in our theoretical framework - preferences and beliefs - seem to have sufficient independent variation.

2.5 How Fairness Shapes Support for Welfare Policies

This section studies whether fairness preferences and beliefs can explain why some people strongly support welfare policies, while others strongly oppose welfare policies. For this purpose, we test central predictions from our theoretical framework in a cross-sectional analysis. First, we test whether the fairness preference types - revealed through two simple transfer choices in our spectator game - differ in their support for welfare policies (2.5.1). Second, we test how beliefs about the causes of inequality shape the policy support of meritocrats (2.5.2). Third, we test whether beliefs about the causes of inequality have a larger effect on the policy support of meritocrats than on the other types (2.5.3).

2.5.1 Fairness Preferences Predict Policy Preferences

Based on our theoretical framework and the preferred transfer levels revealed in the spectator game, we derive the first theoretical prediction:

Prediction 1: Egalitarians should have a higher support for welfare policies than libertarians. This follows from $F_{Ega} = 1 > F_{Lib} = 0$ for all beliefs.

To test this prediction, we present a set of simple OLS regressions which regress policy preferences on dummies for each fairness preference type while controlling for a number of observable characteristics. A "controlling for observables" strategy is the standard approach in the literature studying the relationship between experimental measures of distributive preferences and policy preferences (Enke et al., 2023; Epper et al., 2020; Kerschbamer and Müller, 2020).[24] We include all N=1975 observations in our regression analysis to maximize the power and precision of estimates. Standard errors in all regressions are clustered at the individual level. Figure 2.3 depicts the estimated coefficients of the fairness preference type dummies from an OLS regression that just controls for socio-demographics and economic background characteristics. The outcome variable is support for welfare policies, which is the standardized first principal component of all our policy preferences. The reference category in this regression are subjects with "other" fairness preferences.

Figure 2.3 (a) reveals that egalitarians have the highest support for welfare policies, while libertarians have the strongest opposition against welfare policies. The policy preferences of meritocrats lie in between the attitudes of egalitarians and libertarians. Hence, the rank ordering of support for welfare policies of the fairness preference types is in line with the prediction. The sizes of the estimated coefficients are such that libertarians are on average 0.80 SD less in support of welfare policies compared to egalitarians ($p<0.001$). Meritocrats are on average 0.41 SD more in support of welfare policies than libertarians ($p<0.001$) and 0.39 SD less in support of welfare policies than egalitarians ($p<0.001$).

The estimated coefficients are not only statistically significant but also remarkably large compared to relevant benchmarks: for example, the mean difference in support for welfare policies between subjects in the lowest and the highest income bracket (<\$20,000 vs >\$150,000) is 0.46 SD, between the full-time employed and unemployed it is 0.35 SD (see Appendix A.7 for a comparison).[25]

In Table 2.2 we show that differences in support for welfare policies between the fairness preference types remain large and significant even when we control for people's beliefs about the causes of inequality as well as their left-right political ideology. Controlling for left-right political ideology is seen as a critical robustness check in Enke et al. (2023) and Epper et al. (2020). Note, however, that people's left-right political ideology could itself

[24] This is also more generally the case for the literature on the external validity of experimental measures of economic preferences, such as risk or time preferences (e.g. Charness et al., 2020; Schneider and Sutter, 2020).

[25] The differences in support for redistributive policies between income brackets in our sample are remarkably similar to those in Epper et al. (2020).

FIGURE 2.3: Fairness Preference Types Predict Support for Welfare Policies

(A) Fairness Preferences (B) Beliefs among Meritocrats

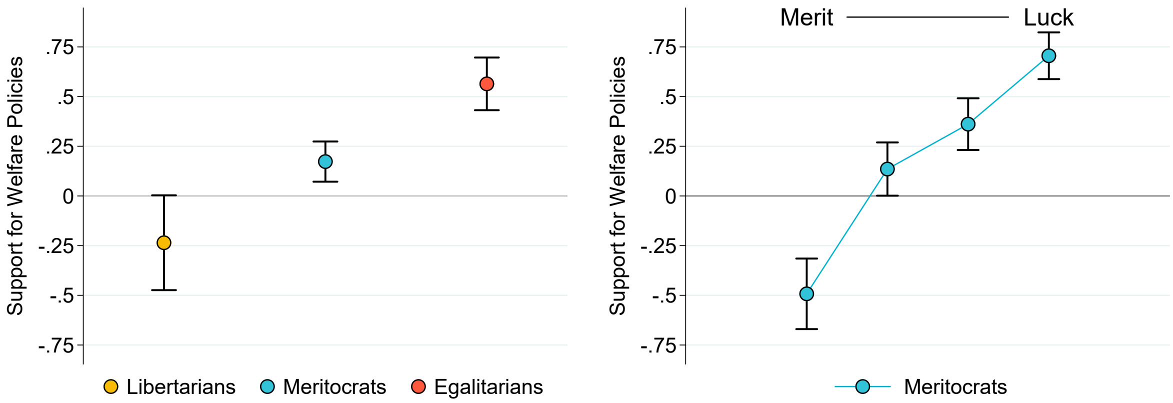

Notes: The figure shows estimated coefficients and 95% confidence intervals from OLS regressions explaining support for welfare policies (standardized first principal component of all policy preferences). Figure (a) shows coefficients of the egalitarian, libertarian, and meritocratic fairness preference type dummies (see Section 2.4 for details on the classification). In Figure (b) meritocrats are classified into four subtypes using quartiles of the belief distribution from high beliefs in merit (left coefficient) to high beliefs in luck (right coefficient) (see Appendix A.1.1.6 for details on the classification). Regressions control for socio-demographics (age, gender, race/ethnicity, education) and economic background characteristics (income bracket and employment status). Subjects classified as "other" serve as the reference category in both regressions, corresponding to the grey horizontal line at 0. Robust standard errors are clustered at the individual level.

be shaped by people's fairness preferences, so that the coefficients of fairness preferences may be biased towards zero once we control for political ideology.[26]

Panel A in Table 2.2 presents six OLS regressions that explain support for welfare policies in our three different policy domains: redistribution, universal health care, and the pandemic support package. In columns (1), (3) and (5) we control for socio-demographics, economic background characteristics and beliefs about the causes of inequality, while in columns (2), (4) and (6) we additionally control for left-right political ideology. In all models, the egalitarian type serves as the reference category. All regressions confirm the pattern depicted in Figure 2.3. When political ideology is included in the regression, the estimated coefficients of fairness preference types decrease by about a third but remain large and jointly significant. When comparing the sizes of the *Libertarian* coefficient across policy domains, we observe that it is larger for economic redistribution than in the other two domains, but statistically significant in all domains and specifications ($p < 0.015$). Based on these analyses, we derive our first result confirming Prediction 1:

[26]Research in political science has shown that left-right political ideology aggregates a wide variety of beliefs, preferences and behavioral motives that shape people's general political attitudes, especially their attitudes towards inequality (Jost et al., 2009). Controlling for political ideology arguably also allows to rule out that people's behavior in our experimental game is just the result of their political ideology. If this was the case, then we should observe no correlation between fairness preferences and policy preferences once we control for political ideology.

Result 1: Egalitarians have a higher support for welfare policies than libertarians, and meritocrats lie in between these two extreme types.

Robustness One may still be concerned about other potential confounders that are not captured by people's socio-demographics, economic background characteristics, beliefs about the causes of inequality, or their left-right political ideology. Epper et al. (2020), for example, show that selfish subjects (14.8% in their sample) are around 0.3 SD less in support of redistributive policies than altruistic and inequality-averse subjects. To identify selfish subjects in our sample, we can make use of our transfer choice with self-interest measured in a slightly adapted version of the classical dictator game. In line with Epper et al. (2020), we find that selfish subjects (13.4% in our sample who do not give any money to another participant) are less in support of welfare policies (0.28 SD), but importantly, differences between the fairness preference types remain large and statistically significant once we control for selfishness (see Appendix Table A.8).

In Appendix Table A.8, we also show that Result 1 is robust if we additionally control

Table 2.2: OLS: Fairness Preferences Predict Policy Preferences

	Redistribution		Univ. Health Care		Pandemic Support	
	(1)	(2)	(3)	(4)	(5)	(6)
Libertarian	-0.63***	-0.47***	-0.49***	-0.29**	-0.40***	-0.25*
	(0.11)	(0.10)	(0.12)	(0.10)	(0.11)	(0.10)
Meritocrat	-0.29***	-0.25***	-0.14*	-0.08	-0.20***	-0.15**
	(0.06)	(0.05)	(0.06)	(0.05)	(0.06)	(0.05)
Other	-0.30***	-0.18**	-0.20**	-0.05	-0.33***	-0.22***
	(0.07)	(0.06)	(0.07)	(0.06)	(0.07)	(0.06)
Beliefs in Merit	-0.19***	-0.11***	-0.17***	-0.08***	-0.15***	-0.08***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Left-Right Political Ideology		-0.15***		-0.18***		-0.14***
		(0.01)		(0.01)		(0.01)
p (Egalitarian = Libertarian)	< 0.001	< 0.001	< 0.001	0.005	< 0.001	0.015
p (Egalitarian = Meritocrat)	< 0.001	< 0.001	0.026	0.116	0.001	0.003
p (Meritocrat = Libertarian)	0.001	0.009	0.002	0.025	0.046	0.293
p (joint test)	< 0.001	< 0.001	< 0.001	0.017	< 0.001	0.004
Socio-Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Income Bracket	Yes	Yes	Yes	Yes	Yes	Yes
Employment Status	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1975	1975	1975	1975	1975	1975
R^2	0.265	0.388	0.186	0.378	0.167	0.280

Notes: OLS estimates with robust standard errors clustered at the individual level in parentheses. Dependent variables: policy preferences as standardized z-scores. Omitted category: "Egalitarians". Socio-Demographics include age, gender, race, ethnicity and education dummies. The "joint test" tests the hypothesis that there is no difference between any of the three fairness preference types (egalitarians, libertarians and meritocrats). *** p<0.001, ** p<0.01, * p<0.05

for trust in government, for people's perceived closeness to other people in the US, for liberal-conservative political ideology on economic policy issues (measured on a 5-point Likert scale) and for voting behavior in the 2016 and 2020 presidential elections.

Last, we show that there are large and robust differences in support for welfare policies between meritocrats with strong and weak meritocratic preferences in line with theoretical predictions (see Appendix Figure A.6 and Table A.9).

2.5.2 Beliefs Predict Policy Preferences Among Meritocrats

Next, we turn to the role of beliefs about the causes of inequality in US society for the policy preferences of meritocrats:

> **Prediction 2:** Meritocrats who believe that inequalities in US society are caused by luck should be more in support of welfare policies than meritocrats who believe that inequalities are caused by merit. This prediction follows from $F'_{Mer}(b_i) > 0$ by definition of the meritocratic type.

Figure 2.3 (b) depicts coefficients from an OLS regression that is identical to the one for Figure 2.3 (a), except that meritocrats are split into four different subtypes according to quartiles of the belief distribution, ranging from strong beliefs in merit (left coefficient) to strong beliefs in luck (right coefficient). Only the coefficients of the meritocratic subtypes are depicted, again, relative to subjects with "other" fairness preferences. Figure 2.3 (b) shows that the higher the belief in merit as a cause of inequality, the lower the support for welfare policies among meritocrats ($p<0.006$ for all pairwise comparisons between meritocratic subtypes). The finding that beliefs about the causes of inequality predict support for redistributive policies is in line with an entire strand of literature (e.g. Alesina and Giuliano, 2011; Alesina and La Ferrara, 2005; Fong, 2001) and therefore unsurprising. What is remarkable, however, when comparing Figures 2.3 (a) and (b), are the relative sizes of the coefficients of fairness preferences and beliefs:

> **Result 2:** Meritocrats who believe that inequalities in US society are mainly caused by luck are, on average, as supportive of welfare policies as egalitarians, while meritocrats who believe that inequalities in US society are mainly caused by merit are, on average, at least as opposed to welfare policies as libertarians.

Therefore, differences in support for welfare policies between individuals with opposing fairness preferences (libertarians and egalitarians) are comparable in size to differences between meritocrats with opposing beliefs (merit vs luck). In terms of their estimated coefficient sizes, fairness preferences and beliefs about the causes of inequality hence prove to be similarly important for explaining disagreements about welfare policies at the individual level.

2.5.3 Fairness Preferences and Beliefs Interact

Last, we test the central theoretical prediction that fairness preferences determine to what extent beliefs about the causes of inequality matter for fairness views - and thus for people's policy preferences. This is clearly the most demanding test of our theoretical framework and of our individual-level classification of fairness preference types.

> **Prediction 3:** The policy preferences of meritocrats should depend more strongly on their beliefs about the causes of inequality than the policy preferences of egalitarians, libertarians and others. This prediction follows from: $F'_{Mer}(b_i) > 0$ by definition of the meritocratic type, $F'_{Ega}(b_i) = F'_{Lib}(b_i) = 0$ by definition of the egalitarian and libertarian type, and the observation that the mean $F'_{Oth}(b_i) \approx 0$ among "Others"[27].

Table 2.3 establishes that there are, in fact, meaningful interaction effects between fairness preferences and beliefs. In Table 2.3 we present OLS estimates that are identical to Table 2.2, except that we now interact the variable "Beliefs in Merit" with dummies for each fairness preference type (meritocrats serve as the omitted category). The "Beliefs in Merit" coefficient in the first row indicates the association between beliefs and policy preferences among meritocrats: this association is strong and significant across policies, confirming the result presented in Figure 2.3 (b).

The interaction terms in Table 2.3 show that the estimated effect of beliefs on policy preferences is substantially weaker for egalitarians and individuals with "other" fairness preferences than for meritocrats. The interaction effect is statistically significant across columns (1) to (6) for egalitarians ($p< 0.006$) and for individuals with "other" fairness preferences ($p< 0.024$). In column (2), for example, a one-point increase in beliefs in merit is associated with a 0.15 SD decrease in support for redistribution among meritocrats ($p<0.001$), but only with half of that effect among egalitarians (0.07 SD, $p=0.002$) and "others" (0.07 SD, $p<0.001$). In most specifications, the estimated effects of beliefs on the policy preferences of egalitarians and individuals with "other" fairness preferences are not zero but instead tend to be negative and significant at conventional levels, though much smaller in size than for meritocrats. At the same time, the estimated interaction effects for libertarians are not consistent with our prediction: the estimated effect of beliefs is not significantly weaker for libertarians than for meritocrats in any specification.[28]

While our results are in line with Prediction 3 in 94.7% of our sample, the absence of an interaction effect for libertarians still warrants conducting additional empirical tests for an

[27] On average, subjects with "Other" fairness preferences transfer \$1.33 when inequalities are caused by luck and \$1.39 when caused by merit. Thus, the mean difference between T_{Luck} and T_{Merit} among "Others" (\$0.06) is much smaller than among meritocrats (\$1.10).

[28] This finding may however be due to a lack of power because we have the lowest number of observations for libertarians (just 5.3% of our sample, N=104) so the coefficient of beliefs is estimated less precisely for libertarians compared to the other fairness preference types.

TABLE 2.3: OLS: Interaction between Fairness Preferences and Beliefs

	Redistribution		Univ. Health Care		Pandemic Support	
	(1)	(2)	(3)	(4)	(5)	(6)
Beliefs in Merit	-0.24***	-0.15***	-0.22***	-0.11***	-0.19***	-0.11***
	(0.02)	(0.01)	(0.02)	(0.01)	(0.02)	(0.02)
Egalitarian	-0.25*	-0.18	-0.45***	-0.37***	-0.22	-0.15
	(0.12)	(0.11)	(0.11)	(0.10)	(0.12)	(0.11)
Egalitarian × Beliefs in Merit	0.10***	0.08***	0.12***	0.09***	0.08**	0.06**
	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)
Libertarian	0.10	0.10	0.09	0.08	0.28	0.28
	(0.25)	(0.22)	(0.26)	(0.23)	(0.26)	(0.24)
Libertarian × Beliefs in Merit	-0.06	-0.05	-0.06	-0.04	-0.07	-0.06
	(0.04)	(0.03)	(0.05)	(0.04)	(0.04)	(0.04)
Other	-0.74***	-0.44***	-0.70***	-0.32**	-0.66***	-0.37**
	(0.13)	(0.13)	(0.13)	(0.12)	(0.13)	(0.13)
Other × Beliefs in Merit	0.12***	0.09***	0.11***	0.06**	0.09***	0.05*
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Political Ideology	No	Yes	No	Yes	No	Yes
Socio-Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Income Bracket	Yes	Yes	Yes	Yes	Yes	Yes
Employment Status	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1975	1975	1975	1975	1975	1975
R^2	0.282	0.396	0.202	0.384	0.179	0.285

Notes: OLS estimates with robust standard errors clustered at the individual level in parentheses. Dependent variables: policy preferences as standardized z-scores. Omitted category: "Meritocrats". Socio-Demographics include age, gender, race, ethnicity and education dummies.
*** p<0.001, ** p<0.01, * p<0.05

interaction effect between fairness preferences and beliefs. First, we confirm that beliefs have a stronger effect on the policy preferences of meritocrats than of non-meritocrats if we pool all three other types (see Appendix Table A.10). Second, we show that there is also a meaningful interaction between beliefs and the strength of meritocratic preferences among meritocrats (see Appendix Table A.11). In line with what theory would predict, the negative relationship between beliefs and policy preferences proves to be much stronger for meritocrats with strong meritocratic preferences compared to meritocrats with weak meritocratic preferences - across policy domains and specifications. Taken together, our data therefore largely confirm Predictions 3:

> **Result 3:** Fairness preferences and beliefs interact. The policy preferences of *Meritocrats* depend more strongly on beliefs about the causes of inequality than those of *Egalitarians* or *Others*.

Robustness The interaction effect of preferences and beliefs among meritocrats also holds once we drop the type-based approach and use a continuous measure of meritocratic preferences (T_{Luck} - T_{Merit}) to test for the interaction (see Appendix Table A.12 Panel A). Similarly, if we drop the type-based approach in our entire sample and just interact a continuous measure of meritocratic preferences (T_{Luck} - T_{Merit}) with beliefs, the interaction is also statistically significant (see Appendix Table A.12 Panel B). Hence, those subjects whose transfer choices in the spectator game are more sensitive to the causes of inequality, also seem to put more weight on beliefs about the causes of inequality when forming their policy preferences.

2.6 The Stability of Fairness Views in Times of Crises

This section studies the stability of fairness preferences, beliefs about the causes of inequality, and support for welfare policies over time. We thus advance beyond the cross-sectional analysis of the previous section and investigate whether shock to fairness views in times of societal crises can explain shifts in support for welfare policies over time. First, we discuss theoretical mechanisms, our data, and our empirical strategy. We then present our evidence on the stability of fairness preferences (2.6.1) and on the stability of beliefs about the causes of inequality (2.6.2). In the last step, we investigate whether shocks to fairness views can explain shifts in support for welfare policies over time (2.6.3).

Theoretical Mechanisms: Preference Shocks versus Belief Shocks How can societal crises shock people's fairness views? And through which mechanism: through shocks to fairness preferences or through shocks to beliefs about the causes of inequality?

In societal crises, people often experience economic shocks that cause job losses and business closures, or natural catastrophes that threaten people's property and health. These exogenous shocks typically generate new economic inequalities within society that are caused by factors outside individual control. Societal crises also cause personal experiences in which people are not in control of their life outcomes. Such personal experiences have, for example, been omnipresent during the coronavirus pandemic: people were hit by job losses, faced the risk of catching a potentially deadly virus, and their economic freedom and civil liberties were restricted by mandatory business closures and governmental lockdowns.

Witnessing exogenous shocks and learning about new inequalities in society may make people more pessimistic about the idea that merit determines economic outcomes, and thus lead to an update in their beliefs about the causes of inequality. People may also update their beliefs in response to personal experiences that they make in a crisis, such as experiencing uncontrollable changes in life circumstances in both the private and professional domain.[29] However, societal crises could also make people reflect on how much

[29]In a seminal model, Piketty (1995) has formulated the dynamics between personal experiences, beliefs about the economic system, and redistributive policies.

and what type of economic inequalities should be considered fair, and thus change their fairness preferences. Throughout the coronavirus pandemic, people have witnessed lively public debates about inequality and social justice. In these debates, Sandel (2020), for example, has prominently criticized the normative basis of the meritocratic fairness ideal. Even without extensive philosophical reflections, people might implement different fairness preferences once their personal circumstances or the societal context changes (see Barr et al., 2016; Cappelen et al., 2021).[30] Therefore in times of a major societal crisis like the coronavirus pandemic, both belief shocks and preference shocks seem plausible.

Panel and Experimental Data We study the stability of fairness preferences, beliefs about the causes of inequality, and support for welfare policies during the pandemic using (i) individual-level panel data collected over a time span of 1.5 years from May 2020 to Fall 2021, and (ii) an experiment within our first wave of data collection, in which subjects write about personal experiences from the first months of the pandemic.[31]

Our panel data spans a period of time in which Americans experienced the unprecedented impacts of the pandemic on US society including more than 750,000 deaths related to COVID-19 (see Appendix Figure A.13).[32] At the same time, the study period also contains many other disruptive societal events besides the pandemic, for example, the Black Lives Matter protests following the death of George Floyd in May 2020, or the January 6th Capitol Hill riots following the 2020 presidential elections. Our panel data therefore provides evidence from a unique time period in which Americans experienced multiple societal crises and historic events that could potentially change their views about economic inequality and their support for policy interventions.

Our experiment complements the panel data because it studies the effect of a specific type of personal experience: a perceived loss of control over one's life outcomes. Such personal experiences were omnipresent during the pandemic, but are also more generally characteristic of many societal and personal crises. In our experiment in Wave 1, each subject (N=745) is randomly assigned to one of three conditions: In our two treatment conditions, *High Control* and *Low Control*, subjects are asked to write about a personal experience from the first months of the pandemic in which they experienced *a lot of control* (*High Control*) or *no control* (*Low Control*) over their lives. In our *Baseline* condition, participants are asked to write about a personal experience that is not related

[30] In a panel study, Barr et al. (2016) show that individuals who become unemployed are less accepting of inequalities due to merit. Cappelen et al. (2021) find in a priming experiment during the first weeks of the pandemic that US Americans become more tolerant of inequalities due to luck (in a survey question) when experimentally reminded of the pandemic, but their beliefs about the causes of inequality do not change.

[31] We pre-registered our experiment in the AEA RCT Registry (AEARCTR-0005856) (`https://doi.org/10.1257/rct.5856-1.0`)

[32] Note that our study period does not cover the first months of the pandemic (February to April 2020), which saw the strongest negative impacts on the US labor market. Instead, it covers the economic recovery from a historic unemployment rate of more than 13.2% in May 2020 to less than 5% in Fall 2021.

to the pandemic. Participants in both treatment conditions are first also provided with information about the severe impacts of the pandemic on US society (the latest COVID-19 case and death counts, and the latest unemployment data). We thereby hold constant the information that subjects have about the impacts of the pandemic on society and then randomize the recall of different types of past personal experiences.

Participants clearly engaged with the writing task and put in considerable effort.[33] A manipulation check shows that the *Low Control* treatment has the desired effect compared to *High Control* measured through self-reported psychological states: subjects have a lower perceived sense of control over their lives and report emotions like negative affect, fear and stress (see Appendix A.1.4.2).

Regarding the mechanism of changes in fairness views, we are mainly interested in the effect of experiencing a loss of control over one's life outcomes, which may be caused by a broad set of experiences in times of crisis. We pre-registered to test this mechanism in our AEA registry prior to the first wave of data collection. In our individual-level panel data, we also test exploratively whether a number of specific personal experiences are related to changes in fairness views: income loss, job loss, spells of unemployment or cases of COVID using a Diff-in-Diff approach (see Appendix A.1.4.7 for details).[34]

Attrition and Balance of Covariates In our panel, we find that panel attrition is not related to fairness views, policy support or political ideology in Wave 1 (see Appendix A.14). Among the socio-demographics, we only find that young subjects are slightly less likely to participate in Wave 2 (see Appendix A.13). Just 4 subjects that started the survey in Wave 2 did not complete it, hence there was no meaningful attrition within the survey.

In our experiment in Wave 1, there was also no differential attrition (just 6 subjects left the survey after being assigned a treatment status, 2 per treatment condition) and randomization led to a balance of covariates (see Appendix A.13). In Wave 2, in the sample of newly recruited subjects (N=729), we also included an experimental manipulation very similar to Wave 1, in which we unfortunately observed strong differential attrition across treatment conditions.[35] The differential attrition seems to warrant excluding the results

[33]They wrote on average 330 characters in *Low Control*, 365 in *High Control*, and 225 in *Baseline*. A research assistant read and coded all texts that were written as part of our manipulation. All but one text (which was copied from the internet) show that people complied with the task and put in considerable effort.

[34]Of our 499 subjects, 78 report that they or someone emotionally close to them had a severe case of COVID between Wave 1 and Wave 2, and 204 subjects experienced a mild infection with COVID. Besides these widespread health impacts of the pandemic, many individuals experienced substantial changes in their economic circumstances: 89 individuals report to have experienced a sustained loss of household income compared to before the pandemic, and 78 individuals lost a job or main source of income in our study period.

[35]In our experiment in Wave 2, 3 subjects did not complete the survey in *Baseline*, 12 subjects in an *Information* condition, and 42 in *Low Control* (Chi^2-test: $p<0.001$). We can only speculate about the reason for the strong attrition of 16.5% of subjects in *Low Control*. One plausible explanation is that subjects at this later point in time did not want to be reminded of negative personal experiences that they had during the pandemic and hence left the experiment.

from the main part of our paper because we cannot rule out that the results are strongly biased due to an unobserved difference between subjects across treatment conditions. For transparency, the results are presented in Appendix A.1.4.8.

2.6.1 Fairness Preferences are Stable at the Population Level

Panel Data - Population Level Figure 2.4 depicts the distribution of fairness preferences of the same individuals (N=499) in Spring 2020 (Wave 1) and in Autumn 2021 (Wave 2). We observe that the distribution of fairness preferences in our sample does not change significantly over time (Chi^2, p=0.177).[36] We find no evidence that the share of subjects with meritocratic fairness preferences in the spectator game decreases over time (49.3% vs 53.1%, p=0.229), and we also do not observe a strong egalitarian (12.6% vs 15.2%, p=0.235) or libertarian shock (5.4% vs 4.8%, p=0.666) due to the crises. There is only a slight tendency for less subjects with "*Other*" fairness preferences in Wave 2 (32.6% vs 26.9%, p=0.045). When we compare the means and distributions of transfer choices in Wave 1 and Wave 2, we also see that they do not differ significantly (see Appendix A.1.4.3). Hence, we conclude:

Result 4: Fairness preferences are stable at the population level over time.

FIGURE 2.4: Stability of Fairness Preference Types Over Time

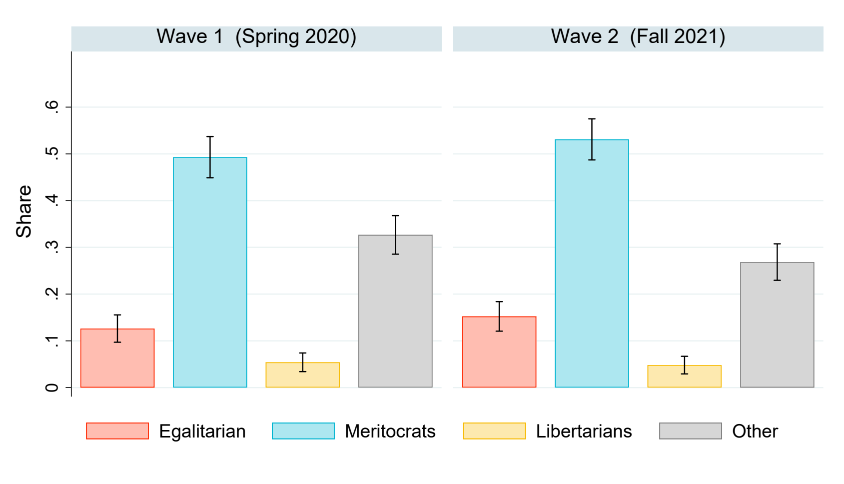

Notes: The figure shows histograms of fairness preferences in Wave 1 and Wave 2 of our panel.

[36]This also holds true if we include the meritocratic subtypes in the distribution (Chi^2, p=0.393) (see Appendix Table (A.16).

Panel Data - Individual Level The stability of fairness preferences at the population level does however not imply preference stability at the individual level. As a first test that confirms some type stability at the individual level over time, we show that an individual's type in Wave 1 is predictive of its type in Wave 2 for all fairness preference types (including "*Others*") (see Appendix Table A.17). At the same time, there is a substantial number of transitions between types (see Appendix Table A.18 for a transition matrix). While almost no transitions take place between the two extreme types (N=3), only half of individuals (47.7%) are classified as the same type in Wave 1 and in Wave 2. A number of empirical tests support the view that these type transitions do not reflect meaningful changes in people's fairness preferences: First, types in Wave 1 are almost as predictive of policy preferences in Wave 2 as types in Wave 2 - even when controlling for types in Wave 2 (see Appendix Table A.19). We would not expect to find this pattern if most type transitions were caused by meaningful changes in fairness preferences. Second, changes in transfer choices are not related to personal experiences such as income shocks, job loss, or COVID cases (see Appendix Table A.30). Third, type transitions are not meaningfully related to changes in policy preferences over time (see Appendix Table A.20). These analyses suggest that most transitions between types reflect measurement error rather than meaningful changes in fairness preferences. At the same time, it may also be the case that fairness preference types are not as discrete as previously thought, which seems to be an interesting avenue for future research.

If type transitions reflect measurement error, then we should be able to leverage the repeated elicitation of fairness preference types to explain policy preferences. In fact, the predictive power of fairness preferences for policy preferences increases substantially once both type classifications are taken into account (in terms of R^2 by about 50 to 70% to then 9% of the total variance, see Appendix Table A.19), even strengthening the results presented in Section 2.5.[37]

At the same time, there is no way to empirically rule out the existence of any meaningful changes in people's fairness preferences in our panel data. Still, given that we do not observe any changes in fairness preferences at the population level, we conclude that changes in fairness preferences can not explain any aggregate changes in people's support for welfare policies over our study period.

Experimental data In short, consistent with our findings from the panel data, we find that neither the distribution of fairness preference types nor the means or distributions of transfer choices change significantly in response to the recall of personal experiences

[37] The total variance in policy preferences that can be explained by heterogeneity in fairness preferences increases to R^2=12% once we also include the meritocratic subtypes. As a benchmark, economic background characteristics (income bracket and employment status) alone can only explain up to 6.5% of the variance, even if we also include measures from both waves.

of *High Control* versus *Low Control* from the first wave of the pandemic (see Appendix A.1.4.4). Hence, the treatments and the associated emotional reactions do not cause meaningful changes in subjects' fairness preferences.[38]

2.6.2 Beliefs Change in Times of Crises

Panel Data - Population Level Figure 2.5b depicts the mean difference in beliefs in merit between Wave 1 and Wave 2 of our panel data with a 95% confidence interval, estimated using a simple random effects panel data model that controls for treatment assignment in Wave 1 (-0.213, $p = 0.042$). Estimates from individual fixed effects models are almost identical in size. Once we additionally control for socio-demographics (including age) and economic background characteristics, the estimated effect increases slightly (-0.239, $p = 0.025$). See Appendix A.1.4.5 for a detailed comparison of the panel data models that we have estimated. Taken together, these panel data models provide evidence for an average decline in beliefs in merit of around 0.1 SD over our study period.

The significance level of the estimated effect cautions against drawing strong conclusions based on our data alone. Therefore, we analyze the most recent wave of the General Social Survey (GSS) from 2022, which also indicates that the belief in a meritocratic society has decreased among US Americans in comparison to the pre-pandemic years.[39] While in 2018, 72.3% of US Americans believed that hard work is most important for getting ahead in life (instead of luck), this share reduced to 63.0% in 2022.[40] This corresponds to a decrease of 0.13 SD between 2018 and 2022 (t-test: p<0.005). The effect is illustrated in Figure 2.5c compared to the long-term trend over the last two decades. Figure 2.5c reveals that the belief in merit as a cause of inequality in the US is now at its lowest level since the 2007/2008 financial crisis.

Panel Data - Individual Level Figure 2.5a depicts the individual level changes in beliefs between Wave 1 and Wave 2. The differences between Wave 1 and 2 tend to be substantial in size: the absolute difference is larger or equal to 1 point for 50.9% of subjects. Also, there is a considerable share of subjects that also increased their belief in merit as a cause of inequality over our study period by 1 or more points (21.6%). In that way, the graphic reveals that there is a lot of variation in beliefs at the individual level in addition to the average decline observed in our entire sample.

[38] In our experiment, we were powered to detect an effect of 0.25 SD (80% power, α=0.05) and can therefore only rule out large effects.

[39] Due to the pandemic, the GSS had to change their survey methodology in 2020/2021 from in-person interviews to online, which required changes in the question-wording. These changes make the 2020/2021 data incomparable to the 2018 data. In 2022, the GSS again conducted in-person interviews using the same question wording as before the pandemic. We use the 2022 in-person sample for our comparison with the 2018 data, reweighted using the standard "wtssnrps" weights.

[40] Moreover in 2018, 15.3% believed that luck and hard work are equally important (2022: 25%), and 12.4% believed that luck is most important (2022: 12%).

FIGURE 2.5: Stability of Beliefs Over Time

(A) Panel: Indiv. Changes (B) Panel: Avg. Change (C) GSS Panel

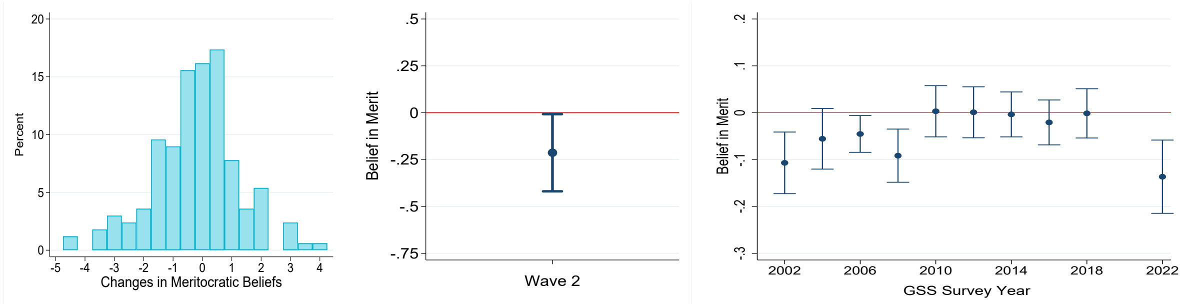

Notes: Figure (a) shows a histogram of the individual level changes in beliefs between wave 1 and wave 2. Figure (b) shows the estimated average change in beliefs over time. Figure (c) shows the mean belief in merit in the General Social Survey over the past 20 years. The question ("getahead"): "Some people say that people get ahead by their own hard work; others say that lucky breaks or help from other people are more important. Which do you think is most important?". Answers are coded on a three point scale: luck most important, hard work most important, both equally important. For 2022 we only use the in-person sample, which uses the same methodology as in the pre-pandemic years (see also Footnote 39).

Understanding the Mechanism of Belief Change Our panel data collection and our experiment were designed to test how personal experiences of high or low control over one's life outcomes impact fairness views. As a first indication for a relationship between personal experiences (e_i) and beliefs about the causes of inequality (b_i), Figure 2.6a shows a binscatter plot indicating a very strong correlation ($\rho = 0.31$, p<0.001) between subjects' own perceived control over their life and their beliefs in a meritocratic society. A lower perceived control over one's own life is associated with the belief that luck rather than merit causes inequalities in society.

Can our experimental manipulation provide credible causal evidence for the hypothesis that lower perceived control over one's own life leads to a higher belief that luck rather than merit causes inequalities in society? Figure 2.6b shows the estimated treatment effects of the *High Control* and *Low Control* treatment in our experiment (N=745). The estimated difference in beliefs between *High Control* and *Low Control* is 0.40 points on the 10-point Likert scale (p=0.027). Moreover, we can see that the treatment effects relative to the Baseline treatment are symmetric, that is, roughly of equal size. In combination with the manipulation check, which shows that subjects on average reported a significantly lower perceived control over their life in the *Low Control* treatment compared to the *High Control* treatment (diff=-0.47 on a 10-point Likert scale, t-test, p=0.014), our experiment provides causal evidence that a personal experience of loss of control - even when just recalled - can reduce beliefs in merit as a cause of inequality in society. The observed treatment effects are however just transitory, which should also be expected given that the manipulation builds on recall and salience of past experiences. In Wave 2 of our panel

Figure 2.6: Mechanism: Beliefs and Personal Experiences of Low Control

(a) Correlation (b) Experiment - Treatments (c) Panel - Subgroups

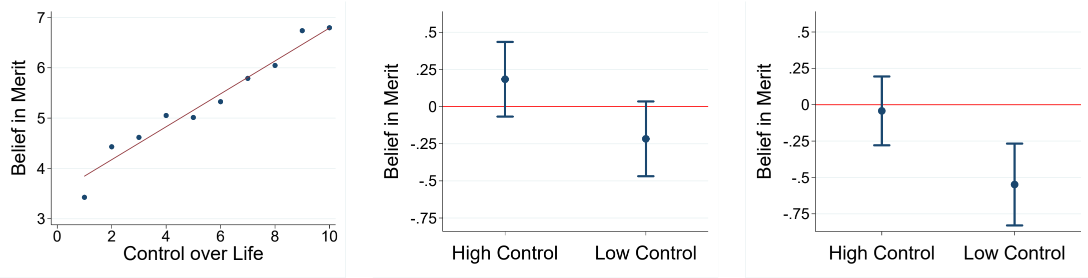

Notes: Figure (a) shows a binscatter plot for subjects own sense of control over their lives and their beliefs for the whole sample. Figure (b) shows the treatment effects or our experimental manipulation of recalling personal experiences from the pandemic of "High Control" or "Low Control" relative to the baseline condition. Figure (c) shows changes in beliefs over time for two subgroups: people who report the same or a higher control over their lives in wave 2 compared to wave 1 ("High Control") and those who report a lower control ("Low Control").

data, the beliefs of subjects do not differ according to their treatment status in Wave 1.[41]

To test whether the observed changes in our panel data can also be plausibly explained by personal experiences of low control, we split our sample into two subgroups: those who report a lower sense of control over their lives in Wave 2 compared to Wave 1 ("Low Control", 30.0% of subjects) and those who report the same or a higher sense of control in Wave 2 ("High Control", 70.0% of subjects). Figure 2.6c shows changes in beliefs over time for these two subgroups. The corresponding fixed effects panel data model, which controls for treatment status in Wave 1, reveals that subjects who reported a lower sense of control reduce their beliefs in meritocracy on average by half a point (-0.55, p=0.001), while all other subjects on average do not change their beliefs (+0.04, p=0.595).[42] Thus, consistent with our experimental result, personal experiences of loss of control also seem to reduce beliefs in merit as the main cause of economic inequality in our panel data.

> **Result 5:** Beliefs in merit as a cause of inequality can decline in times of crises when people make personal experiences in which they lose control over their life outcomes.

[41] In Wave 2, the mean belief of subjects that had been in Baseline is 5.68, which is indistinguishable from the beliefs of subjects who had been in High Control (mean 5.65, t-test p=0.926) or Low Control (mean 5.62, t-test p=0.808). See Appendix Figure A.12 for a graphical overview.

[42] Almost identical estimates are derived when we restrict the analysis to subjects that were assigned to the baseline condition in Wave 1: a reduction in beliefs of -0.56 point (p=0.020) for those who lost control, compared to no change among all other subjects (-0.04, p=0.777). Thus, this subgroup effect cannot be explained by the treatments in Wave 1.

2.6.3 Explaining Changes in Support for Welfare Policies

Can the observed changes in beliefs about the causes of inequality explain changes in policy preferences? First, we can test whether our experimental manipulation of loss of control, which changed beliefs about the causes of inequality, also caused a change in support for welfare policies. When we test for treatment effects using OLS regressions that once control for treatment dummies and once for treatment dummies and control variables in our experimental sample (N=745), we do not find any significant treatment effects on people's support for welfare policies (see Appendix Table A.25). Our findings in this regard are therefore in line with many papers in the literature that use similar experimental manipulations, such as the provision of information, which often find a first-stage effect on beliefs (see e.g. Alesina et al., 2018; Kuziemko et al., 2015), but no significant effects on people's policy preferences.[43]

However, our individual-level panel data allows us to investigate the association between changes in beliefs and changes in policy preferences at the individual level over time. Table 2.4 shows that changes in beliefs in merit over time are indeed associated with changes in support for welfare policies. In columns (1), (3) and (5) we present estimates from two-way fixed effects models (wave and individual fixed effects). The fixed effects models establish that there is a significant association between changes in beliefs about the causes of inequality and changes in support for welfare policies at the individual level: subjects who increase their belief in luck as a cause of inequality in US society between Wave 1 and Wave 2 are also significantly more in support of welfare policies in Wave 2. This pattern holds across policy preferences towards redistribution, universal health care, and the pandemic support package.

> **Result 6:** Changes in beliefs about the causes of inequality over time are associated with changes in support for welfare policies at the individual level.

The estimates in columns (1), (3), and (5) may not be cleanly causally identified, because another time-varying confounder may have caused both, the changes in beliefs and the changes in support for welfare policies. In columns (2), (4), and (6), we control for economic background characteristics (income bracket and employment status) to account for the potentially confounding influence of changes in economic self-interest. Controlling for changes in economic self-interest does not move the estimates by much. In sum, these individual-level panel data results therefore provide more convincing evidence for a

[43] In our experiment, the absence of treatment effects on policy preferences could plausibly be due to an unobserved confounding mechanism that offsets the effect on beliefs: for example, it seems plausible that recalling *Low Control* experiences from the first wave of the pandemic also reduces trust in the US government, because the government did not manage to protect its citizens from the impacts of the pandemic or because the government was responsible for the lockdown policies. Unfortunately, we did not measure trust in government in our experimental sample in Wave 1 to be able to shed light on this mechanism.

causal effect of beliefs about the causes of inequality on policy preferences than a simple correlation shown in most of the existing literature.

Table 2.4: Panel Data: Explaining Changes in Policy Preferences

	Redistribution		Univ. Health Care		Pandemic Support	
	(1)	(2)	(3)	(4)	(5)	(6)
Beliefs in Merit	-0.07**	-0.07**	-0.04*	-0.04*	-0.09**	-0.09**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Income Bracket	No	Yes	No	Yes	No	Yes
Employment Status	No	Yes	No	Yes	No	Yes
p (Belief)	0.007	0.008	0.017	0.017	0.004	0.003
Observations	998	998	998	998	998	998
Clusters	499	499	499	499	499	499

Notes: Coefficients from two-way fixed effects panel data models. Robust standard errors clustered at the individual level in parentheses. Dependent variables: policy preferences as standardized z-scores.
* $p<0.05$, ** $p<0.01$, *** $p<0.001$

2.7 Conclusion

Large strands of literature in political economy have focused on beliefs as the main source of heterogeneity in people's fairness views. In this paper, we show how taking heterogeneity in fairness preferences into account can systematically advance our understanding of people's policy preferences. In our sample, there are large and robust differences in support for welfare policies between individuals with egalitarian, libertarian and meritocratic fairness preferences. Moreover, fairness preferences predict how much policy preferences depend on beliefs about the causes of inequality. These insights about the fundamental properties of fairness views seem relevant in a wide variety of economic settings, ranging from wage setting in firms to support for affirmative action, in which people may demand fair institutions and fair policies.

At the same time, our data also highlight that fairness preferences are rather stable over time, so that changes in policy preferences over time are rather caused by changes in beliefs about the causes of inequality or by economic self-interest. Our results suggest that one relevant mechanism through which beliefs about the causes of inequality change is through personal experiences. Personal experiences in societal crises thus may be an important driver in the formation of people's policy preferences and political ideologies.

A key question for future research is whether the declining belief in a meritocratic society over the course of the pandemic, observed in our panel data and the GSS data, generalizes beyond the US context, and to study whether it has long-term consequences for the public support for welfare policies. If indeed citizens will support, and demand, more expansive welfare policies, we might witness changes in the institutional design of the US welfare state. One important aspect, which we have neglected in this paper, is that - in theory - declining beliefs in merit should also have consequences for labor supply and for investments in human capital. In that way, a decrease in meritocratic beliefs over the course of the pandemic could be related to the "*Great Resignation*" in the US labour market following the pandemic.

Appendix to Chapter 2

A.1 Additional Analyses

A.1.1 Descriptives

A.1.1.1 Histograms of Transfer Choices

FIGURE A.1: Histograms of Transfer Choices

(A) Transfer Luck

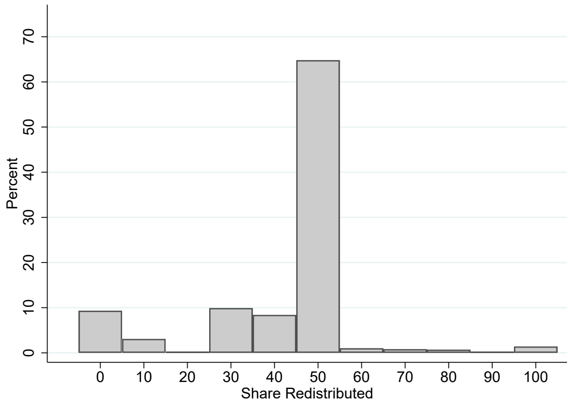

(B) Transfer Effort

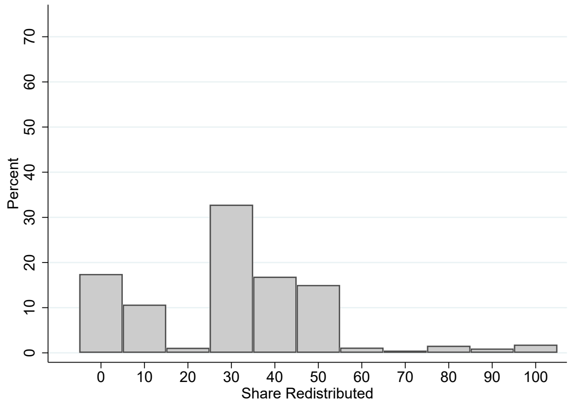

(C) Transfer Luck with Self-Interest

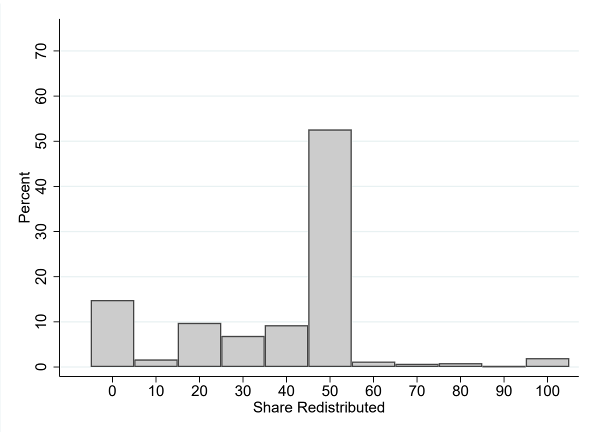

Notes: The figure shows histograms of transfer choices in the spectator game on inequalities due to luck (a), in the spectator game on inequalities due to merit (b) and in the transfer decision on inequalities due to luck with self-interest (c).

A.1.1.2 Type Classification: Order Effects

Table A.1: Type Classification: Order Effects

	Between		Within	
	1st Choice	2nd Choice	1st Choice Merit	1st Choice Luck
Fairness Preference Type	(1)	(2)	(3)	(4)
Egalitarians	15.7%	12.6%	11.4%	12.4%
Libertarians	8.4%	9.9%	6.3%	4.3%
Meritocrats	46.7%	49.3%	50.6%	48.8%
Other	29.1%	28.2%	31.7%	34.5%
N	1975	1975	989	986

Notes: The table shows order effects on the distribution of fairness preference types for between- and within-subjects type classifications. All classifications use the same strict definition for meritocrats used in the between subjects classification. Column (1) shows a between-subject type classification using just the first choice of subjects. Column (2) shows a between-subjects type classification using just the second choice of subjects. Column (3) shows the within-subjects type classification for those subjects who were first randomly assigned the merit condition. Column (4) shows the within-subjects type classification for those subjects who were first randomly assigned the luck condition.

A.1.1.3 Type Classification: Comparison to the Literature

Table A.2: Type Classification: Comparison to the Literature

	Almås et al. (2020)	Cohn et al. (2023)	Our study	Our study
Classification:	Between	Between	Within	Between 1st Choice
Types				
Egalitarians	15.3%	17.8%	11.9%	15.7%
Libertarians	29.4%	12.1%	5.3%	8.4%
Meritocrats	37.5%	60.5%	49.7%	46.7%
Other	17.8%	9.6%	33.1%	29.1%
N	1000	417	1975	1975
Size of Choice Set	7	7	40x40	40
Survey Company	Research Now (Dynata)	YouGov	Prolific	Prolific

Notes: The table shows type distributions in representative US samples in the existing literature.

A.1.1.4 Are Fairness Preferences and Beliefs Correlated?

FIGURE A.2: Beliefs in Merit by Fairness Preference Type

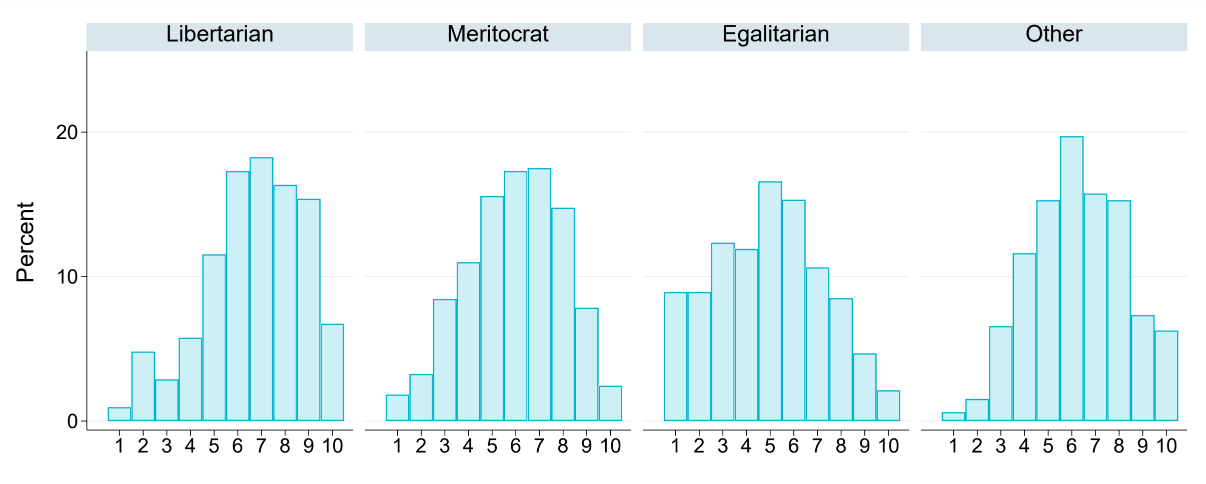

Notes: The figure shows histograms of beliefs in merit by fairness preference type.

A.1.1.5 Subtypes among Meritocrats - Preferences

FIGURE A.3: Classification of Subtypes among Meritocrats - Preferences

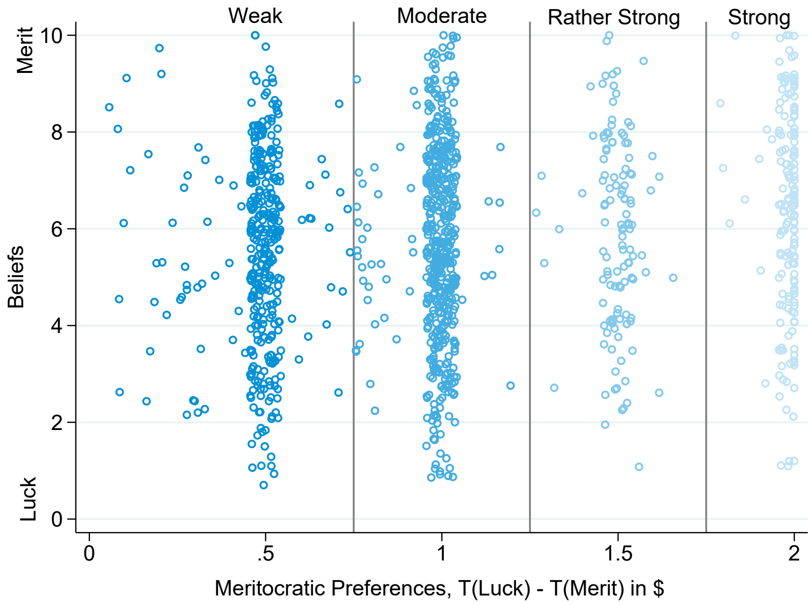

Notes: The figure shows classification of meritocratcs into four subtypes according to their meritocratic preferences using the classification described in Table A.3.

TABLE A.3: Classification of Subtypes among Meritocrats - Preferences

Subtype	Classification	N	Share
Preferences			
Weak	$\$0.00 \leq \mathrm{T_{Luck}} - \mathrm{T_{Merit}} < \$0.75$	249	25.4%
Moderate	$\$0.75 \leq \mathrm{T_{Luck}} - \mathrm{T_{Merit}} < \$1.25$	446	45.4%
Rather Strong	$\$1.25 \leq \mathrm{T_{Luck}} - \mathrm{T_{Merit}} < \$1.75$	107	10.9%
Strong	$\$1.75 \leq \mathrm{T_{Luck}} - \mathrm{T_{Merit}} \leq \$2.00$	180	18.3%
Total		982	100%

Notes: The table reports the classification of meritocratcs into four subtypes according to their meritocratic preferenes.

A.1.1.6 Subtypes among Meritocrats - Beliefs

FIGURE A.4: Classification of Subtypes among Meritocrats - Beliefs

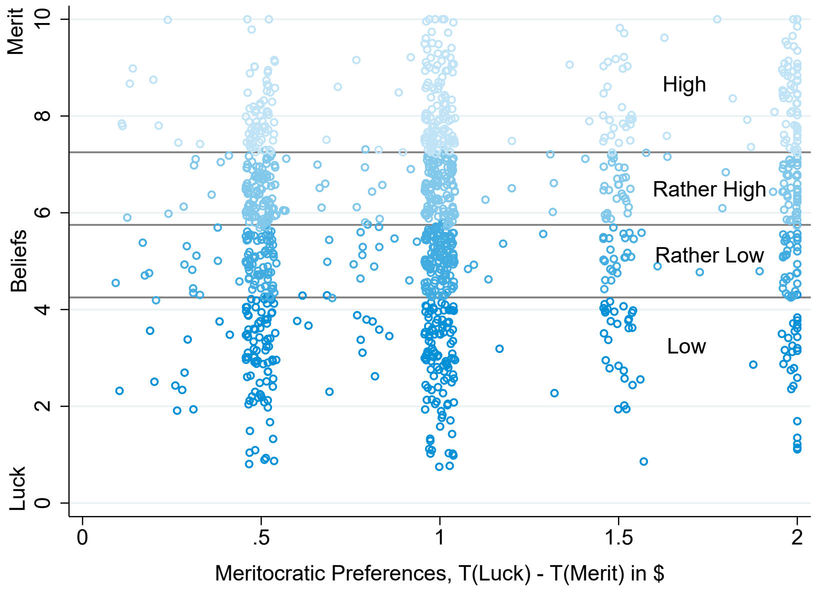

Notes: The figure shows classification of meritocrats into four subtypes according to their beliefs about the causes of inequalities using the classification described in Table A.4.

TABLE A.4: Classification of Subtypes among Meritocrats - Beliefs

Subtype	Classification	N	Share
Beliefs			
Low	$1.0 \leq$ Beliefs ≤ 4.0	241	24.5%
Rather Low	$4.5 \leq$ Beliefs ≤ 5.5	239	24.3%
Rather High	$6.0 \leq$ Beliefs ≤ 7.0	256	26.1%
High	$7.5 \leq$ Beliefs ≤ 10	246	25.1%
Total		982	100%

Notes: Table reports the classification of meritocratcs into four subtypes according to their beliefs about the causes of inequalities.

A.1.1.7 Support for Welfare Policies - Principal Components

Table A.5 shows that the first principal component can explain 62% of the total variance in policy preferences. Moreover, the first principal component is the only component with an eigenvalue large than one, as shown in A.5. Based on the standard criterion to only use components with eigenvalues larger than 1, subjects support for welfare policies can thus be well described by the first principal component alone. Table A.6 shows that the six policies are assigned almost equal weights to construct the first principal component.

TABLE A.5: Eigenvalues of Components and Proportion of Variance Explained

	Eigenvalue	Proportion
1st Component	3.72	62.0
2nd Component	0.76	12.7
3rd Component	0.53	0.09
4th Component	0.40	0.07
5th Component	0.33	0.05
6th Component	0.26	0.04

FIGURE A.5: Eigenvalues after Principal Component Analysis

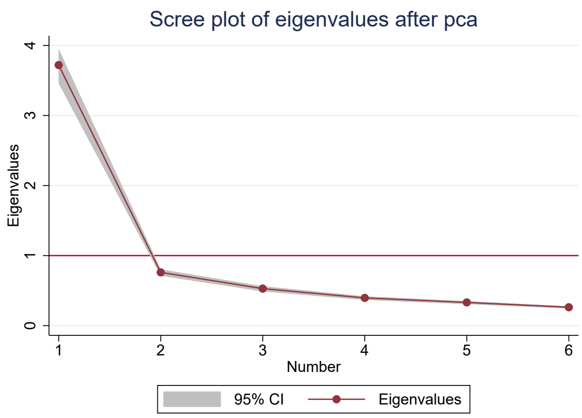

Notes: The figure shows the eigenvalues of the first six principal components with 95% confidence intervals.

TABLE A.6: Principal Components

	1st Comp	2nd Comp	3rd Comp	4th Comp	5th Comp	6th Comp
Redistribution	0.3912	0.5379	-0.2475	0.5209	-0.4132	-0.2332
Univ. Health Care	0.3839	0.6057	0.2914	-0.2571	0.4437	0.3712
E.I. Payments	0.4077	-0.3380	-0.5079	0.2857	0.6163	-0.0103
UE Benefits	0.4318	-0.2393	-0.3194	-0.3894	-0.4909	0.5115
Medicaid	0.4478	-0.1083	0.1502	-0.4904	-0.0166	-0.7242
Paid Sick Leave	0.3825	-0.4008	0.6865	0.4345	-0.1061	0.1472

A.1.2 How Fairness Views Shape Policy Preferences

A.1.2.1 Benchmarking: Fairness Preferences and Income

Table A.7: Benchmarking Fairness Preference Types and Income

	Dep Var: Support for Welfare Policies			
	(1)	(2)	(3)	(4)
Libertarian	-0.88***			-0.80***
	(0.13)			(0.13)
Meritocrat	-0.44***			-0.39***
	(0.06)			(0.06)
Other	-0.59***			-0.56***
	(0.07)			(0.07)
Income (20-35k)		-0.11		-0.12
		(0.09)		(0.09)
Income (35-50k)		-0.19*		-0.20*
		(0.09)		(0.09)
Income (50-75k)		-0.22*		-0.23*
		(0.09)		(0.09)
Income (75-100k)		-0.29**		-0.29**
		(0.10)		(0.10)
Income (100-150k)		-0.47***		-0.43***
		(0.10)		(0.10)
Income (>150k)		-0.42***		-0.41***
		(0.11)		(0.11)
Unemployed		0.19**		0.14*
		(0.06)		(0.07)
Not in Labor Force		-0.06		-0.00
		(0.07)		(0.06)
Age (in decades)			-0.06***	-0.06***
			(0.02)	(0.02)
Female			0.18***	0.12*
			(0.05)	(0.05)
College Degree			-0.03	0.04
			(0.06)	(0.06)
Masters Degree			0.02	0.16*
			(0.07)	(0.07)
Black			0.42***	0.39***
			(0.06)	(0.06)
Asian			0.21*	0.26**
			(0.09)	(0.09)
Race Other			0.19	0.13
			(0.13)	(0.12)
Hispanic			0.18	0.19
			(0.10)	(0.10)
Constant	0.46***	0.24**	0.08	0.72***
	(0.05)	(0.07)	(0.09)	(0.12)
Observations	1975	1975	1975	1975
R^2	0.041	0.028	0.042	0.101

Notes: OLS estimates with robust standard errors clustered at the individual level in parentheses. Dependent variables: support for welfare policies (standardized first principal component of all policy preferences). Omitted category: "Egalitarians".
*** p<0.001, ** p<0.01, * p<0.05

A.1.2.2 Robustness: Fairness Preferences Predict Policy Preferences

TABLE A.8: Robustness: Fairness Preferences Predict Policy Preferences (Wave 2)

	Dep Var: Support for Welfare Policies					
	(1)	(2)	(3)	(4)	(5)	(6)
Libertarian	-0.60***	-0.45**	-0.49***	-0.48***	-0.33**	-0.32**
	(0.15)	(0.15)	(0.14)	(0.14)	(0.12)	(0.12)
Meritocrat	-0.29***	-0.29***	-0.28***	-0.29***	-0.23***	-0.20***
	(0.07)	(0.07)	(0.07)	(0.07)	(0.06)	(0.06)
Other	-0.41***	-0.40***	-0.48***	-0.47***	-0.33***	-0.33***
	(0.08)	(0.08)	(0.08)	(0.08)	(0.07)	(0.07)
Beliefs in Merit	-0.21***	-0.21***	-0.21***	-0.21***	-0.13***	-0.10***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Selfish Type		-0.28**	-0.18*	-0.19*	-0.15	-0.11
		(0.09)	(0.09)	(0.09)	(0.08)	(0.08)
Trust in Government			0.38***	0.39***	0.28***	0.19***
			(0.05)	(0.05)	(0.04)	(0.04)
National Group Aff.				-0.02	0.04	0.07*
				(0.03)	(0.03)	(0.03)
Political Ideology					-0.17***	-0.00
					(0.01)	(0.02)
Liberal						-0.30***
						(0.06)
Moderate						-0.59***
						(0.10)
Conservative						-0.81***
						(0.15)
Very Conservative						-1.31***
						(0.21)
No Vote/Other						-0.24***
						(0.07)
Voted Trump 2020						-0.73***
						(0.10)
No Vote/Other						0.05
						(0.06)
Voted Trump 2016						-0.11
						(0.09)
Socio-Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Income Bracket	Yes	Yes	Yes	Yes	Yes	Yes
Employment Status	Yes	Yes	Yes	Yes	Yes	Yes
p (Egalitarian = Libertarian)	< 0.001	0.002	< 0.001	0.001	0.009	0.009
p (Egalitarian = Meritocrat)	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	0.001
p (Meritocrat = Libertarian)	0.025	0.236	0.111	0.131	0.422	0.308
p (joint test)	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	0.001
Observations	1230	1230	1230	1230	1230	1230
R^2	0.267	0.274	0.320	0.320	0.462	0.526

Notes: OLS estimates with robust standard errors clustered at the individual level in parentheses. Dependent variables: support for welfare policies (standardized first principal component of all policy preferences). Omitted category: "Egalitarians". Socio-Demographics include age, gender, race, ethnicity and education dummies. The "joint test" tests the hypothesis that there is no difference between any of the three fairness preference types (egalitarians, libertarians and meritocrats).
*** p<0.001, ** p<0.01, * p<0.05

A.1.2.3 Robustness: Meritocratic Preference Subtypes

FIGURE A.6: Meritocratic Preference Subtypes Predict Support for Welfare Policies

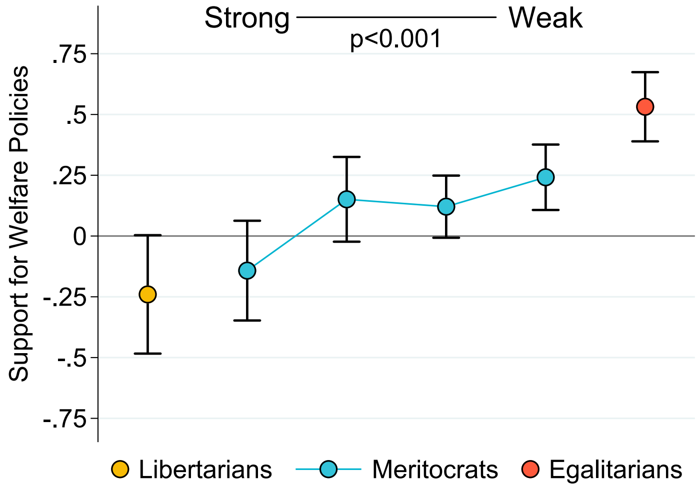

Notes: The figure shows estimated coefficients and 95% confidence intervals from OLS regressions explaining support for welfare policies (first principal component of all policy preferences). The figure shows coefficients for the egalitarian, libertarian and meritocratic subtype dummies. Regressions control for socio-demographics (age, gender, race/ethnicity, education) and economic background characteristics (income bracket and employment status). Subjects classified as “other” serve as the reference group in both regressions. For details on the type classification see Table A.3. Robust standard errors are clustered at the individual level.

Table A.9: Robustness: Meritocratic Preference Subtypes Predict Policy Preferences

	Dep Var: Support for Welfare Policies	
	(1)	(2)
Egalitarian	0.35***	0.21***
	(0.06)	(0.05)
Weak Meritocrat	0.28***	0.12*
	(0.07)	(0.06)
Moderate Meritocrat	0.10	0.03
	(0.06)	(0.05)
Rather Strong Meritocrat	0.17*	0.09
	(0.08)	(0.07)
Strong Meritocrat	-0.07	-0.11
	(0.09)	(0.08)
Libertarian	-0.18	-0.14
	(0.10)	(0.09)
p (Weak = Strong)	0.000	0.006
p (joint test)	0.002	0.041
Political Ideology	No	Yes
Belief - Causes of Inequality	Yes	Yes
Socio-Demographics	Yes	Yes
Income Bracket	Yes	Yes
Employment Status	Yes	Yes
Observations	1975	1975
R^2	0.237	0.394

Notes: OLS estimates with robust standard errors clustered at the individual level in parentheses. Dependent variables: support for welfare policies (standardized first principal component of all policy preferences). Omitted category: "Other" type. Socio-Demographics include age, gender, race, ethnicity and education dummies. The "joint test" tests the hypothesis that there is no difference between any of the meritocratic subtypes. For details on the type classification see Table A.3.
*** p<0.001, ** p<0.01, * p<0.05

A.1.2.4 Robustness: Interaction Effects

Table A.10: Robustness #1: Interaction between Fairness Preferences and Beliefs

	Redistribution		Univ. Health Care		Pandemic Support	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Meritocrats vs Non-Meritocrats						
Beliefs in Merit	-0.24***	-0.15***	-0.22***	-0.11***	-0.19***	-0.10***
	(0.02)	(0.01)	(0.02)	(0.01)	(0.02)	(0.02)
Non-Meritocrat × Beliefs in Merit	0.10***	0.07***	0.09***	0.06**	0.07***	0.04*
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Type FE	Yes	Yes	Yes	Yes	Yes	Yes
p (Non-Meritocrat x B)	<0.001	<0.001	<0.001	0.002	0.001	0.028
Political Ideology	No	Yes	No	Yes	No	Yes
Socio-Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Income Bracket	Yes	Yes	Yes	Yes	Yes	Yes
Employment Status	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1975	1975	1975	1975	1975	1975
R^2	0.275	0.392	0.196	0.381	0.174	0.282

Notes: OLS estimates with robust standard errors clustered at the individual level in parentheses and p-values in square brackets. Dependent variables: policy preferences as standardized z-scores. Omitted category: "Meritocrats".
*** p<0.001, ** p<0.01, * p<0.05

TABLE A.11: Robustness #2: Interaction between Fairness Preferences and Beliefs

	Redistribution		Univ. Health Care		Pandemic Support	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Preferences (Subtypes) Interact w. Beliefs among Meritocrats						
Beliefs in Merit	-0.30***	-0.18***	-0.30***	-0.16***	-0.28***	-0.17***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)
Rather Strong	-0.57	-0.50	-0.45	-0.37	-0.82**	-0.76*
	(0.30)	(0.28)	(0.27)	(0.24)	(0.31)	(0.29)
Rather Strong × Beliefs in Merit	0.15**	0.13**	0.10	0.08	0.16**	0.14**
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Moderate	-0.34	-0.23	-0.48*	-0.35	-0.55*	-0.44*
	(0.21)	(0.19)	(0.22)	(0.20)	(0.23)	(0.22)
Moderate × Beliefs in Merit	0.08*	0.06	0.08*	0.06	0.11**	0.09*
	(0.04)	(0.03)	(0.04)	(0.04)	(0.04)	(0.04)
Weak	-0.21	-0.13	-0.70**	-0.60**	-0.60*	-0.53*
	(0.21)	(0.19)	(0.22)	(0.20)	(0.25)	(0.23)
Weak × Beliefs in Merit	0.10*	0.07*	0.16***	0.13***	0.13**	0.11*
	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)	(0.04)
p (Rather Strong X B)	0.007	0.009	0.057	0.062	0.004	0.005
p (Moderate X B)	0.042	0.075	0.043	0.099	0.009	0.024
p (Weak X B)	0.007	0.009	0.057	0.062	0.004	0.005
p (joint test)	0.024	0.050	0.001	0.005	0.008	0.026
Political Ideology	No	Yes	No	Yes	No	Yes
Socio-Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Income Bracket	Yes	Yes	Yes	Yes	Yes	Yes
Employment Status	Yes	Yes	Yes	Yes	Yes	Yes
Observations	982	982	982	982	982	982
R^2	0.323	0.454	0.244	0.430	0.197	0.311

Notes: OLS estimates with robust standard errors clustered at the individual level in parentheses and p-values in square brackets. Dependent variables: policy preferences as standardized z-scores. Omitted category: "Strong Meritocrats". The "joint test" tests the hypothesis that all interaction terms are equal to zero.
*** p<0.001, ** p<0.01, * p<0.05

TABLE A.12: Robustness #3: Interaction between Fairness Preferences and Beliefs

	Redistribution		Univ. Health Care		Pandemic Support	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Meritocratic Preferences (Continuous) Interact w. Beliefs among Meritocrats						
Meritocratic Preferences	0.17	0.14	0.42**	0.40**	0.39**	0.36*
	(0.13)	(0.12)	(0.13)	(0.12)	(0.15)	(0.14)
Beliefs in Merit	-0.15***	-0.07**	-0.11***	-0.00	-0.09**	-0.01
	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)
Meritocratic Preferences × Beliefs in Merit	-0.06*	-0.05*	-0.09***	-0.08***	-0.08**	-0.07**
	(0.02)	(0.02)	(0.03)	(0.02)	(0.03)	(0.03)
p (Meritocratic Preferences x B)	0.012	0.020	0.001	0.001	0.004	0.008
Political Ideology	No	Yes	No	Yes	No	Yes
Socio-Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Income Bracket	Yes	Yes	Yes	Yes	Yes	Yes
Employment Status	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1205	1205	1205	1205	1205	1205
R^2	0.288	0.423	0.214	0.408	0.173	0.295
Panel B: Meritocratic Preferences (Continuous) Interact w. Beliefs in Full Sample						
Meritocratic Preferences	0.19***	0.05	0.28***	0.12*	0.25***	0.12*
	(0.05)	(0.06)	(0.05)	(0.06)	(0.05)	(0.06)
Beliefs in Merit	-0.18***	-0.11***	-0.16***	-0.07***	-0.15***	-0.08***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Meritocratic Preferences × Beliefs in Merit	-0.04***	-0.02*	-0.05***	-0.02**	-0.04***	-0.02**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
p (Meritocratic Preferences x B)	<0.001	0.014	<0.001	0.009	<0.001	0.009
Political Ideology	No	Yes	No	Yes	No	Yes
Socio-Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Income Bracket	Yes	Yes	Yes	Yes	Yes	Yes
Employment Status	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1975	1975	1975	1975	1975	1975
R^2	0.262	0.388	0.189	0.378	0.167	0.279

Notes: OLS estimates with robust standard errors clustered at the individual level in parentheses. Dependent variables: policy preferences as standardized z-scores. Meritocratic preferences are a continuous measure equal to the difference between the two transfer choices on luck and merit in $: ($T_{Luck}$ - T_{Merit}). For meritocrats, meritocratic preferences take values in (0, 2], see also A.3. In the full sample, meritocratic preferences take values in [-4, 2]. Socio-Demographics include age, gender, race, ethnicity and education dummies.
*** p<0.001, ** p<0.01, * p<0.05

A.1.3 Determinants of Fairness Preferences and Beliefs

A.1.3.1 Determinants of Fairness Preferences

In Figure A.7 we restrict the sample to subjects holding one of the three distinct fairness ideals (egalitarians, meritocrats, libertarians) and exclude subjects who cannot be classified ("others"). In that way, we show the determinants of fairness ideals conditional on being classified as one of the three fairness preference types.

FIGURE A.7: Determinants of Fairness Preferences

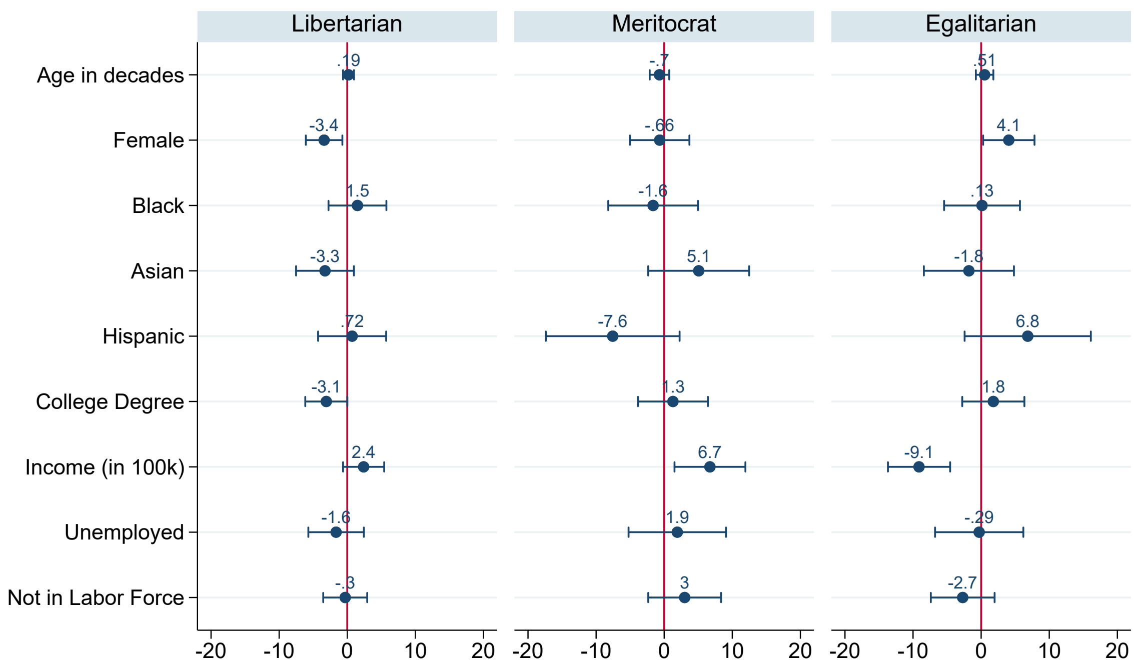

Notes: The figure shows coefficients of three OLS regressions of each fairness preference dummy on our standard set of socio-demographics (age, female, education, income, race/ethnicity dummies). The regressions exclude "type others".

A.1.3.2 Determinants of Beliefs

Figure A.8: Determinants of Beliefs

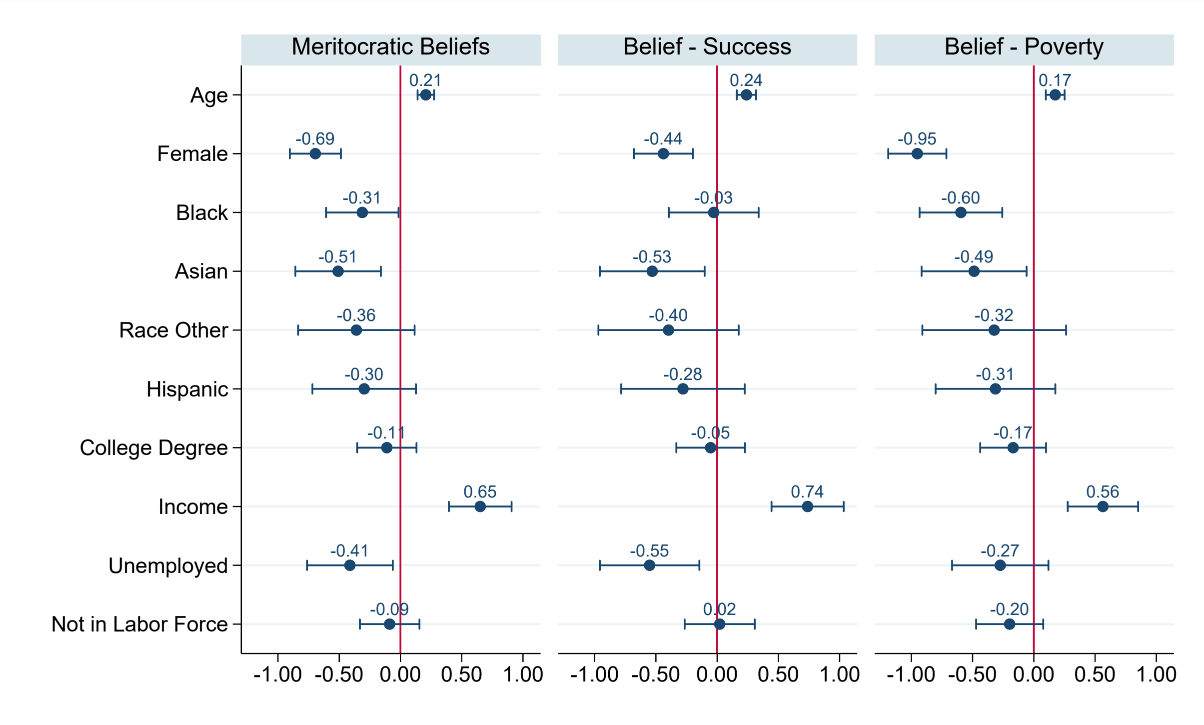

Notes: The figure shows coefficients of three OLS regressions of meritocratic beliefs and each of the two subitems on our standard set of socio-demographics (age, female, education, income, race/ethnicity dummies).

A.1.4 Stability of Fairness Preferences and Beliefs

A.1.4.1 Panel Attrition and Balance of Covariates

Table A.13: Panel Attrition and Balance of Covariates in Experiment

	Panel Attrition	Treatment Status		
Dependent Variable:	Resampled	High Control	Low Control	Baseline
	(1)	(2)	(3)	(4)
Female	-0.004	0.024	0.059	-0.083*
	(0.035)	(0.037)	(0.036)	(0.036)
Age (in decades)	0.078***	-0.011	-0.001	0.012
	(0.012)	(0.013)	(0.013)	(0.013)
Black	0.078	-0.051	0.038	0.013
	(0.049)	(0.052)	(0.054)	(0.053)
Asian	-0.003	-0.104	-0.055	0.161*
	(0.069)	(0.068)	(0.070)	(0.079)
Race Other	-0.073	-0.005	0.026	-0.022
	(0.084)	(0.085)	(0.086)	(0.083)
Hispanic	0.003	-0.006	0.000	0.003
	(0.066)	(0.072)	(0.073)	(0.074)
College Degree	-0.020	-0.011	0.007	0.003
	(0.040)	(0.041)	(0.041)	(0.042)
Masters Degree	-0.083	0.052	0.009	-0.060
	(0.053)	(0.056)	(0.054)	(0.053)
Income (in 100k)	-0.031	-0.036	0.042	-0.007
	(0.042)	(0.045)	(0.044)	(0.045)
Unemployed	-0.081	0.043	-0.042	-0.000
	(0.056)	(0.058)	(0.056)	(0.056)
Not in Labor Force	-0.037	0.051	-0.087*	0.038
	(0.046)	(0.046)	(0.043)	(0.046)
Midwest	0.041	-0.084	0.163**	-0.080
	(0.053)	(0.055)	(0.055)	(0.054)
South	0.050	-0.047	0.039	0.007
	(0.044)	(0.047)	(0.043)	(0.046)
West	0.010	-0.059	0.089	-0.028
	(0.058)	(0.060)	(0.060)	(0.059)
Observations	745	745	745	745
joint significance (p-value)	p<0.001	p=0.815	p=0.227	p=0.311

Notes: Average marginal effects from logit models. Standard errors in parentheses. *** p<0.001, ** p<0.01, * p<0.05

Table A.14: Does Panel Attrition Depend on Outcomes in Wave 1?

	Dep Var: Resampled in Wave 2		
	(1)	(2)	(3)
Meritocrat	-0.017 (0.078)		
Egalitarian	0.095 (0.088)		
Other	-0.021 (0.080)		
Beliefs in Merit	0.002 (0.009)		
Political Ideology		-0.001 (0.006)	
Support for Welfare Policies			0.001 (0.010)
Observations	745	745	745

Notes: Average marginal effects from logit models. Standard errors in parentheses.
*** $p<0.001$, ** $p<0.01$, * $p<0.05$

A.1.4.2 Experiment: Manipulation Check

Figure A.9: Manipulation Check: High Control vs. Low Control

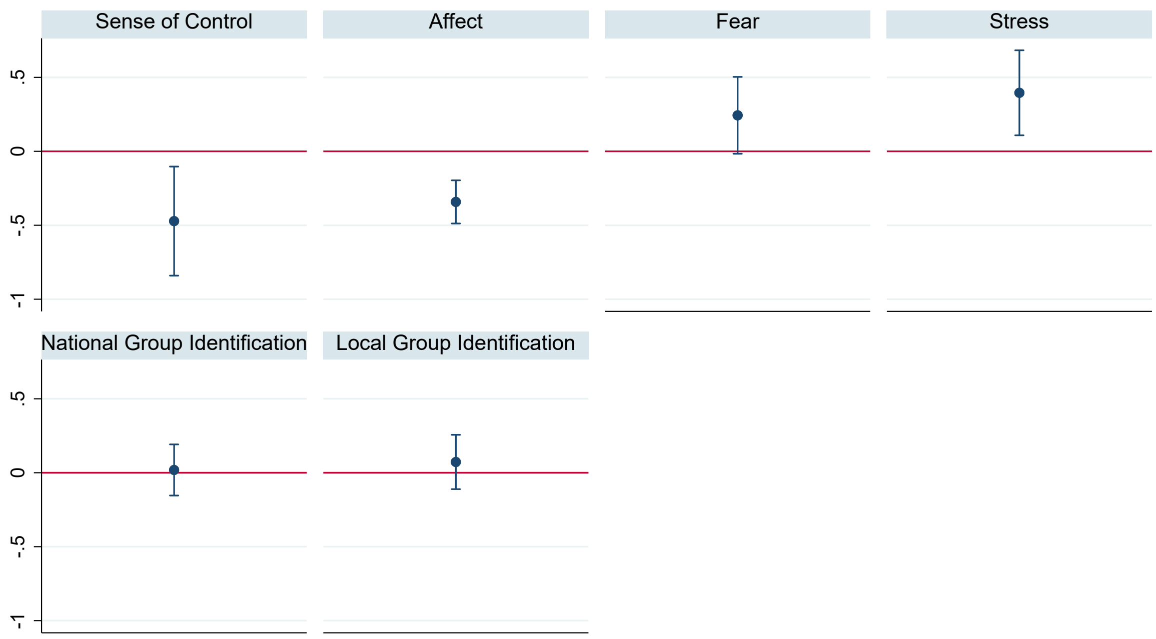

Notes: The figure shows estimated effects with 95% CIs of the Low Control treatment compared to High Control (Red Line). Sense of Control is measured on a 10-point Likert scale. Fear and Stress are measured on 7-point Likert scales. Affect, National and Local Group Affiliation are measured on 5-point Likert scales.

A.1.4.3 Stability of Fairness Preferences: Panel Data

Panel: Population Level Analysis

Figure A.10: Histograms of Transfer Choices W1 vs W2

(a) Luck (b) Merit

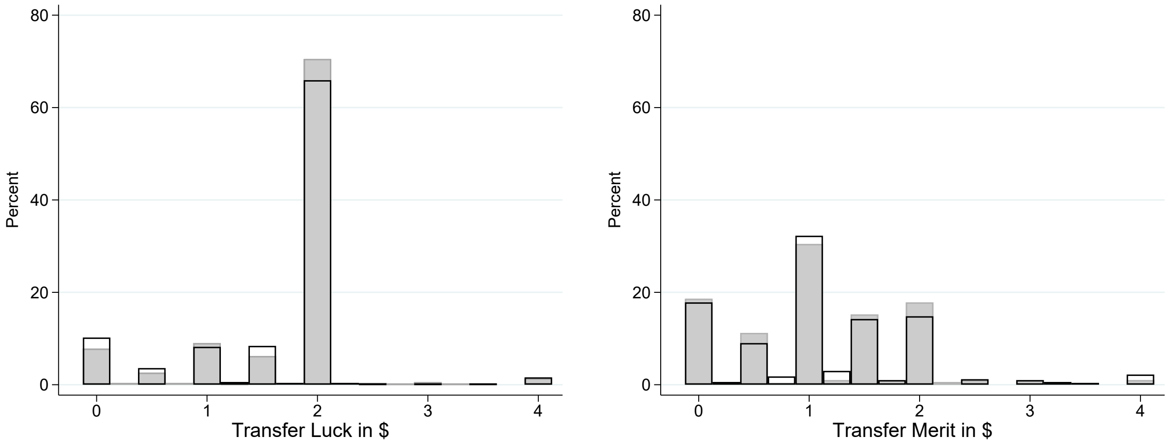

Notes: The figure shows histograms of transfer choices in our panel data (N=499) in wave 1 and wave 2. Gray shaded bars correspond to Wave 2.

Table A.15: Panel: Changes in Transfer Choices over Time

	Transfer Merit		Transfer Luck	
	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Wave 2	-0.041	-0.078	0.058	0.140
	(0.048)	(0.080)	(0.041)	(0.072)
Constant	1.172***	1.246***	1.594***	1.431***
	(0.071)	(0.148)	(0.061)	(0.133)
Individual FE	Yes	Yes	Yes	Yes
Treatment in W1 FE	No	Yes	No	Yes
Observations	998	998	998	998
Panel B: Excluding Outliers (Transfers >$2)				
Wave 2	0.031	-0.041	0.063	0.118
	(0.039)	(0.065)	(0.037)	(0.068)
Constant	0.960***	1.105***	1.534***	1.424***
	(0.058)	(0.120)	(0.056)	(0.126)
Observations	951	951	965	965
Individual FE	Yes	Yes	Yes	Yes
Treatment in W1 FE	No	Yes	No	Yes

Notes: Columns (1) to (4) report fixed-effects panel data estimates with robust standard errors (clustered at the individual level) in parentheses. In Panel B, outliers (transfer choices >2$) are excluded from the analysis. Reference category: Wave 1. In columns (2) and (4) we control for treatment status in Wave 1 (one dummy for the High Control and one for the Low Control condition). In Wave 2, both treatment dummies take the value 0 for all subjects.
*** p<0.001, ** p<0.01, * p<0.05

TABLE A.16: Panel: Fairness Preference (Sub-)Types by Wave

	Wave 1	Wave 2	Total
Egalitarian	63	76	139
	12.6	15.2	13.9
Weak Meritocrat	58	66	124
	11.6	13.2	12.4
Moderate Meritocrat	119	115	234
	23.8	23.0	23.4
Rather Strong Meritocrat	24	29	53
	4.8	5.8	5.3
Strong Meritocrat	45	55	100
	9.0	11.0	10.0
Libertarian	27	24	51
	5.4	4.8	5.1
Other	163	134	297
	32.7	26.9	29.8
Total	499	499	998
	100.0	100.0	100.0
Chi^2: p=0.393			

Notes: The table shows the distribution of fairness preference types in wave 1 and wave 2. The first number in each cell refers to the number of observations, the second to the share in each column.

Panel Data: Individual Level Analysis

Table A.17: Panel: Fairness Preference Types in Wave 1 Predict Types in Wave 2

	Egalitarian W2	Libertarian W2	Meritocrat W2	Type Other W2
	(1)	(2)	(3)	(4)
Egalitarian W1	0.23***			
	(0.05)			
Libertarian W1		0.14***		
		(0.04)		
Meritocrat W1			0.16***	
			(0.04)	
Type Other W1				0.15***
				(0.04)
Constant	0.12***	0.04***	0.45***	0.22***
	(0.02)	(0.01)	(0.03)	(0.02)
p-value	<0.001	<0.001	<0.001	<0.001
Observations	499	499	499	499

Notes: Table reports OLS estimates with standard errors in parentheses. Dependent and independent variables are type dummies.
*** p<0.001, ** p<0.01, * p<0.05

Table A.18: Panel: Transition Matrix between Types

Wave 1 / Wave 2	Libertarian	Meritocrat	Egalitarian	Type Other	Total W1
Libertarian	5	10	3	9	27
	18.52	37.04	11.11	33.33	100.00
Meritocrat	10	151	32	53	246
	4.07	61.38	13.01	21.54	100.00
Egalitarian	0	29	22	12	63
	0.00	46.03	34.92	19.05	100.00
Type Other	9	75	19	60	163
	5.52	46.01	11.66	36.81	100.00
Total	24	265	76	134	499
	4.81	53.11	15.23	26.85	100.00

Notes: Table reports transitions in fairness preference types between wave 1 and wave 2. The first number in each cell refers to the number of observations, the second to the share in each row.

Table A.19: Panel: Types Predict Policy Preferences Across Waves

	Support for Welfare Policies W2			Support for Welfare Policies W1		
	(1)	(2)	(3)	(4)	(5)	(6)
Libertarian W2	-0.65**		-0.50*		-0.52**	-0.41*
	(0.24)		(0.24)		(0.18)	(0.18)
Egalitarian W2	0.60***		0.54***		0.42***	0.38***
	(0.15)		(0.15)		(0.11)	(0.11)
Other W2	-0.15		-0.09		-0.12	-0.06
	(0.12)		(0.12)		(0.09)	(0.09)
Libertarian W1		-0.65**	-0.55*	-0.52**		-0.45**
		(0.23)	(0.23)	(0.17)		(0.17)
Egalitarian W1		0.46**	0.32*	0.29*		0.19
		(0.16)	(0.16)	(0.12)		(0.12)
Other W1		-0.24*	-0.22	-0.25**		-0.23**
		(0.11)	(0.11)	(0.08)		(0.08)
p (joint test W1)		<0.001	0.001	<0.001		
p (joint test W2)	<0.001				<0.001	<0.001
p (Lib W1 = Lib W2)			0.882			0.883
p (Ega W1 = Ega W2)			0.357			0.293
p (Oth W1 = Oth W2)			0.493			0.219
R^2	0.061	0.049	0.090	0.054	0.059	0.093
Observations	499	499	499	499	499	499

Notes: OLS Estimates with robust standard errors in parentheses. Omitted category: Meritocrats. In column (3) and (6), the reference category are subjects classified twice as Meritocrats. The joint test tests the hypotheses that all type coefficients in a wave are equal to zero. Constant not reported.
*** p<0.001, ** p<0.01, * p<0.05

Table A.20: Panel: Type Transitions and Policy Preferences

	Support for Welfare Policies				
	(1)	(2)	(3)	(4)	(5)
Egalitarian	0.01	0.01			
	(0.08)	(0.08)			
Libertarian	0.01		0.01		
	(0.15)		(0.14)		
Meritocrat				0.00	
				(0.06)	
Other	-0.01				-0.01
	(0.07)				(0.06)
Wave 2	-0.28***	-0.28***	-0.28***	-0.28***	-0.28***
	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)
Individual FE	Yes	Yes	Yes	Yes	Yes
Observations	998	998	998	998	998

Notes: Table reports coefficients from two-way fixed effects panel data models. Robust standard errors clustered at the individual level in parentheses. Constant not reported.
* p<0.05, ** p<0.01, *** p<0.001

A.1.4.4 Stability of Fairness Preferences: Experiment Wave 1

FIGURE A.11: Experiment Wave 1: Histograms of Transfer Choices by Treatment

(A) Transfer Choices Luck by Treatment

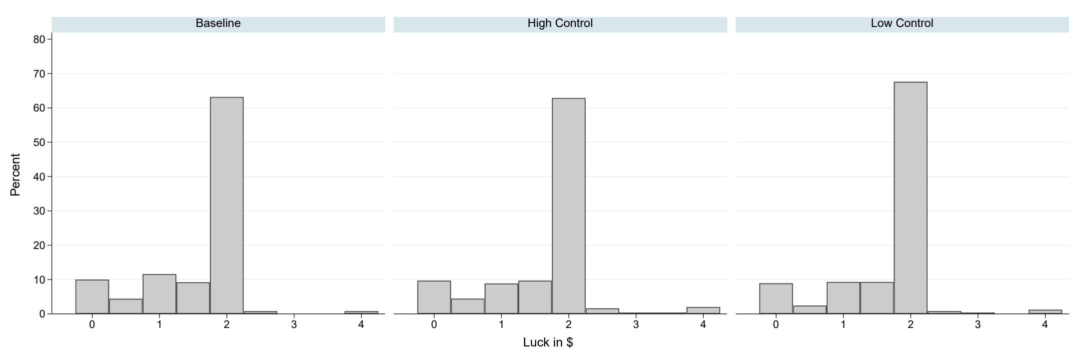

(B) Transfer Choices Merit by Treatment

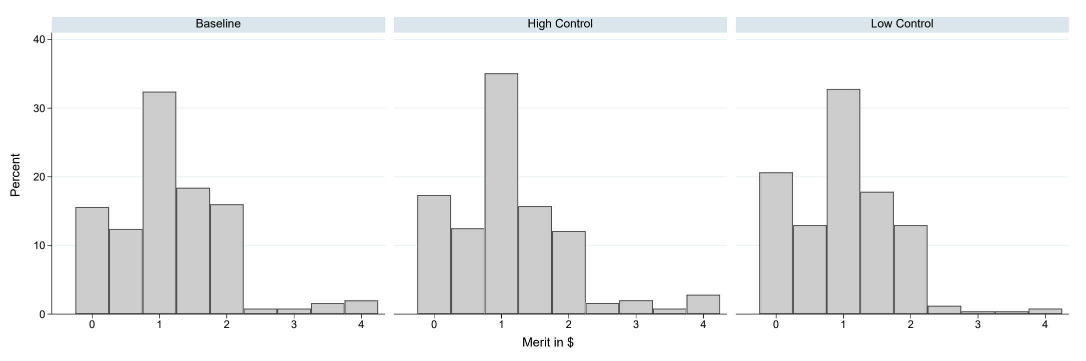

Notes: Figure shows histograms of transfer choices by treatment condition in our experiment in Wave 1 (N=745). Baseline (N=250), High Control (N=248) and Low Control (N=247). Bin size: $0.50.

TABLE A.21: Experiment Wave 1: Fairness Preference Types by Treatment

	Baseline	High Control	Low Control	Total
Libertarian	12	13	15	40
	4.8	5.2	6.1	5.4
Meritocrat	112	127	135	374
	44.8	51.2	54.7	50.2
Egalitarian	36	21	25	82
	14.4	8.5	10.1	11.0
Other	90	87	72	249
	36.0	35.1	29.1	33.4
Total	250	248	247	745
	100.0	100.0	100.0	100.0

Chi^2: p=0.054
Chi^2: p=0.323 (if excluding "Other")

Chi^2: p=0.542 (High Control = Low Control)
Chi^2: p=0.174 (High Control = Baseline)
Chi^2: p=0.092 (Low Control = Baseline)

Notes: Table reports counts and shares of fairness ideals by treatment condition. The first row reports the number of subjects per cell, the second row the share by treatment condition.

Table A.22: Experiment Wave 1: Treatment Effects on Transfer Choices

	Transfer Merit		Transfer Luck		Pr(Transfer Luck=$2)	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample						
Low Control	-0.124	-0.122	0.016	0.020	0.035	0.038
	(0.074)	(0.074)	(0.065)	(0.065)	(0.043)	(0.042)
Baseline	0.022	0.034	-0.070	-0.060	-0.005	-0.006
	(0.074)	(0.074)	(0.065)	(0.065)	(0.043)	(0.043)
Political Ideology	No	Yes	No	Yes	No	Yes
Socio-D.	No	Yes	No	Yes	No	Yes
Observations	745	745	745	745	745	745
Panel B: Excluding Outliers (Transfers >$2)						
Low Control	-0.008	-0.011	0.050	0.056	0.025	0.032
	(0.059)	(0.058)	(0.061)	(0.061)	(0.043)	(0.042)
Baseline	0.076	0.084	-0.018	-0.010	-0.027	-0.024
	(0.059)	(0.058)	(0.061)	(0.060)	(0.044)	(0.043)
Political Ideology	No	Yes	No	Yes	No	Yes
Socio-D.	No	Yes	No	Yes	No	Yes
Observations	706	706	721	721	721	721

Notes: Columns (1) to (4) report OLS estimates with robust standard errors in parentheses. Columns (5) and (6) report average marginal effects from logit models. Outliers (transfer choices >2$) are excluded from the analysis in Panel B. Socio-demographics contain age, gender, race/ethnicity, education, income bracket and employment status. Reference category: High Control.
*** p<0.001, ** p<0.01, * p<0.05

Table A.23: Experiment Wave 1: Treatment Effects on Fairness Preference Types

	Pr(Meritocrat)		Pr(Egalitarian)		Pr(Libertarian)		Pr(Other)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Low Control	0.034	0.035	0.017	0.014	0.008	0.007	-0.059	-0.061
	(0.045)	(0.044)	(0.026)	(0.025)	(0.021)	(0.021)	(0.042)	(0.041)
Baseline	-0.064	-0.072	0.059*	0.061*	-0.004	-0.011	0.009	0.017
	(0.045)	(0.044)	(0.028)	(0.028)	(0.020)	(0.019)	(0.043)	(0.043)
Political Ideology	No	Yes	No	Yes	No	Yes	No	Yes
Socio-D.	No	Yes	No	Yes	No	Yes	No	Yes
Observations	745	745	745	745	745	745	745	745

Notes: Table reports average marginal effects from logit models. Standard errors in parentheses. Socio-demographics contain age, gender, race/ethnicity, education, income bracket and employment status. Reference category: High Control.
*** p<0.001, ** p<0.01, * p<0.05

A.1.4.5 Stability of Beliefs

Panel Data: Changes in Beliefs over Time

Here we present a detailed comparison of the panel data models estimated. All estimates reported in Table A.24 are variants of the following panel data model:

$$\text{Meritocratic Belief}_{i,t} = \alpha_i + \beta_1 \text{Wave}_t + \beta_2 \text{Treatment}_{i,t} + \beta_3 X_{i,t} + \epsilon_{i,t} \tag{A.1.1}$$

In this model we are interested in the coefficient β_1 that indicates the difference in beliefs between Wave 1 and Wave 2. In Panel A, we restrict the analysis to the N=167 subjects in the Baseline condition, as they have not been exposed to any treatment manipulation in Wave 1. In Panel B, we use all N=499 observations, while controlling for treatment assignment in Wave 1. In Wave 2, all subjects are coded as not being exposed to any treatment, which is justified by the observation that treatment effects do not persist over time (see Figure A.12).

FIGURE A.12: Panel: Beliefs by Wave and Treatment

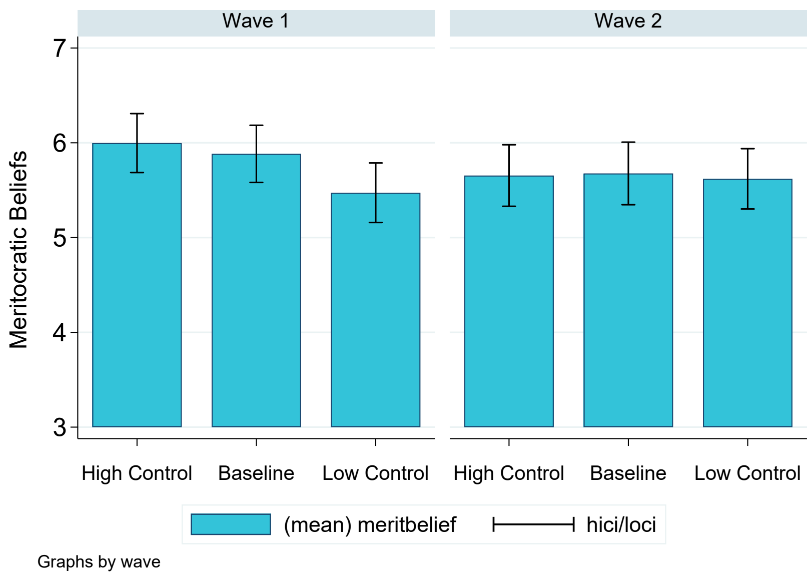

Notes: The figure shows the mean belief by wave and treatment for the N=499 subjects in our panel with 95% CIs.

In Panel B, the random effects model in column (1) and the fixed effects model in column (2) yield almost identical coefficients, even though they use very different types of variation. A Hausman test indicates that there is no statistically significant difference between these estimates (p=0.969) so that using the more efficient random effects model seems appropriate. Another argument for using the random effects model is that it allows to control for socio-demographic characteristics $X_{i,t}$. Most importantly, the random effects models allow to control for an age effect, which seems relevant, because all subjects in our sample have become 1 or 2 years older in between our waves of data collection.

The estimated coefficients for β_1 are similar in size, ranging from -0.21 to -0.26, and significant at the 5% level in all random effects models. Taken together, the results from our panel data provide evidence that beliefs decreased by about a quarter point on a 10-point Likert scale in our sample from Spring 2021 to Fall 2022.

Table A.24: Panel: Changes in Beliefs over Time

	(1)	(2)	(3)	(4)
Panel A: Subjects in Baseline				
Wave 2	-0.21	-0.21	-0.24*	-0.24*
	(0.11)	(0.11)	(0.11)	(0.12)
Age			0.03*	0.02*
			(0.01)	(0.01)
Model	FE	RE	RE	RE
Socio-D.	-	-	-	Yes
Observations	334	334	334	334
Clusters	167	167	167	167
p(Wave 2)	0.07	0.07	0.03	0.04
Panel B: All observations				
Wave 2	-0.21	-0.21*	-0.24*	-0.23*
	(0.11)	(0.11)	(0.10)	(0.11)
High Control	0.14	0.13	0.14	0.16
	(0.16)	(0.15)	(0.15)	(0.15)
Low Control	-0.35*	-0.37*	-0.35*	-0.34*
	(0.17)	(0.15)	(0.15)	(0.15)
Age			0.02***	0.02***
			(0.01)	(0.01)
Model	FE	RE	RE	RE
Socio-D.	-	-	-	Yes
Observations	998	998	998	998
Clusters	499	499	499	499
p(Wave 2)	0.07	0.04	0.02	0.03

Notes: Estimates from panel data models. Robust standard errors clustered at the individual level in parentheses. Socio-demographics contain gender, race/ethnicity, education, and income.
* $p<0.05$, ** $p<0.01$, *** $p<0.001$

A.1.4.6 Explaining Changes in Policy Preferences

Table A.25: Experiment Wave 1: Treatment Effects on Policy Preferences

	Redistribution		Univ. Health Care		Pandemic Support	
	(1)	(2)	(3)	(4)	(5)	(6)
Low Control	0.02	-0.02	0.05	0.01	-0.02	-0.03
	(0.09)	(0.07)	(0.09)	(0.07)	(0.08)	(0.07)
Baseline	0.09	0.05	0.08	0.02	0.02	-0.00
	(0.08)	(0.07)	(0.09)	(0.07)	(0.08)	(0.07)
Political Ideology	No	Yes	No	Yes	No	Yes
Socio-D.	No	Yes	No	Yes	No	Yes
Income Bracket	No	Yes	No	Yes	No	Yes
Employment St.	No	Yes	No	Yes	No	Yes
Observations	745	745	745	745	745	745
R^2	0.002	0.323	0.001	0.389	0.001	0.214

Notes: Table presents OLS estimates with robust standard errors clustered at the individual level in parentheses. Dependent variables: policy preferences as standardized z-scores. Socio-Demographics include age, gender, race, ethnicity, and education dummies.
*** p<0.001, ** p<0.01, * p<0.05

A.1.4.7 Panel Data: Experience Effects

To test whether objective personal experiences are meaningfully related to changes in fairness views, we use a Diff-in-Diff approach. For each personal experience, we specify an indicator variable equal to 1 if an individual made this experience between Wave 1 and Wave 2. As the control group, we specify subjects who did not make the respective personal experience. Where necessary, we restrict our analysis to the subgroup of individuals who has not yet experienced the shock in Wave 1 of the pandemic. These personal experiences are self-reported in our survey.

We estimate these experience effects using a standard two-way fixed effects model with two time periods. The identified effects are not necessarily causal, because these personal experiences may be correlated with other unobserved experiences or time-varying factors that also affect fairness views. Still, the individual-fixed effects can control for any time-invariant unobserved differences between the treated and control groups.

Personal Experience: COVID Case

We classify subjects who stated in Wave 2 that they or someone emotionally close to them had a case of COVID since June 2020 (N=204) as being treated by a *COVID case*. To identify those subjects who have already been treated prior to Wave 1 of our data collection, we use answers to the question in Wave 1, whether they or someone emotionally close to them has been tested positively with COVID (N=20).

Table A.26: Personal Experience: COVID Case

Prior / After W1	Not Treated	Treated	Total
Not Treated	286	193	479
	57.3	38.7	96.0
Already Treated	9	11	20
(Potentially)	1.8	2.2	4.0
Total	295	204	499
	59.1	40.9	100.0

Personal Experience: Severe COVID Case

We classify subjects who stated in Wave 2 that they or someone emotionally close to them had a severe case of COVID since June 2020 (N=78) as being treated by a *severe COVID case*. To identify those subjects who have already been treated prior to Wave 1 of our data collection, we use answers to the question in Wave 1, whether they or someone emotionally close to them has been tested positively with COVID (N=20).

Table A.27: Personal Experience: Severe COVID Case

Prior / After W1	Not Treated	Treated	Total
Not Treated	407	72	479
	81.6	14.4	96.0
Already Treated	14	6	20
(Potentially)	2.8	1.2	4.0
Total	421	78	499
	84.4	15.6	100.0

Personal Experience: Job Loss

We classify subjects who stated in Wave 2 that they or a household member lost a job or main source of income since June 2020 (N=78) as being treated by a *job loss*. To identify those subjects who have already been treated prior to Wave 1 of our data collection, we use answers to the question in Wave 1, whether they or a household member lost a job or main source of income due to the pandemic (N=129).

Table A.28: Personal Experience: Job Loss

Prior / After W1	Not Treated	Treated	Total
Not Treated	338	32	370
	67.7	6.4	74.1
Already Treated	83	46	129
(Potentially)	16.6	9.2	25.9
Total	421	78	499
	84.4	15.6	100.0

Personal Experience: Income Loss

We classify subjects who stated in Wave 2 that in the time period since June 2020 their household lost income compared to before the pandemic (N=89) as being treated by a *Income Loss*. To identify those subjects who have already been treated prior to Wave 1 of our data collection, we use answers to the question in Wave 1, whether their household lost income compared to before the pandemic (N=128).

Table A.29: Personal Experience: Income Loss

Prior / After W1	Not Treated	Treated	Total
Not Treated	333	38	371
	66.7	7.6	74.3
Already Treated	77	51	128
(Potentially)	15.4	10.2	25.7
Total	410	89	499
	82.2	17.8	100.0

Table A.30: Panel: Experience Effects on Fairness Views

	Transfer Merit in \$		Transfer Luck in \$		Beliefs in Merit	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: COVID Case						
Wave 2	0.03	0.04	0.10*	0.12*	-0.19	-0.15
	(0.05)	(0.05)	(0.05)	(0.05)	(0.12)	(0.12)
Treated × Wave 2	-0.01	-0.02	-0.09	-0.11	-0.03	-0.04
	(0.08)	(0.08)	(0.08)	(0.08)	(0.14)	(0.14)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	951	912	965	925	998	958
Panel B: COVID Case Severe						
Wave 2	0.03	0.04	0.07	0.08	-0.17	-0.13
	(0.04)	(0.04)	(0.04)	(0.04)	(0.12)	(0.12)
Treated × Wave 2	-0.01	-0.02	-0.04	-0.05	-0.20	-0.21
	(0.12)	(0.12)	(0.10)	(0.10)	(0.22)	(0.22)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	951	912	965	925	998	958
Panel C: Job Loss						
Wave 2	0.05	0.07	0.08	0.09	-0.22*	-0.18
	(0.04)	(0.05)	(0.04)	(0.05)	(0.11)	(0.13)
Treated × Wave 2	-0.24	-0.26	-0.21	-0.22	0.23	0.19
	(0.16)	(0.16)	(0.14)	(0.14)	(0.30)	(0.30)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	951	705	965	718	998	740
Panel D: Income Loss						
Wave 2	0.03	0.03	0.07	0.07	-0.20	-0.17
	(0.04)	(0.05)	(0.04)	(0.05)	(0.11)	(0.13)
Treated × Wave 2	0.05	0.04	-0.06	-0.07	-0.06	-0.10
	(0.13)	(0.13)	(0.13)	(0.13)	(0.23)	(0.24)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	951	712	965	722	998	742

Notes: Table presents estimates from two-way fixed effects models with robust standard error (clustered at the individual level) in parentheses. Columns (1), (3) and (5) include all observations. Columns (2), (4) and (6) exclude subjects that have potentially already been treated before Wave 1 of our data collection. For details on the classification of subjects see the previous pages. In Columns (1) to (4) outliers with transfer choices large \$2 are excluded. In Column (5) and (6) we additionally control for treatment status in wave 1 due to the observed treatment effect on beliefs in our experiment.
*** $p<0.001$, ** $p<0.01$, * $p<0.05$

A.1.4.8 Experiment in Wave 2 - Analysis

Experimental Design We run a baseline treatment in which subjects get no information and do not write about any personal experience (*Baseline*), one treatment in which subjects are provided with information about the impacts of the pandemic on US society (*Information*) and one treatment in which they are provided with the information and are asked to write about a personal experience of low control from the pandemic (*Low Control*). The information provided is the same as in the treatments in Wave 1 but with up-to-date statistics about the pandemic's health and labor market impacts.

Attrition and Balance Of the 786 subjects that are randomized into one of the treatments, 3 subjects do not complete the survey in *Baseline*, 12 subjects in the *Information* treatment, but in *Low Control* 42 subjects dropped out of the survey (Chi2-test: $p<0.001$). Most of these drop-outs in *Low Control* (32) happened when subjects were asked to write the text about a personal experience. One plausible explanation for this differential attrition is that subjects at this later point of the pandemic do not want to be reminded again of negative personal experiences that they made during the pandemic and hence leave the experiment.

Table A.31: Experiment Wave 2: Differential Attrition

	Baseline	Information	Low Control
Completed	277 98.9%	240 95.2%	212 83.5%
Drop-out	3 1.1%	12 4.8%	42 16.5%
Total	280 100.0%	252 100.0%	254 100.0%
Chi^2: $p<0.001$			

The **differential attrition across treatments** implies that the results of this experiment should be treated with much caution. While the treatment conditions still seem to be reasonably well balanced according to observable socio-demographic characteristics, except for education, (see Table A.32), the results may still be strongly biased due to unobserved differences between subjects across treatments. For that reason, we decided to exclude results of this experiment from the main part of our paper. For completeness, we report the results on the following pages.

Table A.32: Experiment Wave 2: Balance of Covariates

	Treatment Status		
Dependent Variable:	Low Control (1)	Information (2)	Baseline (3)
Female	-0.003	-0.041	0.042
	(0.034)	(0.035)	(0.037)
Age in decades	0.009	-0.014	0.004
	(0.011)	(0.011)	(0.012)
Black	-0.066	0.050	0.019
	(0.050)	(0.057)	(0.058)
Asian	0.012	-0.130*	0.127
	(0.072)	(0.059)	(0.076)
Race Other	-0.078	0.136	-0.056
	(0.072)	(0.084)	(0.077)
Hispanic	-0.108	0.037	0.065
	(0.070)	(0.066)	(0.068)
College Degree	-0.044	0.085*	-0.042
	(0.046)	(0.042)	(0.046)
Masters Degree	-0.143**	0.135**	0.010
	(0.048)	(0.049)	(0.053)
Income (in 100k)	-0.007	0.025	-0.019
	(0.042)	(0.043)	(0.045)
Unemployed	-0.000	0.039	-0.036
	(0.059)	(0.063)	(0.063)
Not in Labor Force	0.003	0.075	-0.077
	(0.041)	(0.045)	(0.044)
Midwest	-0.062	0.015	0.051
	(0.052)	(0.056)	(0.056)
South	0.000	-0.038	0.037
	(0.047)	(0.048)	(0.048)
West	-0.064	-0.057	0.121*
	(0.052)	(0.054)	(0.056)
Observations	729	729	729
joint significance (p-value)	p=0.123	p=0.113	p=0.440

Notes: Average marginal effects from logit models. Standard errors in parentheses.
*** p<0.001, ** p<0.01, * p<0.05

Table A.33: Experiment Wave 2: Fairness Preference Types by Treatment

	Baseline	Info	Low Control	Total
Libertarian	14	14	12	40
	5.1	5.8	5.7	5.5
Meritocrat	130	102	110	342
	46.9	42.5	51.9	46.9
Egalitarian	30	22	24	76
	10.8	9.2	11.3	10.4
Other	103	102	66	271
	37.2	42.5	31.1	37.2
Total	277	240	212	729
	100.0	100.0	100.0	100.0
Chi^2: p=0.344				
Chi^2: p=0.968 (if excluding "Other")				
Chi^2: p=0.088 (Info = Low Control)				
Chi^2: p=0.577 (Info = Baseline)				
Chi^2: p=0.576 (Low Control = Baseline)				

Notes: Table reports counts and shares of fairness ideals by treatment condition. The first row reports the number of subjects per cell, the second row the share by treatment condition.

Table A.34: Experiment Wave 2: Treatment Effects on Transfer Choices and Beliefs

	Transfer Merit		Transfer Luck		Beliefs in Merit	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample						
Info	-0.066	-0.087	-0.001	-0.002	0.015	0.013
	(0.080)	(0.077)	(0.068)	(0.068)	(0.191)	(0.175)
Low Control	-0.038	0.007	0.088	0.102	0.044	0.035
	(0.082)	(0.080)	(0.070)	(0.070)	(0.198)	(0.181)
Political Ideology	No	Yes	No	Yes	No	Yes
Socio-D.	No	Yes	No	Yes	No	Yes
Observations	729	729	729	729	729	729

Notes: Columns (1) to (6) report OLS estimates with robust standard errors in parentheses. Socio-demographics contain age, gender, race/ethnicity, education, income bracket and employment status. Reference category: Baseline. *** p<0.001, ** p<0.01, * p<0.05

A.2 Additional Materials

A.2.1 Timeline of Data Collection

FIGURE A.13: Timeline of Data Collection

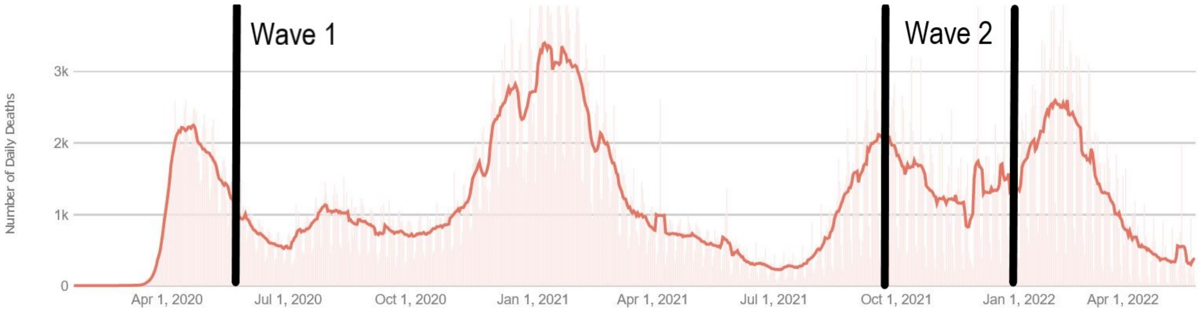

Notes: The figure shows a timeline of our data collection relative to the number of deaths related to COVID-19 in the US.

A.2.2 Instructions

Below we provide instructions for all three rounds of data collection: Wave 1, Wave 2 (Panel), Wave 2 (New Sample). Instructions shown to all participants are presented in black text. Differences between waves and additional questions in Wave 2 are presented in blue text. Instructions only shown to participants in Wave 1 are presented in red text. Instructions only shown to participants in Wave 2 (New Sample) are presented in green text. A dashed line indicates a page break. *Cursive words* at the beginning of a question refer to the variable names in our data set.

- -

Dear participant,

welcome to this research study! Please review the following consent form before proceeding with our survey.

DESCRIPTION: You will be asked questions about yourself, personal experiences and your opinions in relation to the coronavirus. Also, you can take decisions in two economic games. The survey will take approximately 12 minutes to complete.

PAYMENT: You will receive a guaranteed participation compensation of $1.40. Additionally, you will earn a bonus of $0 to $1.20, depending on the actions that you and other participants take. Please make sure that you click through to the end of the survey to be redirected to Prolific. We can only recompense participants who give answers to all questions and complete the last page of the study.

RISK AND BENEFITS: The risk to your participation in this online study are those associated with basic surveys including the recall of pleasant or unpleasant past experiences, such as illness and job loss, and mild stress. The benefit to you is the learning experience from participating in a research study. The benefit to society is the contribution to scientific knowledge.

SUBJECT'S RIGHTS: Your participation is voluntary. You have the right to see or withdraw your data at any time. Your responses will be recorded in a completely anonymous way. To secure the transparency of scientific findings, the completely anonymized data set will be published and made available to other researchers.

WARNING: This survey uses a protocol to check that you are responding from inside the U.S. and not using a Virtual Private Server (VPS), Virtual Private Network (VPN) or proxy to hide your country. In order to take this survey, please turn off your VPS/VPN/proxy if you are using one and also any ad blocking applications. Failure to do this might prevent you from completing the study. For more information on why we are requesting this, see this post from TurkPrime (https://goo.gl/WD6QD4).

YOU ARE NOT ALLOWED TO USE YOUR MOBILE PHONE.

If you have any questions about this project or if you have a research-related problem, you may contact the principle investigator: Maj-Britt Sterba, by email.
Please indicate, in the box below, that you are at least 18 years old, have read and understand this consent form, and you agree to participate in this research study.

I agree to participate in this research study

- -

[*Wave 2*: Please enter your Prolific ID in case it is not automatically displayed.]

- -

Thank you for your participation in this study! Please read the instructions carefully. You will not be able to go back after you have exited a page.

- -

Please answer the following questions about yourself. This information will only be used for statistical purposes. All your responses are anonymous.

Gender What is your gender? [Male; Female; Other]

Age What is your age? [Open textfield]

State In which state do you currently reside? [Drop-down menu]

County In which county or city county do you currently reside? [Open textfield]

Race Choose one or more races that you consider yourself to be: [White; Black or African American; American Indian or Alaska Native; Asian; Native Hawaiian or Pacific Islander; Other]

Ethnicity Are you Spanish, Hispanic, or Latino or none of these? [Yes; None of these]

Education What is the highest level of education you have completed?
[Less than High School; High School / GED; College Degree; Master's Degree]

Employment status What is your current employment status?
[Employed full-time (35+ hours a week); Unemployed and currently looking for work; Unemployed and not currently looking for work; Student; Retired; Homemaker; Self-employed; Other]

Income What was your family's gross household income in 2019 [*Wave 2*: 2020] in US dollars? [Less than $20,000; $20,000 to $34,999; $35,000 to $49,999; $50,000 to $74,999; $75,000 to $99,999; $100,000 to $149,999; More than $150,000]

- -

Political Ideology Please answer the following question about your political orientation by moving the slider below. In general I am,
[Slider between Liberal left and Conservative right]

- -

-- **Wave 1: Experiment** --

High Control and Low Control:

We would now like you to read this short text carefully.

The coronavirus continues to spread in the United States. Until today, there have been at least 1,300,000 cases and more than 84,000 deaths, according to data from Johns Hopkins University. All federal states have declared a state of emergency. Forty states closed down all non-essential businesses. Nationwide more than 36,000,000 people have lost their job since mid-March, according to the Department of Labour.

- -

High Control:

We are interested in your experience during the corona pandemic.

Please take one minute of time to write about a personal experience in the last two months in which you felt that you had **control over some aspect of your life.**

For example, did you perform a daily routine or exercise on a regular basis? Did you work on your home or garden? Did you take preventive measures to protect yourself?

Please describe the experience in as much detail as possible.

- -

Low Control:

We are interested in your experience during the corona pandemic.

Please take one minute of time to write about a personal experience in the last two months in which you felt that you had **no control or choice over what happened to you.**

For example, have you been restricted performing your job or going about your daily activities? Did you have to cancel important plans?

Please describe the experience in as much detail as possible.

- -

Baseline:

We would now like you to read this short text carefully.

Did you hear? The genome of the banana has been sequenced, an important development in scientist's efforts to produce better bananas. A look at that genome has revealed curious things, said Mat Peslop-Harrison, a plant geneticist at the University of Leicester in England who was a coauthor of the report published in the journal Nature. For example, there are regions of the banana genome that make them extra sweet and nutritious.

- -

Baseline:

We are interested in your experience with bananas.

Please take one minute to write about your last experience with eating bananas.

Please describe the experience in as much detail as possible.

------------------------------------- **Wave 1: End of Experiment** -----------------------------

---------------------------- **Wave 2 (New Sample): Experiment** ------------------

Low Control:

We would now like you to read this short text carefully.

The coronavirus continues to spread in the United States. More than 900,000 U.S. citizens were newly infected with the coronavirus in the past week, according to data from Johns Hopkins University. Many federal states have still declared a state of emergency. Since the start of the pandemic, more than 41,700,000 coronavirus cases have been recorded and more than 670,000 people in the U.S. have died after contracting the coronavirus.

Due to the outbreak of the pandemic, the U.S. labour market experienced the highest job losses since the Great Depression. Nationwide more than 20,000,000 people lost their job in the first months of the pandemic, according to the Department of Labor.

- -

Low Control:

We are interested in your experiences during the corona pandemic.

Please take one minute to write about a personal experience during the pandemic in which you felt that you had no control or choice over what happened to you.

Please describe the experience in as much detail as possible.

You can proceed to the next page after one minute has passed.

- -

Baseline: Passive control group (no text displayed)

------------------------- **Wave 2 (New Sample): End of Experiment** ----------------

Next we would like you to tell us how you feel right now.

Negative affect Which of these pictures best describes your current mood?

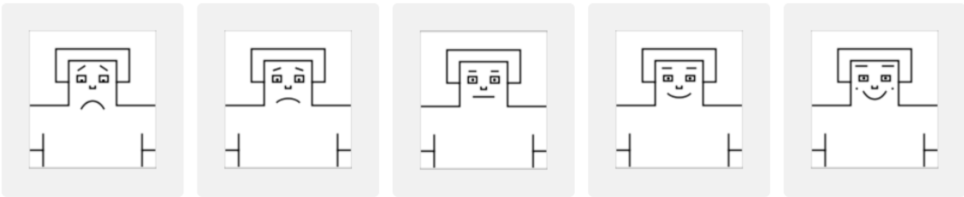

Stress To what extent are you feeling stressed at the moment? [Not at all 1; 2; 3; 4; 5; 6; Very much 7]

Fear To what extent are you experiencing the emotion fear at the moment? [Not at all 1; 2; 3; 4; 5; 6; Very much 7]

How close do you currently feel to:
Close 1 People in your country [5-item scale: Not close at all 1 – Very close 5]
Close 2 People in your local community [5-item scale: Not close at all 1 – Very close 5]

Now we would like you to answer some questions about your attitudes regarding personal and societal issues.

Control over life Sometimes people feel they have completely free choice and control over their lives, while at other times they feel that what they do has no real effect on what happens to them.

Please use this scale to indicate how much freedom of choice and control you feel you currently have in your life.
[10-item scale: No control at all 1 – A great deal of control 10]

Beliefs success How would you place your views on the following scale?

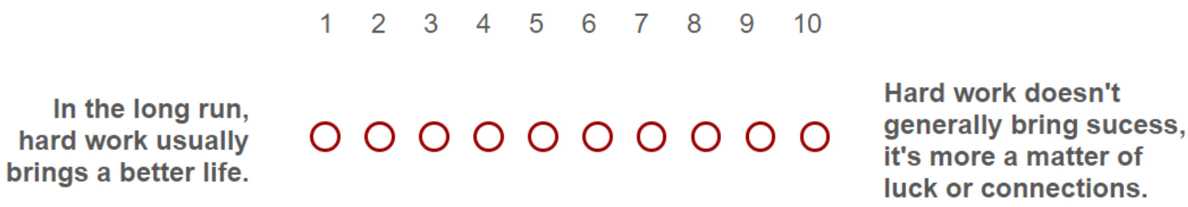

Note: 1 means you completely agree with the statement at the left; 10 means you completely agree with the statement at the right; and if your views fall somewhere in between you can pick any number in between.

Beliefs poverty How would you place your views on the following scale?

Note: 1 means you completely agree with the statement at the left; 10 means you completely agree with the statement at the right; and if your views fall somewhere in between you can pick any number in between.

[*Wave 2*:] In the next section, you will take decisions in two types of economic games, let's call them **Task 1** and **Task 2**. You will learn more about the tasks as you proceed with the survey. Please now proceed to the description of **Task 1**.

TASK 1:

In contrast to traditional survey questions that are often about hypothetical situations, your following decisions can have real consequences.

You will now take decisions that can change the earnings of other participants of this research study.

I understand that my decisions can change the earnings of other participants.

- -

A few days ago two individuals, let's call them "workers", have been recruited online to work on a tedious assignment. Both received a fixed participation compensation of $0.50.

After completing the assignment, they were told that their additional earnings for the assignment would be determined by one of three [*Wave 2:* two] rules. According to all three [*Wave 2:* both] rules one worker earns $4 and the other worker earns $0.

They were not informed about their outcome nor which rule applies. However, they were told that a third person would be informed about the assignment and the rules, and would be given the opportunity to redistribute the earnings and thus determine how much they would actually be paid for the assignment.

- -

You are the third person and we now want you to choose whether to redistribute the earnings between the workers.

You will take three [*Wave 2:* two] decisions, one for each rule that could apply. Each of the three [*Wave 2:* two] rules applies with equal probability. With 25% chance one of your decisions will be implemented.

Note: Your decisions are completely anonymous. The workers will receive their payment within a few days, but will not receive any further information.

- -

Rule #1 [randomized order]**:**

The workers' earnings are determined by their productivity. The more productive worker earns $4, and the other worker earns $0.

Merit If Rule #1 applies, how much of the earnings from the worker that earned $4 do you want to give to the worker that earned $0?

0 0.5 1 1.5 2 2.5 3 3.5 4

in $

Rule #2 [randomized order]**:**

The workers' earnings are determined by a lottery. The worker winning the lottery earns $4, and the other worker earns $0.

Luck If Rule #2 applies, how much of the earnings from the worker that earned $4 do you want to give to the worker that earned $0?

0 0.5 1 1.5 2 2.5 3 3.5 4

in $

[**Wave 1: Ambiguity Rule**]

Rule #3

With a certain probability, the workers' earnings are determined by a lottery. If there is a lottery, the worker winning the lottery earns $4, and the other worker earns $0. If there is no lottery, the more productive worker earns $4, and the other worker earns $0.

If Rule #3 applies, how much of the earnings from the worker that earned $4 do you want to give to the worker that earned $0?

- -

Probability When taking your previous decision: What probability that the earnings were determined by a lottery did you have in mind? [Slider from 0% to 100%]

- -

Before we continue, we would like you to briefly recall the experience that you have been writing about at the beginning of this study. To what degree did you feel in control in that situation?

Recall Control Please tell us: to what degree did you feel in control in that situation? [10-item scale: Not at all 1 – Very much 10]

- -

TASK 2:

You are now matched with one other participant of this study. Both of you have so far filled out the very same survey and will receive a fixed participation compensation of $1.40.

In this task you will take one decision that can change your bonus payment and the bonus payment for the other participant that you are matched with.

I understand that my decision can change my bonus payment and that of the other participant.

- -

You have been matched with another participant. One of you will get a bonus of $1.20, the other one will get no bonus. Who gets the bonus is determined by a lottery.

The one with the bonus of $1.20 can decide whether to give some amount of the bonus to the participant with no bonus.

Altruism In case that you win the lottery: How much of the bonus of $1.20 do you want to give to the other participant with no bonus?

0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1 1.1 1.2

in $

Note: Again, your decision is completely anonymous. The bonuses will be payed out within a few days. You will be informed whether your decision is implemented at the end of the survey.

- -

Now, we would like you to answer some questions about your attitudes in relation to the current coronavirus pandemic, but also to more general questions.

- -

Control beliefs How much control would you say people have over: [5-item scale: No control at all 1 – A great deal of control 5]

- Falling sick to the coronavirus
- Losing their job or main source of income due to the outbreak of the coronavirus
- Their health status in general
- Their financial situation in general
- Their life in general

- -

We will now ask for your attitudes towards measures the federal government and state governments have taken to address the outbreak of the coronavirus policy.

Pandemic support To what degree do you approve or disapprove of the following measures? [5-item scale: Strongly disapprove 1 – Strongly approve 5]

- Economic Impact Payment of 1200$ per Person
  [*Wave 2:* Economic Impact Payments]
- Increase and Expansion of Unemployment Benefits
- Expansion of Medicaid
- Paid Sick Leave

- -

Redistribution Generally, to what degree do you approve of economic redistribution? [Strongly Disapprove; Rather Disapprove; Neutral; Rather Approve; Strongly Approve]

Universal Health Care Generally, to what degree do you approve of universal health care? [Strongly Disapprove; Rather Disapprove; Neutral; Rather Approve; Strongly Approve]

- -

[**Wave 1: Experiences**]

Stay-at-Home Order To what degree do you approve of stay at home orders? [Strongly Disapprove; Rather Disapprove; Neutral; Rather Approve; Strongly Approve]

- -

In this section we would like you to fill out some information on your current financial situation. Remember, all your answers are anonymous.

- -

Job Loss Have you or a member of your household lost a job or main source of income due to the outbreak of the coronavirus? [Yes; No]

- -

Job Loss Peers How many people do you personally know (family, friends, neighbours, colleagues) that have lost a job or main source of income due to the outbreak of the coronavirus? [None; 1-2; 3-5; 6-10; More than 10]

- -

Income Loss How many percent of your regular household's income do you expect to lose this month due to the outbreak of the coronavirus compared to February 2020? [0%; 0% to 20%; 20% to 40%; 40% to 60%; 60% to 80%; 80% to 100%]

- -

In this last section we would like to ask about your personal exposure to the coronavirus (COVID-19).

- -

Covid Case Have you or a member of your family been tested positively for COVID-19? [Yes; No]

- -

Symptoms Have you or a member of your family showed symptoms associated with COVID-19 like fever, cough or difficulty breathing in the last two months? [Yes; Somewhat; No]

- -

High Risk Do you have a serious underlying medical condition that puts you at higher risk for severe illness from COVID-19? [Yes; No]

- -

Symptoms Peers How many people do you personally know (family, friends, neighbours, colleagues) that showed symptoms associated with COVID-19 like fever, cough or difficulty breathing in the last two months? [None; 1-2; 3-5; 6-10; More than 10]

- -

News How frequently have you been consuming information about the outbreak of the coronavirus? [More than 5 times a day; 4-5 times a day; 2-3 times a day; Once a day; Every other day; Once a week; Less than once a week]

- -

[**Wave 2: Experiences and Additional Variables**]

- -

Dear participant, on top of your fixed payment of $1.40 and your potential bonus from Task 2 we pay you a [$1] / [$0.50] bonus for your time and effort. In the following questions, we will ask about personal experiences that you made during the COVID-19 pandemic. You greatly help our research by providing us with this valuable information. Of course, as in the first part of the survey, all your answers are anonymous.

- -

In this section, we would like to learn about your financial situation.

- -

Job Loss 1 Please remember the first months of the coronavirus pandemic. Have you or has a member of your household lost a job or main source of income in the time period from February 2020 to May 2020? [Yes; No]

- -

Job Loss 2 Please now remember the time after the fist wave of the coronavirus pandemic. Have you or has a member of your household lost a job or main source of income since June 2020? [Yes; No]

- -

Income Change How did your gross monthly household income change in the following time periods compared to your gross household income in Februrary 2020 (before the COVID-19 pandemic)?

Note: Your gross household income includes any type of income before taxes (e.g. wages, self-employment income, rental income, retirement income) but excludes government transfers (e.g. unemployment benefits).

- March 2020 until May 2020 [Increased; Stayed the same; Decreased (by up to 20%; Decreased strongly (by 20% to 40%); Decreased drastically (by more than 40%)]

- June 2020 until today [Increased; Stayed the same; Decreased (by up to 20%; Decreased strongly (by 20% to 40%); Decreased drastically (by more than 40%)]

Unemployment Benefits Did your household receive **unemployment benefits** in the following time periods?

- March 2020 until May 2020 [Yes; No]
- June 2020 until today [Yes; No]

Transfers Please try to estimate roughly **how much** government transfers your household received in total in the following time periods.

Note: Government transfers include unemployment benefits, economic impact payments and any other public assistance or welfare payments.

- March 2020 until May 2020 [None; Less than 2,500$; 2,500$ to 5,000$; 5,000$ to 10,000$; More than 10,000$]
- June 2020 until today [None; Less than 2,500$; 2,500$ to 5,000$; 5,000$ to 10,000$; More than 10,000$]

Job Loss Peers Has someone you are emotionally close to (but who is not a member of your household) **permanently** lost his or her main source of income due to the COVID-19 pandemic? [Yes; No]

In this section, we would like to learn about your personal exposure to COVID-19.

Covid Case Have you or has someone emotionally close to you been *tested positively* for COVID-19? [Yes; No]

If *Covid Case* = yes:

Severe Covid Case Did you or someone emotionally close to you have a *severe case* of COVID-19? [Yes; No]

If *Covid Case* = yes:

Covid Case 2 You stated that you or someone emotionally close to you was tested positively for COVID-19. Were any of these tests carried out *since June 2020*? [Yes; No]

If *Severe Covid Case* = yes:

Severe Covid Case 2 You stated that you or someone emotionally close to you had a severe case of COVID-19. Did any of these severe cases of COVID-19 happen *since June 2020*? [Yes; No]

Overall exposure What would you say, how much have you been affected by the outbreak of the coronavirus compared to the average American? [Much less; Somewhat less; Somewhat more; Much more]

Now we would like to ask some more questions about yourself and your opinions on politics and society.

Marital status What is your current marital status? [Married; Living with a partner; Widowed; Divorced/Separated; Never been married]

Household size How many people currently live in your household? [1; 2; 3; 4; 5; 6 or more]]

- -

Economic Orientation On economic policy matters, where do you see yourself on the liberal/conservative spectrum? [Very liberal; Liberal; Moderate; Conservative; Very Conservative]

- -

Trust in Government How much of the time do you think you can trust the government to do what is right? [Never; Only some of the time; Most of the time; Always]

- -

Vote 2016 Whom did you vote for in the 2016 presidential elections? [Donald J. Trump; Hillary Clinton; Other; I did not vote]

- -

Vote 2020 Whom did you vote for in the 2020 presidential elections? [Donald J. Trump; Joseph Biden; Other; I did not vote]

- -

Social expansion During the COVID-19 pandemic the US has increased spending on social security. Do you think that the US should permanently increase spending on social security, that is, even after the pandemic? [Yes; No; No opinion]

- -

Taxing Rich Do you approve of increasing taxes for rich households to pay for a permanent expansion of the social security system? [Yes; No; No opinion]

- -

[Wave 2 (Panel):]

Covid Statements Now, we give you a few statements about how the COVID-19 pandemic might have changed people's views.

Please read through the statements and select all statements that you agree with. You can select as many statements as you like.

The COVID-19 pandemic made me realize that...
... economic inequalities and injustices are inevitable.
... it is unfair if people are in economic need due to no fault of their own.
... I might be in need of financial support at some point in the future.
... it is important to support one another in times of economic need.
... the government can't do much to reduce inequality.
... other [...]

Feedback: Is there anything you would like to tell us? This could relate to the topic of the survey, the ease of understanding of the questions or emotional strain that you felt while completing this survey.

All of your feedback is highly appreciated and helps us improve our research.

[*Open text box*]

Please Click the "Next" Button.

- -

Thank you very much for your participation in this research study!

[In Task 2 you have not won the lottery. Your bonus payment ($0 to $1.20) depends on the decision of another participant.] / [In Task 2 you have won the lottery and decided to give $X of you $1.20 bonus to the other participant.]
[*Wave 2:* You also earned a [$1] / [$0.50] bonus for your time and effort.]
You will receive your bonus payment within a few days.

You have to click the "Next" button at the bottom of this screen for your survey to be counted and to be redirected to Prolific.

If you have questions about this research, please contact the principle investigator, Maj-Britt Sterba, via email at sterba@coll.mpg.de.

Sincerely,
Maj-Britt Sterba
Max Planck Institute for Research on Collective Goods

Official Information on COVID-19: You can find official information from the US government here: https://www.usa.gov/coronavirus

Information about how to stay safe is provided by the Center for Disease Control and Prevention: https://www.cdc.gov/coronavirus

For frequently asked questions see: https://faq.coronavirus.gov

Chapter 3

Revealing Good Deeds: Disclosure of Social Responsibility in Competitive Markets

This chapter is based on the paper "*Revealing Good Deeds: Disclosure of Social Responsibility in Competitive Markets*" by Sören Harrs, Bettina Rockenbach, and Lukas Wenner, published in *Experimental Economics* in 2022, Vol. 25, pp. 1349–1373.[1]

As this paper is published, it has been replaced with a bibliographical reference in the printed version of this dissertation. The paper can be accessed online via the following link: https://doi.org/10.1007/s10683-022-09752-z

Abstract We experimentally study competitive markets with socially responsible production. Our main focus is on the producers' decision whether or not to reveal the degree of social responsibility of their product. Compared to two benchmark cases where either full transparency is enforced or no disclosure is possible, we show that voluntary and costless disclosure comes close to the full transparency benchmark. However, when the informational content of disclosure is imperfect, social responsibility in the market is significantly lower than under full transparency. Our results highlight an important role for transparent and standardized information about social externalities.

[1]Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1– 390838866 and by the Center for Social and Economic Behavior (C-SEB). We thank Jana Friedrichsen, Dorothea Kübler, Adriaan Soetevent and conference participants at ESA Europe 2017 (Vienna) for helpful comments. Viet A Nguyen and Hanan Iqbal provided valuable help in running the experiments. The replication material for the study is publicly available at `https://doi.org/10.17605/OSF.IO/KZ2FH`.

Chapter 4

Identity and Voluntary Efforts for Climate Protection

This chapter is based on the paper "Identity and Voluntary Efforts for Climate Protection" by Marvin Gleue, Sören Harrs, Christoph Feldhaus, and Andreas Löschel, published in the *Journal of Economic Behavior & Organization* in 2024, Vol. 221, pp. 436-476.[1]

As this paper is published, it has been replaced with a bibliographical reference in the printed version of this dissertation. The paper can be accessed online via the following link: https://doi.org/10.1016/j.jebo.2024.03.029

Abstract Can voluntary contributions to public goods be motivated by identity concerns? In a theory-driven field experiment, we test how positive and negative shocks to subjects' environmental identity beliefs affect voluntary efforts for climate protection. In a real-effort task, subjects can generate donations that offset carbon emissions. Prior to the task, we manipulate subjects' beliefs about their environmental identity either positively or negatively compared to a control group. A negative shock to identity ('identity threat') increases effort by about 17 % compared to our control group. This effect is largest for subjects that had a strong prior environmental identity belief. We find no evidence that a positive shock to identity does affect behavior. Our results are in line with some of the main predictions from the belief-based model of identity by Bénabou and Tirole (2011). They also have implications for policymakers and NGOs that want to encourage voluntary contributions to climate protection by leveraging people's identity concerns.

[1]We gratefully acknowledge the support by the organizers of the German Protestant Church Assembly and thank Yan Chen, Bernd Irlenbusch, Arno Riedl, Bettina Rockenbach, Matthias Sutter, Sebastian Tonke, and Peter Werner for very helpful comments. We are indebted to our research assistants for essential help when conducting the experiment. The experiment has been pre-registered in the AEA Social Science Registry as AEARCTR-0004335 https://www.socialscienceregistry.org/trials/4335 including a pre-analysis plan. This work was supported by a grant from the Federal Ministry of Education and Research, Germany, as part of the project "NostaClimate" (project number 01LA1813E). Sören Harrs acknowledges support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany´s Excellence Strategy – EXC 2126/1– 390838866. Declarations of interest: none. Replication files for this paper are publicly available https://doi.org/10.17605/OSF.IO/BDMVT.

Chapter 5

How Narratives Influence Economic Decision-Making: Experimental Evidence

This chapter is based on the paper "*How Narratives Influence Economic Decision-Making: Experimental Evidence*" by Lara Marie Berger, Sören Harrs, and Bettina Rockenbach. Earlier versions of this paper have been published as ECONtribute Working Paper No. 091 and via SSRN under different titles (Link).[1]

Abstract. The strategic provision of information is a powerful way to influence economic decisions. Standard theories explain the persuasiveness of information with its effects on people's beliefs and expectations. In this paper, we provide experimental evidence that information can influence economic decisions through a second mechanism: through short-term effects on people's risk and time preferences. In an information provision experiment, subjects read news articles that either contain an optimistic, a pessimistic, or a balanced narrative about the exogenous shock of the coronavirus pandemic. Reading a more pessimistic narrative leads to more pessimistic expectations about the pandemic and the stock market. However, it also strongly increases subjects' risk aversion and reduces their patience in incentivized experimental games (by around 30-40% of a standard deviation). Our results provide evidence that narratives can influence economic decision-making through two persuasion mechanisms: belief-based persuasion and preference-based persuasion. Understanding these mechanisms seems relevant for managers, investors, and politicians that engage in persuasive communication.

[1]We thank Matthias Sutter, Christopher Roth, Johannes Münster, Sebastian Tonke, Eugenio Verrinia, Susanna Grundmann, Lukas Reinhardt and participants at several conferences for very helpful comments. The experiment has been preregistered in the AEA Registry as AEARCTR-0005795 (https://doi.org/10.1257/rct.5795-1.0). Replication files for this paper are publicly available (https://doi.org/10.17605/OSF.IO/NJ2SQ). We gratefully acknowledge that this project has received funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany´s Excellence Strategy – EXC 2126/1– 390838866. IRB approval has been obtained from the German Association for Experimental Economic Research. Declarations of interest: none.

5.1 Introduction

The strategic provision of information is a powerful way to influence economic decisions. Standard theories in most fields of economics explain the persuasiveness of information with its effects on people's beliefs and expectations. Thus, they are theories of "belief-based" persuasion (DellaVigna and Gentzkow, 2010). Models of media markets, for example, describe how news outlets select and supply information to change the beliefs that news consumers have about the world (e.g. Gentzkow et al., 2015; Mullainathan and Shleifer, 2005).[2]

To test these theories, a recent and rapidly growing empirical literature in economics makes use of information provision experiments (see Haaland et al., 2023, for an excellent review). In information provision experiments, economists typically provide subjects at random with objective statistics, and then study their causal effects on people's beliefs and behaviors.[3] In the news media, in business communication, and in politics, however, much information is not provided in statistics, but rather in qualitative stories and narratives. Yet, studies that experimentally investigate the influence of stories and narratives are still very scarce (Haaland et al., 2023). This gap in the literature raises two major questions: First, whether narratives have an empirically relevant influence on economic-decision making, as prominently argued by Shiller (2017). And second, whether narratives influence economic decision-making through the same mechanism as statistics.

In this paper we address these two questions by providing experimental evidence on how narratives influence economic decision-making. In our main experiment, subjects read naturalistic news articles that either contain an optimistic, a pessimistic, or a balanced narrative about the exogenous shock of the coronavirus pandemic. These articles do not contain any statistics. Instead, they contain narratives, that is, qualitative stories that describe a series of events and explain their causal relationship.[4] We then study the effects of these narratives on two distinct sets of outcomes: (i) on subjects' forward-looking expectations about the pandemic and the stock market (measured in an incentivized forecasting task), and (ii) on risk and time preferences (measured in incentivized experimen-

[2] Models in industrial organization, on the other hand, describe how firms strategically provide information to change what consumers belief about their products (e.g. Nelson, 1974; Stigler, 1961). Models of persuasive communication in general focus on the strategic choice of an information signal to change the beliefs of the audience (e.g. Galperti, 2019; Kamenica and Gentzkow, 2011).

[3] For example, Armantier et al. (2016) and Coibion et al. (2019) provide subjects with statistics about inflation rates. In the context of the coronavirus pandemic, papers have investigated how providing statistical information about the coronavirus, e.g. its transmission rate, impacts macroeconomic expectations (Binder, 2020) or expectations about the pandemic (Fetzer et al., 2021; Rafkin et al., 2021).

[4] While the economic literature has not yet converged on one consistent definition of narratives, most definitions share the view that a narrative is a "story that describes a series of events and explains their causal relationship" (e.g. Barron and Fries, 2023; Eliaz and Spiegler, 2020; Schwartzstein and Sunderam, 2021; Shiller, 2017). According to Shiller (2017), for example, a narrative is a story or explanation of events that can be used to "stimulate the emotions or concerns of others".

tal games). This allows us to test whether narratives influence economic-decision making through a "belief-based persuasion" mechanism, or potentially also through "preference-based persuasion".[5]

Our main experiment (N=423) took place in Germany in May 2020, that is, after the first wave of infections and the first lockdown of the pandemic. Subjects are provided with one of four narratives. The optimistic narrative explains that after opening-up the economy, new COVID-infections will continue to decrease and then remain at a low level. It further raises the expectation that the economy will quickly recover once the political restrictions are lifted, and thus describes a V-shaped recovery. The pessimistic narrative instead explains that opening-up the economy will cause a second wave of infections, which will prove much more deadly than the first one. It also raises the concern that a second wave will cause a renewed lockdown with severe impacts on the economy, and thus describes a W-shaped recovery.[6] The balanced narrative combines elements of the optimistic and the pessimistic narrative. In a baseline condition subjects read a science-related article that is unrelated to the pandemic.

As our first main result, we show that exposure to these narratives about the pandemic changes people's forward-looking expectations. Subjects who read the pessimistic narrative expect on average 6.8% more COVID-related deaths over the following two months than subjects who read the optimistic narrative. Hence, the narratives change expectations about the future course of the pandemic. But can these narratives also change important economic expectations, such as expectations about the development of the stock market? Our data show that subjects in the pessimistic treatment expect that the German stock market index DAX has on average a 478 points (or 4.2%) lower value in two months than subjects in the optimistic treatment. This corresponds to a relatively large effect size for an information provision experiment of around 25% of a standard deviation. Hence, narratives about exogenous shocks prove to be very powerful in influencing important economic expectations.

As our second main result, we find that narratives also have a strong influence on subjects' risk and time preferences measured in incentivized economic games. The more pessimistic the narrative, the more risk-averse and impatient subjects become in our experiment. Subjects who read the optimistic narrative, on average, decide in a risk-neutral way (their average certainty equivalent for a lottery that pays out 4 EUR and 0 EUR

[5] According to DellaVigna and Gentzkow (2010) the term "preference-based persuasion" refers to any form of persuasion that affects behavior independent of beliefs; It can refer to cases where a persuasive message directly affects the utility function or to psychological models of persuasion, where for example emotions may play a role.

[6] For examples of similar news articles at that time, see The New York Times on the future course of the pandemic (https://www.nytimes.com/2020/05/08/health/coronavirus-pandemic-curve-scenarios.html) and Forbes Magazine on shapes of the economic recovery (https://www.forbes.com/sites/phillipbraun/2020/05/28/the-shape-of-economic-recovery/) (both accessed on December 6, 2021).

with 50% probability is EUR 2). Subjects who read the pessimistic narrative, in contrast, reveal a substantial level of risk aversion (their average certainty equivalent is EUR 1.71). This difference in risk preferences is sizeable and corresponds to 41.1% of a standard deviation. The influence of narratives also extends to the domain of time preferences. Subjects in the pessimistic treatment are on average much more likely to choose a payment of 2 EUR today rather than a higher payoff in the future. The difference in time preferences corresponds to 30.6% of a standard deviation, and is thus comparable to the difference in risk preferences. To illustrate this difference: the share of subjects that prefer a payoff of 2 EUR today to a payoff of 3.32 EUR in two months is 18 percentage points higher in the pessimistic treatment than in the optimistic treatment. Hence, after reading the pessimistic narratives, subjects are substantially more risk averse and impatient. Taken together, our results thus provide evidence that exposure to narratives has the potential to cause short-term changes in people's risk and time preferences.

We then provide evidence why narratives can influence people's risk and time preferences. Our data show that narratives can provoke strong emotional reactions in the audience and can instill a general sense of optimism or pessimism in people. This is in line with the argument by Shiller (2017) that one of the main reasons why narratives are often used in persuasive communication is because they can "trigger the emotions or concern" of the audience. When exposed to a more pessimistic narrative about the pandemic, subjects in our experiment feel much more afraid, upset, and nervous, and they are less optimistic about the development of their own personal circumstances. Emotions and optimism are correlated with risk aversion and impatience in our economic games. This suggests that the affective reactions towards narratives are responsible for the observed effects on risk and time preferences, and thus are necessary to activate the "preference-based" persuasion mechanism of narratives.

To provide additional evidence about the mechanism and the generalizability of our results, we run a pre-registered follow-up experiment (N=393). The motivation for this follow-up is to test whether all narratives persuade through the belief-based and the preference-based mechanism, or only those narratives that cause strong affective reactions in the audience. To do so, we provide subjects with either an optimistic or a pessimistic narrative about how a technology shock (the development of ChatGPT) will impact the business model of a company (Google). We chose these narratives because we anticipated that they would not cause strong affective reactions in our subject pool. We find that subjects who read the optimistic narrative have 6.1% higher forward-looking expectations about Google's stock market value in two months (which corresponds to 36% of a standard deviation). But, as pre-registered, these narratives do neither cause strong affective reactions nor any systematic differences in subjects' risk and time preferences. Thereby, the results from our follow-up support the interpretation that narratives only persuade through the "preference-based" mechanism if they cause affective reactions in the audience.

Our paper mainly contributes to two strands of literature. One the one hand, we contribute novel empirical evidence to the emerging literature on narratives in economics (Aina, 2023; Andre et al., 2023, 2022; Barron and Fries, 2023; Eliaz and Spiegler, 2020; Schwartzstein and Sunderam, 2021; Shiller, 2017).[7] Among the related empirical papers, Andre et al. (2023, 2022) document the mental models and narratives that people have on top of their minds about the macroeconomy and about inflation.[8] Barron and Fries (2023) study the strategic provision of narratives in an abstract laboratory game. Our paper, in contrast, studies experimentally how narratives in naturalistic news articles influence economic decision-making. Our experimental results, first of all, confirm that narratives have an important influence on economic decision-making as prominently argued by Shiller (2017). Our data also provide experimental evidence for the main persuasion mechanism in existing theoretical models on narratives: narratives are an effective tool of persuasion because they can change people's beliefs and expectations (Aina, 2023; Eliaz and Spiegler, 2020; Schwartzstein and Sunderam, 2021). However, as our main contribution, we show that narratives can be an effective tool of persuasion through a second - yet neglected - mechanism: through preference-based persuasion.

On the other hand, we contribute to a growing body of literature in economics on the (in)stability of risk and time preferences (e.g. Chuang and Schechter, 2015; Cohn et al., 2015; Schildberg-Hörisch, 2018).[9] Our paper contributes to this literature by providing experimental evidence for one mechanism that can help explain where instabilities in risk- and time preferences can come from: from being exposed to narratives about major exogenous shocks in the media. Our results also imply that managers, investors, and politicians could strategically provide narratives to make their audience behave in a more (or less) risk averse and impatient way. Hence, narratives may be used as a persuasive tool through which opinion leaders could try to strategically influence people's risk and time preferences.

Our paper is structured as follows: In Section 5.2, we describe our experimental design and the data of our main experiment. Section 5.3 presents our main results. Section 5.4 discusses our results, and Section 5.5 concludes.

[7] One strand of literature also investigates the role of narratives in the political and moral domain (Barron et al., 2020; Bénabou et al., 2020; Bursztyn et al., 2020; Verrina and Hillenbrand, 2020).

[8] In a new experimental part of the paper, Andre et al. (2023) also provide subjects with narratives about the causes of high inflation rates. In contrast to our study, the authors study backward-looking narratives about a past event and focus on expectations about inflation as the main outcome.

[9] In this literature, a number of papers have for example documented instabilities in risk- and time preferences over the course of the pandemic using panel data (Harrison et al., 2022; Huber et al., 2021; Li et al., 2022; Shachat et al., 2021a,b).

5.2 Experimental Design and Data

Our main experiment was conducted in Germany during the first wave of the COVID-19 pandemic with N=423 subjects, recruited from the subject pool of the Cologne Laboratory for Economic Research (CLER) via ORSEE (Greiner, 2015). The experiment was implemented with the survey software Qualtrics. The median time for completing the experiment was 15 minutes. Subjects were paid dependent on their economic decisions with an average of EUR 6.21. Payments were made via PayPal. The experiment has been pre-registered in the AEA Social Science Registry as AEARCTR-0005795.[10]

5.2.1 Setting

When we conducted our main experiment, on May 4th 2020, Germany had just lived through six weeks of strict political measures to combat the spread of COVID-19. The set of political measures that were in place since March 23rd 2020 contained, among others, the closure of schools, kindergartens and all non-essential businesses, strict rules of social distancing in public spaces and the prohibition of public gatherings of more than two persons living in different households. It was a wide-spread consensus that these measures had caused the reduction in the number of daily new cases in the weeks prior to the experiment (see Appendix Figure D.5 for a timeline of the pandemic in Germany). Since mid April, a public discussion about lifting the restrictions and re-opening the economy had started in the media and among scientists and politicians.

5.2.2 Experimental Procedures

Figure 5.1 provides a graphical overview of the experimental procedures.[11] Numbers in brackets in this section refer to the stages of the experiment depicted in Figure 5.1. At the beginning of the experiment, subjects were exposed to an article and were incentivized to memorize it as good as possible within two minutes (2).[12] Later in the experiment, subjects faced three questions about the content of the article (6,8,10): for each correct answer subjects were payed EUR 0.50. By incentivizing the careful reading of the article, we made sure that our subjects were sufficiently exposed to our manipulation. Our experimental manipulation is on the provided article. We study four different articles: containing either an optimistic, a pessimistic or a balanced narrative about the pandemic or a science related baseline article not related to COVID-19. Each subject saw and was aware of only one article. See Section 5.2.3 for details on the manipulation.

[10] You can find the pre-registration here: `https://doi.org/10.1257/rct.5795-1.0`.

[11] Transcripts of the instructions are provided in Appendix D.2.4.

[12] After two minutes, subjects were automatically directed to the next page. Subjects could not proceed to the next page independently.

Figure 5.1: Experimental Procedures

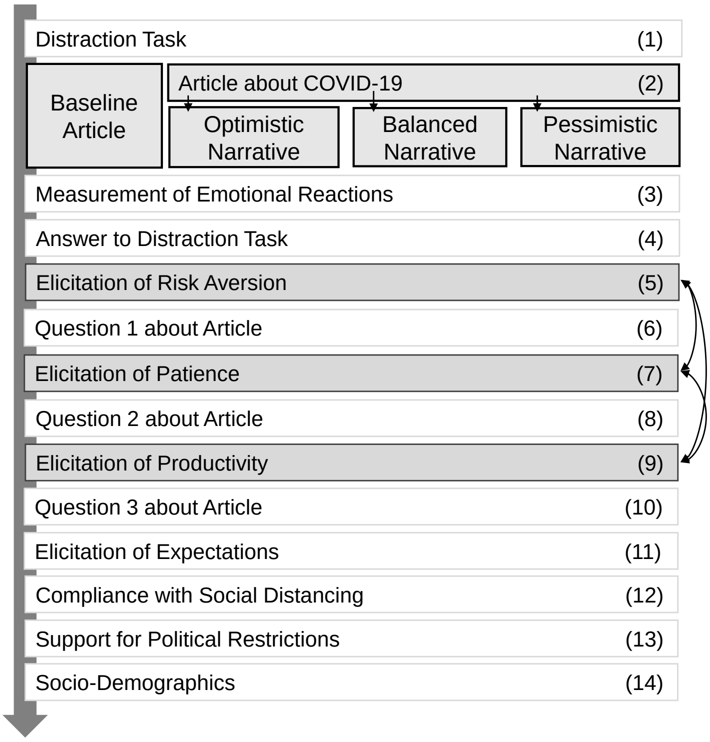

Notes: Figure 5.1 gives an overview of the experimental procedures. The numbers on the right side refer to the different stages of the experiment. The manipulation and the main outcomes are shaded in grey. The order of the elicitation of risk aversion, patience and productivity was randomized.

Immediately after the manipulation, we measured the emotional reactions of subjects (3). Next, we elicited our behavioral outcomes risk aversion (5), patience (7) and productivity[13] (9) in three decision blocks. At the end of the experiment, one of the three decision blocks was randomly drawn for each subject and became payoff relevant. We randomized the order of the three behavioral outcomes to be able to control for order effects. After the behavioral outcomes, we elicited subjects' forward-looking expectations for the pandemic, their personal circumstances and the stock market (11). The final part of the experiment included questions on compliance with social distancing (12) and on support for political restrictions (13). The experiment concluded with collecting the socio-demographic characteristics of subjects (14).

At the very beginning of the experiment, we implemented a distraction task to preempt concerns about experimenter demand effects. Subjects were asked to memorize two phone numbers (1) which they had to recall (4) before we elicited the main outcomes. In case

[13] We included productivity as a third behavioral outcome because of anecdotal reports about reduced productivity in the pandemic. A reduction in productivity may provide one alternative behavioral mechanism through which narratives could impact the economic decision-making of households. We hypothesized that subjects may not be able to focus on the financially incentivized real-effort task (or lack the motivation to do so), after being exposed to bad news about the pandemic in the form of the pessimistic narrative.

some subjects did try to anticipate what our study was about, this task (together with the text memory task in our manipulation) should have created the impression that this study was most likely about working memory ability. This design feature was implemented even though experimenter demand effects have been shown to be only a modest concern in a variety of settings (de Quidt et al., 2018). Concerns about experimenter demand effects should also be mitigated by the design choice to incentivize the elicitation of expectations and of the behavioral outcomes. In that way, it would be quite costly for subjects to deviate from their own expectations or preferred choices in order to comply with experimenter demands.

5.2.3 Manipulation

Participants were randomly assigned to one of four conditions: subjects in the three treatment conditions read an article that provides an optimistic, a balanced or a pessimistic narrative about the COVID-19 pandemic in Germany; subjects in the baseline condition read a science-related article not related to COVID-19. As far as possible, all articles were designed symmetrically regarding their content, length, structure and grammatical style (see Appendix D.2.2 for the transcripts). All narratives about COVID-19 follow a common structure (see Appendix Figure D.6 for details). The narratives describe how opening-up the economy will have causal impacts on the pandemic, the health care system and the economy.

The optimistic narrative explains that after opening-up, new cases will decrease and then remain low (a "one wave" model of the pandemic). It further raises the expectation that the economy will quickly recover once the political restrictions are lifted (V-shaped recovery). The pessimistic narrative instead explains that opening-up will cause a second wave of infections, which will prove much more deadly than the first one (a "two wave" model of the pandemic). It also raises concerns that a second wave will cause a renewed lockdown with severe impacts on the economy (W-shaped recovery). Thereby, the narratives describe different models of the pandemic and of the shape of the economic recovery.[14] The balanced narrative combines elements of the optimistic and the pessimistic narrative. The baseline article covers a story about outer space, structured in an analogous fashion.

All articles in the three treatment conditions are complemented with a figure that sketches the future development of daily new infections in line with the respective narrative. Such epidemic curves and curve scenarios have been widely used in news reporting

[14] For examples of similar news coverages at that time, see The New York Times on pandemic models (https://www.nytimes.com/2020/05/08/health/coronavirus-pandemic-curve-scenarios.html) and Forbes Magazine on shapes of the economic recovery (https://www.forbes.com/sites/phillipbraun/2020/05/28/the-shape-of-economic-recovery/) (both accessed on December 6, 2021).

in countries around the world to visualize the outbreak of the pandemic.[15] These curves of daily new infections were arguably the most important graphics driving sentiments of the general public towards the pandemic. Therefore, we decided to make use of them in our experimental manipulation. The baseline article uses a similar figure unrelated to the pandemic.

5.2.4 Measurement

Expectations We elicit incentivized 2-month forward-looking expectations on the German stock market index DAX, the total number of COVID-19 cases and the total number of deaths related to COVID-19. To anchor our subjects' estimates, we provide official data for each of these variables from the previous day. We incentivize the expectations in the following way: for each variable three subjects are randomly selected and are paid depending on the accuracy of their expectations (with up to EUR 20). We did not disclose the exact payment formula in more detail to subjects. Each subject receives at most a payoff for one of the expectations. This incentive scheme has two noteworthy properties: (i) subjects cannot hedge risk between expectations, and (ii) the game is non-strategic (the expected payoff is independent from the expectations of the other subjects).

Risk aversion is measured as the certainty equivalent for a lottery that pays EUR 4 with 50% probability and EUR 0 with 50%. The certainty equivalent is elicited using a staircase method for risk preferences similar to Falk et al. (2018). Subjects face five consecutive choices between a fixed payment and a lottery that pays EUR 4 with 50% probability and EUR 0 with 50% probability. The amount offered as fixed payment changes from decision to decision: if a subject chose the lottery (the safe payment), the safe payment offered in the next round is increased (reduced). The game tree is provided in Appendix Figure D.7. One of the five decisions is randomly chosen for payment. The certainty equivalent can take 32 values ranging from EUR 0.10 to EUR 3.20.

Patience is measured with an equivalent staircase method for time preferences similar to (Falk et al., 2018). The outcome variable for patience is the future value. The future value indicates the point at which subjects are indifferent between receiving a payment of EUR 2 today and receiving a payment of the future value in 60 days. This time subjects take five consecutive decisions between a payment of EUR 2 today and a payment in 60 days. The payment in 60 days changes from decision to decision: if a subject chose the payment today (in 60 days), the payment in 60 days offered in the next decision round is increased (reduced). The game tree is provided in Appendix Figure D.8. Again, one of

[15]See, for example, this article of The New York Times that uses curve scenarios in a similar way https://www.nytimes.com/2020/05/08/health/coronavirus-pandemic-curve-scenarios.html

the five decisions is randomly chosen for payment. The future value elicited in this game can take 32 values ranging from EUR 2.08 to EUR 4.56.

We measure risk aversion and patience with the staircase method because it allows to elicit fine-grained certainty equivalents and future values in a much more time efficient way compared to classical Multiple Price Lists. Further, it prevents inconsistent choices (multiple switching points) by design and it does not require extensive instructions.[16] As a robustness check, we show in Appendix Table D.3 that our results reproduce if we just analyze the first choice of subjects in each staircase method.

Productivity is measured in a real-effort task: subjects have to count the digit "1" in lines of twelve to fourteen symbols. Subjects have two minutes time to complete as many lines as possible (up to 37). For each correct line subjects are paid EUR 0.10. The design of the task is inspired by a concentration test.

Emotions and Optimism Emotions are measured with 6-items of the i-PANAS-sf scale (Thompson, 2007), which is widely used in psychological research. We elicit three items for positive affect (attentive, determined, inspired) and three items for negative affect (upset, afraid, nervous). Subjects are asked to state the intensity with which they currently experience the respective emotion on a 5-point Likert scale (1 "not at all" to 5 "very much") for each of the six items. Affect is then constructed as the sum of the positive items minus the negative items. As a measure of subjects' *personal optimism* we ask subjects to indicate how they expect their personal circumstances to develop over the next weeks on an 11-point Likert scale from (-5 "very negative" to +5 "very positive").

5.2.5 Sample and Randomization Check

Of the 425 participants that started the experiment only two did not complete it. Hence, there was no considerable attrition. A table of sample characteristics by treatment condition is provided in Appendix Table D.1. We present tests for the pairwise balance of covariates between any two treatment conditions in Appendix Table D.2. For each covariate we conduct either t-tests or Chi^2 tests. Among the 21 tests conducted between the optimistic, pessimistic and balanced treatment, just one test is significant at the 5% level, as should be expected by chance. The imbalance stems from a slightly higher share of non-students in the optimistic treatment (11.4%) compared to the pessimistic treatment (3.8%). Note that this slight imbalance can only be due to chance as we randomized by computer and there was close to no attrition. We address this imbalance as follows: in

[16] In contrast to the Dual Multiple Price Lists of Andersen et al. (2008) and the Convex Time Budget method of Andreoni and Sprenger (2012) the staircase method does not allow for the straight-forward estimation of parameters in the utility function, which is however not necessary to answer the research question at hand.

the main part of this paper, we present results for the full sample while controlling for our set of covariates including student status. As a robustness check, we show in Appendix D.1.6 that all results reproduce in a restricted sample of N=396 subjects that excludes all non-students.

5.2.6 Empirical Strategy

We test for treatment effects by comparing outcomes in the optimistic and the pessimistic treatment as the treatment effects are expected to be largest between these two conditions. The balanced treatment provides a critical consistency check for the hypothesis that the degree of optimism of the narratives drives the treatment effects. If the treatment effects are driven by the degree of optimism of the narratives, then we should observe that the means of the outcomes in the balanced treatment lie between the means in the optimistic and the pessimistic treatments. The baseline condition is included in the design to provide an article that is unrelated to the pandemic as another benchmark.

5.3 Results

We present our results in the following order: first, we present the effects of narratives on expectations, the central variable of interest in existing information provision experiments (see Fuster and Zafar, 2023; Haaland et al., 2023). We then extend the analysis to our novel behavioral outcomes.

5.3.1 Narratives Impact Expectations

Expectations Figure 5.2 depicts the mean forward-looking expectations of subjects for (a) the total number of deaths related to COVID-19 in Germany and (b) the German stock market index DAX by treatment condition. The dashed line in Figure 5.2 indicates the mean in the baseline condition. Across both measures, subjects in the pessimistic treatment are more pessimistic compared to the optimistic treatment. Subjects in the pessimistic treatment expect 634 more people to have died related to COVID-19 within the next two months (+6.8%). They also expect the DAX to close on average 478 points lower in two months than subjects in the optimistic treatment (-4.2%).

Mann-Whitney U tests confirm that the differences between the optimistic and pessimistic treatment are significant for expectations about deaths related to COVID-19 (p=0.024) and for DAX expectations (p=0.025).[17] For expectations about the number of COVID-19 cases, there is no significant difference between the optimistic and pessimistic treatment (Mann-Whitney U test: p=0.170). Hence, our narratives seem to have a strong effect on beliefs about the severity of the pandemic (COVID-19 deaths), yet not so much

[17]Mann-Whitney U tests are our preferred test for treatment effects on expectations as they are robust to outliers.

FIGURE 5.2: Treatment Effects on Expectations

(A) COVID-19 Deaths (B) DAX

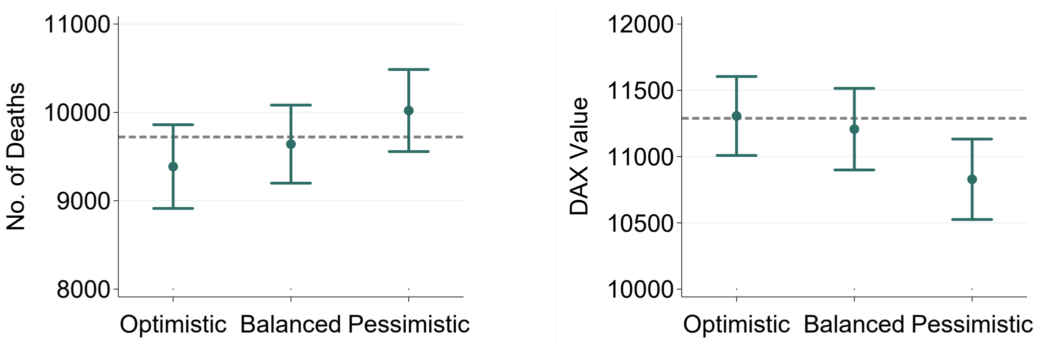

Notes: Figure 5.2 shows means and corresponding 95% confidence intervals for (a) COVID-19 related deaths and (b) expectations about the stock market index DAX in the three treatment conditions. The dashed line indicates the mean in the baseline condition. Based on OLS estimates reported in Appendix Table D.8.

about incidences (COVID-19 cases).[18] Corresponding OLS estimates are presented in Appendix Table D.8.

> **Result 1:** When confronted with a more pessimistic narrative about the coronavirus pandemic, subjects hold more pessimistic expectations about the pandemic and the stock market.

5.3.2 Narratives Impact Behavioral Outcomes

Risk Aversion Figure 5.3a shows the average certainty equivalent elicited for the lottery (50% EUR 0, 50% EUR 4) by treatment condition. The average certainty equivalent in the pessimistic treatment is substantially lower than in the optimistic treatment (EUR 1.71 in pessimistic versus EUR 2.00 in optimistic). On average, subjects in the optimistic treatment act risk neutral, so that they maximize expected earnings, while subjects in the pessimistic treatment show a considerable level of risk aversion. Figure 5.3b depicts histograms of the certainty equivalent in the optimistic treatment relative to the pessimistic treatment.

In Table 5.1 we provide our main regression analyses. In column (1) we report the result of an OLS regression that regresses the certainty equivalent on the treatment dummies with the optimistic treatment serving as the reference group. In column (2) we additionally control for our set of covariates. Column (1) shows that the treatment effect on risk aversion (0.41 standard deviations) is highly significant ($p = 0.002$). The coeffi-

[18] Note that the data on expectations for COVID-19 cases turn out to be more noisy than the other two measures as they contain a number of implausible answers (see Appendix D.1.3 for details), which may explain why we do not detect an effect on this measure.

FIGURE 5.3: Treatment Effects on Risk Aversion

(A) Certainty Equivalent - Means 95% CI (B) Certainty Equivalent - Distributions

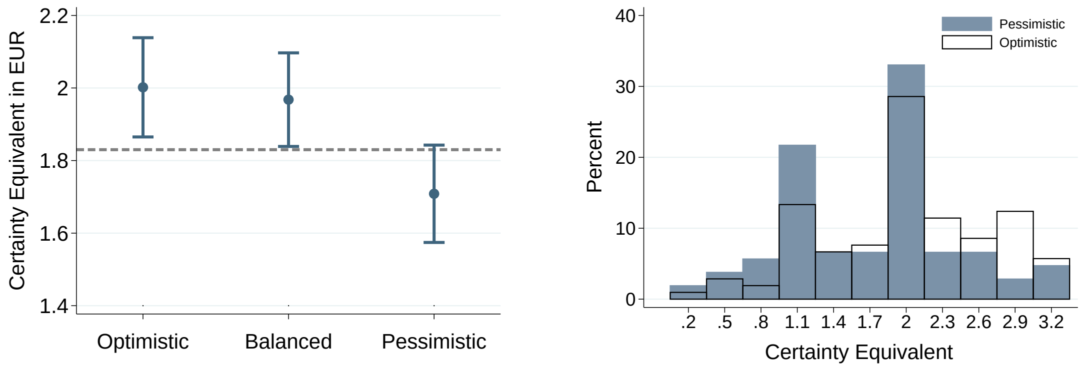

Notes: Figure (a) displays the means and 95% confidence intervals by treatment condition. The dashed line in Figure (a) indicates the mean in the baseline condition. Figure (b) display histograms of the certainty equivalent in the optimistic and pessimistic treatment.

cient of the pessimistic treatment dummy remains highly significant when adding controls in column (2) ($p = 0.004$).

Patience Figure 5.4a depicts the mean future value of a EUR 2 payment today by treatment condition. A higher future value implies a higher individual discount rate and hence less patient behavior.[19] Subjects in the pessimistic treatment act less patient than in the optimistic treatment (future value of EUR 3.14 versus EUR 2.85). Across treatment conditions, subjects show sizeable individual discount rates over a rather short time period of two months. Such high individual discount rates are however common in the literature using similar elicitation procedures (Ericson and Laibson, 2019; Frederick et al., 2002). Figure 5.4b depicts histograms of the future value in the optimistic treatment relative to the pessimistic treatment.

In columns (3) and (4) in Table 5.1 we present equivalent regressions to columns (1) and (2) with the future value as the dependent variable. In column (3) the treatment effect on patience (0.31 standard deviations) is significant at the 5 percent level ($p = 0.030$). It remains significant when adding controls in column (4) ($p = 0.026$).

Consistent with the hypothesis that the degree of optimism of the narratives causes the treatment effects, the means in the balanced and baseline treatment lie in between the optimistic and the pessimistic treatment for both risk aversion and patience.

[19] When confronted with the choice between 2 EUR today and X EUR in two months, the future value indicates the value X for which the subject is indifferent between receiving 2 EUR today and X EUR in two months. For example, a subject with a low future value of 2.10 EUR would choose the later payment for all X>2.10 EUR . A subject with a high future value of 4 EUR would choose 2 EUR today for all X between 2 EUR and 4 EUR and only choose the later payment for X>4 EUR. Also see Andreoni and Sprenger (2015, footnote 4).

Figure 5.4: Treatment Effects on Patience

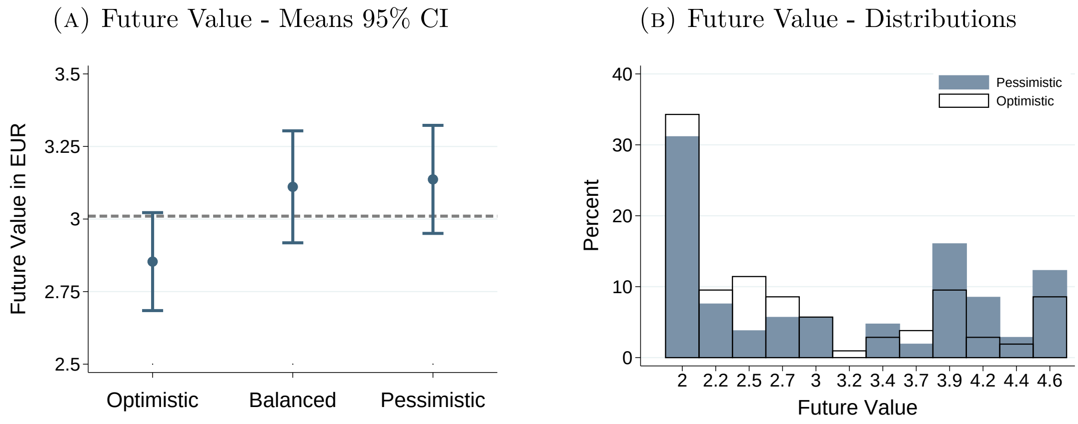

Notes: Figure (a) displays the means and 95% confidence intervals by treatment condition. The dashed line in Figure (a) indicates the mean in the baseline condition. Figure (b) display histograms of the future value in the optimistic and pessimistic treatment. Note that a higher future value implies a higher individual discount rate and hence less patient behavior (also see footnote 19).

Result 2: When confronted with a more pessimistic narrative about the coronavirus pandemic, subjects behave more risk averse and less patient.

Productivity Regarding the productivity in our real-effort task we find that the mean of correctly solved tasks does not differ between the optimistic treatment and the pessimistic treatment (optimistic: 16.2 versus pessimistic 16.3; t-test, p=0.896). Moreover, coefficients and standard errors of the treatment indicators presented in columns (5) and (6) in Table 5.1 indicate that there are no significant differences between any two treatment conditions.

Result 3: Exposure to narratives about the coronavirus pandemic does not affect productivity in a short real-effort task.

Robustness In Table 5.1 we provide p-values adjusted for multiple hypothesis testing as we test the same treatment on three behavioral outcomes. In Appendix Table D.3 we further show with logit models that the treatment effects on risk aversion and patience can already be detected when focusing the analysis on the first of the five decisions in the staircase method, which should address any concerns about using the staircase method as an elicitation method. As a further robustness check, we present tobit models that account for censoring of the outcome variables in Appendix Table D.3. To complete our robustness analysis, we show in Appendix Table D.4 that there are no significant order effects and that there is no significant heterogeneity in treatment effects depending on the order of elicitation.

TABLE 5.1: Average Treatment Effects on Behavioral Outcomes: OLS Estimates

	Certainty Equivalent		Future Value		Productivity	
	(1)	(2)	(3)	(4)	(5)	(6)
Pessimistic	-0.29***	-0.29***	0.28**	0.30**	-0.07	0.23
	(0.10)	(0.10)	(0.13)	(0.13)	(0.53)	(0.52)
Balanced	-0.03	-0.03	0.26**	0.26**	0.52	0.57
	(0.10)	(0.10)	(0.13)	(0.13)	(0.53)	(0.52)
Baseline	-0.18*	-0.17*	0.16	0.21	-0.03	0.10
	(0.10)	(0.10)	(0.13)	(0.13)	(0.53)	(0.52)
Age		0.00		0.00		-0.14***
		(0.01)		(0.01)		(0.04)
Female		-0.06		-0.02		-0.40
		(0.07)		(0.10)		(0.39)
Income		0.02		-0.21*		0.79
		(0.09)		(0.12)		(0.48)
Education		-0.07		0.09		-0.07
		(0.05)		(0.07)		(0.27)
Econ Student		-0.07		-0.21**		0.17
		(0.08)		(0.10)		(0.40)
No Student		-0.06		-0.00		1.88**
		(0.16)		(0.21)		(0.84)
Political Orientation		-0.00		0.02		-0.09
		(0.02)		(0.03)		(0.12)
Risk Group		-0.02		0.06		-1.89***
		(0.12)		(0.16)		(0.63)
Constant	2.00***	2.05***	2.85***	3.02***	16.30***	19.40***
	(0.07)	(0.19)	(0.09)	(0.25)	(0.37)	(0.99)
Observations	423	423	423	423	423	423
R^2	0.028	0.035	0.014	0.038	0.004	0.076
Initial p-values:						
Pessimistic	$p = 0.002$	$p = 0.004$	$p = 0.030$	$p = 0.026$	$p = 0.896$	$p = 0.664$
Adjusted p-values (Romano-Wolf):						
Pessimistic	$p = 0.009$	$p = 0.013$	$p = 0.054$	$p = 0.045$	$p = 0.891$	$p = 0.662$

Notes: Table reports OLS estimates with standard errors in parentheses. The optimistic treatment is the reference group. Adjusted p-values for multiple hypothesis testing are calculated using the Romano-Wolf step-down procedure as described in Clarke et al. (2019). We control for the fact that we test the same treatment on three behavioral outcomes. The adjusted p-values are separately derived for the specification without covariates (columns (1), (3) and (5)) and for the specification with covariates (columns (2), (4) and (6)) using 5000 bootstrap replications. *** p<0.01, ** p<0.05, * p<0.1

Subgroup Analysis Overall we find little evidence that the treatment effects of our manipulation differ systematically across socio-demographic subgroups by gender, age, education or income (see Appendix Table D.5 and Table D.6). One informative observation that supports the external validity of our findings is that the behavioral effects of optimistic and pessimistic narratives on economic decision making persist for economics students (38.3% of the sample), a subgroup with very high financial education. Contrary to what one might expect, the behavioral effects of narratives are, if anything, more pronounced among economics students (see Appendix Table D.6).

We further observe that subjects with a high level of news consumption prior to the experiment update their expectations about the DAX to a smaller degree than subjects with a low level of news consumption (see Appendix Table D.9). We do however not find such a pattern for the behavioral outcomes (see Appendix Table D.6). Hence the behavioral effects of narratives on risk aversion and patience do not decrease in the level of previous news exposure in our sample. This observation may provide a rationale for why opinion leaders repeatedly provide the same narrative to their audience: narratives may be persuasive through their behavioral effects on decision making even if the narratives contain no new information.

5.4 Discussion

5.4.1 Mechanism: Why Do Narratives Cause Behavioral Effects?

The novel and potentially surprising behavioral effects of narratives that we document in our paper raise questions about the underlying mechanism and the generalizability of these effects. Why do narratives about the pandemic cause such strong behavioral effects on risk aversion and patience? And should we expect to observe similar behavioral effects for narratives about other topics?

Shiller (2017) has already argued that one of the main reasons why narratives are often used in persuasive communication is that narratives can have very strong emotional effects on the audience. Moreover, a number of papers in economics have shown that general affect and emotions such as fear or stress can change people's risk aversion and patience (Cahlíková and Cingl, 2017; Cohn et al., 2015; Guiso et al., 2018; Ifcher and Zarghamee, 2011; Meier, 2022). Therefore, it seems plausible that the behavioral effects on risk aversion and patience in our experiment may be caused by affective reactions towards the narratives.

In our experimental data we do in fact find that narratives about the pandemic have very strong emotional effects: subjects in the pessimistic treatment show lower general affect than subjects in the optimistic treatment (t-test, $p=0.004$), which is driven by subjects feeling more afraid (t-test, $p<0.001$), upset (t-test, $p<0.001$) and nervous (t-test, $p=0.004$) after exposure to the pessimistic narrative (see Appendix D.1.4 for details).

We also find that subjects in the pessimistic treatment are much less optimistic about the development of their own personal circumstances in the weeks after the experiment

(Diff: 0.41 SD; t-test: p = 0.007) (see Appendix Figure D.3 for details). This highlights just how strongly subjects expect the pandemic to impact their own lives. These treatment effects on subjects' personal optimism may in turn explain why the narratives had such strong emotional reactions: because subjects anticipated positive or negative personal consequences for their own lives. This conjecture is supported by very strong correlations between subjects personal optimism about their lives and their reported emotional state (see Appendix Table D.11).

In a set of regression analyses we show that subjects' personal optimism and emotional state is correlated with risk taking and patience in our experiment (see Appendix Tables D.12 and D.13). Hence, our data are consistent with the interpretation that the affective effects of narratives cause the treatment effects on risk and time preferences.

The proposed mechanism has an important implication for the generalizability of our results: narratives should only cause effects on risk aversion and patience if the narratives cause strong affective reactions in the audience. Narratives that do not cause any affective reactions, in contrast, should not be able to persuade through the preference-based mechanism.

Follow-up Experiment To test the above hypothesis, we conducted a follow-up experiment in June 2023. We recruited another N=393 subjects from the Cologne Laboratory for Economic Research (CLER). In the follow-up experiment, we provided subjects with either an optimistic or a pessimistic narrative about how a technology shock - the technological breakthrough in the development of generative AI - will impact the business model of Google. Based on the observation that only a small minority of subjects in our sample reported to own Google stocks (7.8%), these narratives should not be of great personal relevance for subjects.[20] We hypothesized that these narratives would change expectations about the development of the Google stock but would neither cause strong affective reactions nor any behavioral effects on risk aversion or patience. We pre-registered these hypotheses in our pre-registration.[21] For a detailed description of the follow-up experiment see Appendix D.3.

In line with our pre-registered hypotheses, we find strong effects on expectations about the Google stock (Diff: 6.1%, SD=0.36, MWU-test: $p < 0.001$), but neither on emotions and optimism, nor on risk aversion and patience (see Appendix D.3.3 for all details).

Taken together, the data from our main experiment and our follow-up experiment provide evidence for the hypothesis that narratives only cause behavioral effects on decision-making if they cause affective reactions in the audience. Based on our results, we would

[20] All of our results hold if we exclude these 7.8% of subjects from our analysis, for which the narratives may be of more personal relevance. The number of subjects that owned Google stocks (N=31) is however too small to conduct a well-powered subgroup analysis.

[21] You can find the pre-registration for the follow-up experiment here: `https://doi.org/10.17605/OSF.IO/T4JPN`.

expect that narratives about exogenous shocks that only hit a specific company or industry cause no strong behavioral effects in the general population. But when the very same optimistic or pessimistic narratives are provided to the workforce of that company, for example by a manager or an investor, they could cause strong affective reactions and lead to shifts in their risk and time preferences. Regarding the public communication between politicians and the general population, narratives about major exogenous shocks that impact most individuals in society, like pandemics, wars, or natural catastrophes, seem most likely to have an influence on economic decision-making through preference-based persuasion.

5.4.2 Economic Relevance and External Validity

We should still discuss whether it is likely that narratives have an economically relevant effect on economic decision-making outside of our controlled experimental context. We believe that numerous arguments support the view that the influence of narratives documented in this paper are economically relevant.

First, our results show that narratives impact fundamental determinants of economic decision-making. Forward-looking expectations are key variables in central models of investor behavior (Lucas and Sargent, 1981; Sims, 2003) and recent information provision experiments have confirmed their causal impact on a variety of economic behaviors (Bailey et al., 2019; Laudenbach et al., 2021; Roth and Wohlfart, 2020). Moreover, almost all economic decisions involve risk or intertemporal trade-offs. Hence, changes in risk aversion and patience should impact a wide variety of economic decisions, from portfolio choices, to savings decisions, or insurance choices.

Second, the effects of narratives on risk and time preferences are at least persistent in the short-term: we do not find that performing a mentally challenging two-minute real-effort task (our productivity measure) prior to the elicitation of risk aversion and patience reduces treatment effects (see Appendix Table D.4). Given that people today are exposed to narratives at high frequency via the internet and social media, even short-term treatment effects on risk aversion and patience would imply a meaningful influence on a wide range of economic choices.

Third, the belief-based and preference-based persuasion mechanisms identified in this paper could interact and amplify each other. Changes in expectations, risk aversion and patience can all separately - but also jointly - influence many economic decisions, for example, investments in the stock market. Thus a narrative could influence an economic decision through both mechanisms at the same time.

Last, our experiment did take place in the same setting in which people today frequently consume news and take a large share of their economic decisions: at home in front of their computers. We hence believe that the effects observed in our experiment translate comparatively well into behavior outside of our experimental context and are economically relevant.

Our experimental results may, however, also exaggerate the effects of narratives on people's behavior. In contexts in which the audience does not trust the sender of a narrative, the effects of narratives on expectations and preferences could be much weaker (or even non-existent). This may also be the case in contexts in which the audience understands the motives of a sender to choose a narrative strategically as a tool of persuasive communication.

5.5 Conclusion

This paper studies how narratives influence economic decision-making. In information experiments, subjects read naturalistic news articles that contain optimistic or pessimistic narratives about exogenous shocks. We then document the effects of these narratives on three fundamental determinants of economic decision-making: forward-looking expectations, risk preferences, and time preferences. Our data show that reading narratives about a major exogenous shock (the coronavirus pandemic) causes changes in expectations about the stock market, but also short-term changes in people's risk and time preferences.

Our results have implications for understanding how the media and persuasive communication influence economic decisions. Standard theories and existing information provision experiments have focused on one particular persuasion mechanism: belief-based persuasion. Our paper, in contrast, provides evidence that narratives have the potential to influence economic decision-making through a second - yet neglected - mechanism: through preference-based persuasion.

Understanding these mechanisms seems important for individuals and organizations that engage in persuasive communication. Managers, investors, and politicians, could try to spread narratives that influence the risk taking behavior and intertemporal decision-making of their audience. Experts and journalists, in contrast, may want to use narratives that only persuade through the belief-based mechanism. How opinion leaders design and strategically spread narratives seem fascinating questions for future research.

Appendix to Chapter 5

D.1 Additional Analyses

D.1.1 Sample Characteristics and Randomization Check

Table D.1: Balance Table

	Optimistic (1)	Pessimistic (2)	Balanced (3)	Baseline (4)	Full Sample (5)
Age	25.81	26.98	25.93	26.92	26.41
	(4.69)	(7.69)	(5.46)	(7.51)	(6.47)
Female	65.7%	62.3%	62.3%	61.3%	62.9%
Income	901.19	893.87	930.42	985.85	927.90
	(448.92)	(423.21)	(483.69)	(484.68)	(460.66)
Education					
High School	47.6%	41.5%	41.5%	47.2%	44.4%
Bachelor	41.9%	43.4%	38.7%	34.0%	39.5%
Master	10.5%	15.1%	19.8%	18.9%	16.1%
Student Status					
Non Econ	57.1%	52.8%	61.3%	50%	55.3%
Econ	31.4%	43.4%	33.0%	45.3%	38.3%
No Student	11.4%	3.8%	5.7%	4.7%	6.4%
Political Orientation	0.78	0.74	0.52	0.33	0.59
	(1.43)	(1.45)	(1.65)	(1.39)	(1.49)
Risk Group COVID-19	9.5%	10.4%	8.5%	9.4%	9.4%
Observations	105	106	106	106	423

Notes: Income: disposable income per month in EUR os; Political Orientation: scale from right (-3) to left (3) with the German parties assigned to values as follows. AFD: -3, FDP: -2, CDU/CSU: -1, SPD: 1, Bündnis90/Grüne: 2, Die Linke: 3; unaffiliated participants were assigned the value 0; Risk Group COVID-19: belonging to a group at high risk for a severe case of COVID-19.

Table D.2: Tests for the Balance of Covariates: p-values

	Opt=Pess (1)	Opt=Bal (2)	Pess=Bal (3)	Pess=Base (4)	Opt=Base (5)	Bal=Base (6)
Age	0.183	0.859	0.254	0.957	0.198	0.273
Female	0.602	0.602	1.00	0.888	0.507	0.888
Income	0.903	0.650	0.559	0.143	0.190	0.406
Education	0.510	0.165	0.618	0.359	0.182	0.693
Student Status	0.043**	0.324	0.278	0.889	0.048**	0.188
Political Orientation	0.820	0.219	0.311	0.038**	0.021**	0.369
Risk Group COVID-19	0.836	0.793	0.638	0.818	0.982	0.810
Observations	211	211	212	211	212	212

Notes: The table reports p-values for the following tests: for age, income, and political orientation the p-values of a t-test; for female, education, student status and risk group the p-values of a Chi2-test. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

D.1.2 Behavioral Outcomes

D.1.2.1 Distribution of Behavioral Outcomes

Figure D.1: Distribution of Behavioral Outcomes

(a) Certainty Equivalent

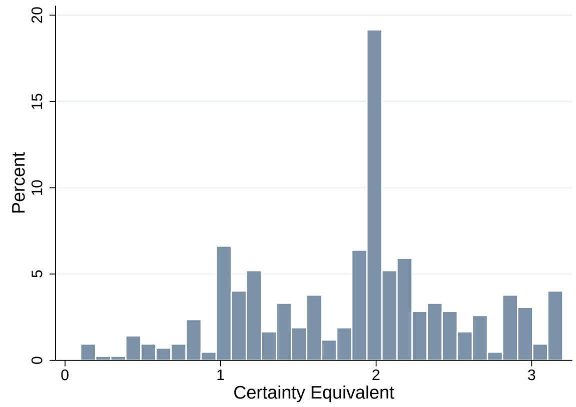

(b) Future Value

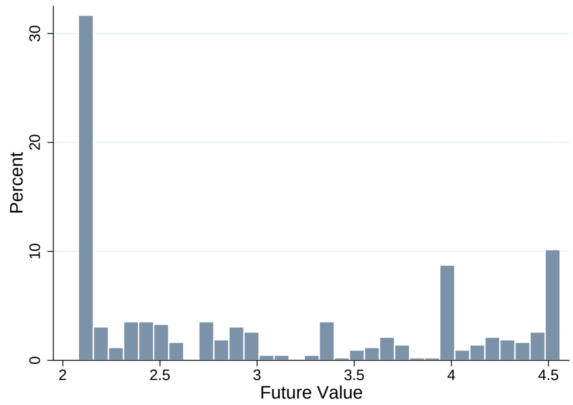

(c) Productivity

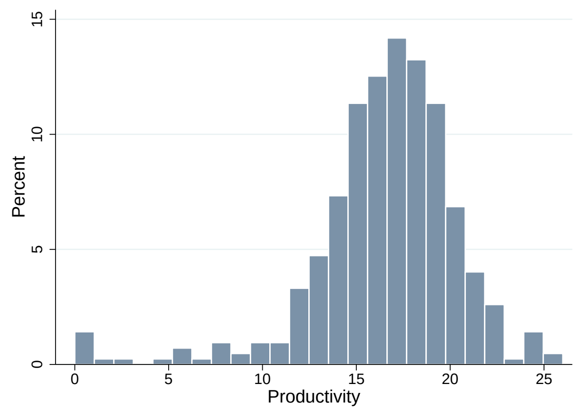

Notes: Histograms for (a) the certainty equivalent, (b) the future value and (c) productivity.

D.1.2.2 Robustness Checks: Treatment Effects on Behavioral Outcomes

Table D.3: Robustness Check: Logit and Tobit Models

	Risk Aversion				Patience			
	Chose Lottery		Certainty Equivalent		Chose 2 € Today		Future Value	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pessimistic	-0.12*	-0.13**	-0.30***	-0.29***	0.17**	0.19***	0.30**	0.32**
	(0.06)	(0.06)	(0.10)	(0.10)	(0.07)	(0.07)	(0.14)	(0.14)
Balanced	0.01	0.01	-0.04	-0.03	0.13**	0.14**	0.28*	0.28*
	(0.07)	(0.07)	(0.10)	(0.10)	(0.07)	(0.07)	(0.14)	(0.14)
Baseline	-0.06	-0.08	-0.18*	-0.17*	0.11	0.14**	0.16	0.20
	(0.07)	(0.07)	(0.10)	(0.10)	(0.07)	(0.07)	(0.14)	(0.14)
Age		0.00		0.00		-0.00		0.00
		(0.00)		(0.01)		(0.00)		(0.01)
Female		-0.05		-0.05		0.00		-0.06
		(0.05)		(0.08)		(0.05)		(0.11)
Income		0.06		0.02		-0.06		-0.22*
		(0.06)		(0.09)		(0.06)		(0.13)
Education		-0.04		-0.07		0.04		0.10
		(0.03)		(0.05)		(0.04)		(0.08)
Econ Student		0.05		-0.09		-0.15***		-0.24**
		(0.05)		(0.08)		(0.05)		(0.11)
No Student		-0.09		-0.08		0.01		-0.03
		(0.10)		(0.16)		(0.11)		(0.23)
Political Orientation		0.00		-0.00		0.01		0.03
		(0.02)		(0.02)		(0.02)		(0.03)
Risk Group		-0.10		-0.02		-0.03		0.07
		(0.08)		(0.12)		(0.08)		(0.17)
Observations	423	423	423	423	423	423	423	423

Notes: Columns (1), (2), (5) and (6) report average marginal effects from logit models on the first decision in the respective elicitation procedure (see Appendix D.2.3). Columns (3), (4), (7) and (8) report coefficients from tobit models that account for censoring from above of the outcome variables. Standard errors in parentheses.*** p<0.01, ** p<0.05, * p<0.1

D.1.2.3 Order Effects: Behavioral Outcomes

Table D.4: Order Effects: Behavioral Outcomes

	Certainty Equivalent		Future Value		Productivity	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Order Effects						
Order = 2	-0.01	-0.02	0.07	0.08	-0.13	-0.27
	(0.08)	(0.08)	(0.11)	(0.11)	(0.46)	(0.45)
Order = 3	-0.12	-0.14*	0.15	0.16	0.13	-0.02
	(0.08)	(0.08)	(0.11)	(0.11)	(0.45)	(0.45)
Controls	No	Yes	No	Yes	No	Yes
Observations	423	423	423	423	423	423
Panel B: Heterogeneous Treatment Effects by Order						
Pessimistic	-0.33**	-0.28*	0.07	0.22	-0.38	0.23
	(0.17)	(0.15)	(0.22)	(0.20)	(0.95)	(0.81)
Order = 2	0.04	-0.02	-0.09	0.06	-0.19	-0.09
	(0.17)	(0.10)	(0.22)	(0.13)	(0.96)	(0.52)
Order = 3	-0.31*	-0.14	-0.14	0.13	-0.46	-0.19
	(0.17)	(0.10)	(0.23)	(0.13)	(0.90)	(0.52)
Pessimistic X Order = 2	-0.06	-0.00	0.27	0.09	-0.36	-0.74
	(0.23)	(0.19)	(0.31)	(0.26)	(1.35)	(1.06)
Pessimistic X Order = 3	0.19	-0.01	0.39	0.15	1.11	0.61
	(0.24)	(0.20)	(0.32)	(0.27)	(1.28)	(1.02)
Controls	No	Yes	No	Yes	No	Yes
Observations	423	423	423	423	423	423
Panel C: Heterogeneous Treatment Effects by Order (After Productivity)						
Pessimistic	-0.35***	-0.34**	0.20	0.29*		
	(0.13)	(0.14)	(0.18)	(0.16)		
Order After RET	-0.19	-0.18	-0.06	0.14		
	(0.14)	(0.14)	(0.19)	(0.11)		
Pessimistic X Order After RET	0.12	0.10	0.18	0.01		
	(0.19)	(0.20)	(0.26)	(0.21)		
Controls	No	Yes	No	Yes		
Observations	423	423	423	423		

Notes: Table reports OLS estimates with standard errors in parentheses. In Panel B and C regressions include dummies for the balanced and baseline treatment and their interactions with the order dummies. In all panels controls include our standard set of covariates (age, female, income, education, econ student, no student, political orientation, risk group). Constants not reported. *** p<0.01, ** p<0.05, * p<0.1

D.1.2.4 Subgroup Analysis

TABLE D.5: Heterogeneous Treatment Effects on Behavioral Outcomes #1

	Certainty Equivalent		Future Value	
	(1)	(2)	(3)	(4)
Panel A: Gender				
Pessimistic	-0.16	-0.15	0.51**	0.55**
	(0.16)	(0.16)	(0.22)	(0.22)
Female	0.05	0.03	0.21	0.16
	(0.14)	(0.15)	(0.19)	(0.20)
Pessimistic × Female	-0.20	-0.21	-0.36	-0.41
	(0.20)	(0.20)	(0.27)	(0.27)
Controls	No	Yes	No	Yes
Observations	423	423	423	423
Panel B: Age				
Pessimistic	-0.29**	-0.29**	0.23	0.26
	(0.13)	(0.13)	(0.17)	(0.17)
Age(>25)	0.03	0.06	-0.11	-0.10
	(0.14)	(0.14)	(0.19)	(0.20)
Pessimistic × Age(>25)	0.00	0.00	0.14	0.10
	(0.20)	(0.20)	(0.27)	(0.27)
Controls	No	Yes	No	Yes
Observations	423	423	423	423
Panel C: Education				
Pessimistic	-0.36**	-0.36**	0.33*	0.36*
	(0.14)	(0.15)	(0.20)	(0.20)
Education (≥Bachelor)	0.02	0.02	0.13	0.16
	(0.14)	(0.14)	(0.19)	(0.19)
Pessimistic × Education (≥Bachelor)	0.11	0.12	-0.09	-0.12
	(0.19)	(0.20)	(0.26)	(0.27)
Controls	No	Yes	No	Yes
Observations	423	423	423	423

Notes: Table reports OLS estimates with standard errors in parentheses. All regressions also include dummies for the balanced and baseline treatment and their interactions with the respective covariate. Controls include our standard set of covariates (age, female, income, education, econ student, no student, political orientation, risk group) excluding the covariate that is interacted with the treatment dummies in the respective regression. Constants not reported. *** p<0.01, ** p<0.05, * p<0.1

Table D.6: Heterogeneous Treatment Effects on Behavioral Outcomes #2

	Certainty Equivalent		Future Value	
	(1)	(2)	(3)	(4)
Panel A: Income				
Pessimistic	-0.30**	-0.29**	0.05	0.07
	(0.14)	(0.14)	(0.19)	(0.19)
Income(≥875)	-0.06	-0.03	-0.28	-0.25
	(0.14)	(0.14)	(0.18)	(0.19)
Pessimistic × Income(≥875)	0.02	0.00	0.45*	0.46*
	(0.19)	(0.20)	(0.26)	(0.26)
Controls	No	Yes	No	Yes
Observations	423	423	423	423
Panel B: Econ Students				
Pessimistic	-0.18	-0.17	0.13	0.11
	(0.12)	(0.12)	(0.16)	(0.17)
Econ Student	0.21	0.20	-0.29	-0.29
	(0.15)	(0.15)	(0.20)	(0.20)
Pessimistic × Econ Student	-0.32	-0.34*	0.42	0.44
	(0.20)	(0.20)	(0.27)	(0.27)
Controls	No	Yes	No	Yes
Observations	423	423	423	423
Panel C: News Consumption				
Pessimistic	-0.25	-0.23	0.26	0.25
	(0.15)	(0.15)	(0.20)	(0.20)
News Consumption (≥Often)	0.00	0.01	0.19	0.20
	(0.14)	(0.14)	(0.19)	(0.19)
Pessimistic ×	-0.08	-0.10	0.04	0.08
News Consumption (≥Often)	(0.20)	(0.20)	(0.26)	(0.27)
Controls	No	Yes	No	Yes
Observations	423	423	423	423

Notes: Table reports OLS estimates with standard errors in parentheses. All regressions also include dummies for the balanced and baseline treatment and their interactions with the respective covariate. Controls include our standard set of covariates (age, female, income, education, econ student, no student, political orientation, risk group) excluding the covariate that is interacted with the treatment dummies in the respective regression. Constants not reported. *** p<0.01, ** p<0.05, * p<0.1

D.1.3 Expectations

Note on data cleaning Expectations about the DAX, COVID-19 related deaths and COVID-19 cases were elicited with an open text box, so that subjects could enter any value. Therefore, the data set contains a number of implausible values and outliers. One noteworthy data cleaning step was performed on these three measures: we recoded values that were unreasonably low and were most likely meant to be in thousands. For example, an entry of 12.5 for the DAX Value was recoded as 12500 and an entry of 10.2 for COVID-19 deaths was recoded as 10200.

In Table D.7 we show with Mann-Whitney U tests that treatment effects on expectations are significant irrespective of performing this data cleaning step. Mann-Whitney U tests are our preferred test for treatment effects on expectations as they are robust to outliers.

Table D.7: Treatment Effects on Expectations: Robustness to Data Cleaning

	DAX (1)	COVID-19 Deaths (2)	COVID-19 Cases (3)
P-values from Mann-Whitney U tests (Optimistic = Pessimistic)			
Prior to Cleaning	0.025	0.021	0.209
After Cleaning	0.025	0.024	0.170
N Implausible Prior to Cleaning	10	23	51
N Cleaned	10	11	22
N Implausible After Cleaning	0	12	29
N total	423	423	422 [a]

Notes: Table reports p-values from Mann-Whitney U tests, the number of observations recoded as part of the data cleaning and the number of observations that are still implausible after data cleaning. [a]One observation for COVID-19 cases is dropped as the subject entered "improved" instead of a number.

After data cleaning, the expectations about COVID-19 deaths and cases still contain a number of implausibly low values (lower than the initial value of COVID-19 deaths/ cases in Germany on May 3rd). This issue is most severe for COVID-19 cases, which is therefore our most noisy measure among the four forward-looking expectations. Figure D.2 depicts the distribution of expectations after data cleaning. We account for outliers in our regression analysis by winsorizing the expectations about the DAX, COVID-19 related deaths and COVID-19 cases. In that way we do not drop any observation from our analysis.

Figure D.2: Distribution of Expectations

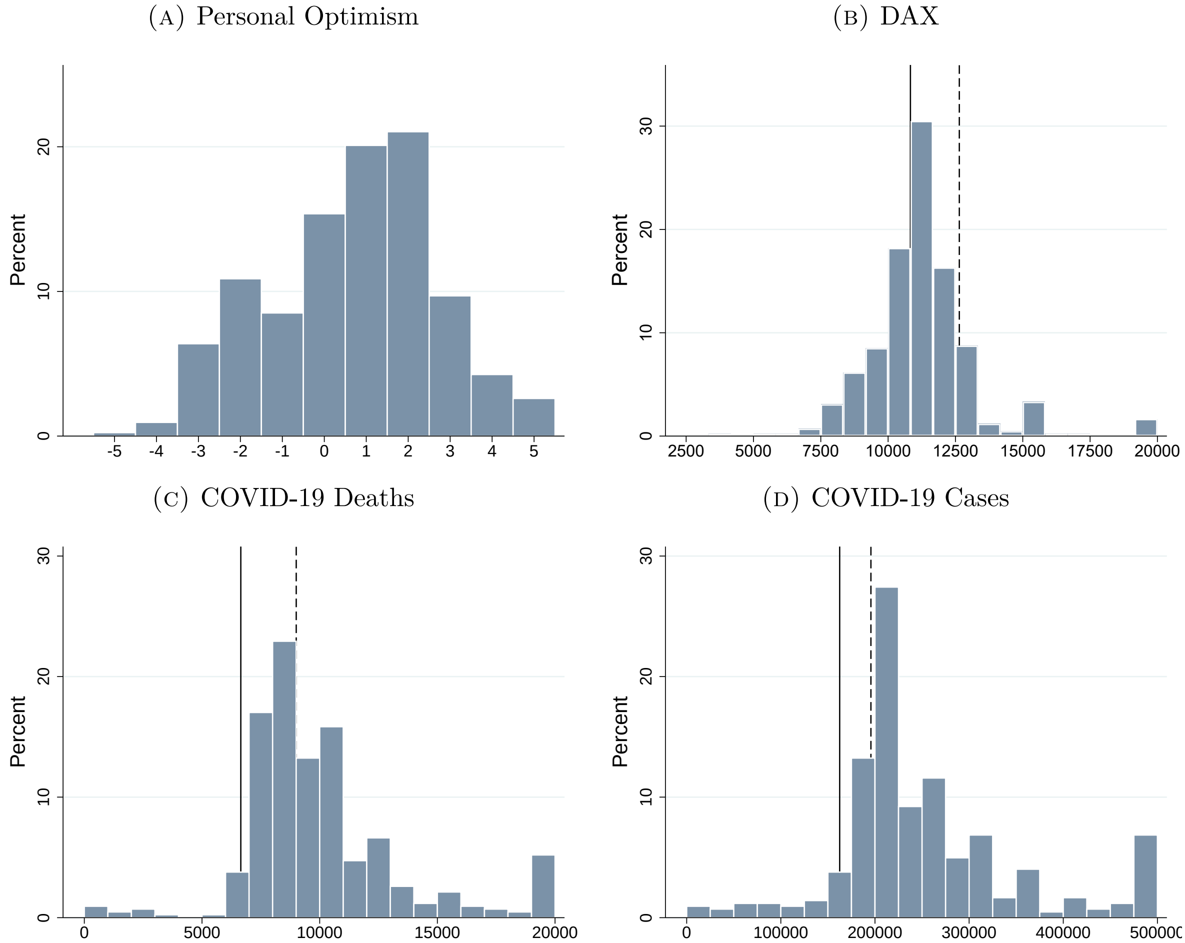

Notes: Histograms for (a) personal optimism (expectations about personal circumstances), (b) expectations about the DAX Value, (c) expectations about COVID-19 deaths and (d) expectations about COVID-19 cases. The solid line indicates the initial value on May 3rd. The dashed line indicates the realized value on July 3rd. Note that all values below the initial value in (c) and (d) are implausible values as the total number of COVID-19 deaths or cases cannot decrease. The histogram for the Dax Value has been winsorized at 20,000 points. The histogram for COVID-19 deaths has been winsorized at 20,000 deaths and the histogram for COVID-19 cases at 500,000 cases.

It was not necessary to perform any data cleaning on the qualitative measure of personal optimism (measured on an 11-point Likert scale) which showed significant treatment effects of our manipulation using both Mann-Whitney U tests ($p = 0.007$) and OLS regressions ($p < 0.005$) (see Table D.8 and Figure D.3).

FIGURE D.3: Treatment Effects on Personal Optimism

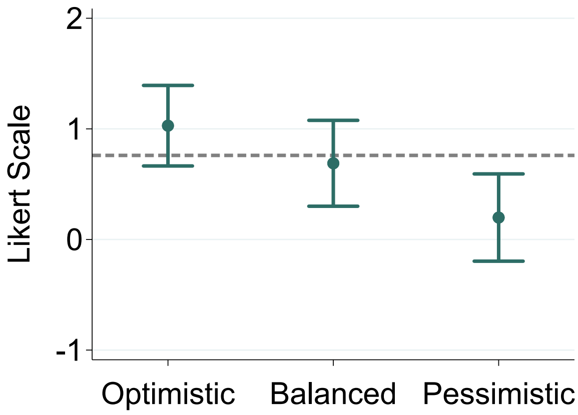

Notes: Means and 95% confidence intervals for personal optimism by treatment condition. The dashed line indicates the mean in the baseline condition.

TABLE D.8: Treatment Effects on Expectations: OLS Estimates

	Personal Optimism		DAX		COVID-19 Deaths		COVID-19 Cases	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pessimistic	-0.83***	-0.77***	-477.81**	-491.35**	634.29*	597.05*	10977.00	9583.35
	(0.28)	(0.28)	(216.74)	(221.44)	(329.10)	(331.66)	(12359.53)	(12596.51)
Balanced	-0.34	-0.38	-99.40	-114.99	254.12	216.54	-1603.95	-1249.13
	(0.28)	(0.28)	(216.74)	(219.43)	(329.10)	(328.66)	(12359.53)	(12481.81)
Baseline	-0.26	-0.31	-17.15	-50.77	335.85	371.06	-169.79	1612.31
	(0.28)	(0.28)	(216.74)	(222.02)	(329.10)	(332.54)	(12388.78)	(12651.73)
Age		-0.03		20.82		-33.44		45.71
		(0.02)		(15.33)		(22.96)		(872.26)
Female		-0.22		66.17		-716.39***		-14408.24
		(0.21)		(167.37)		(250.68)		(9530.51)
Income		0.59**		-48.36		-123.92		-16999.88
		(0.26)		(204.69)		(306.58)		(11644.51)
Education		0.12		8.66		186.64		1898.16
		(0.15)		(116.64)		(174.70)		(6654.85)
Econ Student		0.04		-82.87		-64.90		1520.77
		(0.21)		(169.77)		(254.27)		(9680.29)
No Student		0.64		28.60		-561.35		-6993.36
		(0.45)		(357.08)		(534.82)		(20311.45)
Political Orientation		-0.10		-66.50		141.73*		4188.61
		(0.07)		(53.02)		(79.40)		(3016.61)
Risk Group		-0.65*		-86.96		230.07		-1586.47
		(0.34)		(268.10)		(401.54)		(15250.10)
Constant	1.03***	1.27**	11306.66***	10847.06***	9387.10***	10667.16***	253448.94***	273065.70***
	(0.20)	(0.53)	(153.62)	(421.87)	(233.26)	(631.86)	(8760.19)	(24032.10)
Observations	423	423	423	423	423	423	422	422
R^2	0.022	0.068	0.015	0.027	0.009	0.047	0.003	0.020

Notes: Table reports OLS estimates with standard errors in parentheses. In columns (3) to (8) we control for outliers by winsorizing the outcome variables. In columns (3) and (4) the DAX values have been winsorized at 7,500 and 15,000 points. In columns (5) and (6) the COVID-19 related deaths have been winsorized at the initial value on May 3rd (6,649) and at 15,000 deaths. In columns (7) and (8) COVID-19 cases have been winsorized at the initial value on May 3rd (162,496) and at 500,000 cases. *** p<0.01, ** p<0.05, * p<0.1

TABLE D.9: Heterogeneous Treatment Effects on Expectations by News Consumption

	Personal Optimism		DAX		COVID-19 Deaths		COVID-19 Cases	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pessimistic	-1.20***	-1.20***	-950.43***	-982.75***	848.39*	890.42*	7080.18	8530.55
	(0.43)	(0.42)	(335.22)	(337.90)	(509.05)	(506.99)	(19189.17)	(19307.59)
News Consumption (≥Often)	0.21	0.00	-172.68	-190.71	578.29	587.31	1773.91	4086.59
	(0.40)	(0.40)	(310.99)	(315.24)	(472.26)	(472.98)	(17802.09)	(18012.46)
Pessimistic × News Consumption (≥Often)	0.64	0.75	809.20*	859.43*	-369.94	-513.99	6650.31	1755.60
	(0.56)	(0.56)	(439.06)	(445.94)	(666.74)	(669.08)	(25133.38)	(25480.51)
Observations	423	423	423	423	423	423	422	422
R^2	0.035	0.079	0.027	0.039	0.020	0.056	0.007	0.023

Notes: Table reports OLS estimates with standard errors in parentheses. Subjects with high news consumption (≥ Often) make up 59.6% of our sample. All regressions also include dummies for the balanced and baseline treatment and their interactions with the high news consumption dummy. Controls include our standard set of covariates (age, female, income, education, econ student, no student, political orientation, risk group). In columns (3) to (8) we control for outliers by winsorizing the outcome variables. In columns (3) and (4) the DAX values have been winsorized at 7,500 and 15,000 points. In columns (5) and (6) the COVID-19 related deaths have been winsorized at the initial value on May 3rd (6,649) and at 15,000 deaths. In columns (7) and (8) COVID-19 cases have been winsorized at the initial value on May 3rd (162,496) and at 500,000 cases. *** p<0.01, ** p<0.05, * p<0.1

D.1.4 Emotions

Figure D.4: Treatment Effects on Emotions

(a) Upset (b) Afraid (c) Nervous

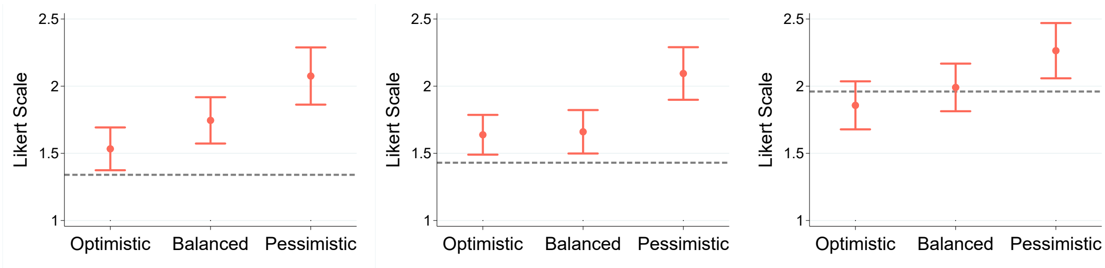

Notes: Figure D.4 shows means and corresponding 95% confidence intervals for the emotional state of subjects across treatment conditions: for feeling (a) upset, (b) afraid and (c) nervous. The dashed line indicates the mean in the baseline condition. Emotions are measured on a 5-point Likert scale (1 "not at all" to 5 "very much").

Table D.10: Treatment Effects on Affect and Emotions

	Affect	Upset	Afraid	Nervous	Attentive	Determined	Inspired
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Pessimistic	-1.23***	0.54***	0.46***	0.41***	0.00	0.02	0.16
	(0.41)	(0.12)	(0.12)	(0.14)	(0.11)	(0.14)	(0.14)
Balanced	-0.08	0.21*	0.02	0.13	0.13	0.09	0.07
	(0.41)	(0.12)	(0.12)	(0.14)	(0.11)	(0.14)	(0.14)
Baseline	0.83**	-0.18	-0.20*	0.11	0.03	0.21	0.31**
	(0.41)	(0.12)	(0.12)	(0.14)	(0.11)	(0.14)	(0.14)
Constant	4.30***	1.53***	1.64***	1.86***	3.72***	3.10***	2.50***
	(0.29)	(0.09)	(0.09)	(0.10)	(0.08)	(0.10)	(0.10)
Observations	423	423	423	423	423	423	423
R-squared	0.06	0.08	0.07	0.02	0.00	0.01	0.01

Notes: Table reports OLS estimates with standard errors in parentheses. Affect is constructed as the sum of the positive items (attentive, determined, inspired) minus the negative items (upset, afraid, nervous) *** p<0.01, ** p<0.05, * p<0.1

D.1.5 Understanding the Mechanism

In Table D.11 we provide pairwise correlation coefficients between subjects' personal optimism, emotions and expectations in our sample.

In Table D.12 and Table D.13 we regress the behavioral outcomes (risk aversion and patience) separately on each potential mediator while controlling for our standard set of socio-demographic variables. We find that subjects who report a lower personal optimism in life tend to behave more risk averse (see column (5) in Table D.12). Further, subjects who feel more afraid and upset tend to act significantly more impatient (see column (2) and (3) in Table D.13). The direction of the correlation between feeling afraid and risk aversion, while not significant (p=0.133), is broadly in line with the literature (Cohn et al., 2015; Guiso et al., 2018; Meier, 2022), that is, more afraid subjects tend to show a higher level of risk aversion.

Expectations about the pandemic and the stock market are not significantly correlated with risk aversion. They are weakly correlated with patience (see column (6) and (7) in Table D.13), but if these were causal relationships, then the treatment effects on these expectations would actually bias against the treatment effect on patience. Hence, both treatment effects on risk aversion and patience cannot be well explained by a change in these forward-looking expectations.

TABLE D.11: Correlation Coefficients between Expectations and Emotions

	Personal Optimism	DAX	COVID-19 Deaths	COVID-19 Cases	Affect	Upset	Afraid	Nervous
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Personal Optimism	1.000							
DAX	0.236***	1.000						
COVID-19 Deaths	-0.164***	-0.152***	1.000					
COVID-19 Cases	-0.104**	-0.079	0.594***	1.000				
Affect	0.283***	0.126***	-0.012	0.016	1.000			
Upset	-0.236***	-0.107**	0.043	0.037	-0.585***	1.000		
Afraid	-0.166***	-0.071	-0.002	0.012	-0.622***	0.391***	1.000	
Nervous	-0.201***	-0.051	-0.013	-0.038	-0.599***	0.385***	0.599***	1.000

Notes: Table reports pairwise Pearson correlation coefficients. DAX values have been winsorized at 7,500 and 15,000 points, COVID-19 related deaths at the initial value on May 3rd (6,649) and at 15,000 deaths, COVID-19 cases at the initial value on May 3rd (162,496) and at 500,000 cases.*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table D.12: Behavioral Mechanism for Risk Aversion

	Certainty Equivalent							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Affect	0.015							
	(0.011)							
Upset		0.004						
		(0.037)						
Afraid			-0.058					
			(0.039)					
Nervous				0.011				
				(0.035)				
Personal Optimism					0.033*			
					(0.017)			
DAX						0.027		
						(0.022)		
COVID-19 Deaths							-0.023	
							(0.015)	
COVID-19 Cases								-0.034
								(0.039)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	423	423	423	423	423	423	423	422
R^2	0.014	0.010	0.015	0.010	0.018	0.013	0.016	0.009

Notes: Table reports OLS estimates with standard errors in parentheses. Expectations about the DAX are winsorized at 7,500 and 15,000 points and transformed in thousands. COVID-19 related deaths are winsorized at the initial value (6,649) and at 15,000 deaths and transformed in thousands. COVID-19 deaths are winsorized at the initial value (162,496) and at 500,000 cases and are transformed in hundred thousands. Controls include our standard set of covariates (age, female, income, education, econ student, no student, political orientation, risk group). Constant not reported. *** p<0.01, ** p<0.05, * p<0.1

TABLE D.13: Behavioral Mechanism for Patience

	Future Value							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Affect	-0.021							
	(0.015)							
Upset		0.094*						
		(0.050)						
Afraid			0.108**					
			(0.052)					
Nervous				0.059				
				(0.046)				
Personal Optimism					-0.007			
					(0.023)			
DAX						0.048		
						(0.029)		
COVID-19 Deaths							-0.036*	
							(0.020)	
COVID-19 Cases								-0.057
								(0.052)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	423	423	423	423	423	423	423	422
R^2	0.029	0.033	0.035	0.028	0.025	0.031	0.032	0.028

Notes: Table reports OLS estimates with standard errors in parentheses. Expectations about the DAX are winsorized at 7,500 and 15,000 points and transformed in thousands. COVID-19 related deaths are winsorized at the initial value (6,649) and at 15,000 deaths and transformed in thousands. COVID-19 deaths are winsorized at the initial value (162,496) and at 500,000 cases and are transformed in hundred thousands. Controls include our standard set of covariates (age, female, income, education, econ student, no student, political orientation, risk group). Constant not reported. *** p<0.01, ** p<0.05, * p<0.1

D.1.6 Robustness Checks: Restricted Sample

Table D.14: Restricted Sample: Average Treatment Effects with OLS

	Certainty Equivalent		Future Value		Productivity	
	(1)	(2)	(3)	(4)	(5)	(6)
Pessimistic	-0.33***	-0.31***	0.30**	0.31**	-519.65**	-545.29**
	(0.10)	(0.10)	(0.14)	(0.14)	(225.63)	(228.26)
Balanced	-0.10	-0.09	0.26*	0.27*	-175.77	-213.90
	(0.10)	(0.10)	(0.14)	(0.14)	(226.71)	(229.28)
Baseline	-0.21**	-0.20*	0.14	0.19	-82.96	-130.03
	(0.10)	(0.10)	(0.14)	(0.14)	(226.16)	(230.37)
Age		-0.00		0.00		24.06
		(0.01)		(0.01)		(18.42)
Female		-0.07		-0.02		114.07
		(0.08)		(0.10)		(173.42)
Income		0.04		-0.22*		3.39
		(0.10)		(0.13)		(215.38)
Education		-0.08		0.07		34.95
		(0.06)		(0.08)		(125.58)
Econ Student		-0.08		-0.20**		-67.55
		(0.08)		(0.10)		(171.02)
No Student		0.00		0.00		0.00
		(.)		(.)		(.)
Political Orientation		-0.00		0.02		-55.89
		(0.02)		(0.03)		(54.98)
Risk Group		-0.07		0.08		-95.63
		(0.13)		(0.17)		(283.14)
Constant	2.04***	2.15***	2.85***	2.97***	11336.62***	10715.56***
	(0.07)	(0.21)	(0.10)	(0.29)	(163.19)	(483.69)
Observations	396	396	396	396	396	396
R^2	0.030	0.041	0.015	0.039	0.016	0.027
Initial p-values:						
(Pessimistic)	0.001	0.002	0.029	0.023	0.813	0.629
Adjusted p-values (Romano-Wolf):						
(Pessimistic)	0.003	0.006	0.054	0.046	0.808	0.627

Notes: Table reports OLS estimates with standard errors in parentheses. Adjusted p-values for multiple hypothesis testing were calculated using the Romano-Wolf step-down procedure as described in Clarke et al. (2019). We control for the fact that we test the same treatment on three different outcomes. The adjusted p-values were separately derived for the specification without covariates (columns (1), (3) and (5)) and for the specification with covariates (columns (2), (4) and (6)). *** p<0.01, ** p<0.05, * p<0.1

Table D.15: Restricted Sample: Robustness Checks with Logit and Tobit Models

	Risk Taking				Patience			
	Chose Lottery		Certainty Equivalent		Chose EUR 2 Today		Future Value	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pessimistic	-0.14**	-0.14**	-0.34***	-0.32***	0.18***	0.20***	0.32**	0.33**
	(0.07)	(0.07)	(0.10)	(0.10)	(0.07)	(0.07)	(0.15)	(0.15)
Balanced	-0.03	-0.02	-0.10	-0.09	0.14**	0.15**	0.28*	0.28*
	(0.07)	(0.07)	(0.10)	(0.10)	(0.07)	(0.07)	(0.15)	(0.15)
Baseline	-0.08	-0.09	-0.21**	-0.20*	0.10	0.13*	0.14	0.19
	(0.07)	(0.07)	(0.10)	(0.10)	(0.07)	(0.07)	(0.15)	(0.15)
Age		0.00		-0.00		0.00		0.00
		(0.01)		(0.01)		(0.01)		(0.01)
Female		-0.04		-0.06		-0.00		-0.06
		(0.05)		(0.08)		(0.05)		(0.11)
Income		0.00		0.00		-0.00		-0.00*
		(0.00)		(0.00)		(0.00)		(0.00)
Education		-0.06		-0.08		0.02		0.08
		(0.04)		(0.06)		(0.04)		(0.08)
Econ Student		0.05		-0.10		-0.15***		-0.23**
		(0.05)		(0.08)		(0.05)		(0.11)
Political Orientation		0.00		0.00		0.01		0.02
		(0.02)		(0.03)		(0.02)		(0.04)
Risk Group COVID-19		-0.15*		-0.07		-0.02		0.09
		(0.08)		(0.13)		(0.09)		(0.19)
Observations	396	396	396	396	396	396	396	396

Notes: Columns (1), (2), (5) and (6) report average marginal effects from logit models on the first decision in the respective game tree (see Appendix D.2.3). Columns (3), (4), (7) and (8) report coefficients from tobit models that account for censoring from above of the outcome variables. Standard errors in parentheses.*** p<0.01, ** p<0.05, * p<0.1

D.2 Additional Materials

D.2.1 Timeline of the COVID-19 Pandemic in Germany

Figure D.5: Daily New Infections in Germany

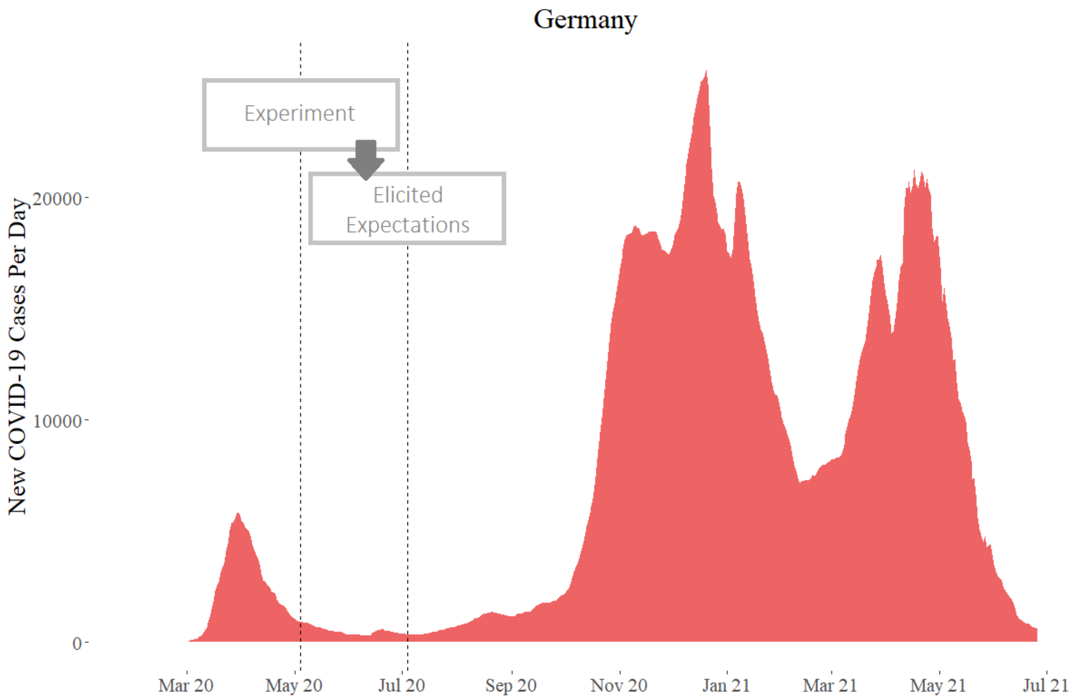

Notes: The graph illustrates the numbers of daily new infections in Germany reported to the Robert Koch Institute from March 2020 to July 2021. Our experiment was conducted on May 4th 2020. Participants had to state their expectations about the course of the pandemic until July 3rd 2020.

D.2.2 Structure and Content of the Pandemic Narratives

D.2.2.1 Structure of Narratives

Figure D.6 illustrates the common structure of all narratives about COVID-19 provided as our manipulation. All narratives consist of five paragraphs covering the same aspects of the COVID-19 pandemic as shown in Figure D.6. The numbers on the right side of Figure D.6 refer to the sentences within the respective paragraph. The corresponding sentences can be found in the transcripts of the narratives provided in sections D.2.2.2 to D.2.2.4.

Figure D.6: Structure of Narratives

Paragraph 1	Implications of “Opening Up”	(1.1 - 1.2)
	Statement of Chancellor Merkel	(1.3)
Paragraph 2	Model of the Pandemic	(2.1 - 2.2)
Paragraph 3	Impacts on the Health Care System	(3.1 - 3.3)
Paragraph 4	Impacts on the Economy	(4.1 - 4.3)
Paragraph 5	Current State of Research	(5.1)
	Predictions about Vaccine	(5.2)

Notes: This figure depicts the common structure of all narratives about COVID-19 used as our experimental manipulation. The numbers on the right side refer to the sentences within the respective paragraph.

The information provided in the narratives was spread in this or in a very similar way in news articles and in public communication in the weeks prior to our experiment. The statements of chancellor Angela Merkel were made during a press conference on 20th April 2020.[22]

The baseline text followed a similar structure. In the baseline text, a quote of Galileo Galilei was used instead of a statement of Angela Merkel and a story about outer space was provided instead of a narrative about the COVID-19 pandemic. The transcript of the baseline text is available in section D.2.2.5.

[22] The transcript of the press conference is available under https://www.bundeskanzlerin.de/bkin-de/aktuelles/pressekonferenz-von-bundeskanzlerin-merkel-1745362 (accessed on April 27th, 2021)

D.2.2.2 Transcript Optimistic Narrative

In Germany the measures to contain the spread of the coronavirus are currently being relaxed. (1.1) Now more and more people move around in public and many shops are reopening. (1.2) Due to its discipline, the population has made great achievements in the last weeks, chancellor Angela Merkel praised in a speech. (1.3)

Many of those currently infected with the coronavirus are expected to recover within the next days. By now, many have already recovered. (2.1) Day by day, the number of new infections decreases compared to previous weeks. This trend is expected to continue. (2.2)

So far the German health care system has not come close to reaching its capacity limit. (3.1) In comparison to Italy or Spain, the situation in Germany has almost always been under control. (3.2) Many physicians in Germany were even less occupied than usual as a lot of non-urgent interventions have been postponed. (3.3)

Due to the relaxation of restrictions, the economy picks up again. (4.1) Customers go shopping more frequently, which stimulates sales for many business. Some people are even starting to make plans for summer holidays. (4.2) It seems that all the effort of the last weeks eventually pays off. (4.3)

Meanwhile scientists around the world are constantly working on better understanding the novel coronavirus. (5.1) A vaccine might soon be found. (5.2)

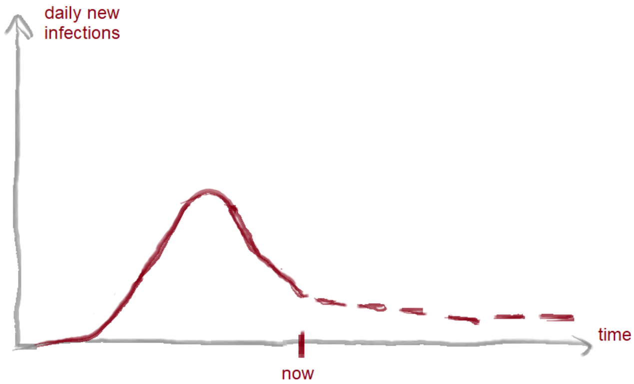

Note: Narratives were provided in German and did not contain the numbers in gray which are included as a reference to the common structure of all narratives (see Figure D.6).

D.2.2.3 Transcript Pessimistic Narrative

In Germany the measures to contain the spread of the coronavirus might be relaxed too soon. (1.1) If more and more people move around in public, a second wave of infections becomes likely. (1.2) The population should not for a second lull itself into a false sense of security, chancellor Angela Merkel warned in a speech. (1.3)

It is expected that in a second wave of infections significantly more elderly will be infected with the coronavirus. (1.2) A second wave would thus turn out to be a lot deadlier. (2.2)

In a second wave the German health care system might collapse. (3.1) Germany could then face conditions like in Italy or Spain, where the situation spiraled out of control. (3.2) Physicians had to decide which patients to treat and whom to let die – the so-called triage. (3.3)

If the virus starts to spread faster and faster again, the economy faces the threat of a second, likely more severe, shutdown. (4.1) A second shutdown would mean final bankruptcy for a lot of businesses. (4.2) In that case all the effort of the last weeks would be lost. (4.3)

Meanwhile, many fundamental questions about the novel coronavirus remain unanswered. So far the infection rate and the most common transmission paths have not been identified. (5.1) Most likely it will take until next year until a vaccine is available. (5.2)

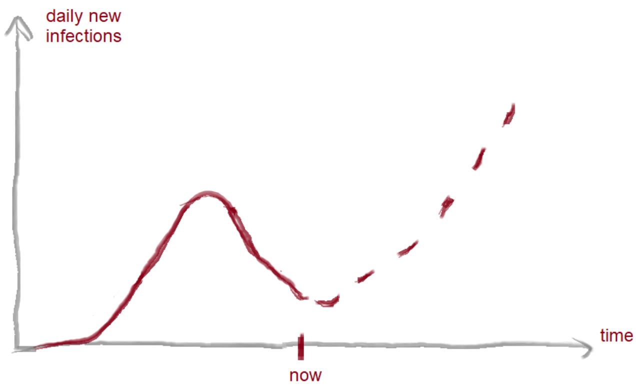

Note: Narratives were provided in German and did not contain the numbers in gray which are included as a reference to the common structure of all narratives (see Figure D.6).

D.2.2.4 Transcript Balanced Narrative

In Germany the restrictive policies to contain the spread of the coronavirus are slowly being relaxed. (1.1) That is good news for people and the economy, but increases the risk of a second wave of infections. (1.2) The population has made great achievements, but should not lull itself into a false sense of security, chancellor Angela Merkel said in a speech. (1.3)

Currently, daily new infections are decreasing. In some regions and age groups there have yet been almost no deaths. (2.1) A second wave of infections could, however, turn out to be a lot deadlier. (2.2)

So far the German health care system has not reached its capacity limit. (3.1) In comparison to Italy and Spain, the situation in Germany has been relatively well under control. (3.2) In some cases physicians in Italy and Spain had to decide whom to treat and whom to let die. (3.3)

Due to the relaxation of restrictions, customers go shopping more frequently. This is good for many businesses. (4.1) A second shutdown could, however, be more severe than the first one. A second shutdown could mean final bankruptcy for a number of businesses. (4.2) Therefore, it remains to be seen if the efforts of the last weeks will eventually pay off. (4.3)

Meanwhile scientists are constantly working on open questions regarding the novel coronavirus. (5.1) It is however hard to predict when a vaccine will be available. (5.2)

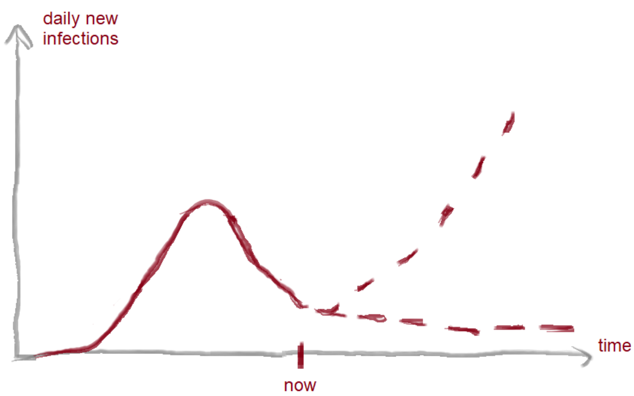

Note: Narratives were provided in German and did not contain the numbers in gray which are included as a reference to the common structure of all narratives (see Figure D.6).

D.2.2.5 Transcript Baseline Text

As early as in the 18th century German scientists dreamt of sailing in outer space. Already the astronomer Johannes Kepler wrote in a letter to Galileo Galilei: “Provide ships or sails that are suitable for the breeze of heaven”. This dream came true last year. A mission showed that objects in outer space can be moved only by the force of a sail.

To a layperson, such a project may seem absurd. There is no air in outer space and hence no wind to blow into an ordinary sail. But apparently it is possible to sail with solar radiation. This is made possible as there is extremely little frictional resistance in outer space.

Previously, many similar cosmic sailing projects have failed. On a recent mission, however, it worked – with the use of a very light space probe and a comparatively large sail.

This was the second time it could be shown that such a mechanical propulsion can work. If the mission continues without any problems, the efforts of the ancient thinkers might finally pay off.

Meanwhile a lot of questions about outer space remain unanswered. A mechanical propulsion that is independent of rocket engines could help lead scientists to many new insights. However, no one can predict if and when this will be the case.

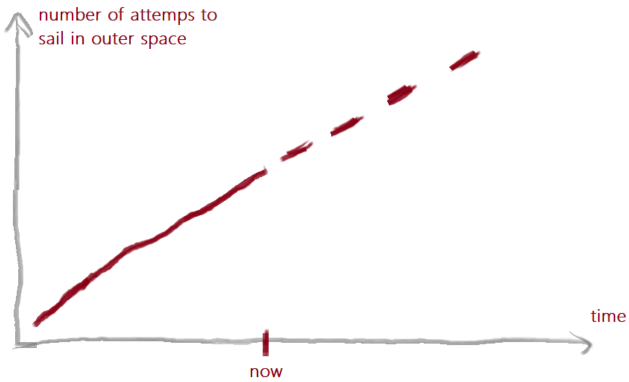

Note: The text was provided in German.

D.2.3 Elicitation of Behavioral Outcomes

D.2.3.1 Risk Aversion

Risk aversion is elicited with the staircase method for risk preferences from Falk et al. (2018) with adjusted payoffs. Subjects take five consecutive decisions, each time facing the following question:

> "*Do you want to receive a safe payment of EUR X or play a lottery with 50 percent chance for EUR 4 and 50 percent chance for EUR 0?*
>
> - *EUR X as safe payment.*
> - *A lottery with 50% chance for EUR 4 and 50% chance for EUR 0.*"

X is replaced with the corresponding value at each decision node in the game tree (see Figure D.7). The starting value for X is 1.65. In the game tree shown in Figure D.7 the action A refers to choosing the lottery, while the action B refers to choosing the safe payment of X. The value at the next decision node is then inserted as X in the subsequent question. The outcome of the game is the certainty equivalent (CE) used for analysis which can take 32 values ranging from EUR 0.10 to EUR 3.20.

D.2.3.2 Patience

Patience is elicited with the staircase method for time preferences from Falk et al. (2018) with adjusted payoffs. Subjects take five consecutive decisions, each time facing the following question:

> "*Do you want to receive EUR 2 euros today or EUR X in two months?*
>
> - *EUR 2 today.*
> - *EUR X in two months.*"

X is replaced with the corresponding value at each decision node in the game tree (see Figure D.8). The starting value for X is 3.32. In the game tree shown in Figure D.8 the action A refers to choosing EUR 2 today while the action B refers to choosing the payment of X in two months. The value at the next decision node is then inserted as X in the subsequent question. The outcome of the game is the future value (FV) used for analysis which can take 32 values ranging from EUR 2.08 to EUR 4.56.

Figure D.7: Game Tree of the Staircase Method For Risk Aversion

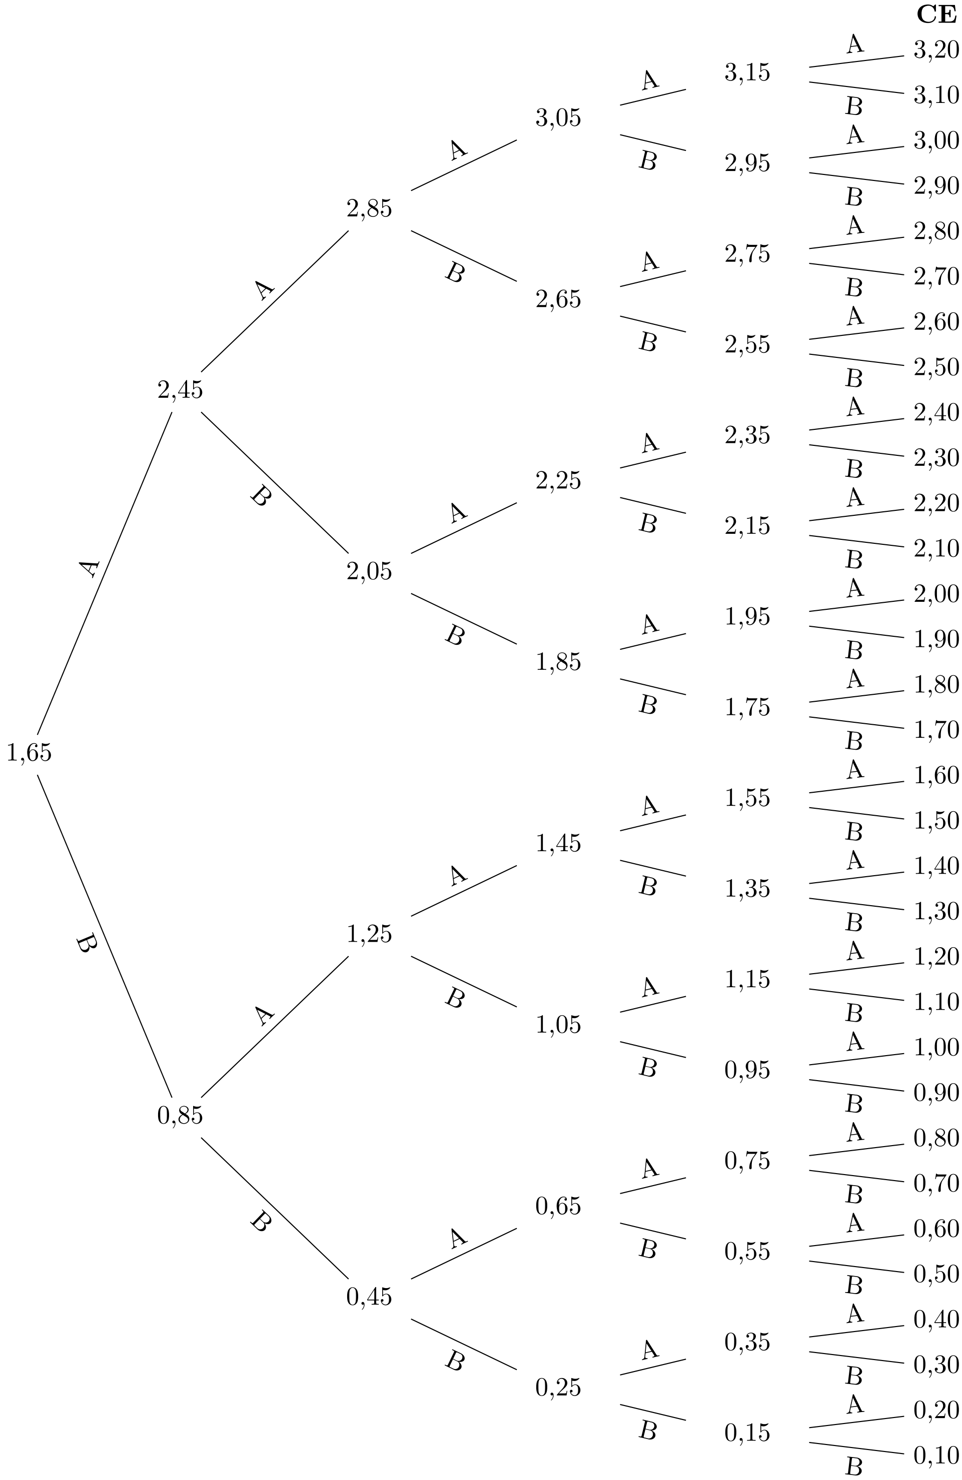

Notes: Participants take five decisions between a lottery with 50% chance for EUR 4 and 50% chance for EUR 0 (A) or EUR X euros as a safe payment (B). X is replaced with the value at each decision node with 1.65 as the initial value. The CE indicates the outcome of the game, that is the certainty equivalent used for analysis.

Figure D.8: Game Tree of the Staircase Method for Patience

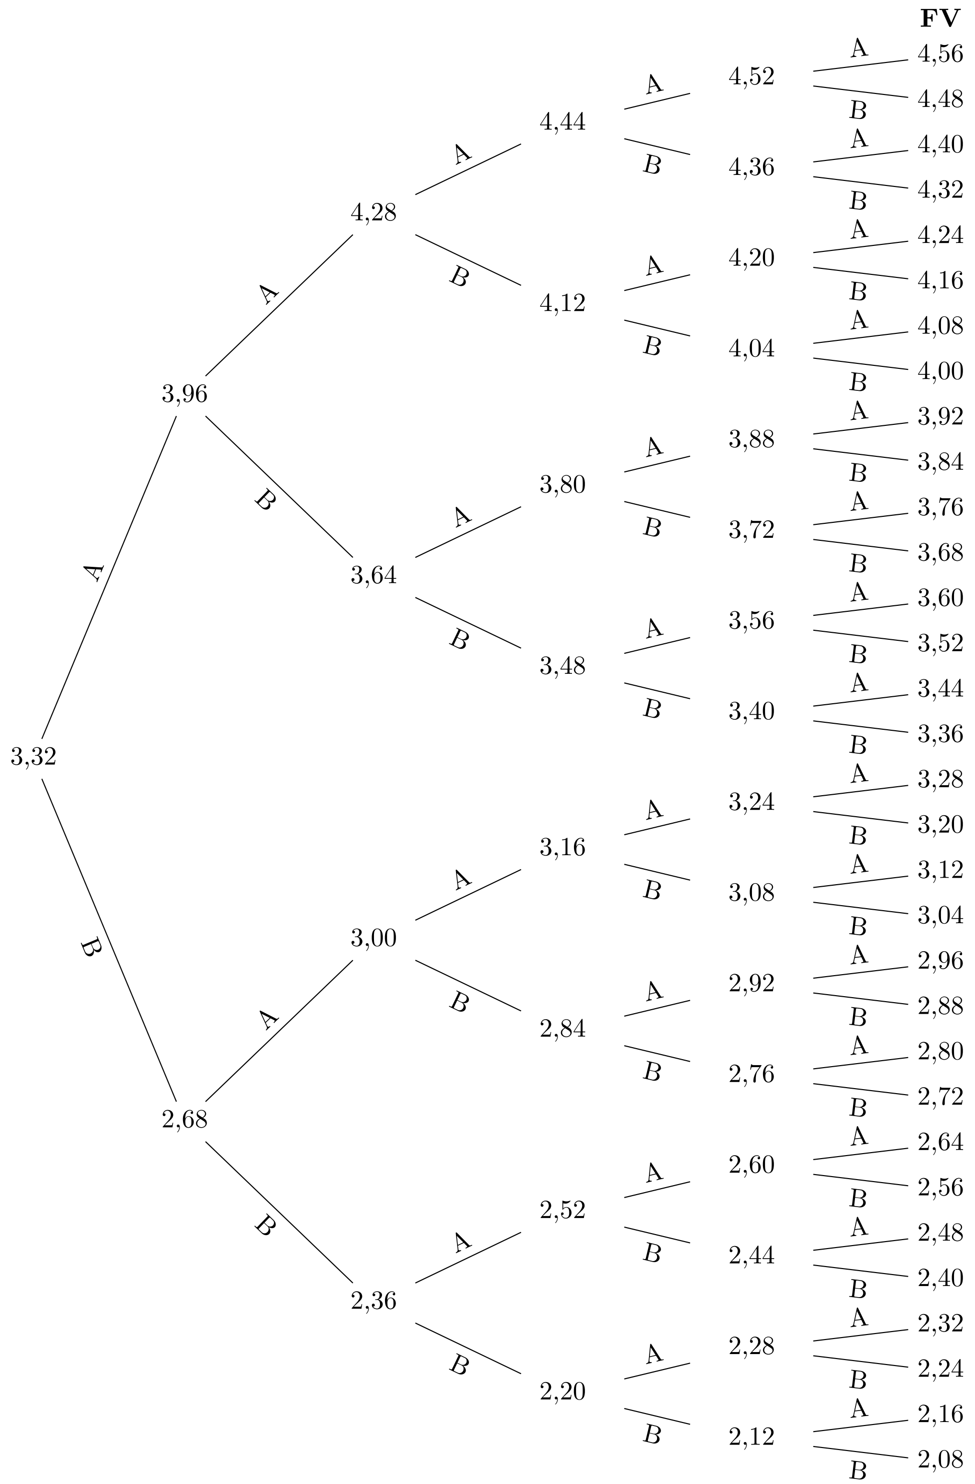

Notes: Participants take five decisions between a payment of EUR 2 today (A) or EUR X in two months (B). X is replaced with the value at each decision node with 3.32 as the initial value. The FV indicates the outcome of the game, that is the future value used for analysis.

D.2.3.3 Productivity

Productivity is measured in a real-effort task: subjects have to count the digit "1" in lines of twelve to fourteen symbols. Subjects have two minutes time to complete as many lines as possible (up to 37). For each correct line subjects are paid EUR 0.10. The lines were presented to participants in sequential order. Subjects could not go back to the previous line to revise their answers. After two minutes, all participants were forwarded and had to stop solving the task. The remaining time was displayed throughout the real effort task (see Figure D.9).

The design of the task is inspired by a concentration test.[23] We calibrated the task so that entering random numbers is not a profitable strategy. Entering random numbers would lead to just 3-4 correct answers in expectation - much less than the productivity of all subjects in a pilot study.

FIGURE D.9: Screenshot of the Productivity Task

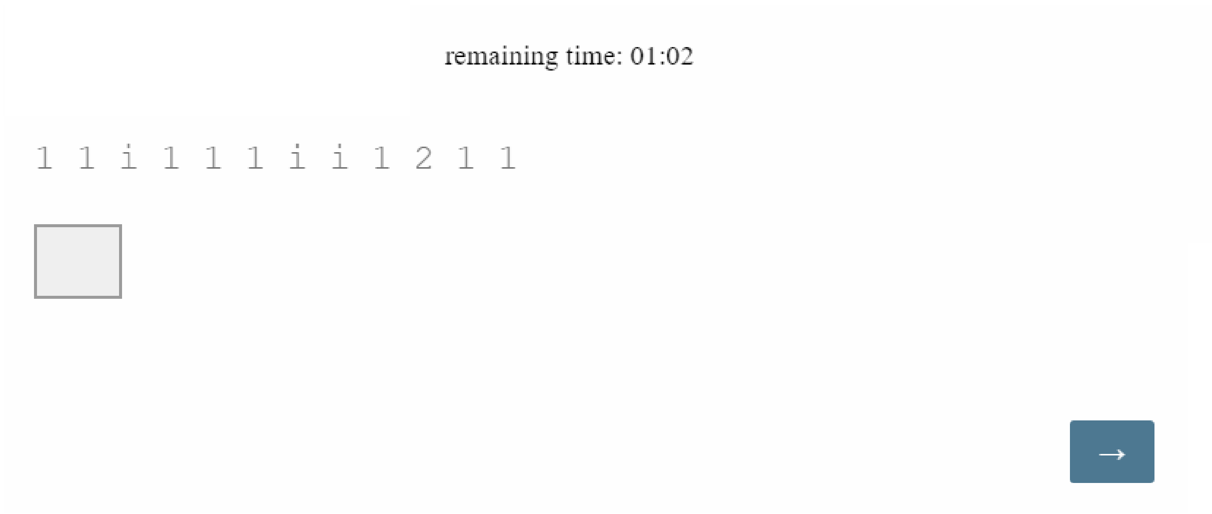

[23]See the KONT-P concentration test, https://www.psychomeda.de/online-tests/konzentrationstest.html

D.2.4 Instructions Main Experiment

Participants received experimental instructions in German. The original instructions are publicly available in our data repository on OSF (https://osf.io/bx396). Below we provide an English transcript. A dashed line indicates the next page of the survey. Explanatory comments (which were not displayed to participants) are indicated by blue, italic text.

Welcome to this online experiment! You will receive €2.50 as a show-up fee. Depending on the decisions you take during this study, you can earn an additional payment. As described in the invitation, your payment will be transferred to your PayPal account. Therefore, you will be asked to provide the email address of your PayPal account at the end of this survey. Please make sure you know the email address of you PayPal account before you begin. Your participation will take approximately 15 minutes.

O I consent to the above conditions.

Participants could only continue when they gave their consent.

- -

As start of this study two telephone numbers will be displayed. Please try to memorise the numbers. You will have 20 seconds to do so.

You will be asked to recognise the two numbers at a later point in time.

- -

05454/444-54 08421/792-65

Participants were automatically forwarded after 20 seconds. A timer indicated the remaining time participants had on this page.

- -

On the next page a topical text will be displayed. Please try to memorise as much of the content as possible. You will have two minutes to do so.

At a later point in time you will be asked to answer three questions about the content of the text. You will earn €0.50 per correct answer.

- -

One of the four treatment manipulations was randomly selected and displayed. Participants could not leave this page independently (skip the text). Subjects were automatically forwarded after two minutes. A timer indicated the remaining time they had left to read the text. The treatment texts are available in full length in Appendix A2.

- -

We would like to know how you feel right now.
The following words describe different feelings and sensations. Read every word, then indicate the intensity with which you experience the respective emotion at the moment. You can choose between five gradations.

	not at all	a little	somewhat	much	very much
upset	○	○	○	○	○
attentive	○	○	○	○	○
afraid	○	○	○	○	○
determined	○	○	○	○	○
nervous	○	○	○	○	○
inspired	○	○	○	○	○

At the beginning of the study you were shown two telephone numbers. Which of the following telephone numbers?

O 02235/679-89 O 0721/972-56
O 08421/792-65 O 05454/444-54

Now you will make decisions in three blocks. At the end of the study, one of the three blocks will be randomly selected. Only the decisions made in the selected block will be relevant for the variable part of your payment.

This means that every decision that you will make can potentially influence the payment you receive. You should therefore take all decisions as if they would be implemented.

In-between the decision blocks we will ask questions about the text that you have read in the beginning.

Decision block 1 starts now.

The order of the elicitation of risk taking, patience and productivity was randomized. Thus, block 1 could contain any of the three behavioral outcomes. As an example, we are presenting the elicitation of risk aversion here.

In this block you will take five decisions. You will always have the choice between a guaranteed payment and a lottery which pays €4 with 50 percent chance and €0 with 50 percent chance.

In this block one of your five decisions is randomly selected to be considered for payment.

Do you want to receive a guaranteed payment of €X or play a lottery with 50 percent chance for €4 and 50 percent chance for €0?

O €X as guaranteed payment   O lottery with 50% chance for €4 and 50% chance for €0

This question was displayed five times with different values for X. The first value for X was €1.65 and subsequent values depended on the previous decisions. Figure A3 in Appendix A3 shows the game tree.

Next is a question on the text that you read in the beginning.

- -

The following statement was made or was contained in the text in this or in a similar fashion:

– sentence –

O True O False

Instead of – sentence – a sentence from the treatment text was displayed to participants. This sentence was the same independently of which outcome was elicited in block 1. The correct answer for all statements is "True". The sentences displayed were:

Pessimistic: The population should not for a second lull itself into a false sense of security, chancellor Angela Merkel warned in a speech.

Balanced: The population has made great achievements, but should not lull itself into a false sense of security, chancellor Angela Merkel said in a speech.

Optimistic: Due to its discipline the population has made great achievements in the last weeks, chancellor Angela Merkel praised in a speech.

Baseline: Already the astronomer Johannes Kepler wrote in a letter to Galileo Galilei: Provide ships or sails that are suitable for the breeze of heaven.

- -

Decision block 2 starts now.

- -

Here one of the two remaining behavioral outcomes was randomly elicited. As an example we are presenting the elicitation of patience here.

In this block you will take five decisions. You always have the choice between a payment you receive directly after your participation in this study and a payment you receive in 2 months (in exactly 60 days). In both cases the money will be transferred to your PayPal account.

One of your five decisions is randomly selected to be considered for payment.

- -

Do you want to receive €2 today or €X in two months?

O €2 today O €X in two months

This question was displayed five times with different values for X. The first value for X was €3.32 and subsequent values depended on the previous decisions. Figure A4 in Appendix A3 shows the game tree.

- -

Next is a question on the text that you read in the beginning.

- -

The following statement was made or was contained in the text in this or in a similar fashion:

– sentence –

O True O False

Instead of – sentence – a sentence from the treatment text was displayed to participants. This sentence was the same independently of which outcome was elicited in block 2. The correct answer for all statements is "True". The sentences displayed were:

Pessimistic: In a second wave of infections, Germany could face conditions like in Italy or Spain.

Balanced: In Germany, the situation has not yet developed like in Italy or Spain. In the worst case, this might change with a second wave of infections.

Optimistic: If the numbers continue to develop in such a positive way, the situation in Germany will not unfold like in Italy or Spain.

Baseline: If a current sailing mission in outer space continues to be successful, the efforts of the ancient thinkers might pay off.

- -

Decision block 3 starts now.

- -

Here, the remaining behavioral outcome was elicited. As an example, we are presenting the productivity task here.

In this block your task is to count how often the digit ‘1’ appears in a line of symbols. For each correct answer you receive €0.10. You have two minutes to solve as many lines as possible.

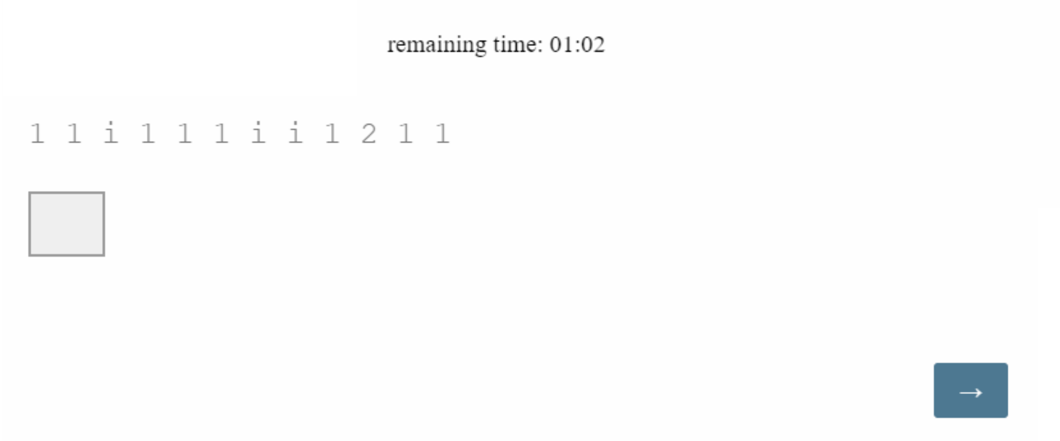

What do you think:

- How many lines did you complete? [*open text box*]
- How many lines did you answer correctly? [*open text box*]

Next is a question on the text that you read in the beginning.

The following statement was made or was contained in the text in this or in a similar fashion:

– sentence –

O True O False

Instead of – sentence – a sentence from the treatment text was displayed to participants. This sentence was the same independently of which outcome was elicited in block 3. The correct answer for all statements is "True". The sentences displayed were:

Pessimistic: Most likely it will take until next year until a vaccine is available.

Balanced: It is hard to predict when a vaccine will be available.

Optimistic: A vaccine might soon be found.

Baseline: A mechanical propulsion that is independent of rocket engines could help lead scientists to many new insights. It is however hard to predict, if and when this will be the case.

- -

The three decision blocks are completed. You now have an opportunity to earn an additional variable payment by making a number of predictions.

- -

You are now asked to make three predictions about the development of key figures regarding the current pandemic until the 3rd of July 2020 (this is in exactly 60 days). Three participants will be selected randomly for each question and will be paid depending on the accuracy of their predictions. The closer the prediction is to the realized value, the higher the payment will be. You can win up to €20 with your predictions.

Note: Your payment is independent of what other participants predict. You should therefore state the value which you regard as most likely for each figure. For the selection of the winners, only one of your predictions will be considered. Therefore it is not possible to spread your risk across predictions and you cannot win multiple times.

We will use official data from the Robert Koch Institute (RKI) and the German stock exchange to evaluate the predictions.

- What do you think: How many confirmed coronavirus cases will there be in Germany on 3rd July 2020 (in 60 days)? [*open text box*] On 3rd of May 2020 the RKI reported 162,496 confirmed coronavirus cases in Germany.
- What do you think: How many confirmed deaths due to the coronavirus will there be in Germany on 3rd July 2020 (in 60 days)? [*open text box*] On 3rd of May 2020 the RKI reported 6,649 confirmed deaths due to coronavirus in Germany.
- What do you think: With how many points will the Dax close on 3rd of July 2020 (in 60 days)? [*open text box*] On 3rd of May the Dax closed with a value of 10,828 points.

- -

Think about your personal circumstances in the next weeks. To what extent do you expect things to develop positively or negatively?

Participants had to answer the above question on an 11-point Likert scale: from very negatively (-5) to very positively (+5).

- -

Think of the upcoming days. How likely is it that ...

1. ... you only make trips that are absolutely unavoidable (e.g. to the pharmacy or supermarket)?
2. ... you always wear a face mask in the public?
3. ... you attend private parties or meet up with more than one person (who do/does not live in the same household)?
4. ... you use public transport?
5. ... you meet or visit persons who are part of a risk-group for the coronavirus?

Participants had to answer the above question on an 5-point Likert scale: very unlikely (1), rather unlikely (2), indecisive (3), rather likely (4), very likely (5). The index for compliance is then constructed based on the five answers as follows: $[(1) + (2) - (3) - (4) - (5)]/5$.

- -

How often did you inform yourself about the impacts of the coronavirus in the last days?

Participants had to answer the above question on an 5-point Likert scale: never (1), seldom (2), sometimes (3), often (4), very often (5).

In your opinion, should the current political measures to contain the spread of the coronavirus be loosened or tightened?

Participants had to answer the above question on an 5-point Likert scale: strongly loosened (-2), rather loosened (-1), neither nor (0), rather tightened (+1), strongly tightened (+2).

Thank you! Finally, a few questions about you:

- How old are you? [ ]
- Which gender do you identify with? [male / female / diverse]
- What is your subject of studies? (If more than one: Major) [all subjects that can be studied at the University of Cologne]
- What is your highest educational achievement? [No formal degree / Secondary Modern School / Junior High School / A-levels / Master Craftsmen / Bachelor / Diploma or Magister / Master / State Examination / PhD]
- How much money do you have at your disposal monthly? (net) [less than 500 euros / 500 euros - 750 euros / 750 euros - 1000 euros / 1000 euros - 1250 euros / 1250 euros - 1500 euros / 1500 euros - 1750 euros / 1750 euros - 2000 euros / more than 2000 euros]
- Which political party do you identify most with? [CDU-CSU / SPD / AfD / FDP / Die Linke / Bündnis90-Die Grünen / other / none]
- In case you would fall sick with the coronavirus: Do you belong to a group of people with an increased risk for a severe case? [yes / no / I don't know]

Thank you for your participation in this study. We need the email address of your PayPal account to be able to transfer the money you earned. As soon as the payment is completed, your email address will be deleted. All data will be stored in an anonymous way.

[*box to enter email address*]

On the next page you will be informed about the exact amount you earned today.

Thanks again for your participation.

As announced, you will earn a guaranteed show-up fee of €2.50. Furthermore, your payment is composed of the following parts:

Out of the three questions about the text you read in the very beginning you answered X questions correctly. This results in an additional payment of €X.

In addition, block X was randomly chosen for your payment. There decision number x was randomly picked to be relevant for you. You decided to X.

Therefore, you will receive a total payment of €X on your PayPal account today and a total payment of €X in exactly 60 days.

The results for the predictions will be published on the 4th of July 2020 on the homepage of the chair for Experimental and Behavioral Economics (https://behavecon.uni-koeln.de). The winners will be paid via PayPal.

If you have questions about the study or your payment please contact harrs@wiso.uni-koeln.de.

Instead of the Xs participants were shown the respective values that applied to them.

D.3 Follow-Up Experiment

We ran a follow-up experiment to provide additional evidence about the underlying mechanisms and examine the generalizability of our results to a different context. In particular, we were interested in the effects of narratives in the context of a different exogenous shock and in the role of emotions as a mechanism for the behavioral effects of narratives (i.e. the effects on risk-taking and patience that we observed in the main experiment). We kept the experimental design as close as possible to the original study, only implementing changes where it was needed in order to allow us to gain the desired insights. The main differences between our main experiment and the follow-up are the narratives used as experimental manipulation and the type of expectations that we measured. The following paragraphs describe the design of the follow-up experiment and its results in detail.

D.3.1 Experimental Design and Data Description

The follow-up experiment was conducted online with N=393 subjects, which were again recruited from the subject pool of the Cologne Laboratory for Economic Research (CLER) via ORSEE (Greiner, 2015). We excluded participants of our first study from this follow-up. Subjects were invited to participate in designated time slots on June 28th, June 30th or July 3rd 2023. Like the main experiment, also the follow-up was implemented with the survey software Qualtrics. The median time for completing the survey was 12 minutes and subjects were paid dependent on their economic decisions with an average of EUR 6.13. Payments were made via PayPal. The follow-up experiment and our hypothesis regarding the results have been pre-registered on OSF.[24]

D.3.1.1 Setting

In contrast to the main experiment, the follow-up included narratives about how a technology shock impacts the business model of a company. In particular, it used narratives about the influence of the technological breakthrough in the development of generative artificial intelligence on the future profitability of Google (Alphabet Inc. A). The development of AI is a setting - just as the appearance of COVID-19 - of high economic and societal relevance and that comes along with high uncertainty about its future impacts. Since ChatGPT - the large language model of OpenAI, that became available from November 2022 onward - was quickly widely used, it was a very salient topic in the news media. In contrast to the COVID-19 setting, this new context - future profitability of Google - however is unlikely to personally affect many of our participants in a similar way the COVID pandemic and the far-reaching restrictions that came with it did. Based on the observation that only a small minority of subjects in our sample reported to own

[24] You can find the pre-registration here: https://doi.org/10.17605/OSF.IO/T4JPN.

Google stocks (7.8%), this setting allows us to study the influence of narratives on financial behavior in a context that is much less personally relevant for participants and thus less emotionally loaded. The idea is to test whether the behavioral effects of narratives found in the main study persist even if emotional reactions are absent.

D.3.1.2 Experimental Procedures

Figure D.10 provides a graphical overview of the experimental procedures of the follow-up.[25] Numbers in brackets in this section refer to the stages of the experiment depicted in Figure D.10. Just as in the main experiment, subjects were exposed to an article at the beginning and were incentivized to memorize it as good as possible within two minutes (2). In contrast to our main study, we only study two types of articles here: one containing an optimistic and one a pessimistic narrative about the impact of generative AI on the profitability of Google. To further mitigate the risk of experimenter demand effects, we emphasized that the content of the articles does not necessarily reflect the view of the experimenters before exposing them to the narrative. Each subject saw and was aware of only one article. More details on the manipulation are provided in section D.3.1.3.

Immediately after the manipulation, we again measured the emotional reactions of subjects (3). Next, we elicited our behavioral outcomes risk aversion (6) and patience (8) in two decision blocks. As we did not find any effects on productivity in the main study we decided not to include this measure in the follow-up. At the end of the experiment, one of the two decision blocks was randomly drawn for each subject and became payoff relevant. Again, we randomized the order of the behavioral outcomes. Then we elicited subjects' expectations about the stock market and their personal circumstances (10) (see section D.3.1.4 for details). The experiment concluded with collecting the socio-demographic characteristics of subjects (11) and eliciting subjects' trust in the provided article and their perceptions of how much they considered the article to be factual or speculative (12).

D.3.1.3 Manipulation

Participants are randomly assigned to one of two conditions: they either read an article that provides an optimistic or a pessimistic narrative about the impacts of the recent technological breakthrough in the development of generative AI on Google. Just as in the main study, both articles were designed symmetrically regarding their content, length, structure and grammatical style as far as possible (see Appendix D.3.4 for the transcripts). The content of the articles is based on a news report from the news agency Reuters about

[25] Transcripts of the instructions are provided in Appendix D.3.4. The original instructions in German are publicly available in our data repository on OSF (`https://osf.io/4d72x`).

FIGURE D.10: Follow-up: Experimental Procedures

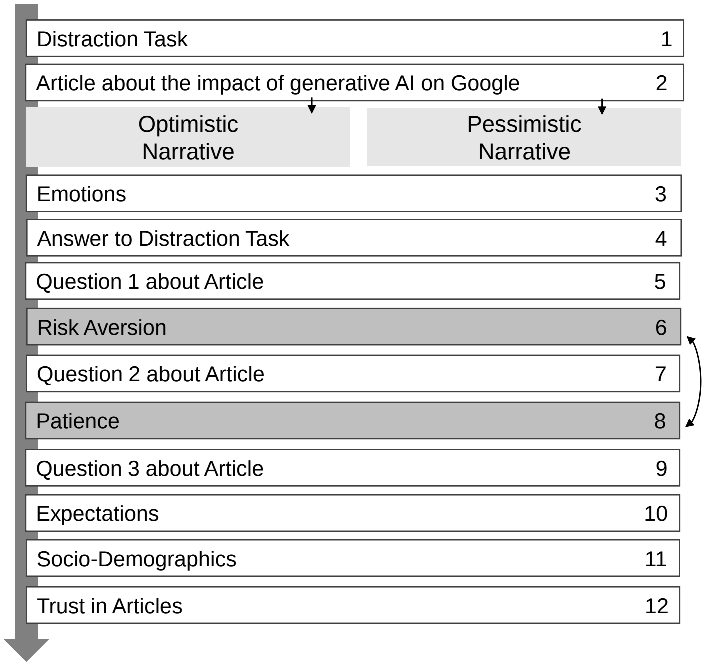

Notes: Figure D.10 gives an overview of the experimental procedures. The numbers on the right side refer to the different stages of the experiment. The manipulation and the main outcomes are shaded in grey. The order of the elicitation of risk aversion and patience was randomized.

the impacts of generative AI on the business model of Google from May 2023.[26] Both articles are complemented with a figure of the historic development of the Google stock price over the past year.

D.3.1.4 Measurement of Expectations

We elicit incentivized 2-month forward-looking expectations on the Google stock (Alphabet Inc. A) and the stock market index S&P500, which includes the shares of the 500 largest publicly traded US companies. To anchor our subjects' estimates, we provide official figures for each of these variables from the previous day as well as the historic development of the previous year. Just as in the main study, we incentivize the expectations in the following way: for each variable three subjects are randomly selected and are paid depending on the accuracy of their expectations (with up to EUR 20) and each subject receives at most a payoff for one of the expectations.

[26]See Reuters https://www.reuters.com/technology/google-expected-unveil-its-answer-microsofts-ai-search-challenge-2023-05-10/ (accessed on August 10th 2023).

D.3.1.5 Measurement of Other Variables

We further elicit our behavioral outcomes (risk aversion and patience), emotions, personal optimism and socio-demographics. Here we stick to the design of the main study and measure all of these previously measured outcomes in exactly the same way. In terms of socio-demographics, we do not ask subjects about whether or not they belong to a risk-group with respect to COVID infections, but instead whether or not they own Google stocks, other stocks or exchange traded funds. We further ask about news consumption habits with respect to news about ChatGPT (instead of COVID).

D.3.1.6 Sample Description and Randomization Check

There was no considerable attrition: Of the 394 participants that started the experiment only one did not complete it. Further, our sample is mostly balanced with respect to covariates. We test the pairwise balance of covariates using either t-tests or Chi^2 tests. Only one of the socio-demographic measures turns out to be different in the two groups on the 5% significance level. There are slightly fewer females in the pessimistic condition (68% vs. 54%). Note that this small imbalance can only be due to chance as we randomized by computer and there was close to no attrition. We address this issue by presenting results where we control for our set of covariates, including gender.

D.3.1.7 Empirical Strategy

We keep our statistical analysis as close as possible to the analysis of the main experiment in order to make the results comparable. In particular, we analyze the behavioral outcomes, personal expectations, and emotions with linear regressions, while our preferred test for the stock market expectations is a Mann-Whitney U test (as it is robust to outliers). Whenever we run regressions with control variables, we use our standard set of controls from the initial experiment with the exception that we do not control for subjects' risk status regarding COVID-19 here, but for owning Google stocks.

D.3.2 Hypotheses

We pre-registered our hypotheses regarding the outcomes of the experiment on OSF (https://osf.io/t4jpn). In our initial experiment, we found evidence that the emotional reactions are the underlying mechanism for the behavioral effects. As the development of the Google stock should be of low personal relevance for most of our participants, we hypothesize that the narratives about Google do change subjects' expectations about the development of the Google stock, but do not to change their emotional reactions, behavioral outcomes or expectations about their personal life. We test the following hypotheses:

- **H1 Google Expectations:** Subjects who read the pessimistic narrative expect the future stock price of Google to be lower than subjects who are exposed to the optimistic narrative.
- **H2 Behavioral Outcomes:** There is no difference in risk aversion and patience of subjects in both treatment groups.
- **H3 Mechanism:** There is no difference in personal optimism and emotions of subjects in both treatment groups.

D.3.3 Results

We present our results in the following order: First, we present the effects of narratives on stock market expectations and then extend the analysis to our behavioral outcomes and the measures that are relevant to our proposed mechanism.

D.3.3.1 Hypothesis 1: Google Expectations

Google Expectations Figure D.11 depicts the mean forward-looking expectations of subjects for the Google stock price in US dollars in two months by treatment condition. As hypothesized, subjects in the pessimistic treatment express lower expectations about Google's performance on the stock market. Specifically, subjects in the pessimistic treatment expect the price of the Google stock to be 7.65 US dollars lower (-6.1%). A Mann-Whitney U-test confirms that the difference in means between the optimistic and pessimistic treatment is significantly different from zero ($p < 0.001$). Corresponding OLS estimates are presented in columns (1) and (2) of Table D.23. This supports Hypothesis 1.

FIGURE D.11: Follow-up: Treatment Effects on Google Stock Expectations

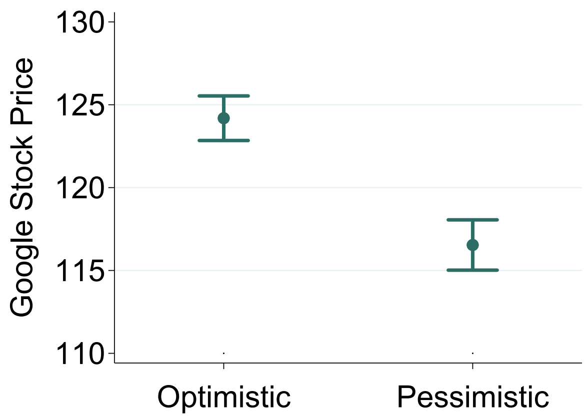

Notes: Figure D.11 shows means and corresponding 95% confidence intervals for expectation on the price of the Google stock (Alphabet Inc. A) in US-dollars in the two treatment conditions.

TABLE D.23: Follow-up: ATE on Stock Market Expectations

	Google Stock				S&P500 Index	
	(1)	(2)	(3)	(4)	(5)	(6)
Pessimistic	-7.65***	-7.92***	-6.68***	-6.90***	-46.85**	-51.51***
	(1.03)	(1.05)	(1.00)	(1.02)	(17.66)	(18.03)
Age		0.07		0.04		1.66
		(0.09)		(0.08)		(1.60)
Female		-1.47		-0.71		-37.96**
		(1.08)		(1.01)		(18.66)
Income		-0.03		-0.18		7.45
		(0.29)		(0.26)		(4.83)
Education		-0.21		0.07		-14.53
		(0.80)		(0.77)		(13.53)
Econ Student		0.66		0.29		18.99
		(1.13)		(1.08)		(18.97)
No Student		0.89		1.16		-13.66
		(1.84)		(1.66)		(33.48)
Political Orientation		-0.27		-0.32		2.61
		(0.35)		(0.33)		(5.80)
Own Google Stock		4.16**		3.29*		44.01
		(1.96)		(1.88)		(33.07)
S&P500 Expectations			0.02***	0.01***		
			(0.00)	(0.00)		
Constant	124.18***	123.00***	34.17**	37.35***	4377.81***	4329.03
	(0.68)	(2.51)	(13.74)	(14.05)	(12.39)	(43.92)
Observations	393	393	393	393	393	393
R^2	0.122	0.148	0.230	0.243	0.017	0.053

Notes: Table reports OLS estimates with robust standard errors in parentheses. The optimistic treatment is the reference group. The measures for the stock market expectations are winsorized at the 95th percentile to minimize the influence of outliers. The Google Stock is measured in US-dollars. The S&P500 index is measured in base points. *** p<0.01, ** p<0.05, * p<0.1

S&P500 Expectations We also examine the effects of our narratives on expectations for the S&P500 index. We find that subjects in the pessimistic treatment expect the S&P500 to close on average 46.85 base points lower in two months than subjects in the optimistic treatment (-1.1%). A Mann-Whitney U-test confirms that the difference in means between the optimistic and pessimistic treatment is significantly different from zero ($p = 0.006$). Thus, we find a much smaller effect of our narratives on the S&P500 index than on the Google stock price. This is what one should expect, as the narratives mainly address the influence on Google, but changing their expectations about Google's future might indirectly impact subjects' expectations about the entire industry.[27] Corresponding OLS estimates are presented in columns (5) and (6) of Table D.23.

As a robustness check, we test whether the effect on the Google stock price expectations

[27]Further, the Google stock is part of the S&P500.

persists even if we control for subjects' general stock market expectations. To do this, we repeat our regression analysis on the Google stock price, but additionally control for the S&P500 expectations. This reveals that even if the general stock market expectations are held constant, subjects in the pessimistic treatment expect the Google stock price to be 6.68 US dollars lower in two months than subjects in the optimistic treatment ($p < 0.001$). Corresponding OLS estimates are available in columns (3) and (4) of Table D.23.

D.3.3.2 Hypothesis 2: Behavioral Outcomes

Risk Aversion Figure D.12a shows the average certainty equivalent elicited for the lottery (EUR 0 with 50% and EUR 4 with 50%) by treatment condition. The average certainty equivalent in the pessimistic treatment is not statistically different from that in the optimistic treatment (EUR 1.93 in pessimistic versus EUR 1.84 in optimistic treatment, $p = 0.210$). This finding, therefore, supports the first part of our hypothesis 2. Corresponding OLS estimates are presented in columns (1) and (2) of Table D.24.

FIGURE D.12: Follow-up: Treatment Effects on Behavioral Outcomes

(A) Certainty Equivalent - Means 95% CI

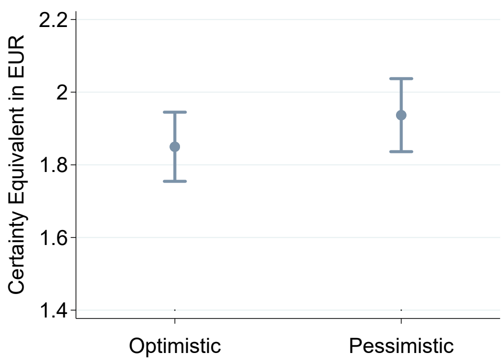

(B) Future Value - Means 95% CI

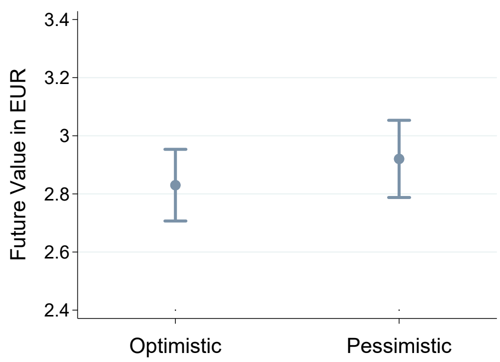

Notes: Figure (a) displays the means and 95% confidence intervals of the certainty equivalent by treatment condition. Figure (b) displays the means and 95% confidence intervals of the future value by treatment condition.

Patience Figure D.12b depicts the mean future value of a EUR 2 payment today by treatment condition. Subjects in the pessimistic treatment do not act statistically significantly more or less patient than in the optimistic treatment (future value of EUR 2.92 versus EUR 2.83, $p = 0.329$). This supports the second part of hypothesis 2. Corresponding OLS estimates are available in columns (3) and (4) of Table D.24.

D.3.3.3 Hypothesis 3: Mechanism

Personal Optimism In contrast to our initial experiment, we do not detect a statistically significant difference in our measure for personal optimism. In the pessimistic treatment, this measure is 0.24 points lower on an 11-point Likert scale ($p = 0.205$) than in the optimistic treatment condition. Corresponding OLS estimates are available in columns (7) and (8) of Table D.24. This finding supports the first part of hypothesis 3.

Emotions While we do find that subjects in our initial experiment were more afraid, upset and nervous in the pessimistic treatment condition, we do not detect any statistically significant differences in terms of affect in this follow-up experiment. This is true for the index on overall affect as well as for each of its components (attentiveness, determination, inspiration, anger, fear and nervousness) individually. Figure D.13 illustrates the means and corresponding 95%-confidence intervals for each of the measured emotional dimensions. OLS estimates for the overall affect index can be found in columns (5) and (6) of Table D.24. Taken together, these findings support the second part of our hypothesis 3.

FIGURE D.13: Follow-up: Treatment Effects on Emotions

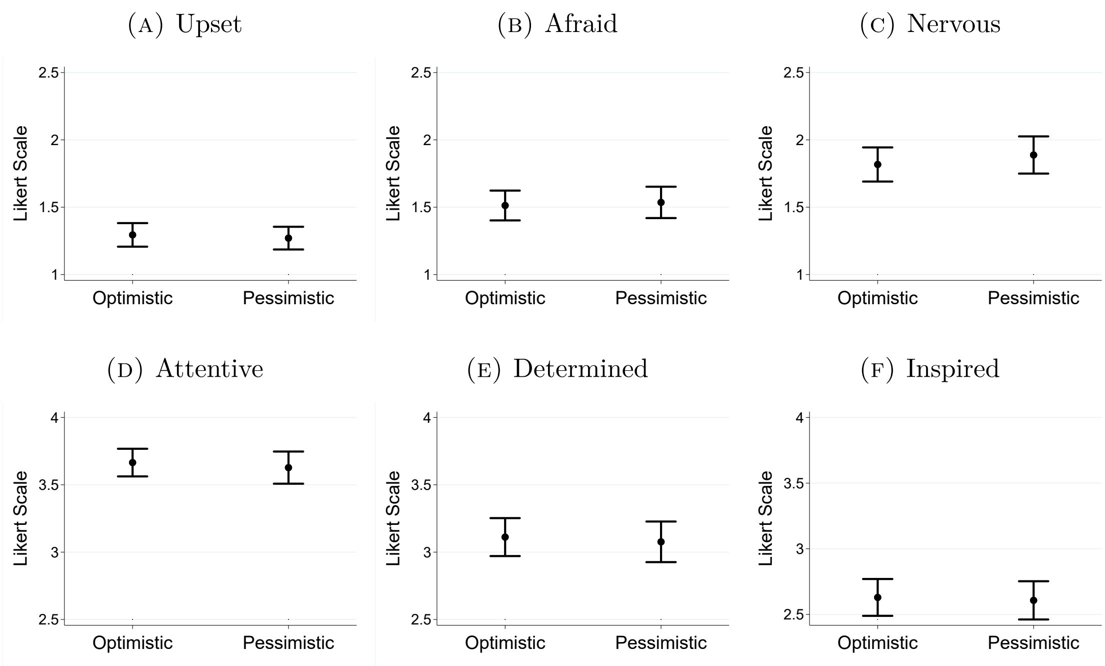

Notes: Figure D.13 shows means and corresponding 95% confidence intervals for the emotional state of subjects across treatment conditions: for feeling (a) upset, (b) afraid, (c) nervous, (d) attentive, (e) determined and (f) inspired. Emotions are measured on a 5-point Likert scale (1 “not at all” to 5 “very much”).

Table D.24: Follow-up: ATE on Behavioral Outcomes, Emotions, and Optimism

	Certainty Equivalent		Future Value		Affect Index		Personal Optimism	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pessimistic	0.08	0.09	0.09	0.12	-0.16	-0.19	-0.24	-0.27
	(0.07)	(0.07)	(0.09)	(0.09)	(0.29)	(0.30)	(0.19)	(0.19)
Age		0.00		0.01		0.00		-0.00
		(0.00)		(0.00)		(0.02)		(0.01)
Female		0.04		0.13		-0.53*		-0.31
		(0.07)		(0.09)		(0.30)		(0.20)
Income		0.02		-0.03		0.09		0.04
		(0.01)		(0.02)		(0.07)		(0.05)
Education		-0.17***		0.03		0.19		-0.02
		(0.05)		(0.06)		(0.07)		(0.21)
Econ Student		0.18**		-0.07		0.16		-0.10
		(0.07)		(0.10)		(0.33)		(0.33)
No Student		0.09		-0.40**		0.35		0.34
		(0.13)		(0.17)		(0.52)		(0.35)
Political Orientation		-0.03		-0.02		-0.07		-0.01
		(0.02)		(0.03)		(0.10)		(0.06)
Own Google Stock		-0.28**		-0.15		0.47		-0.04
		(0.13)		(0.16)		(0.57)		(0.30)
Constant	1.84***	1.60***	2.83***	2.64***	4.78***	4.55***	1.51***	1.63***
	(0.04)	(0.20)	(0.06)	(0.21)	(0.19)	(0.65)	(0.12)	(0.52)
Observations	393	393	393	393	393	393	393	393
R^2	0.003	0.063	0.002	0.032	0.000	0.032	0.004	0.018

Notes: Table reports OLS estimates with robust standard errors in parentheses. The optimistic treatment is always the reference group. *** p<0.01, ** p<0.05, * p<0.1

D.3.3.4 Additional Analyses

Trustworthiness and Speculativeness Our data on the trustworthiness of the provided narratives suggest that most subjects trust the articles. 73.5 percent of them say they found the article either “rather trustworthy” or “very trustworthy”. At the same time, they seem to be aware that they are not provided with pure hard facts. When asked how speculative (as opposed to factual) they perceived the article, 64.6 percent state that it was “rather speculative” or “very speculative”. These perceptions of the articles do not differ between treatment conditions: 73.1% in optimistic state “rather” or “very trustworthy” versus 74.0% in pessimistic (Chi^2-test: p=0.843), and 64.5% in optimistic state “rather” or “very speculative” versus 64.8% in pessimistic, (Chi^2-test: p=0.946).

D.3.4 Instructions Follow-Up

Just as in the initial experiment, participants received experimental instructions in German. The original instructions are publicly available in our data repository on OSF (https://osf.io/4d72x). Below we provide an English transcript. A dashed line indicates the next page of the survey. Explanatory comments (which were not displayed to participants) are indicated by blue, italic text.

Welcome to this online experiment! You will receive €2.50 as a show-up fee. Depending on the decisions you take during this study, you can earn an additional payment. As described in the invitation, your payment will be transferred to your PayPal account. Therefore, you will be asked to provide the email address of your PayPal account at the end of this survey. Please make sure you know the email address of you PayPal account before you begin. Taking part in this experiment is only possile via a computer or laptop. Your participation will take approximately 15 minutes.

O I consent to the above conditions.

Participants could only continue when they gave their consent and if they did not use smartphone browsers (i.e. Safari iPhone).

- -

As start of this study two telephone numbers will be displayed. Please try to memorise the numbers. You will have 20 seconds to do so.

You will be asked to recognise the two numbers at a later point in time.

- -

05454/444-54 08421/792-65

Participants were automatically forwarded after 20 seconds. A timer indicated the remaining time participants had on this page.

- -

On the next page a topical text will be displayed. This text is based on a recent story by the news agency *Reuters*. The text does not necessarily reflect the view of the experimenter. Please try to memorise as much of the content as possible. You will have two minutes to do so. At a later point in time you will be asked to answer three questions about the content of the text. You will earn €0.50 per correct answer.

- -

One of the two treatment manipulations was randomly selected and displayed. Participants could not leave this page independently (skip the text). Subjects were automatically forwarded after two minutes. A timer indicated the remaining time they had left to read the text.

Optimistic

AI Revolution: Google Shapes the Future with Generative AI

Google recently announced that it will integrate generative artificial intelligence (AI) into all of its products, which is expected to contribute to future growth of the tech giant. The development of generative AI will enhance the Google search engine by providing intuitive answers to open-ended questions.

CEO Sundar Pichai showcased this exciting project at the annual Google I/O developer conference. Google aims to redefine all its products through the integration of AI. The inclusion of generative AI is also expected to improve services like Gmail and Google Photos, reinforcing Google's position as a leading technology company.

The initiative is a response to the success of the chatbot ChatGPT, created by the startup OpenAI. Google's AI offensive opens up unique business opportunities by integrating the new technology into existing applications and making it accessible to a large user base.

The decision to integrate AI into all its products will revolutionize the way users interact with Google's services. Google aims to expand its dominance in the billion-dollar online advertising market.

Google's innovative move in the field of generative AI demonstrates the adaptability of the company. Analysts believe that this adaptability and the willingness to innovate will secure the search giant's future market power.

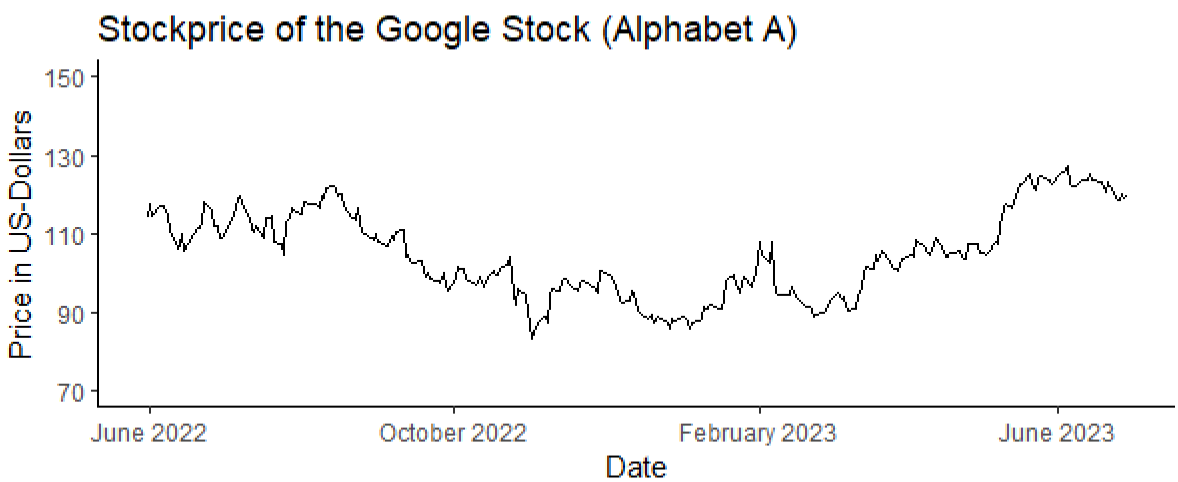

Pessimistic

AI Revolution: Generative AI Threatens Google's Business Model

Google recently announced that it will integrate generative artificial intelligence (AI) into all of its products, which has however sparked concerns about the tech giant's future profitability. The development of generative AI, which provides intuitive answers to open-ended questions, threatens to fundamentally challenge Google's business model.

Despite assurances from CEO Sundar Pichai at the annual Google I/O developer conference, the integration of generative AI into Google's product range, including Gmail and Google Photos, poses significant risks. In particular, the high costs of large AI models could harm Google's profitability.

The initiative is a response to the success of the chatbot ChatGPT, created by the startup OpenAI. Google's belated entry into generative AI and the technological lead of competitors threaten Google's dominance.

Google's decision to alter its main source of income - Google Search - also carries substantial risks. Changes to Google Search could disrupt the existing business model and cost market shares in the billion-dollar online advertising market.

Google's belated entry into the field of generative AI is thus fraught with considerable uncertainties. Analysts are already warning that the technological revolution in AI endangers Google's financial prospects.

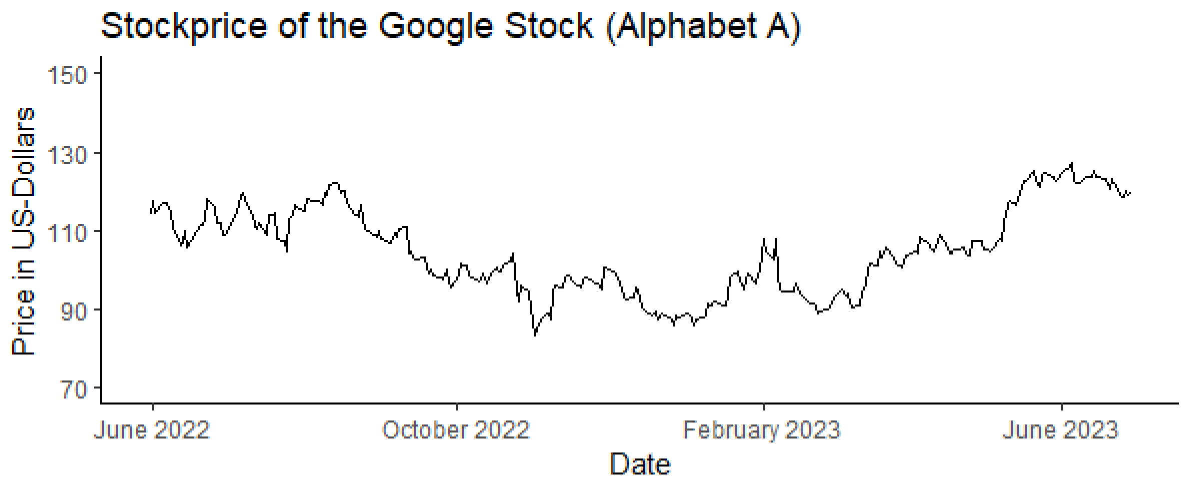

We would like to know how you feel right now.

The following words describe different feelings and sensations. Read every word, then indicate the intensity with which you experience the respective emotion at the moment. You can choose between five gradations.

	not at all	a little	somewhat	much	very much
upset	○	○	○	○	○
attentive	○	○	○	○	○
afraid	○	○	○	○	○
determined	○	○	○	○	○
nervous	○	○	○	○	○
inspired	○	○	○	○	○

At the beginning of the study you were shown two telephone numbers. Which of the following telephone numbers?

O 02235/679-89 O 0721/972-56

O 08421/792-65 O 05454/444-54

Now you will make decisions in three blocks. At the end of the study, one of the three blocks will be randomly selected. Only the decisions made in the selected block will be relevant for the variable part of your payment.

This means that every decision that you will make can potentially influence the payment you receive. You should therefore take all decisions as if they would be implemented.

In-between the decision blocks we will ask questions about the text that you have read in the beginning.

Next is a question on the text that you read in the beginning.

The following statement was made or was contained in the text in this or in a similar fashion:

– sentence –

O True O False

Instead of – sentence – a sentence from the treatment text was displayed to participants. This sentence was the same independently of which outcome was elicited in block 1. The correct answer for all statements is "True". The sentences displayed were:

Pessimistic: Google recently announced that it will integrate generative artificial intelligence (AI) into all of its products, which has however sparked concerns about the tech giant's future profitability.

Optimistic: Google recently announced that it will integrate generative artificial intelligence (AI) into all of its products, which is expected to contribute to future growth of the tech giant.

Decision block 1 starts now.

The order of the elicitation of risk taking and patience was randomized. Thus, block 1 could contain any of the two behavioral outcomes. As an example, we are presenting the elicitation of risk aversion here.

In this block you will take five decisions. You will always have the choice between a guaranteed payment and a lottery which pays €4 with 50 percent chance and €0 with 50 percent chance.

In this block one of your five decisions is randomly selected to be considered for payment.

Do you want to receive a guaranteed payment of €X or play a lottery with 50 percent chance for €4 and 50 percent chance for €0?

O €X as guaranteed payment   O lottery with 50% chance for €4 and 50% chance for €0

This question was displayed five times with different values for X. The first value for X was €1.65 and subsequent values depended on the previous decisions. Figure A3 in Appendix A3 shows the game tree.

Next is a question on the text that you read in the beginning.

The following statement was made or was contained in the text in this or in a similar fashion:

– sentence –

O True   O False

Instead of – sentence – a sentence from the treatment text was displayed to participants. This sentence was the same independently of which outcome was elicited in block 1. The correct answer for all statements is "True". The sentences displayed were:

Pessimistic: Google's belated entry into generative AI and the technological lead of competitors threaten Google's dominance.

Optimistic: Google's AI offensive opens up unique business opportunities by integrating the new technology into existing applications and making it accessible to a large user base.

Decision block 2 starts now.

Here the remaining behavioral outcome was elicited. As an example we are presenting the elicitation of patience here.

In this block you will take five decisions. You always have the choice between a payment you receive directly after your participation in this study and a payment you receive in 2 months (in exactly 60 days). In both cases the money will be transferred to your PayPal account.

One of your five decisions is randomly selected to be considered for payment.

Do you want to receive €2 today or €X in two months?

O €2 today O €X in two months

This question was displayed five times with different values for X. The first value for X was €3.32 and subsequent values depended on the previous decisions. Figure A4 in Appendix A3 shows the game tree.

Next is a question on the text that you read in the beginning.

The following statement was made or was contained in the text in this or in a similar fashion:

– sentence –

O True O False

Instead of – sentence – a sentence from the treatment text was displayed to participants. This sentence was the same independently of which outcome was elicited in block 2. The correct answer for all statements is "True". The sentences displayed were:

*Pessimistic:*Analysts are already warning that the technological revolution in AI endangers Google's financial prospects.

Optimistic: Analysts believe that this adaptability and the willingness to innovate will secure the search giant's future market power.

- -

The two decision blocks are completed. You now have an opportunity to earn an additional variable payment by making a number of predictions.

- -

You are now asked to make two predictions about the development of figures regarding the stock market until the 2nd of September 2023 (this is in exactly 60 days). Three participants will be selected randomly for each question and will be paid depending on the accuracy of their predictions. The closer the prediction is to the realized value, the higher the payment will be. You can win up to €20 with your predictions.

Note: Your payment is independent of what other participants predict. You should therefore state the value which you regard as most likely for each figure. For the selection of the winners, only one of your predictions will be considered. Therefore it is not possible to spread your risk across predictions and you cannot win multiple times.

We will use official data from the New York Stock Exchange to evaluate the predictions.

Question 1:

What do you think: What will be the value in US dollars of the Google stock (Alphabet Inc A) on 2nd of September 2023 (in 60 days)?

[*open text box*]

On 2nd of July the stock value of the Google stock (Alphabet Inc A) was 120 US dollars.

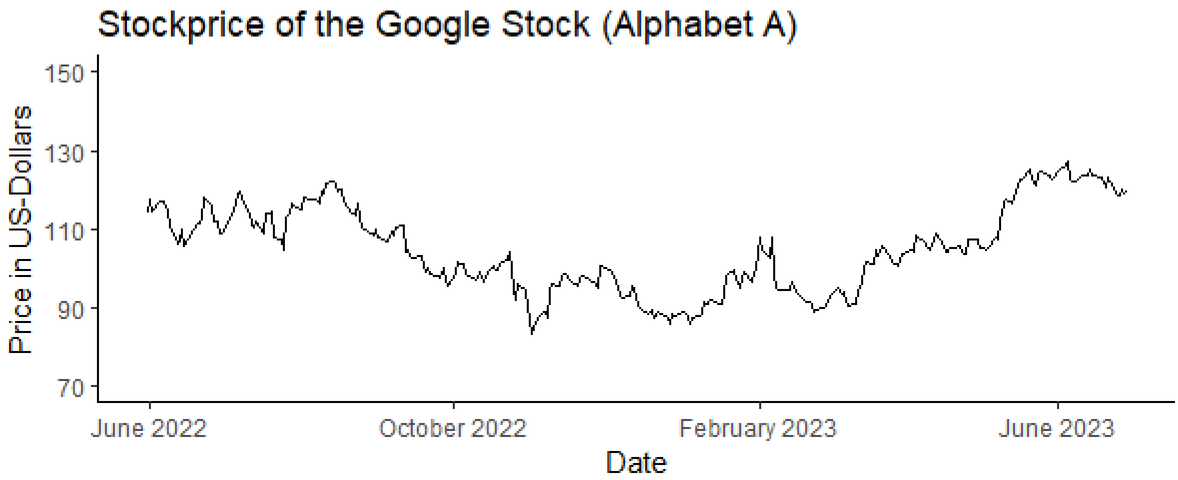

- -

Question 2:

What do you think: What will be the value of the S&P500 index, which includes the 500 biggest US companies, on 2nd of September 2023 (in 60 days)?

[*open text box*]

On 2nd of July the point value of the S&P500 index was 4450 points.

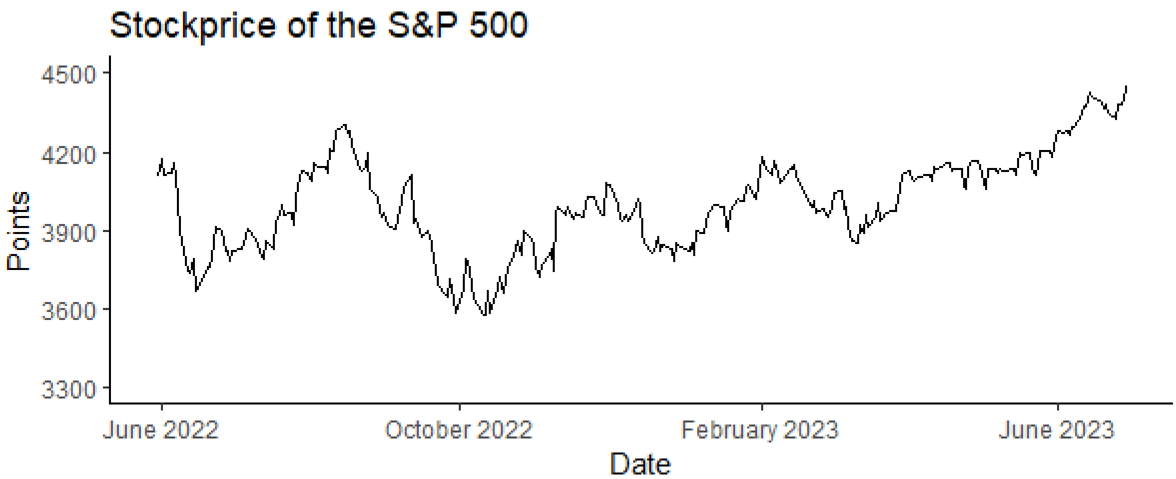

Think about your personal circumstances in the next weeks. To what extent do you expect things to develop positively or negatively?

Participants had to answer the above question on an 11-point Likert scale: from very negatively (-5) to very positively (+5).

Thank you! Now we have a few questions about you:

- How old are you? [ ]
- Which gender do you identify with? [male / female / diverse]
- What is your subject of studies? (If more than one: Major) [all subjects that can be studied at the University of Cologne]
- What is your highest educational achievement? [No formal degree / Secondary Modern School / Junior High School / A-levels / Master Craftsmen / Bachelor / Diploma or Magister / Master / State Examination / PhD]

- How much money do you have at your disposal monthly? (net) [less than 500 euros / 500 euros - 750 euros / 750 euros - 1000 euros / 1000 euros - 1250 euros / 1250 euros - 1500 euros / 1500 euros - 1750 euros / 1750 euros - 2000 euros / more than 2000 euros]
- Which political party do you identify most with? [CDU-CSU / SPD / AfD / FDP / Die Linke / Bündnis90-Die Grünen / other / none]

Do you own any stocks or index funds (ETFs)? [Yes / No]

Do you own stocks of Alphabet Inc (Google)? [Yes / No]

How often did you inform yourself about ChatGPT the last months?

Participants had to answer the above question on an 5-point Likert scale: never (1), seldom (2), sometimes (3), often (4), very often (5).

Now a few questions about the text on Google that you read at the beginning of the experiment:

To what extent has the text raised or lowered your expectations about the future price trend of Google stock (Alphabet Inc A)?

Participants had to answer the above question on an 5-point Likert scale: strongly lowered (-2), rather lowered (-1), neither nor (0), rather raised (+1), strongly raised (+2).

How trustworthy do you consider the text you read about Google at the beginning of the experiment?

Participants had to answer the above question on an 4-point Likert scale: very untrustworthy (1), rather untrustworthy (2), rather trustworthy (3), very trustworthy (4).

How fact-based or speculative do you consider the text you read about Google at the beginning of the experiment?

Participants had to answer the above question on an 4-point Likert scale: very speculative (1), rather speculative (2), fact-based (3), very fact-based (4).

Thanks again for your participation.

As announced, you will earn a guaranteed show-up fee of €2.50. Furthermore, your payment is composed of the following parts:

Out of the three questions about the text you read in the very beginning you answered X questions correctly. This results in an additional payment of €X.

In addition, block X was randomly chosen for your payment. There decision number x was randomly picked to be relevant for you. You decided to X.

Therefore, you will receive a total payment of €X on your PayPal account today and a total payment of €X in exactly 60 days.

The results for the predictions will be published on the 4th of July 2020 on the homepage of the chair for Experimental and Behavioral Economics (https://behavecon.uni-koeln.de). The winners will be paid via PayPal.

If you have questions about the study or your payment please contact harrs@wiso.uni-koeln.de.

Next, you will be redirected to the homepage of the Cologne Laboratory for Economic Research, where you will receive your payout code and can enter your PayPal address. Please do this immediately after this experiment, so that we can carry out your payout.

Instead of the Xs participants were shown the respective values that applied to them.

Bibliography

Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin (2021). Institutional change and institutional persistence. In: *The Handbook of Historical Economics*. Elsevier, pp. 365–389.

Aina, Chiara (2023). Tailored Stories. *Working Paper*.

Akerlof, George A. and Rachel Kranton (2000). Economics and Identity. *The Quarterly Journal of Economics* 115(3), pp. 715–753.

Alesina, Alberto and George-Marios Angeletos (2005). Fairness and redistribution. *American Economic Review* 95(4), pp. 960–980.

Alesina, Alberto, Guido Cozzi, and Noemi Mantovan (2012). The evolution of ideology, fairness and redistribution. *The Economic Journal* 122(565), pp. 1244–1261.

Alesina, Alberto and Paola Giuliano (2011). Preferences for redistribution. *Handbook of Social Economics* 1, pp. 93–132.

Alesina, Alberto and Eliana La Ferrara (2005). Preferences for redistribution in the land of opportunities. *Journal of Public Economics* 89, pp. 897–931.

Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso (2018). Intergenerational mobility and preferences for redistribution. *American Economic Review* 108(2), pp. 521–554.

Almås, Ingvild, Alexander Cappelen, and Bertil Tungodden (2020). Cutthroat capitalism versus cuddly socialism: Are Americans more meritocratic and efficiency-seeking than Scandinavians? *Journal of Political Economy* 128(5).

Almås, Ingvild, Alexander W Cappelen, Erik Ø Sørensen, and Bertil Tungodden (2010). Fairness and the development of inequality acceptance. *Science* 328(5982), pp. 1176–1178.

Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutström (2008). Eliciting Risk and Time Preferences. *Econometrica* 76(3), pp. 583–618.

Andre, Peter (2024). Shallow meritocracy: An experiment on fairness views. *The Review of Economic Studies* Forthcoming.

Andre, Peter, Ingar Haaland, Christopher Roth, and Johannes Wohlfart (2023). Narratives About the Macroeconomy. *CESifo Working Paper No. 10535*.

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart (2022). Subjective Models of the Macroeconomy: Evidence From Experts and Representative Samples. *The Review of Economic Studies* 89(6), pp. 2958–2991.

Andreoni, James and Charles Sprenger (2012). Estimating time preferences from convex budgets. *American Economic Review* 102(7), pp. 3333–3356.

Andreoni, James and Charles Sprenger (2015). Risk preferences are not time preferences: Reply. *American Economic Review* 105(7), pp. 2287–2293.

Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar (2016). The price is right: Updating inflation expectations in a randomized price information experiment. *Review of Economics and Statistics* 98(3), pp. 503–523.

Bailey, Michael, Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel (2019). House Price Beliefs And Mortgage Leverage Choice. *The Review of Economic Studies* 86(6), pp. 2403–2452.

Barr, Abigail, Luis Miller, and Paloma Ubeda (2016). Moral consequences of becoming unemployed. *Proceedings of the National Academy of Sciences of the United States of America* 113(17), pp. 4676–4681.

Barron, Kai and Tilman Fries (2023). Narrative persuasion. *CESifo Working Paper No. 10206*.

Barron, Kai, Heike Harmgart, Steffen Huck, Sebastian O Schneider, and Matthias Sutter (2020). Discrimination, narratives and family history: An experiment with Jordanian host and Syrian refugee children. *The Review of Economics and Statistics*, pp. 1–34.

Bartling, Björn, Alexander W Cappelen, Henning Hermes, and Bertil Tungodden (2023). Free to fail? Paternalistic preferences in the United States. *NHH Dept. of Economics Discussion Paper No. 09/2023*.

Bénabou, Roland, Armin Falk, and Jean Tirole (2020). Narratives , Imperatives , and Moral Persuasion. *IZA Discussion Paper No. 11665*.

Bénabou, Roland and Jean Tirole (2006). Belief in a just world and redistributive politics. *The Quarterly Journal of Economics* 121(2), pp. 699–746.

Bénabou, Roland and Jean Tirole (2011). Identity, morals, and taboos: Beliefs as assets. *The Quarterly Journal of Economics* 126(2), pp. 805–855.

Binder, Carola (2020). Coronavirus fears and macroeconomic expectations. *Review of Economics and Statistics* 102(4), pp. 721–730.

Bursztyn, Leonardo, Ingar K Haaland, Aakaash Rao, and Christopher Roth (2020). Disguising prejudice: Popular rationales as excuses for intolerant expression. *NBER Working Paper No. 27288*.

Cahlíková, Jana and Lubomír Cingl (2017). Risk preferences under acute stress. *Experimental Economics* 20(1), pp. 209–236.

Cappelen, Alexander W., Cornelius Cappelen, and Bertil Tungodden (2023). Second-best fairness: The trade-off between false positives and false negatives. *American Economic Review* 113 (9), pp. 2458–2485.

Cappelen, Alexander W., Ranveig Falch, Erik Ø. Sørensen, and Bertil Tungodden (2021). Solidarity and fairness in times of crisis. *Journal of Economic Behavior & Organization* 186, pp. 1–11.

Cappelen, Alexander W, Sebastian Fest, Erik Ø Sørensen, and Bertil Tungodden (2022a). Choice and personal responsibility: What is a morally relevant choice? *Review of Economics and Statistics* 104(5), pp. 1110–1119.

Cappelen, Alexander W., Astri D. Hole, Erik Ø. Sørensen, and Bertil Tungodden (2007). The pluralism of fairness ideals: An experimental approach. *American Economic Review* 97(3), pp. 818–827.

Cappelen, Alexander W., James Konow, Erik Ø. Sørensen, and Bertil Tungodden (2013). Just luck: An experimental study of risk-taking and fairness. *American Economic Review* 103(4), pp. 1398–1413.

Cappelen, Alexander W., Johanna Mollerstrom, Bjørn-Atle Reme, and Bertil Tungodden (2022b). A meritocratic origin of egalitarian behaviour. *The Economic Journal* 132 (646), pp. 2101–2117.

Charness, Gary, Thomas Garcia, Theo Offerman, and Marie Claire Villeval (2020). Do measures of risk attitude in the laboratory predict behavior under risk in and outside of the laboratory? *Journal of Risk and Uncertainty* 60, pp. 99–123.

Chuang, Yating and Laura Schechter (2015). Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results. *Journal of Development Economics* 117, pp. 151–170.

Clarke, Damian, Joseph Romano, and Michael Wolf (2019). The Romano-Wolf Multiple Hypothesis Correction in Stata. *IZA Discussion Paper No. 12845.*

Cohn, Alain, Jan Engelmann, Ernst Fehr, and Michel Maréchal (2015). Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals. *American Economic Review* 105(2), pp. 860–885.

Cohn, Alain, Lasse J Jessen, Marko Klašnja, and Paul Smeets (2023). Wealthy Americans and redistribution: The role of fairness preferences. *Journal of Public Economics* 225, p. 104977.

Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele (2019). Inflation Expectations and Firm Decisions: New Causal Evidence. *The Quarterly Journal of Economics* 135(1), pp. 165–219.

de Quidt, Jonathan, Johannes Haushofer, and Christopher Roth (2018). Measuring and Bounding Experimenter Demand. *American Economic Review* 108(11), pp. 3266–3302.

DellaVigna, Stefano and Matthew Gentzkow (2010). Persuasion: Empirical Evidence. *Annual Review of Economics* 2, pp. 643–69.

DellaVigna, Stefano and Devin Pope (2018). What motivates effort? Evidence and expert forecasts. *The Review of Economic Studies* 85(2), pp. 1029–1069.

Douglas, Benjamin D, Patrick J Ewell, and Markus Brauer (2023). Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. *Plos one* 18(3), e0279720.

Eliaz, Kfir and Ran Spiegler (2020). A Model of Competing Narratives. *American Economic Review* 110(12), pp. 3786–3816.

Enke, Benjamin, Ricardo Rodríguez-Padilla, and Florian Zimmermann (2023). Moral universalism and the structure of ideology. *The Review of Economic Studies* 90(4), pp. 1934–1962.

Epper, Thomas, Ernst Fehr, and Julien Senn (2020). Social Preferences and Redistributive Politics. *CESifo Working Paper No. 9545.*

Ericson, Keith Marzilli and David Laibson (2019). Intertemporal choice. In: *Handbook of Behavioral Economics*. Vol. 2. Elsevier, pp. 1–67.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde (2018). Global Evidence on Economic Preferences. *The Quarterly Journal of Economics* 133(4), pp. 1645–1692.

Fehr, Ernst, Helen Bernhard, and Bettina Rockenbach (2008). Egalitarianism in young children. *Nature* 454(7208), pp. 1079–1083.

Fehr, Ernst and Gary Charness (2023). Social preferences: Fundamental characteristics and economic consequences. *CESifo Working Paper No. 10488*.

Fehr, Ernst, Daniela Glätzle-Rützler, and Matthias Sutter (2013). The development of egalitarianism, altruism, spite and parochialism in childhood and adolescence. *European Economic Review* 64, pp. 369–383.

Fetzer, Thiemo, Lukas Hensel, Johannes Hermle, and Christopher Roth (2021). Coronavirus perceptions and economic anxiety. *Review of Economics and Statistics* 103(5), pp. 968–978.

Fisman, Raymond, Pamela Jakiela, and Shachar Kariv (2017). Distributional preferences and political behavior. *Journal of Public Economics* 155, pp. 1–10.

Fong, Christina (2001). Social preferences, self-interest, and the demand for redistribution. *Journal of Public Economics* 82(2), pp. 225–246.

Fong, Christina M and Panu Poutvaara (2019). Redistributive politics with target-specific beliefs. *Available at SSRN*.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature* 40(2), pp. 351–401.

Fuster, Andreas and Basit Zafar (2023). Survey experiments on economic expectations. In: *Handbook of Economic Expectations*. Elsevier, pp. 107–130.

Galperti, Simone (2019). Persuasion: The Art of Changing Worldviews. *American Economic Review* 109(3), pp. 996–1031.

Gentzkow, Matthew, Jesse M Shapiro, and Daniel F Stone (2015). Media bias in the marketplace: Theory. In: *Handbook of Media Economics*. Vol. 1. Elsevier, pp. 623–645.

Giuliano, Paola. and Antonio Spilimbergo (2014). Growing up in a recession. *The Review of Economic Studies* 81(2), pp. 787–817.

Greiner, Ben (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association* 1(1), pp. 114–125.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2018). Time varying risk aversion. *Journal of Financial Economics* 128(3), pp. 403–421.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart (2023). Designing information provision experiments. *Journal of Economic Literature* 61(1), pp. 3–40.

Harrison, Glenn W, Andre Hofmeyr, Harold Kincaid, Brian Monroe, Don Ross, Mark Schneider, and J Todd Swarthout (2022). Subjective beliefs and economic preferences during the COVID-19 pandemic. *Experimental Economics* 25(3), pp. 795–823.

Huber, Christoph, Jürgen Huber, and Michael Kirchler (2021). Market Shocks and Professionals' Investment Behavior - Evidence from the COVID-19 Crash. *Journal of Banking & Finance* 133, pp. 106–247.

Hvidberg, Kristoffer B., Claus T. Kreiner, and Stefanie Stantcheva (2023). Social position and fairness views on inequality. *The Review of Economic Studies* 90 (6), pp. 3083–3118.

Ifcher, John and Homa Zarghamee (2011). Happiness and time preference: The effect of positive affect in a random-assignment experiment. *American Economic Review* 101(7), pp. 3109–3129.

Jost, John T., Christopher M. Federico, and Jaime L. Napier (2009). Political ideology: Its structure, functions, and elective affinities. *Annual Review of Psychology* 60(1), pp. 307–337.

Kamenica, Emir and Matthew Gentzkow (2011). Bayesian Persuasion. *American Economic Review* 101(6), pp. 2590–2615.

Kerschbamer, Rudolf and Daniel Müller (2020). Social preferences and political attitudes: An online experiment on a large heterogeneous sample. *Journal of Public Economics* 182, p. 104076.

Konow, James (2000). Fair shares: Accountability and cognitive dissonance in allocation decisions. *American Economic Review* 90(4), pp. 1072–1091.

Kosse, Fabian, Thomas Deckers, Pia Pinger, Hannah Schildberg-Hörisch, and Armin Falk (2020). The formation of prosociality: causal evidence on the role of social environment. *Journal of Political Economy* 128(2), pp. 434–467.

Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *American Economic Review* 105(4), pp. 1478–1508.

Laudenbach, Christine, Annika Weber, and Johannes Wohlfart (2021). Beliefs About the Stock Market and Investment Choices: Evidence from a Field Experiment. *Available at SSRN*.

Levitt, Steven D and John A List (2007). What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic perspectives* 21(2), pp. 153–174.

Li, King King, Ying-yi Hong, Bo Huang, and Tony Tam (2022). Social preferences before and after the onset of the COVID-19 pandemic in China. *Judgment and Decision Making* 17(6), pp. 1313–1333.

Lucas, Robert E and Thomas J Sargent (1981). Rational Expectations and Econometric Practice. Vol. 2. University of Minnesota Press.

Malmendier, Ulrike (2021). FBBVA Lecture 2020 Exposure, experience, and expertise: Why personal histories matter in economics. *Journal of the European Economic Association* 19(6), pp. 2857–2894.

Meier, Armando N (2022). Emotions and risk attitudes. *American Economic Journal: Applied Economics* 14(3), pp. 527–558.

Meltzer, Allan H and Scott F Richard (1981). A rational theory of the size of government. *Journal of Political Economy* 89(5), pp. 914–927.

Mullainathan, Sendhil and Andrei Shleifer (2005). The market for news. *American Economic Review* 95(4), pp. 1031–1053.

Nelson, Phillip (1974). Advertising as information. *Journal of Political Economy* 82(4), pp. 729–754.

Piketty, Thomas (1995). Social mobility and redistributive politics. *The Quarterly Journal of Economics* 110(3), pp. 551–584.

Rafkin, Charlie, Advik Shreekumar, and Pierre-Luc Vautrey (2021). When guidance changes: Government stances and public beliefs. *Journal of Public Economics* 196, p. 104319.

Roemer, John E and Alain Trannoy (2015). Equality of Opportunity. In: *Handbook of Income Distribution*. Vol. 2. Elsevier, pp. 217–300.

Roth, Christopher and Johannes Wohlfart (2018). Experienced inequality and preferences for redistribution. *Journal of Public Economics* 167, pp. 251–262.

— (2020). How do expectations about the Macroeconomy affect personal expectations and Behavior? *Review of Economics and Statistics* 102(4), pp. 731–748.

Sandel, Michael (2020). The tyranny of merit: What's become of the common good? Penguin Books Ltd.

Schildberg-Hörisch, Hannah (2018). Are risk preferences stable? *Journal of Economic Perspectives* 32(2), pp. 135–154.

Schneider, Sebastian O and Matthias Sutter (2020). Higher order risk preferences: New experimental measures, determinants and field behavior. *MPI Collective Goods Discussion Paper No. 2020/22*.

Schwartzstein, Joshua and Adi Sunderam (2021). Using Models to Persuade. *American Economic Review* 111(1), pp. 276–323.

Shachat, Jason, Matthew J. Walker, and Lijia Wei (2021a). How the onset of the Covid-19 pandemic impacted pro-social behaviour and individual preferences: Experimental evidence from China. *Journal of Economic Behavior & Organization* 190, pp. 480–494.

— (2021b). The Impact of an Epidemic: Experimental Evidence on Preference Stability from Wuhan. *AEA Papers and Proceedings* 111, pp. 302–306.

Shiller, Robert J. (2017). Narrative Economics. *American Economic Review* 107(4), pp. 967–1004.

Sims, Christopher A (2003). Implications of rational inattention. *Journal of Monetary Economics* 50, pp. 665–690.

Snowberg, Erik and Leeat Yariv (2021). Testing the waters: Behavior across participant pools. *American Economic Review* 111(2), pp. 687–719.

Stantcheva, Stefanie (2021). Understanding tax policy: How do people reason? *The Quarterly Journal of Economics* 136(4), pp. 2309–2369.

— (2023). How to run surveys: A guide to creating your own identifying variation and revealing the invisible. *Annual Review of Economics* 15, pp. 205–234.

Stigler, George J (1961). The economics of information. *Journal of Political Economy* 69(3), pp. 213–225.

Thompson, Edmund R. (2007). Development and Validation of an Internationally Reliable Short-Form of the Positive and Negative Affect Schedule (PANAS). *Journal of Cross-Cultural Psychology* 38(2), pp. 227–242.

Verrina, Eugenio and Adrian Hillenbrand (2020). The Differential Effect of Narratives on Prosocial Behavior. *MPI Collective Goods Discussion Paper No. 2018/16.*

Winter, Nicholas, Tyler Burleigh, Ryan Kennedy, and Scott Clifford (2019). A simplified protocol to screen out VPS and international respondents using Qualtrics. *Available at SSRN.*

List of Applied Software

- Stata 15.1: used to perform statistical analyses for all chapters.
- Qualtrics: used for data collection for chapters 2, 4, and 5.
- zTree and zTree (unleashed): used for data collection for chapter 3.
- ORSEE: used to recruit participants for chapters 3 and 5
- Prolific: used to recruit participants for chapter 2.
- GPower3: used for power calculations.
- Overleaf: used to compile the papers and the final thesis.

Data and Code Availability

The data, code, and materials used to produce the results in each chapter are stored at the chair of Prof. Rockenbach, and are available upon request. In addition, they are publicly available as follows:

- **Chapter 2:** The data, code, and materials will be made publicly available upon publication of the paper in an academic journal.
- **Chapter 3:** The data, code, and materials are publicly available at the Open Science Framework under `https://doi.org/10.17605/OSF.IO/KZ2FH`
- **Chapter 4:** The data, code, and materials are publicly available at the Open Science Framework under `https://doi.org/10.17605/OSF.IO/BDMVT`
- **Chapter 5:** The data, code, and materials are publicly available at the Open Science Framework under `https://doi.org/10.17605/OSF.IO/NJ2SQ`