

An assemblage of discrimination and social preferences

Citation for published version (APA):

Gagnon, N. (2020). *An assemblage of discrimination and social preferences*. [Doctoral Thesis, Maastricht University]. Maastricht University. <https://doi.org/10.26481/dis.20200909ng>

Document status and date:

Published: 01/01/2020

DOI:

[10.26481/dis.20200909ng](https://doi.org/10.26481/dis.20200909ng)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

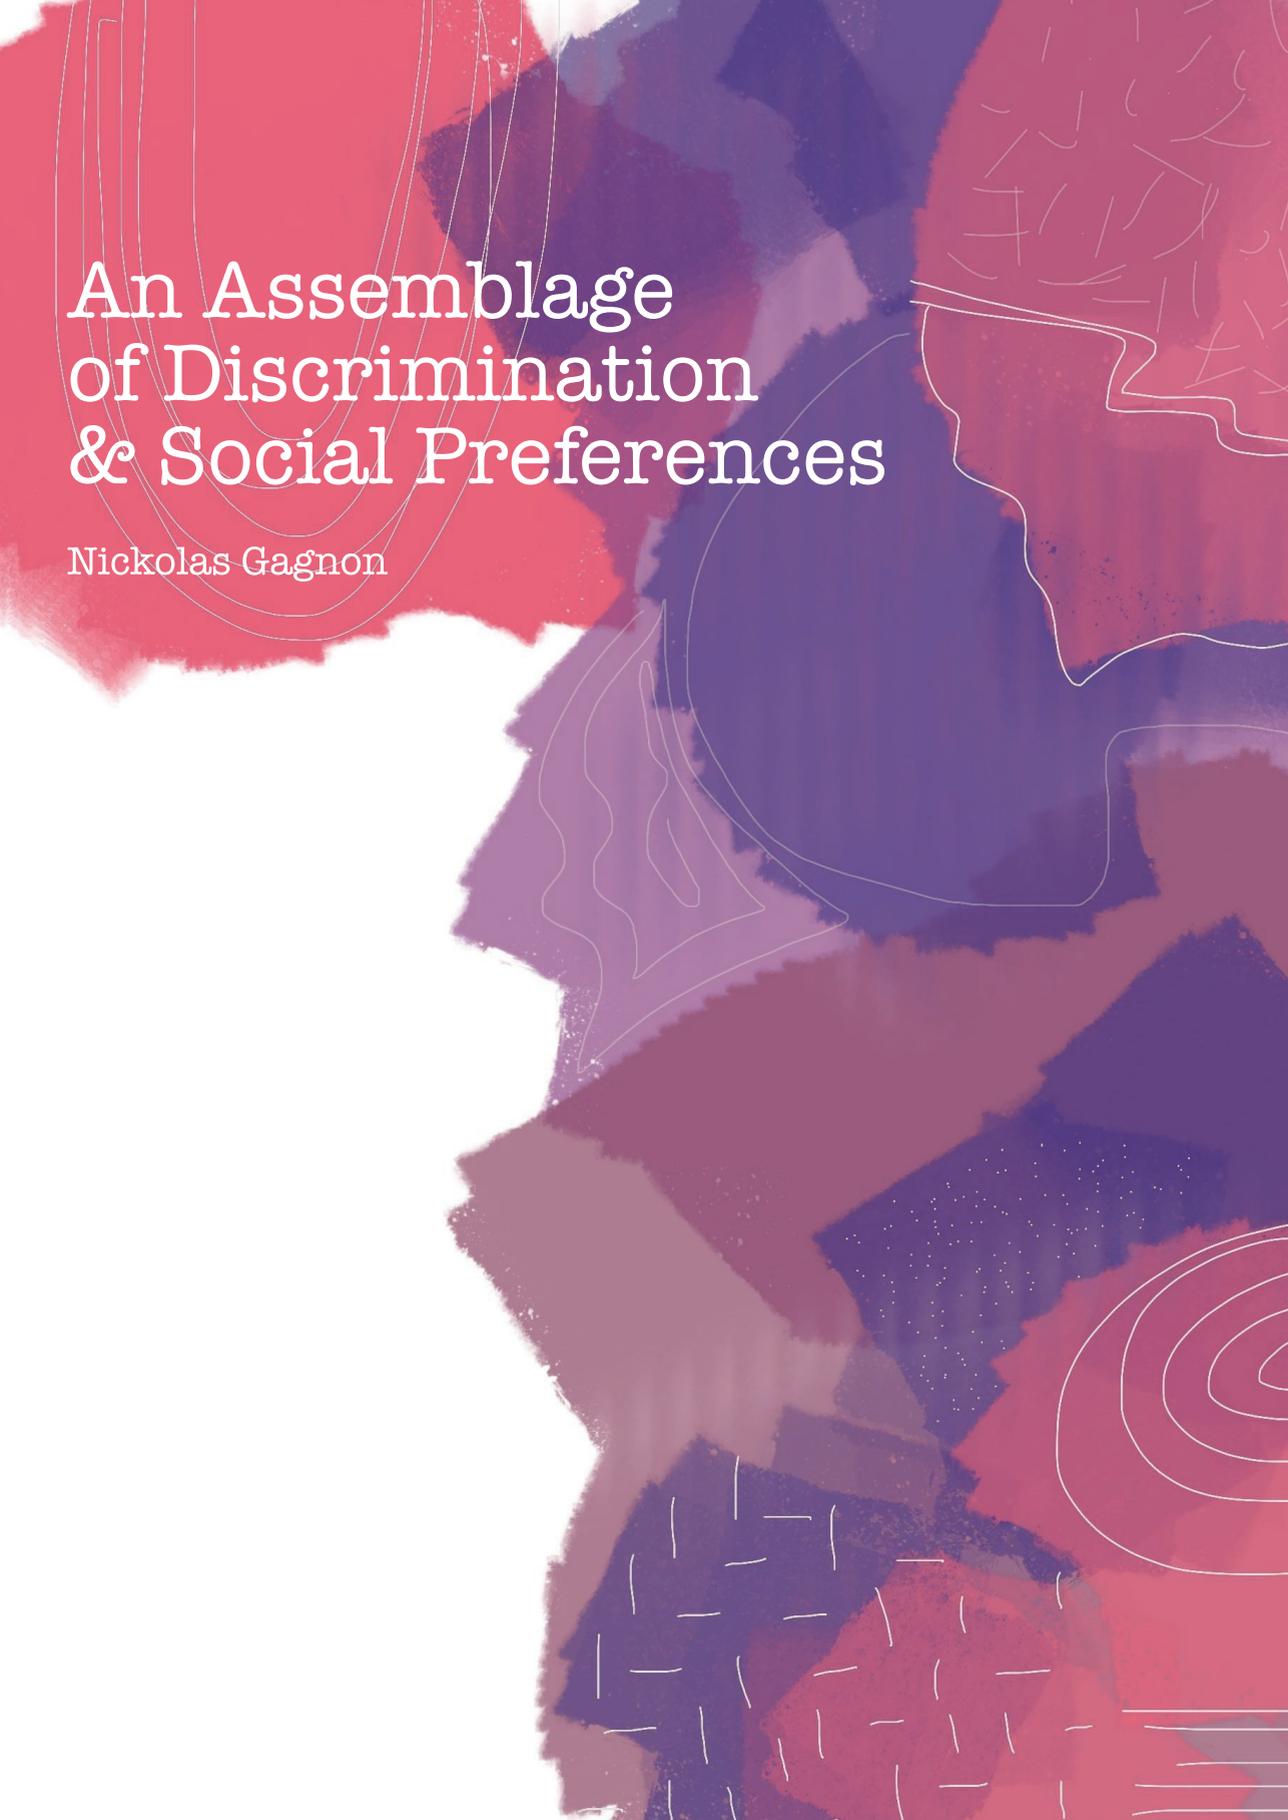
www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.



An Assemblage of Discrimination & Social Preferences

Nickolas Gagnon

© Nickolas Gagnon, Maastricht 2020

All rights reserved. No part of this publication may be reproduced, stored in an automated data system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior permission of the author.

The author acknowledges financial support from Maastricht University and the Fonds de recherche du Québec - Société et culture. Funding was also received from the University of Arizona and Tilburg University.

This book was typeset by the author using \LaTeX . The cover art is the work of Frédérique Surprenant.

An Assemblage of Discrimination and Social Preferences

Dissertation

to obtain the degree of Doctor at Maastricht University
on the authority of the Rector Magnificus Prof. dr. Rianne M. Letschert,
in accordance with the decision of the Board of Deans,
to be defended in public on
Wednesday, 9 September 2020, at 10:00 hours.

by

Nickolas Gagnon

Supervisor:

Prof. Dr. Arno Riedl

Co-Supervisor:

Dr. Kristof Bosmans

Assessment Committee:

Prof. Dr. Thomas Dohmen (Chair)

Prof. Dr. Dorothea Kübler, Technical University of Berlin

Prof. Dr. Sigrid Suetens, Tilburg University

Dr. Alexander Vostroknutov

Acknowledgments

The thanks, they are a-coming—many. I had the opportunity to learn from three great mentors for this thesis; they guided my entrance into academia and I am grateful to each of them, for their counsels, effort, constructive criticism, and unyielding trust and support. Kristof Bosmans welcomed me in Maastricht, and he advised me at every step, patiently discussed, read and corrected, more times than he had to. He taught me to discipline my writing—he partly succeeded.

Arno Riedl came right after in Maastricht. He taught me to pursue one worthwhile idea, ever carefully, rather than overextending in distant projects. He also encountered partial success with me. Charles N. Noussair, whom I knew from my masters' days, generously offered me another academic base, in Tucson, for large swaths of my PhD. He taught me how enthusiastic one could be about research, and this obviously was a success.

I shuttled between three homes during the past few years, sometimes many times a year, spanning the Sonoran desert, La Petite Nation's fields and forests, and the hills of Limburg. This continuous journey means that I owe a lot to an abundance of colleagues, friends as well as to my family members, who contributed to this thesis, in smaller and larger ways, by discussing ideas, providing feedback, revising designs and articles, and supporting me in my endeavors. For the sake of brevity, I restrict myself to mentioning only some of you.

I am most thankful to Marcello Negrini and Henrik Zaunbrecher for the countless hours they spent discussing and commenting everything in this thesis. I am thankful to many other current and former Maastricht University PhD candidates, staff, and faculty members, among whom are Andrzej Baranski, Matteo Bonetti, Fortuna Casoria, Alexandra de Gendre, Jona Linde, Lina Lozano, Elke Lucas, Letícia Micheli, Maria Polipciuc, Christina Rott, Riccardo Saulle, Caterina Schiavoni, Martin Strobel, Frauke Stehr, Eveline Vandewal, Matthias Wibrál, Peter Werner, and Anna Zseleva. I thank the many University of Arizona PhD candidates, staff, and faculty members with whom I interacted over the years, among whom are Richard Kiser, Kristen Lane, Elaine Rhee, and Bohan Ye. I also thank Jan Potters for his supervision during my time in Tilburg, and Mánuel Mágó for his help during that very same period.

Finally, three to four months a year were usually spent at my cousin's house in Ripon, Québec. Beside the beautiful scenery, it sometimes provided a most unusual sight: me writing pieces of this thesis with Victor or Alice on my lap. I am thankful to my entire family and especially to Émilie, Vincent, Sébastien, and Sylvain.

Contents

1	Introduction	1
2	The Effect of Unfair Chances and Gender Discrimination on Labor Supply	5
I.	Introduction	6
II.	Related Literature	10
II.A	<i>Unequal Wages</i>	10
II.B	<i>Unfair Chances</i>	11
II.C	<i>Gender Discrimination</i>	11
III.	Theoretical Framework	12
IV.	Study Design	14
IV.A	<i>Experiment</i>	14
IV.B	<i>Online Labor Markets</i>	17
IV.C	<i>Procedures and Demographics</i>	19
V.	Hypotheses	20
VI.	Results	23
VI.A	<i>Descriptive Statistics</i>	23
VI.B	<i>Test of Hypotheses</i>	24
VI.C	<i>Exploratory Analyses</i>	29
VII.	Discussion	32
VII.A	<i>Unequal Wages and Labor Supply</i>	32
VII.B	<i>Fairness of Chances</i>	33
VII.C	<i>Discrimination, Labor Supply, and the Gender Gap in Earnings</i>	34
VII.D	<i>Possible Extensions</i>	35
VIII.	Conclusion	36
2.A.	Additional Tables	39
2.B.	Additional Figures and Tables	39
2.C.	Instructions of the Experiment	50
2.C.A	Invitation Email	50
2.C.B	Instructions	51

3	On Your Own Side of the Fence	55
I.	Introduction	56
II.	Experimental Design	60
II.A	Wage and Ethnicity	61
II.B	Interaction	61
II.C	Task	62
II.D	Redistribution Decision	63
II.E	Choosing Inter-Ethnic and Inter-Income Interactions	63
II.F	Friendships	66
II.G	Additional Procedural Details	67
III.	Hypotheses	68
IV.	Results	68
IV.A	Overview Redistribution Decisions	68
IV.B	Testing the Hypotheses	69
IV.C	Exploratory Analyses—Interaction Decisions	71
V.	Discussion and Conclusion	75
3.A.	Additional Tables	79
3.B.	Instructions	83
3.C.	Task	86
4	Reciprocity under Brief and Long Time Delays	91
I.	Introduction	92
II.	Experimental Design	96
II.A	Procedures Common to all Experiments	96
II.B	Procedures Specific to Each Experiment	98
III.	Hypotheses	100
IV.	Results	102
IV.A	Summary of the Data	102
IV.B	Evaluation of Hypotheses	103
V.	Conclusion	110
4.A.	Supplemental Information Regarding Wage and Effort	113
4.B.	Analysis of the Relationship between Wages, Emotions, and Effort, using FaceReader Software	116
4.C.	Analysis of Employers' Predictions	120
4.D.	Instructions (for Tucson-H)	121
4.E.	Email Reminders for Long-term Workers	127
5	Unequal Chances: Ex Ante Fairness and Individual Control	131
I.	Introduction	132

II. Methods	135
III. Hypotheses	138
IV. Results	138
V. Discussion	143
5.A. Instructions Fairness	149
5.B. Instructions Individual Control	151
5.C. Additional Information Concerning the Reverse Block	151

List of Figures

2.1 Presentation of the Procedure Determining the Wages and Presentation of Wages	16
2.2 Mean Labor Supply per Payment Scheme	24
2.3 Mean Labor Supply per Gender and Payment Scheme	30
2.4 Task	39
4.1 Scatter Plots of Short- and Long-term Efforts	105

List of Tables

2.1 Wages and Chances of Two Workers in a Pair for each Payment Scheme	15
2.2 Demographic Characteristics of Workers	20
2.3 Predicted Labor Supply in Each Payment Scheme	21
2.4 Labor Supply per Payment Scheme	23
2.5 Tobit Regressions of Labor Supply on Payment Schemes	25
2.6 <i>P</i> -values of Predicted Differences in Labor Supply between Payment Schemes	25
2.7 Tobit Regression of Labor Supply on Payment Schemes with Both Types of Workers	27

2.8	<i>P</i> -values of Predicted Differences in Labor Supply Effect of Payment Schemes between Low-Wage Workers and High-Wage Workers	27
2.9	Labor Supply of Low-Wage Men and Women in UNEQUNFAIR and UNEQDISCR	28
2.10	Tobit Regression of Labor Supply in UNEQUNFAIR and UNEQDISCR, for Low-Wage Men and Women	29
2.11	Labor Supply per Gender and Payment Scheme	30
2.12	Tests of Inequalities in Labor Supply between Payment Schemes, for Men and Women	31
2.13	<i>P</i> -Values of Six Predicted Differences in Labor Supply between Payment Schemes	40
2.14	Tobit Regression of Labor Supply on Payment Schemes, for Men and Women	41
2.15	<i>P</i> -Values of Six Predicted Differences in Labor Supply between Payment Schemes, for Men and Women	42
2.16	<i>P</i> -Value of Test that Positive Discrimination of Women does not Affect their Labor Supply	42
2.17	Hurdle Model (Labor Supply on Payment Schemes), for Low-Wage Workers	43
2.18	<i>P</i> -Values of Six Predicted Differences in Labor Supply between Payment Schemes, Hurdle Model	43
2.19	Tobit Regression of Labor Supply on Payment Schemes, for Young (age ≤ 36) and Old (age > 36)	44
2.20	<i>P</i> -Values of Six Predicted Differences in Labor Supply between Payment Schemes, for Young (age ≤ 36) and Old (age > 36)	45
2.21	Tobit Regression of Labor Supply on Payment Schemes, for Employed Full Time and Not Employed Full Time	46
2.22	<i>P</i> -Values of Six Predicted Differences in Labor Supply between Payment Schemes, for Employed Full Time and Not Employed Full Time	47
2.23	Tobit Regression of Time Spent in the Experiment on Payment Schemes	48
2.24	<i>P</i> -Values of Six Predicted Differences in Time Spent in the Experiment between Payment Schemes	48
2.25	Labor Supply per Payment Scheme, for Workers not Beating the Odds and Workers Beating the Odds	49
2.26	<i>P</i> -Values of Tests on Difference in Labor Supply between Payment Schemes, for Workers Beating the Odds	49

3.1	Treatments—Types of Contact with the Poor by the Rich	62
3.2	Redistribution Rates and Resulting Distributions of Earnings	64
3.3	Summary of Redistribution Rates Chosen by the Rich	69
3.4	Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution	70
3.5	Percentage of Participants Choosing to be paired with a Cau- casian in Each Rank, for Caucasians and Minorities	72
3.6	OLS Regression Estimates—Difference in Number of Caucasians in First and Last Four Ranks, for Caucasians and Minorities	73
3.7	Choices of Contact with the Poor, and Tests of Equality of Choices, for the Poor and the Rich	75
3.8	Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution—Including All Participants with a Friend in the Session	79
3.9	Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution—Including Partici- pants with a Friend who Interact with their Friend in their Con- tact with the Poor	80
3.10	Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution—Including Partici- pants with a Friend who do not Interact with their Friend in their Contact with the Poor	80
3.11	Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution—Subgroup of the Caucasian Rich who Choose more often to Interact with Caucasians	81
3.12	Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution—Subgroup of the Rich who Choose to Avoid Contact with the Poor	82
4.1	Procedural Details	99
4.2	Summary of Participants' Average Decisions	103
4.3	Average Effort by Wage Range, for Short- and Long-Term Workers	104
4.4	Estimates of Tobit Regression (4.1)	106
4.5	Spearman Correlation between Wage and Effort	107
4.6	Spearman Correlation between Wage and Effort as well as be- tween the Unfairness Indicator I and Effort	108
4.7	Estimates of Tobit Regressions (4.2)	109
4.8	Average Effort by Wage Range, for Short- and Long-Term Work- ers	114

4.9	Estimates of Tobit Regression of Effort on Wage and Delay	115
4.10	Spearman Correlation between Wage and Effort for Both Types of Worker in Each Experiment	116
4.11	Estimates of Tobit Regressions of Effort on Emotions (4.1)	118
4.12	Effect of Wage on Emotional Changes in Workers	120
4.13	Tobit Regression of Employers' Effort Predictions on Short- and Long-term Wage Offered, and OLS Regression of Employers' Dif- ference in Effort Predictions on the Wage Offered	120
5.1	Criteria for receiving higher chances in the Regular and Reverse Blocks	138
5.2	Mean Fairness and Wilcoxon Signed-Rank Test Comparisons . . .	140
5.3	Asymmetries and Wilcoxon Signed-Rank Test Comparisons . . .	142
5.4	Mean Control and Wilcoxon Signed-Rank Test Comparisons . . .	142
5.5	Mean Fairness and Wilcoxon Signed-Rank Test Comparisons— Reverse Block	152
5.6	Mean Control and Wilcoxon Signed-Rank Test Comparisons— Reverse Block	152

Chapter 1

Introduction

The topics that bind this thesis, apart from the methodology and conceptual framework—experimental and behavioral economics—are discrimination and social preferences. They also represent my main research interests for the foreseeable future.

The Nature of Prejudice is social psychologist Gordon Allport's 1954 masterpiece. It complements the large body of work that economists have accumulated on discrimination (e.g., from taste-based (Becker, 1957) and statistical (Phelps, 1972; Arrow, 1973) reasons, to more recent work, based on language (Lang, 1986), implicit discrimination (Bertrand et al., 2005), attention (Bartoš et al., 2016) or stereotypes (Bordalo et al., 2016)). Aside from its dissection of prejudice, the book considers the behavior of individuals facing discrimination, a fascinating topic that the reader will recognize in chapter 2. It sheds light upon several ageless issues rarely scrutinized by economists. Furthermore, two influential articles form the starting point of chapter 3: Luttmer (2001) and Dahlberg et al. (2012). They both provide empirical evidence from economics, correlational for the first and quasi-causal for the second, that living close to the poor of another ethnicity decreases support for redistributive policies in a society.

Social preferences stand at the forefront of each chapter that I present here. The most successful models of those preferences are arguably inequality or inequity aversion (e.g., Fehr and Schmidt 1999; Bolton and Ockenfels 2000) and reciprocity (e.g., Rabin, 1993; Dufwenberg and Kirchsteiger, 2004). Chapter 2 does explore workplace consequences of inequality aversion, but focuses more on its less-studied cousin, the preference for a fair chance (e.g., Diamond, 1967), and also examines the specific dislike for gender discrimination. Chapter 3 delves into the interplay of inequality aversion and ethnic discrimination for redistribution decisions. Chapter 4 is about whether the passing of time

perturbs reciprocal behavior, once more in the context of labor markets.

Experiments are my methodology of choice; two chapters make use of the laboratory and one recruits workers on an online labor platform. I find that the most attractive part of experiments is the opportunity (and challenge) of designing the “perfect” environment for what one aims to study. The tight control of the experimental environment, including random assignment to treatments, offers the ability to draw clean causal conclusions, often an arduous or impossible task with observational data. A natural question that arises for an experimentalist is the one of external validity, i.e., how related are behavior inside and outside of the experiment, especially when using the laboratory? On this account, a number of articles (e.g., Karlan 2005; Benz and Meier 2008; Baran et al. 2010; Dohmen et al. 2011; Franzen and Pointner 2013; Cohn et al. 2015; Buser et al. 2014; Herbst and Mas 2015; Potters and Stoop 2016; Riedl and Smeets 2017; Fisman et al. 2017; Falk et al. 2018) provide evidence that simple experiments do offer very useful information for understanding behavior across diverse domains, from donations, lying behavior, loan repayment, work behavior and career choices, to political votes and socially-responsible investments. Moreover, Snowberg and Yariv (2018) show that, for a range of common experimental measures, behavior is qualitatively similar in the laboratory with undergraduate students, on online labor platforms, and using a sample representative of the population.

The Effect of Unfair Chances and Gender Discrimination on Labor Supply is the straightforward title of chapter 2, co-authored with Kristof Bosmans and Arno Riedl, and the most recent work collected for this thesis. Using an experiment on an online labor platform, we present the first investigation of the causal effects of unfair chances in receiving a high or low wage and of gender discrimination (through gender-discriminatory unfair chances) on labor supply decisions of workers. Interestingly, its results provide a new mechanism for the lower labor supply of women, which is a prominent explanation for the gender earnings gap.

On Your Own Side of the Fence is the single-authored chapter 3. Its connotative title comes from Flannery O'Connor’s 1965 short story *Everything That Rises Must Converge*, staged in the American South during the Civil Rights period. Using a laboratory experiment, this study explores the causal mechanism behind the existing literature’s finding that living close to poor individuals of another ethnicity decreases support for income redistribution. It also examines whether those inter-ethnic and inter-income interactions actually occur when

individuals are given the choice. This additional feature allows me to provide a distinct interpretation of the effect of inter-ethnic proximity with the poor on redistribution. The chapter is a much-reshaped version of a design I created several years ago, and the sessions were conducted during my Arizona visits.

Reciprocity under Brief and Long Delays is my first research project, which resulted in chapter 4, co-authored with Charles N. Noussair and appearing in the July 2020 issue of *Economic Inquiry*. With a laboratory experiment, we study whether time erodes the positive wage-effort relationship commonly found in the short term in incomplete contract settings. This work starts with a question I had in Charles' course during my research master's degree, concerning one of experimental economics' most influential articles: Fehr, Kirchsteiger, and Riedl (1993).

A supplementary contribution to this thesis appears in chapter 5. It is a short work co-authored with Leticia Micheli and served as a stepping stone for chapter 2. Using a survey, we investigate empirically the role of different sources of unequal chances in shaping fairness perceptions of those chances.

Without further ado, the research.

Chapter 2

The Effect of Unfair Chances and Gender Discrimination on Labor Supply

Abstract

Labor market opportunities and wages may be unfair for various reasons, and how workers respond to different types of unfairness can have major economic consequences. Using an online labor platform, where workers engage in an individual task for a piece-rate wage, we investigate the causal effect of neutral and gender-discriminatory unfair chances on labor supply. We randomize workers into treatments where we control relative pay and chances to receive a low or a high wage. Chances can be fair, unfair based on an unspecified source, or unfair based on gender discrimination. Unequal pay reduces labor supply of low-wage workers, irrespective of whether the low wage is the result of fair or unfair chances. Importantly, the source of unfair chances matters. When a low wage is the result of gender-discriminatory chances, workers matched with a high-wage worker substantially reduce their labor supply compared to the case of equal low wages (-22%). This decrease is twice as large as those induced by low wages due to fair chances or unfair chances coming from an unspecified source. In addition, exploratory analysis suggests that in response to unequal pay, low-wage male workers reduce labor supply irrespective of the source of inequality, whereas low-wage female workers reduce labor supply only if unequal pay is due to gender-discriminatory chances. Our results concerning gender discrimination indicate a new reason for the lower labor supply of women, which is a prominent explanation for the gender gap in earnings.

This chapter is co-authored with Kristof Bosmans and Arno Riedl.

I. Introduction

Chances are a pervasive feature of labor market activities and outcomes. Workers are confronted with them when they face hiring, bonus payment, promotion, and dismissal decisions. Those chances might be fair, but can also be unfair for many reasons, ranging from favoritism and nepotism to discrimination. For instance, recent work suggests that women evaluated for tenure in economics departments face lower chances than men of equal ability (Sarsons et al., 2019). In this paper, using an experiment, we present the first investigation of the causal impact of unfair chances in receiving a high or low wage on labor supply decisions of workers. Moreover, we provide the first examination of the causal effect of gender discrimination (through gender-discriminatory unfair chances) on labor supply decisions.

Recent work has shown that unequal pay can have major effects on workers' behavior, such as reduced labor supply and productivity as well as increased job separations (Bracha et al., 2015; Breza et al., 2018; Dube et al., 2019). At the same time, a lasting idea in economics is that unfair chances influence equity judgments (e.g., Diamond, 1967), and empirical studies have provided support for this claim (e.g., Bolton et al., 2005; Brock et al., 2013; Cappelen et al., 2013).¹ Combining the results of those two strands of literature suggests that unequal chances on labor markets may also affect workers' behavior significantly, although this has yet to be studied. Moreover, that inequality comes from a specific source such as gender discrimination may generate especially strong responses. In spite of this, whether gender-discriminatory inequality does provoke reactions that are distinct from more neutral inequality has thus far not been examined.

To investigate the causal effect of unfair chances and gender discrimination on labor supply, we conduct a controlled experiment. We use an online labor platform, where we hire workers who individually engage in a task at a fixed piece-rate wage. We can draw causal conclusions by fully controlling both the chances leading to pay inequality between workers and the source of these chances. Workers first learn the procedure that will determine their wage as well as the wage of another worker, which can be high or low. Then, workers are informed of their own and the other worker's resulting wage, after which they decide individually how much to work.

¹Psychologists have suggested that procedures regarded as unfair can engender undesirable work behavior, such as decreased productivity and retaliation (e.g., Leventhal, Karuza, and Fry 1980, Skarlicki and Folger 1997).

To cleanly isolate the effect of chances and their source on labor supply, we implement a design that rules out peer interactions, risk, and reciprocity considerations. Providing less labor is costly to workers because it reduces their own earnings, but has no other effects. An important feature of our design is that we implement an explicit gender-discriminatory policy, which has been approved by the local ethical review committee.² Moreover, the use of an online labor platform provides us with a relatively large number of participants who are arguably more representative than student participants. The anonymity of the online platform also helps in avoiding peer effects from observing other workers quitting the task.³

In our experiment we measure labor supply in the real effort task and implement five different payment schemes (treatments) where we create wage inequality through fair or unfair chances and also vary the source of unfair chances. We randomize each worker into one of the five schemes. In each scheme, a worker is anonymously matched with another worker engaged in the same individual task. In two baseline schemes, called EQLOW and EQHIGH, both workers receive the low wage or receive the high wage, respectively. In these schemes, we do not mention the procedure that determines the wages. In the other three schemes workers in a pair receive unequal wages: one receives the high wage and one receives the low wage. In the payment scheme UNEQ-FAIR, both workers have a fair chance (50%) of receiving the high wage. By contrast, in the schemes called UNEQUNFAIR and UNEQDISCR one worker has a 25% chance of receiving the high wage and the other has a 75% chance. In the former scheme no specific source for the unfair chances is mentioned. In the latter scheme, pairs of workers consist of a man and a woman, and they are informed that their chances depend on their gender. One worker has a 25% chance of receiving the high wage explicitly because she is woman or he is a man, whereas the other worker has a 75% chance because he is man or she is a woman.

To derive our hypotheses regarding labor supply in the different schemes, we provide a theoretical framework which extends the framework on inequality aversion and work morale by Card, Mas, Moretti, and Saez (2012) and Breza, Kaur, and Shamdasani (2018) with the model of social preferences incorporating chances by Saito (2013). The model assumes that workers engaged in the same work dislike wage inequality, and more so if this wage inequality is the

²Ethical Review Committee Inner City of Maastricht University (ID: ERCIC-054-15-11-2017).

³In Section IV.B we discuss potential drawbacks of online experiments and how we minimize them.

result of unfair chances. Both unfair chances and wage inequality increase the marginal disutility of working, thus reducing labor supply. In addition, we posit that gender discrimination as the source of unfair chances creates an extra psychological cost that further increases workers' marginal disutility from unequal treatment.

We compare the labor supply of workers at a given wage across payment schemes, and test four pre-registered hypotheses. Our first and second hypothesis states that for both low-wage workers and high-wage workers, unequal wages, unfair chances, and gender discrimination each incrementally decrease labor supply. Our third hypothesis says that these effects are stronger for low-wage workers than for high-wage workers. Finally, our fourth hypothesis conjectures that men and women may not react equally to gender-discriminatory chances compared to the same unfair chances coming from an unspecified source.

Our main results can be summarized as follows. In line with the first hypothesis, low-wage workers in worker pairs with unequal pay who receive their low wage through fair chances work significantly less (-13%) than low-wage workers who are matched with another low-wage worker. This effect of wage inequality is consistent with results reported in the three previous studies on the topic (Bracha et al., 2015; Breza et al., 2018; Dube et al., 2019). Interestingly, the fact that in our case fair chances determine the unequal pay does not remove this negative effect. Further, in contrast to the first hypothesis, low wages resulting from unfair chances based on an unspecified source produce a similar labor supply as low wages based on fair chances. However, low-wage workers do significantly decrease their labor supply if they face unfair gender-discriminatory chances. The average impact of gender discrimination on labor supply is significant and large in economic terms. Low wages resulting from unfair gender-discriminatory chances reduce labor supply by 15% relative to low wages coming from unfair chances based on an unspecified source, and by 22% relative to equal low wages.

In contrast, high-wage workers appear to be immune to the different types of inequality that we employ. In no case is their labor supply significantly different in one payment scheme than in another. Thus, the second hypothesis is not supported.

The third hypothesis, that the adverse effects on labor supply are stronger for low-wage workers than for high-wage workers, holds when comparing the payment scheme with gender-discriminatory chances to the other schemes. In

all other comparisons, low-wage workers and high-wage workers do not significantly differ in their responses.

Finally, we obtain evidence that low-wage women respond more strongly to gender-discriminatory unfair chances than low-wage men do, supporting our fourth hypothesis. Moreover, an exploratory comparison across all five payment schemes reveals a distinct gender difference among low-wage workers. Men decrease their labor supply in response to any type of disadvantageous inequality, whereas disadvantaged women reduce their labor supply only if the low-wage is due to gender-discriminatory chances.

Our study provides three main contributions to the literature. First, it is the first to investigate the causal effect of unfair chances on labor supply decisions of workers. Our finding that the fairness of chances from an unspecified source does not affect labor supply stands in contrast to the empirical literature on unfair chances and income redistribution (e.g., Bolton et al., 2005; Krawczyk and Le Lec, 2010; Brock et al., 2013; Cappelen et al., 2013). Specifically, whereas the distribution of initial chances has been shown to influence redistribution decisions, we find that it has no impact on workers' labor supply decisions. This suggests that the response to unfair chances depends on the context.

Second, our study also provides the first evidence of the causal impact of gender discrimination on labor supply decisions. Our finding that unfair chances based on gender have a large negative impact on labor supply is consistent with literature in medicine and psychology, suggesting that discrimination imposes a specific psychological cost (e.g., Pascoe and Smart Richman, 2009). This adds to the few economic studies showing that (ethnic minority) workers modify their behavior when discriminated (Parsons et al., 2011; Glover et al., 2017), making them appear less productive for employers. Our study is the only one that identifies the effect of discrimination while controlling for changes in monetary incentives resulting from the presence of discrimination.

A third contribution of our study is to the research on gender differences in labor markets more generally (Croson and Gneezy, 2009; Niederle, 2016). Our result that men reduce their labor supply in reaction to unequal wages from fair and unfair chances coming from an unspecified source, whereas women do not, is consistent with the findings of Bracha et al. (2015). They, however, do not study the effect of unfair chances. In addition, we show that also women decrease their labor supply if they face lower wages that are the result of gender discrimination.

Finally, an important implication of our study is that it suggests a novel

and complementary explanation for the gender gap in earnings. Goldin (2014) and Blau and Kahn (2017) review the literature on the earnings gender gap and conclude that the modern gender gap is mostly explained by the lower labor supply of women. The standard explanation for this lower labor supply is that women value temporal flexibility more than men, plausibly because they have to bear greater household responsibilities (e.g., Goldin, 2014; Wiswall and Zafar, 2017). Importantly, we provide evidence that the experience of gender discrimination itself can reduce labor supply. Not accounting for this channel might lead to misjudgment of the impact of discrimination and ill-advised policies. The identified channel may also affect the gender earnings gap in the manner of a self-fulfilling prophecy. Since the labor market offers high returns to long work hours, a lower labor supply can be both a reaction to discrimination and a rationale for employers to pay women less than men.

The rest of the paper is organized as follows. Section II situates our research in the literature, Section III presents the theoretical framework, Section IV describes the study design, Section V advances the hypotheses, Section VI details the results, Section VII discusses our findings and their implications, and Section VIII briefly concludes.

II. Related Literature

II.A. Unequal Wages

A stream of literature suggests that wage differentials perceived as unfair hamper the work morale of workers (Adams, 1965; Akerlof and Yellen, 1990; Pfeffer and Langton, 1993; Bewley, 1999). Empirical evidence shows that unequal wages for similar work indeed negatively affect several labor outcomes. For instance, wage inequality decreases work satisfaction and increases job searches among disadvantaged workers (Card et al., 2012) and hurts their productivity (Gächter and Thöni, 2010; Cohn et al., 2014; Ockenfels et al., 2014).⁴

We discuss here in more detail those studies that investigate the effect of unequal wages on labor supply, which is our variable of interest. Bracha et al. (2015) report a laboratory experiment where workers are paid piece-rate wages in an individual task and have to decide for how long to work. If no justification is provided, unequal wages decrease labor supply of male but not of female low-pay workers. High-pay workers' labor supply does not respond to

⁴However, individuals might accept or demand inequality on the grounds of equity (Konow, 2000). For example, wage differentials could be viewed as equitable if they reflect observable productivity differentials (Abeler et al., 2010; Breza et al., 2018). This could even be the case if differences in productivity are possible, but unobserved (Charness and Kuhn, 2007).

wage inequality, irrespective of gender. Breza et al. (2018) conduct a field experiment with male workers in an Indian firm. Workers work individually in small teams in which wage inequality is manipulated. If workers do not observe other workers' productivity, then inequality decreases labor supply of low-pay and high-pay workers as well as the productivity of low-pay workers. However, wage inequality has no effect if it reflects observable productivity differences. Dube et al. (2019) exploit a natural experiment caused by changes in the wage structure of a large American firm. They find that workers arbitrarily receiving a low relative pay after the change were more likely to quit, whereas workers with a higher relative pay did not change their behavior. Those important studies are informative about the effect of wage inequality on labor supply decisions. However, they do not examine the role of chances or discrimination as a reason for wage differences and how it may affect labor supply.

II.B. Unfair Chances

Economists have long considered the welfare implications of assessing inequality in terms of ex ante chances and ex post outcomes (Harsanyi, 1955; Diamond, 1967; Hammond, 1981; Epstein and Segal, 1992; Fleurbaey, 2010). Models of social preferences have also recently incorporated a dislike for unequal chances, usually referred to as a concern for ex ante or procedural fairness (Karni and Safra, 2002; Bolton et al., 2005; Trautmann, 2009; Krawczyk, 2011; Saito, 2013). A number of laboratory experiments have lent empirical support to the notion that individuals take into consideration the fairness of chances when making distributive decisions (Bolton et al., 2005; Karni et al., 2008; Krawczyk and Le Lec, 2010; Krawczyk, 2010; Brock et al., 2013; Cappelen et al., 2013; Andreoni et al., 2016; Grimalda et al., 2016; Trautmann and van de Kuilen, 2016; Cettolin and Riedl, 2016; Miao and Zhong, 2018). However, how ex ante chances influence the labor decisions of workers has not been investigated.⁵

II.C. Gender Discrimination

Women face a gender gap in earnings, have lower promotion chances, are less present in high-paid jobs, work less hours, work more part time, and have a

⁵Organizational psychologists have studied a related concept referred to as procedural justice. According to one prominent form of procedural justice, procedures are fair to the extent that decisions are "consistent" and without "bias" (see, e.g., Leventhal et al., 1980; Skarlicki and Folger, 1997).

lower labor participation (for overviews, see, Altonji and Blank, 1999; Riach and Rich, 2002; Goldin, 2014; Blau and Kahn, 2017). A number of studies suggest that demand-side gender discrimination plays an important role in explaining women’s disadvantaged labor market position (see, e.g., Bertrand and Duflo, 2017; Neumark et al., 1996; Goldin and Rouse, 2000; Sarsons et al., 2019).⁶

Interestingly, potential supply-side effects of discrimination are much less studied. Parsons et al. (2011) present evidence from American baseball showing that minority players change their behavior in response to discrimination by officials. Glover et al. (2017) show that ethnically-biased managers in a large French grocery chain decrease minority workers’ productivity and labor supply. A field experiment of Ibañez and Riener (2018) examines some aspects of gender discrimination (Affirmative Action for women) on job applications. In those studies, discrimination changes the monetary incentives for workers who are discriminated, so that the response to discrimination is entangled with the change in incentives.

Studies in medicine and psychology show that discrimination is correlated with serious negative consequences for physical and mental well-being. In a meta-analytic review, Pascoe and Smart Richman (2009) link discrimination to a range of psychological issues, such as anger, stress, anxiety, distress, and low general wellbeing, all of which we can reasonably expect to considerably lower one’s work satisfaction. However, the reviewed studies do not investigate the impact on workers’ labor decisions.

III. Theoretical Framework

We adapt the framework of Card et al. (2012) and Breza et al. (2018) to model how workers may react to unequal wages, unfair chances, and gender discrimination. In the original model, wage inequality between workers engaged in the same work decreases work satisfaction or morale, which translates into lower marginal utility from work and thus into lower labor supply. Our model is an extension that also accounts for chances in the process leading to unequal wages. That is, unfair chances are also assumed to decrease marginal utility from work. To explore the role of gender discrimination, we assume that individuals are more averse to unfair chances caused by gender discrimination

⁶In this literature, taste-based discrimination (Becker, 1957) and statistical discrimination (Phelps, 1972; Arrow, 1973) are the two most discussed forms of discrimination. Other forms include language discrimination (Lang, 1986), implicit discrimination (Bertrand et al., 2005), attention discrimination (Bartoš et al., 2016), and stereotypes (Bordalo et al., 2016).

than to unfair chances coming from an unspecified source.⁷ We use this model to derive most of our hypotheses, which are formulated in *Section V*.

Consider two workers, i and j , engaged in the same work receiving piece-rate wages w_i and w_j , which are known to both workers. There is no interaction between the two workers. A worker, say i , chooses labor supply l_i by taking into account his or her own wage, the wage of the other worker j , the chances that lead to their respective wages, and the cost of providing labor.

The modeling of marginal disutility created by ex post wage inequality is inspired by Fehr and Schmidt (1999) and is also used in Breza et al. (2018). It is denoted P_i and given by

$$P_i(w_i, w_j) = \alpha_i \max\{w_j - w_i, 0\} + \beta_i \max\{w_i - w_j, 0\}, \quad (2.1)$$

where the first term on the right-hand side measures the marginal disutility from disadvantageous wage inequality and the second term the marginal disutility from advantageous wage inequality, with $\alpha_i > \beta_i > 0$.⁸ That is, wage inequality produces a marginal disutility, and this marginal disutility is greater for disadvantageous inequality than for advantageous inequality.

The marginal disutility created by unfair chances, denoted by A_i for ex ante inequality, is inspired by Saito (2013)⁹ and takes the form

$$A_i(Ew_i, Ew_j) = \alpha'_i \max\{Ew_j - Ew_i, 0\} + \beta'_i \max\{Ew_i - Ew_j, 0\}, \quad (2.2)$$

where Ew_i and Ew_j denote expected wages. Similar to equation (2.1), here the first term on the right-hand side reflects the marginal disutility from disadvantageous expected wage inequality, and the second term that from advantageous expected wage inequality. As above we assume that $\alpha'_i > \beta'_i > 0$ but we allow for $\alpha_i \neq \alpha'_i$ and $\beta_i \neq \beta'_i$. That is, the disutility weights placed on wage inequality and unfair chances may differ.

We embed the aversion to unequal wages and the aversion to unfair chances described in equations (2.1) and (2.2) in the labor supply decision in the following way. A worker i chooses labor supply l_i in order to maximize the utility

⁷Unlike Breza et al. (2018) and in line with our experimental implementation, we rule out moral hazard and assume that work effort is fully contractible.

⁸This assumption is based on the empirical evidence reported in Bracha et al. (2015), Breza et al. (2018) and Dube et al. (2019) which suggests that, on average, $\alpha_i > \beta_i = 0$ or $\alpha_i > \beta_i > 0$. The literature on social preferences often makes the weaker assumption $\alpha_i \geq \beta_i \geq 0$ (e.g., Fehr and Schmidt, 1999). In our model, if $\alpha_i = \beta_i > 0$, then the effect of unequal wages on marginal disutility is the same for advantaged and disadvantaged workers. If $\alpha_i = \beta_i = 0$, then the problem collapses to standard selfish preferences and unequal wages do not affect the morale of workers.

⁹For earlier theoretical work combining social preferences and the effect of (un)fair chances, see Bolton et al. (2005) and Trautmann (2009).

function

$$U_i(w_i, w_j, l_i) = w_i l_i - P_i(w_i, w_j) l_i - A_i(Ew_i, Ew_j) l_i - \frac{l_i^2}{2}. \quad (2.3)$$

In equation (2.3), the first term on the right-hand side corresponds to the utility of monetary earnings derived from working, the second term is the disutility created by wage inequality, and the third term reflects the disutility created by unfair chances. The final term is the utility cost of providing labor.¹⁰

Assuming an interior solution, the optimal labor supply is given by

$$l_i^* = w_i - P_i(w_i, w_j) - A_i(Ew_i, Ew_j). \quad (2.4)$$

Given our assumptions about $\alpha_i, \beta_i, \alpha'_i$ and β'_i , the term $P_i(w_i, w_j)$ is strictly positive when wages are unequal and the term $A_i(Ew_i, Ew_j)$ is strictly positive when chances are unfair. Thus, unequal wages and unfair chances both reduce the optimal labor supply.¹¹ Disadvantageous inequality reduces the optimal labor supply more than advantageous inequality does because of $\alpha_i > \beta_i$ and $\alpha'_i > \beta'_i$.

Regarding gender discrimination, we posit that it translates into further marginal disutility from unequal wages or unfair chances. As before, we assume that the marginal disutility caused by discrimination is greater for disadvantaged than for advantaged workers. That is, gender discrimination in chances would increase α'_i and β'_i , and would increase α'_i more than β'_i . Therefore, unfair chances based on gender discrimination reduce the optimal labor supply more than unfair chances based on an unspecified source and the labor supply reduction is greater for disadvantaged than for advantaged workers.

Before translating these theoretical considerations into testable hypotheses in Section V, we first present our study design and experiment.

IV. Study Design

IV.A. Experiment

We hired workers on an online labor platform for performing a real effort task, which consisted of entering lines of random characters.¹² Each worker was

¹⁰For simplicity, we assume a quadratic cost function. The hypotheses derived from the model stay qualitatively the same when assuming any other strictly increasing and strictly convex cost function.

¹¹Equation (2.4) assumes an interior solution. If the optimal labor supply is a corner solution—zero or maximum labor supply—altering the inequality of wages or the unfairness of chances may not affect the optimal labor supply.

¹²Here we describe the design of the experiment. The online labor market and the pros and cons of using it for research purposes are described and discussed in more detail in Section IV.B below.

assigned the same task and carried it out individually, entering one line at a time. The payment was on a piece-rate basis, that is, a worker received a fixed payment per correctly entered line. If a mistake was made when entering a line, the worker was informed and had to correct it before proceeding to the next line. The length of the lines increased with the number of lines completed, which made the task increasingly harder over time.¹³ Each worker decided individually how many lines to enter. A worker could stop working at any time by leaving the experiment. Workers were informed about this and that they could not reenter the experiment once they had left. They were also instructed that they could work for at most 65 minutes. The number of lines entered is our measure of labor supply.

Each worker was randomly assigned to one payment scheme and anonymously paired with another worker in the same scheme who was engaged in the exact same task on the platform. Each worker in a pair was first informed about the procedure that would lead to his or her own wage and the wage of the other worker. Thereafter, each of the two workers was informed about his or her own resulting wage and the wage of the other worker. Thus, workers only started working after any uncertainty about their own and the other worker’s wage was resolved. Figure 2.1(a) provides a screenshot example of what workers saw when they were informed about the procedure leading to the wages, and Figure 2.1(b) shows a screenshot example of what they saw when informed about the resulting wages. To emphasize the piece-rate nature of the wage, it was described as a “payment per line.”

Table 2.1: Wages and Chances of Two Workers in a Pair for each Payment Scheme

Payment Scheme	Wage of Worker, Wage of Other Worker	Chance of Worker, Chance of Other Worker	Source of Chances
EQLow	£0.03, £0.03	-	-
EQHIGH	£0.06, £0.06	-	-
UNEQFAIR	£0.03, £0.06	50%, 50%	Unspecified
UNEQUNFAIR	£0.03, £0.06	25%, 75%	Unspecified
UNEQDISCR	£0.03, £0.06	25%, 75%	Gender Discrimination

Note: Chances are to receive the high wage (£0.06).

Table 2.1 summarizes the five payment schemes (treatments) we employ. The schemes determine, within a worker pair, the wages and the procedure creating these wages. In the schemes EQLow and EQHIGH, no chances are

¹³The number of characters contained in a line ranged from 10 at the start to 26 at the end. We implemented this to mimic an increasing and convex cost of labor supply. There was a maximum of 85 lines and workers were not informed about this beforehand (see Figure 2.4 in Section 2.B of the Appendix for two screenshots of the task). Further details on the task can be found in the instructions of the experiment, which are provided in Section 2.C of the Appendix.

Payment per Line

In this part you can earn a payment with each line you correctly enter. You will now be informed about the amount of this payment per line.

We also inform you about the payment per line of another participant in this economic experiment. This other participant is a real person, participates in the exact same experiment as you do, and has been recruited from the same on-line work platform as you. The other participant also finished the Practice Part and faces the exact same task as you. This means that the other participant faces the exact same lines as you in the task.

Your payment and the payment of the other participant are **decided through a computerized lottery**. One of you will receive a payment per line of £0.06, the other will receive a payment per line of £0.03.

You are a woman, therefore you have:
 a **25% chance** of receiving the payment per line of £0.06 and
 a **75% chance** of receiving the payment per line of £0.03.

The other participant is a man, therefore he has:
 a **75% chance** of receiving the payment per line of £0.06, and
 a **25% chance** of receiving the payment per line of £0.03.

The software will inform you about the outcome of the lottery on the next page.

Once you have learned your payment per line, you can start the task.

(a) Presentation of the Procedure Determining the Wages (Payment Scheme UNEQDISCR)

The lottery outcome is the following:

Your payment per line is **£0.03**.
 The other participant's payment per line is **£0.06**.

(b) Presentation of the Wages (Payment Scheme UNEQDISCR)

Figure 2.1: Presentation of the Procedure Determining the Wages and Presentation of Wages

involved and both workers receive either the low piece-rate wage of £0.03 or the high piece-rate wage of £0.06. These treatments serve as controls for the labor supply effect of receiving a low or a high wage when inequality in wages and procedures is absent. In the three other schemes, the two workers in a pair face chances to obtain the high or the low wage in a lottery. The lottery draw is dependent such that one worker receives the high wage and the other worker receives the low wage. In UNEQFAIR both workers have a fair chance of 50% to receive the high wage, whereas in UNEQUNFAIR one worker faces a low chance of 25% and the other worker faces a high chance of 75% to receive

the high wage. In both of these schemes we do not inform workers about the reason why the chances are allocated in this manner. In contrast, in the scheme UNEQDISCR the unfair chances explicitly discriminate one gender over the other. That is, one worker in the pair is informed that she (he) faces a 25% chance of receiving the high wage because she (he) is a woman (man), and that the other worker faces a 75% chance of receiving the high wage because he (she) is a man (woman). In this treatment, in half of the cases men face higher chances and in the other half women face higher chances.¹⁴

At the start of the experiment, a worker electronically signed an informed consent form and then read the instructions. Each worker had to correctly answer nine comprehension questions and go through a practice phase to become familiar with the task. Only thereafter was the worker assigned to a payment scheme and informed about the procedure leading to the wages in his or her worker pair. After having learned his or her own wage and the wage of the other worker in the pair, a worker could start working on the task.

IV.B. Online Labor Markets

We recruit workers on the UK-based online platform Prolific (www.prolific.ac). The use of such online labor markets for experiments has gained in popularity among economists in recent years.¹⁵ For our research the use of an online platform provides important advantages over a laboratory experiment or a field experiment inside a firm, but it also has some potential shortcomings. In what follows we discuss these advantages and how we deal with potential disadvantages.

The first advantage is that the online platform greatly reduces the possibility of peer effects, because workers on the platform do not interact with each other in any way during the task and can quit working without other workers noticing it. On the platform there is no channel through which workers could commu-

¹⁴Workers were not informed about this balance.

¹⁵Examples include Pallais (2014) on inexperienced workers, Kuziemko, Norton, Saez, and Stantcheva (2015) on redistribution preferences, Gilchrist, Luca, and Malhotra (2016) on employer-employees relationships, Pallais and Sands (2016) and Horton (2017a) on labor market referrals and recommendations, Bordalo et al. (2016) on stereotypes, Horton (2017b) on minimum wages and employment, Lyons (2017) on diversity and production in teams, Coffman, Coffman, and Ericson (2016) on anti-gay sentiments, Coffman et al. (2017) and Sarsons et al. (2019) on gender discrimination, and De Quidt, Haushofer, and Roth (2018) on experimenter demand effects. Horton, Rand, and Zeckhauser (2011), Arechar, Gächter, and Molleman (2018), and Snowberg and Yariv (2018) report that common economic games and elicited behavior in online and laboratory experiments provide qualitatively similar results. Bohren et al. (2018) also use an online scientific platform to study gender discrimination, where individuals are not paid but volunteer. Furthermore, several studies use existing labor data from online platforms, e.g., Ghani, Kerr, and Stanton (2014), Stanton and Thomas (2015), and Dube, Jacobs, Naidu, and Suri (2018). Finally, see Chen and Konstan (2015) for a survey of several experiments on different platforms.

nicate with each other and to our knowledge also no informal website exists through which this happened. There are also restrictions on accounts and on participation per IP address, which we discuss below. Moreover, only a subset of registered potential workers meeting our criteria are invited to participate, which limits the probability that registered participants who know each other are invited.¹⁶ In a post-experiment questionnaire, 95.4% of workers reported that they did not discuss the task with someone else when deciding whether or not to participate, and a similar number (97.4%) declared that they completed the task without the help of someone else. In a field experiment inside a firm, in contrast, workers might communicate with others, observe how much others work, and news might spread that there exist different payment schemes. Also, in the laboratory it is likely that participants receive cues about the behavior of others because the typical experiment has multiple participants inside the same room.¹⁷

The second advantage is that it enables us to recruit a relatively large number of workers at reasonable costs, which increases the statistical power to detect differences in labor supply across payment schemes. Moreover, the online platform provides access to a pool of workers with diverse demographic backgrounds, which arguably increases external validity relative to a sample of undergraduate students. The third advantage is that, with approval from an ethical committee, we were allowed to engage in gender discrimination on the platform. This would have been difficult to implement in a field experiment inside a firm for legal and other reasons.

We took precautions to minimize potential problems that are associated with conducting experiments using online platforms. As participants have to read and understand the instructions of the experiment without support, it might be

¹⁶Our criteria were: UK is the country of residence, registration as a man or woman, and an approval rate of at least 80% for previous participation in other studies. The number of registered individuals meeting our criteria was greater than 6,000. Workers meeting the selection criteria could register for our experiment without receiving an invitation email if they logged in on the website and selected our experiment, provided that our required number of workers had not been attained. The fact that our experiment was almost fully conducted within 24 hours limits this possibility. Importantly, there is no gain for workers from discussing or working with someone else during our experiment due to the nature of the task.

¹⁷We have considered to run the experiment in the laboratory, but decided against it because it would have been very difficult and expensive to avoid peer effects. For instance, any worker who stops working and leaves the laboratory is likely to be noticed by other workers. An alternative would have been to have only one worker at a time in the laboratory, but this causes at least three problems: (1) it is extremely time consuming to collect a large enough number of observations, (2) it may open the door to session effects (Fréchette, 2012), and (3) information regarding the experiment can spread among potential participants, because the experiment would take place over a long time period.

that they do not read them carefully enough or do not fully understand them. To ensure proper reading and understanding, we required participants to correctly answer nine exhaustive comprehension questions about the instructions. Participants who failed a question three times were automatically excluded and did not participate in the experiment. Another issue might be that participants do not fully trust that the instructions are truthful because online platforms do not necessarily have the reputation to be deception-free. To minimize this possibility, we made clear in the invitation to the experiment and again in the instructions that we do not use deception and that this is the standard in economic experiments. In a post-experiment questionnaire, participants report that they understood the instructions well and that they largely trusted that the instructions were truthful.¹⁸

Another concern might be that workers participate more than once, because it is not possible to directly verify the identity of participants. To minimize this possibility, the platform employs a number of measures to prevent duplicate accounts. These measures include limiting participation in an experiment to once per account, limiting the number of accounts per IP address, limiting participation in an experiment to once per IP address, requiring a unique non-voice over IP phone number per account, and limiting accounts to one per Paypal or Circle account for payment.¹⁹ The platform also forbids the use of VPNs and tracks changes in the country of connection and other suspicious participation patterns.

IV.C. Procedures and Demographics

Workers on the online platform were invited via email in January 2018 and freely decided whether or not to participate.²⁰ In total 1,271 workers success-

¹⁸The comprehension questions also prevented automatic programs (robots) from entering our experiment by passing as human workers. Importantly, exclusion of subjects who failed was independent of the payment schemes, because they were excluded from the experiment before they were allocated to a scheme. Participants reported how well they understood the instructions on a Likert scale from 1 (Not at All) to 7 (Very Well) with a mean answer of 6.27 (SD = 1.04, $N = 1,254$) and reported how much they trusted the instructions on a Likert scale from 1 (Not at All) to 7 (Completely) with a mean answer of 5.94 (SD = 1.42, $N = 1,254$).

¹⁹Our study was registered using two separate experiments on the platform, one only accessible to men and one only accessible to women. This is a feature of the platform, which requests that filtering by gender be done in this manner. This means that participation in the study was limited to one man and one woman per IP address (using gender reported on the platform). A total of 6% of participants had the same IP address at the time of their participation as another participant of the other gender. The address can be the same for different reasons, e.g., workers participate from the same house, public space or workplace. Note that Paypal and Circle also take steps to prevent duplication of accounts.

²⁰The invitation email can be found in Section 2.C of the Appendix.

fully completed the comprehension questions and participated in the experiment. On average, those workers spent 26.35 (SD = 15.56) minutes in the experiment, and were paid 2.64 (SD = 1.53) pounds. Table 2.2 summarizes their demographic characteristics.

Table 2.2: Demographic Characteristics of Workers

Demographic Characteristic	Mean (SD) or Percentage
Age	38 (12)
Task Experience on Platform	141 (176)
Woman	50%
Student	16%
UK National	93%
Caucasian/White	88%
Employed Full-Time	50%
Employed Part-Time	20%
Job Seeker	18%
Not in Paid Work	6%
Other Work Situation	5%

Note: N varies between 1,263 and 1,271 by characteristic because we could not obtain some characteristics from the platform for a few workers.

Not everyone who logged into the experiment actually participated. Specifically, the sample does not include the following individuals. First, the software automatically prevented 281 individuals who did not complete the comprehension questions from being assigned to a payment scheme and starting the task. Second, eight individuals quit at the end of the comprehension questions or during the practice phase. Third, 48 individuals were excluded because they exceeded the time limit.²¹ Fourth, seven individuals were removed because the reported gender in the experiment did not correspond to their gender in the platform database.

V. Hypotheses

Our four hypotheses were registered before the execution of the experiment²² and all refer to workers who do not beat the odds. These are all low-wage workers who faced a payment scheme with (i) equal low wages without chances, (ii) fair chances to receive a high wage, or (iii) low chances to receive a high wage, and all high-wage workers who faced a payment scheme with (i) equal high wages without chances, (ii) fair chances to receive a high wage, or (iii) high chances to receive a high wage. We only consider those workers because too

²¹These individuals went over the time limit of 65 minutes despite being explicitly forbidden from doing so in the experiment description on the platform, and being provided with a time countdown from 65 to 0 minutes during their work to remind them of the time limit.

²²American Economic Association's Randomized Control Trials Registry (ID: AEARCTR-0002655).

few workers beat the odds for an informative statistical analysis. The first three hypotheses are based on the optimal labor supply derived in the theoretical framework presented in Section III.²³

Table 2.3: Predicted Labor Supply in Each Payment Scheme

Low-Wage Worker i	
Payment Scheme	
EQLOW	$l_i^E(w_l) = w_l$
UNEQFAIR	$l_i^F(w_l) = w_l - P_i(w_l, w_h)$
UNEQUNFAIR	$l_i^U(w_l) = w_l - P_i(w_l, w_h) - A_i(Ew_l, Ew_h)$
UNEQDISCR	$l_i^D(w_l) = w_l - P_i(w_l, w_h) - A_i^D(Ew_l, Ew_h)$
High-Wage Worker i	
Payment Scheme	
EQHIGH	$l_i^E(w_h) = w_h$
UNEQFAIR	$l_i^F(w_h) = w_h - P_i(w_h, w_l)$
UNEQUNFAIR	$l_i^U(w_h) = w_h - P_i(w_h, w_l) - A_i(Ew_h, Ew_l)$
UNEQDISCR	$l_i^D(w_h) = w_h - P_i(w_h, w_l) - A_i^D(Ew_h, Ew_l)$

Note: The predicted labor supply of a worker i is given by Equation (2.4) in Section III: $l_i^* = w_i - P_i(w_i, w_j) - A_i(Ew_i, Ew_j)$, where the term $A_i^D(Ew_i, Ew_j)$ indicates the presence of gender-discriminatory chances ($A_i^D(Ew_i, Ew_j) > A_i(Ew_i, Ew_j)$). Predictions are for workers who do not beat the odds.

Table 2.3 shows the predicted labor supply in each payment scheme, in the upper part for low-wage workers, and in the lower part for high-wage workers. Recall that, for unequal wages and unfair chances, P_i and A_i take on positive values, and for unfair chances, A_i takes on smaller values than A_i^D . From the table, it then follows straightforwardly that, for a low-wage worker i ,

$$l_i^E(w_l) > l_i^F(w_l) > l_i^U(w_l) > l_i^D(w_l),$$

and for a high-wage worker i ,

$$l_i^E(w_h) > l_i^F(w_h) > l_i^U(w_h) > l_i^D(w_h).$$

This leads to our first two hypotheses.

HYPOTHESIS 1, LOW-WAGE WORKERS: For low-wage workers, labor supply ranks across payment schemes as follows: EQLOW > UNEQFAIR > UNEQUNFAIR > UNEQDISCR.

HYPOTHESIS 2, HIGH-WAGE WORKERS: For high-wage workers, labor supply ranks across payment schemes as follows: EQHIGH > UNEQFAIR > UNEQUNFAIR > UNEQDISCR.

²³The theoretical optimal labor supply is based on the interior solution. We conducted a pilot study in advance to ensure that the parameters of the experiment (e.g., piece-rate wages, length of lines, duration) do not produce too many corner outcomes where workers do not work at all or finish all tasks. As we will see later in Section VI, workers in the experiment indeed overwhelmingly choose an interior outcome.

Recall that we assume that being in a disadvantageous position (low wage or low chance of receiving the high wage) creates a larger marginal disutility from inequality than being in an advantageous position (high wage or high chance of receiving the high wage).²⁴ We also assume that gender-discriminatory unfair chances are disliked more than unfair chances from an unspecified source, and that negative discrimination is worse than positive discrimination. This implies that, at a given wage, unequal wages, unfair chances, and gender discrimination each decrease the optimal labor supply more for a low-wage worker than for a high wage worker. More formally, it holds that²⁵

$$\begin{aligned} l_i^E(w_l) - l_i^F(w_l) &> l_i^E(w_h) - l_i^F(w_h), \\ l_i^F(w_l) - l_i^U(w_l) &> l_i^F(w_h) - l_i^U(w_h), \text{ and} \\ l_i^U(w_l) - l_i^D(w_l) &> l_i^U(w_h) - l_i^D(w_h). \end{aligned}$$

Our third hypothesis is therefore as follows.

HYPOTHESIS 3, DISADVANTAGE VS. ADVANTAGE: For each of the following comparisons, the labor supply decrease is greater for low-wage workers than for high-wage workers: UNEQFAIR vs. EQLOW/EQHIG, UNEQUNFAIR vs. UNEQFAIR, and UNEQDISCR vs. UNEQUNFAIR.

As gender discrimination is generally experienced by women rather than by men in society, discriminating against women might have a different effect on labor supply than discriminating against men. However, the a priori direction of the difference is not clear. On the one hand, discrimination against women worsens existing inequalities and may be especially painful for them, and men might perceive discrimination against them as a justified compensation for everyday discrimination of women. This might lead to a strong negative labor supply reaction by women but a positive or neutral one by men. On the other hand, women may have weaker negative labor supply reactions because they are used to discrimination, and men may be habituated to higher chances so that their new experience of lower chances may be especially frustrating and thus strongly decrease their labor supply. Therefore, our fourth hypothesis, which concerns this possible gender difference, is not directed and we state the null hypothesis.

²⁴See Section III, Equations (2.2) and (2.3) together with $\alpha_i > \beta_i$ and $\alpha'_i > \beta'_i$.

²⁵To obtain these three inequalities we use Table 2.3. For the first inequality we have $l_i^E(w_l) - l_i^F(w_l) = P_i(w_l, w_h)$ and $l_i^E(w_h) - l_i^F(w_h) = P_i(w_h, w_l)$. Because $\alpha_i > \beta_i$, it holds that $P_i(w_l, w_h) > P_i(w_h, w_l)$. The two other inequalities are obtained in the same manner.

HYPOTHESIS 4, GENDER AND NEGATIVE DISCRIMINATION: The difference in labor supply between UNEQUNFAIR and UNEQDISCR is equal for both genders.

VI. Results

In this section we first report descriptive statistics regarding labor supply under the different payment schemes followed by tests of our four hypotheses. Thereafter, we present some exploratory analyses on potential gender differences beyond our hypotheses and on labor supply responses at the extensive and intensive margins. Recall that we measure labor supply as the number of lines entered during the experiment and that workers could stop working at any time.

VI.A. Descriptive Statistics

Figure 2.2 and Table 2.4 summarize labor supply by type of worker and payment scheme.²⁶ The left part of the figure and table present the labor supply of low-wage workers, and the right part that of high-wage workers. For low-wage workers, the scheme EQLOW generates the largest mean labor supply, followed by UNEQUNFAIR and UNEQFAIR (with little difference between them), and then by UNEQDISCR. For high-wage workers, all schemes produce a comparable mean labor supply.²⁷

Table 2.4: Labor Supply per Payment Scheme

	Low-Wage Workers					High-Wage Workers				
	Mean	SD	Min.	Max.	<i>N</i>	Mean	SD	Min.	Max.	<i>N</i>
EQLOW/EQHIG	43.20	27.63	.04	.16	128	44.03	28.93	.04	.20	128
UNEQFAIR	37.44	29.16	.06	.17	125	43.23	29.64	.03	.20	127
UNEQUNFAIR	39.22	27.76	.04	.15	143	43.50	28.85	.02	.20	143
UNEQDISCR	33.62	26.41	.08	.11	143	45.74	26.72	.02	.16	145

Note: Labor supply is measured by the number of lines completed and ranges from 0 to 85. Min. and Max. indicate the percentage of workers completing the minimum and maximum number of lines.

²⁶Workers beating the odds are excluded unless we state otherwise (see Section V).

²⁷A comparison between low and high wages shows a significantly larger labor supply when a high wage is paid only if gender discrimination is involved (EQLOW/EQHIG, $p = 0.72$; UNEQFAIR, $p = 0.12$; UNEQUNFAIR, $p = 0.20$; UNEQDISCR, $p < 0.001$; two-sided t -tests). Non-parametric Mann-Whitney tests qualitatively lead to the same results. The implied wage elasticities of labor supply for our task are 0.02 in EQLOW/EQHIG, 0.16 in UNEQFAIR, 0.11 in UNEQUNFAIR, and 0.36 in UNEQDISCR. Overall, these elasticities are in keeping with those estimated on online labor platforms. For instance, Dube et al. (2018) estimate the market-wide elasticity on Amazon Turk to be around 0.10.

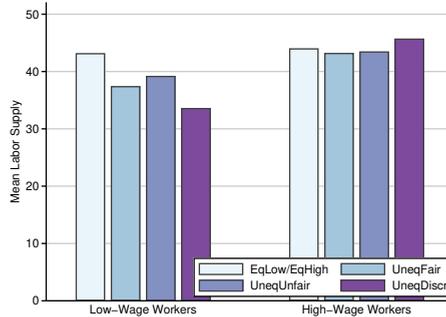


Figure 2.2: Mean Labor Supply per Payment Scheme

Note: Labor supply is measured by the number of lines completed and ranges from 0 to 85. N ranges from 127 to 145 workers per payment scheme.

VI.B. Test of Hypotheses

To test our hypotheses we employ non-parametric rank tests as well as Tobit regressions to account for lower-bound and upper-bound censoring of the dependent variable. For the few cases where the two techniques lead to different results in terms of statistical significance, we give priority to the non-parametric tests.²⁸ Table 2.5 shows coefficient estimates of the Tobit regressions of labor supply on dummies for the payment schemes, separately for low-wage workers and high-wage workers. The payment schemes EQLOW and EQHIGH, respectively, serve as baselines.

To test Hypotheses 1 and 2, we start by testing the null hypothesis of equality of all schemes, and then conduct pairwise comparisons between the schemes using non-parametric Dunn’s tests as well as the regression estimates.²⁹ Both non-parametric Kruskal-Wallis (KW) tests and parametric Wald tests (on the restriction that the three scheme coefficients from the regression are jointly equal to zero) confirm that labor supply differs across the four schemes for low-wage workers, but not for high-wage workers (low-wage workers: $p = 0.033$ KW, $p = 0.006$ Wald; high-wage workers: $p = 0.825$ KW, $p = 0.901$ Wald;

²⁸We do so because non-parametric tests do not assume that error terms are normally distributed. For the Tobit regressions, we use robust standard errors because we find evidence of heteroscedasticity in our data. In the regressions, we also include control variables (age, gender, ethnicity, student status, employment status, experience on the platform, an index reflecting the percentage of approved participation in tasks on the platform, and day and time of participation).

²⁹The Dunn’s test (Dunn, 1964) allows us to conduct multiple pairwise comparisons, and is considered to be the correct test after a Kruskal-Wallis test. In the regressions, removing the controls does not change the results qualitatively.

Table 2.5: Tobit Regressions of Labor Supply on Payment Schemes

Scheme	Low-Wage Workers		High-Wage Workers	
	(1)	(2)	(3)	(4)
UNEQFAIR	-4.144 (4.337)	-1.309 (4.547)		
UNEQUNFAIR	-4.008 (4.115)	-1.607 (4.397)		
UNEQDISCR	-13.610**** (4.051)	1.146 (4.220)		
Controls	Yes	Yes		
Prob > F	0.001	0.010		
Pseudo R^2	0.011	0.008		
N	533	542		

Note: EQLOW/EQHIGH serves as baseline. Standard errors are in parentheses. Two-sided p -values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

two-sided tests). For the pairwise comparisons, we use one-sided tests because our hypotheses are directional. Table 2.6 presents the p -values of the tests for the three main comparisons contained in each of the two hypotheses, with and without the Benjamini-Hochberg (BH) correction for multiple hypothesis testing within each hypothesis.³⁰

Table 2.6: P -values of Predicted Differences in Labor Supply between Payment Schemes

Predicted Inequality	Low-Wage Workers				High-Wage Workers			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Technique		Technique		Technique		Technique	
BH Correction	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit
	No	No	Yes	Yes	No	No	Yes	Yes
EQLOW/EQHIGH > UNEQFAIR	0.017	0.170	0.050	0.255	0.371	0.387	1.000	1.000
UNEQFAIR > UNEQUNFAIR	0.878	0.513	0.878	0.513	0.780	0.473	1.000	0.709
UNEQUNFAIR > UNEQDISCR	0.021	0.006	0.031	0.019	0.886	0.752	0.886	0.752
N	539	533	539	533	543	542	543	542

Note: One-sided p -values in the direction predicted. BH corrections account for multiple hypothesis testing.

For low-wage workers, Table 2.6 shows that UNEQFAIR significantly decreases labor supply compared to EQLOW if we use non-parametric Dunn's tests, in line with Hypothesis 1. However, in contrast to this hypothesis, unfair chances have no additional negative effect, as labor supply in UNEQFAIR and UNEQUNFAIR are not significantly different. Finally, UNEQDISCR significantly reduces labor supply compared to UNEQUNFAIR, as predicted.³¹ The

³⁰Table 2.13 in the Appendix presents the same for the six possible pairwise comparisons; the results are the same. The BH correction (Benjamini and Hochberg, 1995) is a common False Discovery Rate procedure, which controls for the probability of false positives among significant results. It differs from Family-wise Error Rate procedures such as Holm-Bonferroni (Holm, 1979), which control for the probability of at least one false positive among significant results.

³¹Table 2.13 in the Appendix shows that UNEQDISCR also reduces labor supply compared to

detected effects are also economically significant, which can be assessed using Table 2.4. Unfair chances due to gender discrimination decreases mean labor supply from 43.20 in EQLOW to 33.62 in UNEQDISCR, corresponding to a reduction of 22%. Mean labor supply in UNEQUNFAIR is 39.22, indicating that gender-discriminatory unfair chances reduce mean labor supply by 15% compared to unfair chances from an unspecified source. Moreover, mean labor supply in UNEQFAIR is 37.44, representing a decrease of 13% compared to EQLOW.³²

For high-wage workers, Table 2.6 reveals that none of the predicted inequalities in Hypothesis 2 hold. That is, high-wage workers provide similar labor supply across payment schemes. We summarize our first two results as follows.

RESULT 1, LOW-WAGE WORKERS: *(a) Gender-discriminating unfair chances (UNEQDISCR) lower labor supply, compared to each other scheme. (b) Unfair chances from an unspecified source (UNEQUNFAIR) do not decrease labor supply compared to fair chances from an unspecified source (UNEQFAIR). (c) UNEQFAIR and UNEQUNFAIR reduce labor supply compared to EQLOW.*

RESULT 2, HIGH-WAGE WORKERS: *All payment schemes (UNEQDISCR, UNEQUNFAIR, UNEQFAIR, EQHIGH) produce similar labor supply.*

We now turn to Hypothesis 3, which states that the labor supply decrease is larger for low-wage workers than for high-wage workers, when comparing the payment schemes UNEQFAIR with EQLOW/EQHIGH, UNEQUNFAIR with UNEQFAIR, and UNEQDISCR with UNEQUNFAIR, respectively. We evaluate the hypothesis with a Tobit regression using dummy variables for UNEQFAIR, UNEQUNFAIR, and UNEQDISCR as well as their interactions with a dummy variable for high-wage workers. We also include a set of controls that is common to both low-wage and high-wage workers. Table 2.7 presents the coefficient estimates. The overall null hypothesis, that there is no difference in labor supply between low-wage and high-wage workers in all payment scheme comparisons simultaneously, is represented by Restriction I at the bottom of the table. A Wald test marginally rejects this restriction (two-sided p -value = 0.073).³³

Table 2.8 presents p -values from Wald tests conducted separately for each of the three inequalities that compose Hypothesis 3. It shows that the decreases

EQLOW and UNEQUNFAIR. Moreover, UNEQUNFAIR reduces labor supply compared to EQLOW.

³²In terms of pooled standard deviations, UNEQDISCR decreases labor supply by 0.35 standard deviations compared to EQLOW, and by 0.21 standard deviations compared to UNEQUNFAIR. The scheme UNEQFAIR reduces labor supply by 0.20 standard deviations relative to EQLOW. We estimate that our design can detect differences between schemes of approximately 0.20 standard deviations with medium statistical power ($\geq 50\%$).

³³Without controls, the p -value is 0.154.

Table 2.7: Tobit Regression of Labor Supply on Payment Schemes with Both Types of Workers

Scheme	All Workers
UNEQFAIR	-4.636 (4.385)
UNEQUNFAIR	-4.432 (4.098)
UNEQDISCR	-13.133*** (4.058)
UNEQFAIR \times HighWage	2.030 (6.318)
UNEQUNFAIR \times HighWage	4.066 (6.050)
UNEQDISCR \times HighWage	13.307** (5.790)
HighWage	2.291 (4.352)
Controls	Yes
Prob $> F$	0.000
Pseudo R^2	0.008
N	1075
Restriction I	
UNEQFAIR \times HighWage = 0	
UNEQUNFAIR \times HighWage = UNEQFAIR \times HighWage	
UNEQDISCR \times HighWage = UNEQUNFAIR \times HighWage	
Wald Test (two-sided p -value) = 0.073	

Note: EQLOW serves as baseline. Standard errors are in parentheses. Two-sided p -values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

Table 2.8: P -values of Predicted Differences in Labor Supply Effect of Payment Schemes between Low-Wage Workers and High-Wage Workers

Predicted Inequality	All Workers	
	(1)	(2)
	Tobit	Tobit
	No	Yes
UNEQFAIR \times HighWage > 0	0.374	0.561
UNEQUNFAIR \times HighWage $>$ UNEQFAIR \times HighWage	0.503	0.503
UNEQDISCR \times HighWage $>$ UNEQUNFAIR \times HighWage	0.023	0.070
N	1075	1075

Note: One-sided p -values are presented, in the direction predicted. BH corrections account for multiple hypothesis testing.

in labor supply caused by UNEQFAIR compared to EQLOW/EQHIG and by UNEQUNFAIR compared to UNEQFAIR are not significantly larger for low-wage workers than for high-wage workers, rejecting the first and second inequality of the hypothesis. However, as predicted, the labor supply reduction caused by UNEQDISCR relative to UNEQUNFAIR is (marginally) significantly larger for low-wage workers. Overall, our analysis provides evidence only in favor of the discrimination part of Hypothesis 3. We state our third result as follows.

RESULT 3, DISADVANTAGE VS. ADVANTAGE: (a) *The decrease in labor supply caused by gender-discriminating unfair chances (UNEQDISCR) relative to unfair chances from an unspecified source (UNEQUNFAIR) is larger for low-wage workers than for high-wage workers.* (b) *The decreases caused by unfair chances from an unspecified source (UNEQUNFAIR) relative to fair chances from an unspecified source (UNEQFAIR) and by UNEQFAIR relative to EQLOW/EQHIGH are similar for both types of workers.*

Finally, we evaluate whether low-wage men and women respond differently to gender-discriminating unfair chances (UNEQDISCR) relative to unfair chances coming from an unspecified source (UNEQUNFAIR), as stated in Hypothesis 4. Table 2.9 presents descriptive statistics regarding labor supply for the two schemes, separately by gender. The right part of the table shows that the mean labor supply of women differs considerably between the two schemes (42.37 in UNEQUNFAIR and 30.38 in UNEQDISCR; roughly 28%). In contrast, as can be seen in the left part of the table, the mean labor supply of men is essentially equal in both schemes (36.13 in UNEQUNFAIR and 36.92 in UNEQDISCR).

Table 2.10 presents the estimates from a Tobit regression for low-wage workers in the two schemes with UNEQUNFAIR serving as the baseline, a dummy variable for women, and an interaction term of UNEQDISCR with the dummy variable for women. For men, discrimination does not significantly alter labor supply compared to unfair chances, as the coefficient of UNEQDISCR is insignificant ($p = 0.805$). However, the interaction term is negative and significant ($p = 0.040$), indicating that the labor supply decrease caused by gender-discriminatory chances relative to the same unfair chances without gender discrimination is stronger for women than for men.³⁴ Our fourth result is as follows.

RESULT 4, GENDER AND NEGATIVE DISCRIMINATION: *The decrease in labor supply caused by UNEQDISCR relative to UNEQUNFAIR is greater for low-wage women*

³⁴The interaction term is also significant without controls ($p = 0.043$).

Table 2.9: Labor Supply of Low-Wage Men and Women in UNEQUNFAIR and UNEQDISCR

	Low-Wage Men					Low-Wage Women				
	Mean	SD	Min.	Max.	<i>N</i>	Mean	SD	Min.	Max.	<i>N</i>
UNEQUNFAIR	36.13	27.44	.04	.14	72	42.37	27.92	.04	.15	71
UNEQDISCR	36.92	29.27	.06	.15	71	30.38	23.00	.11	.06	72

Note: Labor supply is measured by the number of lines completed and ranges from 0 to 85. Min. and Max. indicate the percentage of workers completing the minimum and maximum number of lines.

Table 2.10: Tobit Regression of Labor Supply in UNEQUNFAIR and UNEQDISCR, for Low-Wage Men and Women

Scheme	Low-Wage Workers
UNEQDISCR	-1.363 (5.510)
UNEQDISCR \times Woman	-16.105** (7.805)
Woman	1.736 (6.077)
<hr/>	
Controls	Yes
Prob $> F$	0.039
Pseudo R^2	0.012
N	283

Note: UNEQUNFAIR serves as baseline. Standard errors are indicated in parentheses. Two-sided p -values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

than for low-wage men.

In summary, we find partial support for our hypotheses. In payment schemes involving chances to receive a low wage or a high wage, low-wage workers reduce labor supply relative to a scheme in which the low wage payment is the same for everyone. This effect is especially pronounced when chances are gender-discriminatory. Interestingly, there is no difference between fair and unfair chances when these come from an unspecified source. For high-wage workers, labor supply appears to be unaffected by the nature of the payment scheme. There is also partial evidence that the decrease in labor supply is stronger for low-wage workers than for high-wage workers, especially when gender-discriminatory chances are involved. Finally, women respond more strongly to gender-discriminatory chances than men do.

VI.C. Exploratory Analyses

Here we report some additional analyses that go beyond our pre-registered hypotheses. We first delve into other potential differences between genders and we then examine the effect of payment schemes on labor supply at the extensive and intensive margins.

Further Gender Differences. Figure 2.3 and Table 2.11 show descriptive statistics regarding the labor supply of workers in each payment scheme, separately for men and women. A comparison of the behavior of low-wage workers shows that the labor supply reactions to the different types of payment schemes differ considerably between genders. Men lower their labor supply in response to any of the three payment schemes with unequal wages, whereas women

decrease their labor supply only in response to unequal wages resulting from gender-discriminatory chances.

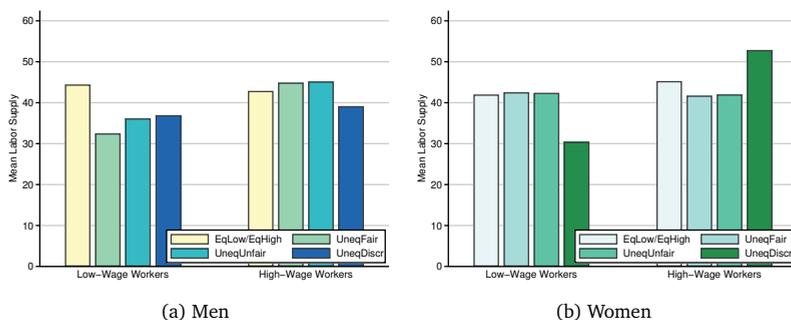


Figure 2.3: Mean Labor Supply per Gender and Payment Scheme

Note: Labor supply is measured by the number of lines completed and ranges from 0 to 85. N ranges from 62 to 75 workers per payment scheme.

Table 2.11: Labor Supply per Gender and Payment Scheme

	Low-Wage Workers					High-Wage Workers				
	Mean	SD	Min.	Max.	N	Mean	SD	Min.	Max.	N
Men										
EQLOW/EQHIG	44.40	27.20	.02	.18	65	42.83	27.95	.06	.16	64
UNEQFAIR	32.46	28.70	.10	.13	63	44.85	29.35	.03	.19	62
UNEQUNFAIR	36.13	27.44	.04	.14	72	45.13	30.05	.01	.23	69
UNEQDISCR	36.92	29.27	.06	.15	71	39.12	25.08	.04	.09	75
Women										
EQLOW/EQHIG	41.95	28.24	.06	.14	63	45.23	30.05	.02	.25	64
UNEQFAIR	42.50	28.97	.03	.21	62	41.68	30.05	.03	.20	65
UNEQUNFAIR	42.37	27.92	.04	.15	71	41.97	27.80	.03	.16	74
UNEQDISCR	30.38	23.00	.11	.06	72	52.83	26.77	.03	.23	70

Note: Labor supply is measured by the number of lines completed and ranges from 0 to 85. Min. and Max. indicate the percentage of workers completing the minimum and maximum number of lines.

We test for differences in labor supply across payment schemes separately for men and women, using non-parametric Dunn's tests and Tobit regression estimates.³⁵ Table 2.12 reports the p -values of the tests for the three inequalities predicted by our first two hypotheses, now separately for men and for women. The tests corroborate the impression one has from looking at the descriptive statistics. Low-wage male workers significantly decrease their labor supply in UNEQFAIR compared to EQLOW, but do not further decrease labor supply in

³⁵For the regression, we use the same specification as in Table 2.5, but include a dummy variable for women and interact this variable with the payment schemes. These Tobit regression estimates can be found in Table 2.14 of the Appendix.

response to the additional inequalities contained in UNEQFAIR and UNEQDISCR. In contrast, low-wage female workers significantly decrease their labor supply only in UNEQDISCR relative to UNEQUNFAIR. Moreover, relative to EQLOW, low-wage men do significantly decrease their labor supply in reaction to any of the three types of inequalities (p -values between 0.002 and 0.106, depending on test and BH correction), whereas low-wage women significantly decrease their labor supply only when facing gender-discriminatory chances ($p \leq 0.009$ and $p \geq 0.548$ otherwise). In contrast, workers who receive a high wage do not significantly reduce their labor supply in response to the different payment schemes, irrespective of their gender.³⁶

Table 2.12: Tests of Inequalities in Labor Supply between Payment Schemes, for Men and Women

Predicted Inequality	Low-Wage Workers				High-Wage Workers				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
	Technique	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit
BH Correction	No	No	Yes	Yes	No	No	Yes	Yes	
Men									
EQLOW/EQHIG > UNEQFAIR	0.002	0.016	0.006	0.049	0.818	0.783	0.818	0.783	
UNEQFAIR > UNEQUNFAIR	0.920	0.716	0.920	0.716	0.248	0.508	0.372	0.762	
UNEQUNFAIR > UNEQDISCR	0.238	0.398	0.358	0.598	0.063	0.123	0.190	0.369	
<i>N</i>	271	268	271	268	270	270	270	270	
Women									
EQLOW/EQHIG > UNEQFAIR	0.767	0.732	0.767	0.732	0.104	0.206	0.311	0.619	
UNEQFAIR > UNEQUNFAIR	0.751	0.305	1.000	0.457	0.801	0.517	1.000	0.775	
UNEQUNFAIR > UNEQDISCR	0.003	0.001	0.009	0.004	0.994	0.981	0.994	0.981	
<i>N</i>	268	265	268	265	273	272	273	272	

Note: One-sided p -values are presented, in the direction predicted. BH corrections account for multiple hypothesis testing.

For high-wage female workers the descriptive statistics suggest an interesting pattern opposite to Hypothesis 2, namely an increase in labor supply in response to positive discrimination (see, e.g., Panel (b) of Figure 2.3). We compare UNEQDISCR and UNEQUNFAIR as well as UNEQDISCR and EQHIGH without imposing the predicted direction on our tests. We find that high-wage women in UNEQDISCR significantly increase their labor supply relative to UNEQUNFAIR ($p \leq 0.038$, two-sided). The increase is also marginally significant relative to

³⁶More precisely, high-wage men reduce their labor supply in UNEQDISCR compared to UNEQFAIR, but only at marginally significant levels without the BH correction. For details on those comparisons, please see Table 2.15 in the Appendix.

EQHIGH if we use non-parametric tests ($p = 0.055$, two-sided).³⁷

Extensive and Intensive Margins. To evaluate whether different payment schemes affect labor supply of low-wage workers at the extensive or the intensive margin, we use a two-equation hurdle model with a lower bound. We find that relative to UNEQUNFAIR, the scheme UNEQDISCR reduces both the probability that low-wage workers start to work at all ($p = 0.026$ without and $p = 0.079$ with BH correction) and the labor supply of those who do decide to work ($p = 0.029$ without and $p = 0.086$ with BH correction). The negative effect on labor supply of UNEQFAIR relative to the scheme EQLOW appears to affect only the intensive margin (intensive margin: $p = 0.044$ without and $p = 0.087$ with BH correction; extensive margin: $p \geq 0.228$).³⁸

VII. Discussion

VII.A. Unequal Wages and Labor Supply

Bracha et al. (2015), Breza et al. (2018), and Dube et al. (2019) (henceforth *BGL*, *BKS*, and *DGL*) study the effect of unequal wages on labor supply, but do not investigate the role of chances and discrimination. *BGL* conduct a laboratory experiment, *BKS* employ a field experiment, and *DGL* exploit a natural experiment. Here, we briefly compare our results concerning unequal wages resulting from fair chances or unfair chances with an unspecified source (that is, UNEQFAIR or UNEQUNFAIR) to the results of these three studies.³⁹

In line with these three studies, we obtain the result that unequal wages significantly decrease labor supply of low-wage workers. In addition, our finding that unequal wages do not affect labor supply of high-wage workers is also found by *BGL* and *DGL*, although *BKS* report evidence that such workers may reduce their labor supply. Furthermore, regarding gender differences, we

³⁷Table 2.16 in the Appendix reports p -values of comparisons and tests.

³⁸ P -values are one-sided in the predicted direction. Appendix Tables 2.17 and 2.18, respectively, show the results of the estimation and the p -values for the pairwise comparisons between payment schemes. In the Appendix we also provide further supplementary analyses. Tables 2.19–2.22 detail labor supply reactions to the payment schemes for four additional demographic groups: the young and the old, and the full-time and part-time employed or unemployed. Furthermore, our pre-registered measure of labor supply is the number of lines, but other measures of labor supply are conceivable. Tables 2.23 and 2.24 show workers' responses to the different payment schemes if we use time spent in the experiment as the labor supply measure. Note that this measure does not only include time spent working on the task. As such, it is most likely not an appropriate measure of labor supply. Nevertheless, the results are very similar to our pre-registered measure. Lastly, for completeness, Tables 2.25 and 2.26 describe how the relatively few workers who beat the odds respond to the different payment schemes.

³⁹*BGL* and *BKS* also consider cases where unequal wages may be justified, e.g., by productivity differences. We do not discuss these results.

find that only low-wage men negatively respond to wage differences, whereas women do not. This is consistent with the result reported in *BGL* that only men respond to wage inequality.⁴⁰

Thus, our findings corroborate the existing results on the negative labor supply effects of wage inequality and add evidence that this may be mainly driven by low-wage men.

VII.B. Fairness of Chances

One motivation of our research is to explore how the (un)fairness of initial chances between two workers affects their labor supply decisions once wages are known. To the best of our knowledge, no other study investigates this question. However, a number of scholars have analyzed the effect of ex ante fairness more generally. Closest to our research are the studies examining whether and how fairness of ex ante chances between individuals influences their equity judgments.

Starting with Diamond (1967), many have argued that an unequal outcome is more acceptable if it is generated by fair chances than if it is generated by unfair chances (see Section II.B). However, there may be a tension between ex ante and ex post fairness (e.g., Myerson, 1981; Fudenberg and Levine, 2012; Saito, 2013). Nevertheless, experimental studies have shown that individuals are more likely to accept an unequal outcome that results from fair chances than one that results from unfair chances (e.g., Bolton et al., 2005; Grimalda et al., 2016). In stark contrast, we find that behavior of workers is insensitive to initial chances when they are generated by an unspecified source, as labor supply is almost identical under the payment schemes *UNEQFAIR* and *UNEQUNFAIR*.⁴¹ However, workers in our study do respond to unfair chances coming from gender discrimination, a prominent form of procedural unfairness. This complements findings that workers' reaction to unequal wages depends on the reason behind wage inequality (e.g., workers may accept or even demand wage inequality if it reflects productivity differences, Abeler et al., 2010; Breza et al., 2018, or if some other justification is provided, Bracha et al., 2015).

Our research carries a potential implication for managerial policy. Managers

⁴⁰*BKS* only sample men and *DGL* do not report effects by gender. There are also other demographic differences between the four studies. For instance, *BGL* employ American university students, *DGL* use young American part-time workers, and *BKS* use Indian temporary workers who may be older. We use online UK workers with a median age of 36, most of whom are either full-time or part-time employed.

⁴¹We note that, as in the aforementioned experiments, in *UNEQFAIR* and *UNEQUNFAIR* chances are assigned without a specific reason.

might be tempted to believe that offering initially fair chances will mitigate the negative effects of unequal wages in bonus or promotion schemes. However, our results suggest that, *ceteris paribus*, unequal wages resulting from fair chances will still exert a detrimental effect on disadvantaged workers' labor supply.

VII.C. Discrimination, Labor Supply, and the Gender Gap in Earnings

Goldin (2014) and Blau and Kahn (2017) report that the most important determinant of the modern gender earnings gap is that women exhibit lower labor supply. Women are less present in high-pay occupations, which usually demand long working hours (e.g., lawyer, manager, professor), and women work less and earn less *within the same occupation*, which typically offer rapidly-rising returns to working hours. The main explanation for the lower labor supply put forward in the literature is that women prefer temporal flexibility at work, notably working less hours, because they have to bear a much greater share of household responsibilities (e.g., Bertrand et al., 2010; Flabbi and Moro, 2012; Goldin, 2014; Goldin and Katz, 2016; Wiswall and Zafar, 2017; Cortés and Pan, 2019).⁴²

Our result that women decrease their labor supply in response to gender-discriminatory chances offers a complementary explanation for women's lower labor supply and lower earnings. Consider the following example. A young female lawyer starts at a law firm, and is as willing as anyone to work extra hours. The older partners promote similar young male lawyers in her place because they wrongly believe that, being a woman, she will be less willing to work extra hours. Our result then suggests that, unsatisfied at being discriminated, she chooses not to work extra hours whenever she can, thereby decreasing her value to the firm. Observing her behavior over time, the older partners are comforted in their initial choice, because their (formerly wrong) belief is now confirmed. Thus, the discrimination that women experience in labor markets may decrease their willingness to enter potential high-income occupations and may also reduce their willingness to work long hours for a given salary within an occupation. In this case, the effect of discrimination on labor supply can

⁴²In a recent field experiment Mas and Pallais (2017) find that women indeed have greater preference for temporal flexibility, but that this is not enough to explain the gender earnings gap. There is evidence supporting other explanations for the modern gender earnings gap, including discrimination (e.g., Neumark et al., 1996; Goldin and Rouse, 2000; Reuben et al., 2014; Sarsons et al., 2019), differences in bargaining behavior and competitiveness (e.g., Niederle and Vesterlund, 2007; Buser et al., 2014; Card et al., 2015), differences in productivity (e.g., Mulligan and Rubinstein, 2008), social norms (e.g., Coffman, 2014; Bertrand et al., 2015; Bursztyn et al., 2017), and stereotypes (e.g., Bernaldo et al., 2019).

become a self-fulfilling prophecy, and breaking this cycle may therefore prove to be no small feat.

Our result that gender-based unfair chances reduce female labor supply also challenges a commonly made assumption in the estimation of discrimination. Discrimination is regularly estimated by measuring the difference in earnings between men and women with the same characteristics—e.g., age, education, experience, hours worked and occupation—under the assumption that these characteristics are themselves unaffected by discrimination (see Oaxaca, 1973; Blinder, 1973, and Fortin et al., 2011 for an overview of the method). However, finding that men earn higher wages because they work longer hours—and that, hence, men and women would earn similar wages if they worked the same hours—can hide that women work less hours (and incur disutility) exactly because of discrimination (see our argument above). That is, the impact of gender discrimination is likely underestimated.

Finally, our data also provide suggestive evidence that positive gender discrimination may increase labor supply of women. Specifically, the labor supply of women is greater under unequal high wages with positive discrimination of women (UNEQDISCR) than under unequal high wages with unfair chances from an unspecified source (UNEQFAIR) and than under equal high wages (EQHIGH). This is consistent with a few studies finding beneficial effects of positive discrimination in other settings (Balafoutas and Sutter, 2012; Niederle et al., 2013; Ibañez and Riener, 2018), even though our study controls for monetary incentives. However, we are careful in drawing definite conclusions on this matter and leave this important question for future research.

VII.D. Possible Extensions

We have chosen an experiment design that allows us to explore the effects of unfair chances and gender discrimination in the cleanest possible way. Naturally, several extensions are conceivable to account for features of labor decisions that are not part of our design. The following four extensions appear particularly interesting. First, since employers often cannot fully contract effort, and it has been shown that employees tend to reciprocate higher wages with higher effort (e.g., Fehr et al., 1993, 1998b), one could study how the presence of discrimination affects reciprocal effort. It is conceivable that the possibility of gift-exchange—giving workers the possibility to negatively (positively) respond to an employer who negatively (positively) discriminates—by changing work effort will reveal even greater effects of discrimination on work

behavior than those observed in our setting of complete contracts. Second, one could augment our design to account for team production, and study how discrimination impacts the ability of teams to cooperate, coordinate, and produce. In such a setting discrimination may generate negative external effects on non-discriminated workers. Third, since promotion chances often take the form of multi-period tournaments where effort influences chances of promotion (e.g., Lazear and Rosen, 1981), one could study the effect of discrimination on labor supply in a more complex tournament-style design. Fourth, it may be interesting to study whether other common forms of discrimination, such as ethnic discrimination, generate similar effects as those we find for gender.

VIII. Conclusion

We provide the first causal evidence regarding how unfair chances to receive a low or high wage—stemming from an unspecified source of from gender discrimination—affect labor supply decisions. We find that, at a given wage, explicit gender discrimination in chances considerably reduces the labor supply of disadvantaged workers compared to an equal-wage setting (−22%). This is the case even though workers only hurt themselves by working less. Moreover, low wages stemming from gender-discriminatory chances reduce labor supply almost twice as much as low wages resulting from fair chances (−13%). Interestingly, in the absence of discrimination, low-wage workers are insensitive to whether unequal wages result from fair or unfair chances. Advantaged workers are unresponsive to any type of inequality that we examine.

Moreover, we find suggestive evidence that men and women respond differently to unfair chances and discrimination. Men reduce their labor supply when they are in a disadvantaged wage position, irrespective of the reason behind the disadvantage, whereas women decrease their labor supply only when their low wage can be attributed to gender-based discrimination in chances.

Our findings provide a novel supply-side effect of gender discrimination in labor markets, and offer a complementary way to account for the lower supply of women and the gender gap in earnings. More broadly, our study opens new avenues for research on the reactions of workers who face discrimination.

Appendix

2.A. Additional Tables

We provide additional figures and tables in Section 2.B and the instructions of the experiment as seen by the workers in Section 2.C. Each figure and table is referred to in the main text by its number.

2.B. Additional Figures and Tables

Line 37:
2zvqte548q2co4vk

Please select "Type this Line" to type the line or select "Leave" to stop typing lines and finish the experiment.

Type this Line Leave

Next Page >>

Line 37:
2zvqte548q2co4vk

Please enter the line:

Next Page >>

Figure 2.4: Task

Note: In the upper screenshot, the worker sees a line of characters, and decides whether to type the line or to leave the experiment. In the lower screenshot, if choosing not to leave, the worker is required to enter the line.

Table 2.13: *P*-Values of Six Predicted Differences in Labor Supply between Payment Schemes

Predicted Inequality	Low-Wage Workers				High-Wage Workers			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Technique	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	Dunn
BH Correction	No	No	Yes	Yes	No	No	Yes	Yes
EQLow/EqHigh > UNEQFAIR	0.017	0.170	0.050	0.212	0.371	0.387	1.000	1.000
EQLow/EqHigh > UNEQUNFAIR	0.058	0.165	0.087	0.248	0.426	0.357	1.000	1.000
EQLow/EqHigh > UNEQDISCR	0.001	0.001	0.006	0.002	0.853	0.607	1.000	0.910
UNEQFAIR > UNEQUNFAIR	0.878	0.513	0.878	0.513	0.780	0.473	1.000	0.946
UNEQFAIR > UNEQDISCR	0.081	0.013	0.102	0.025	0.905	0.720	0.905	0.900
UNEQUNFAIR > UNEQDISCR	0.021	0.006	0.031	0.019	0.886	0.752	1.000	0.752
<i>N</i>	539	533	539	533	543	542	543	542

Note: One-sided *p*-values are presented, in the direction predicted. BH corrections account for multiple hypothesis testing.

Table 2.14: Tobit Regression of Labor Supply on Payment Schemes, for Men and Women

Scheme	Low-Wage Workers			High-Wage Workers		
	Men	Women	Men & Women	Men	Women	Men & Women
	(1)	(2)	(3)	(4)	(5)	(6)
UNEQFAIR	-12.788** (5.950)	3.919 (6.311)	-12.682** (5.909)	4.947 (6.307)	-5.478 (6.674)	3.238 (6.351)
UNEQUNFAIR	-9.366* (5.585)	0.715 (5.938)	-8.896 (5.669)	5.069 (6.368)	-5.213 (6.347)	3.342 (6.273)
UNEQDISCR	-10.802* (5.952)	-15.378*** (5.561)	-11.092* (5.831)	-1.698 (5.640)	7.046 (6.390)	-3.679 (5.649)
UNEQFAIR × Woman			17.242** (8.605)			-8.769 (9.269)
UNEQUNFAIR × Woman			9.878 (8.163)			-9.287 (8.848)
UNEQDISCR × Woman			-5.021 (7.966)			10.430 (8.533)
Woman			-5.400 (6.212)			5.317 (6.854)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Prob > F	0.053	0.000	0.000	0.045	0.000	0.003
Pseudo R ²	0.014	0.019	0.013	0.013	0.010	0.010
N	268	265	542	270	272	542

Note: EQLOW/EQHIG serves as baseline. Standard errors are indicated in parentheses. Two-sided p -values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

Table 2.15: *P*-Values of Six Predicted Differences in Labor Supply between Payment Schemes, for Men and Women

		Men							
		Low-Wage Workers				High-Wage Workers			
Predicted Inequality	Technique	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
BH Correction		No	No	Yes	Yes	No	No	Yes	Yes
EQ _{LOW} /EQ _{HIGH} >	UNEQ _{FAIR}	0.002	0.016	0.012	0.098	0.818	0.783	0.818	0.979
EQ _{LOW} /EQ _{HIGH} >	UNEQ _{UNFAIR}	0.021	0.047	0.041	0.095	0.818	0.787	1.000	0.787
EQ _{LOW} /EQ _{HIGH} >	UNEQ _{DISCR}	0.018	0.035	0.055	0.106	0.111	0.382	0.222	0.764
	UNEQ _{FAIR} >	0.920	0.716	0.920	0.716	0.248	0.508	0.372	0.762
	UNEQ _{UNFAIR} >	0.912	0.630	1.000	0.787	0.065	0.122	0.196	0.731
	UNEQ _{UNFAIR} >	0.238	0.398	0.358	0.598	0.063	0.123	0.380	0.369
	DISCR _{HANCE}								
	<i>N</i>	271	268	271	268	270	270	270	270

		Women							
		Low-Wage Workers				High-Wage Workers			
Predicted Inequality	Technique	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
BH Correction		No	No	Yes	Yes	No	No	Yes	Yes
EQ _{LOW} /EQ _{HIGH} >	UNEQ _{FAIR}	0.767	0.732	0.958	0.732	0.104	0.206	0.621	1.000
EQ _{LOW} /EQ _{HIGH} >	UNEQ _{UNFAIR}	0.768	0.548	0.768	0.685	0.139	0.206	0.418	0.618
EQ _{LOW} /EQ _{HIGH} >	UNEQ _{DISCR}	0.005	0.003	0.009	0.006	0.972	0.864	1.000	1.000
	UNEQ _{FAIR} >	0.751	0.305	1.126	0.457	0.801	0.517	1.000	1.000
	UNEQ _{UNFAIR} >	0.004	0.001	0.012	0.004	0.996	0.977	0.996	1.000
	UNEQ _{UNFAIR} >	0.003	0.001	0.018	0.004	0.994	0.981	1.000	0.981
	UNEQ _{DISCR}								
	<i>N</i>	268	265	268	265	273	272	273	272

Note: One-sided *p*-values are presented, in the direction predicted. BH corrections account for multiple hypothesis testing.

Table 2.16: *P*-Value of Test that Positive Discrimination of Women does not Affect their Labor Supply

Inequality	High-Wage Women	
	(1)	(2)
Technique	Dunn	Tobit
EQ _{HIGH} ≠ UNEQ _{DISCR}	0.055	0.271
UNEQ _{UNFAIR} ≠ UNEQ _{DISCR}	0.012	0.038
<i>N</i>	273	272

Note: Two-sided *p*-values are presented. The Tobit estimates come from specification (5) of Table 2.14.

Table 2.17: Hurdle Model (Labor Supply on Payment Schemes), for Low-Wage Workers

Scheme	Low-Wage Workers
Extensive Margin	
UNEQFAIR	-4.417 (5.927)
UNEQUNFAIR	-4.406 (5.413)
UNEQDISCR	-15.442*** (5.715)
Intensive Margin	
UNEQFAIR	-0.516* (0.302)
UNEQUNFAIR	-0.304 (0.316)
UNEQDISCR	-0.773** (0.307)
Controls	Yes
Prob > F	0.032
Pseudo R^2	0.013
N	533

Note: EQLOW serves as baseline. Standard errors are in parentheses. Two-sided p -values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

Table 2.18: P -Values of Six Predicted Differences in Labor Supply between Payment Schemes, Hurdle Model

Predicted Inequality	Low-Wage Workers			
	(1)	(2)	(3)	(4)
	Extensive		Intensive	
Margin	No	Yes	No	Yes
BH Correction				
EQLOW/EQHIG > UNEQFAIR	0.228	0.285	0.044	0.087
EQLOW/EQHIG > UNEQUNFAIR	0.208	0.312	0.168	0.210
EQLOW/EQHIG > UNEQDISCR	0.003	0.021	0.006	0.035
UNEQFAIR > UNEQUNFAIR	0.501	0.501	0.787	0.787
UNEQFAIR > UNEQDISCR	0.039	0.078	0.138	0.206
UNEQUNFAIR > UNEQDISCR	0.026	0.079	0.029	0.086
N	533	533	533	533

Note: One-sided p -values are presented, in the direction predicted. BH corrections account for multiple hypothesis testing.

Table 2.19: Tobit Regression of Labor Supply on Payment Schemes, for Young (age ≤ 36) and Old (age > 36)

Scheme	Low-Wage Workers		High-Wage Workers	
	Young (1)	Old (2)	Young (3)	Old (4)
UNEQFAIR	-5.633 (6.245)	-4.065 (6.183)	-9.100 (7.093)	5.953 (6.122)
UNEQUNFAIR	-10.527* (6.099)	0.962 (5.712)	-5.550 (6.463)	1.650 (6.260)
UNEQDISCR	-18.440*** (6.241)	-9.274* (5.360)	-0.442 (6.158)	3.300 (5.940)
Controls	Yes	Yes	Yes	Yes
Prob $> F$	0.034	0.088	0.000	0.625
Pseudo R^2	0.013	0.014	0.011	0.007
N	260	273	290	252

Note: Median age is 36 in the sample. EQLOW/EQHIG serves as baseline. Standard errors are in parentheses. Two-sided p -values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

Table 2.20: *P*-Values of Six Predicted Differences in Labor Supply between Payment Schemes, for Young (age ≤ 36) and Old (age > 36)

Young									
Predicted Inequality	Low-Wage Workers				High-Wage Workers				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
	Technique	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit
BH Correction	No	No	Yes	Yes	No	No	Yes	Yes	
EQLow/EQHigh > UNEQFAIR	0.086	0.184	0.108	0.230	0.054	0.100	0.321	0.602	
EQLow/EQHigh > UNEQUNFAIR	0.047	0.043	0.094	0.086	0.106	0.196	0.318	0.587	
EQLow/EQHigh > UNEQDISCR	0.003	0.002	0.018	0.010	0.751	0.471	1.000	0.943	
UNEQFAIR > UNEQUNFAIR	0.187	0.207	0.187	0.207	0.851	0.712	1.000	1.000	
UNEQFAIR > UNEQDISCR	0.036	0.017	0.108	0.052	0.952	0.918	0.952	0.918	
UNEQUNFAIR > DISCRHANCE	0.058	0.088	0.087	0.132	0.901	0.823	1.000	1.000	
<i>N</i>	263	260	263	260	291	290	291	290	

Old									
Predicted Inequality	Low-Wage Workers				High-Wage Workers				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
	Technique	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit
BH Correction	No	No	Yes	Yes	No	No	Yes	Yes	
EQLow/EQHigh > UNEQFAIR	0.032	0.256	0.096	0.384	0.877	0.834	0.877	0.834	
EQLow/EQHigh > UNEQUNFAIR	0.191	0.567	0.287	0.709	0.791	0.604	1.000	1.000	
EQLow/EQHigh > UNEQDISCR	0.032	0.042	0.190	0.127	0.863	0.711	1.000	0.888	
UNEQFAIR > UNEQUNFAIR	0.950	0.792	0.950	0.792	0.162	0.244	0.974	1.000	
UNEQFAIR > UNEQDISCR	0.755	0.186	0.944	0.371	0.229	0.322	0.687	0.966	
UNEQUNFAIR > UNEQDISCR	0.049	0.023	0.099	0.140	0.821	0.607	1.000	0.911	
<i>N</i>	276	273	276	273	252	252	252	252	

Note: One-sided *p*-values are presented, in the direction predicted. BH corrections account for multiple hypothesis testing.

Table 2.21: Tobit Regression of Labor Supply on Payment Schemes, for Employed Full Time and Not Employed Full Time

Scheme	Low-Wage Workers		High-Wage Workers	
	Full-Time (1)	Not Full-Time (2)	Full-Time (3)	Not Full-Time (4)
UNEQFAIR	-7.091 (6.546)	0.891 (5.955)	-1.010 (6.384)	-3.379 (6.724)
UNEQUNFAIR	-7.809 (6.109)	-0.410 (5.661)	2.169 (6.374)	-5.549 (6.327)
UNEQDISCR	-10.922* (6.074)	-15.172*** (5.578)	5.921 (5.872)	-5.965 (6.239)
Controls	Yes	Yes	Yes	Yes
Prob > F	0.301	0.002	0.003	< 0.001
Pseudo R^2	0.009	0.017	0.004	0.017
N	254	279	296	246

Note: Workers who are full-time employed account for 51% of the sample. Those who are not full-time employed are: part-time employed (20%), unemployed (6%), not in a paid job (18%) or other (5%). EQLOW/EQHIG serves as baseline. Standard errors are in parentheses. Two-sided p -values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

Table 2.22: *P*-Values of Six Predicted Differences in Labor Supply between Payment Schemes, for Employed Full Time and Not Employed Full Time

Employed Full Time									
Predicted Inequality	Low-Wage Workers				High-Wage Workers				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
Technique	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	
BH Correction	No	No	Yes	Yes	No	No	Yes	Yes	
EQLOW/EQHIGH > UNEQFAIR	0.033	0.140	0.066	0.280	0.183	0.437	1.000	1.000	
EQLOW/EQHIGH > UNEQUNFAIR	0.029	0.101	0.086	0.304	0.752	0.633	1.000	1.000	
EQLOW/EQHIGH > UNEQDISCR	0.016	0.037	0.094	0.220	0.944	0.843	1.000	1.000	
UNEQFAIR > UNEQUNFAIR	0.244	0.455	0.244	0.455	0.821	0.694	1.000	1.000	
UNEQFAIR > UNEQDISCR	0.183	0.256	0.275	0.384	0.973	0.889	0.973	0.889	
UNEQUNFAIR > UNEQDISCR	0.187	0.280	0.234	0.349	0.947	0.749	1.000	1.000	
<i>N</i>	254	254	254	254	296	296	296	296	

Not Employed Full Time									
Predicted Inequality	Low-Wage Workers				High-Wage Workers				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
Technique	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	
BH Correction	No	No	Yes	Yes	No	No	Yes	Yes	
EQLOW/EQHIGH > UNEQFAIR	0.089	0.559	0.133	0.559	0.249	0.308	0.249	0.616	
EQLOW/EQHIGH > UNEQUNFAIR	0.216	0.471	0.270	0.589	0.221	0.191	0.286	0.572	
EQLOW/EQHIGH > UNEQDISCR	0.004	0.004	0.023	0.021	0.165	0.170	0.990	1.000	
UNEQFAIR > UNEQUNFAIR	0.889	0.415	0.889	0.622	0.222	0.361	0.278	0.452	
UNEQFAIR > UNEQDISCR	0.041	0.005	0.082	0.009	0.167	0.338	0.502	0.507	
UNEQUNFAIR > DISCRHANCE	0.006	0.004	0.018	0.011	0.190	0.471	0.380	0.471	
<i>N</i>	279	279	279	279	246	246	246	246	

Note: One-sided *p*-values are presented, in the direction predicted. BH corrections account for multiple hypothesis testing.

Table 2.23: Tobit Regression of Time Spent in the Experiment on Payment Schemes

Scheme	Low-Wage Workers		High-Wage Workers	
	(1)	(2)	(3)	(4)
UNEQFAIR	-1.478 (1.935)	-0.729 (2.002)		
UNEQUNFAIR	-0.720 (1.803)	-0.062 (1.946)		
UNEQDISCR	-4.621*** (1.760)	2.060 (1.923)		
Controls	Yes	Yes		
Prob > F	0.002	0.004		
Pseudo R^2	0.009	0.008		
N	533	542		

Note: Time spent in the experiment is measured in minutes. Mean times are 22.65 (SD = 14.96) minutes for low-wage workers, and 28.19 (SD = 16.02) minutes for high-wage workers. Note that our time measure is not necessarily time worked since it starts when workers begin reading the instructions and ends when they quit the experiment or at 65 minutes if they have not quit by then. Moreover, this measure does not account for the breaks that workers can take. Therefore, it is most likely not an appropriate measure of labor supply. EQLOW/EQHIG serves as baseline. Standard errors are in parentheses. Two-sided p -values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

Table 2.24: P -Values of Six Predicted Differences in Time Spent in the Experiment between Payment Schemes

Predicted Inequality	Low-Wage Workers				High-Wage Workers			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit	Dunn	Tobit
BH Correction	No	Yes	No	Yes	No	Yes	No	Yes
EQLOW/EQHIG > UNEQFAIR	0.028	0.223	0.057	0.334	0.148	0.358	0.889	1.000
EQLOW/EQHIG > UNEQUNFAIR	0.180	0.345	0.225	0.431	0.247	0.487	0.741	1.000
EQLOW/EQHIG > UNEQDISCR	0.006	0.004	0.037	0.027	0.921	0.858	1.000	1.000
UNEQFAIR > UNEQUNFAIR	0.949	0.657	0.949	0.657	0.852	0.639	1.000	1.000
UNEQFAIR > UNEQDISCR	0.137	0.042	0.205	0.083	0.907	0.929	0.970	0.929
UNEQUNFAIR > UNEQDISCR	0.013	0.010	0.039	0.029	0.926	0.879	1.000	1.000
N	539	533	539	533	543	542	543	542

Note: One-sided p -values are presented, in the direction predicted. BH corrections account for multiple hypothesis testing.

Table 2.25: Labor Supply per Payment Scheme, for Workers not Beating the Odds and Workers Beating the Odds

Low-Wage Workers		Mean	SD	Min.	Max.	<i>N</i>
Workers <i>not</i> Beating the Odds						
	EQLOW	43.20	27.63	.04	.16	128
	UNEQFAIR	37.44	29.16	.06	.17	125
	UNEQUNFAIR	39.22	27.76	.04	.15	143
	UNEQDISCR	33.62	26.41	.08	.11	143
Workers Beating the Odds						
	UNEQUNFAIR(BEATODDS)	37.98	27.58	.02	.16	50
	UNEQDISCR(BEATODDS)	38.15	32.10	.04	.24	46
High-Wage Workers		Mean	SD	Min.	Max.	<i>N</i>
Workers <i>not</i> Beating the Odds						
	EQHIGH	44.03	28.93	.04	.20	128
	UNEQFAIR	43.23	29.64	.03	.20	127
	UNEQUNFAIR	43.50	28.85	.02	.20	143
	UNEQDISCR	45.74	26.72	.02	.16	145
Workers Beating the Odds						
	UNEQUNFAIR(BEATODDS)	42.55	28.34	.02	.17	47
	UNEQDISCR(BEATODDS)	40.72	27.79	.02	.15	46

Note: Labor supply is measured by the number of lines completed and ranges from 0 to 85. Min. and Max. indicate the percentage of workers completing the minimum and maximum number of lines. Low-wage workers beating the odds had high chances, and high-wage workers beating the odds had low chances.

Table 2.26: *P*-Values of Tests on Difference in Labor Supply between Payment Schemes, for Workers Beating the Odds

Predicted Inequality	Low-Wage Workers		High-Wage Workers	
	(1)	(2)	(3)	(4)
Technique	Dunn	Tobit	Dunn	Tobit
UNEQFAIR > UNEQUNFAIR(BEATODDS)	0.181	0.316	0.241	0.454
UNEQUNFAIR(BEATODDS) > UNEQDISCR(BEATODDS)	0.811	0.658	0.192	0.279
<i>N</i>	221	219	220	219

Notes: One-sided *p*-values are presented, in the direction predicted by the model.

2.C. Instructions of the Experiment

2.C.A. Invitation Email

[Notes: Potential participants see the time limit that they have to complete the experiment]

We would like to invite you to participate in an online economic experiment about decision making. You will be paid a reward of £0.70 for about 5 minutes of participation. Thereafter, as will be explained in the instructions, you can earn more money with the decisions you make by participating in this experiment for a longer time.

IMPORTANT: All information provided will be collected and stored ANONYMOUSLY.

You receive this invitation because you are registered at Prolific. Please consult the Prolific website in case you want your data to be removed from the platform.

2.C.B. Instructions

[Notes: Participants are provided with a countdown from 65 to 0 minutes.]

Instructions

Welcome to this economic experiment,

You can earn a considerable amount of money with the decisions you make. Please read these instructions carefully. Importantly, unlike experiments in some other social sciences, economic experiments employ a strict non-deception policy. This means that all information you receive is truthful.

The only way to leave this economic experiment and be paid is to click on the button “Leave” and go to the next page. Once you do this, you will see a message that the experiment is now over and that you can close your browser page. You will not be paid if you leave at any moment by closing your browser window without clicking on the button “Leave” and going to the next page that tells you that the experiment is over.

This economic experiment consists of a Practice Part, where you cannot yet earn money, followed by a Task Part where you can earn money. The Practice Part consists of these Instructions, some comprehension questions, and a practice exercise. It is important that you answer the comprehension questions correctly by yourself. Please do not consult other people when answering these questions. In case you do not answer a question correctly, you will have two more chances to correct your answer. If you do not answer all questions correctly after these two additional chances, you will not be able to participate in the Task Part and the experiment ends for you. In that case you will be paid £0.45. When you have answered all comprehension questions correctly you can participate in the Task Part. In the Task Part, you can earn money by working on a task. You can stop working on the task whenever you prefer.

Recall, that to leave this economic experiment and to **be paid** you need to click on the “Leave button” and go to the next page. Once you have done this, you will see a message that the experiment is over for you and that you can close your browser page. You will **not be paid** if you leave the experiment without following the described procedure.

Note that you cannot leave the experiment and be paid before you finish the Practice Part (which lasts about 5 minutes). Thereafter you can leave the experiment at any time.

After you leave the experiment using the “Leave” button, the money you have earned will be paid to you through Prolific.

Task Part

In the Task Part of this experiment, you can earn money by working on a task. You can decide how much of the task you want to complete. The task is to enter preset lines of random numbers and/or letters on your computer. You will receive a payment for each line you copy correctly. Nobody else than yourself will derive any earnings from your work, including the experimenters. The lines of numbers and/or letters you enter have no further use for anyone.

You will see one line at the time. Once you have entered a line correctly, you can go to the next page to see the next line. Each time you see a new line, you can decide whether you would like to type this line or leave the experiment.

In case you make a mistake when entering the line, the software will tell you so. You will need to correct this mistake before you can proceed to the next line.

The length of the sequences of random numbers and/or letters will increase as you complete more lines. That is, lines will be relatively short at the beginning but get longer over time.

You will be informed of your payment per line at the beginning of the Task Part.

In the Task Part you may also receive anonymous payment information regarding another participant.

Leave the Experiment

You can stop entering lines at any moment. Note, however, that the only way to stop and to be paid is to click on the “Leave” button and then go to the next page. You will then see a message that the experiment is over for you, that you need to click on a Prolific link to validate your participation, and that you can then close your browser window.

You will see the “Leave” button whenever you are presented a new line. If you decide to leave, you will not be able to start working again. That is, **once you leave the experiment you cannot go back.**

Payment

When you leave the experiment according to the described procedure you will receive a **payment per line you entered correctly.** You will be informed about the amount of the payment per line when you see the first line to be entered. In addition, you will also receive a **fixed amount of £0.70,** irrespective of the number of lines entered.

Decision

The decision you make in this economic experiment is to choose how much of the task you want to complete. You are the only one deciding how much you work. Your decision only affects your own earnings.

Practice Part

In the Practice Part, we ask you to correctly answer a number of comprehension questions. It is important that you answer these comprehension questions by yourself. For each question you will have three chances. If you do not correctly answer all comprehension questions you will not be able to participate in the Task Part. In this case the experiment will end for you and you will be paid £0.45.

After having correctly answered all comprehension questions, you will enter two practice lines to make you familiar with the task. Neither can you earn payments per line with these practice lines, nor will these practice lines affect the Task Part in any way.

This is the end of the instructions.

I confirm that I have read the instructions carefully and I am ready to start the Practice Part. I will not be able to go back to the instructions once I go to the next page.

Chapter 3

On Your Own Side of the Fence

Abstract

This paper explores the causal mechanism behind the literature's finding that living close to poor individuals of another ethnicity decreases support for income redistribution. We employ a laboratory experiment to extract causal relationships for some of the plausible forces at play. Highly-paid participants ("the rich") are randomly assigned to interact with poorly-paid participants ("the poor") of different ethnicities or with other highly-paid participants. They then propose redistribution schemes affecting a larger group of rich and poor. In contrast to previous studies, the results show that proximity with the poor increases redistribution—regardless of the poor's ethnicity. Moreover, our design allows us to examine the counterfactual possibility that participants avoid certain types of interactions if they are not exogenously enforced. We report that, given the choice, Caucasians tend to avoid interactions with minorities. This suggests that inter-ethnic proximity with the poor does not directly affect preferences for redistribution; rather, compared to intra-ethnic proximity, it can decrease contact with the poor nearby, which in turn reduces redistribution.

I. Introduction

*They should rise, yes, but on their own side of the fence. —Mrs. Chestny, Everything That Rises Must Converge (O'Connor, 1965)*¹

Economists and political scientists alike have long recognized that ethnic differences feature prominently in redistributive policy decisions. Woodward (1955) documents that Reconstruction-era Democrats successfully framed redistribution proposals as transfers from Caucasians to African-Americans. Both Becker (1957) and Orr (1976) formalize the idea that dislike of minority welfare recipients can lower redistribution levels. More recently, Alesina et al. (2001) and Alesina and Glaeser (2004) argue that dislike for minorities shape welfare differences between the ethnically heterogeneous U.S. and relatively homogeneous European countries.² Among the important contributions to this literature are Luttmer (2001) and subsequently Dahlberg et al. (2012). Both studies provide evidence that inter-ethnic proximity with the poor is a specific channel by which ethnic differences decrease the support for redistribution. The causal support, however, is restricted and it is unclear how proximity might actually alter support for redistribution. This paper delve into the causal mechanism behind this effect at a fine-grained level—using the controlled environment provided by a laboratory experiment.

Using U.S. data, Luttmer (2001) shows that Caucasians (African-Americans) living near African-American (Caucasian) welfare recipients exhibit lower support for welfare measures, whereas those living close to Caucasian (African-American) welfare recipients have no such tendency, after controlling for correlates of proximity. Dahlberg et al. (2012) offer causal support for this first result in Sweden: cities receiving more (quasi-randomly-allocated) immigrants from outside of the OECD experience a decrease in support for redistribution, especially among their richer inhabitants. These findings are striking and con-

¹While this highlights animosity in contact, another fictional character, Harry Lime, played by Orson Welles in *The Third Man* (Reed et al., 1949), illustrates instead that remoteness might inhibit social conscience: *I never feel comfortable [meeting them, the] victims? Look down there. Tell me. Would you really feel any pity if one of those dots stopped moving forever?*

²There is a large literature connecting ethnic divisions and lower redistribution levels as well as lower provision of public services with a redistributive component, among which are Gilens (1995, 1996); Easterly and Levine (1997); Poterba (1997); Goldin and Katz (1999); Alesina et al. (1999); Roemer and Van der Straeten (2005); Miguel and Gugerty (2005); Glaeser (2005); Lee and Roemer (2006); Roemer and Van der Straeten (2006); Lind (2007); Roemer et al. (2007); Shayo (2009); Baldwin and Huber (2010); Bandiera and Levy (2011); Swee (2015); Kuziemko and Washington (2018). See the overviews by Alesina and La Ferrara (2005) and Stichnoth and Van der Straeten (2013).

sistent with the proposition in sociology that immediacy between ethnic groups competing for status, economic or political power worsens inter-group dislike (Blumer, 1958).³

However, the literature has yet to establish causality for other ethnic groups than relatively recent immigrants in Sweden. At least two streams of literature that have examined inter-group interactions at more micro levels also stand to some extent at odds with the results. The first is Contact Theory (Allport, 1954), a prominent social psychology concept that has been influential in American court rulings regarding desegregation and Affirmative Actions (e.g., *Brown v. Board of Education* in 1954, *Fisher v. University of Texas* in 2016). It contends that, at least under favorable conditions, inter-group contact should *improve* inter-group attitudes.⁴ Although studies in this area do not investigate the effect of proximity with the poor of another ethnicity on redistributive decisions, the theory might predict a positive effect. The second is the Identifiable Victim Effect (Schelling, 1968), which states that identifying the person affected by one's choices leads to *heighten* care for this person, although this literature does not explore the possible role of ethnic differences.⁵

Moreover, it is unclear how exactly proximity might influence support for redistribution as the process by which it happens is still somewhat of a black box. An obvious explanation is that proximity with poor minorities makes salient that welfare recipients are likely to be from another ethnic group and that this triggers intra-group self-interest. Nonetheless, other explanations are plausible. An increase in heterogeneity can engender changes in neighborhood dynamics. For instance, it could decrease interactions with the poor of a different ethnicity and lower social capital (as suggested, among others, by Alesina and La Ferrara, 2005; Putnam, 2007; and Algan et al., 2016), which may in turn affect

³Also see Bobo and Hutchings (1996). For a discussion of this literature, see Bobo (1999).

⁴Theoretical contributions to Contact Theory include Pettigrew (1998) and Brown and Hewstone (2005). For a recent overview of the evidence, see Pettigrew and Tropp (2006) and chapter 9 of Brown (2011). The theory is applied to several types of groups, including ethnic and income groups. It has recently attracted interests from economists, who report positive effects of contact that is much longer in nature than the interactions we study here (e.g., Boisjoly et al., 2006; Camargo et al., 2010; Burns et al., 2015; Carrell et al., 2015; Finseraas et al., 2016; and Rao, 2019). Notably, Rao (2019) finds a positive effect of contact with the poor on benevolence toward them.

⁵Bjorvatn and Cappelen (2003) and Lupu and Pontusson (2011) model this intuition at the societal level, proposing that increased distance between groups decreases support for redistribution by the rich. A number of studies in experimental economics and psychology have found a positive effect on generosity of divulging information about potential recipients or allowing short-term interactions between donors and recipients (e.g., Bohnet and Frey, 1999; Batson et al., 2002; Small and Loewenstein, 2003; Charness and Gneezy, 2008; and Andreoni and Rao, 2011).

support for redistribution.⁶ These changing features are difficult to account for in uncontrolled settings and it is therefore at least as difficult to identify their potential effects on support for redistributive policies.

In the present laboratory experiment, we study the effect of intra-ethnic and inter-ethnic proximity with the poor on the rich's income redistribution decisions. In our setting, proximity generates social interactions and redistribution decisions are monetarily incentivized.⁷ The experiment is conducted at a large public university in the American Southwest, which readily provides the required ethnic diversity. Participants are assigned fixed unequal wages for completing an individual task. Proximity of highly-paid participants ("the rich") with poorly-paid participants ("the poor") as well as ethnic differences between the rich and the poor are manipulated.⁸ After interacting with another individual, a participant proposes a redistribution scheme affecting a larger group of 10 participants. This larger group consists of all individuals taking part in an experimental session, including the participant, and has a fixed total income. We measure the within-subject effects of contact with poor Caucasians and poor minorities on the redistribution scheme proposed by the rich, compared to interacting with another rich. In addition, using a methodology that prevents statistical discrimination based on ability, we examine the avoidance of interactions with two types of participants: those of a different ethnicity, and those who are assigned to be poor.⁹

Four factors play in favor of a laboratory experiment as a methodology of choice here. First, it allows us to investigate the causality in the link between proximity and support for redistribution by exogenously varying intra- and inter-ethnic proximity. Second, it allows us to measure redistribution decisions. Third, we can ensure that proximity translates into interactions. Fourth, it also enables us to examine the counterfactual possibility that, given the opportunity

⁶For example, Alesina and La Ferrara (2005) find that U.S. residents of more unequal and more ethnic- or income-diverse Metropolitan Statistical Areas participate less in social and community activities. At a smaller level, Algan et al. (2016) show that almost-randomly-allocated residents of nationality-diverse social housing blocks in France report to interact less with their neighbors, which in turn appears to reduce their ability to maintain the local public good that is housing quality.

⁷Distributional choices of participants in economic experiments have been shown to correlate with behavior outside of those experiments, such as loan repayments (Karlan, 2005), donations and other pro-social behaviors (Benz and Meier, 2008; Baran et al., 2010; Franzen and Pointner, 2013), voting decisions (Finan and Schechter, 2012; Fisman et al., 2017), work productivity (Cohn et al., 2015), and socially-responsible investments (Riedl and Smeets, 2017).

⁸The "rich" and "poor" wording is used for convenience.

⁹In this respect, an experimental design related to ours is Hedegaard and Tyran (2018), which measures the extent to which individuals with Muslim-sounding names are avoided in a Danish workplace.

to initiate contact, individuals actually avoid certain types of interactions.¹⁰

The main findings are as follows. First, both intra-ethnic and inter-ethnic contact with the poor increase the redistribution levels proposed by the rich. The average increase corresponds to approximately 1/4 of a standard deviation. We find no statistical differences between the effects produced by intra-ethnic and inter-ethnic contact. Second, when facing the choice to interact with either Caucasian or minority participants, Caucasian participants tend to favor interactions with other Caucasians. That is, they select on average approximately 25% more Caucasians for their preferred interactions. Minority participants do not exhibit significant ethnic inclination regarding whom they interact with. Third, the majority of the rich avoid contact with the poor. In fact, a mere 5% of the rich elect to interact with the poor, and 29% are indifferent to interacting with fellow rich or with the poor.

Our first result contrasts with the negative inter-ethnic effect of proximity reported in Luttmer (2001) and Dahlberg et al. (2012). However, our second finding suggests a possible rationale for this difference. Our result regarding proximity arises after interactions in the experiment are enforced, whereas, when given the choice, we find that Caucasians tend to avoid inter-ethnic interactions. Consequently, inter-ethnic proximity with the poor can affect support for redistribution indirectly, by decreasing contact with the poor nearby. Specifically, relative to proximity with the poor of one's own ethnicity, the decreased interactions with the poor stemming from proximity with the poor of another ethnicity reduce the rich's preferred levels of redistribution. As such, our experiment offers a new and distinct interpretation of the negative effect of inter-ethnic proximity with the poor on redistribution.

Our third result, that the rich avoid the poor, is consistent with a recent literature regarding moral behavior and social preferences, which argues that many individuals avoid to place themselves in situations where they become generous (e.g., Dana et al., 2006; DellaVigna et al., 2012 and Andreoni et al., 2017). For example, individuals are more generous when a person engages them and asks them to donate for a cause, but often actively avoid such a person. Based on our data, preserving some distance with the poor is a successful strategy for the rich who prefer to be selfish, but somehow find it difficult to control their generosity (or to justify their selfishness) once they face the poor.

¹⁰Moreover, the negative effect observed in Luttmer (2001) and Dahlberg et al. (2012) is at least as strong for the rich as for the poor in the first and stronger for the rich than for the poor in the second. As a university degree is a common characteristic of high-income earners, the students at a large (non-elite) public university in this study should be somewhat representative of high-income earners who graduated in the near past.

We recognize that this study does not capture all channels that might drive the effects of intra-ethnic and inter-ethnic contacts with poorer neighbors on redistribution. Rather, the goal of this experiment is to cleanly examine some of the plausible forces at play that are difficult to assess with other methodologies. In doing so, it also joins a variety of experiments crafted to evaluate ethnic discrimination in other settings (e.g., Fershtman and Gneezy, 2001; List, 2004; Bertrand and Mullainathan, 2004; Castillo and Petrie, 2010; Hedegaard and Tyran, 2018; Cettolin and Suetens, 2019).

The paper proceeds as follows. Section II details the experimental design. Section III provides the hypotheses regarding the effect of contact. Section IV tests our hypotheses and analyses choices regarding inter-ethnic and inter-income interactions. Section V briefly discusses possible ramifications of our findings.

II. Experimental Design

The experimental design recreates in a highly-controlled fashion the situation of neighborhood proximity and adds a redistribution decision. There are three identical rounds in the experiment, only one of which is randomly chosen to be paid. At the start of the session, each participant is assigned one wage to be paid for completing an individual task during a round, and this wage stays unchanged across rounds. The task consists in copying lines of characters. Moreover, in each round, before they complete the individual task, two participants are assigned to interact with each other in a pair. There are 10 participants in a session, which allows us to vary the income and ethnic composition of the pairs across rounds, while keeping the total income within a session fixed. After completing the task, each participant proposes a redistribution scheme affecting the distribution of earnings among the 10 participants.

To examine preferences for interacting with others with a different ethnicity or a different income, we add another stage at the start of every round. In this stage, participants select which other participants they prefer to interact with. To create incentives, there is a small probability that their choice is implemented.

At the very beginning of the session, participants are given a written copy of the entire instructions. They then already read the first section, which provides a comprehensive overview of the experiment. The instructions are available in the Appendix.

II.A. Wage and Ethnicity

We use real ethnicity, which we simplify by dividing participants into two overarching groups: Caucasians and minorities. We invite approximately the same number of participants from both groups in every session. This ethnic composition is natural for participants, as it resembles the composition of the university's student body. At the start of the session, six rich participants are privately assigned the high wage (20.50 USD), and four poor participants the low wage (7 USD).¹¹ Wages are provided for completing the individual task described in Section II.C in a round. A participant's wage stays the same for each of the three rounds. Participants know the distribution of wages, but not which participant has which wage. No reason is provided for how wages are distributed. Unknown to the participants, wages are randomized separately for Caucasians and for minorities, i.e., the wages of individuals inside each group are drawn from a distribution of three high wages and two low wages.¹²

II.B. Interaction

For the duration of a round, participants are matched in pairs in areas that are physically separated by partitions. The members of a pair are seated next to each other, at individual desks that are roughly one meter apart. They can see each other and be seen by the experimenter, but they cannot observe participants in other areas. Participants are re-matched at the start of the next round and cannot be paired again with the same participant. Participants are not told how the matching is made. We generally paired a rich participant with the following three neighbors throughout the session, in random order: (i) a poor Caucasian, (ii) a poor minority, (iii) a rich participant of any ethnicity. These

¹¹The terms "low-wage participant" and "high-wage participant" are employed to designate participants during the experiment.

¹²There are obviously different ethnicities among minority participants. However, it would have been impossible to recruit enough participants of a specific group to conduct this study. We catalogue six categories within minorities; the resulting pool is diverse: "Caucasians/Whites" account for 53%, "Asians/Asian-Americans" for 23%, "Latino(a)s/Hispanics" for 14%, "Blacks/African-Americans" for 8%, and "Members of one of the Native Nations of North America", "Mixed", and "Others" for 3%. The actual ratios are more often 4:6, 6:4, 3:7 or 7:3 than 5:5 due to no-shows. We also tried to balance ethnic groups by gender in the invitations. Furthermore, upon the arrival of participants, the experimenter had to classify them into Caucasians or minority, because the information regarding the ethnicity of invited students could not be accessed at that moment. Classification was based on names as well as visual appearance. Participants also reported their ethnicity at the end of the session. In the few cases in which there was a difference between the two classifications, self-reported ethnicity was then used in the data analysis. When the Caucasians to minorities ratio differed from 5:5, then the extra Caucasian (minority) participants received wages drawn from the minority (Caucasian) participants' distribution.

Table 3.1: Treatments—Types of Contact with the Poor by the Rich

Treatment	Description
NO CONTACT	Rich interacts with another rich of any ethnicity
INTRA-ETHNIC CONTACT	Rich Caucasian (minority) interacts with a poor Caucasian (minority)
INTER-ETHNIC CONTACT	Rich Caucasian (minority) interacts with a poor minority (Caucasian)

Note: Treatments indicate the type of contact with the poor made by the rich.

pairs correspond to our three treatments, which indicate the type of contact with the poor faced by the rich in a round and which we detail in Table 3.1: NO CONTACT, INTRA-ETHNIC CONTACT, and INTER-ETHNIC CONTACT. The treatments are implemented within subject in order to increase statistical power.¹³

The interaction is short and consists in the disclosure of personal information, which participants had consented before the experiment to communicate to others. Participants in a pair disclose the following to one another in pre-determined turns: name, age, academic major, and year of study. They also reveal their own wage to each other, by showing a card with their wage written on it.

II.C. Task

After the interaction, each participant individually completes the task for his or her own wage. We emphasized from the start of the experiment that the completion of the task is strictly individual; there are no possibilities of cooperation or competition. The task is to copy by hand a set of lines from a text in Latin. The lines are the same for everyone, but the rich copy the full set of lines, while the poor copy two-thirds of them. The rich work more in order to make the situation reflects the common belief that effort plays at least some role in the higher income of the rich relative to the poor. This feature is constant across treatments.¹⁴ Participants earn their assigned wage only if they complete the

¹³Our design requires five Caucasian and five minority participants per session for these interactions to occur for all rich. Because of no-shows, some rich interacted twice with a Caucasian or with a minority poor. We note that our within-subject design might generate inertia in the participants' choices, which would prevent us from observing differences between treatments. Nevertheless, we do obtain within-subject variation in redistribution choices. Moreover, the point estimates of the effects of intra- and inter-ethnic contacts on redistribution decisions (Section IV.B) would be similar if we were instead to analyze only the first round to simulate a between-subject design.

¹⁴We provide the task in the Appendix. Participants are given the sheet with the entire set of lines, on which the part for the poor to complete is indicated. The poor cannot copy a greater part of the text to earn more. The experimenter inspects the work of each participant in the first minutes to confirm that they do it correctly, and inspects it again at the end of the time limit to ensure that the task has been properly completed.

task within a given time. The task was piloted to be easy enough for participants to complete with reasonable effort (and every participants did complete the task in every round on time).

II.D. Redistribution Decision

At the end of the round each participant individually and anonymously chooses the income redistribution scheme to be implemented for all 10 participants in the session. This was known to all participants from the beginning of the session. The redistribution schemes are presented in the form of redistribution rates corresponding to earning distributions for the six rich and for the four poor of the session. Table 3.2 shows the distribution resulting from each rate. It is made clear that, irrespective of the rate chosen, there are minimal deviations in the total earnings of the 10 participants in the session—i.e., income redistribution does not decrease efficiency.¹⁵

Participants can select any redistribution rates from -10 to 115 , in increments of 5 . A higher rate increases the earnings of the poor; a lower rate increases those of the rich. Note that the rate of 115 makes the poor richer than the rich, 100 equalizes earnings, 0 leaves initial wages unchanged, and -10 makes the rich even richer. The rates correspond to the percentage of each participant's income that is taken from them, and then redistributed equally among all participants. As such, it replicates a feature of many redistributive public policies: services are often paid by a tax corresponding to a percentage of income, and offered more or less equally to members of a society, e.g., education, healthcare, and public transport. For convenience, our income tax rate is flat.

A single redistribution rate from one of the three rounds is randomly chosen to be implemented for the session. Participants receive no feedback between rounds.

II.E. Choosing Inter-Ethnic and Inter-Income Interactions

In order to observe decisions to engage in inter-ethnic and inter-income interactions, we insert an additional stage at the start of every round. In order to do that, at the beginning of the experimental session, participants are given large

¹⁵Minimal deviations occur because we round earnings to the nearest quarter. The exact table presented to participants is part of the instructions in the Appendix. The choice to round earnings to the nearest quarter was motivated by the (repeated, but anecdotal) observation that students are unhappy with small coin denominations and can even return them. The redistribution decision is entered on a computer. It was computerized using z-Tree software (Fischbacher, 2007).

Table 3.2: Redistribution Rates and Resulting Distributions of Earnings

Redistribution Rate	Earnings of the Poor	Earnings of the Rich	Redistribution Rate	Earnings of the Poor	Earnings of the Rich
-10	6.25	21.00	55	11.50	17.50
-5	6.50	20.75	60	11.75	17.25
0	7.00	20.50	65	12.25	17.00
5	7.50	20.25	70	12.75	16.75
10	7.75	20.00	75	13.00	16.50
15	8.25	19.75	80	13.50	16.25
20	8.50	19.50	85	14.00	16.00
25	9.00	19.25	90	14.25	15.75
30	9.50	19.00	95	14.75	15.25
35	9.75	18.50	100	15.00	15.00
40	10.25	18.25	105	15.50	14.75
45	10.75	18.00	110	16.00	14.50
50	11.00	17.75	115	16.25	14.25

Note: Earnings are in USD. Each participant chooses a redistribution rate. The experiment always comprises four poor and six rich in a session. The rates change the distribution of earnings between rich and poor, but the total earnings of the 10 participants is virtually constant.

identification numbers that they wear visibly for the session.¹⁶ At the start of the round, the 10 participants are seated around a large table. They clearly see all other participants and their identification numbers. There, each participant individually chooses other participants to be paired and interact with for the round. They write down their choices on a form behind a ballot-style box that prevents any other participants from observing their answers. They make two types of selection choices.

The first type of selection choices, which we use to measure preferences for inter-ethnic interactions, proceeds as follows. Participants rank all other participants from the one they most prefer to interact with to the one they least prefer to interact with. The ranking is made using the identifying numbers. The ethnicity of the participant indicated in each rank is then employed to examine the ethnicity of those selected in the most and least preferred ranks. If this method is chosen to count, the pairs are made by randomly drawing a first participant and creating the pair with his or her favorite participant with whom

¹⁶The identifying numbers, ranging from 1 to 10, are given in the following manner. The group of Caucasians and the group of minorities each randomly receives a set of five consecutive numbers. For instance, in a session Caucasians receive numbers 3, 4, 5, 6, 7, and minorities receives 8, 9, 10, 1, 2. In another session, Caucasians receive 9, 10, 1, 2, 3, and minorities receive 4, 5, 6, 7, 8. Within each ethnic group, numbers are handed to participants on their arrival, and were not perfectly randomized. Moreover, as discussed before, since most sessions did not have five Caucasians and five minorities, some Caucasians were given numbers intended for minorities, and vice-versa, depending on no-shows. As a robustness check in Section-IV.C, where we analyze decisions to engage in inter-ethnic interactions, we provide evidence that the identification numbers themselves have no impact on the results.

he or she was not paired in a previous round. Then, a second participant is drawn and paired with his or her favorite available participant whom he or she was not paired with in a previous round, and so on. Ranking all other participants in the session is therefore only incentive compatible in the first round.¹⁷

The second type of selection choices is used to measure preferences for contact with the poor. Participants choose one of the three available options: (1) CONTACT: to be paired and interact with a poor participant, (2) NO CONTACT: to be paired and interact with a rich (instead of a poor) participant, or (3) INDIFFERENT: to be randomly paired and interact with a poor or rich participant.¹⁸ If this method is selected to count, participants are randomly drawn one by one and their choices are implemented as long as those choices are available.

Each round, a random draw is made and there is a 2% probability that one of the two types of choices is used to form the pairs. There is an equal probability that each type is selected. The experimenter does not reveal whether the selection choices have been used to create the pairs. While this technique introduces a small amount of endogenous interactions in the experiment, it leaves almost all interactions exogenous.¹⁹

At the moment of making selection choices, participants know how the rest of the experiment will proceed. Importantly, it has been stressed to them that they will never cooperate or compete, which is necessary to avoid statistical discrimination based on beliefs about ability in the task. Participants make selection choices in each of the three rounds, but we only analyze those from the first round. This is because: (1) those choices are untainted by experience, and (2) for decisions regarding inter-ethnic interactions, those choices are the only incentive-compatible ones, as previously explained.

¹⁷For instance, in the second round, participants can write down that their most preferred participant is the one they were paired with during the first round. This choice, however, cannot be implemented. We also note that, by mistake, two participants were matched together twice in a session where choices were implemented.

¹⁸The options are presented to participants as: paired with a low-wage participant, paired with a high-wage participant, and indifferent. Since each session contains six rich and four poor, choosing the indifferent option also gives one higher chances to be paired with a rich. We still treat the option as real indifference. Counting indifference choices as avoidance of the poor would only reinforce our result regarding inter-income interactions in Section IV.C.

¹⁹Since there is only a small probability that their selection choices are implemented, participants might not take those choices seriously and take the simplest route by randomizing their answers. In our design, random choices would introduce a bias against finding any selection preferences in terms of ethnicity or wage. Furthermore, the small amount of endogenous selection does not affect the result regarding the effect of contact (see footnote 24 in Section IV.B).

II.F. Friendships

The structure of our invitations for the experiment minimizes, but does not eliminate, the possibility of having friends in a session. To begin with, we invited potential participants from a large pool of approximately 1,000 students registered to be invited for economic experiments, across two separate academic years.²⁰ Then, we invited a subset of students in the pool to register for very few sessions at a time (generally one or two sessions), and waited for the session(s) to be full before sending out invitations for new sessions. Most of the time, registration was completed quickly, limiting possible coordination on a session. The first sessions were typically full within a few minutes, the middle ones within a few hours or so, and the last ones within a day or two.

Nevertheless, ignoring that some might register with a friend can bias our results of the effect of contact with the poor on redistribution. In our design, one's friend is more likely to be rich than to be poor—there are six rich and four poor inside a session. If a participant is assigned to be rich and has one friend in the session, the friend has a 5/9 probability to also be rich. The times that a participant meet his or her friend in one of the two interactions with the poor reveal that the friend is poor with certainty. Then, contact with the poor friend can lead one to support more redistribution.²¹ In contrast, not meeting one's friend in the interaction with the poor slightly increases the probability that one's friend is rich, which might then reduce one's preferred levels of redistribution. Similarly, by the same argument, meeting one's friend in the only encounter with the rich can decrease support for redistribution, and not meeting one's friend in the interaction with the rich increase it.

Hence, whether an aggregate bias arises in the effect of contact on redistribution, and whether it exhibits a positive or negative sign depends on participants' relative reactions to those possible events and their ability to correctly update their beliefs. For this reason, we first collect data on whether participants had a friend in the session.²² A non-negligible minority of rich participants (26%) report to have at least one friend in the session—almost always one friend.²³ We leave out those participants with a friend in our analysis

²⁰The pool consists of students from the University of Arizona, which has 44,000 students. A majority comes from the faculty of business (which has around 6,000 students), but some also come from other faculties such as engineering, sciences, and social sciences. The pool changes from year to year, as graduating students leave the pool and new students enter.

²¹Goeree et al. (2010) show that individuals in economic experiments are indeed more generous with friends.

²²Self-reported friendships are commonly employed in economics (e.g., Leider et al., 2009; Branas-Garza et al., 2010; Goeree et al., 2010; Camargo et al., 2010).

²³Note that, in one session, we resorted to use Qualtrics software instead of z-Tree to ask end-

of contact. The evidence from our data is that, on average, including those participants with a friend creates a negative bias on the effect of both intra- and inter-ethnic contacts. That is, the point estimates of both effects would decrease.

Second, we also collected data on whether one was paired with a friend for each round. This allows us to report the bias for those who interacted with a friend when meeting a poor participant separately from those who interacted with a poor participant who is a stranger. Although the aggregate bias on the effect of contact is negative, the few rich who interacted with a poor friend indeed exhibit a positive bias. Moreover, no rich met a friend when interacting with another rich. We return to this issue further in footnote 26 of Section IV.B.

In Section IV.C, we explain how we account for possible biases created by the presence of a friend when we analyze decisions to engage in inter-ethnic interactions.

II.G. Additional Procedural Details

We conducted the experiment at the Economic Science Laboratory of the University of Arizona, with 240 participants in 24 sessions, in 2016 and 2017. Due to a crash of the recruiter software, 17 individuals participated twice, but we identified them afterwards, and removed their second participation from the data. This leaves us with 223 participants. Selection choices were implemented in one round of two sessions. Participants were all students at the university. They were paid their after-redistribution wages as well as a 3 USD show-up fee. The resulting mean earning was 19.30 USD. The sessions lasted for 90 minutes. Participants read a comprehensive overview of the experiment at the very beginning of the session. Then, for the first round, they also read a short detailed section for every part of the experiment right before it occurred. The experimenter read aloud each section with the participants, and encouraged them to ask questions.

of-session questions, and the software did not let participants record that they had no friends in the session. In response, almost all 10 participants entered the highest possible number, which was higher than the maximum number of friends possible in a session. We treated those participants as having zero friends. Moreover, in one session, one participant reported to have 10 friends in the session, which we took as a mistake since it is impossible, and corrected the number to one.

III. Hypotheses

The first hypothesis concerns the effect of intra-ethnic contact with the poor on the redistribution decisions of the rich. As discussed in the introduction, experiments related to the Identifiable Victim Effect have shown that identification of a recipient increases benevolence of potential donors, although without considering whether ethnic similarity plays a role. On a larger scale, Luttmer (2001) reports that, in the U.S., intra-ethnic proximity with the poor in the United States is not significantly correlated with support for redistribution. We still propose the following hypothesis.

HYPOTHESIS 1: Intra-ethnic contact with the poor increases the rich's support for redistribution, relative to interacting with other rich.

The second hypothesis pertains to the effect of inter-ethnic contact with the poor on the redistribution decisions of the rich. Luttmer (2001) and Dahlberg et al. (2012) provide empirical evidence that proximity with the poor from another ethnic group is correlated with lower support for redistribution in the U.S., and causally decreases support for redistribution in Sweden, respectively. We test the following hypothesis in our own causal setting with Americans.

HYPOTHESIS 2: Inter-ethnic contact with the poor decreases the rich's support for redistribution, relative to interacting with other rich.

In addition, in exploratory analyses, we examine patterns of decisions to engage in inter-ethnic and in inter-income interactions. Specifically, we evaluate whether Caucasian and minority participants are more likely to choose to interact with Caucasian or with minority participants. Similarly, we analyze whether the rich and the poor are more likely to select interactions with the rich or with the poor.

IV. Results

IV.A. Overview Redistribution Decisions

Table 3.3 provides the means and standard deviations of redistribution rates proposed by the rich, separated by the type of contact they experience. In addition, it divides data for Caucasian and minority participants. We have a total of

Table 3.3: Summary of Redistribution Rates Chosen by the Rich

Treatments	Redistribution Rates								
	All			Caucasians			Minorities		
	Mean	SD	N	Mean	SD	N	Mean	SD	N
NO CONTACT	7.23	24.01	103	5.20	22.69	51	9.23	25.29	52
INTRA-ETHNIC CONTACT	11.60	31.43	100	8.13	27.82	48	14.81	34.38	52
INTER-ETHNIC CONTACT	16.22	34.11	94	13.06	31.57	54	20.50	37.26	40

Note: Treatments indicate the type of contact with the poor. We show mean, standard deviation and number of participants per treatment. Redistribution rates range from -10 to 115 . There are 99 rich participants (51 Caucasians and 48 minorities). A few participants encountered the same treatment twice, so that they did not experience all three treatments.

99 rich participants (51 Caucasians and 48 minorities).²⁴ Contact with the poor appears to increase redistribution, for Caucasians as well as for minorities. Crucial to our investigation, participants do not respond negatively to either types of contact. Although not shown, there is also substantial lower bound censoring in the redistribution rates chosen: 45% of the rich's decisions correspond to the minimum rate of -10 , which maximizes their own earnings.

IV.B. Testing the Hypotheses

To analyze our data, we employ a linear random-effects regression (RE), as well as a Tobit random-effects regression (Tobit RE) accounting for lower bound censoring of the redistribution rates at -10 . The Tobit RE is our preferred regression considering the amount of censoring in our data. Let RedistRate_{it} be the redistribution rate proposed by a rich individual i at the end of round t . We employ two specifications. The first uses a single term for contact; the second distinguishes between intra-ethnic and inter-ethnic contacts. Those are

$$\begin{aligned} \text{RedistRate}_{it} &= \text{Contact}_{it}\omega + \mathbf{X}_i\boldsymbol{\alpha} + \delta_t + \delta_i + \epsilon_{it}, \text{ and} \\ \text{RedistRate}_{it} &= \text{IntraContact}_{it}\omega_{intra} + \text{InterContact}_{it}\omega_{inter} + \mathbf{X}_i\boldsymbol{\alpha} \\ &\quad + \delta_t + \delta_i + \epsilon_{it}. \end{aligned}$$

The variable Contact_{it} takes value 1 if i has contact with any poor in round t , and 0 otherwise. The variable IntraContact_{it} has value 1 if a Caucasian (minority) i has contact with a Caucasian (minority) poor in round t , and 0 otherwise.

²⁴This is after removing participants with a friend, as explained in Section II.F. Moreover, we use data from all 24 sessions, but Result 1 would be the same if we exclude the two sessions that contain one round in which pairing choices were implemented. Note again that participants are not told whether their pairing decisions are implemented. Furthermore, we ignore the redistribution decisions of the poor because they do not face any treatments (and neither of the two hypotheses regarding the effect of contact concerns them).

Table 3.4: Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution

	Redistribution Rate	
	RE	Tobit RE
ω	5.920** (2.474)	9.429*** (3.570)
ω_{intra}	4.224* (2.373)	6.837 (4.238)
ω_{inter}	7.606** (3.133)	11.767*** (4.118)
Wald Test	.183	.260
N	297	297

Note: We report two-sided p -values of Wald tests evaluating the difference between ω_{intra} and ω_{inter} . There are three observations (N) per individual. We employ robust standard errors for the RE regression. Clustering standard errors at the session level in the RE regression provides the same results. Using the wild cluster bootstrap (see Cameron et al. (2008), calculation package by Roodman et al. (2019)) with a linear fixed-effect regression (FE) would also not change the results. p -values: * $p < .10$, ** $p < .05$, *** $p < .01$.

Similarly, the variable $InterContact_{it}$ is 1 if a Caucasian (minority) i experiences contact with a minority (Caucasian) poor in round t , and becomes 0 otherwise. Consequently, the effect of contact with any poor is measured by ω . The effects of intra-ethnic and inter-ethnic contacts reside in ω_{intra} and ω_{inter} , respectively. We control for two demographic characteristics in X_i .²⁵ The regression also includes a trend coefficient, δ_t , and an individual constant, δ_i . The error term is denoted by ϵ_{it} .

Table 3.4 shows the estimates of the effect of contact on redistribution levels. In both regressions, the estimates reveal that contact with the poor increases redistribution (ω , RE $p = .017$, Tobit RE $p = .008$). This is also true once we separate intra- and inter-ethnic contacts, although the effect of intra-ethnic contact is very close to marginal significance using our preferred regression and marginally significant otherwise (ω_{intra} , RE $p = .075$, Tobit RE $p = .107$; ω_{inter} , RE $p = .015$, Tobit RE $p = .004$). Wald tests show that intra-ethnic and inter-ethnic contacts produce statistically indistinguishable effects.²⁶

²⁵Those controls are included for their potential to influence altruistic behavior in this context, but Result 1 stays the same if we exclude them.

²⁶Table 3.8 shows the estimates from the same regressions, but adding rich participants with a friend to the sample. As discussed in Section II.F, including those participants can bias our estimate of the effect of contact. Our data suggests that, relative to Table 3.4, their inclusion creates a negative aggregate bias on the effect of both intra-ethnic and inter-ethnic contacts. Moreover, the direction of the bias varies as predicted based on whether or not those with a friend encountered that friend in their interactions with the poor. Table 3.9 show the regression estimates if we include the rich who interact with a poor friend. Relative to Table 3.4, their inclusion positively biases the effect of both intra-ethnic and inter-ethnic contacts. Finally, Table

We return to Table 3.3 to estimate the size of the effect. Contact approximately increases the mean redistribution rate by 6.00, starting from 7.00. To interpret this effect, we provide two measures. First, the effect of contact corresponds to an increase in redistribution rates close to 1/4 of a standard deviation. Second, using Table 3.2, which converts redistribution rates into after-redistribution earnings, the effect translates in a reduction of the rich's income from approximately 2.6 to 2.4 times the poor's income.²⁷ In sum, the very short-term contact that we implement has a moderate, but positive effect on the redistribution levels chosen by the rich.

We write our first result, which supports Hypothesis 1 and goes against Hypothesis 2, as follows.

RESULT 1: Both intra-ethnic and inter-ethnic contacts with the poor increase income redistribution by the rich. Taken together, their effect is significant; individually, the inter-ethnic effect is significant, the intra-ethnic effect marginally significant or close to.

In addition, as a robustness check, we estimate the coefficients from the same regressions, but accounting for a few different minority groups. We do not report the analysis here, but the results show again that intra-ethnic and inter-ethnic contacts do not produce significantly different effects. Moreover, both effects are qualitatively positive and only the effect of inter-ethnic contact is significant. This is unsurprising because accounting for multiple minority groups reduces the number of observations classified as intra-ethnic contact.

IV.C. Exploratory Analyses—Interaction Decisions

We examine the interaction choices of participants in light of the ethnicity and of the wage of the participants they decide to interact with. We include choices of rich and poor participants in our analyses, since both face the same interaction choices. As explained in Section II.E, we only use decisions made in the first round, to prevent that experience plays a role and to maintain incentive compatibility.

3.10 shows the estimates if we include the rich with a friend, but who do not encounter him or her in their interactions with the poor. Their presence engenders a negative bias.

²⁷We make this comparison by using the change in earnings resulting from a change of redistribution rate from 10 to 15 in Table 3.2. In dollar terms, the income distribution changes from \$7.75 for the poor and \$20.00 for the rich, to \$8.25 for the poor and \$19.75 for the rich.

Table 3.5: Percentage of Participants Choosing to be paired with a Caucasian in Each Rank, for Caucasians and Minorities

Preferred Rank	1	2	3	4	5	6	7	8	9
Caucasians	60.0%	54.4%	52.2%	46.7%	44.4%	44.4%	37.8%	37.8%	49.4%
<i>N</i>	90	90	90	90	90	90	90	90	85
Minorities	54.4%	50.0%	48.9%	58.9%	53.3%	52.2%	47.8%	52.2%	44.3%
<i>N</i>	90	90	90	89	90	90	90	90	79

Note: We indicate the percentage of participants choosing to be paired with a Caucasian (instead of a minority) participant in each rank, separately for Caucasians and minorities. The lower the rank, the more likely an interaction is to occur. We exclude participants who write their own identification number in a rank or who write a number twice. A few observations are missing for some ranks because participants did not select someone in those ranks.

Inter-Ethnic Interactions

Table 3.5 presents the percentage of participants choosing to interact with a Caucasian in each of the nine ranks (1 being the most preferred rank, 9 the least preferred). The upper part of the table presents the data for Caucasians and the lower part the data for minorities. The ranking shows that Caucasians choose more often to be paired with Caucasians in their first than in their last ranks. Minorities exhibit less clear-cut preferences in this regard.

The variable Δ_i is our measure of whether participant i exhibit a preference to be paired with Caucasians. It denotes the number of Caucasian participants that i selects in the first four ranks minus the number of Caucasian participants that i selects in the last four ranks. Therefore, Δ_i ranges from -4 to 4 . It takes value -4 if i chooses no Caucasians in the first four ranks and only Caucasians in the last four ranks. Similarly, it takes value 4 if i chooses only Caucasians in the first four ranks and none in the last four ranks.

We use OLS regressions to test whether participants exhibit a preference to interact with Caucasians, separately for Caucasians and minorities. The dependent variable is Δ_i and the independent variable of interest is the constant term. The constant captures the mean number of additional Caucasian participants chosen in the first four ranks relative to the last four ranks. Table 3.6 presents the coefficient estimates for the constant, with and without control variables. Those controls are for gender, for whether one has a friend in the session, and for ID numbers.²⁸

²⁸For the analysis, we exclude participants who write their own identification number in a rank, who write some number twice or who write nothing for a rank. The controls are as follows. First, we account for one's gender and the difference in the gender of participants selected in the first and last four ranks. We do this by adding a dummy variable for gender, a variable measuring the number of men in the first four ranks minus the number of men in the

Table 3.6: OLS Regression Estimates—Difference in Number of Caucasians in First and Last Four Ranks, for Caucasians and Minorities

	Δ_i							
	Caucasians				Minorities			
Constant	0.424** (0.171)	0.492** (0.208)	0.644** (0.274)	0.839** (0.363)	0.165 (0.203)	0.254 (0.216)	-0.044 (0.272)	-0.251 (0.428)
Controls								
Friend	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Gender	No	No	Yes	Yes	No	No	Yes	Yes
ID number	No	No	No	Yes	No	No	No	Yes
$F_{controls} > 0$.509	.854	.682		.343	.389	.500
N	85	85	85	85	79	79	79	79

Notes: The dependent variable Δ_i is the difference in the number of Caucasians that i selected in the first four ranks relative to the last four ranks. We conduct separate regressions for Caucasian and minority participants. F -tests indicate the joint significance levels of control variables. We employ robust standard errors. Clustering standard errors at the session level provides the same results (with and without the wild cluster bootstrap). p -values: * $p < .10$, ** $p < .05$, *** $p < .01$.

The estimates of the constant confirm the patterns observed in the descriptive data. On average, Caucasian participants are significantly more likely to select other Caucasians in the first relative to the last four ranks. Moreover, this is not the case for minority participants. The control variables are never jointly significant, and do not alter the results. The difference for Caucasians is substantial: they choose on average roughly 0.5 more Caucasians among their four preferred individuals to interact with than among their four least preferred. Using descriptive data from Table 3.5, this translates into Caucasians choosing approximately 25% more Caucasians for their preferred than for their least preferred interactions.

Our second result is thus as follows.

RESULT 2: *Caucasian participants elect to interact with Caucasians more often than with minorities. Minority participants exhibit no significant ethnic prefer-*

last four ranks, and an interaction term between those two variables. We do this to capture the possibility that participants select members of their own or the other gender more frequently. Second, we correct for friendships by including a dummy variable for those participants with a friend in the session. The inclusion of this variable prevents that our constant captures a choice to interact more frequently with those of the same ethnicity in the preferred ranks because of the presence of a friend. Third, we account for the identification number (ID) that we employ (see Section II.E). We include a dummy variable for being assigned a number in the upper half of the number range. We also add a variable for the difference between the sum of the numbers for participants selected in the first four ranks and in the last four ranks. That is, we cover the possibility that being assigned a low or high number or having a tendency to select low or high numbers in the preferred ranks could drive our findings. Moreover, a handful of participants choose the numbers following their own number in ascending order as their preferred four ranks, and controlling for them with a dummy variable in the regression would not affect the results.

ences in their interactions.

In addition, we report that the effect of inter-ethnic contact with the poor on the redistribution proposals of those Caucasian rich who choose more often to interact with other Caucasians appears similar as for our entire sample of rich. That is, the effect of inter-ethnic contact is positive and significant at the 5% level using Tobit RE, and positive but not significant using RE. Table 3.11 provides the estimates from the same regressions as the ones used for Table 3.4, but for this specific subgroup. We do note, however, that this analysis essentially relies on one quarter of our full sample of rich (leaving us with only 26 participants), which decreases our statistical power.

Inter-Income Interactions

The left part of Table 3.7 presents the percentage of participants choosing each of the three options regarding contact with the poor. Recall that the options are: (1) CONTACT: interaction with the poor, (2) NO CONTACT: interaction with the rich (instead of the poor), and (3) INDIFFERENT: whether interaction is with the poor or the rich is randomized. The data is separated for the rich and the poor. As for the choices of inter-ethnic interactions, we include participants with a friend in the session.

The right part of the table shows the p -values from tests evaluating whether participants choose CONTACT and NO CONTACT with the same frequency, as well as whether participants choose CONTACT/INDIFFERENT and NO CONTACT with the same frequency. We employ Logit regressions to obtain those tests. The dependent variable is a dummy variable that takes value 1 if a participant chooses NO CONTACT, and 0 otherwise. The independent variables only include a constant term. The p -values presented are for the constant.

Overall, the table shows that participants rarely choose CONTACT. Noticeably, the rich almost never select CONTACT. Both the poor and the rich prefer NO CONTACT significantly more often than CONTACT. Moreover, the rich select NO CONTACT significantly more often than CONTACT and INDIFFERENT joined together. We also note that the results would be the same if we were to control for having a friend in the session.

Our third result reflects the behavioral patterns regarding the choices of inter-income interactions.

RESULT 3: *Neither the rich nor the poor seek contact with the poor; most of the*

Table 3.7: Choices of Contact with the Poor, and Tests of Equality of Choices, for the Poor and the Rich

	Choice of Contact with the Poor (<i>percentage</i>)			Test Equality of Choices (<i>p-value</i>)		
	CONTACT	INDIFFERENT	NO CONTACT	CONTACT \neq NO CONTACT	CONTACT/INDIFFERENT \neq NO CONTACT	
Poor (<i>N</i> =90)	20.0%	28.9%	51.1%	Poor <i>N</i>	<.001 64	.834 90
Rich (<i>N</i> =133)	5.3%	28.6%	65.4%	Rich <i>N</i>	<.001 95	<.001 133

Note: There are three options regarding contact with the poor: (1) CONTACT: interaction with the poor, (2) INDIFFERENT: interaction with the poor or the rich is randomized, and (3) NO CONTACT: interaction with the rich. The last two columns show *p*-values from Logit regressions. We employ robust standard errors. Clustering standard errors at the session level provides the same results (with and without the wild cluster bootstrap). One participant did not provide an answer and is therefore excluded from the analysis.

rich avoid contact with the poor.

Moreover, the effect of contact with the poor on redistribution is positive and significant for the specific subgroup of rich who choose to avoid contact with the poor. Table 3.12 reproduces the regression estimates from Table 3.4, but for that specific subgroup. If anything, the point estimates are qualitatively larger than for the full sample. This means that those rich who choose NO CONTACT increase their support for redistribution when they are exogenously assigned to interact with the poor.

V. Discussion and Conclusion

Previous influential studies (Luttmer, 2001; Dahlberg et al., 2012) have documented that inter-ethnic proximity with the poor decreases support for income redistribution. Nevertheless, this effect is somewhat at odds with some other literature described in the introduction, and the channel by which inter-ethnic proximity with the poor influences redistributive choices is still unknown. Using a controlled laboratory experiment, we investigate the causal mechanism behind this effect. We fail to find that inter-ethnic proximity with the poor (poorly-paid) participants decreases the amount of redistribution preferred by the rich (highly-paid) participants. In fact, we find that both intra-ethnic and inter-ethnic contacts with the poor similarly increase redistribution by the rich.

Our experiment is also designed to measure participants' avoidance of certain types of interactions. Among our participants, on average, Caucasians tend to avoid interactions with minorities, even though our design rules out the possibility of statistical discrimination based on ability. This finding enables us to

speak to a possible indirect effect of ethnic heterogeneity on support for income redistribution: that it decreases interactions with the poor, which can then lower redistribution. Transposed outside of the laboratory, this means that, relative to living nearby poor of one's own ethnicity, living close to poor of another ethnicity could weaken support for redistributive measures. We also note that our finding that minority participants show no significant average preference for interacting with Caucasians or minorities might result from the heterogeneity of minority groups among our participants.

This additional finding can be useful for public policy aiming to prevent or alleviate a negative effect of ethnic heterogeneity on redistribution. It suggests that such policy should aim to increase inter-ethnic interactions within ethnically-diverse neighborhoods, to create a positive externality on support for redistribution. One could wonder whether those who specifically avoid inter-ethnic interactions can really be positively affected by inter-ethnic interactions with the poor. Our results provide some evidence that this is indeed the case. That is, our subsample of rich Caucasians who choose to interact more frequently with other Caucasians than minorities does not appear to respond any differently than the entire sample.

Moreover, as discussed in the introduction, recent findings have shown that individuals can avoid situations where they give more to others. We show that the majority of rich participants in our experiment avoid interacting with poor participants when given the choice. Some rich are indifferent between interacting with the poor or not, and almost none actively seeks to interact with the poor. Avoiding this contact is not necessarily intended to avoid being generous, but it proves effective to refrain one from giving. For instance, not interacting with the poor can allow enough distance to be stingy with the poor at the same time as preserving a positive self-image (see Bénabou and Tirole (2006) for an example of self-image concerns). This perhaps allows one to avert painful cognitive dissonance (Festinger, 1962; Akerlof and Dickens, 1982). As a caveat, we note that the very strong avoidance of the poor that we find may overestimate avoidance in the field, where affinity for a neighbor can obviously exist on other grounds than income.

The short-lived interactions that we employ enable participants to recognize an individual affected by their redistribution choices, as do interactions with poorer individuals in one's neighborhood.²⁹ As such, the Identified Victim Effect literature suggests that this recognition renders individuals more empa-

²⁹We do not claim that individuals know the exact income of their neighbors as in this experiment, but that they have an idea of whether those neighbors are poor or not.

thetic and altruistic. However, repeated interactions with nearby poorer neighbors may generate a larger positive effect on support for redistribution. For instance, Contact Theory suggests that acquaintance with an out-group member can improve knowledge about and empathy toward other income or ethnic groups.

One limitation of our study is that we employ only one format of interactions between individuals, which is relatively neutral, and that we do not vary the quality of those interactions. For instance, a difference in the effect of intra-ethnic and inter-ethnic interactions with the poor could arise because inter-ethnic interactions are more frequently negative. Another limitation is that we do not vary the source of income inequality; it is always assigned in our design, and sources of inequality and beliefs about those sources have been shown to influence redistribution decisions (e.g., Fong, 2001; Cappelen et al., 2007). Whether interactions with or identification of the poor whose income is generated through a mix of effort and luck has the same effect as the one we find remains to be explored.

Finally, future studies could investigate whether the indirect effect of inter-ethnic proximity that we uncover is indeed found in the field. Future research could also try to assess the importance of this effect relative to other possible forces.

Appendix

3.A. Additional Tables

Table 3.8: Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution—Including All Participants with a Friend in the Session

	Redistribution Rate	
	RE	Tobit RE
ω	3.148 (2.483)	5.734 (3.710)
ω_{intra}	3.051 (2.533)	5.534 (4.430)
ω_{inter}	3.242 (2.969)	5.918 (4.331)
<i>Wald Test</i>	.937	.934
<i>N</i>	399	399

Note: Compared to Table 3.4, the aggregate bias induced by including participants with a friend on the effect of intra-ethnic and inter-ethnic contacts is negative. We report two-sided p -values of Wald tests evaluating the difference between ω_{intra} and ω_{inter} . There are three observations (N) per individual. We employ robust standard errors for the RE regression. Clustering standard errors at the session level in the RE regression provides the same results. Using the wild cluster bootstrap with an FE regression has also no bearing on the results. p -values: * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 3.9: Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich’s Support for Redistribution—Including Participants with a Friend who Interact with their Friend in their Contact with the Poor

	Redistribution Rate	
	RE	Tobit RE
ω	6.581*** (2.383)	11.300*** (3.659)
ω_{intra}	5.163** (2.350)	9.221** (4.312)
ω_{inter}	8.026*** (2.979)	13.241*** (4.234)
<i>Wald Test</i>	.242	.365
<i>N</i>	318	318

Note: Compared to Table 3.4, as predicted, the bias induced by including participants who interacted with their poor friend is positive. We report two-sided p -values of Wald tests evaluating the difference between ω_{intra} and ω_{inter} . There are three observations (N) per individual. We employ robust standard errors for the RE regression. Clustering standard errors at the session level in the RE regression provides the same results. Using the wild cluster bootstrap with an FE regression has also no bearing on the results. p -values: * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 3.10: Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich’s Support for Redistribution—Including Participants with a Friend who do not Interact with their Friend in their Contact with the Poor

	Redistribution Rate	
	RE	Tobit RE
ω	2.469 (2.567)	3.988 (3.677)
ω_{intra}	2.231 (2.593)	3.442 (4.422)
ω_{inter}	2.693 (3.091)	4.478 (4.285)
<i>Wald Test</i>	.854	.824
<i>N</i>	378	378

Note: Compared to Table 3.4, as predicted, the bias induced by including participants with a friend who did not encounter their friend in their interactions with the poor is negative. We report two-sided p -values of Wald tests evaluating the difference between ω_{intra} and ω_{inter} . There are three observations (N) per individual. We employ robust standard errors for the RE regression. Clustering standard errors at the session level in the RE regression provides the same results. Using the wild cluster bootstrap with an FE regression has also no bearing on the results. p -values: * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 3.11: Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution—Subgroup of the Caucasian Rich who Choose more often to Interact with Caucasians

	Redistribution Rate	
	RE	Tobit RE
ω	6.418 (4.525)	11.550* (6.608)
ω_{intra}	5.507 (4.185)	6.785 (8.275)
ω_{inter}	7.101 (5.219)	14.484** (7.291)
<i>Wald Test</i>	.611	.337
<i>N</i>	78	78

Note: We use the subsample of the 26 Caucasian rich participants choosing more Caucasians in their preferred four ranks than in their least preferred four ranks of participants to interact with. We report two-sided p -values of Wald tests evaluating the difference between ω_{intra} and ω_{inter} . There are three observations (N) per individual. We employ robust standard errors for the RE regression. Clustering standard errors at the session level in the RE regression provides the same results. Using the wild cluster bootstrap with an FE regression has also no bearing on the results. p -values: * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 3.12: Estimates of the Effect of Intra- and Inter-Ethnic Contact with the Poor on the Rich's Support for Redistribution—Subgroup of the Rich who Choose to Avoid Contact with the Poor

	Redistribution Rate	
	RE	Tobit RE
ω	9.349*** (3.394)	14.125*** (5.074)
ω_{intra}	7.034** (3.098)	11.587* (5.991)
ω_{inter}	11.758** (4.591)	16.541*** (5.901)
<i>Wald Test</i>	.223	.428
<i>N</i>	189	189

Note: We report two-sided p -values of Wald tests evaluating the difference between ω_{intra} and ω_{inter} . There are three observations (N) per individual. We employ robust standard errors for the RE regression. Clustering standard errors at the session level in the RE regression provides the same results. Using the wild cluster bootstrap with an FE regression has also no bearing on the results. p -values: * $p < .10$, ** $p < .05$, *** $p < .01$.

3.B. Instructions

Part 1 – General Instructions³⁰

Welcome. In this experiment, the instructions are simple and you can earn a considerable amount of money by following them. You are not allowed to talk with other participants except when instructed otherwise. If anything is unclear, please do not hesitate to ask the experimenter. You can do so at any time by raising your hand. At the end of the experiment, the earnings you have made will be paid to you in cash on top of a \$3 show-up fee.

You will go through 3 rounds of a game called the Neighborhood Game. Each round, you have to complete a task within 8 minutes. The task is done individually. It is to copy some text by hand. It can be done on time with moderate effort. You receive a wage for completing the task. There are 6 participants who receive a high wage of \$20.50 and 4 participants who receive a low wage of \$7. The length of the text to copy is smaller for participants with a low wage. This wage will be determined for each of you in a moment. This wage does not change throughout the rounds. If you are assigned a low or high wage now, you will keep this same wage for the 3 rounds. Hide your wage in your pocket and do not reveal it to others except when instructed otherwise.

After completing the task, each of you proposes a redistribution rate for the group. The redistribution rates proposed are anonymous. After the end of the experiment, one of the 3 rounds is randomly selected to count for payment. One redistribution proposal from the selected round is randomly chosen and implemented for everyone. Each participant is paid his/her after-redistribution wage for the selected round. Participants are not paid for the 2 rounds that are not selected.

Every round, each of you will be paired with another participant. Before engaging in the task, which is performed individually, you will be asked to reveal the following information to the participant you are paired with: name, age, major, year of study, and your wage for the task. After doing so, you can start the task and may not communicate with the participant you are paired with or anyone else. More information will be provided once you are seated.

³⁰Note to the reader: Participants were knowingly videotaped for a part of the session. We subsequently did not use the videos because of an error in the recording software.

Part 2 – Selection

Before each round, you have a chance to select which participant you would prefer to be paired with. You will have a new chance to do this before each round. This selection will never be revealed to other participants. You make a selection in two ways. First, you can choose whether you would prefer to be paired with a high or a low wage participant. Second, you can rank the other participants in order of preference to be paired with. To do so, you can use the identification number of each participant. No participants can be paired together twice.

There is a 2% chance that one of the two ways of pairing is selected. Each way has an equal chance of being selected. Therefore, each way has a 1% chance of being selected.

If one of the ways is selected, one participant is drawn randomly first and given his/her first choice. If the participant has already been matched with the participant of his/her choice, he/she will be paired with his/her next preferred choice. Then, another participant is drawn and given his/her preferred choice that is still available. This goes on until there is no participant left. This process is conducted by the experimenter and is not revealed to any participant.

If neither the first way nor the second way is randomly selected, the experimenter allocates participants to others, without taking into account participants' choices in the selection questions.

Moreover, to ensure the privacy of your answers with respect to other participants, it will not be revealed if one of the ways by which you can select others yourself has been picked to assign participants. The experimenter will simply bring everyone's answers to another room and come back with the seat number of each participant.

The experimenter will call you to go to a table to make your choices without others being able to see. There will be a paper on the table to write down your choices. There is a reminder of what you should do on this paper. Once you are done with writing down your choices, please fold the paper in two and place it

in the box on the table.

Part 3 – Communication

Before we go to the task, we kindly ask you to introduce yourself to the participant seated next to you. To do so, please reveal the following information: your name, age, major, and year of study. Please also reveal your wage by showing the card your wage is written on. We ask you not to communicate anything else. The participant on the aisle side can start.

Part 4 – Task and Redistribution Proposal

You are asked to complete individually a task to earn your wage. The wage cannot be earned without completing the task on time. The task is to copy the text on your desk onto the empty sheet of paper, also on your desk. It is simple and can be completed by all participants on time with moderate effort. Each of you has the exact same text. Participants with the high wage of \$20.50 have to copy the full text. Participants with the low wage of \$7 have to copy 2/3 of the text. Only the part of the text in gray has to be copied by participants with a low wage. Participants with a high wage have to copy the part in gray and the part in black.

You will be given 8 minutes to complete the task. You will be reminded when there are 4 minutes left and when there are 2 minutes left. The experimenter will also come to each of you in the first few minutes to make sure that you are doing the right thing. Even if all participants finish earlier, everyone will still have to wait for the 8 minutes to pass.

Once the time is over, you are asked to propose a redistribution rate. This proposal will not be revealed to other participants. At the end of the experiment, one of the 3 rounds will be selected randomly. One redistribution rate proposed in the round selected will be randomly chosen to count for everyone's payment. Only this selected rate will be revealed and not who proposed it. If you do not complete the task, you can neither propose a redistribution rate for the round nor receive payment for the round.

The redistribution rate can be any rate from -10% to 115% , in any increment

of 5%. A higher redistribution rate makes the after-redistribution wages more equal up to the rate of 100%. A redistribution rate of 0% leaves them as they were originally. The after-redistribution wages will be paid for the round selected at the end of the experiment.

Please keep in mind that each participant's wage will be the same in each of the 3 rounds. Please also keep in mind that there are 6 participants with a high wage and 4 with a low wage. The following table gives you the after-redistribution wages in dollars of both types of participants for all possible redistribution rates:

Redistribution Rate (%)	After-Redistribution Wage		Total Group Earnings (\$)
	Low-Wage Participant (\$)	High-Wage Participant (\$)	
-10	6.25	21.00	151
-5	6.50	20.75	150.5
0	7.00	20.50	151
5	7.50	20.25	151.5
10	7.75	20.00	151
15	8.25	19.75	151.5
20	8.50	19.50	151
25	9.00	19.25	151.5
30	9.50	19.00	152
35	9.75	18.50	150
40	10.25	18.25	150.5
45	10.75	18.00	151
50	11.00	17.75	150.5
55	11.50	17.50	151
60	11.75	17.25	150.5
65	12.25	17.00	151
70	12.75	16.75	151.5
75	13.00	16.50	151
80	13.50	16.25	151.5
85	14.00	16.00	152
90	14.25	15.75	151.5
95	14.75	15.25	150.5
100	15.00	15.00	150
105	15.50	14.75	150.5
110	16.00	14.50	151
115	16.25	14.25	150.5

The experimenter will give you a moment to examine the table. Understanding the table is crucial to understanding how the redistribution rate you propose can affect your payment and the payments of others. If you have any questions, please do not hesitate to ask.

3.C. Task

Low-wage participants copy only the part of the text in blue and high-wage participants copy the whole text.³¹

³¹Note for the reader: the three tasks have the same length, and correspond to the task in round 1–3. The text for the low-wage participants is highlighted in gray. The text consists of the

TASK 1

- (1) Arma virumque cano, Troiae qui primus ab
- (2) oris Italiam, fato profugus, Laviniaque venit
- (3) litora, multum ille et terris iactatus et alto
- (4) vi superum saevae memorem Iunonis ob iram;
- (5) multa quoque et bello passus, dum conderet urbem,
- (6) inferretque deos Latio, genus unde Latinum,
- (7) Albanique patres, atque altae moenia Romae.
- (8) Musa, mihi causas memora, quo numine laeso,
- (9) quidve dolens, regina deum tot volvere casus
- (10) insignem pietate virum, tot adire labores
- (11) impulerit. Tantaene animis caelestibus irae?

TASK 2

- (1) Urbs antiqua fuit, Tyrii tenuere coloni,
- (2) Karthago, Italiam contra Tiberinaque longe
- (3) ostia, dives opum studiisque asperrima belli;
- (4) quam Iuno fertur terris magis omnibus unam
- (5) posthabita coluisse Samo; hic illius arma,

first 33 lines of Virgil's Aeneid.

- (6) hic currus fuit; hoc regnum dea gentibus esse,
- (7) si qua fata sinant, iam tum tenditque fovetque.
- (8) Progeniem sed enim Troiano a sanguine duci
- (9) audierat, Tyrias olim quae verteret arces;
- (10) hinc populum late regem belloque superbum
- (11) venturum excidio Libyae: sic volvere Parcas.

TASK 3

- (1) Id metuens, veterisque memor Saturnia belli,
- (2) prima quod ad Troiam pro caris gesserat Argis-
- (3) necdum etiam causae irarum saevique dolores
- (4) exciderant animo: manet alta mente repostum
- (5) iudicium Paridis spretaeque iniuria formae,
- (6) et genus invisum, et rapti Ganymedis honores.
- (7) His accensa super, iactatos aequore toto
- (8) Troas, reliquias Danaum atque immitis Achilli,
- (9) arcebat longe Latio, multosque per annos
- (10) errabant, acti fati, maria omnia circum.
- (11) Tanta molis erat Romanam condere gentem!

Chapter 4

Reciprocity under Brief and Long Time Delays

Abstract

We report the results from three experiments embedded in the same overarching design, which extends the Gift Exchange paradigm for the study of worker-employer relationships. We focus on the effect of the length of the delay, between the time at which workers learn their wage and when they choose an effort level, on the relationship between wage and effort. We compare effort choices made within a few hours after workers learn their wage, with those made several weeks afterward. Our two effort choices are single “cold” choices, but the varying delay allows us to identify any dissipation of reciprocal behavior. We find that the strength of the wage-effort relationship decreases over time, and this change appears to be driven by those workers who receive low wages. We extend our study of delays to the wage-effort relationship below self-reported fair wages. Our findings suggest that workers’ inclination to punish stingy employers weakens over time.

A version of this chapter, co-authored with Charles N. Noussair, appears in the July 2020 issue of *Economic Inquiry*.

I. Introduction

Experimental research on the behavior of labor markets in which worker effort is not contractible, beginning with Fehr et al. (1993), has yielded a number of important insights. Even in one-shot interactions, the average wages paid to workers considerably exceed the competitive market-clearing level. In addition, worker effort tends to be greater than the minimum effort possible, despite the fact that exerting effort is costly to the worker. A robust relationship of workers reciprocating higher wages with greater effort is observed. These findings have spawned a large literature (e.g., Fehr et al., 1996, 1997, 1998a,b; Fehr and Falk, 1999; Gächter and Falk, 2002; Hannan et al., 2002; Brandts and Charness, 2004; Brown et al., 2004; Charness, 2004). See Fehr et al. (2009), Charness and Kuhn (2011), Casoria and Riedl (2013), and Cooper and Kagel (2016) for surveys.

Although worker-employer relationships may be short-term or long-term in nature, the laboratory studies cited above consider immediate behavior, in the sense that all wage and effort decisions are made within one laboratory session, which is typically under two hours in duration. Several recent studies report field experiments focused on whether the reciprocal relationship between wage and effort is durable over longer time horizons. Reciprocity, in the form of increased effort by workers, is commonly found within a few hours following a wage increase. In contrast, the few studies analyzing behavior over a longer timeframe provide some evidence that the tendency toward reciprocation can survive for at least one day. The results are consistent with the notion that reciprocity fades over horizons of days or weeks.

Specifically, several experimental studies have considered worker reciprocity over a time frame of several hours. Kube et al. (2012) observe higher effort for the three hours they monitor after a one-time non-monetary gift, but not for a comparable monetary gift. Gilchrist et al. (2016) detect a similar effect for four hours following an unexpected wage increase, but not when the wage increase is expected. Kube et al. (2013) find no reciprocal response to a wage increase, but do observe a sharp decrease in productivity following a wage decrease, over the six hours of their experiment. In a study spanning two work days, Gneezy and List (2006) observe that reciprocity lasts only for a few hours after an increase in workers' hourly wage. A few studies have observed an effect for a longer time span, but these studies are not conclusive. Bellemare and Shearer (2009) find reciprocity lasting for one day after a one-off bonus payment, and thereafter workers' effort comes back down to previous levels. Cohn

et al. (2015) report higher productivity following increases in hourly wages for the entire time horizon of their study, which was several days. However, they vary workers' wages more than once, which may renew the stimulus. Hossain and List (2012) find an increase in productivity lasting for four weeks after a one-time monetary payment, though they rely on a relatively small sample of workers.¹

Those studies differ from each other in many aspects. The time horizons analyzed differ, ranging from a few hours to four weeks. Effort decisions are also typically repeated, and thus some workers might exhibit a decline in effort because they feel like they have already reciprocated for any prior employer wage decisions. Furthermore, worker behavior might depend on the context (List, 2006; Hennig-Schmidt et al., 2010; Kube et al., 2012, 2013; Gilchrist et al., 2016), the perception of the fairness of the baseline wage (Fehr et al., 2009; Cohn et al., 2015), prior history (Bellemare and Shearer, 2009), or anticipated future interaction between worker and employer (Hossain and List, 2012).

In the experiments reported here, we consider whether reciprocal behavior is long-term in nature, and we do so in a very simple one-shot design. This feature rules out any effect of past or anticipated future interaction. We compare the decisions of "short-term" workers choosing their effort within three hours of learning their wage, with those of "long-term" workers making their choice four weeks after receiving their wage information. Before long-term workers are required to make their effort decision, they are reminded of the wage decision their employer made four weeks previously. This is similar to the typical situation of an employee or contractor who learns her wage when hired, and is then reminded about her wage at the time she actually starts the job later on.² Our short-term delay represents the case of workers who learn their wage shortly before their work begins, and then choose how much effort to exert. Under the "hot vs. cold" terminology, both of our effort decisions are "cold" decisions, made after a period of reflection.³ The longer delay allows us to measure the dissipation of the reciprocity motive between the two time horizons. There is

¹In most of these studies, it is unclear whether the observed wage-effort relationship consists of positive or negative reciprocity. This is because reference wages, which are needed to classify reciprocity as positive or negative, are not elicited.

²To conduct this experiment, it is imperative to remind long-term workers of the wage that they were offered previously. Otherwise, many would forget the wage that they were offered, and we could not be sure that their effort decisions were taken with the knowledge of the wage that they had received.

³"Hot" decisions are made instantaneously, whereas "cold" decisions allow for more reflection. See, for example, the articles by Brandts and Charness (2000) and for a discussion of differences in the two settings.

exactly one decision made by each worker. This allows us to clearly identify any and all reciprocal behavior that the worker exhibits.⁴ The experiment is decontextualized to the extent that is typical in laboratory experiments. Interactions are anonymous so that prior or future interactions between participants are not relevant to their decisions.

We conduct three separate experiments, where we extend the laboratory Gift Exchange paradigm for the study of employer-worker relationships to long time horizons. The experiments are embedded in the same overarching design that allows scope for a reciprocal worker-employer relationship to emerge. The second experiment improves possible censoring issues faced by the first, and the third experiment employs a richer design that allows us to consider reciprocity from additional perspectives. To take advantage of the entire database (266 workers), we analyze our three experiments together whenever possible. In all three experiments, we consider whether workers choose effort levels to reciprocate employers' wage decisions when wages were set either less than three hours or one month previously. Subjects are divided into groups of three: one employer and two workers (a short-term and a long-term worker).⁵ The employer chooses one wage level that applies to both workers. The short-term worker then chooses an effort level after leaving the laboratory, but at most three hours afterward. The long-term worker submits the effort decision four weeks after the experimental session. At the time their decision is due, long-term workers receive a reminder informing them of the wage they received and the payoff structure in effect.⁶

We structure our three experiments to allow application of the Partial Gift Exchange version of the Efficiency Wage hypothesis of Akerlof (1982) to the data. Our third experiment is designed to allow a test of a key part of a later

⁴One can ask whether reciprocal behavior decreases beyond a certain threshold precisely because workers feel that they have reciprocated a past wage increase. While this is a very interesting question, we believe it is important to first ask whether reciprocity decreases over time when only a single decision is taken. This is because finding that reciprocity decreases in a repeated setting could simply reflect the fact that workers are less inclined to reciprocate with the passing of time.

⁵Maximiano et al. (2007) show that the reciprocity found in Gift Exchange laboratory experiments is also observed when multiple workers are matched with one employer who offers the same wage to all workers. Other studies considering the case of multiple workers and one employer in Gift Exchange settings include Güth et al. (2001), Charness and Kuhn (2007), Abeler et al. (2010), Gächter and Thöni (2010), Cohn et al. (2014), and Charness et al. (2016). Most of these studies investigate social comparison effects when workers receive different wages. In our experiments, we circumvent these effects by constraining employers to offer the same wage to both of their workers.

⁶The reminder creates roughly the typical level of information a worker has when one starts a job (though the wage may have been set considerably in advance, workers typically have not forgotten how much they are getting paid when work begins on a project).

refinement of the hypothesis, the Fair Wage-Effort hypothesis of Akerlof and Yellen (1988) and Akerlof and Yellen (1990), in each of the two timeframes. The two models describe the worker-employer relationship in an environment in which effort is not contractible. In the first model, Akerlof (1982) assumes that higher wages offered by the employer are reciprocated with greater effort on the part of workers. In the second model, Akerlof and Yellen (1990) assume that workers form a belief regarding what the fair wage is, and that they exhibit an asymmetric response to the wages about this fair wage. Wages below those that workers perceive as fair are punished with low effort, but those in excess of fair levels do not result in additional effort. Cohn et al. (2015) recently obtained support for this second model, finding that the sustained reciprocity they uncover in their field experiment is driven by the subset of workers who felt that they were paid unfairly before a wage increase. That is, only those workers who felt that they were paid an unfair wage before increased their productivity in response to the increase, at least over the timespan of several days that they study. We test whether we observe reciprocity when wages are below fair wages (which we refer to as negative reciprocity) as well as when wages are above fair wages (which we call positive reciprocity), and whether one type of reciprocity dissipates more than the other with time.⁷

One reason that reciprocity may be weaker in the long term than in the short term is the role of emotions. In justifying their Fair Wage-Effort hypothesis, Akerlof and Yellen (1990) propose that wages below those considered as fair trigger a reaction of anger on the part of workers, and in turn the feeling of anger triggers low effort. There exists empirical evidence in economics and social psychology that anger is correlated with immediate (over seconds or a few minutes) negatively reciprocal behavior ((Bosman and Van Winden, 2002); Offerman, 2002; Sanfey, Rilling, Aronson, Nystrom, and Cohen, 2003; Xiao and Houser, 2005; (Ben-Shakhar et al., 2007); Reuben and van Winden, 2008; Hopfensitz and Reuben, 2009; Reuben and van Winden, 2010, (Bolle et al., 2014); van Leeuwen, Noussair, Offerman, Suetens, van Veelen, and van de Ven, 2018; Harth and Regner, 2017) in contexts other than the Gift Exchange paradigm. Kirchsteiger, Rigotti, and Rustichini, (2006) observe a similar relationship in the Gift Exchange setting. Psychological research has also

⁷That is, wages lower (higher) than the fair level are assumed to lie in a negative (positive) domain for the worker, with the potential to trigger negative (positive) reciprocity. Several authors have used fairness concerns to differentiate positive and negative reciprocity. In addition to Akerlof and Yellen (1988) and Akerlof and Yellen (1990), which we discuss below, we note the work of Rabin (1993) and Dufwenberg and Kirchsteiger (2004), where negative deviations from fair or equitable behavior that are perceived as intentional generate negative reciprocity.

highlighted the role of anger in aggression and retaliation more generally, and that anger differs among individuals facing the same situation, i.e., it is influenced by one's personality traits as well as norms (e.g., Averill, 1983; Frijda, 1986; Eisenberger et al., 2004). Within economics, Battigalli et al. (2019) have recently modelled anger and aggression, using the tools of psychological game theory. Unlike Akerlof and Yellen (1990), they consider the possibility that anger subsides over time and that the desire for retaliation wanes. In Appendix 4.A, we report some suggestive evidence that anger plays a role in the behavior that we observe.

We obtain two main results in this study. First and foremost, the strength of the reciprocal wage-effort relationship is weaker after several weeks than it is within three hours. In particular, workers receiving low wages choose greater effort after the long delay, indicating that workers are more inclined to punish stingy employers in the short-term than later on. Second, we fail to find evidence in favor of the Fair Wage-Effort hypothesis. For short-term workers, we observe that reciprocity is not significantly stronger for wages below self-reported fair wages than for those above. There is a similar lack of asymmetry in the effect of wages below and above fair levels for long-term workers.

The rest of this paper is organized as follows: section II describes the experimental design, section III presents the hypotheses, section IV reports the results, and section V offers our conclusions.

II. Experimental Design

In this section we describe the structure of each of the three experiments. The first two experiments were conducted in the Netherlands and the third one in the United States. The first one is called Tilburg-L (Tilburg Low Efficiency), the second is called Tilburg-H (Tilburg High Efficiency), and the third is called Tucson-H (Tucson High Efficiency). Section-II.A describes the procedures that were identical in the three experiments. Section-II.B then covers the aspects that differed among them. It explains why three experiments were conducted and how the last experiment, Tucson-H, allows testing of some hypotheses beyond those evaluated in the first two experiments.

II.A. Procedures Common to all Experiments

Participants are assigned randomly to one of three roles: 1) employer, 2) short-term worker, or 3) long-term worker. The roles are private information. Groups are formed, each consisting of one employer and two workers, one of each

type. We describe the two workers as short- or long-term, based on when they make their effort decision. The participants are informed that the groups are randomly constituted and anonymous. It is emphasized to them that they will play the game exactly once.

The game has two stages. In the first stage, the employer, endowed with wealth a , decides on a wage w to pay to her two workers. The wage must be equal for the two workers. The wage is costly to the employer and benefits the workers. In the second stage, each worker, with initial endowment d , observes his wage and chooses an effort level e . The worker's effort is costly to himself, with marginal cost c , and benefits the employer by the productivity parameter b .

The short-term worker must submit an effort level, e_I within three hours after the end of the session. The long-term worker chooses an effort level e_D within a five-day interval beginning four weeks after the session.⁸ The workers, therefore, do not choose their effort in the laboratory. They send their choice to an email address provided to them. Participants are informed during the session that the long-term workers would receive an email reminder, containing the wage they have been awarded and a copy of the instructions.⁹ Participants are also informed that if a worker does not send an email with her effort choice, she would not be paid anything beyond the show-up fee and the employer would be paid back the wage offered to that worker.¹⁰ The earnings of workers are sent to them on the day after the receipt of their effort choice. The earnings of the employer are sent on the day after the receipt of the effort choice of the long-term worker she is matched with.

The payoffs to the three types of participants are given by:¹¹

$$\text{Employer: } a - 2 \times w + b \times (e_I + e_D)$$

⁸The time intervals for Tilburg-L differ slightly. Short-term workers were instructed to send their effort choice within two hours, and long-term workers during a 24-hour window that began one month after the session. This was not enough time for some workers to send in their effort choice. We accepted effort choices from a few short-term workers after the two-hour deadline. We informed long-term workers that they would have additional days to choose their effort, and accepted choices up to eight days after the original deadline.

⁹The instructions are provided in Appendix 4.D and the email reminders are given in Appendix 4.E.

¹⁰Paying back the wage to employers when workers do not choose an effort level does not change the fact that the least costly action of workers is choosing an effort of zero, and the action of workers that earns the employer the most is choosing maximum effort.

¹¹The payoff structure was presented to subjects in terms of these formulas, with the actual values in effect, rather than the variables indicated. See Charness et al. (2004) for a discussion of the effects of different formats of presentation.

Short-Term worker: $d + w - c \times e_T$

Long-Term worker: $d + w - c \times e_D$

The ranges of possible wages, effort levels, effort costs and output differ among the three experiments. The values are given in Table 4.1 along with other procedural details. Wages and effort can be chosen in increments of 10 cents and 0.1 units, respectively. Note that the ratio b/c , denoting the benefit of effort to the employer over the cost of effort to the worker, is constant within each experiment (Gift Exchange experiments often feature a diminishing b/c ratio). We use a constant ratio for simplicity, which is especially desirable since participants only play the game once. Sessions generally had 9 or 12 participants and lasted for 40 minutes. Participants are 51.5% male and their mean age is 22.1 (SD=3.0).

A session proceeds in the following manner. Participants arrive at the laboratory and are seated individually at a computer. They are given a written copy of the instructions, which the experimenter reads aloud. Participants are forbidden from communicating with others. Before making their choices, they must complete two practice exercises to confirm their understanding. They are encouraged to ask questions and their individual answers are verified by the experimenter. Help is provided if needed, as each participant is required to fill in the correct answers to proceed further.¹²

II.B. Procedures Specific to Each Experiment

The Tilburg-L and Tilburg-H Experiments

The procedures for the Tilburg-L experiment are described fully in Section-II.A. The Tilburg-H experiment differs from Tilburg-L in one main aspect: it reduces the potential censoring of effort from above for low wages in Tilburg-L. That is, unlike in Tilburg-L and as in standard Gift Exchange experiments, the effort choice of workers in Tilburg-H is not limited by the wage offered. In Tilburg-L, workers with low wages are restricted to choose low wages both in the short and long term. This means that we cannot observe a possible effect of the time delay on the effort choices of workers with low wages in Tilburg-L. Moreover, the parameters in effect in Tilburg-H also ensure greater productivity of effort. This is captured in the higher b/c ratio. We increased

¹²The game was programmed using the z-Tree software (Fischbacher, 2007).

Table 4.1: Procedural Details

Parameters, Range of Variables, and other Details	Full Sample	Tilburg-L	Tilburg-H	Tucson-H
a		12 EUR	8 EUR	11 USD
b		2.5	2.5	2.5
c		1	0.25	0.25
d		4 EUR	5 EUR	4 USD
e		[0, w]	[0, 2.4]	[0, 4]
w		[0 EUR, 4 EUR]	[0 EUR, 2 EUR]	[0 USD, 4 USD]
Show-up Fee		0 EUR	0 EUR	8 USD
Avg Earnings Employer		10 EUR	8 EUR	22 USD
Avg Earnings Short-Term Worker		6 EUR	6 EUR	14 USD
Avg Earnings Long-Term Worker		6 EUR	6 EUR	14 USD
Payment Method		Bank Transfer	Bank Transfer	Cash and Mail
Time Period Conducted		March–May, 2015	April–June, Nov.–Dec., 2015	Jan.–March, 2016
Location		CentER Lab (Tilburg, NL)	CentER Lab (Tilburg, NL)	ESL (Tucson, US)
N Participants	431	138	99	194
Attrition Short-Term Workers	4%	0%	3%	8%
Attrition Long-Term Workers	10%	9%	9%	13%
Net N Participants	410	134	95	181
Net N Employers	144	46	33	65
Net N Short-Term Workers	138	46	32	60
Net N Long-Term Workers	128	42	30	56

Note: (a) Bank transfers are the most common method of payment in the Netherlands, even for small payments. (b) One subject in the role of long-term worker participated twice. We removed the second observation. (c) The wages are similar (i) for short- and long-term workers who do not respond, and (ii) for workers who do not respond and workers who do respond. This suggests that attrition is unlikely to drive differences between the behavior of short- and long-term workers. First, as explained in section IV, we normalize wages by dividing the wage in each experiment by the maximum feasible wage offer in this experiment. The mean normalized wages of short- and long-term workers who did not respond are 0.54 (SD=0.31, N=6) and 0.54 (SD=0.32, N=15), slightly above half of the maximum possible wage and close to the mean normalized wage of 0.60 (SD=0.30, N=266) for workers who did respond. Second, we find no significant relationships between attrition and wage offers for either short-term or long-term workers using an OLS regression.

the ratio to reduce censoring from below, in order to be better able to observe possible changes in effort over time.¹³ At the same time that he chooses a wage, the employer also requests the same non-binding effort level from each worker. The effort request is included in the reminder to long-term workers.

The Tucson-H Experiment

This experiment extends the design of Tilburg-H to allow us to consider reciprocity in greater detail. Although the key ratio b/c remains the same, there are other differences in the parameters, as indicated in Table 4.1. Moreover, workers are asked to state what they think would be a fair wage before learning their actual wage. It is emphasized to them that this fair wage would not be revealed to any other participant, including their employer, and therefore could

¹³Kessler (2013) found that a relatively high ratio tends to reduce the censoring of the effort data at the lower bound.

not affect the wage they receive. Participants receive a show-up fee of 8 USD at the time of the session on top of their earnings from the game, which are paid at a later time.

III. Hypotheses

The first two hypotheses apply to all three experiments, and concern the effect of long vs. short time delays on effort choices. They originate from the results of prior experimental work. Hypotheses 3 and 4 apply exclusively to our third experiment, Tucson-H. They emerged after the first two experiments and guided the design of Tucson-H. These hypotheses posit asymmetries in worker behavior between cases in which the wage is perceived as fair and as unfair.

We first consider the relationship between effort and time delay, conditional on wage. That is, we test whether delay has any effect on the level of effort or on the sensitivity of effort to wages received. In the regression specification that we employ, these correspond to the intercept and slope coefficients, respectively. We allow for (1) long-term workers choosing a different effort level than short-term workers irrespective of the wage, and/or (2) long- and short-term workers exhibiting different sensitivities of effort to wage. Prior evidence that the level of effort may be affected by time delays comes from the dictator game literature (Kovarik, 2009; Dreber et al., 2016), investigating the relationship between altruism and time-lags. They find that when dictators have to split money that is payable in the future, they keep more for themselves. However, our environment is very different, and it is unclear whether effort levels would differ at the two time lags, controlling for wage level. Furthermore, as we discussed in the introduction, while the current literature suggests that the wage-effort relationship is weaker in the long term, the evidence is not extensive. Therefore, we take a conservative approach and propose the following null hypothesis.

HYPOTHESIS 1: For a given wage, short-term workers exert the same effort as long-term workers ($e_I|w = e_D|w$).

Hypothesis 2 specifically considers whether a positive relationship between wage and effort appears with similar force within three hours and one month after workers learn their wages. As discussed in the introduction, the overall evidence from field experiments supports the contention that there is a decrease in reciprocity after a long delay. The Gift Exchange version of the Efficiency

Wage Hypothesis (Akerlof, 1982), on which most of the reciprocity labor literature is based, does not consider whether time should affect reciprocal behavior. In laboratory studies comparing reciprocity under immediate response and a short time delay in other types of interaction, the results are mixed, but generally find a decrease in reciprocity. Bosman et al. (2001) find that anger is associated with rejections in the ultimatum game and that a one-hour delay does not decrease rejections. However, in their study, only a handful of participants reject offers so that a decrease would be difficult to detect. More recently, Grimm and Mengel (2011) find that a 10-minute delay reduces rejections in the ultimatum game and Neo et al. (2013) observe a reduction of rejections in the ultimatum game with a similar delay, but no effect of delay in the investment game. Also related is the Internet-based experiment of Oechssler et al. (2015), which finds that an opportunity to revise one's decision 24 hours after having made a decision to accept or reject an offer in an ultimatum game decreases rejection in only one of their two conditions. Thus, our second null hypothesis is as follows.¹⁴

HYPOTHESIS 2: The sensitivity of effort to wage is the same for short- and long-term workers ($de_I/dw = de_D/dw$).

Hypothesis 3 is a test of the Fair Wage-Effort Hypothesis (Akerlof and Yellen, 1988, 1990). The hypothesis relies on the assumption that workers respond to a higher wage with higher effort only up to the point where the wage is considered fair. Workers do not reciprocate for an increase in wage above the fair wage with greater effort. Cohn et al. (2015) are, to our knowledge, the only authors who test whether there is an asymmetric response about the self-reported fair wage. Because the study differs from ours, notably in terms of one-shot vs. repeated effort decisions and of the timeframe employed, we maintain as our null hypothesis that the wage-effort relationship is the same for workers whose wage is below vs. above their belief about what constitutes a fair wage.¹⁵

HYPOTHESIS 3: Short-term workers receiving a wage below their self-reported fair wage exhibit the same relationship between wage and effort as those receiving

¹⁴Logically, Hypothesis 2 is nested in Hypothesis 1, which asserts that effort, conditional on wage, is identical for short-term and long-term workers. However, Hypothesis 2 is of special interest because it captures any reciprocal behavior of the workers.

¹⁵Considering that negative reciprocity has been repeatedly found to be stronger than positive reciprocity with reference measures other than the self-reported fair wage, our hypothesis could also be one-sided.

a wage above their fair level ($de_I/dw_L = de_I/dw_H$ for all w_L and w_H , where $w_H > w_F > w_L$ and w_F is the fair wage).

Hypothesis 4 concerns whether any asymmetry in reciprocity below vs. above the fair wage for short-term workers continues to exist for long-term workers. It is possible that negative reciprocity dissipates over time while positive reciprocity does not. This could be the case if, for example, negative reciprocity is caused by a strong negative emotion, such as anger, that would decrease over time. In the absence of any prior evidence, our null hypothesis is that any asymmetry in reciprocity is unaffected by a longer delay.

HYPOTHESIS 4: Short-term and long-term workers exhibit the same relationship between wage and effort if (a) the wage they receive exceeds their reported fair wage, and if (b) the wage they receive is less than their reported fair wage. (That is, (a) $de_I/dw = de_D/dw$ for $w > w_F$, and (b) $de_I/dw = de_D/dw$ for $w < w_F$).

IV. Results

This section is organized in two parts. The first part presents an overview of the data, and the second reports the tests of our hypotheses.

IV.A. Summary of the Data

In Table 4.2, we summarize the average values of participants' decision variables for the full sample, as well as for the three individual experiments. We normalize wages and efforts in our full sample by dividing wages and efforts by their maximum possible values in each experiment.¹⁶ For the full sample and each individual experiment, the average effort of short- and long-term workers is close to the midpoint of the range of possible wage levels, which is well in excess of the minimum of zero. There are no significant differences in the short- and long-term average efforts in the full sample (t -test, p -value is 0.27). In Tucson-H where fair wages are reported, the wage workers view as fair on average is roughly three times as far from the minimum as the maximum possible wage, and exceeds the actual average wage. The fair wage is slightly, though significantly, higher for long-term than short-term workers (t -test, p -value is 0.04).

In Table 4.3, we present the average effort levels for each half of the wage

¹⁶We normalize the data by dividing wages and efforts by 4 in Tilburg-L and Tucson-H, and wages by 2 and efforts by 2.4 in Tilburg-H.

Table 4.2: Summary of Participants' Average Decisions

Experiment	Full Sample	<i>N</i>	Tilburg-L	<i>N</i>	Tilburg-H	<i>N</i>	Tucson-H	<i>N</i>
Wage	0.600 (0.303)	144	2.067 (1.288)	46	1.218 (0.683)	33	2.617 (1.014)	65
Short-Term Effort	0.232 (0.285)	138	0.343 (0.595)	46	0.538 (0.727)	32	1.395 (1.217)	60
Long-Term Effort	0.274 (0.331)	128	0.343 (0.596)	42	0.610 (0.833)	30	1.704 (1.397)	56
Short-Term Fair Wage							3.134 (0.801)	65
Long-Term Fair Wage							3.413 (0.716)	64

Note: Effort and wage range from 0 to 1 in the Full Sample, and their range differs between individual experiments (see Table 4.1 for details). In Tilburg-L, effort levels are limited above at the wage received. Standard deviations are in parentheses.

range, for both types of workers.¹⁷ In the full sample, for short-term workers, an increasing relationship between wage and short-term effort appears. For long-term workers, the wage-effort relationship seems to be weaker. Specifically, wages in the first (lower) half of the wage range lead to greater effort for long- than for short-term workers. Moreover, the picture is similar for each individual experiment. This tendency for greater long-term effort in the lower half is less pronounced in Tilburg-L, where the effort choices of workers are bounded above by the wage offered, so that is impossible for workers with low wages to select high effort levels (e.g., workers with a wage of 0 can only choose an effort of 0). We can make similar observations from Figure 4.1, which illustrates the relationship between the wages workers receive and the effort they choose. The graphs contain the normalized data from the full sample, for short- and long-term workers separately. We highlight, in pale orange rectangles, the effort at wages 0 and 0.1 in Tilburg-L, which are bounded above at 0 and 0.1, respectively, in both the short and long term. We include the fitted regression lines in the figure.

IV.B. Evaluation of Hypotheses

Hypotheses 1 and 2

The first two hypotheses concern differences in effort between short- and long-term workers. To test for differences, we pool the normalized data of workers in our three experiments and estimate the following Tobit specification.¹⁸

¹⁷Appendix 4.A shows the same table for each quarter of the wage range; the pattern of effort is similar.

¹⁸We use a Tobit specification because there is a fair amount of lower censoring (35% of observations) for the effort variable at 0, and some censoring above (6%) at the maximum effort choice, as can be seen in Figure 4.1. We assume normalized effort to be censored below 0 and

Table 4.3: Average Effort by Wage Range, for Short- and Long-Term Workers

Wage Range	Short-Term Workers		Long-Term Workers	
	Average Effort	<i>N</i>	Average Effort	<i>N</i>
Full Sample				
0.00-0.50	0.085 (0.149)	53	0.211 (0.286)	50
0.51-1.00	0.324 (0.310)	85	0.314 (0.353)	78
Tilburg-L				
0.0-2.0	0.175 (0.439)	24	0.362 (0.571)	21
2.1-4.0	0.527 (0.694)	22	0.324 (0.634)	21
Tilburg-H				
0.0-1.0	0.145 (0.157)	11	0.782 (1.001)	11
1.1-2.0	0.743 (0.823)	21	0.511 (0.730)	19
Tucson-H				
0.0-2.0	0.622 (0.812)	18	1.128 (1.114)	18
2.1-4.0	1.726 (1.218)	42	1.976 (1.447)	38

Note: The numbers presented are averages. For the full sample, wage and effort are normalized so that both range from 0 to 1. For the individual experiments, the variables are not normalized. In Tilburg-L and Tucson-H, wage and effort range from 0 to 4, but effort is limited above by the wage in Tilburg-L. In Tilburg-H, wage ranges from 0 to 2 and effort from 0 to 2.4. The standard deviations are in parentheses.

$$e_j = \alpha + \beta wage_j + \theta wage_j D_j + \alpha_{TilbH} TilbH_j + \alpha_{TucH} TucH_j + TilbH_j + TucH_j + \gamma D_j + \epsilon_j. \quad (4.1)$$

In this equation, e_j denotes the effort of worker j , $wage_j$ is the wage that j has received, and D_j indicates a long delay (equals 0 for short-term workers and 1 for long-term workers). To simplify, we will usually write “delay” instead of “long delay”. There is also an interaction term between wage and the delay included in the specification. We include two dummies, for the Tilburg-H and Tucson-H experiments. The estimates are presented in column 1 of Table 4.4. In column 2, we include interaction terms between the wage and Tilburg-H and the wage and Tucson-H. In column 3, we add additional interaction terms between the wage and delay that are specific to Tilburg-H and Tucson-H.

The first hypothesis is that the time delay has no effect on effort for a given wage. That is, in terms of a regression of effort on wage, the delay affects neither the intercept (effort for the minimum wage) nor the slope (the wage-effort sensitivity). We assess this proposition by testing whether the coefficients of Wage x Delay and Delay in column 1 are both equal to zero. Note that our first hypothesis differs from claiming that the average effort is the same for both types of worker (as reported in the previous subsection, the average effort does not differ). This is because similar average effort can hide a pattern in which a long time delay causes effort to change in opposite directions for workers who

above 1 for simplicity in analyzing the data, even though in Tilburg-L effort is actually censored at the wage offered. We use robust standard errors in all regressions.

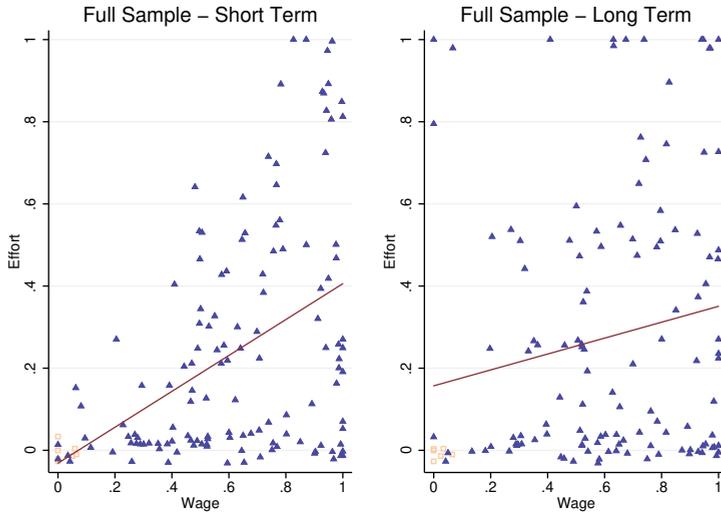


Figure 4.1: Scatter Plots of Short- and Long-term Efforts

Pale orange rectangles indicate efforts at wages 0 and 0.1 from Tilburg-L, which are limited above at 0 and 0.1, respectively

receive low wages and workers who receive high wages. We reject the hypothesis (F -test, p -value is 0.037). We obtain the same result for the coefficients of columns 2 or 3. This yields our first result.

RESULT 1: For a given wage, short- and long-term workers make significantly different effort choices.

We now consider the nature of the effect of the time delay on the behavior of short- and long-term workers. Our second hypothesis is that the wage-effort relationship does not change with delay. The coefficient of Wage is positive and significant (columns 1–3, p -values are < 0.001), replicating the common finding that higher wages lead to higher effort on the part of workers.¹⁹ The

¹⁹Though not shown here, short-term workers providing their effort earlier vs. later within the permitted time interval exhibit the same wage-effort relationship. That is, short-term workers responding within 15 or 30 minutes after the session show no differences in their wage-effort relationship compared to short-term workers responding later in the 3-hour interval. This suggests that they recall the wage they receive equally well later in the three hour interval as they do in the first few minutes. Because the long-term workers are reminded of their wage at the beginning of the window in which they make their decision, we take this as confirmatory evidence that the information that short-and long-term workers have at the time they choose their effort

Table 4.4: Estimates of Tobit Regression (4.1)

	Effort		
Wage	0.680**** (0.121)	0.672**** (0.169)	0.690**** (0.171)
Wage × D	-0.478** (0.209)	-0.480** (0.200)	-0.525** (0.223)
D	0.355** (0.138)	0.358*** (0.130)	0.349*** (0.129)
Wage × TilbH		-0.287 (0.242)	-0.247 (0.254)
Wage × TucH		0.305 (0.218)	0.235 (0.226)
Wage × D × TilbH			-0.082 (0.208)
Wage × D × TucH			0.161 (0.177)
TilbH	0.229*** (0.074)	0.410** (0.167)	0.408** (0.166)
TucH	0.416**** (0.059)	0.211 (0.139)	0.207 (0.137)
Constant	-0.529**** (0.093)	-0.522**** (0.115)	-0.514**** (0.114)
Prob. > F	0.000	0.000	0.000
Pseudo R ²	0.200	0.218	0.223
N	266	266	266

Note: N denotes the sample size, which includes both types of workers, split almost evenly. p -values: *0.10, **0.05, ***0.01, ****0.001.

coefficient of Wage × Delay in column 1 is, however, negative and significant (p -value is 0.02). As column 2 shows, this is robust to including wage terms for Tilburg-L and Tucson-H in the regression. This indicates that the long time delay decreases the marginal effect of the wage on effort provided. Moreover, we do not find significant reciprocity in the long term (for the F -tests of the restriction that Wage + Wage × Delay = 0 in columns 1–3, the p -values are > 0.23). In column 3, we can see that there are no significant differences in the effect of time delay on the wage-effort relationship among the three experiments (Wage × Delay × Tilburg-H, p -value is 0.70; Wage × Delay × Tucson-H, p -value is 0.36; p -value from the restriction that both equal zero is 0.44).

We supplement this parametric analysis with non-parametric tests. Table 4.5 presents the Spearman correlations between wage offered and effort chosen, for both short- and long-term workers, in the full sample. We perform Spearman’s rank correlation tests and indicate the p -values in square brackets. The short-term wage-effort correlation is positive and significant, and, unlike for the parametric analysis, the long-term wage-effort correlation is marginally

is similar.

Table 4.5: Spearman Correlation between Wage and Effort

	Full Sample	<i>N</i>
Short-Term Workers	0.467 [<0.001]	138
Long-Term Workers	0.173 [0.051]	128

Note: We indicate the two-sided *p*-values from Spearman’s rank correlation tests in square brackets.

significant. Nevertheless, in line with the rest of the analysis, the correlation is larger in magnitude in the short than in the long term (the long-term correlation amounts to 37% of the short-term one).

We report, in Appendix 4.A, that the wage-effort relationship also weakens over time if we use an approach based on wage ranges instead of the wages themselves. As the descriptive statistics suggest, we show that this decrease is driven by workers who receive wages in the lower half of the range of possible wages. That is, long-term workers offered low wages provide more effort than short-term workers offered the same low wages.²⁰ This pattern provides the basis for our second result.

RESULT 2: The wage-effort relationship is stronger within three hours, than four weeks after, workers learn their wage. Specifically, low wages generate less shirking in the long term.

Hypotheses 3 and 4

Our two remaining hypotheses apply only to our third experiment, Tucson-H, and concern the asymmetry in reciprocity about the fair wage. Our sample of workers consequently decreases from 266 to 116 for the analysis. There is therefore a relatively small number of workers falling in the categories that we compare to one another.

We analyze the behavior of workers as a function of their wage and an unfairness indicator variable *I* that denotes whether their wage is below their self-reported fair wage. The indicator takes value 1 if the wage is below fair levels, and 0 otherwise. In order to give an idea of the wage-effort relationship as well as the unfairness indicator-effort relationship, Table 4.6 shows the Spearman correlations between *I* and effort, separately for short- and long-term work-

²⁰We note that the equitable payoff—which we take as an equal payoff to the employer and to each employee, once the employer has decided on a wage—is achieved here when the employer offers a standardized wage of approximately 0.5. This corresponds to offering half of the maximum possible wage, which is 1. This means that those short- and long-term workers who are offered a wage below the wage generating the equitable payoff are the ones choosing different effort levels. That is, long-term workers treated inequitably provide greater effort than short-term workers.

Table 4.6: Spearman Correlation between Wage and Effort as well as between the Unfairness Indicator I and Effort

	Wage	I	N
Short-Term Workers	0.463 [<0.001]	-0.188 [0.151]	60
Long-Term Workers	0.298 [0.026]	-0.210 [0.121]	56

Note: The unfairness indicator I takes value 1 if Fair Wage $>$ Wage and value 0 otherwise. We indicate the two-sided p -values from Spearman's rank correlation tests in square brackets.

ers. We conduct Spearman's rank correlation tests and present the p -values in square brackets. We make two observations. First, the wage-effort correlation is positive and significant for both short- and long-term workers, although the magnitude of the short-term correlation is greater (the long-term correlation corresponds to 64% of the short-term one). Second, for both short- and long-term workers, the unfairness indicator-effort correlation is negative, indicating that offering wages below the fair wage decreases effort, but the relationship is not statistically significant.

To take into account both the effect of the absolute wage and the effect of the wage being below the fair level, we estimate the following Tobit regression.

$$e_j = \alpha + \beta_1 wage_j + \theta_1 wage_j D_j + \alpha + \beta_2 I_j + \theta_2 I_j D_j + \gamma D_j + \epsilon_j. \quad (4.2)$$

The unfairness indicator variable I_j captures whether worker j receives a wage below the fair wage.²¹ We include a variable for the wage, and to detect differences between the short and the long term, we interact the wage and the indicator with the time delay D_j .²²

The estimates are presented in columns 1–3 of Table 4.7. Column 1 provides estimates from the regression including only the wage variables. Column 2 includes the two variables containing the unfairness indicator, but excludes the wage variables. Column 3 employs the full regression specification.

²¹Instead, we could analyze the behavior of workers as a function of their wage and whether their wage is below their reported fair wage by a sufficiently large amount. The cutoff that corresponds to the Fair Wage-Effort hypothesis is 0 USD, i.e., any wage that is below fair levels. For the purpose of our analysis, choosing higher cutoffs might allow us to decrease the collinearity between the wage and the indicator variable for whether the wage is below the fair wage. However, the Pearson correlation between the wage and whether the wage is below fair levels remains near 0.65 for a wide range of cutoffs (i.e., 0.50, 1.00, 1.50 or 2.00 USD). Consequently, choosing cutoffs of 0.50, 1.00, 1.50 or 2.00 USD would provide the same Results 3 and 4 as the cutoff of 0 USD. Note that, among short-term workers, 42 receive a wage below the fair wage and 23 get a wage above, and among long-term workers, 44 have a wage below the fair wage and 20 have a wage above. Moreover, 26 short-term workers receive a wage at least 1 USD below the fair wage and 39 have a wage above, and 31 long-term workers have a wage at least 1 USD below the fair wage, while 33 get a wage above.

²²Note that we do not include a variable for the self-reported fair wage among the explanatory variables because it would be highly correlated with our unfairness indicator.

Table 4.7: Estimates of Tobit Regressions (4.2)

	Effort		
Wage	0.918**** (0.233)	1.274**** (0.329)	
Wage \times D	-0.312 (0.355)	-0.743* (0.444)	
I		-0.638 (0.453)	1.011 (0.613)
I \times D		-0.365 (0.784)	-1.254 (0.888)
D	1.240 (0.916)	0.647 (0.677)	3.207* (1.664)
Constant	-1.213** (0.596)	1.647**** (0.385)	-2.812** (1.213)
Prob. $> F$	0.000	0.149	0.000
Pseudo R ²	0.055	0.016	0.062
<i>N</i>	116	116	116

Note: The unfairness indicator I takes value 1 if Fair Wage $>$ Wage and value 0 otherwise. Standard errors are in parentheses. p -values: *0.10, **0.05, ***0.01, ****0.001.

In column 1, we find a significant positive wage-effort relationship (wage, p -value is $<$ 0.001), but, unlike for the full sample, the decrease in this relationship over time is not significant (wage \times D, p -value is 0.382). This difference can be explained by the loss in statistical power compared to the full sample.²³ In column 2, the effect of a wage below the fair wage is negative but insignificant (I, p -value is 0.162), and the attenuation of the effect is also not significant (I \times D, p -value is 0.643).

To evaluate Hypothesis 3—that short-term workers would exhibit the same wage-effort relationship for wages below and above their self-reported fair wage—we use column 3, as it allows us to isolate the effect of the wage being below fair levels. We find that the wage has a significant impact (wage, p -value is $<$ 0.001), but that having a wage below the fair wage in itself does not decrease effort significantly (the effect of I) column 3 is actually positive, though only close to marginally significant, p -value is 0.102).²⁴

We refer again to column 3 in order to examine Hypothesis 4—that the asymmetry between negative and positive reciprocity would be the same for short- and long-term workers. We see that the signs of the two coefficients

²³For instance, simulations show that if we draw 10,000 random subsamples of 116 workers (58 short-term and 58 long-term workers, without replacement) from our full sample, we detect a decrease in reciprocity in 27% of them at the 5% significance level, and in 42% of them at the 10% significance level.

²⁴In view of the fact that negative reciprocity has been found to be stronger than positive in several other studies, using a one-sided p -value instead would also be reasonable here, and this would make the difference very close to significant at the 5% level.

of the interactions with time delay suggest an attenuation in the effect of the wage, as well as the effect of the wage being below fair levels. The decrease in the effect of the wage is marginally significant ($\text{Wage} \times \text{D}$, p -value is 0.097), but the change over time in the effect of having a wage below what is considered fair is not significant ($\text{I} \times \text{D}$, p -value is 0.161). Results 3 and 4 describe our findings concerning hypotheses 3 and 4.

RESULT 3: For short-term workers, the sensitivity of effort to wage is similar for wages below fair levels than for those above.

RESULT 4: The difference in the sensitivity of effort to wage between wages below and above fair levels does not differ significantly between short-term and long-term workers.

In Appendix 4.B, we present an exploratory analysis regarding the relationships between workers' wages, their emotions, and their subsequent effort choices, using data on emotional state obtained with Facereading software. Our results are suggestive of a negative relationship between workers' anger and their effort levels in the short term, though not the long term.²⁵

V. Conclusion

Akerlof (1982) formalizes the Gift-Exchange version of the Efficient Wage hypothesis, in which employers offer wages that are higher than the market-clearing wage and workers reciprocate higher wages with greater effort. While much support for the existence of reciprocal provision of effort by workers has been gathered in short-term interactions, studies of whether reciprocal behavior persists over longer time periods have generally found that it persists within a work day, but decreases afterwards (Gneezy and List, 2006; Bellemare and Shearer, 2009; Hossain and List, 2012; Kube et al., 2012, 2013; Cohn et al., 2015; Gilchrist et al., 2016). Here, we consider whether the reciprocity workers exhibit in their effort choices differs under a brief vs. a long delay (within three hours vs. after four weeks) from the time at which they learn their wage. We test whether a positive wage-effort relationship is observed in an environment that is more decontextualized (in the laboratory) and simpler (a one-shot interaction) than previous studies concerned with reciprocity over long time

²⁵In addition, Appendix 4.C details employers' expectations of short-term and long-term effort choices.

periods. This allows us to clearly isolate any decrease in reciprocal behavior engendered by the passing of time over a long period.

In our third, more extensive experiment, we also consider one aspect of the Fair Wage-Effort hypothesis of Akerlof and Yellen (1988) and Akerlof and Yellen (1990) that might differ across timeframes: the contention that workers shirk in retaliation for wages below the level that they perceive as fair, but do not reward wages greater than the fair level with more effort. The existence of this asymmetry has been supported by the field experiment reported by Cohn et al. (2015). We examine whether this asymmetry exists in our setting, both in the short and long term.

Our main finding is that the strength of the reciprocal wage-effort relationship observed for short-term workers (who choose their effort within three hours of learning the wage) decreases for long-term workers (who wait one month before choosing their effort). We also find that workers offered low wages are the ones exhibiting different behavior over time. That is, long-term workers with low wages choose greater effort than short-term workers with the same low wages.

We obtain two additional findings from our third experiment that fail to support the Fair Wage-Effort hypothesis. The first is that, once we take into account the effect of the wage, workers do not reciprocate significantly more strongly for a wage that falls below the self-reported fair wage. The second is that this failure of the Fair-Wage effort hypothesis is present under both the short and long delays. Because of the reduced sample size on which these last two findings rest, they should not be taken as absolutely definitive. The signs of the estimated coefficients point in the direction of stronger reciprocation for wages below fair levels, as well as an attenuation of this asymmetry for long term workers, but the effects are below conventional thresholds of significance. Future research can shed more light on the presence of these two relationships.

Overall, our findings add to the evidence gathered by previous studies suggesting that workers' reciprocity weakens over long time periods. Moreover, we report that the magnitude of the wage appears to influence the effect of delays: wages need to be low enough for reciprocity to change over time. For employers, our results imply that examination of the timing of the task at hand relative to when wages are determined can benefit those trying to induce high effort from their employees.

Appendix

4.A. Supplemental Information Regarding Wage and Effort

Table 4.8 provides the average effort levels for each quarter of the wage range, for both types of workers. The pattern mirrors the one in Table 4.3, which reports the same exercise using halves instead of quarters. The wage-effort relationship appears to weaken over time, mostly because long-term workers with wages in the first (lowest) quarter of the wage range choose greater efforts. Note again that in Tilburg-L the effort choices of workers are bounded above by the wage offered, so that workers with very low wages cannot select high effort levels.

We present additional analyses regarding the effect of delay on the wage-effort relationship, to supplement the analyses we present in the main text for Result 2. In column 1 of Table 4.9, we show estimates from a Tobit regression that is similar to regression (4.1) in Table 4.4. The specification has separate terms for each experiment, and we do not use normalized data because separate estimates are made for each experiment.²⁶ For robustness, and to further identify which wages result in different effort choices over time, we present estimates from Tobit regressions using bins for different wage ranges in columns 2 and 3. In these regressions, we employ the normalized data. In column 2, we use a dummy variable for a wage being in the upper (2nd) half of the wage range, as well as two terms interacting separately the wage in the lower (1st) half and the wage in the upper (2nd) half of the wage range with the long delay. We do the same with quartiles of the wage range in column 3. This allows us to identify which wages are affected by delay.

The estimates allow us to make a number of observations. First, in column 1, we can see the individual effect of time delay on the wage for each experiment. Compared to our full sample, this substantially reduces the number of workers in each analysis (88 in Tilburg-L, 62 in Tilburg-H, and 116 in Tucson-H). While

²⁶Consequently, for column 1, unlike for the other columns, we only account for censoring below 0, as non-normalized efforts are censored above at different values across experiments.

Table 4.8: Average Effort by Wage Range, for Short- and Long-Term Workers

Wage Range	Short-term Workers		Long-term Workers	
	Average Effort	<i>N</i>	Average Effort	<i>N</i>
Full Sample				
0.00-0.25	0.022 (0.056)	25	0.203 (0.341)	23
0.26-0.50	0.142 (0.182)	28	0.218 (0.236)	27
0.51-0.75	0.267 (0.246)	44	0.280 (0.342)	40
0.76-1.00	0.385 (0.361)	41	0.351 (0.364)	38
Tilburg-L				
0.0-1.0	0.015 (0.038)	13	0.018 (0.041)	11
1.1-2.0	0.364 (0.607)	11	0.740 (0.648)	10
2.1-3.0	0.480 (0.718)	10	0.130 (0.200)	10
3.1-4.0	0.567 (0.702)	12	0.500 (0.834)	11
Tilburg-H				
0.0-0.5	0.100 (0.127)	6	1.380 (1.226)	5
0.6-1.0	0.200 (0.187)	5	0.283 (0.382)	6
1.1-1.5	0.491 (0.592)	11	0.278 (0.559)	9
1.6-2.0	1.020 (0.976)	10	0.720 (0.828)	10
Tucson-H				
0.0-1.0	0.167 (0.408)	6	1.000 (1.000)	7
1.1-2.0	0.850 (0.880)	12	1.209 (1.221)	11
2.1-3.0	1.443 (0.947)	23	1.871 (1.420)	21
3.10-4.0	2.068 (1.435)	19	2.106 (1.513)	17

Note: The numbers presented are averages. In the full sample, wage and efforts both are normalized to range from 0 to 1. In the individual experiments, wage and effort are not normalized. In Tilburg-L and Tucson-H, wage and effort range from 0 to 4, but effort is limited above by the wage in Tilburg-L. In Tilburg-H, wage ranges from 0 to 2 and effort from 0 to 2.4. Standard deviations are in parentheses.

the three Wage \times Delay coefficients are negative, the effect is only significant in Tilburg-H (*p*-values are 0.25 in Tilburg-L, 0.02 in Tilburg-H, and 0.28 in Tucson-H).²⁷

Second, we consider the effect of delay on effort in the different wage bins. In column 2, we observe that a wage in the upper half of the wage range increases effort (*p*-value < 0.001). Delay also increases effort for wages in the lower half (*p*-value is 0.01), but does not decrease effort for wages in the upper half (*p*-value is 0.57). Moreover, the two changes are significantly different from each other (*F*-test, *p*-value is 0.02), confirming that reciprocity is weaker in the long-term. In column 3, we see that a wage in any of the three upper quarters of the wage range increases effort compared to a wage in the first (lowest) quarter (*p*-values are all < 0.01). The time delay increases effort for wages in the first quarter (*p*-value is 0.02), but leads to no significant differences in effort for wages in any other quarter. An *F*-test does not quite reject the restriction that the four coefficients for delay are all equal (*p*-value is 0.09), so that

²⁷We only find significant reciprocity for long-term workers in Tucson-H (*F*-tests on the restriction Wage + Wage \times Delay = 0, separately for each experiment; *p*-value for Tucson-H is 0.02).

Table 4.9: Estimates of Tobit Regression of Effort on Wage and Delay

	Effort		Effort		Effort
Wage	0.494****	Wage 2nd Half	0.340****	Wage 2nd Quarter	0.314****
× TilbL	(0.145)		(0.069)		(0.116)
Wage	0.770***	Wage 1st Half	0.221**	Wage 3rd Quarter	0.459****
× TilbH	(0.251)	× Delay	(0.086)		(0.108)
Wage	0.799****	Wage 2nd Half	-0.038	Wage 4th Quarter	0.610****
× TucH	(0.192)	× Delay	(0.068)		(0.118)
Wage × Delay	-0.250			Wage 1st Quarter	0.363**
× TilbL	(0.219)			× Delay	(0.155)
Wage × Delay	-1.121**			Wage 2nd Quarter	0.135
× TilbH	(0.472)			× Delay	(0.097)
Wage × Delay	-0.310			Wage 3rd Quarter	-0.023
× TucH	(0.289)			× Delay	(0.085)
Delay	0.614			Wage 4th Quarter	-0.057
× TilbL	(0.512)			× Delay	(0.103)
Delay	1.379**				
× TilbH	(0.643)				
Delay	1.153				
× TucH	(0.755)				
TilbH	0.670		0.236***		0.238***
	(0.483)		(0.073)		(0.073)
TucH	0.492		0.426****		0.422****
	(0.606)		(0.060)		(0.059)
Constant	-1.369****		-0.329****		-0.518****
	(0.362)		(0.074)		(0.110)
Prob. > F	0.000		0.000		0.000
Pseudo R ²	0.132		0.176		0.203
N	266		266		266

Note: N denotes the sample size containing both types of workers, split almost evenly by type (see Table 4.1 for more detail). Standard errors are in parentheses. p -values: *0.10, **0.05, ***0.01, ****0.001.

we cannot clearly detect a change in reciprocity over time with this regression specification. Overall, however, our more powerful specifications show that the effect of wages on worker's effort decreases with the long delay, and this is due to higher effort for wages in the lower half of the wage range.

In addition, Table 4.10 provides the non-parametric Spearman correlations between wage offered and effort chosen, for both short- and long-term workers, in each experiment separately. We perform Spearman's rank correlation tests on each of the correlation coefficients and indicate the p -values in square brackets. The patterns are similar as those in the full sample. The short-term wage-effort correlations are positive and significant in each of the individual experiments. The long-term wage-effort correlation is significant only in Tucson-H. Again, the correlations are always smaller in magnitude in the long than in the short term.

Table 4.10: Spearman Correlation between Wage and Effort for Both Types of Worker in Each Experiment

	Tilburg-L	<i>N</i>	Tilburg-H	<i>N</i>	Tucson-H	<i>N</i>
Short-Term Workers	0.420 [0.004]	46	0.438 [0.012]	32	0.463 [<0.001]	60
Long-Term Workers	0.162 [0.306]	42	-0.078 [0.681]	30	0.298 [0.026]	56

Note: We indicate the two-sided p -values from Spearman's rank correlation tests in square brackets.

4.B. Analysis of the Relationship between Wages, Emotions, and Effort, using FaceReader Software

Another feature of the design of Tucson-H compared to the two other experiments is that we videotape participants for the entire session, with their consent. The videotapes are then analyzed with Noldus FaceReader™ software, which tracks facial expressions and analyzes the emotions they express. The software tracks facial movements and associates specific muscle movements to the six basic universal emotions catalogued by Paul Ekman and his colleagues (e.g. Ekman et al., 1987; Ekman and Friesen, 2003). The emotions are happiness, fear, anger, disgust, surprise, and sadness. FaceReader also measures how closely a facial expression conforms to a neutral state and generates an overall measure of emotional valence. FaceReader has been employed in a number of experimental economic studies focusing on emotions (Nguyen and Noussair, 2014; Breaban and Noussair, 2018; Fiala and Noussair, 2017; Van Leeuwen et al., 2018), and has also been used in marketing (Teixeira et al., 2012; Lewinski et al., 2014) and psychology research (Chentsova-Dutton and Tsai, 2010; Gadea et al., 2015).

We consider whether there is a relationship between wages and resulting worker emotions, as well as between worker emotion and effort provision, and whether these relationships differ between the two types of worker. Specifically, we are concerned with anger and overall emotional valence (net positivity). Emotional valence is calculated by the FaceReader software as the intensity of happiness minus the intensity of the strongest negative emotion, i.e., the maximum intensity score of anger, disgust, fear, and sadness. To calculate the change in an emotion in response to learning one's wage, we take the difference between the average intensity of an emotion over the 45 seconds after a worker learns her wage and the average intensity in the 45 seconds before learning the wage. We use this time interval because many participants take some time before they look and then read the text revealing their wage on their computer. We drop any participant whose video recording malfunctioned due to mechanical problems or whose video cannot be analyzed by the FaceReader

software for at least 40% of the 45 seconds both before and after learning the wage. Doing otherwise would result in several participants having relatively few data points from which to create an average. We thus have to leave out 54% of the participants. This leaves 53 observations: 26 short-term workers and 27 long-term workers.²⁸

To test whether changes in emotions are related to effort, while accounting for the multiple emotions and censoring of effort at 0, we estimate the following Tobit regression.

$$e_j = \alpha + \beta M_j(1 - D_j) + \theta M_j D_j + \gamma D_j + \epsilon_j. \quad (4.1)$$

where M_j includes either changes in the intensity of emotional valence or individual emotions. The regression includes all workers, but allows for a distinct effect of each emotion on short- and long- term workers' effort. The estimates are given in Table 4.11. Columns 1–3 report estimates including only emotions, while 4-5 contain anger and the wage as well. Demographic controls such as age, gender, and nationality might all be related to emotional reactions. We have estimated the same equations including these demographic variables but do not present the estimates here, because they usually do not affect the significance of the anger coefficient. However, we do point out when including the demographic variables changes the results.²⁹

Column 1 shows that the estimate of the effect of emotional valence on both short- and long-term effort is not significant when we include it in the regression without any of the specific emotions. In column 2, we show that, when anger is the only emotion in the regression, its coefficient is negative and significant for short-term workers (p-value is 0.02), and near zero and insignificant

²⁸We used the MuCap software package (Doyle and Schindler, 2015) to merge the z-Tree and FaceReader output for analysis. Several participants could not be analyzed since their cameras stopped recording because of malfunctions. Others did not conform to the FaceReader requirements for analysis (usually that the face must be in full view). Many participants started the experiment seated such that the camera could record their whole face, but then leaned down or slouched until a part of their face was out of view. A few others were covering their face with their hands at the crucial moment at which they received their wage. We also excluded five workers who did not choose an effort level. The 53 workers whose videos could be analyzed by the software were offered the same average wage (2.61, SD=1.03) as the 63 workers whose videos could be analyzed (2.63, SD=0.97; *t*-test-p-value is 0.89). The two groups were also not significantly different in terms of age and gender.

²⁹Non-parametric Spearman correlations between wages, anger and effort levels show similar patterns as those observed using the regressions. That is, short-term workers become angrier the lower the wage they receive, and for them, greater anger is associated with lower subsequent effort. For long-term workers, lower wages do not generate significantly more anger, and anger is not associated with subsequent effort.

Table 4.11: Estimates of Tobit Regressions of Effort on Emotions (4.1)

		Effort				
Short-Term Workers						
	Valence	1.152				
		(1.367)				
	Anger	-3.594**	-4.736**	-2.092	-13.703**	
		(1.491)	(1.941)	(1.318)	(5.463)	
	Sadness		-1.013			
			(1.362)			
	Disgust		3.613			
			(3.080)			
	Fear		5.506			
			(5.802)			
	Surprise		1.858			
			(1.820)			
	Happiness		-18.232			
			(12.737)			
	Wage			0.980**	0.641	
				(0.444)	(0.400)	
	Wage × Anger				4.732**	
					(2.106)	
Long-Term Workers						
	Valence	2.283				
		(2.393)				
	Anger	0.644	-0.405	1.093	-13.882	
		(4.847)	(6.553)	(4.259)	(11.780)	
	Sadness		-2.022			
			(5.175)			
	Disgust		-1.754			
			(3.804)			
	Fear		-12.037			
			(7.752)			
	Surprise		-0.860			
			(3.223)			
	Happiness		3.871			
			(9.493)			
	Wage			0.469	0.269	
				(0.432)	(0.460)	
	Wage × Anger				4.991	
					(4.281)	
	Delay	0.625	0.252	0.173	1.650	1.200
		(0.635)	(0.621)	(0.697)	(1.624)	(1.594)
	Constant	1.276***	1.443****	1.827****	-1.199	-0.080
		(0.388)	(0.347)	(0.452)	(1.093)	(0.923)
	Prob. > F	0.505	0.130	0.303	0.039	0.066
	Pseudo R ²	0.014	0.022	0.064	0.056	0.076
	N	53	53	53	53	53

Note: Each emotion has a separate intensity score from 0 to 1, and valence ranges from -1 to 1. Standard errors are in parentheses. *p*-values: *0.10, **0.05, ***0.01, ****0.001.

for long-term ones (*p*-value is 0.90). We cannot, however, statistically distinguish the two coefficients from each other (*F*-test, *p*-value is 0.41). This might be due to the high standard errors for the estimates of anger for long-term workers. A similar picture appears in column 3 when we include all emotions:

the coefficient of anger is negative and significant only for short-term workers (p-value is 0.02 for short-term workers and 0.95 for long-term workers), and no other emotion has a significant effect. In column 4, including wage and anger together yields a significant wage coefficient only for short-term workers, and renders short-term anger insignificant (p-value is 0.12).³⁰ This suggests that at least some of the relationship between anger and effort in other specifications is a consequence of the effect of wages on the anger level. We include the interaction of wage and anger in column 5. In the short term, wage is insignificant (p-value is 0.12), anger is significant (p-value is 0.02), and the interaction term is positive and significant (p-value is 0.03).³¹ In the long term, none of the three coefficients is significant.

Although the number of observations is limited, this analysis indicates that in the short term, the detrimental effect of low wages on effort is stronger if workers become angry when learning their wage. Furthermore, we clearly observe that among the emotions we measure, only anger has a consistent association with effort decisions. The effect only appears for short term workers.

We also study the effect of wages on emotions. To do this, we regress each emotion on wage individually using OLS.³² The estimates are shown in Table 4.12. The main finding is that the wage has a significantly negative impact on the anger level of short-term workers (p-value is < 0.01) and no impact for long-term workers (p-value is 0.77). The difference between the two is significant (F -test, p-value is 0.04), and long-term workers are also less angry when offered a wage of zero (p-value is 0.04). These two last results do become marginally insignificant, however, when demographics are included. Furthermore, a higher wage significantly increases the valence of both types of workers and the happiness of long-term workers.

³⁰Including demographics makes short-term anger significant (p-value is 0.04). The anger variable is subject to potential endogeneity problems when included in the same equation as the wage, as both effort and anger are correlated with the wage offered. Though not presented here, using the wage as an instrument for anger in a Two-Stage Least Squares regression also yields a significant negative coefficient of anger for short-term workers, which is about twice as large as the coefficient reported in column 2 of Table 4.11.

³¹If we include demographic controls, the interaction term becomes insignificant (p-value is 0.09) and wage and anger are unaffected. However, using a one-sided p-value instead is reasonable in this instance, as we expect anger to either have no effect or reinforce the sensitivity to the wage, and this would make the effect significant at the $p < .05$ level.

³²We use a dummy variable for long-term workers, and separate the emotions of short- and long-term workers, to obtain a separate estimate of the effect of each emotion on the effort of each type of worker workers. We do not include demographics in the specification.

Table 4.12: Effect of Wage on Emotional Changes in Workers

	Anger	Sadness	Disgusted	Fear	Surprised	Happiness	Valence
Wage	-0.089***	-0.024	-0.041	0.003	0.045	0.007	0.132***
× Short Term	(0.032)	(0.024)	(0.038)	(0.005)	(0.042)	(0.005)	(0.045)
Wage	-0.006	0.010	-0.052	-0.009	-0.046	0.018**	0.082**
× Long Term	(0.021)	(0.022)	(0.033)	(0.007)	(0.049)	(0.009)	(0.036)
Long Term	-0.246**	-0.086	0.073	0.054*	0.286	-0.054	0.089
	(0.116)	(0.090)	(0.141)	(0.029)	(0.207)	(0.033)	(0.170)
Constant	0.304***	0.057	0.087	-0.011	-0.183	-0.011	-0.383**
	(0.105)	(0.073)	(0.110)	(0.009)	(0.130)	(0.014)	(0.144)
Prob. > F	0.064	0.761	0.148	0.328	0.569	0.075	0.002
R ²	0.125	0.013	0.099	0.048	0.064	0.134	0.187
N	53	53	53	53	53	53	53

Note: Each emotion has a separate intensity score from 0 to 1, and valence ranges from -1 to 1. Standard errors are in parentheses. *p*-values: *0.10, **0.05, ***0.01, ****0.001.

4.C. Analysis of Employers' Predictions

Table 4.13: Tobit Regression of Employers' Effort Predictions on Short- and Long-term Wage Offered, and OLS Regression of Employers' Difference in Effort Predictions on the Wage Offered

	Predicted Short-term Effort	Predicted Long-term Effort		Predicted Long- Minus Short-term Effort
Wage	0.541****	0.397****	Wage	-0.144**
	(0.085)	(0.110)		(0.069)
Tucson-H	0.046	0.017	Tucson-H	-0.023
	(0.045)	(0.058)		(0.040)
Constant	0.144**	0.227***	Constant	0.070
	(0.064)	(0.076)		(0.044)
Prob. > F	0.000	0.002	Prob. > F	0.040
Pseudo R ²	1.353	0.213	R ²	0.046
				0.050
N	98	98	N	98

Note: Standard errors are in parentheses. *p*-values: *0.10, **0.05, ***0.01, ****0.001.

In Tilburg-H and Tucson-H, employers are asked to predict the choices of both the short- and the long-term workers whom they are matched with. One of the two predictions is selected to count toward their payment. In Tilburg-H, they earn 1 EUR if the prediction selected to count is within 0.2 units of the actual effort chosen. In Tucson-H, they receive 1 USD if their selected prediction is within 0.4 of the actual effort.³³

We explore whether employers correctly anticipate the weaker wage-effort relationship for long-term workers in the Tilburg-H and Tucson-H experiments. First, we regress the effort predictions of employers on the wages offered by the employers using a Tobit regression, for short- and long-term workers separately.³⁴ We include a dummy for the Tucson-H experiment. We present the

³³This methodology is not incentive compatible because participants have an incentive to hedge. It is also not incentive compatible for beliefs close to the minimum or maximum. We use it because it gives us a simple, although rough idea of the predictions of employers.

³⁴We use normalized wages and effort predictions.

estimates in columns 1-2 of Table 4.13. Employers who offer higher wages anticipate significantly higher effort, both in the short and long term (p-values are <0.001), although the relationship is flatter in the long-term. Second, we regress the difference between the long-term and short-term effort predictions of employers (predicted long-term effort minus predicted short-term effort) on the wages that employers offered, using an OLS regression. The estimates are shown in columns 3 and 4. We include a dummy for Tucson-H in column 4. The wage term is negative and significant (p-values are 0.04 in column 1 and 0.047 in column 2). This means that, as wage offers increase, expected long-term effort decreases relative to short-term effort.

4.D. Instructions (for Tucson-H)

We present the instructions for the Tucson-H experiment. The instructions are very similar for the other two experiments. They are available on request.

Introduction

Welcome. The camera on top of your computer screen is currently turned on and you will be filmed for the entire session. The video will be seen only by the experimenters and will be destroyed once data have been gathered from it. The content of the video will not affect your payment. If you disagree with being filmed, you are free to leave at any time, and you will keep the show-up fee, but you will not receive additional payment.

In this experiment, the instructions are simple and if you read them carefully you can earn a considerable amount of money. Everything you earn is on top of a 5 dollar show-up fee and a 3 dollar participation fee. You are forbidden to communicate with other participants at any time. At the end the experiment, the earnings you have made will be paid to you by sending you cash or a check in the mail. Take a moment now to write down your name and your mailing address on the envelope in front of you. Please be careful, as we will use this envelope to send your payment. This experiment takes around 40 minutes in the lab. In addition, depending on the role you are randomly assigned, you might have to take 5 minutes outside the lab to send an email, either in the 3 hours after the session today or in 4 weeks. If anything is unclear during the experiment, please raise your hand and the experimenter will come to you.

Roles

In this experiment, participants are randomly assigned, with equal chance, to 1 of the 3 roles: 1) employer, 2) employee choosing now or 3) employee choosing later. You will be placed in a group of 3 made of 1 employer, 1 employee choosing now, and 1 employee choosing later. The identity of the members of each group will be kept secret from all participants. Your role will be now revealed to you on your screen. Do not reveal your role to others and do not mention your role aloud to the experimenter.

General Instructions

This is an experiment about employer-employees interaction. First, the employer chooses a wage between 0 and 4 dollars for his/her 2 employees and asks them to choose a specific effort level between 0 and 4. The wage is the same for each employee. The employer cannot force the employees to choose a specific effort level. Second, the employees learn their wage and each employee chooses an effort level between 0 and 4. The employee choosing now will choose an effort level within 3 hours after the end of today's session and the employee choosing later will choose an effort level in 4 weeks. A higher effort level costs an employee more but benefits the employer more. Each unit of effort chosen costs the employee 25 cents and increases the earnings of the employer by 2.50 dollars. Employees receive the wage, no matter what effort level they choose. In addition, the employer is asked to make a prediction about the effort of each employee. Also, before learning his/her wage, each employee is asked to state the wage he/she thinks would be fair to receive.

Specific Instructions for the Employer

The employer chooses a wage between 0 and 4 dollars for each of the two employees. The wage can be chosen in any increment of 10 cents. The wage is the same for both employees. The employer has to pay the wage twice because there are two employees, one choosing an effort today and one choosing an effort in 4 weeks. The wages are costly for the employer and benefit the employee. The wages are deducted from the employer's earnings. The wage for one employee is added to that employee's earnings. The employer also asks for a specific effort level from each employee. The specific effort asked is the same for each employee. The employer cannot force an employee to choose a spe-

cific effort. Each employee is free to choose an effort level that is the same or different from the level that is asked. Each unit of effort chosen by an employee will cost him/her 25 cents and pay the employer 2.50 dollars. Your payment calculation is described in the section Payment Calculation for All Roles.

In addition, the employer is asked to predict the effort level of each employee. This prediction will not be revealed to any participant. Once both employees have chosen their effort level, the prediction for one employee will be randomly chosen. If this prediction is within 0.4 of the effort chosen by the employee, the employer will receive an additional 1 dollar.

The payment to the employer will be sent by mail on the day after the receipt of the effort level chosen by the employee choosing in 4 weeks. The employer has no email to send. If an employee does not choose an effort level, the employer is paid back the wage that was given to the employee.

Specific Instructions for the Employee Choosing Now

The employer will offer each employee a wage between 0 and 4 dollars and ask for a specific effort level between 0 and 4. The wage chosen is deducted from the employer's earnings. The employer cannot force an employee to choose a specific effort, he/she can only ask for it. The employee will then have to choose an effort between 0 and 4. The effort can be chosen in any increment of 0.1. The effort is costly for the employee and benefits the employer. Each unit of effort chosen by an employee will cost him/her 25 cents and pay the employer 2.50 dollars. Your payment calculation is described in the section Payment Calculation for All Roles.

The employee choosing now chooses his/her effort level within the 3 hours after the end of today's session. He/she does not choose the effort level during the session. He/she chooses the effort level by sending an email to the experimenter. The email address of the experimenter is written on the last page of these instructions. You can leave with the last page of the instructions after today's session. Thus you can also write anything you might need on the last page. The envelope with your payment will be mailed on the day after the receipt of your email. Please respect the 3 hour deadline as it will be enforced.

If you send your effort choice more than 3 hours after the end of this session, you will only be paid your show-up and participation fees.

In addition, before learning his/her wage, each employee will be asked what is the wage that he/she would consider fair to receive. The wage can be any-

thing between 0 and 4 dollars, in any increment of 10 cents. This fair wage will not be revealed to any participant in the experiment. It cannot influence the wage that an employee receives. Please simply try to be truthful even if it does not influence your earnings.

Specific Instructions for the Employee Choosing Later

The employer will offer each employee a wage between 0 and 4 dollars and ask for a specific effort level between 0 and 4. The wage chosen is deducted from the employer's earnings. The employer cannot force an employee to choose a specific effort, he/she can only ask for it. The employee will then have to choose an effort between 0 and 4. The effort can be chosen in any increment of 0.1. The effort is costly for the employee and benefits the employer. Each unit of effort chosen by an employee will cost him/her 25 cents and pay the employer 2.50 dollars. Your payment calculation is described in the section Payment Calculation for All Roles.

The employee choosing later chooses his/her effort level only in 4 weeks after today's session. He/she does not choose the effort level during the session. He/she chooses the effort level by sending an email to the experimenter. The email address of the experimenter is written on the last page of these instructions. You can leave with the last page of the instructions after today's session. Thus you can also write anything you might need on the last page. The date in 4 weeks is also written on the last page. You will have 5 days to send your effort then, with the starting day being the date in 4 weeks. The envelope with your payment will be mailed on the day after the receipt of your email. Please respect the 5 day period beginning in 4 weeks. If you send your choice of effort before or after the 5 day period, you will only be paid your show-up and participation fees. An email reminder will be sent to you on the starting day in 4 weeks. The wage, the effort asked by the employer, and a copy of these instructions will be in the email reminder.

In addition, before learning his/her wage, each employee will be asked what is the wage that he/she would consider fair to receive. The wage can be anything between 0 and 4 dollars, in any increment of 10 cents. This fair wage will not be revealed to any participant in the experiment. It cannot influence the wage that an employee receives. Please simply try to be truthful even if it does not influence your earnings.

Payment Calculation for All Roles

Earnings in dollars for participants in each of the 3 roles are described by these formulas:

Employer: $11 - 2 \times \text{wage} + 2.5 \times (\text{effort of employee choosing now} + \text{effort of employee choosing later})$

Employee choosing now: $4 + \text{wage} - 0.25 \times \text{effort of employee choosing now}$

Employee choosing later: $4 + \text{wage} - 0.25 \times \text{effort of employee choosing later}$

Note that the 5 dollar show-up fee, the 3 dollar participation fee, and the potential payment for the employer's predictions are not included. Please take a minute to study these formulas carefully as your earnings depend on understanding them.

Practice

During the experiment, you can choose only once. Before you start, please practice with 2 examples on the last page. The examples will not affect your earnings. When you have completed them, please raise your hand and the experimenter will verify your answers. This is the best time to ask questions to the experimenter to be sure that you understand. Please do not hesitate to ask. Your earnings depend on your understanding of the instructions.

At the end of the session, when you are done with the questionnaire, please stay seated and raise your hand. The experimenter will come to you and ask you to sign a sheet before you can leave. Please do not communicate with others during that time and leave your envelope at your seat. You are free to leave with the last page of the instructions, on which the email address is written.

You can leave with this page after today's session. Feel free to write anything on it.

Example 1: The employer chooses a wage of U , one employee chooses effort T , and the other employee chooses effort L . Do not replace the variables with numbers. Please write down the earnings for each role:

Employer:

Employee choosing effort T :

Employee choosing effort L :

Example 2: The employer chooses a wage of Z , one employee chooses effort Q , and the other employee chooses effort S . Do not replace the variables with numbers. Please write down the earnings for each role:

Employer:

Employee choosing effort Q :

Employee choosing effort S :

Date today: xx—xx—2016

Date in 4 weeks: yy—yy—2016

Wage can be anything between 0 and 4, in any increment of 0.1.

Effort level can be anything between 0 and 4, in any increment of 0.1.

Email to send choice of effort: XXXXXXXXXXX@XXXXX

4.E. Email Reminders for Long-term Workers

For conciseness, we present here only the email reminder for long-term workers in the Tucson-H experiment. A copy of the original instructions was attached. The email reminders are similar across experiments. The reminders for the two other experiments are available on demand.

Dear participant,

4 weeks ago, you were assigned to the role of **Employee Choosing Later** in an experiment at the Economic Science Lab of the University of Arizona. You had to send an email to this address with your choice of effort level 4 weeks after the session.

This is the **reminder** that the date 4 weeks after is today, **[date]**. You have now 5 days to **send your choice of effort**. These 5 days include today.

During the session, you were placed in a group of 3 participants consisting of an employer, an employee choosing now, and an employee choosing later. The employer had to choose a wage between 0 and 4 dollars, in any increment of 10 cents, for each employee. The wage was the same for both employees. The wages are deducted from the employer's earnings.

The employer offered you a wage of \$[wage] during the session and asked for an effort level of [effort]. You now have to choose your effort level. Note that the employer cannot force you to choose a specific effort level.

You can choose an effort level between 0 and 4, in any increment of 0.1. Each unit of effort that you choose will cost you 25 cents, and pay the employer 2.50 dollars.

Payment Calculation (in dollars):

Employer: $11 - 2 \times \text{Wage} + 2.5 \times (\text{Effort of Employee Choosing Now} + \text{Effort of Employee Choosing Later})$

Employee Choosing Now: $4 + \text{Wage} - 0.25 \times \text{Effort of Employee Choosing Now}$

Employee Choosing Later: $4 + \text{Wage} - 0.25 \times \text{Effort of Employee Choosing Later}$

If anything is unclear, please contact me. I also attach a **copy of the instructions** for the experiment. The envelope with your payment will be sent by mail on the day after the receipt of your email.

You must respect the deadline to be paid.

Chapter 5

Unequal Chances: Ex Ante Fairness and Individual Control

Abstract

Unequal financial outcomes often originate from unequal chances. Yet, compared to outcomes, little is known about how individuals perceive unequal distributions of chances. We investigate empirically the role of different sources of unequal chances in shaping inequality perceptions. Importantly, we do so from an ex ante perspective—i.e., before the chances are realized—which has rarely been explored. In an online survey, we asked uninvolved respondents to evaluate ex ante the fairness of unequal allocations of chances. We varied the source of inequality of chances, using a comprehensive range of factors which resemble several real world situations. Respondents also evaluated how much control individuals hold over the distribution of chances. Results show that the source of inequality has considerable effects on the ex ante perception of fairness and control. Unequal chances based on socioeconomic and biological factors, such as gender, family income and ethnicity, are evaluated to be both unfair and beyond one's control. Allocations based on factors judged to be under higher individual control, such as effort and knowledge, are perceived as fairer. Luck is the exception, ranking as high in fairness as effort and knowledge, but similarly low in individual control as ethnicity, family income, and gender.

This short chapter is co-authored with Leticia Micheli.

I. Introduction

Unequal chances abound in the pursuit of financial rewards; erecting the landscape of the financial outcomes likely to be achieved by different groups of individuals. The striking recent study of Chetty et al. (2018) reports that in the United States “controlling for parental income, black boys have lower incomes in adulthood than white boys in 99% of Census tracts.” Policies such as additional investments for low-income children, diversity quotas in firms, and Affirmative Action in university admissions strive to redress these uneven chances before they translate into unequal outcomes. They attempt to “level the playing field” as the saying goes. Such policies can be observed across the globe. Brazil invests additional funds in the education of low-income children, through its Bolsa Família cash-transfer program. Norway as well as other European countries require a gender quota on corporate boards. Universities in India, South Africa, and the United States increase the admission chances of disadvantaged groups through Affirmative Action. Nevertheless, our understanding of how individuals perceive the unequal chances that these policies address is minimal. In fact, there is a noticeable empirical gap in the existing literature on redistribution: we have limited knowledge about which types of unequal chances are considered to be fair from the ex ante standpoint—i.e., before chances are realized and incomes are known. We employ ourselves to bridge this gap.

Using a survey, we investigate the fairness perceptions of different procedures used to distribute unequal chances, from the ex ante perspective of an uninvolved spectator. We use a large set of procedures, which include factors often serving as bases for unequal chances in society such as effort, knowledge, luck, gender, ethnicity, and family income. Those factors are more comprehensive than those that the literature on fairness has employed, enabling a more fine-grained exploration of differences among a range of situations that can generate unequal chances. Moreover, we examine the perception of individual control over these procedures and the relationship between control and ex ante fairness. This is important because work in distributive justice has emphasized the divide between factors perceived as within and beyond one’s control in the equity perception of different types of inequality. Lastly, we touch upon Affirmative Action, which raises the odds for disadvantaged groups over advantaged groups. We do this by testing whether it is fairer to provide higher chances to groups that are generally disadvantaged than to groups that are generally advantaged. Our hope is that understanding how individuals perceive unequal chances originating from different factors can provide insights to policymakers

in terms of which unequal chances are perceived as justifiable and which may call for rectification.

The economic literature investigating inequality and redistribution decisions draws a line between *ex post* and *ex ante* fairness. While *ex post fairness* concerns the distribution of incomes between individuals once these incomes are known (Harsanyi, 1955), *ex ante fairness* applies to the distribution of chances that are to generate incomes (Diamond, 1967). The literature in psychology and management (Cohen-Charash and Spector, 2001; Tyler, 2000; Leventhal et al., 1980; Thibaut and Walker, 1975) also studies the procedure generating outcomes separately from the outcomes themselves and uncovers crucial factors that influence individuals' appraisal of procedures, such as respect, an opportunity to be involved or voice their concerns, and the absence of biases.

The empirical literature on redistribution focuses mostly on *ex post* fairness, once incomes are already known. Studies show that a substantial proportion of individuals display concern for equality of incomes, exhibiting a willingness to decrease their own income to improve the conditions of others through redistribution (Bolton and Ockenfels, 2000; Charness and Rabin, 2002; Dawes et al., 2007; Fehr and Schmidt, 1999; Forsythe et al., 1994). There is also empirical evidence that individuals care about *ex ante* inequality, although the topic is less studied. For instance, recent economic experiments have presented their participants with unequal chances to win at a lottery, and an opportunity to redistribute chances before the realization of that lottery. Many individuals indeed show concern for equality when redistributing chances, even if it reduces their own chances of winning (Brock et al., 2013; Krawczyk and Le Lec, 2010). Chances also influence the redistributive decisions of uninvolved spectators (Cappelen et al., 2013; Cettolin and Riedl, 2016). Moreover, there is some evidence that individuals do consider the distribution of initial chances in their redistribution decisions even after the realization of these chances (Bolton et al., 2005; Grimalda et al., 2016). Finally, theoretical work on social preferences for obtaining a fair chance has been the focus of a number of papers in the past twenty years (Bolton et al., 2005; Karni and Safra, 2002; Krawczyk, 2011; Saito, 2013; Trautmann, 2009).

Studies investigate the fairness perception and the distribution decisions of incomes that have been generated through different procedures, but they do so almost exclusively from an *ex post* perspective. For example, they examine whether the patterns of fairness perceptions and income distribution decisions differ if income inequality is caused by different factors, such as effort, knowl-

edge, and luck (Cappelen et al., 2007; Gächter and Riedl, 2005; Schokkaert and Devooght, 2003; Konow, 2000). Typically, a sizeable proportion of individuals accept more inequality caused by effort and knowledge than luck. The prominent explanation for this type of behavior is the view that individuals should be held accountable for factors under their control, such as effort and knowledge, but not for others which they cannot influence, such as luck. This is the position taken by the liberal egalitarian fairness ideal (Roemer, 1998) and the accountability principle (Konow, 2000). Several earlier equality concepts in philosophy also hold individuals responsible for their choices, the difference often being to which extent one is responsible for certain factors and not for others, which is notoriously difficult to determine (Arneson, 1989; Cohen, 1989; Dworkin, 1981).

In contrast, how different factors behind the distribution of chances affect the ex ante fairness perception of these chances has scarcely been explored by empirical studies. We believe that this missing step is socially valuable because many policies—such as the Affirmative Action described in the first paragraph—intend not to correct incomes once unequal chances have been realized, but to correct different types of unequal chances before they are realized. Furthermore, whether the positive relationship between perceived control and fairness of inequality documented ex post is typical from an ex ante perspective has not been researched much. Studying ex ante perceptions is essential because our knowledge regarding ex post fairness views does not necessarily inform us about ex ante fairness views. For instance, a society's fairness evaluation of a procedure based on luck may differ ex ante and ex post: providing equal chances of obtaining either a high or low income to everyone might be judged as fair at the outset, only to be considered unfair once unequal incomes have been randomly allocated.

The only study that has empirically investigated the role of different factors in ex ante fairness views compares how distributions of chances based on height versus a combination of effort and knowledge impact individuals' ex ante choice of future redistribution (Krawczyk, 2010). The study finds that, ex ante, individuals are more prone to support future redistribution when height influence chances than when effort and knowledge does, which provides some evidence that liberal egalitarianism is present ex ante. We add to this research by investigating a wider range of factors underlying unequal distribution of chances, which are mostly chosen to resemble real-world situations of inequality. Moreover, we study individuals' fairness perceptions in combination with control

perceptions, rather than redistribution decisions. We note that the previous study also included questions regarding fairness perceptions of procedures, but did not find significant differences between the reported fairness of the height and the effort and knowledge procedures. However, as the author reports, its use of a between-subject design might generate too much noise to distinguish fairness differences between the procedures. Our within-subject design is likely to alleviate this issue.

II. Methods

We recruited 82 undergraduate students from the online recruitment system for economic experiments (ORSEE) (Greiner, 2015) of the Behavioral and Experimental Economics Laboratory (BEELab) at the School of Business and Economics of Maastricht University (The Netherlands) to take part in an online Qualtrics survey. We obtained informed consent from participants to take part in the study. They completed the survey in April and May 2017 and received 5 EUR for their participation. They were 41% male and on average 21.5 years old. The study was conducted following the peer-approved procedures established by Maastricht University's Behavioral and Experimental Economics Laboratory (BEELab). Our study was approved by the Maastricht University's Behavioral and Experimental Economics Laboratory at a public ethics review and project proposal meeting that is mandatory for all scholars who wish to use the BEELab facilities. All methods were performed in accordance with the relevant guidelines and regulations of the BEELab and Maastricht University for studies with human participants.

Participants considered a hypothetical scenario in which unequal chances to win a prize of 10 EUR are allocated between two other individuals. Participants could not win the prize themselves, and, for simplicity, we framed chances as lottery tickets. Chances are always allocated in the following manner: one individual has a 90% chance, and the other has a 10% chance. Only one of the two individuals wins the prize. Participants were asked questions concerning different procedures determining who receives which chance. The scenario was described to participants in the following manner:

Consider a situation in which lottery tickets are distributed between two persons. Generally we ask you to assume that these two persons have a similar social, economic and educational background. In some cases this will not be so and you will be explicitly informed about it. Each ticket gives a 10% chance to win the

lottery. There can be only one winner in the lottery, who receives 10 EUR as a prize. The ticket allocation is as follows: one person gets 9 tickets and thus has a 90% chance of winning the lottery prize, whereas the other person gets 1 ticket and thus has a 10% chance of winning the lottery prize. The allocation of tickets is based on various procedures.

We then presented 11 different procedures in random order. Although all procedures presented led to the same consequences (i.e., one person receiving a 10% chance to win a single lottery prize and the other receiving a 90% chance of winning that same prize), these procedures differed with respect to the factor that determines who receives which chances. The factor presented was one of the following: benevolence, knowledge, effort, luck, student identification number (ID), height, weight, nationality, gender, ethnicity (skin color), and family (parental) income.

All procedures were presented once in our “Regular” Block, where the higher chance was allocated to individuals who were benevolent, more knowledgeable, exerted more effort, lucky, with a higher ID, tall, heavy, male, German, light-skinned, or from a high-income family. We also presented participants with the same procedures in our “Reverse” Block, this time allocating higher chances to the opposite individuals (e.g., individuals who are non-benevolent, dark-skinned, women or from a low-income family). We are more interested in the “Regular” Block because the procedures in this block described cases that generally occur more often in real life, such as advantaging men, light-skinned individuals, or individuals from a high-income family. For our analysis, we are only interested in a select group of these procedures to evaluate asymmetries in fairness, which we explain below. The order of blocks was randomized.

Table 5.1 summarizes the criteria for receiving higher chances, in each procedure of each block. For benevolence, we presented the case of an individual who decided to help a person in need and another individual who decided to *not* help a person in need. In this specific case, higher chances are allocated to the benevolent person in the “Regular” Block, while higher chances are allocated to the non-benevolent person in the “Reverse” Block. For knowledge, we used a general knowledge quiz, where the person with a higher number of correct answers would receive higher chances in the “Regular” Block and lower chances in the “Reverse” Block. For effort we employed a slider task (Gill and Prowse, 2012), a common effort task in experimental economics, which is explained in the example below. Again, the person exerting more effort would re-

ceive higher chances in the “Regular” Block and lower chances in the “Reverse” Block. Then, for luck we adopted a 6-side die roll where one person would roll a number equal or smaller than 3 and receive higher chances, whereas the other would receive lower chances. In the ID procedure, we used greater and smaller student ID numbers (which are generally allocated by order of enrolment, although this might be unknown to students). For height we contrasted taller and shorter individuals, while for weight we contrasted heavy and light persons and for gender we employed women versus men. For nationality, we made use of Germans versus non-Germans (German is the most common nationality in the participants pool). For ethnicity, we contrasted light-skinned to dark-skinned individuals. Finally, for family income we considered individuals coming from higher or lower income families.

Participants first rated the fairness of all procedures for both the Regular and Reverse Blocks. After that, participants were exposed a second time to both blocks, again in random order, and rated the same procedures in terms of how much individual control one has over chances received. Responses were given on a 7-point Likert scale ranging from *Very Unfair* to *Very Fair* and from *No Control* to *Full Control*. The full instructions are provided in the Appendix. The following two examples regarding effort and gender describe how the procedures were presented to participants. Both examples belong to the Regular Block.

Effort Example: Chances are allocated based on the results of a task that depends on effort. In the task, people see dots appearing randomly on horizontal bars on their computer screen. Their task is to position as many dots as possible in the middle of the horizontal bars. The person who places more dots in the correct position in a given time receives a 90% chance to win, and the person who places less dots in the correct position receives a 10% chance to win.

Gender Example: Chances are allocated based on gender. The two persons have different genders. The man receives a 90% chance to win, and the woman receives a 10% chance to win.

Table 5.1: Criteria for receiving higher chances in the Regular and Reverse Blocks

Procedure	Regular Block	Reverse Block
Benevolence	Choosing to <i>help</i> someone in need	Choosing <i>not to help</i> someone in need
Knowledge	<i>Higher</i> results in a knowledge quiz	<i>Lower</i> results in a knowledge quiz
Effort	Correctly position <i>more</i> dots in the task	Correctly position <i>less</i> dots in the task
Luck	Roll of 6-faced die gives <i>3 or less</i>	-
ID	<i>Higher</i> student ID number	<i>Lower</i> student ID number
Height	<i>Taller</i> person	<i>Shorter</i> person
Weight	<i>Heavier</i> person	<i>Lighter</i> person
Gender	<i>Male</i>	<i>Female</i>
Nationality	<i>German</i>	<i>Non-German</i>
Ethnicity	<i>Lighter</i> skin color	<i>Darker</i> skin color
Family Income	<i>Higher</i> parental income	<i>Lower</i> parental income

Note: We did not reverse Luck because this would simply mean changing the winning numbers of the die roll, which should have no effect on its fairness and control ratings.

III. Hypotheses

We test whether there are differences in fairness among procedures, and whether fairness is a correlate of the perception of control. Our first hypothesis is that procedures will be perceived differently in terms of fairness and individual control based on the factors used to decide who would receive higher and lower chances.

Our second hypothesis is that procedures with low individual control are perceived as more unfair, in line with liberal egalitarianism.

Furthermore, we test for asymmetries in fairness ratings when chances favor one group over the other for the following procedures: benevolence, knowledge, effort, gender, ethnicity, and family income. We expect that asymmetries in these cases will arise for two reasons. First, previous research shows that individuals reward “merit” in the case of effort and knowledge, and this might extend to benevolence. This can lead to individuals finding it fairer to allocate higher chances to those who were benevolent, who provided more correct answers in the knowledge quiz and who exerted more effort than to those with opposite characteristics. Second, one might find it fairer to favor generally disadvantaged groups because this compensates existing inequality. Hence, our third hypothesis is that participants rate procedures as fairer when females, individuals with dark skin and low parental income receive higher chances than males and individuals with white skin and high parental income.

IV. Results

RESULT 1, EX ANTE FAIRNESS RATINGS: *Regular-block procedures advantaging individuals who are benevolent, knowledgeable, exert effort and are lucky, score*

high on fairness, whereas procedures advantaging individuals based on socioeconomic or biological factors (e.g., gender, height, skin color or parental income) score low on fairness.

Table 5.2 presents the mean fairness ratings of Regular-block procedures. We use non-parametric Wilcoxon signed-rank tests to analyze pairwise differences between procedures in their distribution of fairness ratings ($N=82$). Stars in a cell indicate that the procedure in the column is fairer than the corresponding row. P-values are corrected for multiple comparisons using the Holm-Bonferroni method (Holm, 1979) and are indicated by the number of stars. For completeness, Table 5.5 in the Appendix also shows fairness ratings of the procedures in the Reverse Block.

This result indicates that not all chance inequalities are considered equally fair from the ex ante viewpoint, which is in line with what other studies have found for the ex post perspective. Indeed, despite the fact that all procedures led to an unequal distribution of chances between individuals, some procedures were considered fairer than others. Unlike for the ex post perspective, however, procedures giving higher chances to those who exert effort or are knowledgeable are not significantly fairer than the luck procedure, which simply randomly allocates chances. Procedures based on socioeconomic and biological criteria are the ones which really differ from luck, by being substantially less fair.

Overall, our first result suggests that individuals may be prone to accept, ex ante, certain inequality of chances, e.g., when they come through effort, knowledge, or even luck, while endorsing policies that rectify others, e.g., when they rest on biological and socioeconomic factors. Next, we test whether individuals' judgments of ex ante fairness differ depending on who is advantaged by the unequal distribution of chances.

RESULT 2, ASYMMETRIES IN EX ANTE FAIRNESS RATINGS: *It is fairer to advantage individuals who are benevolent, dark-skinned, knowledgeable, and exert more effort, compared to advantaging individuals with the opposite characteristics. It is not fairer to advantage women over men.*

In most cases, as hypothesized, results show that ex ante fairness ratings differ depending on who is favored with higher chances. The exception is gender, for which, unlike what we hypothesized, we did not find significant asymmetries.

Table 5.2: Mean Fairness and Wilcoxon Signed-Rank Test Comparisons

Procedure	Benev.	Effort	Luck	Know.	ID	Height	Weight	Nat.	Inc.	Gender	Ethn.
Benevolence (Benev.)	5.38 (1.47)										
Effort		5.04 (1.53)									
Luck	*		4.63 (1.77)								
Knowledge (Know.)	***	**		4.38 (1.88)							
ID	***	***	***	***	2.17 (1.67)						
Height	***	***	***	***	**	1.37 (.75)					
Weight	***	***	***	***	**		1.35 (.60)				
Nationality (Nat.)	***	***	***	***	***			1.33 (.85)			
Family Income (Inc.)	***	***	***	***	***				1.29 (.60)		
Gender	***	***	***	***	***					1.21 (.46)	
Ethnicity (Ethn.)	***	***	***	***	***	**	**				1.13 (.34)

Note: Standard deviations are indicated in parentheses. Stars in a cell indicate that the procedure in the column is fairer than the procedure in the corresponding row. Holm-Bonferroni-corrected p -values: * $p < .05$; ** $p < .01$; *** $p < .001$.

Table 5.3 reports the mean fairness for the procedures advantaging one group (Regular Block) versus the other group (Reverse Block). Non-parametric Wilcoxon signed-rank tests ($N=82$) are used to compare procedures in the Regular and Reverse Blocks. Stars in the table indicate whether the distributions of the procedures are significantly different. The asymmetries that we identify can occur for different reasons. First, asymmetries in fairness among procedures that involve some sort of merit (benevolence, effort, knowledge) can occur because individuals find it fair to distribute chances in a meritocratic manner. Second, individuals may consider it fairer to give higher chances to disadvantaged groups similar to what Affirmative Action policies often do.

We make two observations. First, finding no asymmetries in the case of gender might be due to the specific setting that we study—a Dutch business school with a large share of women among its student body, which for students is likely not highly unequal in terms of gender. Second, although we do find asymmetries in fairness, procedures that favor disadvantaged groups are still generally considered unfair compared to procedures that advantage those who are benevolent, exert effort, or are knowledgeable. This result reflects to some extent the mixed support affirmative action can receive. Even though individuals recognize that giving more chances to the disadvantaged is fairer, such procedures are still considered relatively unfair. Alternatively, the low relative fairness of the affirmative action procedures that we use might be

due to the somewhat egalitarian environment of the students. Another reason for this result is that these procedures create a large amount of inequality by distributing very unequal chances instead of slightly advantaging the generally-disadvantaged group.

RESULT 3, CONTROL AND FAIRNESS: *Individuals generally consider procedures that one has control over to be fairer. Interestingly, luck stands out as the exception: it ranks high in fairness but low in control.*

Table 5.4 details the mean control ratings of the Regular procedures. We calculate the non-parametric Spearman correlation between fairness and control for each individual ($N=79$), for the 11 procedures contained in the Regular Block. Note that we cannot calculate the correlation for 3 participants because there are no variations in their fairness or control answers. For the 79 remaining participants, the Spearman ρ is positive for 95% of them, at least of medium strength ($\rho \geq .30$) for 92%, and strong ($\rho \geq .60$) for 63%. The ρ is significantly greater than zero at the .05 significance level for 62% of participants, and at the .10 level for 71%, while it is significantly smaller than zero at the .10 level for none. For the interested reader, Table 5.6 in the Appendix also displays the mean control ratings for the Reverse Block.

Our third result indicates that a substantial share of participants exhibit a positive correlation between fairness and control when they make ex ante judgments about unequal distributions of chances. This is in line with what other studies have found for ex post judgments of unequal outcomes. That is, the judgments of many individuals appear broadly consistent with the liberal-egalitarian fairness ideal, whereby individuals find unequal chances based on factors over which one has control to be fair. However, the fact that the procedure randomly allocating high chances is as fair ex ante as the procedures advantaging those who exert effort or are knowledgeable—while at the same time being judged to be under no control—is inconsistent with the liberal-egalitarian fairness ideal.

Table 5.3: Asymmetries and Wilcoxon Signed-Rank Test Comparisons

Higher Chance to:	Mean Fairness	Difference
Benevolent	5.38 (1.47)	***
Non-Benevolent	1.93 (1.44)	
More Effort	5.04 (1.53)	***
Less Effort	2.52 (1.65)	
More Knowledgeable	4.38 (1.88)	***
Less Knowledgeable	2.17 (1.39)	
High Family Income	1.29 (.60)	***
Low Family Income	2.74 (1.79)	
Man	1.21 (.46)	
Woman	1.32 (.83)	
Ethnicity—Light Skin	1.13 (.34)	**
Ethnicity—Dark Skin	1.44 (.98)	

Note: Standard deviations are indicated in parentheses. Stars in a cell indicate a significant difference in fairness depending on who is advantaged. Holm-Bonferroni-corrected p -values: * $p < .05$; ** $p < .01$; *** $p < .001$.

Table 5.4: Mean Control and Wilcoxon Signed-Rank Test Comparisons

Procedure	Benev.	Effort	Know.	Weight	Luck	Inc.	ID	Gender	Nat.	Height	Ethn.
Benevolence (Benev.)	5.91 (1.51)										
Effort	***	5.22 (1.41)									
Knowledge (Know.)	***		4.98 (1.49)								
Weight	***	***	***	2.95 (1.65)							
Luck	***	***	***	***	1.55 (1.00)						
Family Income (Inc.)	***	***	***	***		1.50 (1.03)					
ID	***	***	***	***			1.45 (.72)				
Gender	***	***	***	***				1.28 (.76)			
Nationality (Nat.)	***	***	***	***					1.28 (.86)		
Height	***	***	***	***						1.27 (.74)	
Ethnicity (Ethn.)	***	***	***	***							1.26 (.70)

Note: Standard deviations are indicated in parentheses. Stars in a cell indicate that the procedure in the column is under more control than the procedure in the corresponding row. Holm-Bonferroni-corrected p -values: * $p < .05$; ** $p < .01$; *** $p < .001$

V. Discussion

This study addresses a gap in the existing literature on redistribution regarding how individuals evaluate the ex ante fairness of distributions of chances based on different factors. We exposed uninvolved participants to different procedures used to distribute unequal chances to win a monetary prize between two individuals. Before the chances were realized, participants were asked to rate the fairness of each procedure as well as to rate how much individual control each procedure allowed for.

Our results indicate that, as other studies have found for ex post inequality, the fairness of ex ante inequality depends on the causes that generate this inequality. Procedures advantaging those who are benevolent, knowledgeable, exert more effort or advantaging those who are simply lucky are rated similar in terms of fairness by uninvolved spectators. The assessment of luck differs from the ex post perspective, in which procedures that advantage individuals based on luck are often considered to be less fair than meritocratic procedures that advantage those who exert more effort or who are more knowledgeable (Cappelen et al., 2007; Gächter and Riedl, 2005; Konow, 2000).

A clear demarcation line between the different procedures only appears when we compare the fairness of the group of aforementioned procedures with the fairness of the group of procedures that are based on biological and socioeconomic factors. That is, the group of procedures based on gender, height, weight, nationality, ethnicity or family income are substantially less fair than the other procedures, which are based neither on biological nor socioeconomic factors. The difference between the observed two groups of procedures is in line with what we regularly witness in societies: unequal chances based on biological or socioeconomic characteristics are denounced as unfair and can elicit strong reactions, such as social movements against racism and sexism demanding equal chances for members of different groups. Generally, this result also reinforces the need to differentiate inequality and fairness (Starmans et al., 2017).

Moreover, the asymmetry observed in the fairness of certain procedures is consistent with individuals supporting, on average, Affirmative Action for usually disadvantaged ethnic and socioeconomic groups. Even in our setting of a simple lottery among university students, it is indeed fairer to favor certain disadvantaged groups. We observe, however, that increasing the chances of certain ethnic and socioeconomic groups is still considered unfair in this context relative to advantaging those who exert more effort or are more knowledge-

able. We do not find that it is fairer to give higher chances to one gender over the other, although we note that this might be the result of using students in a relatively gender-balanced environment. Overall, asymmetries in perceived fairness provide a potential source of expansion for current research.

Furthermore, procedures allocating chances in a fairer manner rate higher in individual control. This is broadly consistent with the liberal-egalitarian fairness ideal that a large share of individuals have been shown to hold in terms of ex post fairness (Cappelen et al., 2007). Interestingly, unlike ex post, liberal egalitarianism fails to hold ex ante in the case of luck, which is considered both fair and outside of one's control. On the one hand, the procedure randomly allocating higher chances is considered as fair as the procedures advantaging those who exert effort or are knowledgeable. On the other hand, whereas procedures based on effort or knowledge are judged to be under individuals' control, luck is obviously considered to be under no control, as are procedures based on biological or socioeconomic factors. This reveals a sharp difference between ex ante and ex post judgments.

How is it possible for the random allocation procedure to be fair from an ex ante perspective, and to then suddenly become unfair ex post? Previous theoretical work suggests that individuals can be time-inconsistent when making ex ante and ex post fairness judgements (Myerson, 1981; Trautmann and Wakker, 2010). For instance, they can first support policies that aim at ex ante chance equality, but lead to unequal incomes, and, once these unequal incomes are realized, they then support policies that aim at ex post income equality. This problem can arise when, at the moment of making an ex ante choice, individuals do not take into account possible ex post inequality. Our result on this matter provides indirect support for time-inconsistency, by showing that a random allocation procedure is ex ante fair, whereas other studies have found it to be unfair ex post. A recent working paper also reports behavior consistent with this phenomenon: individuals often make an ex ante fair choice before outcomes are realized, only to then switch to an ex post fair choice once outcomes are known (Andreoni et al., 2016).

In addition, we note that it can be puzzling that individuals evaluate chances allocated based on luck as fairer than chances allocated based on other factors that could be easily rationalized as coming from luck, such as height. We note that luck being perceived as fair ex ante and unfair ex post may provide an explanation for this phenomenon. That is, on the one hand, an allocation of chances based on height may be seen through ex post lenses, as the luck that

could determine height has already been realized. On the other hand, an allocation of chances based on luck may be seen as still not realized—i.e., as part of a compound lottery—and therefore be perceived through *ex ante* lenses. Another possibility is that the exception of luck might have arisen through social conventions. That is, although die rolls cannot be controlled, they could be perceived as fair because individuals often voluntarily opt to use them to settle indivisible claims, whereas individuals rarely settle claims based on height.

Overall, our results contribute to the literature on distributive justice in the following ways. First, we extend the investigation of how different factors causing inequality are perceived to the *ex ante* domain. By doing so, we indicate similarities and differences regarding how fairness ideals are perceived in the *ex ante* and *ex post* domains. Secondly, by exploring a range of factors that often serve as the basis for unequal chances in society, we show which inequalities in chances are perceived as fairer than others. Lastly, by comparing who is favored by unequal chances (i.e., groups that are generally advantaged or disadvantaged in society), we point out asymmetries in fairness perceptions which are consistent with the notion of Affirmative action.

Our findings carry two potential implications for policymakers. First, the causes of unequal distributions of chances need to be carefully considered if one wants to make chance-equalizing policies that correspond to the fairness perceptions of individuals. Second, considered with the results of other studies, our study suggests that individuals exhibit time-inconsistency regarding the fairness of different redistribution policies. That is, they may support policies that grant equal chances, but create income inequalities, only to later switch their support to policies that aim at equal incomes. A way to avoid this reversal is perhaps to conscientiously consider the *ex post* inequalities that equal chances may generate, whenever there are public debates on equal-chance versus equal-income policies.

We note that we do not directly investigate how the fairness perceptions that we observe translate into redistribution behavior. For some of our conclusions, we assume that they translate into such decisions. We leave it for future research to explore whether the many differences in fairness that we find actually transform into different redistribution decisions involving monetary stakes. The starting point for such an investigation is the work of (Krawczyk, 2010), discussed in the introduction, which finds that individuals redistribute more when chances are assigned based on height than on a combination of effort and knowledge. We note that this study explicitly interprets chances based

on height as being the same as chances based on luck—and then concludes that individuals redistribute differently if chances are based on luck or on effort/knowledge. Our results hint, however, that height might not be a correct proxy for luck. This is because we find that chances based on height are judged as much more unfair than chances based on luck, effort or knowledge, all of which are judged to be similarly fair.

Finally, we realize that while in our study participants were clearly aware of the reason behind inequality in chances, this is often not the case in real life. Individuals' misconceptions of the reasons why someone might have higher or lower chances might considerably affect their fairness judgments, and consequently their support for redistributive policies. It would be interesting to measure the contribution of these misconceptions to the opposition to and support for public policies that aim to establish ex ante equality, such as Affirmative Action.

Appendix

5.A. Instructions Fairness

Consider a situation in which lottery tickets are distributed between two persons. Generally we ask you to assume that these two persons have a similar social, economic and educational background. In some cases this will not be so and you will be explicitly informed about it. Each ticket gives a 10% chance to win the lottery. There can be only one winner in the lottery, who receives 10 EUR as a prize. The ticket allocation is as follows: one person gets 9 tickets and thus has a 90% chance of winning the lottery prize, whereas the other person gets 1 ticket and thus has a 10% chance of winning the lottery prize. The allocation of tickets is based on various procedures. We would like to know how fair or unfair you consider each of the following procedures to be. Please also note that you cannot win the lottery prize in question.

Questions Fairness—Regular Block

Note for the reader: the titles of each question—displayed in Italics—were not shown to participants.

How fair/unfair would you consider the distribution of chances to win the lottery prize based on the following procedure:

-Knowledge: Chances are allocated based on the results in a general knowledge quiz. The person with the higher results receives a 90% chance to win, and the person with the lower results receives a 10% chance to win.

-Effort: Chances are allocated based on the results of a task that depends on effort. In the task, people see dots appearing randomly on horizontal bars on their computer screen. Their task is to position as many dots as possible in the middle of the horizontal bars. The person who places more dots in the correct position in a given time receives a 90% chance to win, and the person who places less dots in the correct position receives a 10% chance to win.

-Luck: Chances are allocated by throwing a 6-faced die. If the die number is

smaller than or equal to 3, one person receives a 90% chance to win, and the other person receives a 10% chance to win. If the die number is greater than 3, the chances are reversed.

-Height: Chances are allocated based on height. The taller person receives a 90% chance to win, and the shorter person receives a 10% chance to win.

-Weight: Chances are allocated based on weight. The heavier person receives a 90% chance to win, and the lighter person receives a 10% chance to win.

-Gender: Chances are allocated based on gender. The two persons have different genders. The man receives a 90% chance to win, and the woman receives a 10% chance to win.

-ID: Chances are allocated based on student ID number. The person with the higher student ID number receives a 90% chance to win, and the person with the lower student ID number receives a 10% chance to win.

-Nationality: Chances between students at the School of Business and Economics of Maastricht University are allocated based on one's country of origin. The German student receives a 90% chance to win, whereas the non-German student receives a 10% chance to win.

-Ethnicity: Chances are allocated based on one's skin color. The two persons have different skin colors. The person with the lighter skin color receives a 90% chance to win, whereas the person with the darker skin color receives a 10% chance to win.

-Benevolence: Chances are allocated based on one's decision to help. The two persons both had an opportunity to help someone in need and only one decided to help. The person who helped receives a 90% chance to win, whereas the person who did not help receives a 10% chance to win.

-Family Income: Chances are allocated based on one's parental income. The person whose parents have a higher income receives a 90% chance to win, whereas the person whose parents have a lower income receives a 10% chance to win.

5.B. Instructions Individual Control

Consider the same situation in which lottery tickets are distributed between two persons. Generally we ask you to assume that these two persons have a similar social, economic and educational background. In some cases this will not be so and you will be explicitly informed about it. The allocation of tickets is based on various procedures. Now we would like to know how much control you consider people have over the distribution of chances for each of the following procedures. Again, there are no right or wrong answers. We are solely interested in your preferences and opinions. It is important for our research that you answer all questions honestly.

Questions Individual Control—Regular Block

How much control do you think people have over the chances they receive to win the lottery prize, when the distribution of chances is based on the following procedure.

Note for the reader: Procedures were described in the same manner as in the *Fairness Instructions*.

5.C. Additional Information Concerning the Reverse Block

Although we did not present Luck as a procedure in the Reverse Block, we present fairness and control ratings for Luck (obtained in the Regular Block) in the following tables for comparison purposes.

Table 5.5: Mean Fairness and Wilcoxon Signed-Rank Test Comparisons—Reverse Block

Procedure	Luck	Family Inc.	Effort	ID	Know.	Benev.	Nat.	Height	Ethn.	Weight	Gender
Luck	4.63 (1.77)										
Family Income (Inc.)	***	2.74 (1.79)									
Effort	***		2.52 (1.65)								
ID	***			2.33 (1.66)							
Knowledge (Know.)	***				2.17 (1.39)						
Benevolence (Benev.)	***		**			1.93 (1.44)					
Nationality (Nat.)	***	***	**	*	+		1.66 (1.21)				
Height	***	***	***	**	**			1.50 (.92)			
Ethnicity (Ethn.)	***	***	***	**	**				1.44 (.98)		
Weight	***	***	***	***	***					1.35 (.62)	
Gender	***	***	***	***	***	*					1.32 (.83)

Note: Standard deviations are indicated in parentheses. Stars in a cell indicate that the procedure in the column is fairer than the procedure in the corresponding row. Holm-Bonferroni-corrected p -values: + $p < .10$; * $p < .05$; ** $p < .01$; *** $p < .001$.

Table 5.6: Mean Control and Wilcoxon Signed-Rank Test Comparisons—Reverse Block

Procedure	Benev.	Effort	Know.	Weight	Luck	ID	Inc.	Height	Gender	Ethn.	Nat.
Benevolence (Benev.)	5.61 (1.85)										
Effort	***	4.90 (1.70)									
Knowledge (Know.)	***		4.63 (1.70)								
Weight	***	***	***	2.78 (1.66)							
Luck	***	***	***	***	1.55 (1.00)						
ID	***	***	***	***		1.46 (.88)					
Family Income (Inc.)	***	***	***	***			1.32 (.68)				
Height	***	***	***	***				1.29 (.76)			
Gender	***	***	***	***	+				1.23 (.73)		
Ethnicity (Ethn.)	***	***	***	***						1.23 (.65)	
Nationality (Nat.)	***	***	***	***	**	*					1.20 (.74)

Note: Standard deviations are indicated in parentheses. Stars in a cell indicate that the procedure in the column is under more control than the procedure in the corresponding row. Holm-Bonferroni-corrected p -values: + $p < .10$; * $p < .05$; ** $p < .01$; *** $p < .001$.

Valorization Addendum

The chapters contained in this thesis are intended to be useful for understanding the behavior of individuals in the contexts of labor markets and public policy. Chapter 2 will serve firms, unions, courts, and policymakers aiming to understand possible reactions of workers to unfair chances and, especially, to gender discrimination. Chapter 3 proposes a new perspective to policymakers concerned with the negative effect of ethnic diversity on support for redistributive policies. Chapter 4 provides food for thought for managers trying to elicit long-term effort from workers. Finally, chapter 5 is of interest to policymakers with a stake in understanding how individuals perceive different sources of unequal chances in terms of fairness.

Bibliography

- J. Abeler, S. Altmann, S. Kube, and M. Wibral. Gift exchange and workers' fairness concerns: when equality is unfair. *Journal of the European Economic Association*, 8(6):1299–1324, 2010.
- J. S. Adams. Inequity in social exchange. In L. Berkowitz, editor, *Advances in Experimental Social Psychology*, volume 2, pages 267–299. Academic Press, New York and London, 1965.
- G. A. Akerlof. Labor contracts as partial gift exchange. *Quarterly Journal of Economics*, 97(4):543–569, 1982.
- G. A. Akerlof and W. T. Dickens. The economic consequences of cognitive dissonance. *American Economic Review*, 72(3):307–319, 1982.
- G. A. Akerlof and J. L. Yellen. Fairness and unemployment. *American Economic Review*, 78(2):44–49, 1988.
- G. A. Akerlof and J. L. Yellen. The fair wage-effort hypothesis and unemployment. *Quarterly Journal of Economics*, 105(2):255–283, 1990.
- A. Alesina and E. L. Glaeser. *Fighting poverty in the US and Europe: A world of difference*. Oxford University Press, 2004.
- A. Alesina and E. La Ferrara. Ethnic diversity and economic performance. *Journal of Economic Literature*, 43(3):762–800, 2005.
- A. Alesina, R. Baqir, and W. Easterly. Public goods and ethnic divisions. *Quarterly Journal of Economics*, 114(4):1243–1284, 1999.
- A. Alesina, E. Glaeser, and B. Sacerdote. Why doesn't the US have a European-style welfare system? *NBER Working Paper No. 8524*, 2001.
- Y. Algan, C. Hémet, and D. D. Laitin. The social effects of ethnic diversity at the local level: A natural experiment with exogenous residential allocation. *Journal of Political Economy*, 124(3):696–733, 2016.

- G. W. Allport. *The nature of prejudice*. Addison-Wesley Reading, MA, 1954.
- J. G. Altonji and R. M. Blank. Race and gender in the labor market. In O. C. Ashenfelter and D. Card, editors, *Handbook of Labor Economics*, volume 3c, chapter 48, pages 3143–3259. North Holland, Amsterdam, 1999.
- J. Andreoni and J. M. Rao. The power of asking: How communication affects selfishness, empathy, and altruism. *Journal of Public Economics*, 95(7-8):513–520, 2011.
- J. Andreoni, D. Aydin, B. Barton, B. D. Bernheim, and J. Naecker. When fair isn't fair: sophisticated time inconsistency in social preferences. *Working Paper*, 2016.
- J. Andreoni, J. M. Rao, and H. Trachtman. Avoiding the ask: A field experiment on altruism, empathy, and charitable giving. *Journal of Political Economy*, 125(3):625–653, 2017.
- A. A. Arechar, S. Gächter, and L. Molleman. Conducting interactive experiments online. *Experimental Economics*, 21(1):99–131, 2018.
- R. J. Arneson. Equality and equal opportunity for welfare. *Philosophical Studies*, 56(1):77–93, 1989.
- K. Arrow. The theory of discrimination. In O. Ashenfelter and A. Rees, editors, *Discrimination in Labor Markets*, pages 3–33. Princeton University Press, Princeton, 1973.
- J. R. Averill. Studies on anger and aggression: implications for theories of emotion. *American Psychologist*, 38(11):1145, 1983.
- L. Balafoutas and M. Sutter. Affirmative action policies promote women and do not harm efficiency in the laboratory. *Science*, 335(6068):579–582, 2012.
- K. Baldwin and J. D. Huber. Economic versus cultural differences: Forms of ethnic diversity and public goods provision. *American Political Science Review*, 104(4):644–662, 2010.
- O. Bandiera and G. Levy. Diversity and the power of the elites in democratic societies: Evidence from indonesia. *Journal of Public Economics*, 95(11-12): 1322–1330, 2011.

- N. M. Baran, P. Sapienza, and L. Zingales. Can we infer social preferences from the lab? Evidence from the Trust Game. *NBER Working Paper No. 15654*, 2010.
- V. Bartoš, M. Bauer, J. Chytilová, and F. Matějka. Attention discrimination: theory and field experiments with monitoring information acquisition. *American Economic Review*, 106(6):1437–75, 2016.
- C. D. Batson, J. Chang, R. Orr, and J. Rowland. Empathy, attitudes, and action: Can feeling for a member of a stigmatized group motivate one to help the group? *Personality and Social Psychology Bulletin*, 28(12):1656–1666, 2002.
- P. Battigalli, M. Dufwenberg, and A. Smith. Frustration, aggression, and anger in leader-follower games. *Games and Economic Behavior*, 117:15–39, 2019.
- G. S. Becker. *The Economics of Discrimination*. University of Chicago Press, Chicago, 1957.
- C. Bellemare and B. Shearer. Gift giving and worker productivity: Evidence from a firm-level experiment. *Games and Economic Behavior*, 67(1):233–244, 2009.
- G. Ben-Shakhar, G. Bornstein, A. Hopfensitz, and F. Van Winden. Reciprocity and emotions in bargaining using physiological and self-report measures. *Journal of Economic Psychology*, 28(3):314–323, 2007.
- R. Bénabou and J. Tirole. Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–1678, 2006.
- Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B (Methodological)*, pages 289–300, 1995.
- M. Benz and S. Meier. Do people behave in experiments as in the field?—Evidence from donations. *Experimental Economics*, 11(3):268–281, 2008.
- M. Bertrand and E. Duflo. Field experiments on discrimination. In E. Duflo and A. Banerjee, editors, *Handbook of Economic Field Experiments*, volume 1, chapter 8, pages 309–393. North Holland, Amsterdam, 2017.
- M. Bertrand and S. Mullainathan. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013, 2004.

- M. Bertrand, D. Chugh, and S. Mullainathan. Implicit discrimination. *American Economic Review*, 95(2):94–98, 2005.
- M. Bertrand, C. Goldin, and L. F. Katz. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 2(3):228–255, 2010.
- M. Bertrand, E. Kamenica, and J. Pan. Gender identity and relative income within households. *Quarterly Journal of Economics*, 130(2):571–614, 2015.
- T. Bewley. *Why Don't Wages Fall in a Recession*. Harvard University Press, Cambridge, 1999.
- K. Bjorvatn and A. W. Cappelen. Inequality, segregation, and redistribution. *Journal of Public Economics*, 87(7-8):1657–1679, 2003.
- F. D. Blau and L. M. Kahn. The gender wage gap: extent, trends, and explanations. *Journal of Economic Literature*, 55(3):789–865, 2017.
- A. S. Blinder. Wage discrimination: reduced form and structural estimates. *Journal of Human Resources*, 8(4):436–455, 1973.
- H. Blumer. Race prejudice as a sense of group position. *Pacific Sociological Review*, 1(1):3–7, 1958.
- L. Bobo and V. L. Hutchings. Perceptions of racial group competition: Extending Blumer's theory of group position to a multiracial social context. *American Sociological Review*, pages 951–972, 1996.
- L. D. Bobo. Prejudice as group position: Microfoundations of a sociological approach to racism and race relations. *Journal of Social Issues*, 55(3):445–472, 1999.
- I. Bohnet and B. S. Frey. Social distance and other-regarding behavior in dictator games: Comment. *American Economic Review*, 89(1):335–339, 1999.
- J. A. Bohren, A. Imas, and M. Rosenberg. The dynamics of discrimination: Theory and evidence. *PIER Working Paper (18-016)*, 2018.
- J. Boisjoly, G. J. Duncan, M. Kremer, D. M. Levy, and J. Eccles. Empathy or antipathy? The impact of diversity. *American Economic Review*, 96(5):1890–1905, 2006.

- F. Bolle, J. H. Tan, and D. J. Zizzo. Vendettas. *American Economic Journal: Microeconomics*, 6(2):93–130, 2014.
- G. E. Bolton and A. Ockenfels. Erc: A theory of equity, reciprocity, and competition. *American Economic Review*, 90(1):166–193, 2000.
- G. E. Bolton, J. Brandts, and A. Ockenfels. Fair procedures: Evidence from games involving lotteries. *Economic Journal*, 115(506):1054–1076, 2005.
- P. Bordalo, K. Coffman, N. Gennaioli, and A. Shleifer. Stereotypes. *Quarterly Journal of Economics*, 131(4):1753–1794, 2016.
- P. Bordalo, K. Coffman, N. Gennaioli, and A. Shleifer. Beliefs about gender. *American Economic Review*, 109(3):739–73, 2019.
- R. Bosman and F. Van Winden. Emotional hazard in a power-to-take experiment. *Economic Journal*, 112(476):147–169, 2002.
- R. Bosman, J. Sonnemans, M. Zeelenberg, et al. Emotions, rejections, and cooling off in the ultimatum game. *Working Paper, University of Amsterdam*, 2001.
- A. Bracha, U. Gneezy, and G. Loewenstein. Relative pay and labor supply. *Journal of Labor Economics*, 33(2):297–315, 2015.
- P. Branas-Garza, R. Cobo-Reyes, M. P. Espinosa, N. Jiménez, J. Kovářik, and G. Ponti. Altruism and social integration. *Games and Economic Behavior*, 69(2):249–257, 2010.
- J. Brandts and G. Charness. Hot vs. cold: Sequential responses and preference stability in experimental games. *Experimental Economics*, 2(3):227–238, 2000.
- J. Brandts and G. Charness. Do labour market conditions affect gift exchange? Some experimental evidence. *Economic Journal*, 114(497):684–708, 2004.
- A. Breaban and C. N. Noussair. Emotional state and market behavior. *Review of Finance*, 22(1):279–309, 2018.
- E. Breza, S. Kaur, and Y. Shamdasani. The morale effects of pay inequality. *Quarterly Journal of Economics*, 133(2):611–663, 2018.
- J. M. Brock, A. Lange, and E. Y. Ozbay. Dictating the risk: Experimental evidence on giving in risky environments. *American Economic Review*, 103(1):415–37, 2013.

- M. Brown, A. Falk, and E. Fehr. Relational contracts and the nature of market interactions. *Econometrica*, 72(3):747–780, 2004.
- R. Brown. *Prejudice: Its social psychology*. John Wiley & Sons, 2011.
- R. Brown and M. Hewstone. An integrative theory of intergroup contact. *Advances in Experimental Social Psychology*, 37(37):255–343, 2005.
- J. Burns, L. Corno, and E. La Ferrara. Interaction, prejudice and performance. Evidence from south africa. *IFS Working Paper No. W19/03*, 2015.
- L. Bursztyn, T. Fujiwara, and A. Pallais. ‘Acting wife’: marriage market incentives and labor market investments. *American Economic Review*, 107(11): 3288–3319, 2017.
- T. Buser, M. Niederle, and H. Oosterbeek. Gender, competitiveness, and career choices. *Quarterly Journal of Economics*, 129(3):1409–1447, 2014.
- B. Camargo, R. Stinebrickner, and T. Stinebrickner. Interracial friendships in college. *Journal of Labor Economics*, 28(4):861–892, 2010.
- A. C. Cameron, J. B. Gelbach, and D. L. Miller. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3): 414–427, 2008.
- A. W. Cappelen, A. D. Hole, E. Ø. Sørensen, and B. Tungodden. The pluralism of fairness ideals: an experimental approach. *American Economic Review*, 97 (3):818–827, 2007.
- A. W. Cappelen, J. Konow, E. Ø. Sørensen, and B. Tungodden. Just luck: An experimental study of risk-taking and fairness. *American Economic Review*, 103(4):1398–1413, 2013.
- D. Card, A. Mas, E. Moretti, and E. Saez. Inequality at work: the effect of peer salaries on job satisfaction. *American Economic Review*, 102(6):2981–3003, 2012.
- D. Card, A. R. Cardoso, and P. Kline. Bargaining, sorting, and the gender wage gap: quantifying the impact of firms on the relative pay of women. *Quarterly Journal of Economics*, 131(2):633–686, 2015.
- S. E. Carrell, M. Hoekstra, and J. E. West. The impact of intergroup contact on racial attitudes and revealed preferences. *NBER Working Paper No. 20940*, 2015.

- F. Casoria and A. Riedl. Experimental labor markets and policy considerations: Incomplete contracts and macroeconomic aspects. *Journal of Economic Surveys*, 27(3):398–420, 2013.
- M. Castillo and R. Petrie. Discrimination in the lab: Does information trump appearance? *Games and Economic Behavior*, 68(1):50–59, 2010.
- E. Cettolin and A. Riedl. Justice under uncertainty. *Management Science*, 63(11):3739–3759, 2016.
- E. Cettolin and S. Suetens. Return on trust is lower for immigrants. *Economic Journal*, 129(621):1992–2009, 2019.
- G. Charness. Attribution and reciprocity in an experimental labor market. *Journal of Labor Economics*, 22(3):665–688, 2004.
- G. Charness and U. Gneezy. What’s in a name? Anonymity and social distance in dictator and ultimatum games. *Journal of Economic Behavior & Organization*, 68(1):29–35, 2008.
- G. Charness and P. Kuhn. Does pay inequality affect worker effort? Experimental evidence. *Journal of Labor Economics*, 25(4):693–723, 2007.
- G. Charness and P. Kuhn. Lab labor: What can labor economists learn from the lab? In O. Ashenfelter and D. Card, editors, *Handbook of Labor Economics*, volume 4a, chapter 3, pages 229–330. Elsevier, 2011.
- G. Charness and M. Rabin. Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 117(3):817–869, 2002.
- G. Charness, G. R. Frechette, and J. H. Kagel. How robust is laboratory gift exchange? *Experimental Economics*, 7(2):189–205, 2004.
- G. Charness, R. Cobo-Reyes, J. A. Lacomba, F. Lagos, and J. M. Pérez. Social comparisons in wage delegation: Experimental evidence. *Experimental Economics*, 19(2):433–459, 2016.
- Y. Chen and J. Konstan. Online field experiments: a selective survey of methods. *Journal of the Economic Science Association*, 1(1):29–42, 2015.
- Y. E. Chentsova-Dutton and J. L. Tsai. Self-focused attention and emotional reactivity: the role of culture. *Journal of Personality and Social Psychology*, 98(3):507, 2010.

- R. Chetty, N. Hendren, M. R. Jones, and S. R. Porter. Race and economic opportunity in the United States: An intergenerational perspective. *NBER Working Paper No. 24441*, 2018.
- K. B. Coffman. Evidence on self-stereotyping and the contribution of ideas. *Quarterly Journal of Economics*, 129(4):1625–1660, 2014.
- K. B. Coffman, L. C. Coffman, and K. M. M. Ericson. The size of the LGBT population and the magnitude of antigay sentiment are substantially underestimated. *Management Science*, 63(10):3168–3186, 2016.
- K. B. Coffman, C. L. Exley, and M. Niederle. When gender discrimination is not about gender. *Harvard Business School Working Paper*, 2017.
- G. A. Cohen. On the currency of egalitarian justice. *Ethics*, 99(4):906–944, 1989.
- Y. Cohen-Charash and P. E. Spector. The role of justice in organizations: A meta-analysis. *Organizational Behavior and Human Decision Processes*, 86(2): 278–321, 2001.
- A. Cohn, E. Fehr, B. Herrmann, and F. Schneider. Social comparison and effort provision: evidence from a field experiment. *Journal of the European Economic Association*, 12(4):877–898, 2014.
- A. Cohn, E. Fehr, and L. Götte. Fair wages and effort provision: Combining evidence from a choice experiment and a field experiment. *Management Science*, 61(8):1777–1794, 2015.
- D. J. Cooper and J. H. Kagel. Other-regarding preferences. volume 2, chapter 4, page 217–289. Princeton University Press, 2016.
- P. Cortés and J. Pan. When time binds: substitutes for household production, returns to working long hours, and the skilled gender wage gap. *Journal of Labor Economics*, 37(2):351–398, 2019.
- R. Croson and U. Gneezy. Gender differences in preferences. *Journal of Economic Literature*, 47(2):448–74, 2009.
- M. Dahlberg, K. Edmark, and H. Lundqvist. Ethnic diversity and preferences for redistribution. *Journal of Political Economy*, 120(1):41–76, 2012.

- J. Dana, D. M. Cain, and R. M. Dawes. What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes*, 100(2):193–201, 2006.
- C. T. Dawes, J. H. Fowler, T. Johnson, R. McElreath, and O. Smirnov. Egalitarian motives in humans. *Nature*, 446(7137):794–796, 2007.
- J. De Quidt, J. Haushofer, and C. Roth. Measuring and bounding experimenter demand. *American Economic Review*, 108(11):3266–3302, 2018.
- S. DellaVigna, J. A. List, and U. Malmendier. Testing for altruism and social pressure in charitable giving. *Quarterly Journal of Economics*, 127(1):1–56, 2012.
- P. A. Diamond. Cardinal welfare, individualistic ethics, and interpersonal comparison of utility: comment. *Journal of Political Economy*, 75(5):765, 1967.
- T. Dohmen, A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner. Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 9(3):522–550, 2011.
- A. Dreber, D. Fudenberg, D. K. Levine, and D. G. Rand. Self-control, social preferences and the effect of delayed payments. *Working Paper, Washington University in Saint Louis*, 2016.
- A. Dube, J. Jacobs, S. Naidu, and S. Suri. Monopsony in online labor markets. *NBER Working Paper (w24416)*, 2018.
- A. Dube, L. Giuliano, and J. Leonard. Fairness and frictions: the impact of unequal raises on quit behavior. *American Economic Review*, 109(2):620–63, 2019.
- M. Dufwenberg and G. Kirchsteiger. A theory of sequential reciprocity. *Games and Economic Behavior*, 47(2):268–298, 2004.
- O. J. Dunn. Multiple comparisons using rank sums. *Technometrics*, 6(3):241–252, 1964.
- R. Dworkin. What is equality? part 1: Equality of welfare, part 2: Equality of resources. *Philosophy and Public Affairs*, 10(3-4):185–246 and 283–345, 1981.
- W. Easterly and R. Levine. Africa's growth tragedy: policies and ethnic divisions. *Quarterly Journal of Economics*, 112(4):1203–1250, 1997.

- R. Eisenberger, P. Lynch, J. Aselage, and S. Rohdieck. Who takes the most revenge? Individual differences in negative reciprocity norm endorsement. *Personality and Social Psychology Bulletin*, 30(6):787–799, 2004.
- P. Ekman and W. V. Friesen. *Unmasking the face: A guide to recognizing emotions from facial clues*. Ishk, Los Altos, United States, 2003.
- P. Ekman, W. V. Friesen, M. O’sullivan, A. Chan, I. Diacoyanni-Tarlatzis, K. Heider, R. Krause, W. A. LeCompte, T. Pitcairn, P. E. Ricci-Bitti, et al. Universals and cultural differences in the judgments of facial expressions of emotion. *Journal of Personality and Social Psychology*, 53(4):712, 1987.
- L. G. Epstein and U. Segal. Quadratic social welfare functions. *Journal of Political Economy*, 100(4):691–712, 1992.
- A. Falk, A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde. Global evidence on economic preferences. *Quarterly Journal of Economics*, 133(4):1645–1692, 2018.
- E. Fehr and A. Falk. Wage rigidity in a competitive incomplete contract market. *Journal of Political Economy*, 107(1):106–134, 1999.
- E. Fehr and K. M. Schmidt. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3):817–868, 1999.
- E. Fehr, G. Kirchsteiger, and A. Riedl. Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics*, 108(2):437–459, 1993.
- E. Fehr, G. Kirchsteiger, and A. Riedl. Involuntary unemployment and non-compensating wage differentials in an experimental labour market. *Economic Journal*, 106(434):106–121, 1996.
- E. Fehr, S. Gächter, and G. Kirchsteiger. Reciprocity as a contract enforcement device: Experimental evidence. *Econometrica*, 65(4):833–860, 1997.
- E. Fehr, E. Kirchler, A. Weichbold, and S. Gächter. When social norms overpower competition: Gift exchange in experimental labor markets. *Journal of Labor economics*, 16(2):324–351, 1998a.
- E. Fehr, G. Kirchsteiger, and A. Riedl. Gift exchange and reciprocity in competitive experimental markets. *European Economic Review*, 42(1):1–34, 1998b.

- E. Fehr, L. Goette, and C. Zehnder. A behavioral account of the labor market: The role of fairness concerns. *Annual Review of Economics*, 1(1):355–384, 2009.
- C. Fershtman and U. Gneezy. Discrimination in a segmented society: An experimental approach. *Quarterly Journal of Economics*, 116(1):351–377, 2001.
- L. Festinger. *A theory of cognitive dissonance*, volume 2. Stanford university press, 1962.
- L. Fiala and C. N. Noussair. Charitable giving, emotions, and the default effect. *Economic Inquiry*, 55(4):1792–1812, 2017.
- F. Finan and L. Schechter. Vote-buying and reciprocity. *Econometrica*, 80(2):863–881, 2012.
- H. Finseraas, Å. A. Johnsen, A. Kotsadam, and G. Torsvik. Exposure to female colleagues breaks the glass ceiling—Evidence from a combined vignette and field experiment. *European Economic Review*, 90:363–374, 2016.
- U. Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, 2007.
- R. Fisman, P. Jakiela, and S. Kariv. Distributional preferences and political behavior. *Journal of Public Economics*, 155:1–10, 2017.
- L. Flabbi and A. Moro. The effect of job flexibility on female labor market outcomes: estimates from a search and bargaining model. *Journal of Econometrics*, 168(1):81–95, 2012.
- M. Fleurbaey. Assessing risky social situations. *Journal of Political Economy*, 118:649–680, 2010.
- C. Fong. Social preferences, self-interest, and the demand for redistribution. *Journal of Public Economics*, 82(2):225–246, 2001.
- R. Forsythe, J. L. Horowitz, N. E. Savin, and M. Sefton. Fairness in simple bargaining experiments. *Games and Economic Behavior*, 6(3):347–369, 1994.
- N. Fortin, T. Lemieux, and S. Firpo. Decomposition methods in economics. In O. C. Ashenfelter and D. Card, editors, *Handbook of Labor Economics*, volume 4a, chapter 1, pages 1–102. North Holland, Amsterdam, 2011.

- A. Franzen and S. Pointner. The external validity of giving in the dictator game. *Experimental Economics*, 16(2):155–169, 2013.
- G. R. Fréchet. Session-effects in the laboratory. *Experimental Economics*, 15(3):485–498, 2012.
- N. H. Frijda. *The emotions*. Cambridge University Press, Cambridge, United States, 1986.
- D. Fudenberg and D. K. Levine. Fairness, risk preferences and independence: Impossibility theorems. *Journal of Economic Behavior and Organization*, 81(2):606–612, 2012.
- S. Gächter and A. Falk. Reputation and reciprocity: Consequences for the labour relation. *Scandinavian Journal of Economics*, 104(1):1–26, 2002.
- S. Gächter and A. Riedl. Moral property rights in bargaining with infeasible claims. *Management Science*, 51(2):249–263, 2005.
- S. Gächter and C. Thöni. Social comparison and performance: experimental evidence on the fair wage–effort hypothesis. *Journal of Economic Behavior & Organization*, 76(3):531–543, 2010.
- M. Gadea, M. Aliño, R. Espert, and A. Salvador. Deceit and facial expression in children: the enabling role of the “poker face” child and the dependent personality of the detector. *Frontiers in Psychology*, 6:1089, 2015.
- E. Ghani, W. R. Kerr, and C. Stanton. Diasporas and outsourcing: evidence from odesk and india. *Management Science*, 60(7):1677–1697, 2014.
- D. S. Gilchrist, M. Luca, and D. Malhotra. When $3 + 1 > 4$: gift structure and reciprocity in the field. *Management Science*, 62(9):2639–2650, 2016.
- M. Gilens. Racial attitudes and opposition to welfare. *Journal of Politics*, 57(4): 994–1014, 1995.
- M. Gilens. “Race coding” and white opposition to welfare. *American Political Science Review*, 90(3):593–604, 1996.
- D. Gill and V. Prowse. A structural analysis of disappointment aversion in a real effort competition. *American Economic Review*, 102(1):469–503, 2012.
- E. L. Glaeser. The political economy of hatred. *Quarterly Journal of Economics*, 120(1):45–86, 2005.

- D. Glover, A. Pallais, and W. Pariente. Discrimination as a self-fulfilling prophecy: evidence from french grocery stores. *Quarterly Journal of Economics*, 132(3):1219–1260, 2017.
- U. Gneezy and J. A. List. Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments. *Econometrica*, 74(5): 1365–1384, 2006.
- J. K. Goeree, M. A. McConnell, T. Mitchell, T. Tromp, and L. Yariv. The 1/d law of giving. *American Economic Journal: Microeconomics*, 2(1):183–203, 2010.
- C. Goldin. A grand gender convergence: its last chapter. *American Economic Review*, 104(4):1091–1119, 2014.
- C. Goldin and L. F. Katz. Human capital and social capital: the rise of secondary schooling in america, 1910–1940. *Journal of Interdisciplinary History*, 29(4): 683–723, 1999.
- C. Goldin and L. F. Katz. A most egalitarian profession: pharmacy and the evolution of a family-friendly occupation. *Journal of Labor Economics*, 34(3): 705–746, 2016.
- C. Goldin and C. Rouse. Orchestrating impartiality: the impact of “blind” auditions on female musicians. *American Economic Review*, 90(4):715–741, 2000.
- B. Greiner. Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125, 2015.
- G. Grimalda, A. Kar, and E. Proto. Procedural fairness in lotteries assigning initial roles in a dynamic setting. *Experimental Economics*, 19(4):819–841, 2016.
- V. Grimm and F. Mengel. Let me sleep on it: Delay reduces rejection rates in ultimatum games. *Economics Letters*, 111(2):113–115, 2011.
- W. Güth, M. Königstein, J. Kovács, and E. Zala-Mező. Fairness within firms: The case of one principal and multiple agents. *Schmalenbach Business Review*, 53 (2):82–101, 2001.
- P. J. Hammond. Ex-ante and ex-post welfare optimality under uncertainty. *Economica*, 48(191):235–250, 1981.

- R. L. Hannan, J. H. Kagel, and D. V. Moser. Partial gift exchange in an experimental labor market: Impact of subject population differences, productivity differences, and effort requests on behavior. *Journal of Labor Economics*, 20(4):923–951, 2002.
- J. C. Harsanyi. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4):309–321, 1955.
- M. S. Hedegaard and J.-R. Tyran. The price of prejudice. *American Economic Journal: Applied Economics*, 10(1):40–63, 2018.
- H. Hennig-Schmidt, A. Sadrieh, and B. Rockenbach. In search of workers' real effort reciprocity—a field and a laboratory experiment. *Journal of the European Economic Association*, 8(4):817–837, 2010.
- D. Herbst and A. Mas. Peer effects on worker output in the laboratory generalize to the field. *Science*, 350(6260):545–549, 2015.
- S. Holm. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2):65–70, 1979.
- J. J. Horton. The effects of algorithmic labor market recommendations: evidence from a field experiment. *Journal of Labor Economics*, 35(2):345–385, 2017a.
- J. J. Horton. Price floors and employer preferences: Evidence from a minimum wage experiment. *CESifo Working Paper (6548)*, 2017b.
- J. J. Horton, D. G. Rand, and R. J. Zeckhauser. The online laboratory: conducting experiments in a real labor market. *Experimental Economics*, 14(3):399–425, 2011.
- T. Hossain and J. A. List. The behavioralist visits the factory: Increasing productivity using simple framing manipulations. *Management Science*, 58(12):2151–2167, 2012.
- M. Ibañez and G. Riener. Sorting through affirmative action: three field experiments in Colombia. *Journal of Labor Economics*, 36(2):437–478, 2018.
- D. S. Karlan. Using experimental economics to measure social capital and predict financial decisions. *American Economic Review*, 95(5):1688–1699, 2005.
- E. Karni and Z. Safra. Individual sense of justice: a utility representation. *Econometrica*, 70(1):263–284, 2002.

- E. Karni, T. Salmon, and B. Sopher. Individual sense of fairness: an experimental study. *Experimental Economics*, 11(2):174–189, 2008.
- J. B. Kessler. When will there be gift exchange? addressing the lab-field debate with laboratory gift exchange experiments. 2013.
- J. Konow. Fair shares: Accountability and cognitive dissonance in allocation decisions. *American Economic Review*, 90(4):1072–1091, 2000.
- J. Kovarik. Giving it now or later: Altruism and discounting. *Economics Letters*, 102(3):152–154, 2009.
- M. Krawczyk. A glimpse through the veil of ignorance: equality of opportunity and support for redistribution. *Journal of Public Economics*, 94(1-2):131–141, 2010.
- M. Krawczyk. A model of procedural and distributive fairness. *Theory and Decision*, 70(1):111–128, 2011.
- M. Krawczyk and F. Le Lec. 'Give me a chance!' An experiment in social decision under risk. *Experimental Economics*, 13(4):500–511, 2010.
- S. Kube, M. A. Maréchal, and C. Puppe. The currency of reciprocity: Gift exchange in the workplace. *American Economic Review*, 102(4):1644–62, 2012.
- S. Kube, M. A. Maréchal, and C. Puppe. Do wage cuts damage work morale? evidence from a natural field experiment. *Journal of the European Economic Association*, 11(4):853–870, 2013.
- I. Kuziemko and E. Washington. Why did the Democrats lose the South? Bringing new data to an old debate. *American Economic Review*, 108(10):2830–67, 2018.
- I. Kuziemko, M. I. Norton, E. Saez, and S. Stantcheva. How elastic are preferences for redistribution? Evidence from randomized survey experiments. *American Economic Review*, 105(4):1478–1508, 2015.
- K. Lang. A language theory of discrimination. *Quarterly Journal of Economics*, 101(2):363–382, 1986.
- E. P. Lazear and S. Rosen. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864, 1981.

- W. Lee and J. E. Roemer. Racism and redistribution in the united states: A solution to the problem of American exceptionalism. *Journal of Public Economics*, 90(6-7):1027–1052, 2006.
- S. Leider, M. M. Möbius, T. Rosenblat, and Q.-A. Do. Directed altruism and enforced reciprocity in social networks. *Quarterly Journal of Economics*, 124(4):1815–1851, 2009.
- G. S. Leventhal, J. Karuza, and W. R. Fry. Beyond fairness: a theory of allocation preferences. *Justice and Social Interaction*, 3(1):167–218, 1980.
- P. Lewinski, M. L. Fransen, and E. S. Tan. Predicting advertising effectiveness by facial expressions in response to amusing persuasive stimuli. *Journal of Neuroscience, Psychology, and Economics*, 7(1):1, 2014.
- J. T. Lind. Fractionalization and the size of government. *Journal of Public Economics*, 91(1-2):51–76, 2007.
- J. A. List. The nature and extent of discrimination in the marketplace: Evidence from the field. *Quarterly Journal of Economics*, 119(1):49–89, 2004.
- J. A. List. The behavioralist meets the market: Measuring social preferences and reputation effects in actual transactions. *Journal of Political Economy*, 114(1):1–37, 2006.
- N. Lupu and J. Pontusson. The structure of inequality and the politics of redistribution. *American Political Science Review*, 105(2):316–336, 2011.
- E. F. Luttmer. Group loyalty and the taste for redistribution. *Journal of Political Economy*, 109(3):500–528, 2001.
- A. Mas and A. Pallais. Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–59, 2017.
- S. Maximiano, R. Sloof, and J. Sonnemans. Gift exchange in a multi-worker firm. *Economic Journal*, 117(522):1025–1050, 2007.
- B. Miao and S. Zhong. Probabilistic social preference: how Machina’s Mom randomizes her choice. *Economic Theory*, 65(1):1–24, 2018.
- E. Miguel and M. K. Gugerty. Ethnic diversity, social sanctions, and public goods in kenya. *Journal of Public Economics*, 89(11-12):2325–2368, 2005.

- C. B. Mulligan and Y. Rubinstein. Selection, investment, and women's relative wages over time. *Quarterly Journal of Economics*, 123(3):1061–1110, 2008.
- R. B. Myerson. Utilitarianism, egalitarianism, and the timing effect in social choice problems. *Econometrica*, 49(4):883–897, 1981.
- W. S. Neo, M. Yu, R. A. Weber, and C. Gonzalez. The effects of time delay in reciprocity games. *Journal of Economic Psychology*, 34:20–35, 2013.
- D. Neumark, R. J. Bank, and K. D. Van Nort. Sex discrimination in restaurant hiring: an audit study. *Quarterly Journal of Economics*, 111(3):915–941, 1996.
- Y. Nguyen and C. N. Noussair. Risk aversion and emotions. *Pacific Economic Review*, 19(3):296–312, 2014.
- M. Niederle. Gender. In J. H. Kagel and A. E. Roth, editors, *Handbook of Experimental Economics*, volume 2, chapter 8, pages 481–562. Princeton University Press, Princeton and Oxford, 2016.
- M. Niederle and L. Vesterlund. Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*, 122(3):1067–1101, 2007.
- M. Niederle, C. Segal, and L. Vesterlund. How costly is diversity? Affirmative action in light of gender differences in competitiveness. *Management Science*, 59(1):1–16, 2013.
- R. Oaxaca. Male-female wage differentials in urban labor markets. *International Economic Review*, 14(3):693–709, 1973.
- A. Ockenfels, D. Sliwka, and P. Werner. Bonus payments and reference point violations. *Management Science*, 61(7):1496–1513, 2014.
- F. O'Connor. *Everything that rises must converge*. Macmillan, 1965.
- J. Oechssler, A. Roeder, and P. Schmitz. Cooling off in negotiations: Does it work? *Journal of Institutional and Theoretical Economics*, 171(4):565–588, 2015.
- L. L. Orr. Income transfers as a public good: An application to afdc. *American Economic Review*, 66(3):359–371, 1976.

- A. Pallais. Inefficient hiring in entry-level labor markets. *American Economic Review*, 104(11):3565–99, 2014.
- A. Pallais and E. G. Sands. Why the referential treatment? Evidence from field experiments on referrals. *Journal of Political Economy*, 124(6):1793–1828, 2016.
- C. A. Parsons, J. Sulaeman, M. C. Yates, and D. S. Hamermesh. Strike three: discrimination, incentives, and evaluation. *American Economic Review*, 101(4):1410–35, 2011.
- E. A. Pascoe and L. Smart Richman. Perceived discrimination and health: a meta-analytic review. *Psychological Bulletin*, 135(4):531, 2009.
- T. F. Pettigrew. Intergroup contact theory. *Annual Review of Psychology*, 49(1):65–85, 1998.
- T. F. Pettigrew and L. R. Tropp. A meta-analytic test of intergroup contact theory. *Journal of Personality and Social psychology*, 90(5):751, 2006.
- J. Pfeffer and N. Langton. The effect of wage dispersion on satisfaction, productivity, and working collaboratively: evidence from college and university faculty. *Administrative Science Quarterly*, 38(3):382–407, 1993.
- E. S. Phelps. The statistical theory of racism and sexism. *American Economic Review*, 62(4):659–661, 1972.
- J. M. Poterba. Demographic structure and the political economy of public education. *Journal of Policy Analysis and Management*, 16(1):48–66, 1997.
- J. Potters and J. Stoop. Do cheaters in the lab also cheat in the field? *European Economic Review*, 87:26–33, 2016.
- R. D. Putnam. E pluribus unum: Diversity and community in the twenty-first century. *Scandinavian political studies*, 30(2):137–174, 2007.
- M. Rabin. Incorporating fairness into game theory and economics. *American Economic Review*, pages 1281–1302, 1993.
- G. Rao. Familiarity does not breed contempt: Generosity, discrimination, and diversity in delhi schools. *American Economic Review*, 109(3):774–809, 2019.
- C. Reed, A. Korda, D. O. Selznick, and G. Greene. *The Third Man*. British Lion, 1949.

- E. Reuben, P. Sapienza, and L. Zingales. How stereotypes impair women's careers in science. *Proceedings of the National Academy of Sciences*, 111(12): 4403–4408, 2014.
- P. A. Riach and J. Rich. Field experiments of discrimination in the market place. *Economic Journal*, 112(483):F480–F518, 2002.
- A. Riedl and P. Smeets. Why do investors hold socially responsible mutual funds? *Journal of Finance*, 72(6):2505–2550, 2017.
- J. E. Roemer. *Theories of distributive justice*. Harvard University Press, Cambridge, MA, 1998.
- J. E. Roemer and K. Van der Straeten. Xenophobia and the size of the public sector in France: a politico-economic analysis. *Journal of Economics*, 86(2): 95–144, 2005.
- J. E. Roemer and K. Van der Straeten. The political economy of xenophobia and distribution: The case of Denmark. *Scandinavian Journal of Economics*, 108 (2):251–277, 2006.
- J. E. Roemer, W. Lee, and K. van der Straeten. *Racism, xenophobia, and distribution: Multi-issue politics in advanced democracies*. Harvard University Press, 2007.
- D. Roodman, M. Ø. Nielsen, J. G. MacKinnon, and M. D. Webb. Fast and wild: Bootstrap inference in stata using boottest. *Stata Journal*, 19(1):4–60, 2019.
- K. Saito. Social preferences under risk: Equality of opportunity versus equality of outcome. *American Economic Review*, 103(7):3084–3101, 2013.
- H. Sarsons, K. Gërkhani, E. Reuben, and A. Schram. Gender differences in recognition for group work. *Working Paper*, 2019.
- T. C. Schelling. The life you save may be your own. *Problems in public expenditure*, pages 127–162, 1968.
- E. Schokkaert and K. Devooght. Responsibility-sensitive fair compensation in different cultures. *Social Choice and Welfare*, 21(2):207–242, 2003.
- M. Shayo. A model of social identity with an application to political economy: Nation, class, and redistribution. *American Political Science Review*, 103(2): 147–174, 2009.

- D. P. Skarlicki and R. Folger. Retaliation in the workplace: the roles of distributive, procedural, and interactional justice. *Journal of Applied Psychology*, 82(3):434, 1997.
- D. A. Small and G. Loewenstein. Helping a victim or helping the victim: Altruism and identifiability. *Journal of Risk and Uncertainty*, 26(1):5–16, 2003.
- E. Snowberg and L. Yariv. Testing the waters: Behavior across participant pools. *NBER Working Paper (24781)*, 2018.
- C. T. Stanton and C. Thomas. Landing the first job: the value of intermediaries in online hiring. *Review of Economic Studies*, 83(2):810–854, 2015.
- C. Starmans, M. Sheskin, and P. Bloom. Why people prefer unequal societies. *Nature Human Behaviour*, 1(4):0082, 2017.
- H. Stichnoth and K. Van der Straeten. Ethnic diversity, public spending, and individual support for the welfare state: A review of the empirical literature. *Journal of Economic Surveys*, 27(2):364–389, 2013.
- E. L. Swée. Together or separate? Post-conflict partition, ethnic homogenization, and the provision of public schooling. *Journal of Public Economics*, 128:1–15, 2015.
- T. Teixeira, M. Wedel, and R. Pieters. Emotion-induced engagement in internet video advertisements. *Journal of Marketing Research*, 49(2):144–159, 2012.
- J. W. Thibaut and L. Walker. *Procedural justice: A psychological analysis*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1975.
- S. T. Trautmann. A tractable model of process fairness under risk. *Journal of Economic Psychology*, 30(5):803–813, 2009.
- S. T. Trautmann and G. van de Kuilen. Process fairness, outcome fairness, and dynamic consistency: Experimental evidence for risk and ambiguity. *Journal of Risk and Uncertainty*, 53(2-3):75–88, 2016.
- S. T. Trautmann and P. P. Wakker. Process fairness and dynamic consistency. *Economics Letters*, 109(3):187–189, 2010.
- T. R. Tyler. Social justice: Outcome and procedure. *International Journal of Psychology*, 35(2):117–125, 2000.

- B. Van Leeuwen, C. N. Noussair, T. Offerman, S. Suetens, M. Van Veelen, and J. Van De Ven. Predictably angry—facial cues provide a credible signal of destructive behavior. *Management Science*, 64(7):3352–3364, 2018.
- M. Wiswall and B. Zafar. Preference for the workplace, investment in human capital, and gender. *Quarterly Journal of Economics*, 133(1):457–507, 2017.
- C. V. Woodward. *The Strange Career of Jim Crow*. Oxford University Press, 1955.

