

Empirical studies on information, beliefs, and choices in education and work

Citation for published version (APA):

de Koning, B. K. (2022). *Empirical studies on information, beliefs, and choices in education and work*. [Doctoral Thesis, Maastricht University]. Maastricht University. <https://doi.org/10.26481/dis.20220926bk>

Document status and date:

Published: 01/01/2022

DOI:

[10.26481/dis.20220926bk](https://doi.org/10.26481/dis.20220926bk)

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.

**EMPIRICAL STUDIES ON INFORMATION, BELIEFS, AND
CHOICES IN EDUCATION AND WORK**

© Bart Kasper de Koning, Maastricht 2022.

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

Cover Lemono - Edited by Amber Crox and Thijs de Koning
Production Published by ROA || Printed in the Netherlands by Canon
ISBN 978-90-5321-612-5

EMPIRICAL STUDIES ON INFORMATION, BELIEFS, AND CHOICES IN EDUCATION AND WORK

Dissertation

To obtain the degree of Doctor at Maastricht University,
on the authority of the Rector Magnificus, Prof. Dr. Pamela
Habibović,
in accordance with the decision of the Board of Deans,
to be defended in public
on 26 September, 2022, at 16:00 hours

by

Bart Kasper de Koning

Promotors

Prof. Dr. Didier Fouarge

Prof. Dr. Robert Dur (Erasmus University Rotterdam)

Assessment Committee

Prof. Dr. Bart Golsteyn (Chair)

Prof. Dr. Inge de Wolf

Prof. Dr. Anne Gielen (Erasmus University Rotterdam)

Dr. Steffen Altmann (University of Copenhagen)

Acknowledgments

I dreaded writing these acknowledgments. Not because I do not want to thank anyone, but rather because I find it so difficult to properly express how happy I am to have been surrounded by so many amazing people these last five years. Here goes nothing.

Didier and Robert, you are without a doubt the first people I want to thank here. Didier, I did not always enjoy wading through the tracked changes you made, sentences you crossed out, and new phrases you handwrote on my papers. However, the papers got better every time I did. I very much look up to you for how complete of a person you are. Somehow, you manage to combine running a 68-person department with all the cooking, cycling and general life-enjoying you do. Robert, there's very few people I enjoy discussing all things economics with as much as with you. With all you do, it sometimes surprised me that you were able to free up an entire afternoon to discuss our projects, new ideas, and whatever else I came to you with. However, I am grateful that you did, as those were some of the most fun and useful chats I had these past years. Thank you both!

Throughout my Ph.D., I got to work with a number of other co-authors as well. I have learned a great deal from every single one of them, and each chapter in this thesis would have been a lot worse without their input. Annemarie and Steffen, thank you for all your work on our joint paper, as well as all the advice you have given me beyond that. And for the great barbecue dinner! Johannes, based on our profiles, we

Acknowledgments

are not a likely set of co-authors. Yet, working together with you has been a blast, perhaps exactly because of that. Michèle, Paul, Philipp and Sandra, our project has proven to be a serious undertaking, and I am very glad we have such a great team to get it done. Michèle and Philipp, I am very happy that we will get to work together even more closely in Ithaca. Paul, we have already proven that we work well together over Zoom, so that will not be any issue. However, I am sure going to miss our frequent skates together. Despite that, I will keep training, and who knows, maybe I will finally be able to keep up with you when we do get to skate again.

Beyond my co-authors, many others have helped the projects in this paper come to fruition. I would like to also thank them: Nisan Mol, Bastian Schilderink, Pauline Thoolen, Erik Fleur, Jacco Tunzi, Mirjam Bahlmann, Robert Jan van Egmond, Dorothy Pillen-Warmerdam, Frank Tinkelenberg, Koen Weide, Marc van der Steeg, Yvonne Engels and Mario Keer.

With the help of everyone listed above, I managed to complete this thesis. However, I still have to defend it in front of a group of scholars I greatly look up to: Bart Golsteyn, Anne Gielen, Steffen Altmann, Inge de Wolf, Sunčica Vujić and Nico Pestel. Thank you for taking the time to read and assess my thesis, and opposing me during the defense.

While I have directly worked with a lot of people, there are many more people I had the pleasure of calling my colleague. Alina, Alexander, Andries, Anna, Arnold, Babs, Barbara, Carla, Chayenne, Cécile, Davey, Elke, Esther, Evie, Fabienne, Frank, Hala, Harald, Henry, Inge, Jacqueline, Jaime, Jessie, Jim, Joyce, Kars, Katarina, Kim, Lara, Made-

lon, Mantej, Margo, Maria, Maria, Marie-Christine, Mark, Melissa, Mel-
line, Mirte, Mélanie, Nick, Olivier, Pelin, Per, Raymond, Rolf, Ruud,
Sabine, Sabine, Saena, Sander, Sandra, Sanne, Sanne, Stijn, Suzanne,
Suzanne, Tijana, Tim, Timo, Tom, Willemijn and Zola. Thank you!
Because of the COVID-19 pandemic, I have not been able to get to
know some of you as well as I would have liked, but the great thing
about academia is that we are very likely to meet again. Miranda and
Mariëlle, thank you for all the help with getting the thesis ready for
printing!

There are a few colleagues missing from the list above, that I would
like to say a few words about. Alexandra, we only overlapped for
a short while in Maastricht. Yet, you and Nico invited me to come
to Melbourne for a visit and made me feel so welcome. Thank you!
Pascal, you became a colleague in January of 2020. That means you had
a fairly rough start: no email in the first month, and no working in the
office starting from the third. Despite that, we became good friends.
Ich hoffe, dass wir uns bald wiedersehen! Lynn, you and I started our
Ph.D.s on the same day, in the same office, and remained office mates
until the end. It has been very nice to be able to share enthusiasm and
frustration with someone at the exact same stage of the Ph.D. process,
and I am very happy you agreed to be my paranymph. Thank you!
Victor, the first time we talked, it was through a series of very formal
messages on Facebook, in which you stated you intended to apply for
a Ph.D. position in Maastricht. The chat ended with a formal statement
saying there would be ice cold beer waiting for you once you arrived.
That beer was waiting there when you arrived, and we have had many
more since then. Now, I consider you one of my closest friends and am

very happy you are my paranymph. Here's to many more ice cold beers together!

There are a number of other friends that I would like to take a moment to thank. Isabela, in economics we often talk about 'externalities'; the consequence of a decision made by two parties on an uninvolved third party. A nice example: you came to Maastricht for Victor, but I also benefited from that, by now being able to call you my friend! Naina and Mariana, jumping on the bike to go to your place was one of my favorite things to do in Maastricht. In the second place because of the delicious food. In the first place because of the amazing company! Max, being at our place together was not one of my favorite things. That had very little to do with you, and a lot with our landlord. Despite that, I greatly enjoyed the time we spent together and the friendship that came out of it. Thank you! Johannes, we initially bonded over pasta carbonara and a so-so tutor manual. That turned into you taking on the role of Klaus Becker (Boris leider nicht) to help me learn German, and from there we formed a friendship that is incredibly dear to me. I hope we get to take many more trips, eat many more amazing dinners, but most of all, that we get to have many more good talks.

While I lived in Maastricht for the first two and half years of my Ph.D., I could often be found a little further north. I had good reasons for that, some of which I would like to mention as well. Bob, Date, Floris, Jelle, Max, Ricardo, Sjoerd, Sjors, I tried to think of something witty to say here. In the end, however, I decided that I prefer to just tell you that I feel very fortunate to call you my friends. Util, when we first met, we were in the same place, did the same things, and even thought the same thoughts. Then you crossed a couple of oceans, and now I will

cross an ocean of my own. We seem to keep moving apart, but if it is up to me, we will remain as close as ever.

After those two and a half years, I moved away from Maastricht to live in Den Bosch. Luckily, a lot of cool people live around there as well. De Dirkens: Annita, Charlotte, Eline, Jacques, Jakob, Joey, Lukas, and Meyke; Frederique, Nick, Ilse, Ruben, Manon, Sven; and many other friends I got to know through Loes. A little further away, too: oma Lucie, Joep, Jeanine, Joost, Eva, Pim, Cécile, Eric, Jessie, Hellen, Mark, Rik, Kim, Huug, Henk, Petra, Thom, and Ivar. Thank you all!

Dan zijn er nog een aantal mensen die ik graag in het Nederlands wil bedanken. Jan en Hélène, mijn ouders, die een enorme hekel hebben aan dankwoorden. Ik houd het daarom kort: bedankt voor alles dat jullie voor mij gedaan hebben. Mijn broertje, Thijs, die een cruciale rol heeft gespeeld in het tot leven brengen van de visuele elementen in de twee veldexperimenten in dit proefschrift. Bedankt daarvoor en nog veel meer; ik ben je nog een paar tekstbewerkingen verschuldigd, ben ik bang. Amber en Thijs, bedankt voor jullie hulp met de kaff! Marion, Odylle en Rick, we kennen elkaar inmiddels negen (!) jaar. Ik kan me nog goed herinneren hoe zenuwachtig ik was om jullie te ontmoeten. Dat was niet nodig, want jullie zijn een geweldige schoonfamilie!

Dan jij, Loes. Wat ik hier opschrijf zal nooit recht doen aan wat je voor mij betekent. Dus in plaats van naar de juiste woorden te zoeken, ga ik dit document nu afsluiten en voor ons een glaasje champagne inschenken. Ik houd van jou.

Bart Kasper de Koning
's-Hertogenbosch
2022-08-14

Contents

- 1 Introduction** **1**

- 2 Student Satisfaction Scores Affect Higher Education Enrollment** **7**
 - 2.1 Introduction 9
 - 2.2 Institutional Context 14
 - 2.2.1 The Dutch higher education system 14
 - 2.2.2 Informational resources 17
 - 2.3 Data 20
 - 2.3.1 Enrollment 20
 - 2.3.2 Student satisfaction scores 22
 - 2.3.3 Substitutes 24
 - 2.4 The Impact of Satisfaction Scores on Enrollment 29
 - 2.4.1 Basic fixed effects regressions 29
 - 2.4.2 Identification of causal impacts 34
 - 2.5 Conclusion 40
 - Appendix 2.A Additional Figures 42
 - Appendix 2.B Additional Tables 45

- 3 Correcting Beliefs about Job Opportunities and Wages** **51**

3.1	Introduction	53
3.2	Institutional Context	62
3.3	Recruitment and Randomization	64
3.4	Experimental Design	69
3.4.1	Occupation test	69
3.4.2	Elicitation of baseline information	70
3.4.3	Information provision	71
3.4.4	Video	73
3.4.5	Elicitation of posterior beliefs and ranking	73
3.4.6	Alternative occupations	73
3.4.7	Elicitation of posterior intended profile choice	74
3.5	Data	74
3.5.1	Sample	74
3.5.2	Survey data	75
3.5.3	Administrative data	77
3.6	Results	78
3.6.1	Descriptive statistics	78
3.6.2	Treatment effects	82
3.7	Conclusion	95
Appendix 3.A	Recruitment text	97
Appendix 3.B	Additional Figures	99
Appendix 3.C	Additional Tables	100
4	Jobs Reports Affect Personal Job Loss Expectations	121
4.1	Introduction	123
4.2	The Jobs Reports	127
4.3	Data & Methodology	130
4.3.1	Data	130

4.3.2	Methodology	132
4.3.3	Justification of identifying assumptions	138
4.4	Results	143
4.4.1	An intuitive first look	143
4.4.2	Main results	145
4.4.3	Heterogeneous treatment effects	148
4.4.4	Alternative outcomes	151
4.5	Conclusion	152
	Appendix 4.A Additional Survey Questions	154
	Appendix 4.B Additional Tables	156
5	Stimulating Occupational Mobility among Unemployed Job Seekers	157
5.1	Introduction	159
5.2	Experimental Design	164
5.2.1	Institutional context	164
5.2.2	Sample selection	166
5.2.3	Interventions	170
5.2.4	Randomization, data collection and timeline . . .	174
5.2.5	Hypotheses	178
5.3	Descriptive Results	179
5.3.1	How do job seekers search?	181
5.3.2	How well are job seekers informed?	184
5.4	Empirical Analysis	187
5.4.1	Take-up: email opening and clicking statistics . .	187
5.4.2	Experimental analysis	192
5.4.3	Remaining analyses	203
5.5	Conclusion	204

Appendix 5.A Additional Tables	206
Appendix 5.B Additional Figures	213
6 Conclusion	215
Bibliography	221
Impact	233
Summary	237
About the author	241
ROA Dissertation Series	243

List of Figures

Figure 2.1	Studiekeuze123.nl website	18
Figure 2.2	Distribution of satisfaction scores	23
Figure 2.3	Visual representation rounding discontinuity analysis	35
Figure 2.A.1	Daily visitors to Studiekeuze123.nl	42
Figure 2.A.2	Distribution of no. of first-year enrolled students	43
Figure 2.A.3	Distribution of $S_{j,k}$ and $S_{j,k}^w$	44
Figure 3.1	Job opportunities and hourly wages of selected occupations	79
Figure 3.2	Prior belief accuracy by relevant group	81
Figure 3.3	Posterior belief accuracy by relevant group	83
Figure 3.4	Weighted average actual prospects and prior beliefs by profile	90
Figure 3.B.1	Graphical representation of randomization	99

Figure 4.1	Search interest	129
Figure 4.2	Development of expectations and unemployment rate over time	131
Figure 4.3	Visual representation of treatment allocation	134
Figure 4.4	Visual representation of placebo treatment allocation	139
Figure 4.5	Variance of $\hat{\beta}_t$ for real Jobs Report and placebo Jobs Reports	140
Figure 4.6	Correlation between report impacts on expectations about national unemployment rate and personal job loss	144
Figure 5.1	Unemployment and vacancies in the Netherlands	168
Figure 5.2	Job prospects Covid and non-Covid occupations	170
Figure 5.3	Example of information email visualization	172
Figure 5.4	Number of search occupations	181
Figure 5.5	Expected and actual job finding prospects	185
Figure 5.6	Clicks on occupation by rank	188
Figure 5.7	Benefits	193
Figure 5.8	Earnings	195
Figure 5.9	Outflow survey invitations	198
Figure 5.B.1	Number of suggestions initially in search set	213
Figure 5.B.2	Share receiving benefits	214

List of Tables

Table 2.1	Descriptive statistics on enrollment and satisfaction scores	21
Table 2.2	Relationship between satisfaction score, national average, satisfaction score of substitutes & first-year enrollment	33
Table 2.3	Impact of being just above satisfaction score rounding threshold on first-year enrollment	36
Table 2.4	Impact of being just above national average satisfaction score on first-year enrollment	38
Table 2.5	Impact of being rounded up over close substitute on first-year enrollment	40
Table 2.B.1	Substitutes for Econometrics at Erasmus University Rotterdam (3,528 views)	45
Table 2.B.2	Similarity of programs and their substitutes	46
Table 2.B.3	Satisfaction scores & first-year enrollment over time	47
Table 2.B.4	Standardized satisfaction score, national average, satisfaction score of substitutes & first-year enrollment	48

Table 2.B.5	Heterogeneity of relationship between satisfaction score, national average, satisfaction score of substitutes & first-year enrollment	49
Table 2.B.6	Robustness check for size; relationship between satisfaction score, national average, satisfaction score of substitutes & first-year enrollment	50
Table 3.1	Treatment assignment, participation and analysis sample overview	66
Table 3.2	Balance of covariates across treatment groups	76
Table 3.3	Treatment effect on likelihood changing favorite occupation and change in prospects	86
Table 3.4	Treatment effect on profiles considered immediately after intervention	89
Table 3.5	Treatment effect on labor market prospects chosen profiles	92
Table 3.6	Treatment effect on job opportunities and hourly wages chosen study program	94
Table 3.C.1	Balance check survey respondents	100
Table 3.C.2	Job opportunities of selected occupations by treatment group and year	101
Table 3.C.3	Hourly wages of selected occupations by treatment group and year	102
Table 3.C.4	Heterogeneity job opportunities of selected occupations	103
Table 3.C.5	Heterogeneity hourly wages of selected occupations by treatment group and year	104
Table 3.C.6	Heterogeneity in prior beliefs	105

List of Tables

Table 3.C.7	Treatment effect on posterior beliefs job opportunities by prior belief accuracy	106
Table 3.C.8	Treatment effect on posterior beliefs hourly wages by prior belief accuracy	107
Table 3.C.9	Detailed treatment effect on posterior beliefs	108
Table 3.C.10	Sender effect on posterior beliefs	109
Table 3.C.11	Heterogeneous treatment effects on posterior beliefs job opportunities	110
Table 3.C.12	Heterogeneous treatment effects on posterior beliefs hourly wages	111
Table 3.C.13	Long-term treatment effect on posterior beliefs	112
Table 3.C.14	Heterogeneous treatment effect on favorite occupation's job opportunities	113
Table 3.C.15	Heterogeneous treatment effect on favorite occupation's hourly wages	114
Table 3.C.16	Effect of sender on likelihood favorite occupation changing and its prospects	115
Table 3.C.17	Long-term treatment effect on prospects favorite occupation	116
Table 3.C.18	Heterogeneity in change in profiles considered immediately after intervention	117
Table 3.C.19	Impact of occupational information on likelihood choosing related profile	118
Table 3.C.20	Treatment effect on prospects of chosen study program by gender	119
Table 3.C.21	Treatment effect on prospects of chosen study program by QOL score	120

Table 4.1	Correlation of news shocks between groups	141
Table 4.2	Correlation of news shocks between groups - Placebo treatment	142
Table 4.3	Effect of cohort-specific news shock on personal job loss expectations - 50/50 split	146
Table 4.4	Effect of cohort-specific news shock on personal job loss expectations - Jackknife	146
Table 4.5	Effect of Placebo news shock on personal job loss expectations - 50/50 split	147
Table 4.6	Effect of Placebo news shock on personal job loss expectations - Jackknife	148
Table 4.7	Heterogeneous treatment effects - Jackknife	150
Table 4.8	Alternative outcomes - Jackknife	152
Table 4.B.1	Summary statistics of $\hat{\beta}_{i,t}$	156
Table 4.B.2	Expectations and Search	156
Table 5.1	Timeline experimental set-up and sample sizes	176
Table 5.2	Sample descriptives: administrative data	180
Table 5.3	Comparison primary and alternative occupation (sur- vey data)	183
Table 5.4	Clicks on occupation by rank, job finding score, au- tomation risk and no. of transitions	189
Table 5.5	Outflow survey invitations and response rates	197
Table 5.6	Outflow to different occupations	199
Table 5.7	Difference-in-differences analysis survey outcomes: Job search activities	201
Table 5.8	Difference-in-differences analysis survey outcomes: labor market beliefs	202

List of Tables

Table 5.A.1 Selection occupations with low job prospects 206

Table 5.A.2 Comparison of composition of survey-respondents
and rest of sample 207

Table 5.A.3 Survey responses about broader job search 208

Table 5.A.4 Balance table: administrative records 209

Table 5.A.5 Balance table: survey responses 210

Table 5.A.6 Balance table: administrative records non-Covid oc-
cupations only 211

Table 5.A.7 Balance table: survey responses non-Covid occu-
pations only 212

1

Introduction

At its core, economics is a social science concerned with agent behavior under scarcity: how do governments, firms and individuals make **choices** when one option excludes another? Reductively, agents weigh the benefits against the costs. However, benefits and costs are hardly ever certain. Agents therefore base their choice on their **beliefs** about the benefits and costs of each available option. The more accurate the beliefs, the better the decision. For agents to make the right decisions, they thus need **information**. In this thesis, I study how exactly information affects individuals' beliefs and choices about education and work.

To understand why studying information in the context of education and work is so important, one must first understand why these choices

are so important. There is a broad literature that has shown that educational choices affect labor market success (Bleemer & Mehta, 2022; Ketel, Leuven, Oosterbeek, & van der Klaauw, 2016; Kirkeboen, Leuven, & Mogstad, 2016), the likelihood of becoming a criminal (Machin, Marie, & Vujić, 2011), who one marries (Lafortune, 2013) and many more outcomes that shape one's life. The same holds for work. Apart from being people's primary source of income, jobs provide a sense of purpose and meaning (Kaplan & Schulhofer-Wohl, 2018). Not all jobs are the same, however; there are large differences in earnings (U.S. Department of Labor, 2022), non-pecuniary amenities (Clark, Cotofan, & Layard, 2022) as well as perceptions of usefulness (Dur & Van Lent, 2019) between occupations. On the other side of the coin, being unemployed can cause feelings of depression and anxiety (Mandal, Ayyagari, & Gallo, 2011) as well as other health issues (Caroli & Godard, 2016). In short, the degree one chooses to obtain and the job one decides to apply to are potentially life-altering choices.

Economic theory presumes that all of these choices are determined by individuals' beliefs about the benefits and costs of the different options they have. If these beliefs were generally accurate, I would not have written this thesis. The problem is that individuals' beliefs about these topics are often inaccurate. Students often hold erroneous beliefs about the returns to different education programs (Baker, Bettinger, Jacob, & Marinescu, 2018; Hastings, Neilson, & Zimmerman, 2015; Hastings, Neilson, Ramirez, & Zimmerman, 2016; Pekkala Kerr, Pekkariinen, Sarvimäki, & Uusitalo, 2015; Conlon, 2019). Similarly, (unemployed) job seekers are overoptimistic about their chances of finding employment (Spinnewijn, 2015), and remain so over the duration of their employment spell (A. Mueller, Spinnewijn, & Topa, 2021). This

means that people make some of the most important decisions of their lives being misinformed.

This thesis sheds light on how information affects beliefs and choices about education and work. It consists of four chapters, written in collaboration with several co-authors. The first two chapters cover educational choices. The two chapters thereafter are about work. On both topics, one chapter is based on observational data, and another on a field experiment I conducted. Both types of studies come with their distinct advantages and disadvantages. The field experiments provide a lot of control: over the information itself, over who receives the information, and over the collected data. It allows me to cleanly identify the impact of information I deem relevant, on any outcome I care to collect data for. With observational data, it requires a little more effort to credibly identify causal effects, and the perfect dataset is not always available. However, observational studies have higher external validity, providing evidence on the impact of information on a more diverse set of individuals, over a longer time frame. There are further differences between the two. In experimental studies, I can (for the most part) ensure that individuals receive the information. This is not true for the observational studies. Both types of studies therefore answer different questions. The field experiments answer whether receiving information affects beliefs and choices. The observational studies answer whether information being available does so. In short, the two types of studies complement each other; both in terms of validity and in terms of the questions they answer. Together, they provide a comprehensive look into the impact information has on individuals' beliefs and choices about education and work.

Moving to the individual chapters, **Chapter 2** studies the impact of published student satisfaction scores (ranging from 1 to 5) on enrollment of first-year students for the near universe of existing higher education programs in the Netherlands between 2011 and 2019. To determine each programs' closest substitutes, I use pageview data from the largest Dutch educational information website. This allows me to not only analyze the impact of changes in a program's own published student satisfaction score, but also the impact of changes in the student satisfaction scores of its substitutes. I analyze the impact of these satisfaction scores using fixed effects Poisson regressions and exploit rounding discontinuities to identify causal effects. The findings show that student satisfaction scores matter for enrollment. An increase in a program's student satisfaction score leads to higher levels of enrollment, whereas an increase in the student satisfaction scores of substitutes leads to lower levels of enrollment. Point estimates of the impact of a program's student satisfaction score being rounded up to the next tenth on first-year enrollment range between 1.70% and 3.52%, depending on the bandwidth around the threshold we consider. Conditional on being above the rounding threshold, a program being rounded up over at least one of its closest substitutes increases first-year enrollment by up to 4.37%.

Chapter 3 presents the results of a large-scale field experiment in which I provide students at randomly selected schools with information about the job opportunities and hourly wages of a set of occupations they are interested in. The experiment takes place on an online career guidance counseling platform that is widely used in the Netherlands, and involves 28,267 pre-vocational secondary education students in 243 schools over a period of 2 years. I find that the information improves

the accuracy of students' beliefs, both in the short run (for job opportunities and hourly wages) and in the long run (for job opportunities only). Students who receive the information also change their favorite occupation 0.88 to 2.16 percentage points more often, and switch towards an occupation with better labor market prospects when they do so. Last, and most importantly, they select secondary school specializations related to occupations with better labor market prospects (1.5% and 0.3% – €0.05 an hour – higher than the control group mean for job opportunities and wages, respectively) and choose post-secondary education programs with higher expected wages (2.5% – approximately €0.40 an hour – higher than the control group mean).

In **Chapter 4**, I use data from the New York Federal Reserve's Survey of Consumer Expectations to study how the United States Bureau of Labor Statistics' Employment Situation Reports (Jobs Reports) affect individuals' expectations about the likelihood of losing their own job. I do this in two steps. First, I estimate the information shocks of the jobs reports on expectations about the development of the national unemployment rate in the next twelve months. I do this by comparing survey responses shortly before and after publication of the reports. Second, I estimate how these shocks affect individuals' expectations about losing their own job in the same time frame. The results show that when a report is estimated to increase beliefs about the likelihood of the unemployment rate increasing by 1 percentage point, beliefs about the likelihood of personal job loss during that time increase by up to 0.22 percentage points. I further find that the information shock negatively affects individuals' beliefs about the likelihood of finding a new job if they were to lose their current one, but surprisingly has a positive effect on their beliefs about the likelihood of voluntarily leav-

ing their job. The results are robust to the use of different bandwidths around the reports' publication dates and placebo treatments provide reassurance that the information shock is indeed the mechanism driving the result.

Chapter 5 studies the impact of online information provision to job seekers who are looking for work in occupations with relatively poor labor market prospects. I provide the information through a personalized email containing suggestions about suitable alternative occupations and how the prospects of these alternatives compare to the job seekers' current occupation of interest. A second treatment adds a motivational video aimed at addressing the psychological hurdles of switching to a different occupation. I evaluate the interventions using a randomized field experiment with 30,129 unemployed job seekers, and acquire additional descriptive information on beliefs and job search. The results show no impact on received benefits and earnings in the first eight months after the treatment. The findings do show that treated individuals are 1.79 percentage points more likely to have found a job seven months after the intervention, although this difference decreases to 1.19 percentage points four months later. Moreover, treated individuals are between 5 and 6 percentage points more likely to have done so in an occupation different from their initial occupation of interest. This may be promising for their longer-term prospects.

Taken together, Chapters 2 to 5 show that information matters in important ways for beliefs and choices about education and work. Finally, In **Chapter 6**, I reflect on the findings, their policy implications and potential directions for future research.

2

Student Satisfaction Scores Affect Enrollment in Higher Education Programs

Abstract

We study the impact of published student satisfaction scores (ranging from 1 to 5) on enrollment of first-year students for the near universe of higher education programs in the Netherlands between 2011 and 2019. We use pageview data from the largest Dutch educational information website to determine each programs' closest substitutes. This allows us to not only analyze the impact of changes in a program's own published student satisfaction score, but also the impact of changes in the student satisfaction scores of its substitutes. We analyze the impact of these satisfaction scores using fixed effects Poisson regressions and exploit rounding discontinuities to identify causal effects. Our findings show that student satisfaction scores matter for enrollment. An increase in a program's student satisfaction score leads to higher levels of enrollment, whereas an increase in the student satisfaction scores of substitutes leads to lower levels of enrollment. Point estimates of the impact of a program's student satisfaction score being rounded up to the next tenth on first-year enrollment range between 1.70% and 3.52%, depending on the bandwidth around the threshold we consider. Conditional on being above the rounding threshold, a program being rounded up over at least one of its closest substitutes increases first-year enrollment by up to 4.37%.

This Chapter is based on joint work with Annemarie Künn and Steffen Künn. We would like to thank Erik Fleur and Jacco Tunzi of the Dutch Education Executive Agency for their help in acquiring data on enrollment. We would like to thank Pauline Thoolen, Bastian Schilderink and Nisan Mol of Studiekeuze123.nl for their help in acquiring data on satisfaction scores as well as pageview data.

2.1 Introduction

Every year, millions of students worldwide choose different university study programs. An important decision, as education affects many later life outcomes, such as labor market success (Kirkeboen et al., 2016), the likelihood of committing crimes (Machin et al., 2011) and marriage (Lafortune, 2013). However, the value of education does not only come from its returns upon completion. Earlier literature has shown that students enjoy substantial consumption value from their studies (see e.g., Alstadsæter 2011; Oosterbeek and Van Ophem, 2000; Oreopoulos and Salvanes, 2011).

Students decide on what study program to enroll in under imperfect information. Indicators such as rankings, labor market projections and student satisfaction scores are supposed to aid students' decisions. Due to progressing digitalization, students nowadays have easy, user-friendly and instant access to standardized information about universities and programs. The publication of such indicators on university and program quality may lead to better-informed choices, and enhance aggregate human capital and welfare by improving matches (Horstschräer, 2012). At the same time, we observe that markets for higher education have become more competitive, a process which initially started with the Bologna process, aimed at making international degrees more comparable. As such, universities have an incentive to invest in and promote quality indicators in order to attract more and better students (Hazelkorn, Loukkola, & Zhang, 2014).

Given the massive investments by public authorities and universities alike in these indicators, it is crucial to know whether the information provided indeed guides students' educational choices. Based on

a number of observational and experimental studies, we know that labor market information about expected earnings and employment opportunities has a moderate, but generally positive effect on students' decisions to enroll in certain programs (see e.g., Bonilla-Mejía, Bontan, and Ham, 2019; Hastings et al., 2015; Kerr, Pekkarinen, Sarvimäki, and RoopeUusitalo, 2020, and Chapter 3 of this thesis). We also know that university rankings play an important role in attracting students (Brewer, Eide, & Ehrenberg, 1999; Griffith & Rask, 2007; R. E. Mueller & Ruckerbie, 2005). The recent report by the *European University Association* underlines the increasing importance of university rankings and shows how university strategies are driven by these rankings (Rauhvargers, 2013). While labor market projections and rankings inform students on the likely impact of their decision on later life outcomes, they do not provide information on the more short-term consumption value of education. In this Chapter, we study how information about student satisfaction scores – determined by factors such as the quality of teaching, knowledge accumulation, the curriculum, community building, fellow student behavior and institutions' responsiveness and concern (Douglas, Douglas, McClelland, & Davies, 2015; Gibson, 2010) – impact enrollment.

Universities invest significant resources to improve student satisfaction. While student satisfaction may be important to universities in and of itself, they may also expect it to have spillover effects on more tangible metrics, such as student performance, student outcomes, student retention, and enrollment of new students. The evidence on these spillover effects is scant, however. While many studies in educational sciences have shown a strong correlation between student satisfaction and performance, there is no evidence for a causal interpretation. In

one article, Bean and Bradley (1986) argue “that satisfaction had a greater influence on performance than performance had on satisfaction”. However, their methodology would likely not hold up to today’s standards.¹ In a recent study, Britton et al. (2021) show that student satisfaction is not related to student outcomes as measured by graduates’ labor market earnings. This evidence suggests that the only remaining economic argument for universities to invest in student satisfaction is based on the assumption that student satisfaction is key to retain students and attract new ones. However, population-representative, causal evidence on this assumption is missing. To the best of our knowledge, there is only one study, by Horstschräer (2012), who shows a positive impact of student satisfaction score on applications. However, Horstschräer uses time variation to identify the effects of student satisfaction, and acknowledges that endogeneity concerns may still exist. Moreover, this evidence is restricted to one specific program (medicine) with a strong selection of high-ability students. This subpopulation can be assumed to be well informed and hence particularly responsive to variations in publicly available indicators, including student satisfaction.

To our knowledge, we are the first to provide a detailed causal analysis on the role of student satisfaction in determining students’ educational decisions. We rely on administrative records documenting the first-year enrollment figures of almost all university level study programs in the Netherlands between 2011 and 2019. We link these enrollment figures to satisfaction scores (ranging from 1 to 5) published on the largest Dutch educational information website (Studiekeuze123.nl) in

¹They use a 2-stage least squares estimation strategy. They assume that institutional fit and utility influence satisfaction, but not GPA, and that high-school performance will influence GPA, but not satisfaction.

June of the year before, from 2010 to 2018. These satisfaction scores are rounded before being published, providing us with a source of exogenous variation to assess their impact. Satisfaction scores are prominently displayed on the website and therefore likely used to compare study programs to one another. To fully capture their impact on enrollment, one needs to have information on satisfaction score of programs' close substitutes as well. To this end, we construct a substitutability matrix using the Studiekeuze123.nl's pageview data. Specifically, we determine how often two program pages are viewed by the same user. We argue that the more often two pages are viewed in tandem, the closer substitutes they are. These three data sources allow us to provide a detailed look into the impact of programs' student satisfaction scores, as well as the scores of their substitutes, on subsequent enrollment.

We start our analysis by running fixed effects Poisson regressions of first-year enrollment on published satisfaction scores. We find that there is a clear positive relationship between the two, that is surprisingly linear. In a next step, we also include the national average satisfaction score of the same degree program at different universities. It has a clear, negative relationship with enrollment. We further extend our analysis by adding a number of measures of the average satisfaction score of a program's closest substitutes as determined by our substitutability matrix. These measures also show a negative relationship with first-year enrollment. By adding these measures, the negative effect of the degree-level national average moves towards zero. Our interpretation of this finding is that our measure of substitutes' satisfaction scores is more informative than a 'naive' metric such as the degree-level national average. It underlines the importance of properly identifying substitutes.

While the fixed effects Poisson regressions take care of many potential sources of endogeneity, we cannot interpret these results as causal. In a next step, we therefore exploit the fact that satisfaction scores are rounded to the nearest tenth before being published on *Studiekeuze12-3.nl*. These rounding discontinuities allow us to compare programs when they were just below the rounding threshold to when they were just above the rounding threshold. We estimate a positive effect of being rounded up to the next tenth on subsequent first-year enrollment of 1.70% to 3.52%, but the effects are marginally significant and sensitive to the chosen bandwidth. Next, we look at programs whose satisfaction score is very close to the degree-level national average. We compare programs when they were rounded to the same satisfaction score as the national average to when they were rounded up over the national average. We find no impact of being just above the national average, but argue that this may be because programs very close to the national average have, by definition, very average satisfaction scores. Lastly, we study how important it is to receive a higher student satisfaction score than a close substitute. We show that being rounded up over at least one close substitute, conditional on being above the rounding threshold, increases first-year enrollment by up to 4.37%.

We contribute to the literature by providing the first causal evidence of the impact of student satisfaction scores on enrollment. We base this on observational data that include the near universe of existing university programs in an industrialized country. Our findings contribute to our understanding of the impact of information dissemination and provide new insights into the educational choice of students that goes beyond later life returns. The Netherlands is an ideal country to study this question, as there is strong competition between programs, because of

geographical proximity as well as homogeneous university quality.

The rest of this Chapter is structured as follows. Section 2.2 explains the institutional context: the Dutch higher education system and the available informational resources. Section 2.3 describes the data we use and provides a first look into the relationship between satisfaction scores and enrollment. In Section 2.4, we present a more sophisticated data analysis, lay out how we identify the causal impact of satisfaction scores on first-year enrollment and present the results. Section 2.5 concludes.

2.2 Institutional Context

2.2.1 The Dutch higher education system

To understand how student satisfaction scores may impact enrollment, it is important to discuss the higher education market in the Netherlands. It is a highly competitive market on the supply side: tuition fees are fixed at the national level, admission to most programs is open conditional on prerequisites (i.e., no selection), university quality is relatively homogeneous, and the country is geographically small. Students thus have many options available to them, without major differences in selectivity, quality or costs. Relevant information on the expected utility they would enjoy at different available programs is therefore likely to drive their decisions.

In total, there are 37 universities of applied sciences and 12 research universities, which are all public entities. All universities are subject to monitoring by the Ministry of Education, with periodic accreditation rounds. This is to ensure homogeneous educational quality

and standards across all public universities within the Netherlands (Inspectorate of Education, 2022). Students from the European Union enrolled at accredited public universities are eligible to subsidized tuition fees of about €2,000.- per year.² Both universities of applied sciences as well as research universities generally offer a broad range of degree programs, with the exception of three research universities of technology.

Students' level of secondary education is decisive for the type of university students can apply for. Universities of applied sciences are open to graduates of higher general continued education, preparatory scientific education and middle-level applied education. Research universities are only open for graduates of preparatory scientific education and under certain conditions to graduates of universities of applied sciences. Most students entering universities of applied sciences come from higher general continued education. For research universities, most students come from preparatory scientific education (Inspectorate of Education, 2022). Both at higher general continued education and preparatory scientific education tracks, students have to choose a specific educational profile in their third year. An educational profile consists of a number of fixed subjects, that prepare pupils for specific areas of study. There are four different profiles: culture and society, economy and society, nature and health, nature and technology (Nuffic, 2019). The different profiles provide access to different study programs in higher education. In most cases, having the right

²Tuition fees have increased gradually over the years. For the 2010/2011 academic year, the subsidized tuition fee was €1,627 per year (Ministry of Education, Culture and Science, 2010). For the 2018/2019 academic year, it was €2,060 per year (Ministry of Education, Culture and Science, 2018). Students outside of the EU/EEA pay sticker price.

degree level as well as educational profile provides access to a study program. Since 2014, students are required to complete an advisory intake procedure before enrollment, but its results are not binding (Chu & Westerheijden, 2018). Only in exceptional cases, universities can apply a 'numerus fixus', limiting the number of students that can enroll. They can do so because of excessive demand, high costs of education or poor labor market prospects. Some degrees, such as medicine, have a numerus fixus at all universities. Others, such as psychology and international business only have a numerus fixus at particular universities. Once a numerus fixus has been instated, the enrolment process changes as well. Students that apply for a program with a numerus fixus have to write a letter of motivation and are often required to take a test. Aside from that, there is a maximum number of numerus fixus study programs that students can apply for (generally two or three).

The uniform governmental policies have prevented the rise of an institutional hierarchy (Veerman et al., 2010). All Dutch research universities are within the global top 250, but none of them is in the global top 50 according to the Times Higher Education ranking (2022). The fact that these universities all rank fairly high, arguably without any university being 'world class', means that students are unlikely to consider particular universities to be more desirable than others based on prestige. The competition among universities is further strengthened because of the relatively small geographical size of the Netherlands, that makes spatial accessibility high (Sa, Florax, & Rietveld, 2004). Even though the universities are not equally spread throughout the Netherlands, the size of the country enables students to reach any university either by commuting or small distance relocation.

2.2.2 Informational resources

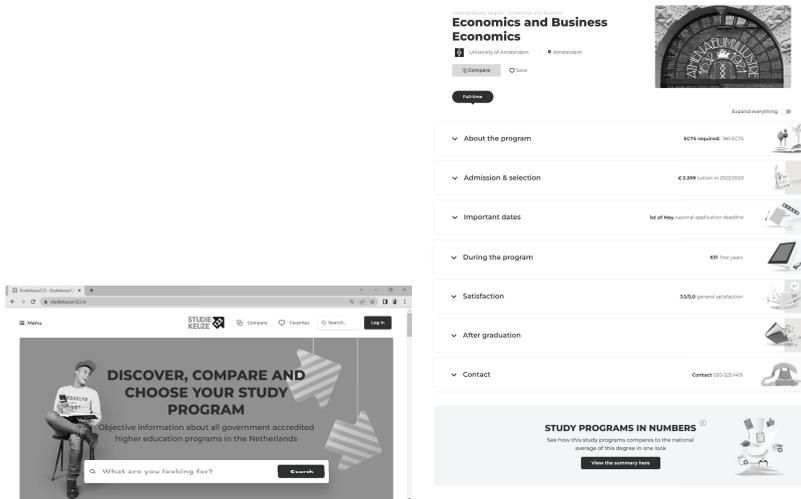
To aid prospective students in their decision making, The Dutch Ministry of Education and Science, together with students and higher education institutions, founded the 'Studiekeuze123' foundation. The aim of this foundation is to provide students with objective information about post-secondary education programs at research universities and universities of applied sciences. To this end, the foundation does two things: it maintains the national student survey and operates the Studiekeuze123.nl website.

The national student survey is an annual survey sent out to all individuals enrolled in higher education in the Netherlands.³ All students enrolled at government funded universities are invited to participate in the survey. More than 200,000 students answered the survey each year between 2010 and 2018. The survey asks students how satisfied they are about the study programs they are enrolled in on a number of dimensions (content, professors, facilities, general atmosphere, etc.). The most important and widely publicized metric is students' general satisfaction with their study program. The relevant question reads "*How satisfied are you with your course programme in general?*". Students can answer on a scale from 1 (very dissatisfied) to 5 (very satisfied).

Studiekeuze123 visualizes and publishes the information on their website. The Studiekeuze123 website is the largest website for higher education information in the Netherlands, boasting 632,884 unique visitors between May of 2019 and June of 2020. Figure 2.A.1 in the Appendix displays the number of daily visitors during this time period.

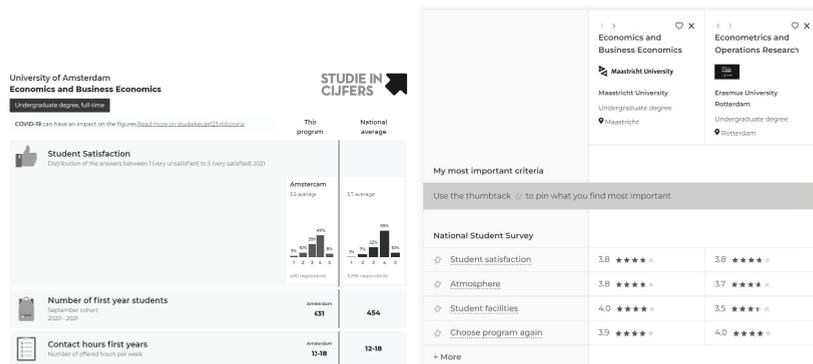
³There was some discussion about the participation of universities of applied sciences in 2020, but this is beyond the scope of our data.

Figure 2.1: Studiekeuze123.nl website



(a) Studiekeuze123.nl homepage

(b) Studiekeuze123.nl program page



(c) 'Study Programs in Numbers' page

(d) Program comparison page

Note: Translated versions of original Studiekeuze123.nl webpages.

Figure 2.1 shows a number of important pages on the Studiekeuze123.nl website. Panel (a) shows the Studiekeuze123.nl homepage. The website is designed in such a way that looking up programs is front and center. It uses a smart search bar, that directs students to the available programs based on the keywords they put in. Once students find a program, they are redirected to the program page, displayed in Panel (b). Students can find a wide range of information on this page: information about the program, requirements for enrollment, important dates for their application, statistics on the number and demographic characteristics of enrolled students, student satisfaction, graduation rates, labor market statistics and contact information. Note that the general satisfaction score of a program is listed and immediately visible as students open the page.

On the program page, students have a number of further options. Two are of particular interest to our study. Firstly, they can open the 'Study Programs in Numbers' sheet for the program. This is a quick summary of all the program's relevant statistics. Panel (c) provides a translated example of (part of) the sheet. The first information students see when they open the sheet is the general student satisfaction score of the program they are looking at, as well as the national average for the degree. The page further provides information on the number of first year students, contact hours, as well as (not displayed) the share of students who continue on to the second year, share of students who receive degrees within four years, further education and labor market information⁴. Another option students have is to compare the program to other programs they are interested in. Panel (d) provides an exam-

⁴If the information pertains to a program at a University of Applied Sciences or a Masters' program.

ple of such a comparison between two programs. Again, the student satisfaction scores of both programs is the first information students see when they navigate to this page. Figure 2.1, combined with the large number of visitors on *Studiekeuze123.nl*, suggests that students are likely to be aware of the satisfaction scores of the programs they are interested in, and could thus have a real impact on enrollment decisions.

2.3 Data

2.3.1 Enrollment

Our outcome of interest is the enrollment numbers of Dutch study programs. For this, we rely on data we received from the Dutch Education Executive Agency (DEEA). On the first of October of each year⁵, the DEEA records the enrollment status of all students in higher education. For each student, they register in which degree program(s) they are enrolled, as well as at which university. We refer to the degree program-university combination as a study program. Aside from the enrollment status of each student, the DEEA records a number of background characteristics, such as gender, migration background and secondary school results. We have access to data from 2011 to 2019. As our analyses are at the study program level, we reshape the data to a panel at that level. For each year, we record the number of (first-year) students enrolled in the program and how many have certain characteristics. We have data on 1,894 programs over a period of 9 years in our sample.

⁵This is about a month after the official start of the academic year.

Table 2.1: Descriptive statistics on enrollment and satisfaction scores

	Mean	Std. dev.	Within std. dev.
Enrollment			
No. of first-year students	71.86	91.14	22.34
Total no. of students	328.93	360.34	70.33
No. of male first-year students	33.58	51.89	12.01
No. of first-year students w/ migration background	25.52	44.60	14.43
Avg. high school grade	6.57	0.31	0.174
Sat. scores			
Student satisfaction score	3.94	0.28	0.167
National avg. sat. score of degree	3.92	0.198	0.103
Avg. sat. score. no. 1 substitute	3.94	0.254	0.148
Avg. sat. score. top 3 substitutes	3.94	0.190	0.103
Avg. sat. score. top 5 substitutes	3.94	0.168	0.088
Avg. sat. score. top 10 substitutes	3.93	0.147	0.075
Avg. sat. score. top 20 substitutes	3.94	0.132	0.066
Weighted avg. sat. score of substitutes	3.81	0.130	0.097
Observations	12821		

Note: data at program-year level. *No. of first-year students, total no. of students, no. of male first-year students, no. of first-year students w/ migration background and avg. high school grade* are based on data from the Dutch Education Executive Agency. *Student satisfaction score, national avg. sat. score of degree, weighted avg. sat. score of substitutes, weighted avg. sat. score top 20 subst.* are based on the ‘Studiekeuzedatabase’ from Studiekeuze123.nl. In case of the latter two variables, the substitutes are determined as described in Section 2.3.3, using Studiekeuze123.nl’s pageview data. Within standard deviation is the standard deviation of the individually (program level) demeaned variables.

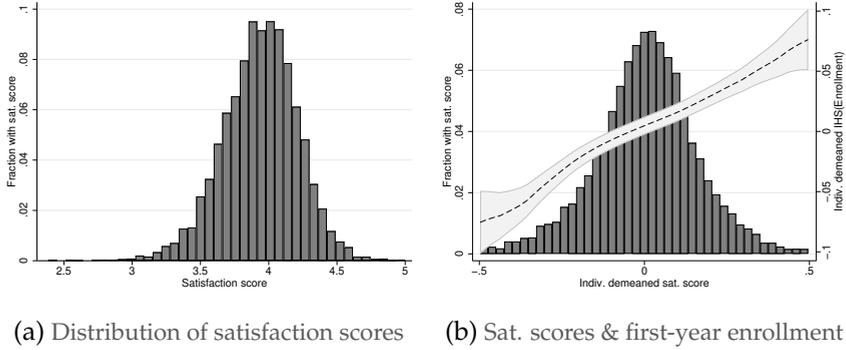
The top part of Table 2.1 provides some descriptive statistics on enrollment. The mean ‘cohort size’ is about 72, but the standard deviation is just above 91. The size of the standard deviation makes clear that there are large differences between programs in the number of first-year students that enroll. Figure 2.A.2 in the Appendix shows the distribution of the number of first-year enrolled students. The distribution is similar to a Power law distribution; there are a large number of small programs and a small number of large programs. The largest programs enroll more than 800 new first-year students on a yearly ba-

sis. The data are similar for the total number of students enrolled, as well as the number of first-year male students and students with a migration background; the standard deviation is large compared to the mean. The average high school grade of enrolled students is 6.57, with a standard deviation of just 0.31. It shows that programs are fairly homogeneous in terms of student ability; likely a product of the open admission policy in Dutch higher education.

2.3.2 Student satisfaction scores

We obtain information on the historical satisfaction scores for the different study programs directly from Studiekeuze123's database records from June of every year. For each study program, we have access to the unrounded student satisfaction score. These scores are rounded to the nearest tenth before they are uploaded to the Studiekeuze123 website. To analyze the impact of the satisfaction scores on subsequent enrollment, we match the satisfaction scores published in June of year $t - 1$ to enrollment in year t . Students likely decide on their study program before June and are even forced to register before the 1st of May since the 2014/2015 academic year. Figure 2.A.1 in the Appendix shows that in 2019 and 2020, the high traffic months on Studiekeuze123.nl are October, January and April.

Figure 2.2: Distribution of satisfaction scores



Note: The x-axis displays the satisfaction scores. Panel (a) displays distribution of raw satisfaction scores. In Panel (b) the satisfaction scores are demeaned at the program level. The right-hand side y-axis of panel (b) shows the inverse hyperbolic sine of enrollment, again demeaned at the individual level. The dashed line displays the polynomial fit, with the gray area displaying the 95% confidence interval.

Panel (a) of Figure 2.2 shows the distribution of student satisfaction scores. The first row of the bottom part of Table 2.1 provides further details. The Figure shows that satisfaction scores are approximately normally distributed, with a mean just below 4.0 (3.94). The standard deviation is equal to 0.28. Panel (b) provides a first look into the relationship between student satisfaction scores and first-year enrollment. The x-axis shows the program level demeaned satisfaction score. The left y-axis shows the distribution of these scores and the right y-axis shows the program-level demeaned inverse hyperbolic sine⁶ of enrollment. The reason we demean at the program level is that the variance

⁶The inverse hyperbolic sine is a transformation similar to the natural logarithm. However, contrary to the natural logarithm, it allows for values of zero.

of enrollment is very high, as discussed in Section 2.3.1. The within-program variance is substantially lower, as shown in Table 2.1. Moreover, by only exploiting within-program variation in our analyses, we take care of any time-invariant endogeneity concerns. For instance, one may expect that programs with larger budgets have higher student satisfaction, but are able to invest more in marketing as well. Time-varying endogeneity concerns, such as budgetary changes, do remain. Keeping these caveats in mind, Figure 2.2 shows a clear positive relationship between enrollment and a program's satisfaction score.

2.3.3 Substitutes

While satisfaction scores may be an important determinant of a prospective student's enrollment decision, their value does not have a clear interpretation. Higher levels of satisfaction scores are unlikely to drive larger number of students to enroll in higher education. It is more likely that students use these numbers to compare different programs they are interested in to each other. Therefore, it is key to identify programs' substitutes for the empirical analysis.

Ideally, we would be able to calculate the *cross-price elasticity* between programs. However, this would require (i) exogenous price changes and (ii) enough data to properly estimate the elasticity. Our setting provides neither. As stated in Section 2.2, the cost of enrollment is the same for almost all study programs. In addition, we only have nine years of enrollment data available to us, for fewer than 2,000 programs.

An alternative to the cross-price elasticity would be to define substitutes by their characteristics. For instance, it is likely that a program's closest substitutes are the same degree programs at different universities close by, or same-sector programs at the same university. However, this would cause us to potentially miss a lot of substitutes that do not fit these criteria. For instance, students interested in the math-heavy degrees in the economics and business sector, may also be interested in engineering degrees. These degrees are not in the same sector, however, and generally not available at the same universities. We thus require a more sophisticated approach.

To determine each program's closest substitutes, we make use of page-view data from *Studiekeuze123.nl* between May of 2019 and June of 2020. These data allow us to observe the behavior of prospective students looking for information about study programs. Apart from providing us with insight into how prospective students use these types of websites, it also allows us to construct a substitutability matrix. We construct the substitutability matrix by looking at which two program pages are viewed by the same user most often. We argue that if students who look at a certain program are highly likely to look at another program, these are likely to be close substitutes. To be more precise, we calculate the following:

$$S_{j,k} = \frac{Views_k|Viewed_j}{Views_j}. \quad (2.1)$$

Here, $S_{j,k}$ is our substitutability index. $Views_k|Viewed_j$ is the number of views to program k 's page by individuals who have also viewed program j 's page. $Views_j$ indicates the number of views that pro-

gram j 's page received in total. The higher this number, the larger the share of individuals who visit program j 's page as well as program k 's page.

A potential issue with this metric is that it is attenuated by individuals who look at a very large number of programs. We may want to assign more weight to individuals looking at just two programs, than to those looking at twenty programs. We therefore construct a weighted substitutability index:

$$S_{j,k}^w = \frac{\sum_{i=1}^N \frac{1}{\text{Page Visits}_i} \text{Viewed}_{i,k} | \text{Viewed}_{i,j}}{\text{Views}_j}. \quad (2.2)$$

Here, $\sum_{i=1}^N \frac{1}{\text{Page Visits}_i} \text{Viewed}_{i,k}$ adds a weight to individual i 's visit that is the inverse of the number of pages individual i visited. For each program, we can then rank the substitutes in descending order of $S_{j,k}^w$.

Table 2.B.1 in the Appendix provides an example for the *Econometrics and Operations Research* degree at Erasmus University Rotterdam; one of the twelve research universities of the Netherlands. Its number one ranked substitute is *Econometrics* at the Free University of Amsterdam. Unsurprisingly, this is the same degree program at a university close by. The substitutes ranked number two and four are *Econometrics* at the University of Amsterdam and *Economics* at Erasmus University Rotterdam, respectively. What's interesting about these programs is that the program page of *Economics* at Erasmus University Rotterdam was viewed more often by people who also viewed *Econometrics* at the university ($S_{j,k}$ is higher than for *Econometrics* at the University of

Amsterdam). However, most of these visitors viewed a large number of programs ($s_{j,k}^w$ is lower than for Econometrics at the University of Amsterdam), indicating that these were perhaps visitors more uncertain about their choice. Moving slightly down the ranks, *Engineering* at Delft University of Technology shows up at the eighteenth place. While not as close of a substitute as the Econometrics and Economics programs, we observe that 6.4% of individuals who visited the program page for Econometrics at Erasmus University Rotterdam, also visit the program page for Engineering at Delft University of Technology, which means they are definitely substitutes to some degree. This makes sense, as both programs require a very high level of mathematical ability to complete. *Medicine* shows up at the twentieth place. Again, not such an obvious substitute, but it is likely that high ability students are interested in both Econometrics – which is regarded as difficult – and Medicine – which is tough to get into without stellar grades.

To further test the credibility of the identified substitutes, Table 2.B.2 in the Appendix shows how similar programs' substitutes are based on their rank. We would expect that close substitutes are often (i) at the same university as the study program considered or (ii) the same degree program at a different university. Columns (1) and (2) of Table 2.B.2 show that 41.9% and 33.4% of the closest substitutes according to our metric are at the same university and the same university location⁷, respectively. This number drops by 6.39 and 4.18 percentage points for rank 2, and continues to drop until rank 20. Still, over 30% of the substitutes at rank 20 are at the same university, and over 25% at the same university location. This makes sense, as students will often

⁷Some universities offer educational programs at more than one location.

have some geographic preference, despite the small country size. Column (3) shows the likelihood of the closest substitute being the same degree program. Again, we see that the likelihood of this being true for the number one ranked substitute is quite high, at 43.6%. It drops when we move down the ranks. The drop is noticeably sharper here, as the number of universities one can study a particular degree program at is often limited. What makes these findings even more striking, is that it is impossible for a study program's substitute to be the same degree program at the same university location, as that is the study program itself. Put differently, Columns (2) and (3) are mutually exclusive. Column (4) combines the two and shows the likelihood of a substitute being either at the same university and location, or the same degree program. The number one ranked substitute of a program is either at the same university and location or the same degree program at a different university in 77.1% of cases. These results clearly show that the substitutes we find make sense. However, while 77.1% of closest substitutes are either at the same university location or the same degree program, neither is the case for 22.9% of study programs; almost a fourth of all programs. The substitutability matrix is particularly valuable for these programs, as it allows us to identify substitutes in an organic manner that does not rely on overlap between topics or locations. Columns (5) and (6) further show the mean value of our substitutability metrics by rank.

2.4 The Impact of Satisfaction Scores on Enrollment

2.4.1 Basic fixed effects regressions

Our main goal is to study the relationship between first-year enrollment and (i) a program's own student satisfaction score, (ii) the national degree average for that program, and (iii) the satisfaction scores of its substitutes. To this end, we use fixed effects Poisson regressions. The fixed effects take care of the large between program variance in first-year enrollment, as well as account for the time-invariant differences between programs. We use Poisson regressions as our outcome measure is a count variable. Specifically, we estimate the following equation

$$\text{First-year enrollment}_{j,t} = \beta \text{Satisfaction Score Metrics}'_{j,t-1} + \gamma \mathbf{T}' + \zeta_j + \varepsilon_{j,t}. \quad (2.3)$$

Here, First-year enrollment $_{j,t}$ indicates the number of students enrolled in the first year of program j at time t . **Satisfaction Score Metrics'** $_{j,t-1}$ is a vector and includes the relevant satisfaction score metrics for program j at time $t-1$. \mathbf{T}' is a vector of time dummies, ζ_j are the program fixed effects and $\varepsilon_{j,t}$ is the idiosyncratic error term.

Table 2.B.3 in the Appendix provides a first look into the impact of a program's satisfaction scores, as well as the development of satisfaction scores over time. The first two columns of the Table show the results of a fixed effects Poisson regression of first-year enrollment on

a program's satisfaction score (as a continuous variable and individual dummies, respectively) as well as the academic year. The Table shows that a program's satisfaction score has a positive relation to first-year enrollment in the subsequent years. Column (2) shows that the relationship is fairly linear, except for at the tail ends of the satisfaction score distribution. Columns (1) and (2) further show that programs have grown substantially between 2011 and 2019. Column (3) shows that there is a clear positive time trend in satisfaction scores as well, having increased by almost 0.15 in 2019 compared to 2011 (scores published in June of 2018 and June of 2010). It underlines how important it is that we control for the academic year in all of our subsequent analyses.

The results from Table 2.B.3 show that a program's satisfaction score is indeed positively correlated with first-year enrollment. However, the model is incomplete, as it does not account for the satisfaction scores of substitutes. Table 2.2 provides a more complete picture. It shows the results of fixed effects Poisson regressions of first-year enrollment on a program's own satisfaction score, the degree's national average, as well as a number of metrics for the satisfaction score of its closest substitutes. All of the coefficients can be approximately interpreted as the %-impact of a 0.1 point increase in the relevant satisfaction score on first-year enrollment. Column (1) confirms the positive relationship between a program's own satisfaction score and first-year enrollment. An increase of 0.1 in the satisfaction score is associated with an increase of approximately 1.9% in first-year enrollment. Column (1) also shows that the national average of all programs that offer the same degree is associated with lower first-year enrollment numbers. As shown in Table 2.B.2, the closest substitutes for a program are often the same

degree programs at different universities. A higher national average therefore makes the outside option more attractive, leading to lower first-year enrollment in the considered study program. Columns (2) to (7) of Table 2.2 all add a different metric of the average satisfaction score of a program's substitute. In Columns (2), (3), (4), (5) and (6), we include the average satisfaction score of the closest one, three, five, ten and twenty substitutes, respectively. As we increase the number of substitutes we consider, the impact of the national average decreases, and the impact of the average satisfaction scores of the substitutes becomes more pronounced. In Column (6), where we consider the average satisfaction score of a program's closest twenty substitutes, we see that the impact of the national average is no longer significant. It is likely that the degree-level national average provides a partial picture of the satisfaction scores of a program's closest substitutes. The average satisfaction score of the identified substitutes looks to contain this information as well, crowding out the impact of the national average. It is a testament to the importance of identifying substitutes in the way we do. It may seem counterintuitive that the impact of substitutes' satisfaction scores increases as we include more substitutes; the satisfaction scores of closer substitutes are likely to be more important. However, the measure of the satisfaction score of the closest substitute is much more noisy, potentially leading to attenuation bias. On top of that, the standard deviation decreases as we include more substitutes, as shown in Table 2.1. Table 2.B.4 shows that the increase remains, but is slightly less pronounced when we standardize the program's satisfaction score, national average and the satisfaction score of substitutes. Column (7) of Table 2.2 shows that when we consider a weighted average of the satisfaction score of all substitutes of a program, there seems to be no effect. This measure of substitutes' satisfaction scores may be

inaccurate, however; as we may assign too much weight to very far-away substitutes because we include data from e.g. scraping bots. The negative impact of the national average remains.

Beyond Table 2.2, it is interesting to see whether any particular programs or type of students drive these results. Table 2.B.5 in the Appendix provides a look into these heterogeneous effects. Columns (2) and (3) split the sample into undergraduate and master's programs. For the undergraduate program's we confirm the impact of a program's own satisfaction score on enrollment. While the impact of the satisfaction score of substitutes remains negative, the estimate is statistically insignificant. For master's programs, both coefficients retain the same sign as in Table 2.2, but lose their significance. We cannot conclude that the effects are driven by either undergraduate or master's programs. Columns (4) and (5) of Table 2.B.5 show the impact of the satisfaction score metrics on the number of first-year enrolled men, and the number of enrolled first-year students with a migration background. Both outcomes show similar results to those found in Table 2.2, indicating that neither students of a specific gender, nor students with (or without) a migration background drive our results.

Another potential concern is that our results are driven by small programs. Since fewer students from these programs fill out the national student survey, these programs naturally have larger variance in satisfaction scores. Moreover, small changes in student numbers may be percentually large. To ensure our results are robust to excluding programs from a certain size, Table 2.B.6 in the Appendix shows the results when we exclude different programs based on their average size over the years. Columns (1) to (4) show that the results hold when we

Table 2.2: Relationship between satisfaction score, national average, satisfaction score of substitutes & first-year enrollment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	W/o substitutes	No. 1 subst.	Top 3. subst.	Top 5 subst.	Top 10 subst.	Top 20 subst.	All substitutes (weighted)
Dep. var: first-year enrollment							
Satisfaction score of program ($\times 10$)	0.0191*** (0.00330)	0.0191*** (0.00329)	0.0193*** (0.00327)	0.0194*** (0.00328)	0.0195*** (0.00328)	0.0198*** (0.00329)	0.0191*** (0.00329)
National average ($\times 10$)	-0.0126** (0.00607)	-0.0130** (0.00611)	-0.0112* (0.00605)	-0.00975 (0.00605)	-0.00876 (0.00622)	-0.00732 (0.00612)	-0.0125** (0.00601)
Satisfaction Score Substitute(s) ($\times 10$)		0.00115 (0.00277)	-0.00465 (0.00488)	-0.0104* (0.00595)	-0.0171** (0.00753)	-0.0315*** (0.0112)	-0.00169 (0.00790)
Year dummies	Yes						
Observations	12821	12821	12821	12821	12821	12821	12821

Note: Regressions are at study program-year level, estimated by fixed effects Poisson regression. All measures of satisfaction scores are multiplied by 10. This means coefficients can be interpreted as the impact of an increase of 0.1. *Satisfaction score of program* denotes a program's own satisfaction score. *National average* denotes the national average of the same degree program at all universities that offer it. *Satisfaction score of substitutes* indicates the average satisfaction score of all substitutes that are ranked above the number indicated in Column titles (2) to (7). Standard errors are clustered at the study program level. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$.

only include the largest 95%, 90%, 75% and 50% in our analysis sample. Columns (5) to (8) show that the same is true for when we only include the smallest 50%, 75%, 90% and 95%. It proves that our results are not driven by programs at the tail ends of the size distribution, but are actually present across different program sizes.

All these results point towards a systematic relationship between programs' satisfaction scores and first-year enrollment. However, we cannot interpret these results as causal. It is likely that endogeneity still plagues these estimates. For instance, an increase in a study program's budget is likely to be spent on both facilities and marketing; the former increasing student satisfaction, the latter first-year enrollment. In the next Section, we will therefore attempt to identify the causal impact of these satisfaction scores by exploiting rounding discontinuities.

2.4.2 Identification of causal impacts

2.4.2.1 Rounding discontinuities

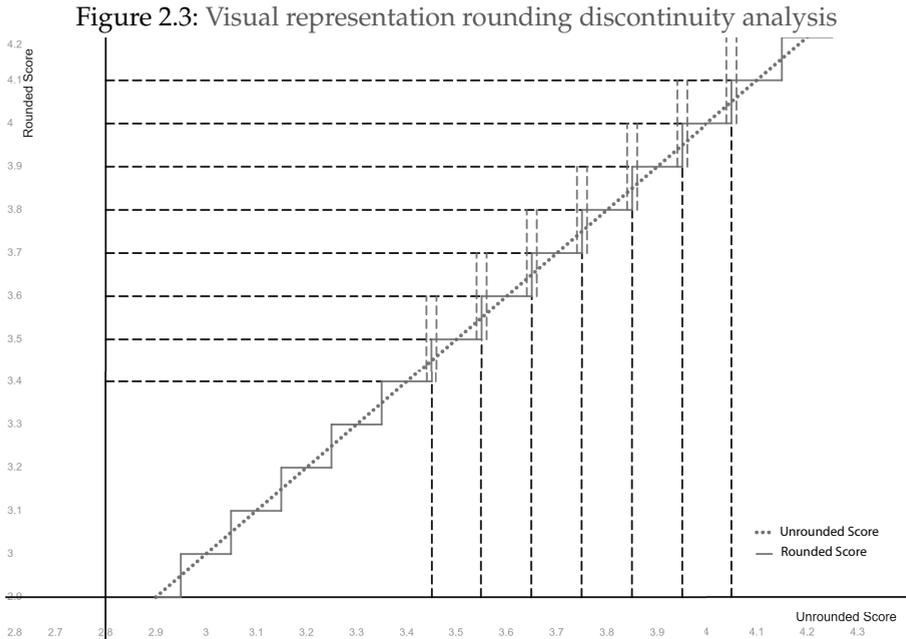
To identify the causal impact of satisfaction scores, we exploit the natural discontinuities that come with the way in which the programs' satisfaction scores are rounded before being published on the Studiekeuze123 website. Satisfaction scores are rounded to the nearest tenth. More formally, the rounded satisfaction score of a program is calculated as follows:

$$\frac{1}{10} \left\lfloor 10 \times (\text{Satisfaction score}) + \frac{1}{2} \right\rfloor, \quad (2.4)$$

where $\lfloor x \rfloor$ denotes the floor function, returning the largest integer that is smaller than, or equal to x .

The fact that Studiekeuze123 rounds these satisfaction scores before publication means that while the raw satisfaction scores of two programs may be very close, they are a tenth apart on the Studiekeuze123 website. We can exploit this rounding discontinuity to study the effect of having a higher published satisfaction score. We compare programs when they ended up just below to when they ended up just above a particular threshold, again using a fixed effects Poisson regression. Figure 2.3 provides a visual representation of the analysis we do. The dotted line represents the unrounded satisfaction score of a program. The solid line represents the rounded satisfaction score. At each $.05^{th}$, the rounded and published satisfaction scores makes a discontinuous jump to the next tenth. We exploit this by comparing programs very close to this discontinuity, indicated by the dashed gray bandwidths.

These programs have very similar unrounded satisfaction scores, but their published satisfaction scores differs. We once again estimate our specifications using fixed effects Poisson regressions. For this analysis, we also include dummies for every threshold, as well as polynomials of the distance to the relevant threshold.



Note: Visual representation of discontinuity analysis. Dotted line indicates unrounded score. Solid line indicates rounded score. Dashed lines indicate bandwidth around the rounding threshold. We consider different bandwidths in our analysis.

Table 2.3 shows the impact of being just above the rounding threshold, compared to just below, on subsequent enrollment of first year students for multiple bandwidths around the rounding thresholds. Our estimates show a positive impact of being just above the threshold

Table 2.3: Impact of being just above satisfaction score rounding threshold on first-year enrollment

	B = 0.01	B = 0.02	B = 0.03	B = 0.04	B = 0.05
Dep. Var: first-year enrollment					
Rounded up	0.0325 (0.0346)	0.0352* (0.0209)	0.0295* (0.0159)	0.0178 (0.0136)	0.0170 (0.0120)
Year dummies	Yes	Yes	Yes	Yes	Yes
Threshold dummies	Yes	Yes	Yes	Yes	Yes
Distance polynomials	Yes	Yes	Yes	Yes	Yes
Observations	1645	4127	6716	9102	10902
Rounded up	879	2131	3409	4601	5474

Note: regressions estimated by fixed effects Poisson regression. Standard errors are clustered at the study program level. *Rounded up* indicates whether a program was rounded up (i.e., above the rounding threshold). Distance polynomials include a linear and squared term of the distance to the closest rounding threshold. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$.

of up to 3.52%, but are imprecise. For bandwidths of 0.02 and 0.03 around the threshold, the effect is marginally significant; for larger bandwidths, the estimates decrease and are statistically insignificant.

2.4.2.2 The degree-level average

The results from Section 2.4.1 indicate that not just a program's own satisfaction score, but also that of competing programs may be important. As discussed in Section 2.2, Studiekeuze123 often shows a program's satisfaction score next to the national average for the degree, making it a very salient metric to judge a program by. In this Section, we expand our analysis of the impact of the national average by estimating a more flexible specification compared to that in Table 2.2. In

addition, we again exploit rounding discontinuities to investigate how being just above the national average impacts enrollment.

Column (2) of Table 2.4 shows the results of the extended specification. We confirm that the national average has a negative impact on enrollment, but do not observe any nonlinearities, or differences in coefficient size depending on whether a program's student satisfaction score is above or below the national average.

To once again exploit the rounding discontinuity, we take a sample of programs whose unrounded satisfaction score is within a certain bandwidth from the unrounded national average. To avoid conflating the effects we find with the effects found in Section 2.4.2.1, we only consider programs that are above the rounding threshold. The main difference between this analysis and that in Section 2.4.2.1, is that the likelihood of being 'treated' (i.e., rounded up over the national average) is not $\frac{1}{2}$. Narrow bandwidths cause both the program's satisfaction score as well as the national average to be well over the rounding threshold most of the time, meaning they will often be equal in publication.⁸ With a larger bandwidth, the likelihood of a program being rounded up over the national average increases exponentially. It is therefore valuable to consider a relatively large bandwidth. We use a bandwidth of 0.05, which means we have 193 treated observations; just over 10% of the sample. Column (3) shows no effect of being rounded up over the degree-level national average. While this may be surprising given the saliency of the degree-level national average, it is important to note that the only programs we consider are programs that are average for

⁸To illustrate this, consider the bandwidth to be .01. Only when the program's satisfaction score is between $x.x5$ and $x.x6$ is there a chance that the national average is just below the rounding threshold.

Table 2.4: Impact of being just above national average satisfaction score on first-year enrollment

	(1)	(2)	(3)
	Base specification	Extended specification	Rounded up over (B = 0.05)
Dep. Var: First Year Enrollment			
Satisfaction score of program (×10)	0.0191*** (0.00330)	0.0235*** (0.00688)	
National average (×10)	-0.0126** (0.00607)	-0.0176** (0.00864)	
Above national average		-0.00398 (0.0134)	0.0138 (0.0480)
Above national average × Distance to national average		-0.00286 (0.0104)	
(Distance to national average) ²		0.00389 (0.0120)	
Above national average × (Distance to national average) ²		-0.0133 (0.0147)	
Year dummies	Yes	Yes	Yes
Observations	12821	12821	1922
N Above National Average			193

Note: regressions estimated by fixed effects Poisson regression. *Satisfaction score of program* denotes a program’s own satisfaction score. *National average* denotes the national average of the same degree program at all universities that offer it. *Above national average* indicates whether program was rounded up over the national average or not. *Distance to national average* is the difference between a program’s own satisfaction score and the degree level average. Its singular term is omitted, as it is collinear with the first two terms. *N Above National Average* indicates how many observations were treated. Distance polynomials include a linear and squared term of the distance to the closest rounding threshold. Standard errors are clustered at the study program level. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$.

their degree. Students will thus often have an outside option that provides a better satisfaction score, potentially muting the results.

2.4.2.3 Rounding discontinuities and close substitutes

Students may not only compare a program’s satisfaction score to the national average, but are likely to also compare satisfaction scores between programs. Here, we re-use the methodology described in Sec-

tion 2.4.2.2, but compare a program's satisfaction score to that of its closest substitutes instead of the national average. We analyze situations in which a program is just above the rounding threshold, and very close to (a number of) its closest substitute(s). The advantage is that this avoids the issue of all considered programs being very average. We once again take a bandwidth of 0.05 for this analysis, as it suffers from the same issue as our analysis of the national average: while the sample size decreases linearly with the bandwidth, the share of treated observations becomes exponentially smaller when the bandwidth decreases.

Table 2.5 shows the impact of being rounded up over at least one close substitute. We find that all of our estimates are positive, but tend to become smaller as we include more substitutes. The point estimate of a program being rounded up over its number one ranked substitute in Column (1) is positive, but not statistically significant. However, we have only 1041 observations, of which a mere 139 are treated. Column (2) shows a clearly positive and significant effect of being rounded up over at least one of the top 3 substitutes of 4.37%. A sizeable effect. Column (3) shows a qualitatively similar result, although the point estimate decreases slightly. The point estimates in Columns (4) and (5) are smaller, and not significant, or only marginally. The reason for this may be that the treatment variable 'close call win' is diluted. Being rounded up over the twentieth substitute is likely to have less of an impact to being rounded up over one of the top three substitutes of any one program.

Table 2.5: Impact of being rounded up over close substitute on first-year enrollment

	(1)	(2)	(3)	(4)	(5)
	No. 1 subst.	Top 3 subst.	Top 5 subst.	Top 10 subst.	Top 20 subst.
Dep. var: first-year enrollment					
Close call win	0.0155 (0.0300)	0.0437** (0.0175)	0.0426*** (0.0160)	0.0157 (0.0108)	0.0148* (0.00876)
Year dummies	Yes	Yes	Yes	Yes	Yes
Threshold dummies	Yes	Yes	Yes	Yes	Yes
Distance polynomials	Yes	Yes	Yes	Yes	Yes
Observations	1041	3090	4552	6496	7840
N at Least one Win	139	476	760	1362	2116

Note: regressions estimated by fixed effects Poisson regression. Only includes programs when they were above the rounding threshold. *Close call win* indicates whether a program was rounded up over at least one close substitute. Columns indicate how many substitutes were considered. Distance polynomials include a linear and squared term of the distance to the closest rounding threshold. Standard errors are clustered at the study program level.

*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$.

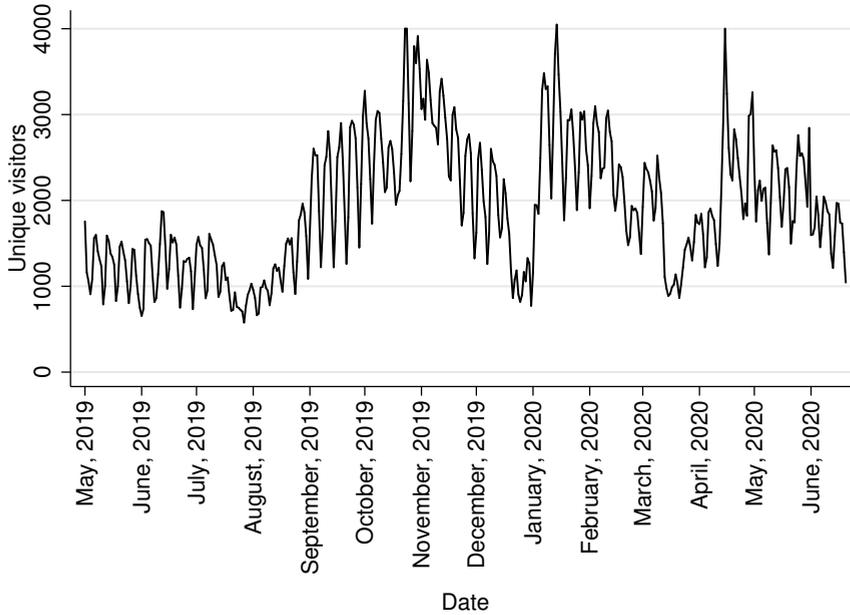
2.5 Conclusion

In this Chapter, we analyzed the impact of published student satisfaction scores on first-year enrollment in university programs. We show that satisfaction scores matter for enrollment: both the satisfaction score of the program itself as well as that of its substitutes. Analyses exploiting rounding discontinuities show that a satisfaction score being rounded up to the next tenth increases first-year enrollment in the subsequent year by 1.70% to 3.52%, although the estimates are imprecise. We further find that a program's relative student satisfaction score is of particular importance: conditional on being rounded up, a program that has a (slightly) higher published satisfaction score than at least one of its top substitutes will see an increase in first-year enrollment in the subsequent year of up to 4.37% on average.

Our findings underline the importance of digital information provision to students. However, it should perhaps also serve as a warning sign of how to present information. The fact that we find that substitutes being on different sides of rounding discontinuities has a significant and sizeable impact on the choices between these programs is not necessarily positive. It implies that students sometimes make a decision on which program to enroll in based on a difference that is more of an artefact of publication than a truly large difference in student satisfaction. Informational websites may thus want to emphasize the role of uncertainty in these types of metrics and urge students to interpret all information with care.

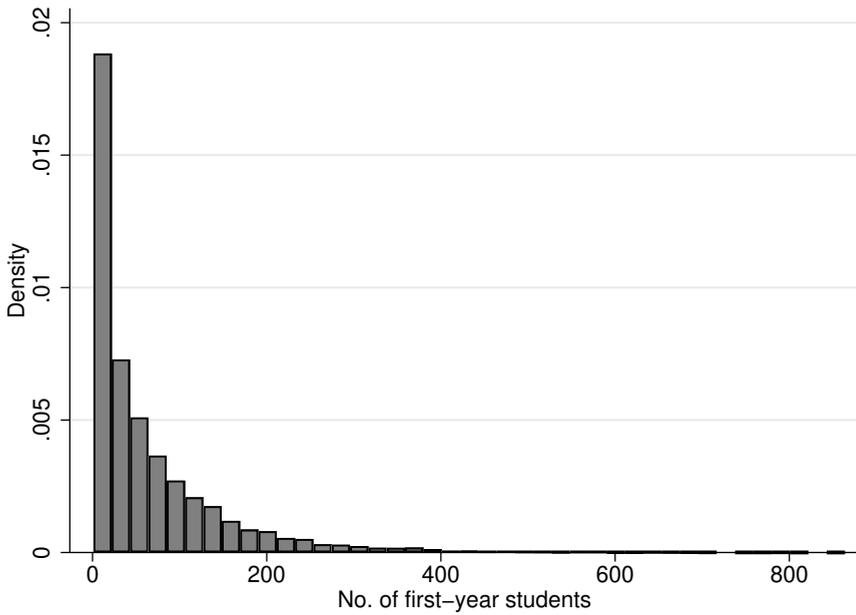
Appendix 2.A Additional Figures

Figure 2.A.1: Daily visitors to Studiekeuze123.nl



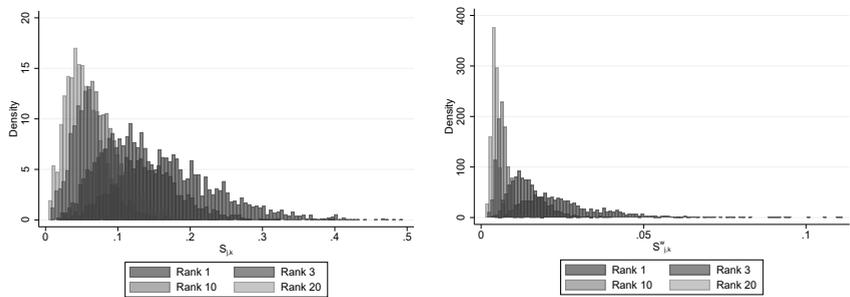
Note: Figure shows number of daily visitors to Studiekeuze123.nl between May of 2019 and June of 2020. Four values were (much) larger than 4,000. They have been winsorized to 4,000 to make the graph more easily readable. Original values: 14th of January, 2020: 4,047 visitors; 24th of October, 2019: 5,404 visitors; 15th of April, 2020: 7,548 visitors; 23rd of October, 2019: 18,111 visitors.

Figure 2.A.2: Distribution of no. of first-year enrolled students



Note: Figure shows distribution of no. of first-year students enrolled in programs. Data at program-year level.

Figure 2.A.3: Distribution of $S_{j,k}$ and $S_{j,k}^w$



(a) Distribution of $S_{j,k}$ by rank (b) Distribution of $S_{j,k}^w$ by rank

Note: Figure shows distribution of substitutability metrics $S_{j,k}$ and $S_{j,k}^w$ as explained in Section 2.3.3 by rank of substitute (based on $S_{j,k}^w$).

Appendix 2.B Additional Tables

Table 2.B.1: Substitutes for Econometrics at Erasmus University Rotterdam (3,528 views)

Substitute	Substitute rank	Times also viewed	$S_{j,k}$	Weighted views	$S_{j,k}^w$
Econometrics - Free University Amsterdam	1	683	0.194	86.02	0.024
Econometrics - University of Amsterdam	2	428	0.121	65.59	0.019
Economics - Erasmus University Rotterdam	4	553	0.157	61.06	0.017
Engineering - Delft University of Technology	18	225	0.064	14.86	0.004
Medicine - Erasmus University Rotterdam	20	150	0.043	13.50	0.004

Note: *Times also viewed* indicates the number of times an individual visitor visited both the program page for Econometrics at Erasmus University Rotterdam and that of its substitute. $S_{j,k}$ displays the share of individuals who visit the substitute's program page, conditional on visiting the program page for Econometrics at Erasmus University Rotterdam. *Weighted views* penalizes a visitor's view by $\frac{1}{\sum_i \text{Views}_i}$ (1 over the number of pages the visitor viewed in total). $S_{j,k}^w$ displays the weighted share of individuals who visit the substitute's program page, conditional on visiting the program page for Econometrics at Erasmus University Rotterdam. Rank is determined by $S_{j,k}^w$.

Table 2.B.2: Similarity of programs and their substitutes

	Same Uni	Same Uni & Loc.	Same Degree	Same Uni & Loc. or Same Degree	$S_{j,k}$	$S_{j,k}^w$
Rank 1 (Constant)	0.420*** (0.00999)	0.334*** (0.00955)	0.436*** (0.0100)	0.771*** (0.00851)	0.176*** (0.00147)	0.0260*** (0.000265)
Rank 2	-0.0639*** (0.0139)	-0.0418** (0.0133)	-0.0516*** (0.0141)	-0.0934*** (0.0127)	-0.0346*** (0.00188)	-0.00810*** (0.000308)
Rank 3	-0.0680*** (0.0139)	-0.0340* (0.0133)	-0.118*** (0.0138)	-0.152*** (0.0130)	-0.0542*** (0.00179)	-0.0117*** (0.000288)
Rank 4	-0.0717*** (0.0139)	-0.0307* (0.0133)	-0.151*** (0.0136)	-0.182*** (0.0131)	-0.0680*** (0.00174)	-0.0140*** (0.000280)
Rank 5	-0.0537*** (0.0140)	-0.00451 (0.0135)	-0.215*** (0.0131)	-0.220*** (0.0132)	-0.0755*** (0.00170)	-0.0155*** (0.000275)
Rank 6	-0.0627*** (0.0139)	-0.0209 (0.0134)	-0.224*** (0.0130)	-0.245*** (0.0132)	-0.0839*** (0.00168)	-0.0166*** (0.000272)
Rank 7	-0.0627*** (0.0139)	-0.0164 (0.0134)	-0.257*** (0.0127)	-0.274*** (0.0132)	-0.0879*** (0.00167)	-0.0175*** (0.000270)
Rank 8	-0.0627*** (0.0139)	-0.0164 (0.0134)	-0.289*** (0.0124)	-0.305*** (0.0132)	-0.0938*** (0.00164)	-0.0182*** (0.000269)
Rank 9	-0.0652*** (0.0139)	-0.0160 (0.0134)	-0.305*** (0.0122)	-0.320*** (0.0132)	-0.0979*** (0.00164)	-0.0188*** (0.000268)
Rank 10	-0.0516*** (0.0140)	-0.00943 (0.0135)	-0.307*** (0.0121)	-0.316*** (0.0132)	-0.102*** (0.00163)	-0.0193*** (0.000268)
Rank 11	-0.0664*** (0.0139)	-0.0119 (0.0134)	-0.333*** (0.0118)	-0.345*** (0.0131)	-0.104*** (0.00162)	-0.0197*** (0.000267)
Rank 12	-0.0754*** (0.0139)	-0.0254 (0.0134)	-0.331*** (0.0118)	-0.356*** (0.0131)	-0.107*** (0.00161)	-0.0200*** (0.000267)
Rank 13	-0.0656*** (0.0139)	-0.0193 (0.0134)	-0.345*** (0.0116)	-0.364*** (0.0131)	-0.110*** (0.00160)	-0.0203*** (0.000267)
Rank 14	-0.0783*** (0.0139)	-0.0336* (0.0133)	-0.359*** (0.0114)	-0.392*** (0.0130)	-0.112*** (0.00161)	-0.0206*** (0.000266)
Rank 15	-0.0902*** (0.0138)	-0.0398** (0.0133)	-0.359*** (0.0114)	-0.399*** (0.0130)	-0.114*** (0.00159)	-0.0209*** (0.000266)
Rank 16	-0.0996*** (0.0138)	-0.0533*** (0.0132)	-0.359*** (0.0114)	-0.413*** (0.0129)	-0.115*** (0.00159)	-0.0211*** (0.000266)
Rank 17	-0.107*** (0.0137)	-0.0631*** (0.0131)	-0.362*** (0.0114)	-0.425*** (0.0129)	-0.117*** (0.00158)	-0.0213*** (0.000266)
Rank 18	-0.109*** (0.0137)	-0.0631*** (0.0131)	-0.370*** (0.0112)	-0.434*** (0.0128)	-0.119*** (0.00158)	-0.0215*** (0.000266)
Rank 19	-0.0959*** (0.0138)	-0.0537*** (0.0132)	-0.372*** (0.0112)	-0.426*** (0.0128)	-0.121*** (0.00157)	-0.0216*** (0.000266)
Rank 20	-0.117*** (0.0137)	-0.0734*** (0.0131)	-0.371*** (0.0112)	-0.444*** (0.0128)	-0.122*** (0.00157)	-0.0218*** (0.000266)
Observations	48800	48800	48800	48800	48800	48800

Note: regressions estimated by ordinary least squares using the full set of identified substitutes. Columns display the dependent variables. The outcome variable Columns (1) to (4) is a dummy indicating whether the substitute at a certain rank is at the same university (1), at the same university and university location (2), the same degree program (3) and either at the same university and location or the same degree program (4). The dependent variables Columns (5) and (6) are the substitutability metrics, as explained in Section 2.3.3. Substitutes are ranked based on $S_{j,k}^w$. The number one ranked substitute is used as the baseline. Standard errors are clustered at the study program level. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$.

Table 2.B.3: Satisfaction scores & first-year enrollment over time

	(1)	(2)	(3)
	Enrollment	Enrollment	Satisfaction score
Satisfaction score of program ($\times 10$)	0.0159*** (0.00311)		
2012	0.00156 (0.00988)	0.00360 (0.0103)	-0.201*** (0.0708)
2013	0.0751*** (0.0117)	0.0729*** (0.0122)	0.207*** (0.0706)
2014	0.0298** (0.0133)	0.0361** (0.0140)	0.235*** (0.0698)
2015	-0.00461 (0.0163)	0.000143 (0.0167)	0.528*** (0.0700)
2016	0.0570*** (0.0162)	0.0609*** (0.0166)	1.179*** (0.0687)
2017	0.112*** (0.0172)	0.116*** (0.0175)	1.465*** (0.0684)
2018	0.140*** (0.0180)	0.145*** (0.0185)	1.552*** (0.0676)
2019	0.151*** (0.0174)	0.156*** (0.0180)	1.259*** (0.0675)
Satisfaction score: 3.6		-0.00447 (0.0222)	
Satisfaction score: 3.7		0.00994 (0.0219)	
Satisfaction score: 3.8		0.0191 (0.0231)	
Satisfaction score: 3.9		0.0392* (0.0230)	
Satisfaction score: 4.0		0.0486* (0.0248)	
Satisfaction score: 4.1		0.0790*** (0.0265)	
Satisfaction score: 4.2		0.116*** (0.0297)	
Satisfaction score: 4.3		0.146*** (0.0329)	
Satisfaction score: 4.4		0.133*** (0.0363)	
Satisfaction score: 4.5		0.228*** (0.0667)	
Constant			38.57*** (0.0527)
Observations	12969	12241	13073

Note: regressions are at study program-year level. Columns (1) and (2) are estimated by fixed effects Poisson regressions. Column (3) by ordinary least squares. *Satisfaction score of program* denotes a program's own satisfaction score. Standard errors are clustered at the study program level. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$.

Table 2.B.4: Standardized satisfaction score, national average, satisfaction score of substitutes & first-year enrollment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	W/o substitutes	No. 1 subst.	Top 3. subst.	Top 5 subst.	Top 10 subst.	Top 20 subst.	All substitutes (weighted)
Dep. var: first-year enrollment							
Satisfaction score of program (standardized)	0.0532*** (0.00918)	0.0531*** (0.00916)	0.0537*** (0.00911)	0.0539*** (0.00913)	0.0544*** (0.00913)	0.0552*** (0.00915)	0.0532*** (0.00915)
National average (standardized)	-0.0375** (0.0180)	-0.0386** (0.0181)	-0.0332* (0.0180)	-0.0290 (0.0180)	-0.0260 (0.0185)	-0.0217 (0.0182)	-0.0372** (0.0179)
Satisfaction Score Substitute(s) (standardized)		0.00296 (0.00711)	-0.00883 (0.00927)	-0.0174* (0.00998)	-0.0251** (0.0110)	-0.0413*** (0.0147)	-0.00246 (0.0115)
Year dummies	Yes						
Observations	12821	12821	12821	12821	12821	12821	12821

Note: regressions are at study program-year level, estimated by fixed effects Poisson regression. All measures of satisfaction scores are standardized. This means coefficients can be interpreted as the impact of an increase of one standard deviation. *Satisfaction score of program* denotes a program's own satisfaction score. *National average* denotes the national average of the same degree program at all universities that offer it. *Satisfaction score of substitutes* indicates the average satisfaction score of all substitutes that are ranked above the number indicated in Column titles (2) to (7). Standard errors are clustered at the study program level. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$.

Table 2.B.5: Heterogeneity of relationship between satisfaction score, national average, satisfaction score of substitutes & first-year enrollment

	(1)	(2)	(3)	(4)	(5)
	Base result	Undergraduate	Master's	Men	Migration background
Dep. var: first-year enrollment					
Satisfaction score of program ($\times 10$)	0.0198*** (0.00329)	0.0231*** (0.00372)	0.00448 (0.00514)	0.0220*** (0.00385)	0.0151*** (0.00440)
National average ($\times 10$)	-0.00732 (0.00612)	-0.00640 (0.00699)	-0.00303 (0.0103)	-0.00209 (0.00667)	-0.0114 (0.00964)
Satisfaction Score Substitute(s) ($\times 10$)	-0.0315*** (0.0112)	-0.0196 (0.0123)	-0.0317 (0.0274)	-0.0329** (0.0131)	-0.0815*** (0.0171)
Year dummies	Yes	Yes	Yes	Yes	Yes
Observations	12821	8919	3902	12783	12793

Note: regressions are at study program-year level, estimated by fixed effects Poisson regression. All measures of satisfaction scores are multiplied by 10. This means coefficients can be interpreted as an increase of 0.1. *Satisfaction score of program* denotes a program's own satisfaction score. *National average* denotes the national average of the same degree program at all universities that offer it. *Satisfaction score of substitutes* indicates the average satisfaction score of the top 20 substitutes. Columns (1) and (2) show results for undergraduate and master's programs, respectively. Columns (3) and (4) take the number of first-year students who are men, and the number of first-year students who have a migration background as the dependent variable, respectively. Standard errors are clustered at the study program level. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$.

Table 2.B.6: Robustness check for size; relationship between satisfaction score, national average, satisfaction score of substitutes & first-year enrollment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Largest 95%	Largest 90%	Largest 75%	Largest 50%	Smallest 50%	Smallest 75%	Smallest 90%	Smallest 95%
Dep. var: first-year enrollment								
Satisfaction score of program ($\times 10$)	0.0199*** (0.00330)	0.0200*** (0.00332)	0.0206*** (0.00348)	0.0220*** (0.00410)	0.0117*** (0.00374)	0.0124*** (0.00336)	0.0159*** (0.00346)	0.0171*** (0.00331)
National average ($\times 10$)	-0.00737 (0.00613)	-0.00735 (0.00622)	-0.00646 (0.00651)	-0.00538 (0.00758)	-0.0142* (0.00787)	-0.00377 (0.00717)	0.00146 (0.00670)	-0.00347 (0.00641)
Satisfaction Score Substitute(s) ($\times 10$)	-0.0314*** (0.0112)	-0.0311*** (0.0112)	-0.0304*** (0.0115)	-0.0295** (0.0127)	-0.0521*** (0.0177)	-0.0618*** (0.0133)	-0.0477*** (0.0117)	-0.0386*** (0.0111)
Year dummies	Yes							
Observations	12373	11832	10020	6757	6046	9420	11460	12133

Note: regressions are at study program-year level, estimated by fixed effects Poisson regression. All measures of satisfaction scores are multiplied by 10. This means coefficients can be interpreted as an increase of 0.1. *Satisfaction score of program* denotes a program's own satisfaction score. *National average* denotes the national average of the same degree program at all universities that offer it. *Satisfaction score of substitutes* indicates the average satisfaction score of the top 20 substitutes. Columns indicate which programs are included in the analysis. Standard errors are clustered at the study program level. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$.

3

Correcting Beliefs about Job Opportunities and Wages: A Field Experiment on Education Choices

Abstract

We run a large-scale field experiment in which we provide information to students at randomly selected schools about the job opportunities and hourly wages of a small set of occupations they are interested in. The experiment takes place on an online career guidance counseling platform that is widely used in the Netherlands, and involves 28,267 pre-vocational secondary education students in 243 schools over a period of 2 years. We find that the information improves the accuracy of students' beliefs, both in the short run (for job opportunities and hourly wages) and in the long run (for job opportunities only). Students who receive the information also tend to change their favorite occupation 0.88 to 2.16 percentage points more often, and switch towards an occupation with better labor market prospects if they do so. Last, and most importantly, they select secondary school specializations related to occupations with better labor market prospects (1.5% and 0.3% – €0.05 an hour – higher than the control group mean for job opportunities and wages, respectively) and choose post-secondary education programs with higher expected wages (2.5% – approximately €0.40 an hour – higher than the control group mean).

This Chapter is based on joint work with Robert Dur and Didier Fouarge. We thank Mirjam Bahlmann, Robert Jan van Egmond, Dorothy Pillen-Warmerdam, Frank Tinkenberg and Koen Weide from Qompas for their cooperation. We also thank Marc van der Steeg. This project received funding from the Dutch Ministry of Education, Culture, and Science. Fouarge further acknowledges the financial support of NRO, UWV, Randstad, S-BB (grant: 405-17-900). This study is registered in the AEA RCT Registry and the digital object identifier (DOI) is: <https://doi.org/10.1257/rct.3220>.

3.1 Introduction

Each year, millions of teenagers around the world face a choice that has far-reaching consequences, both for themselves and for society: the choice of post-secondary education program. This choice is important for themselves, as the program from which they earn a degree is an important determinant of future labor market outcomes (see e.g. Bleemer and Mehta, 2022; Ketel et al., 2016; Kirkeboen et al., 2016). It is also important for society, as it affects future shortages and excess supply of labor in important occupations.

Despite its huge importance, students often decide on their field of study without having accurate information about the labor market prospects of different programs (Baker et al., 2018; Conlon, 2019; Hastings et al., 2015, 2016; Pekkala Kerr et al., 2015) and careers (Arcidiacono, Hotz, & Kang, 2012; Betts, 1996). As a result, many teenagers end up choosing programs that have a bleak outlook, both in terms of job opportunities and wages.

To help students make better choices, several large-scale field experiments have tested whether providing information to students about labor market prospects makes a meaningful difference in students' educational choices. The results of these experiments tend to be sobering. Even though students' choices move in the direction of education programs with better labor market prospects, the size of these effects tends to be limited, not seldomly statistically indistinguishable from zero (see e.g., Bonilla-Mejía et al., 2019; Conlon, 2019; Hastings et al., 2015; Pekkala Kerr et al., 2015). A possible reason might be that the information provided in earlier experiments is too coarse: interventions commonly provide information about the labor market returns

to enrolling in university, or about different majors rather than about occupations. While majors are an important determinant of future earnings (Altonji, Arcidiacono, & Maurel, 2016), subsequent occupational choices explain a large part of the difference in earnings between majors (Altonji, Blom, & Meghir, 2012) and students seem to be well aware of that (Arcidiacono et al., 2012). Hence, a promising next step in this literature is to provide more fine-grained information to students about the labor market prospects of occupations.

In this Chapter, we report the results of a low-cost field experiment in which we provide students in randomly selected schools with personally targeted information about the labor market prospects of a small set of occupations they are interested in. To our knowledge, we are the first to do so. We study whether the information leads students to correct their beliefs about the labor market prospects of these occupations, shifts students' preferences over occupations, and influences their education choices. Our multi-year field experiment involves 28,267 students at 243 different schools for pre-vocational education in the Netherlands. The students are in grades 8 to 10¹ and generally are between 13 and 16 years old.

The field experiment takes place on an online career guidance counseling platform that is market leader among pre-vocational secondary education schools in the Netherlands. On the platform, students do numerous assignments to find out what they like, what they are good at and, ultimately, which occupations would be a good fit for them. As part of one of these assignments, students also take an extensive occupation test. This test results in a short-list of twenty (out of 353)

¹The second to fourth year of pre-vocational secondary education in the Netherlands.

occupations that fit the students' interests best according to the answers they provide. Students take part in our experiment right after this test.

Our experiment proceeds as follows. First, we ask students in which secondary-school specializations (called: "profiles") they are most interested. Next, we show students their shortlist of twenty occupations and ask them to select from it the five that they like most. We then ask them to state their beliefs about the job opportunities and hourly wages for these five occupations, and to rank them based on how much they would like to work in them. Subsequently, we provide students of randomly selected schools with information about the job opportunities and, for a random subset of these schools, hourly wages of the selected occupations. Students at the remaining schools do not receive any information and form our control group. To learn whether it matters who provides the information (the 'sender'), we mention to some students that the information is provided by a labor market research institute, whereas we mention to others that a specific researcher from this institute – who is either male or female and experienced or inexperienced – provides the information. The identity of the sender is randomized within the treatment group.

Next, students in both the treatment and the control group watch a video, get the opportunity to update their stated beliefs, re-rank their preferred occupations and update their interest in the different profiles. These answers are our first set of outcome measures. In addition to these data, we obtain (i) post-experimental survey data (up to one and a half year later) on the above mentioned beliefs and preferences, as well as their post-secondary education choices, and (ii) administra-

tive data at the school level on students' profile choices. All our analyses follow the design we registered in the AEA RCT Registry prior to the start of the experiment, except where indicated.²

In line with the earlier studies cited above, our results show that students have highly inaccurate beliefs about the job opportunities and hourly wages of the occupations that they like. They tend to overestimate both. Moreover, the interest of a student in an occupation is strongly positively correlated with the student's expectations about the occupation's job opportunities and hourly wages. This suggests that students' beliefs about occupations' labor market prospects play a role in students' occupational aspirations, underlining the potential importance of providing accurate information about these prospects.

Our information intervention is effective in correcting beliefs. In the short run, treated students overestimate the job opportunities and hourly wages to a smaller degree, make smaller absolute errors, and are more likely to hold correct beliefs. The improved accuracy is mostly driven by students correcting overestimations. Our post-experimental survey data show that these effects partly persist: those who received the information in their final school year have more accurate expectations about the job opportunities up to seven months later.

We also find evidence that the treatment increases the likelihood that students change their favorite occupation. If students do so, they tend to substitute the initial occupation for one with better job opportunities or hourly wages. We do not find evidence that this ranking persists in the survey fielded after the experiment. However, this may be driven by selection into the survey. The sample of surveyed students differed

²Digital object identifier: <http://doi.org/10.1257/rct.3220>.

from the full sample in the experiment in that the former was less likely to change their favorite occupation for one with better prospects during the experiment than the latter.

In contrast to our predictions, we find no evidence that the treatment impacts students' intended choice of profile right after the experiment. However, administrative data at the school level shows that in treated schools, students select profiles associated with occupations with better job opportunities and higher earnings. A possible interpretation of these two findings is that treated students need some time to make up their mind and discuss the obtained information with their parents or peers before they actually revise their profile choice. Survey data collected from students graduating secondary education shows that students who received information about hourly wages are more likely to choose a study program with higher earning prospects. This further shows that the treatment is indeed effective in altering education choices.

The identity of the sender of the information that is mentioned in the intervention — the labor market research institute or a researcher from this institute, either senior or junior, either female or male — appears inconsequential for the subsequent beliefs or preference ranking of occupations.

Our study contributes to a growing body of literature on the role of labor market expectations in education choices. Studies have invariably found that students have highly noisy beliefs about the labor market returns of different study programs (Baker et al., 2018; Hastings et al., 2015, 2016; Pekkala Kerr et al., 2015; Conlon, 2019) and earnings in different careers (Arcidiacono et al., 2012; Betts, 1996). Students who are

more concerned with the labor market prospects of programs, however, are less likely to overestimate these prospects (Hastings et al., 2016). The differences in concerns about these prospects are large between men and women (Wiswall & Zafar, 2017; Zafar, 2013). Men tend to care more about pecuniary outcomes, whereas women care more about job security and flexibility. Similarly, we find in our data that male students select occupations with better job opportunities and higher hourly wages. However, they are also more likely to overestimate these and make larger absolute errors; contrasting Hastings et al. (2016), but possibly explained by students' awareness of the gender gap in earnings (Reuben, Wiswall, & Zafar, 2017). A number of studies further document that students from low socioeconomic status backgrounds have less accurate expectations (Baker et al., 2018; Hastings et al., 2015, 2016). This could be explained by their parents having less information (Bleemer & Zafar, 2018; Lernetporer, Werner, & Woessmann, 2018), thus making the process of acquiring this information more costly. We indeed confirm that students from lower socioeconomic status neighborhoods make larger absolute errors and are less likely to be correct about the hourly wages of the occupations they select, but this does not hold for the job opportunities. Lastly, students have been shown to be uninformed about programs with good labor market prospects outside of their field of interest (Hastings et al., 2015), and underestimate earnings of their least preferred programs (Zafar, 2011).

A number of field-experimental studies have tested the effects of interventions aimed at improving students' knowledge about the returns to – and costs of – education. Evidence from the Dominican Republic shows that providing students with information about the returns

to attending secondary school increases enrolment (Jensen, 2010). For the general secondary education student population in industrialized countries, providing information about the returns to further education does not seem to influence actual enrollment (Pekkala Kerr et al., 2015; Bonilla-Mejía et al., 2019). There is some evidence that it does increase intended enrollment, particularly for students from low socioeconomic status backgrounds (Oreopoulos & Dunn, 2013; McGuigan, McNally, & Wyness, 2016; Peter & Zambre, 2017). Most closely related to our study are a number of studies that focus on providing information about the returns to specific study programs or institutions. These generally find stronger effects. Some studies show that, after being provided with such information, students are more likely to enroll in more prestigious institutions (Bonilla-Mejía et al., 2019) and higher-return study programs (Hastings et al., 2015). It has also been documented that simply receiving information about a study program makes students more likely to enroll in them (Conlon, 2019).

Despite the differences in context and outcome measures used, it is worthwhile to consider how the effectiveness of our intervention compares to these closely related information interventions. Bonilla-Mejía et al. (2019) find no effect of information provision on student enrollment in higher education. They do find that students who receive information are 0.5 percentage points more likely to enroll in highly selective colleges. Despite high returns to attending a selective college (Hastings, Neilson, & Zimmerman, 2013), the overall impact of Bonilla-Mejía et al.'s (2019) intervention on expected earnings is likely low, since the impact only applies to a small group of students. Conlon (2019) finds no impact of the implemented intervention on expected earnings of chosen majors. Our results are most easily compared to

Hastings et al. (2015). Their main result shows that their intervention increases the expected earnings of chosen degrees by 1.4% of the control group mean. Our survey results show that students who received wage information choose study programs with wages that are 2.5% higher than the control group mean. In addition, we find that, long before their study program choice, students in treated schools select profiles with weighted job opportunities and hourly wages that are 1.5% higher and 0.3% higher than the control group mean, respectively. This does not take into account any changes in occupational preferences within profiles. In short, our intervention – that is relatively cheap to implement – looks to be effective compared to similar studies and shows that information about study programs or occupations can have an impact long before students have to make a decision on their degree program or major.

Our study further draws on, and contributes to, recent work on role models. Porter and Serra (2020) show that female students are more likely to enroll in economics classes when they get to listen to a female role model talk about her experiences in university, as well as her career path and achievements (Porter & Serra, 2020). Moreover, Del Carpio and Guadalupe (2021) ran an experiment studying female enrollment in a 5-month software coding program. They show that removing a ‘success story’ of a female participant from the information page decreases enrollment by four percentage points. Our inclusion of the different ‘information senders’ provides a further look into how the characteristics of a person providing information affects the degree to which it is used.

Our main contribution is that, to the best of our knowledge, we are the

first to present students with information on the labor market prospects of occupations rather than specific study programs. Information about occupations may be more relevant as the true returns to education strongly depend on occupational sorting after graduation. Our setting provides a unique opportunity to do so, as vocational education programs are strongly tied to occupations. The occupations we provide information about are those that students are most interested in, which maximizes the relevance of the information. Furthermore, we do not just treat students who are close to post-secondary education, but also those who still have to decide on their specialization in secondary education. This allows us to analyze what the impact of our information treatment is at different stages of students' educational careers. Lastly, with the exception of Hastings et al. (2015), all field-experimental studies we know of required students to attend a presentation or visit a website they otherwise would not have. Our intervention is designed within an established career guidance platform actually used as part of students' curriculum in school. This provides a more natural environment. The intervention is low-cost and easy to replicate. Based on our field-experimental results, the company that we collaborate with intends to include our intervention on the platform in the near future.

The rest of this Chapter is structured as follows. Section 3.2 explains the institutional context: the Dutch education system and career guidance practice. Section 3.3 shows how we recruited schools and randomized them into treatment groups. Section 3.4 describes the experimental design. Section 3.5 lays out the data specifications and Section 3.6 presents the results. Section 3.7 concludes.

3.2 Institutional Context

In this experiment, we focus on students enrolled in pre-vocational secondary education in the Netherlands. Pre-vocational secondary education is one of the three main tracks of Dutch secondary education³. As the name suggests, it is vocationally-oriented and offers a broad range of subjects. It is also the largest track in terms of student numbers: in the 2017/2018 school year, about 53% of Dutch children in secondary school attended pre-vocational secondary education (Dutch Inspectorate of Education, 2020).

The pre-vocational secondary education program takes four years to complete (Nuffic, 2019). At the end of the second year, students choose a 'learning pathway' and profile. Pre-vocational secondary education is divided into four 'learning pathways': the basic vocational program, advanced vocational program, combined program, and theoretical program (Nuffic, 2019). In the theoretical program, students mostly take general subjects. The combined program drops one general subject in favor of four hours of vocational training, but is otherwise the same. In the basic and advanced vocational programs, students receive approximately 12 hours of vocational training instead of general subjects. General subjects are taught at a lower level compared to the combined and theoretical programs, with the level at the advanced vocational program being slightly above that of the basic vocational program. Within the learning pathways, students also choose a profile⁴. This profile de-

³Pre-vocational secondary education is known as 'vmbo' in Dutch. The two other tracks are higher general secondary education (havo) and pre-university education (vwo).

⁴For the basic vocational, advanced vocational, and mixed program there are ten available profiles: 1. Building, housing and interiors, 2. Engineering, fitting out and energy, 3. Transport and mobility, 4. Media, design and IT, 5. Maritime and technology,

termines what subjects are taught (Government of the Netherlands, n.d.-a). Both the learning pathway and profile a student chooses have important consequences for the opportunities for further education at the time the student graduates, on which we expand below. At the end of the fourth year, students have to decide how to continue their education. Dutch law dictates that students cannot leave education until they are either eighteen years of age or have a 'starting qualification' (i.e., an intermediate vocational education or higher general continued education degree).

As students usually graduate from their initial pre-vocational education program at age sixteen, entering the labor force directly is generally not an option. This leaves them with essentially two options: move on to post-secondary intermediate vocational education or enroll in a different (sub)track of secondary education. Graduates from all learning pathways are eligible to enroll in intermediate vocational education. The exact level at which graduates can enroll depends on the chosen learning pathway. Graduates from the basic vocational program can enroll in qualification level 2 of intermediate vocational education only. Graduates from the other three programs can enroll in levels 2, 3 and 4 (Government of the Netherlands, n.d.-b). Programs in intermediate vocational education generally train students for a specific occupation.

To aid students in navigating these choices, schools are required to provide career guidance counseling. To structure their career guidance

6. Care and welfare, 7. Business and commerce, 8. Catering, baking and leisure, 9. Animals, plants and land and 10. Services and products. For the theoretical program, there are four options: 1. Care and welfare, 2. Engineering and technology, 3. Business and 4. Agriculture

counseling efforts, most schools make use of online platforms. For this experiment, we partner with a company called Qompas, which provides an online platform to schools. The platform consists of a number of assignments that help students get to know more about themselves and the choices they will have to make. While students can access the platform at any time, the idea is that schools use Qompas during their career guidance counseling classes at set times during the week. All assignments the students complete are saved and stored in their personal file, which they are supposed to review periodically. We implement the experiment described in this Chapter within one of Qompas's assignments: the occupation assignment. While the Qompas system has a suggested order for doing the different assignments, schools ultimately decide in which year students do which. Schools usually have students do the occupation assignment in the second, third or fourth year of education. We expand on the assignment in Section 3.4.

3.3 Recruitment and Randomization

We recruited schools to participate in the experiment directly through the Qompas system. At the time of recruitment, 300 schools for pre-vocational secondary education were registered in the Qompas system, which comprises about 15% of all schools of this type in the Netherlands. Of these schools, thirteen were not eligible to participate in the experiment because of missing information.

The 287 remaining schools were informed about the experiment through a system message as well as an email. Qompas informed schools that they, together with a research institute of Maastricht University, were asked by the Ministry of Education, Culture, and Science to do

research into the effects of labor market information on the choices of pre-vocational secondary education students. Qompas further explained to schools that the research would be conducted by way of an experiment within the Qompas career guidance counseling platform. Schools also received contact details of the person responsible for the experiment at Qompas in case they had any questions, complaints or did not want to participate. Appendix 3.A provides the original version as well as an English translation of the message. Only a single school indicated that it did indeed not want to be a part of the experiment. This left us with 286 schools.

To randomize schools, we employed a stratified procedure at the school level. The reason for randomizing at the school level instead of at the student level is twofold. First, it reduces the chance of there being spillover effects between students who receive different treatments. Second, we expected that schools would be less willing to participate if some of their students were to be provided with information, whereas others were not.

Table 3.1: Treatment assignment, participation and analysis sample overview

Treatment Group	Frac. of Schools	Assigned Schools	Participating Schools	Participating Students	Schools in Analysis	Students in Analysis
Control Group	1/3	96	83	12,544	81	9,275
Job Opp. Info by Researcher (Treatment 1)	1/6	47	42	6,917	42	5,117
Job Opp. Info by Research Institute (Treatment 2)	1/6	47	40	6,470	40	5,151
Job Opp. & Wage Info by Researcher (Treatment 3)	1/6	48	38	5,580	38	4,254
Job Opp. & Wage Info by Research Institute (Treatment 4)	1/6	48	43	5,680	42	4,470
Total	1	286	246	37,191	243	28,267

We randomized schools into three main groups of approximately equal size: a control group, a treatment group that receives information about just job opportunities, and a treatment group that receives information about both job opportunities and hourly wages. The latter two groups were randomly assigned to receive information from either a research institute or a specific researcher from this institute. Columns (2) and (3) of Table 3.1 display the exact division of schools assigned over the different groups. We explain the difference between the treatment groups in further detail in Section 3.4.3.

We stratified schools on the basis of three characteristics: the number of broad profiles offered in the school, the number of students who completed the occupation test in the year before the experiment, and the quality of life indicator of neighborhoods the students come from. For the available profiles, we relied on data from Qompas. Qompas also registered the number of students who completed the occupation test in the previous year. However, data was not available for all schools. If no data was available, we predicted the number using the number of newly registered students in the Qompas system and the total number of students in the school itself.⁵ If data on one of the two was not available, we predicted the number using just the available measure. For the quality of life in neighborhoods students came from, we relied on the quality of life indicator developed by the Dutch Ministry of the Interior and Kingdom Relations.⁶ All neighborhoods (defined by their 4-digit postal code) in the Netherlands have a score, ranging

⁵Data on the number of students in the school itself is provided as open data by the Dutch education executive agency; <https://duo.nl/open-onderwijsdata/databestanden/vo/leerlingen/leerlingen-vo-2.jsp>; Retrieved: 22-06-2018.

⁶<https://data.overheid.nl/dataset/leefbaarometer-meting-2018>; Retrieved: 22-06-2018.

from 1 (very low quality of life) to 9 (very high quality of life). For every school, we calculated the weighted average quality of life indicator score of the neighborhoods the school's student body came from.⁷ If no data on the residential location of students was available, we predicted the average quality of life indicator score using the score of the school's neighborhood.

We used a block design to randomize. Because the profile choice is one of our outcome variables and largely determines the variety of occupations the students are likely to be interested in, we first sought balance on this dimension. We divided the schools into three groups: predetermined choice (only one theoretical profile available), limited choice (between one and three theoretical profiles available) and all four profiles available. Within these groups, we subsequently ranked schools based on the number of students who completed the occupation test last year. We split these groups into three more equal groups based on this dimension. As schools vary a lot in size, we hoped to improve balance in terms of sample size in this way. Lastly, within each of the now nine groups, we ranked schools on the basis of the weighted average quality of life indicator score. We then further split these groups into two. Increased balance on this dimension is important as we estimate heterogeneous effects based on the indicator. In the end, we were left with eighteen strata.

Within each stratum, schools were randomly assigned to the different treatment groups according to the division specified in Table 3.1. As not every stratum contained a perfect multitude of six schools, not all schools could be assigned in one go. We dealt with the unassigned

⁷This information is available in the data set referred to in footnote 5.

schools by recreating strata as mentioned above, omitting the division in two based on the weighted average quality of life indicator score. Within each of the now nine strata, schools were again randomly assigned. For unassigned schools arising from this procedure, we repeated the procedure once more, now stratifying only based on the freedom of profile choice. The last ten remaining unassigned schools were sorted based on the freedom of profile choice and then assigned based on a randomly ordered list of the control and treatment groups. Figure 3.B.1 in the Appendix provides a visual representation of the procedures.

3.4 Experimental Design

In this Section, we describe the experimental design in detail. The accompanying [Appendix D \(online\)](#) shows screen captures of the screens students in each of the control and treatment groups get to see in the experiment.

3.4.1 Occupation test

The assignment on occupation choice in the Qompas method consists of two parts: a test and a reflective assignment. Although we make no alterations to the test, we use its results in the experiment. During the test, Qompas asks students to answer 90 questions about themselves and their attitude towards a number of salient occupations (e.g., waiter/waitress, mason, mechanic). The aim of this test is to predict what sort of occupations the student might be interested in. Based on the answers, Qompas calculates a score for each of the 353 occupations in

their system. This score represents how well the various occupations fit the student's preferences and abilities. Qompas subsequently uses the results of this test in the reflective assignment, which contains our intervention.

3.4.2 Elicitation of baseline information

Before the start of the experiment, we establish a baseline of students' preferences and beliefs. To do so, we ask students a number of questions before being exposed to the intervention. The first question we ask is about their intended profile choice, which the second year students still have to make at this point. They can pick multiple options, in case they are not sure yet. We subsequently show students the twenty occupations that fit them best according to the test and ask them to select the five occupations they are most interested in. Students then receive information on the day-to-day activities in these occupations. After they read the information, we ask the students to rank the occupations in order of how much they would like to work in them later in life. Lastly, we ask students to state their beliefs about the job opportunities and gross hourly wages of the five occupations they selected using a slider.⁸ The options for job opportunities are "very poor", "poor", "reasonable", "good", and "very good". The options for the hourly wage range between €10.- and €26.-, with €1.- intervals.

During the first year of the experiment (the 2018/2019 school year),

⁸We ask for gross hourly wage because many youngsters in the Netherlands have a side job, e.g., in a supermarket, and are likely to have a good understanding of what they earn per hour with this job.

the sliders had a default option: “reasonable” for the job opportunities and €18.- for the hourly wages. Qompas removed this default option for the 2019/2020 school year. Moreover, in the 2018/2019 school year, students were able to alter their prior beliefs later on in the experiment by returning to them after receiving the information. Qompas corrected this error for the 2019/2020 school year. Because of these issues, we only consider the students who went through the experiment in the 2019/2020 school year whenever prior beliefs are relevant.

3.4.3 Information provision

After we elicit the baseline preference ranking and beliefs about the labor market prospects, we present treated students with information about the labor market prospects of the occupations they selected. Control group students do not get any labor market information. For treatment groups 1 and 2, we provide information about the forecasted job opportunities. In treatment groups 3 and 4 we add information about the occupations’ median hourly wage levels. Maastricht University’s Research Center for Education and the Labor Market (ROA)⁹ provided us with the information. As part of one of its research programs, ROA develops labor market forecasts for job opportunities of 113 different occupational groups in the next six years.¹⁰ This is what we use to inform students about the job opportunities. ROA also calculated the median hourly wage of intermediate vocational education graduates for these 113 occupations. To this end, they used data from the Dutch

⁹www.roa.nl

¹⁰For information on methods, validity, and the governance of this project, see <https://roa.maastrichtuniversity.nl/research/research-projects/project-onderwijs-arbeidsmarkt-poa>. These forecasts are used by the national unemployment agency and for the accreditation of new study programs.

Labor Force survey, matched to administrative records. We match the Qompas occupations to these occupational groups.

In treatments 1 and 3, we tell students that the information is presented by a researcher affiliated to Research Center for Education and the Labor Market. We divide senders into four groups: inexperienced male researchers, experienced male researchers, inexperienced female researchers, and experienced female researchers. In this context, experience is defined by the seniority of the information sender. We consider a researcher who did not have a Ph.D. (yet) at the time of the experiment's launch to be inexperienced, and consider a researcher with a Ph.D. to be experienced. To ensure understanding, we present senders' experience as either 'beginning researcher' or 'experienced researcher'¹¹. For each sender, we show the name and experience on the screen¹². We do not explicitly mention gender, but the names of all senders are indicative of their gender and the Dutch word for 'researcher' is different for men and women. We do not show pictures of the senders, so as to avoid bias caused by appearance unrelated to status or gender.

In treatment groups 2 and 4, we do not specify a human information sender. Instead, we tell students that the Research Center for Education and the Labor Market will provide them with the information. As we do not provide students in the control group with any information, we do not show them a sender either.

¹¹In Dutch: 'beginnend onderzoek(st)er' and 'ervaren onderzoek(st)er'. We do not present the different statuses as 'junior' and 'senior', respectively, because we are worried about a lack of understanding. 'Beginning' and 'experienced' are more commonly used in the scenario described above in Dutch than in English.

¹²With their consent, we use the actual names of Research Center of Education and the Labor Market employees.

3.4.4 Video

Next, we show students in all treatment and control groups a short video about work in general¹³. The video does not mention any particular occupations or the importance of job opportunities and wages. The main reason to show the video is to create some time between the first and second elicitation of beliefs for the control group. Without the video, students in the control group would be asked to state their beliefs a second time right after the first.

3.4.5 Elicitation of posterior beliefs and ranking

To estimate the effect of the treatment on beliefs and preferences, we elicit the students' ranking and beliefs a second time after the video. We show students their initial ranking and beliefs and ask them if they want to change anything.

3.4.6 Alternative occupations

37.7% of students select only occupations of which the job opportunities are forecasted to be "very bad", "bad" or "reasonable". We suggested to those students a few alternative occupations with better labor market prospects. To treated students, we state that the labor market prospects for their chosen occupations are not very good, and that the proposed alternatives have better prospects. We do not tell control group students why we offer them alternatives. All students receive information on the day-to-day activities of these occupations. If students got to see the alternative occupations, they get the opportunity

¹³<https://www.youtube.com/watch?v=YJ78VDQrO3c>

to include these occupations in their ranking. Initially, we place these alternative occupations at the bottom of the ranking in a randomized order.

Information about the labor market prospects of the alternative occupations was supposed to only be provided to students in the treatment groups. However, due to a programming error, control group students also received information about the job opportunities of the alternative occupations as well as their initial set of occupations. Because of this error, we do not consider the alternative occupations in our analysis at all and remove students who were suggested alternatives from our post-intervention analyses.

3.4.7 Elicitation of posterior intended profile choice

At the end of the experiment, we once again ask students what profile they intend to choose. We show them their initial selection and allow them to alter it.

3.5 Data

3.5.1 Sample

We collected data between September of 2018 and July of 2020, covering the 2018/2019 and 2019/2020 school years. 249 schools actually participated in the experiment, for a total of 40,176 individuals. At the other 37 schools, the part of the platform that included our experiment was not used by any student. As schools could not know their treatment assignment before going through the experiment, this forms no

threat to our internal validity. A small fraction of the individuals who went through the experiment were either first-year students (1,855) or school administrators involved in study guidance (48). We exclude them from the data. Of the remaining group of students, 1,082 did not make a first ranking of their selected occupations. As these students bring no data worth analyzing, we also exclude them. 8,924 students changed their initial preference ranking on a different day than on the day they went through the experiment. This could be because these students went through the experiment multiple times, making the belief and ranking measures unreliable. We therefore remove these students from the sample as well. None of these sample restrictions are related to treatment status. After imposing our restrictions, we are left with 28,267 individuals from 243 schools. Columns (4) to (7) of Table 3.1 show how these numbers relate to the number of assigned schools. Table 3.2 shows that covariates are balanced between the control and treatment groups.

3.5.2 Survey data

In addition to the experimental data, we conducted a survey among graduating students in the 2019/2020 school year. The survey was fielded between the 15th of April and the 20th of May, 2020. The survey was sent to 9,510 students of which 1,061 responded. Again, we impose a number of sample restrictions. In our analysis, we only consider students who went through the experiment, did not change their prior ranking on a different day than they created it, did not see the alternative occupations and were either in the second-to-last year of secondary school in the 2018/2019 school year or the final year of sec-

Table 3.2: Balance of covariates across treatment groups

	Control		Treatment 1		Treatment 2		Treatment 3		Treatment 4		P-value joint sign.
	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	
No. of students	164.14	134.24	175.55	147.33	178.00	121.59	158.21	118.49	150.10	110.95	0.82
No. of profiles available	3.37	0.99	2.93	1.30	3.30	1.14	3.21	1.07	3.14	1.20	0.38
Age	14.04	1.01	14.02	0.99	14.16	0.99	14.06	0.99	14.05	1.04	0.82
Male	0.52	0.50	0.53	0.50	0.56	0.50	0.52	0.50	0.50	0.50	0.43
Grade	2.47	0.64	2.45	0.64	2.56	0.68	2.43	0.61	2.48	0.67	0.84
QOL score	6.57	1.35	6.55	1.31	6.77	1.23	6.65	1.42	6.46	1.46	0.59

Note: no. of students and no. of profiles available are school level variables. Age, male, grade and QOL score are individual level variables. QOL score refers to the quality of life indicator score of the neighborhood the student lives in (see Section 3.3 in this Chapter). The last column of the Table indicates whether a joint significance tests shows a significant difference between the treatment groups and the control group for each of the variables considered.

ondary school in 2019/2020 school year. After we impose our sample restrictions, we are left with 4,389 survey invitees, and 405 respondents. To incentivize responses, we announced that we would raffle off 20 €25.- vouchers for a large Dutch e-tailer among survey respondents. In the survey, we once again ask students to state their beliefs about the labor market prospects of the occupations they selected as well as to rank the occupations based on how much they would like to carry them out later in life. Furthermore, we ask them about their plans for next year. If students indicate they will go on to intermediate vocational education, we ask them to state what study program they enrolled in using a free-form text field. We manually match these study programs to their official study program identifier. All study program identifiers can be linked to the Research Center of Education and the Labor Market's labor market information system, which we use in our analysis. For the analysis of the survey respondents' beliefs and preferences, all above mentioned sample restrictions apply as well. For the analysis of the study program choice, we relax the restriction on the al-

teration of the prior ranking. Naturally, we do not include prior beliefs in these analyses. Table 3.C.1 in the Appendix shows that responding to the survey is not related to treatment. We do observe that older students and male students are less likely to respond to the survey.

3.5.3 Administrative data

To analyze the actual profile choice of students, we use administrative data at the school level¹⁴. These data provide us with information on the number of students that follow a particular profile on the 1st of October of each year from 2018 to 2020.

While these data allow us to analyze the impact of the treatment on students' profile choices, there are a number of caveats. First of all, because the data is at the school level, we are not able to restrict the sample only to those who have not seen the suggested alternative occupations. Our estimations will therefore likely be lower bounds. Secondly, we lose a lot of power by not being able to do the analyses at the individual level. We recently received permission to match our experimental data to administrative records at the individual level. Analyses using individual-level data will be performed and added once the data have become available for us.

¹⁴Available at https://duo.nl/open_onderwijsdata/databestanden/vo/leerlingen/leerlingen-vo-1.jsp.

3.6 Results

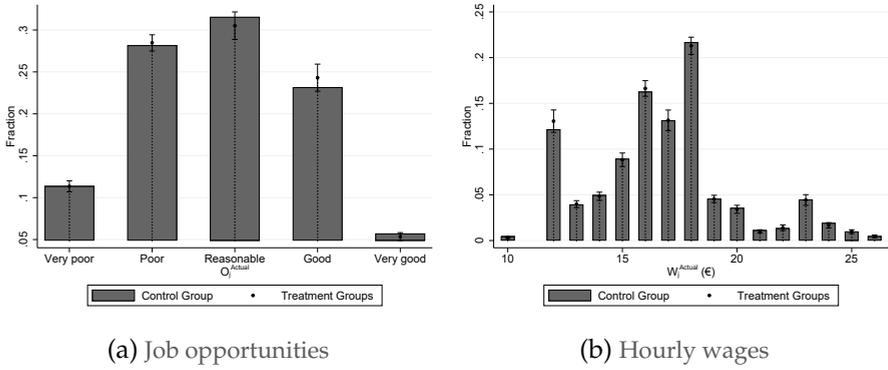
3.6.1 Descriptive statistics

3.6.1.1 Selected occupations

Figure 3.1 shows the job opportunities and hourly wages of the occupations students selected for their top five before the intervention. Most selected occupations have job opportunities that are either poor (category 2), reasonable (3) or good (4). Hourly wages generally range between €12.- and €18.-. The Figure also shows that before the interventions there is no difference between the control and treatment groups in terms of job opportunities and hourly wages for the occupations the students selected for their top five. Tables 3.C.2 and 3.C.3 in Appendix C confirm this. Although Table 3.C.3 does show that students in the first treatment group select occupations with lower hourly wages, the joint significance tests do not allow us to reject that the selection process between the treatment and control groups is the same. These Tables also show that there is no difference in the labor market prospects of the selected occupations between the first and second year of the experiment.

There are some interesting patterns in the selection of the occupations. Tables 3.C.4 and 3.C.5 in Appendix C show that male students generally select occupations with better job opportunities and higher hourly wages than female students do. Students in later years and students from low socioeconomic status neighborhoods choose occupations with higher hourly wages, but no better job opportunities.

Figure 3.1: Job opportunities and hourly wages of selected occupations



Note: graphic representation of multinomial logit estimation. Standard errors clustered at school level. Bars indicate level for control group. Dots and error bars indicate level for treatment group and 95% confidence interval, respectively.

3.6.1.2 Prior beliefs

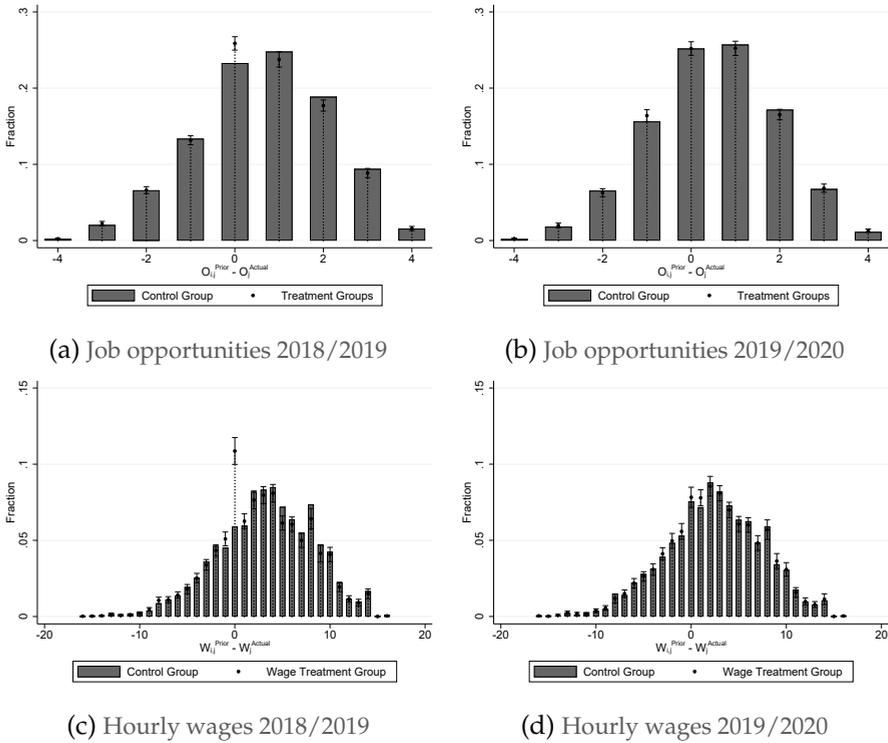
Figure 3.2 shows the prior belief accuracy of the control and treatment groups in the two years of the experiment. We denote the prior beliefs of individual i about the job opportunities of occupation j by $O_{i,j}^{Prior}$, and the actual job opportunities for that occupation by O_j^{Actual} . We apply the same logic to the hourly wages, which we denote as W . To measure belief accuracy, we first consider the difference between individual i 's belief about the prospects of occupation j and its actual prospects: $O_{i,j}^{Prior} - O_j^{Actual}$ and $W_{i,j}^{Prior} - W_j^{Actual}$. These differences, which we report in Figure 3.2, allow us to analyze the degree of over- and underestimation of job opportunities and hourly wages. In the 2018/2019 school year, treated students show significantly more accurate expectations about the job opportunities and hourly wages than

do control group students. This is likely due to the fact they could correct their initial beliefs, as discussed in Section 3.4.3. In the 2019/2020 school year, when the programming error was fixed, there is no difference between the beliefs of control and treatment group students. The figure also shows a left-skewed distribution, which indicates that students tend to overestimate the labor market prospects of their preferred occupations.

When using central tendency measures, errors in beliefs that have opposite directions may cancel each other out. We therefore consider two additional metrics to assess the accuracy of students' beliefs and how these differ by a number of characteristics. First, we analyze the absolute values of the belief errors: $|O_{i,j}^{Prior} - O_j^{Actual}|$ and $|W_{i,j}^{Prior} - W_j^{Actual}|$. The combination of the overestimation and absolute error allows us to infer to what degree errors are caused by overestimation and underestimation. Secondly, we analyze how often beliefs are exactly correct (i.e., $O_{i,j}^{Prior} - O_j^{Actual} = 0$ and $W_{i,j}^{Prior} - W_j^{Actual} = 0$).

Because we bound students' stated expectations by our use of sliders, students cannot overestimate occupations with good job opportunities and high hourly wages to the same degree as occupations that have worse prospects. Since there is heterogeneity in occupational preferences, we have to account for this in our analyses. We do this by adding an occupation fixed effect in our analysis of belief accuracy. This means we compare individuals' belief accuracy conditional on the occupation they selected. Table 3.C.6 in Appendix C shows that male students tend to overestimate both job opportunities and hourly wages to a larger degree. They also make larger absolute errors and are less likely to be correct. Third and fourth year students, who are

Figure 3.2: Prior belief accuracy by relevant group



Note: graphic representation of multinomial logit estimation at occupation level. The x-axis displays the degree of overestimation. For the job opportunities, the numbers indicate the overestimation in categories (i.e., -2 denotes an underestimation of two categories, whereas +2 indicates an overestimation of two categories). For the hourly wages, the overestimation is displayed in Euros. Standard errors clustered at school level. Bars indicate level for control group. Dots and error bars indicate level for treatment group and 95% confidence interval, respectively.

closer to making a decision than second year students do not do much better when it comes to the job opportunities, but make smaller abso-

lute errors for the hourly wages. This might be because of the fact that we present the job opportunities in a categorical manner. Even if students do have a good idea about future job opportunities, they might not agree on the qualifications we assign to them. Students in schools where more profiles are available seem to make smaller absolute errors and are somewhat more likely to be correct about both job opportunities and hourly wages. What's most striking about the Table is the effect the initial ranking of the occupation has on the belief accuracy. Higher ranked occupations are overestimated to a much larger degree. The difference between the number one and number five ranked occupation is almost an entire category for the job opportunities and €1.50 for the hourly wages.

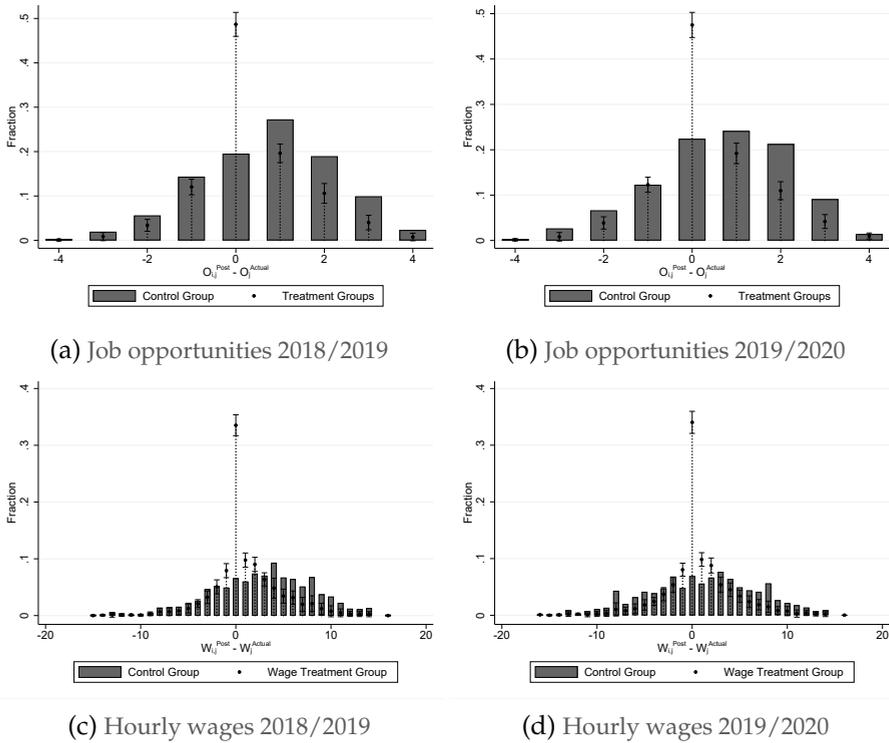
3.6.2 Treatment effects

3.6.2.1 Posterior beliefs

Moving to the effect of the treatment, Figure 3.3 shows the posterior belief accuracy for the control group and relevant treatment groups. We denote the posterior beliefs of individual i about the job opportunities of occupation j by $O_{i,j}^{Post}$ and that of the hourly wages by $W_{i,j}^{Post}$. The graphs show that in both years, students in the treatment groups are much more likely to be correct about the job opportunities and the hourly wages of their selected occupations. This is largely driven by the correction of overestimations. Treated students correct beliefs more often and more strongly than control students. Students who initially underestimated the labor market prospects of their occupations react much less strongly than those who initially overestimated them. Tables

3.C.7 and 3.C.8 in Appendix C confirms this for the 2019/2020 cohort, where we can use students' prior beliefs in the analysis.

Figure 3.3: Posterior belief accuracy by relevant group



Note: graphic representation of multinomial logit estimation at occupation level. The x-axis displays the degree of overestimation. For the job opportunities, the numbers indicate the overestimation in categories (i.e., -2 denotes an underestimation of two categories, whereas +2 indicates an overestimation of two categories). For the hourly wages, the overestimation is displayed in Euros. Standard errors clustered at school level. Bars indicate level for control group. Dots and error bars indicate level for treatment group and 95% confidence interval, respectively.

Table 3.C.9 shows that the treatment is equally effective when a researcher is said to provide the information, compared to an institute. Zooming in on the specific researcher, Table 3.C.10 shows that neither whether a male or a female researcher provides the information, nor whether the sender was an ‘experienced’ or ‘beginning’ researcher matters for the degree to which beliefs are updated. Table 3.C.11 shows that when it comes to job opportunities, third and fourth year students react more strongly to the treatment than second year students. The same holds for male versus female students. Table 3.C.12 shows that treated fourth year students are more often correct than earlier year students, although the change is smaller than for the job opportunities and only marginally significant.

Next, we study how persistent the effects on posterior beliefs are. Table 3.C.13 shows that beliefs about the job opportunities remain more accurate at the time of the survey for students treated in the 2019/2020 school year (that is, up to seven months after treatment). This does not hold for those treated in the 2018/2019 school year (who completed the survey over a year after the treatment). However, we cannot ascribe the difference to time since treatment alone. As we survey graduating students, the students who received the information most recently were also much closer to the end of their secondary school career when they did. Information on job opportunities and hourly wages may become more important as students get closer to their post-secondary education decision. As such, the reason these students better recall the information may be that they paid more attention to it, not that that they received it more recently. With our data, we cannot distinguish between these two mechanisms. For the hourly wages, we find that treated students do not have more accurate beliefs than the control group for both years of the experiment.

3.6.2.2 Rankings

Table 3.3 shows how the treatment affects the likelihood of students changing their favorite occupation between the first and second elicitation. We observe that students in the treatment group indeed change their favorite occupation significantly more often than those in the control group. The effect size is fairly small, however. In the control group, approximately 5.53% of students change their favorite occupation. In the treatment groups, this fraction is 0.88 to 2.16 percentage points higher.

The fact that students in the treatment group change their favorite occupation (slightly) more often does not tell the whole story, however. Table 3.3 also shows whether students in the treatment group switch towards occupations with better labor market prospects. ΔO_j^{Actual} and ΔW_j^{Actual} , respectively, denote the difference in the job opportunities and hourly wages between the number one ranked occupation at first elicitation and the number one ranked occupation at second elicitation. If a student does not change his or her favorite occupation between the first and second elicitation, $\Delta O_j^{Actual} = \Delta W_j^{Actual} = 0$. Columns (2) and (4) show the effect unconditional on actually changing the number one ranked occupation. The job opportunities in the treatment groups rise by anywhere from 0.0190 to 0.0305 categories. For the wage treatments, the hourly wages rise by €0.09. Columns (3) and (5) show the change for students who did change their favorite occupation. For students in the treatment groups, the job opportunities move up by 0.285 to 0.447 categories and hourly wages by €1.12 to €1.20. It is important to note that in both cases, the job opportunities and hourly wages do not move at all for control group students.

Table 3.3: Treatment effect on likelihood changing favorite occupation and change in prospects

	(1)	(2)	(3)	(4)	(5)
	Pr(Fav. Change)	ΔO_j^{Actual}	ΔO_j^{Actual} (Changed)	ΔW_j^{Actual}	ΔW_j^{Actual} (Changed)
Sender: researcher					
Info: job opportunities	0.00877* (0.00454)	0.0190*** (0.00622)	0.295*** (0.101)	0.0130 (0.0174)	0.192 (0.292)
Sender: institute					
Info: job opportunities	0.0126** (0.00529)	0.0305*** (0.00650)	0.447*** (0.102)	0.0303** (0.0149)	0.430* (0.253)
Sender: researcher					
Info: job opp. & wages	0.0216*** (0.00562)	0.0278*** (0.00577)	0.358*** (0.0880)	0.0876*** (0.0215)	1.115*** (0.295)
Sender: institute					
Info: job opp. & wages	0.0189*** (0.00624)	0.0214*** (0.00640)	0.285*** (0.0944)	0.0904*** (0.0199)	1.197*** (0.278)
Constant	0.0553*** (0.00317)	0.000666 (0.00325)	0.0120 (0.0590)	0.00466 (0.0102)	0.0843 (0.186)
Observations	27387	27387	1791	27387	1791

Note: *Constant* refers to the control group estimate. Each row above refers to the incremental estimate for each of the four treatment groups. Pr(Fav. Change) in Column (1) is the chance that a student changed his or her favorite occupation between second elicitation. ΔO_j^{Actual} in Column (2) denotes the difference between the job opportunities of the student's favorite occupation at the second elicitation and the first elicitation. It is equal to 0 if the student did not change. ΔW_j^{Actual} in Column (4) denotes the equivalent for the hourly wages. Columns (3) and (5) only contain observations where the student did switch favorite occupations between first and second elicitation. Standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual level.

Table 3.C.14 in the Appendix shows that treated students in schools with four profiles available are not more likely to change their favorite occupation compared to schools with fewer profiles available, but do switch to occupations with better job opportunities when they do. This may be driven by the fact that these students have a larger set of options to choose from. This Table, together with Table 3.C.15 shows that we find no further evidence for heterogeneous treatment effects. Table 3.C.16 shows there is no effect of the information sender either.

We do not find any evidence that treated students still prefer occupations with better prospects in the survey. However, Columns (1) and (3) of Table 3.C.17 show that the treated students in the survey did not switch to occupations with better prospects directly after the intervention either.

3.6.2.3 Profile & study program choice

Moving to the profile choice, we first study how our intervention affects the intended profile choice of students right after the intervention. Table 3.4 provides no evidence that the treatment impacts second year students' intended profile choice or the number of profiles they consider right after the experiment. Table 3.C.18 in Appendix C shows that there is some heterogeneity based on the number of profiles available in the school, however. Looking at Column (3), the treatment seems to marginally narrow the scope of profiles students in the basic, advanced vocational and mixed programs are willing to consider if they are in schools that offer very few profiles. Table 3.C.19 shows that even learning that an occupation that fits with a certain profile has very good job opportunities or hourly wages does not make students

more likely to include that profile in their choice set immediately after the experiment. While somewhat surprising, it may be that students require some time to process the information they received and adjust their intended choices based on this. Based on administrative data at the school level, we can study how profile choices actually materialized. Unfortunately, this decreases the number of observations and therefore statistical power, especially since not every single profile is available in every school. Because of this, we consolidate our treatment groups to 'Treated' (treatment groups 1 through 4) and 'Wage Information' (treatment groups 3 and 4).

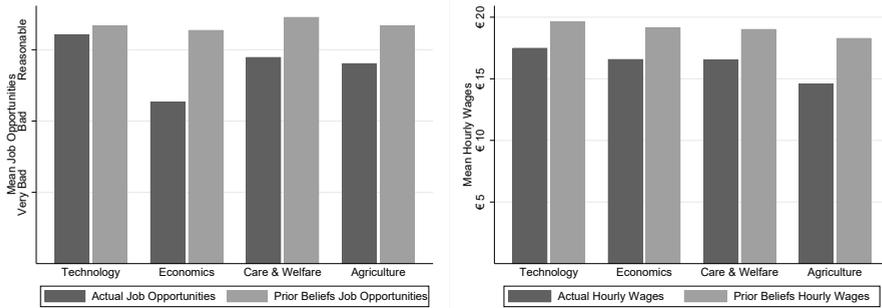
To analyze and quantify the impact our intervention had on the actual profile choice, we assign a value for job opportunities as well as hourly wages to each profile. We do this by taking the average job opportunities and hourly wages of the occupations associated with a certain profile, weighted by how often that occupation was chosen by students. Figure 3.4 shows these metrics by profile, as well as the weighted average prior beliefs for the occupations associated with these profiles. It is worthwhile to note how large the difference between the profiles in terms of the actual job opportunities and hourly wages they provide are. For instance, the dark gray bars in Panel (a) show that the average job opportunities for the technology profile is almost a category higher than that of the economics profile. Likewise, panel (b) shows that the hourly wages for occupations associated with technology profiles are close to €3.- an hour higher than for occupations associated with the agriculture profile. Panels (c) and (d) show a similar picture for the other profiles, with even more striking differences, as these profiles are more fine-grained. The light gray bars, showing students' prior beliefs about the labor market prospects in these occupations, show a much

Table 3.4: Treatment effect on profiles considered immediately after intervention

	(1)	(2)	(3)
	Pr(Same Profile Pre-Post)	No. of Theoretical Profiles	No. of Other Profiles
Sender: researcher			
Info: job opportunities	-0.0395 (0.0264)	-0.0152 (0.0188)	0.0468 (0.0475)
Sender: institute			
Info: job opportunities	-0.0274 (0.0229)	0.00781 (0.0229)	0.00588 (0.0525)
Sender: researcher			
Info: job opp. & wages	-0.0383 (0.0236)	0.00177 (0.0268)	0.00648 (0.0443)
Sender: institute			
Info: job opp. & wages	-0.0359 (0.0311)	-0.00870 (0.0197)	0.0229 (0.0481)
No. of theoretical profiles a priori		0.554*** (0.0190)	
No. of other profiles a priori			0.443*** (0.0333)
Constant	0.692*** (0.0126)	0.169*** (0.0122)	0.287*** (0.0412)
Observations	10671	5901	4772
F-Stat joint sign. of treatments.	1.133	0.312	0.368
P-value F-Stat joint sign. of treatments.	0.342	0.869	0.831

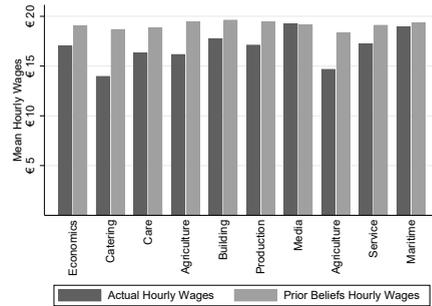
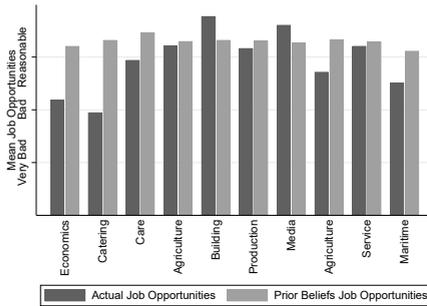
Note: *Constant* refers to the control group estimate. Each row above refers to the incremental estimate for each of the four treatment groups. No. of profiles a priori is a metric of how many profiles a student considered before the intervention. Pr(Same Profile Pre-Post) in Column (1) indicates the likelihood that student did not change his or her profile choice between the two elicitation. No. of Theoretical Profiles and No. of Other Profiles in Columns (2) and (3) denote the number of profiles a student considered at second elicitation, respectively. Only second year students, who did not see alternative occupations are included in this analysis. Standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual level.

Figure 3.4: Weighted average actual prospects and prior beliefs by profile



(a) Job opportunities theoretical profiles

(b) Hourly wages theoretical profiles



(c) Job opportunities other profiles

(d) Hourly wages other profiles

Note: dark gray bars display the average weighted job opportunities and hourly wages of selected occupations by profile, weighted by how often the occupation was selected. Light gray bars display the average prior belief about the job opportunities and hourly wages in these profiles.

flatter picture. While students do seem to judge the occupations associated with some profiles to have better prospects than others, the differences are slim. With such large differences, that students are not a priori aware of, it seems likely that some students may change their preferences over profiles once they learn more about the labor market prospects in the occupations associated with them.

To analyze this, we determine the weighted average job opportunities and hourly wages of the number of third year students enrolled in each profile at the school level. Our baseline period is the 1st of October 2018. As our experiment started in the 2018/2019 school year, the intervention could not have had an impact on the profile choice yet, as students decide on their profile at the end of the previous school year. Table 3.5 shows the results of our analysis. Controlling for the weighted average job opportunities and hourly wages of the chosen profiles in 2018, students in the relevant treatment groups choose profiles with better job opportunities and higher hourly wages. The effects are marginally significant, but point in the same direction. The effect is 0.0424 categories (1,5% of the control group mean) for the job opportunities and approximately €0.05 (0.3% of the control group mean) for the hourly wages. We should note that, because of the way in which we calculate the job opportunities and hourly wages of the profiles, we assume that the selection of occupations is unaffected by the treatment. This is not necessarily the case, as the results from our survey will show.

Table 3.5: Treatment effect on labor market prospects chosen profiles

	(1)	(2)
	Weighted job opportunities	Weighted wage level
Treated	0.0424* (0.0238)	0.00417 (0.0286)
Weighted job opportunities 2018	0.822*** (0.0377)	
Wage information		0.0421 (0.0272)
Weighted wage level 2018		0.961*** (0.0161)
Constant	0.473*** (0.0960)	0.620** (0.257)
Observations	444	444
F-Stat Treated + wage information		3.100
P-value F-Stat Treated + wage information		0.080

Note: *Constant* refers to the control group estimate. *Treated* refers to all four treatment groups. *Wage information* is the incremental estimate for treatment groups 3 and 4, that received wage information. ‘*Weighted job opportunities*’ in Column (1) and the second row refers to the average job opportunities of the profiles students chose, based on how often they chose them. ‘*Weighted wage level*’ refers to the same for the hourly wages. Section 3.6.2.3 explains the metrics in more detail. ‘*F-stat Treated + wage information*’ and the corresponding p-value refer to the test of significance for the estimate of ‘*Treated*’ + the estimate of ‘*Wage information*’, which provides the estimate for the treatment groups that received information about hourly wages as well as about job opportunities. Standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at school-year level.

Lastly, Table 3.6 shows how the treatments affect the study program choice. Column (1) shows that students who receive information about just the job opportunities from a researcher choose study programs with job opportunities that are 0.43 categories better on average. This is a sizeable effect. However, the effect of the other treatments are very close to – and not significantly different from – zero. A joint significance test reveals that the treatments did not lead to students choosing study programs with significantly better job opportunities. Column (2) shows the treatment effect on the hourly wages of the chosen study programs. Compared to control group students, students who receive information about both the job opportunities and the hourly wages from a researcher choose study programs of which graduates earn €0.74 an hour more on average. The other wage treatment shows a slight positive, but insignificant effect. A joint significance test of the two treatments shows that the information about the hourly wages indeed lead students to choose study program with higher earnings prospects.

Table 3.6: Treatment effect on job opportunities and hourly wages chosen study program

	(1)	(2)
	Job Opportunities Chosen Program	Average Hourly Wages Chosen Program
Sender: researcher		
Info: job opportunities	0.432** (0.184)	0.217 (0.291)
Sender: institute		
Info: job opportunities	-0.0544 (0.160)	0.0660 (0.259)
Sender: researcher		
Info: job opp. & wages	0.00253 (0.157)	0.740** (0.296)
Sender: institute		
Info: job opp. & wages	-0.0958 (0.253)	0.0367 (0.276)
Constant	2.361*** (0.0801)	17.43*** (0.154)
Observations	405	405
F-Stat Relevant Treatments	1.678	3.331
P-value F-Stat Relevant Treatments	0.160	0.0393

Note: *Constant* refers to the control group estimate. Each row above refers to the incremental estimate for each of the four treatment groups. Job opportunities chosen program in Column (1) refers to the assigned job opportunities of the student's chosen study program in the Research Centre of Education and Labor Market's research program. Average hourly wages of chosen program refers to the hourly wage level of the program, from the same source. Standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual level.

3.7 Conclusion

In this Chapter, we presented a field experiment aimed at improving the accuracy of Dutch pre-vocational education students' beliefs about the job opportunities and hourly wages of occupations they are interested in. In line with the literature, we find that students' prior beliefs are highly inaccurate. In our sample, both job opportunities and hourly wages are strongly overestimated, particularly for students' favorite occupations. This could be innocuous, and simply the result of students rationalizing their choices. However, since our results indicate that students do indeed attach some value to the labor market prospects of occupations when making educational decisions, another explanation is likely. If students gather noisy information and tend to gravitate towards the occupations for which they learn the labor market prospects are best, these will often be the occupations for which the information was least accurate in a winner's curse fashion. This underlines the importance of providing students with accurate information.

Our results show that providing such information is effective in correcting belief errors in the short term. However, survey data gathered after the experiment shows that these beliefs stick for at most a couple of months and only for the job opportunities. Students who receive information are more likely to change their favorite occupation between the first and second elicitation of the ranking and, if they do so, switch towards occupations with better labor market prospects. We are unable to confirm whether this change in preferences holds in the long term, however. Even though we do not see very strong effects on stated beliefs and preferences in the long term, we do see that students

in treated schools enroll in profiles associated with occupations that have better labor market prospects. Similarly, students who received information about the hourly wages of occupations do enroll in post-secondary education programs that have better earnings prospects. We find no evidence that it matters whether the information is provided by a person or an institute, and in the former case whether this person is experienced, inexperienced, male or female.

A limitation of our experiment is that we either use self-reported measures of beliefs or have to rely on school level data. We intend to repeat our analysis using administrative data at the individual level on profile choice, post-secondary education program choice and degree matriculation. The big advantage of this data for the profile choice is that we can actually link the information students received about occupations associated with a certain profile to their choices. Currently, we have to rely on averages by profile. Degree matriculation is particularly important as well, as it is not something we can currently analyze. If students choose different study programs because of the information, but do not end up finishing these, the net effect of the information may still be negative. Likewise, if the information motivates students to finish their study programs, our findings are an underestimation of the true effect.

Appendix 3.A Recruitment text

Dutch

ROA (Researchcentrum voor Onderwijs en Arbeidsmarkt aangesloten bij Universiteit Maastricht) en Qompas zijn samen door het Ministerie van Onderwijs, Cultuur en Wetenschap (OCW) gevraagd om onderzoek uit te voeren naar de invloed van arbeidsmarktinformatie op de keuze van vmbo-leerlingen voor een studie.

Door middel van een A/B-test in de lesmethode Qompas VMBO/-Mavo gaan we onderzoeken of vmbo'ers bij het maken van hun studiekeuze letten op informatie over baankans en of die informatie ertoe bijdraagt dat zij een betere keuze maken. Met deze informatie kan Qompas haar lesmethode doorontwikkelen om scholieren in de toekomst nog beter te kunnen helpen met hun studiekeuze.

Wij hopen dat uw school meewerkt aan dit onderzoek. Alle gegevens worden anoniem verwerkt. Voor meer informatie kunt u contact opnemen met [REDACTED].

English

ROA (The Research Center for Education and the Labor Market, part of Maastricht University) and Qompas were asked by the Ministry of Education, Culture and Science (OCW) to do research on the influence of labor market information on the education choices of intermediate vocational education students.

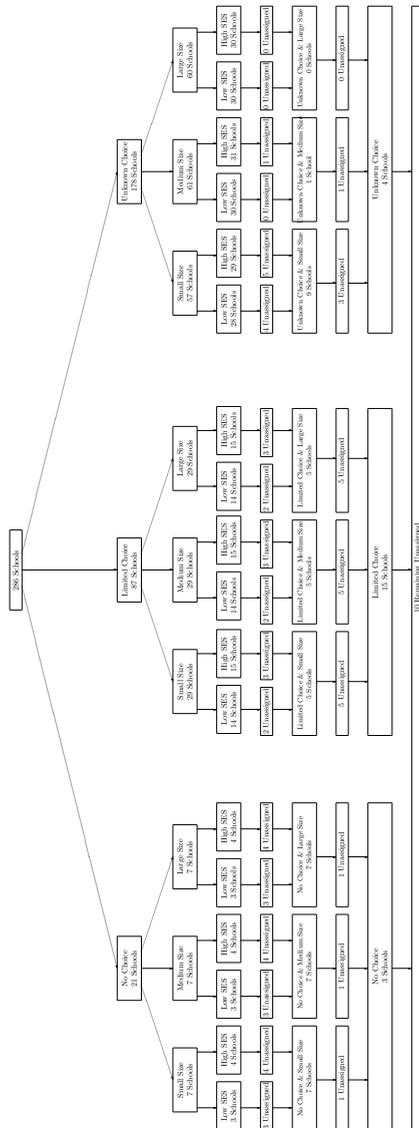
Through an A/B-test in the Qompas system we will research whether intermediate vocational education students take information about job

opportunities into account when making education choices and whether this information helps them make a better choice. With this information, Qompas can improve its platform by being even more able to help students with their education choice.

We hope your school will participate in this study. All details will be processed anonymously. For more information, you can contact [REDACTED].

Appendix 3.B Additional Figures

Figure 3.B.1: Graphical representation of randomization



Appendix 3.C Additional Tables

Table 3.C.1: Balance check survey respondents

	(1)	(2)	(3)	(4)	(5)
	Answered Survey	Age	Grade	Male	QOL score
Sender: researcher					
Info: job opportunities	0.0148 (0.0178)				
Sender: institute					
Info: job opportunities	0.00580 (0.0184)				
Sender: researcher					
Info: job opp. & wages	-0.00294 (0.0157)				
Sender: institute					
Info: job opp. & wages	-0.0116 (0.0125)				
Answered Survey		-0.0840* (0.0500)	-0.0161 (0.0343)	-0.192*** (0.0273)	-0.0251 (0.0680)
Constant	0.0960*** (0.00848)	14.88*** (0.0482)	3.242*** (0.0400)	0.563*** (0.0147)	6.710*** (0.0759)
Observations	4389	4012	4389	4388	4292
F-Stat Treatments	0.637				
P-value F-Stat Treatments	0.637				

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual level.

Table 3.C.2: Job opportunities of selected occupations by treatment group and year

		(1)	(2)	(3)	(4)	(5)	(6)
		Mean value	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Sender: researcher							
Info: job opportunities		-0.00181 (0.0223)	-0.00589 (0.0317)	-0.0207 (0.0288)	-0.0105 (0.0286)	0.0253 (0.0294)	-0.000699 (0.0317)
Sender: institute							
Info: job opportunities		0.0174 (0.0272)	-0.0187 (0.0330)	0.0376 (0.0335)	0.0205 (0.0354)	0.0482 (0.0391)	0.00394 (0.0437)
Sender: researcher							
Info: job opp. & wages		0.00634 (0.0194)	-0.0190 (0.0273)	0.0302 (0.0344)	0.0441* (0.0239)	0.00259 (0.0284)	-0.0236 (0.0296)
Sender: institute							
Info: job opp. & wages		-0.0285 (0.0211)	-0.0288 (0.0365)	0.000159 (0.0288)	-0.0253 (0.0264)	-0.0411 (0.0305)	-0.0507* (0.0281)
2019/2020		0.0160 (0.0145)	0.0450* (0.0250)	0.0575*** (0.0212)	-0.00715 (0.0230)	0.0200 (0.0223)	-0.0298* (0.0173)
Sender: researcher							
Info: job opportunities	× 2019/2020	0.00402 (0.0237)	-0.0156 (0.0389)	-0.0202 (0.0355)	0.0586 (0.0378)	-0.0542 (0.0348)	0.0415 (0.0364)
Sender: institute							
Info: job opportunities	× 2019/2020	-0.0168 (0.0251)	-0.0180 (0.0365)	-0.0629* (0.0370)	-0.0153 (0.0343)	-0.0312 (0.0406)	0.0353 (0.0357)
Sender: researcher							
Info: job opp. & wages	× 2019/2020	0.0395 (0.0272)	0.0508 (0.0426)	-0.0384 (0.0447)	0.0233 (0.0384)	0.0590 (0.0429)	0.0904** (0.0412)
Sender: institute							
Info: job opp. & wages	× 2019/2020	0.00884 (0.0308)	-0.0310 (0.0491)	-0.0408 (0.0406)	0.00896 (0.0452)	0.0160 (0.0395)	0.0853** (0.0345)
Constant		2.826*** (0.0116)	2.875*** (0.0215)	2.798*** (0.0174)	2.813*** (0.0160)	2.814*** (0.0151)	2.843*** (0.0151)
Observations		28267	27805	27811	27801	27715	27598
F-Stat Non-interacted Treatments		0.799	0.230	0.848	1.927	1.289	0.955
P-value F-Stat Non-interacted Treatments		0.527	0.922	0.496	0.107	0.275	0.433
F-Stat Treatments + Treatments × 19/20		1.704	1.856	1.411	1.623	2.367	1.750
P-value F-Stat Treatments + Treatments × 19/20		0.150	0.119	0.231	0.169	0.053	0.140

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual level. Discrepancies in observations are caused by the fact that a few occupations could not be assigned to a level of job opportunities. We do calculate an average score for the other occupations a student selected in this case.

Table 3.C.3: Hourly wages of selected occupations by treatment group and year

	(1)	(2)	(3)	(4)	(5)	(6)
	Mean value	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Sender: researcher						
Info: job opportunities	-0.194** (0.0906)	-0.266** (0.124)	-0.152 (0.107)	-0.236** (0.106)	-0.197** (0.0923)	-0.0812 (0.111)
Sender: institute						
Info: job opportunities	-0.0628 (0.0764)	-0.0616 (0.109)	-0.0123 (0.0866)	0.0110 (0.0951)	-0.141 (0.0916)	-0.106 (0.0824)
Sender: researcher						
Info: job opp. & wages	-0.0284 (0.101)	-0.0886 (0.150)	0.0328 (0.0991)	-0.00867 (0.109)	-0.0544 (0.104)	-0.0343 (0.120)
Sender: institute						
Info: job opp. & wages	-0.0441 (0.119)	-0.0531 (0.161)	0.126 (0.129)	-0.0951 (0.123)	-0.110 (0.118)	-0.0585 (0.127)
2019/2020	-0.0111 (0.0547)	-0.0542 (0.0774)	0.0884 (0.0662)	0.00624 (0.0806)	-0.0298 (0.0615)	-0.0636 (0.0806)
Sender: researcher						
Info: job opportunities × 2019/2020	0.172 (0.108)	0.225* (0.130)	0.0570 (0.145)	0.240 (0.146)	0.109 (0.139)	0.177 (0.133)
Sender: institute						
Info: job opportunities × 2019/2020	0.0426 (0.0737)	0.0362 (0.116)	-0.0325 (0.0993)	0.0129 (0.113)	0.0359 (0.1000)	0.190* (0.115)
Sender: researcher						
Info: job opp. & wages × 2019/2020	-0.0672 (0.107)	0.0798 (0.138)	-0.176 (0.128)	-0.139 (0.136)	-0.0173 (0.125)	-0.0347 (0.153)
Sender: institute						
Info: job opp. & wages × 2019/2020	-0.0786 (0.0939)	-0.0604 (0.139)	-0.396*** (0.122)	-0.0962 (0.127)	-0.0110 (0.112)	0.126 (0.124)
Constant	16.79*** (0.0573)	16.86*** (0.0778)	16.61*** (0.0597)	16.72*** (0.0737)	16.81*** (0.0633)	16.85*** (0.0640)
Observations	28267	27805	27811	27801	27715	27598
F-Stat Non-interacted Treatments	1.224	1.215	1.065	1.974	1.320	0.442
P-value F-Stat Non-interacted Treatments	0.301	0.305	0.374	0.099	0.263	0.778
F-Stat Treatments + Treatments × 19/20	0.372	0.108	1.642	1.178	0.167	0.761
P-value F-Stat Treatments + Treatments × 19/20	0.828	0.980	0.164	0.321	0.955	0.551

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual level. Discrepancies in observations are caused by the fact that a few occupations could not be assigned to a level of hourly wages. We do calculate an average score for the other occupations a student selected in this case.

Table 3.C.4: Heterogeneity job opportunities of selected occupations

	(1)	(2)	(3)	(4)	(5)	(6)
	Mean value	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Age	0.000312 (0.00922)	0.00610 (0.0162)	0.00944 (0.0149)	-0.0195 (0.0157)	0.00515 (0.0156)	-0.00153 (0.0144)
3rd year	-0.0121 (0.0213)	-0.0122 (0.0338)	-0.0239 (0.0390)	0.0106 (0.0354)	-0.0491* (0.0292)	0.0204 (0.0255)
4th year	-0.00236 (0.0306)	0.0173 (0.0543)	-0.00492 (0.0555)	0.0565 (0.0476)	-0.113** (0.0525)	0.0363 (0.0434)
Male	0.212*** (0.0187)	0.0497* (0.0275)	0.178*** (0.0272)	0.282*** (0.0263)	0.291*** (0.0261)	0.295*** (0.0235)
QOL score	-0.000243 (0.00628)	0.0127 (0.0104)	-0.00224 (0.0100)	-0.00870 (0.0100)	-0.00843 (0.00879)	0.00672 (0.00975)
No. of Profiles Available=3	-0.00670 (0.0371)	0.0462 (0.0405)	0.0844* (0.0467)	-0.0768 (0.0805)	0.0316 (0.0514)	-0.126*** (0.0418)
No. of Profiles Available=4	-0.00275 (0.0375)	0.0561 (0.0387)	0.0988** (0.0481)	-0.0457 (0.0824)	0.00898 (0.0490)	-0.130*** (0.0368)
Constant	2.730*** (0.125)	2.658*** (0.229)	2.536*** (0.196)	3.042*** (0.236)	2.665*** (0.230)	2.766*** (0.217)
Observations	8576	8425	8422	8419	8394	8350

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual-occupation level. Only includes control group students. 2nd year, female students in schools where only 1 profile is available are baseline.

Table 3.C.5: Heterogeneity hourly wages of selected occupations by treatment group and year

	(1)	(2)	(3)	(4)	(5)	(6)
	Mean value	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Age	0.0297 (0.0319)	0.0912** (0.0431)	0.0926* (0.0472)	-0.00160 (0.0544)	0.00327 (0.0434)	-0.0401 (0.0556)
3rd year	0.256*** (0.0798)	0.310*** (0.108)	0.198* (0.113)	0.331*** (0.114)	0.160 (0.0967)	0.253** (0.106)
4th year	0.484*** (0.0933)	0.335** (0.148)	0.522*** (0.139)	0.526*** (0.175)	0.433*** (0.142)	0.521*** (0.170)
Male	0.852*** (0.0418)	0.759*** (0.0621)	0.935*** (0.0735)	0.950*** (0.0756)	0.943*** (0.0693)	0.698*** (0.0635)
QOL score	-0.0711*** (0.0203)	-0.0713** (0.0321)	-0.0650*** (0.0244)	-0.103*** (0.0316)	-0.0592** (0.0262)	-0.0486 (0.0324)
No. of Profiles Available=3	0.0234 (0.119)	-0.0399 (0.153)	0.223 (0.156)	-0.0464 (0.161)	0.0768 (0.108)	-0.133 (0.152)
No. of Profiles Available=4	0.185* (0.100)	0.242* (0.143)	0.408*** (0.144)	0.0880 (0.127)	0.0908 (0.0902)	0.0615 (0.145)
Constant	16.16*** (0.448)	15.37*** (0.652)	14.88*** (0.707)	16.76*** (0.767)	16.49*** (0.620)	17.23*** (0.782)
Observations	8576	8425	8422	8419	8394	8350

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual-occupation level. Only includes control group students. 2nd year, female students in schools where only 1 profile is available are baseline.

Table 3.C.6: Heterogeneity in prior beliefs

	(1)	(2)	(3)	(4)	(5)	(6)
	$O_{i,j}^{Prior} - O_j^{Actual}$	$ O_{i,j}^{Prior} - O_j^{Actual} $	$O_{i,j}^{Prior} - O_j^{Actual} = 0$	$W_{i,j}^{Prior} - W_j^{Actual}$	$ W_{i,j}^{Prior} - W_j^{Actual} $	$W_{i,j}^{Prior} - W_j^{Actual} = 0$
Age	0.00421 (0.00885)	0.000280 (0.00744)	0.00645** (0.00317)	0.0371 (0.0328)	0.0268 (0.0255)	0.00108 (0.00224)
3rd year	0.00935 (0.0213)	0.00441 (0.0134)	-0.0109** (0.00547)	-0.130 (0.0875)	-0.252*** (0.0595)	0.00432 (0.00473)
4th year	0.0785*** (0.0287)	-0.0192 (0.0213)	-0.00269 (0.0103)	-0.436*** (0.150)	-0.527*** (0.107)	0.0111 (0.00720)
Male	0.0677*** (0.0147)	0.0471*** (0.0110)	-0.0151*** (0.00529)	0.644*** (0.0798)	0.422*** (0.0487)	-0.0167*** (0.00384)
QOL score	0.00345 (0.00562)	-0.00286 (0.00408)	0.00120 (0.00159)	-0.0400 (0.0315)	-0.0424* (0.0220)	0.00266** (0.00125)
No. of Profiles Available=3	0.0419 (0.0331)	-0.0412** (0.0166)	0.0209** (0.00861)	0.107 (0.234)	-0.119 (0.0927)	0.00500 (0.00537)
No. of Profiles Available=4	0.0419 (0.0271)	-0.0294** (0.0146)	0.0156** (0.00767)	-0.0550 (0.223)	-0.165* (0.0832)	0.0122** (0.00545)
Rank=2	-0.235*** (0.0114)	-0.105*** (0.0113)	0.0271*** (0.00517)	-0.584*** (0.0356)	-0.344*** (0.0323)	0.00623 (0.00398)
Rank=3	-0.425*** (0.0131)	-0.174*** (0.0130)	0.0382*** (0.00647)	-0.902*** (0.0367)	-0.499*** (0.0314)	0.0113** (0.00508)
Rank=4	-0.611*** (0.0155)	-0.211*** (0.0128)	0.0511*** (0.00686)	-1.192*** (0.0500)	-0.546*** (0.0348)	0.0110** (0.00466)
Rank=5	-0.818*** (0.0200)	-0.200*** (0.0153)	0.0438*** (0.00759)	-1.485*** (0.0529)	-0.583*** (0.0427)	0.0187*** (0.00434)
Constant	-0.241 (0.190)	1.138*** (0.129)	0.111 (0.0678)	-3.519*** (0.712)	4.856*** (0.540)	-0.0272 (0.0354)
Observations	41842	41842	41842	41825	41825	41825

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual-occupation level. Regressions include occupation dummies. Only includes control group students.

Table 3.C.7: Treatment effect on posterior beliefs job opportunities by prior belief accuracy

	(1)	(2)
	$ O_{i,j}^{Post} - O_j^{Actual} $	$O_{i,j}^{Post} - O_j^{Actual} = 0$
Treated	0.0319*** (0.00824)	0.0154* (0.00860)
$ O_{i,j}^{Prior} - O_j^{Actual} $	0.911*** (0.00456)	-0.241*** (0.00361)
Treated $\times O_{i,j}^{Prior} - O_j^{Actual} $	-0.112*** (0.00925)	0.00920* (0.00478)
$(O_{i,j}^{Prior} - O_j^{Actual} > 0)$	-0.0242*** (0.00600)	-0.339*** (0.00559)
Treated $\times (O_{i,j}^{Prior} - O_j^{Actual} > 0)$	-0.0574*** (0.0115)	0.0823*** (0.00782)
Wage information	0.0205** (0.00961)	-0.00861 (0.00949)
Wage information $\times O_{i,j}^{Prior} - O_j^{Actual} $	-0.0197 (0.0137)	0.00949* (0.00548)
Wage information $\times (O_{i,j}^{Prior} - O_j^{Actual} > 0)$	-0.0293* (0.0174)	0.0168 (0.0102)
Constant	0.120** (0.0522)	0.557*** (0.0368)
Observations	64579	64579

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual-occupation level. Regressions contain occupation dummies. Treated = All treatment groups. Wage info = Treatments 3 & 4.

Table 3.C.8: Treatment effect on posterior beliefs hourly wages by prior belief accuracy

	(1)	(2)
	$ W_{i,j}^{Post} - W_j^{Actual} $	$W_{i,j}^{Post} - W_j^{Actual} = 0$
Treated	0.188*** (0.0306)	-0.0250*** (0.00909)
$ W_{i,j}^{Prior} - W_j^{Actual} $	0.914*** (0.00406)	-0.0222*** (0.000609)
Treated \times $ W_{i,j}^{Prior} - W_j^{Actual} $	-0.0481*** (0.00660)	0.00294*** (0.000859)
$(W_{i,j}^{Prior} - W_j^{Actual} > 0)$	0.00122 (0.0179)	-0.173*** (0.00466)
Treated \times $(W_{i,j}^{Prior} - W_j^{Actual} > 0)$	-0.0337 (0.0307)	0.0151** (0.00668)
Wage information	0.0232 (0.0468)	0.0928*** (0.0114)
Wage information \times $ W_{i,j}^{Prior} - W_j^{Actual} $	-0.156*** (0.0127)	-0.00325*** (0.00121)
Wage information \times $(W_{i,j}^{Prior} - W_j^{Actual} > 0)$	-0.114** (0.0536)	0.0293*** (0.00998)
Constant	0.458*** (0.159)	0.196*** (0.0227)
Observations	64565	64565

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual-occupation level. Regressions contain occupation dummies. Treated = All treatment groups. Wage info = Treatments 3 & 4.

Table 3.C.9: Detailed treatment effect on posterior beliefs

	(1)	(2)	(3)	(4)	(5)	(6)
	$O_t^{Post} - O_t^{Actual}$	$ O_t^{Post} - O_t^{Actual} $	$O_t^{Post} - O_t^{Actual} = 0$	$W_t^{Post} - W_t^{Actual}$	$ W_t^{Post} - W_t^{Actual} $	$W_t^{Post} - W_t^{Actual} = 0$
Sender: researcher						
Info: job opportunities	-0.103*** (0.0135)	-0.135*** (0.0159)	0.0721*** (0.00821)	-0.173** (0.0722)	-0.127** (0.0585)	0.00334 (0.00265)
Sender: institute						
Info: job opportunities	-0.103*** (0.0171)	-0.139*** (0.0120)	0.0762*** (0.00627)	-0.0627 (0.0726)	-0.0367 (0.0591)	0.00271 (0.00309)
Sender: researcher						
Info: job opp. & wages	-0.135*** (0.0137)	-0.138*** (0.0124)	0.0714*** (0.00563)	-0.870*** (0.0843)	-0.914*** (0.0685)	0.113** (0.00590)
Sender: institute						
Info: job opp. & wages	-0.127*** (0.0155)	-0.163*** (0.0140)	0.0854*** (0.00704)	-0.799*** (0.0758)	-0.920*** (0.0690)	0.110*** (0.00488)
Constant	-0.559*** (0.0706)	1.005*** (0.0539)	0.249*** (0.0309)	-3.247*** (0.287)	4.486*** (0.229)	0.0204 (0.0142)
Observations	136721	136721	136721	136707	136707	136707
F-Stat Researcher; Job Opp. = Institute; Job Opp.	0.000	0.039	0.185			
P-value Researcher; Job Opp. = Institute; Job Opp.	0.999	0.844	0.667			
F-Stat Researcher; Job Opp. & Wage = Institute; Job Opp. & Wage	0.249	2.038	2.824	0.568	0.005	0.219
P-value F-Stat f2 F-Stat Researcher; Job Opp. & Wage = Institute; Job Opp. & Wage	0.618	0.155	0.094	0.452	0.944	0.640

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual-occupation level. Regressions contain occupation dummies.

Table 3.C.10: Sender effect on posterior beliefs

	(1)	(2)	(3)	(4)	(5)	(6)
	$O_{i,j}^{Post} - O_j^{Actual}$	$ O_{i,j}^{Post} - O_j^{Actual} $	$O_{i,j}^{Post} - O_j^{Actual} = 0$	$W_{i,j}^{Post} - W_j^{Actual}$	$ W_{i,j}^{Post} - W_j^{Actual} $	$W_{i,j}^{Post} - W_j^{Actual} = 0$
Sender: Female - Low status	-0.00271 (0.0250)	-0.0100 (0.0198)	-0.00256 (0.00918)	0.388** (0.188)	0.0337 (0.129)	-0.00948 (0.0166)
Sender: Male - High status	-0.00166 (0.0263)	0.0103 (0.0238)	-0.00616 (0.0108)	0.133 (0.122)	-0.0303 (0.119)	0.00212 (0.0154)
Sender: Male - Low status	0.00609 (0.0277)	0.0236 (0.0250)	-0.00504 (0.0106)	0.0675 (0.225)	-0.104 (0.165)	0.0188 (0.0151)
Male	0.100*** (0.0278)	0.105*** (0.0240)	-0.0412*** (0.0104)	0.583*** (0.196)	0.391*** (0.120)	-0.0330** (0.0160)
Sender: Female - Low status × Male	-0.0102 (0.0405)	0.0163 (0.0315)	0.00345 (0.0154)	-0.479* (0.244)	-0.214 (0.153)	0.0294 (0.0232)
Sender: Male - High status × Male	0.0312 (0.0438)	-0.0128 (0.0400)	0.0104 (0.0167)	-0.223 (0.214)	-0.127 (0.185)	0.0227 (0.0206)
Sender: Male - Low status × Male	0.0356 (0.0402)	-0.0196 (0.0376)	0.0195 (0.0162)	0.140 (0.236)	0.0185 (0.215)	-0.000800 (0.0217)
Constant	-0.690*** (0.118)	0.696*** (0.0921)	0.373*** (0.0606)	-3.798*** (0.810)	3.643*** (0.635)	0.101* (0.0587)
Observations	44964	44964	44964	20466	20466	20466

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual-occupation level. Regressions contain occupation dummies. Female students with female high status sender are baseline. Regressions (1), (2) and (3) contain treatments 1 & 3. Regressions (4), (5) and (6) conly contain treatment 3.

Table 3.C.11: Heterogeneous treatment effects on posterior beliefs job opportunities

	(1)	(2)	(3)
	$O_{ij}^{Post} - O_{ij}^{Actual}$	$ O_{ij}^{Post} - O_{ij}^{Actual} $	$O_{ij}^{Post} - O_{ij}^{Actual} = 0$
Treated	-0.0783** (0.0374)	-0.154*** (0.0228)	0.0759*** (0.0117)
Age (demeaned)	0.00340 (0.00877)	0.000383 (0.00725)	0.00659** (0.00313)
Treated × Age (demeaned)	0.00432 (0.0116)	0.00924 (0.00948)	-0.0106* (0.00461)
3rd year	0.00995 (0.0210)	0.00359 (0.0129)	-0.0115** (0.00570)
4th year	0.0902*** (0.0272)	-0.0236 (0.0202)	-0.00475 (0.00975)
Treated × 3rd year	-0.00166 (0.0257)	-0.0449** (0.0194)	0.0368*** (0.00941)
Treated × 4th year	-0.0585 (0.0385)	-0.0897*** (0.0330)	0.0587*** (0.0167)
Male	0.0632*** (0.0145)	0.0452*** (0.0107)	-0.00691 (0.00506)
Treated × Male	0.0511*** (0.0183)	0.0566*** (0.0129)	-0.0287*** (0.00617)
QOL score (demeaned)	0.00438 (0.00609)	-0.00309 (0.00436)	0.000764 (0.00175)
Treated × QOL score (demeaned)	-0.00702 (0.00736)	-0.00333 (0.00549)	0.00244 (0.00239)
No. of Profiles Available=3	0.0462 (0.0370)	-0.0433** (0.0169)	0.0201** (0.00983)
No. of Profiles Available=4	0.0504 (0.0315)	-0.0330** (0.0159)	0.0129 (0.00906)
Treated × No. of Profiles Available=3	-0.0511 (0.0421)	0.0191 (0.0243)	-0.0113 (0.0131)
Treated × No. of Profiles Available=4	-0.0787** (0.0362)	-0.0178 (0.0226)	0.00822 (0.0120)
Constant	-0.676*** (0.0729)	0.972*** (0.0569)	0.257*** (0.0319)
Observations	125647	125647	125647

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual-occupation level. Only includes control group students. Treated = All Treatments. 2nd year, female students in schools where only 1 profile is available are baseline.

Table 3.C.12: Heterogeneous treatment effects on posterior beliefs hourly wages

	(1)	(2)	(3)
	$W_i^{Post} - W_i^{Actual}$	$ W_i^{Post} - W_i^{Actual} $	$W_i^{Post} - W_i^{Actual} = 0$
Wage information	-0.712*** (0.252)	-1.013*** (0.109)	0.117*** (0.00849)
Age (demeaned)	0.0285 (0.0337)	0.0182 (0.0257)	0.00130 (0.00217)
Wage information × Age (demeaned)	-0.0235 (0.0508)	-0.00156 (0.0404)	-0.00640 (0.00433)
3rd year	-0.119 (0.0871)	-0.249*** (0.0553)	0.00474 (0.00471)
4th year	-0.419*** (0.158)	-0.484*** (0.106)	0.00865 (0.00799)
Wage information × 3rd year	0.117 (0.141)	-0.00494 (0.110)	0.0167 (0.0114)
Wage information × 4th year	0.289 (0.220)	0.00619 (0.182)	0.0305* (0.0162)
Male	0.618*** (0.0771)	0.370*** (0.0493)	-0.0173*** (0.00460)
Wage information × Male	-0.0446 (0.0995)	0.0921 (0.0710)	-0.00908 (0.00765)
QOL score (demeaned)	-0.0390 (0.0338)	-0.0389* (0.0227)	0.00185 (0.00123)
Wage information × QOL score (demeaned)	0.0360 (0.0407)	0.0216 (0.0322)	-0.00283 (0.00234)
No. of Profiles Available=3	0.0950 (0.258)	-0.164* (0.0984)	0.0119** (0.00495)
No. of Profiles Available=4	-0.0871 (0.247)	-0.204** (0.0839)	0.0163*** (0.00506)
Wage information × No. of Profiles Available=3	-0.170 (0.280)	0.0753 (0.137)	-0.00454 (0.00995)
Wage information × No. of Profiles Available=4	-0.261 (0.262)	-0.0408 (0.115)	-0.00568 (0.00918)
Constant	-3.743*** (0.397)	4.765*** (0.349)	-0.00183 (0.0213)
Observations	80042	80042	80042

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual-occupation level. Only includes control group students. Wage info = Treatments 3 & 4. Treatments 1 & 2 are excluded from this analysis. 2nd year, female students in schools where only 1 profile is available are baseline.

Table 3.C.13: Long-term treatment effect on posterior beliefs

	(1)	(2)	(3)	(4)	(5)	(6)	
	$O_{i,t}^{Survey} - O_{i,t}^{Actual}$	$ O_{i,t}^{Survey} - O_{i,t}^{Actual} $	$O_{i,t}^{Survey} - O_{i,t}^{Actual} = 0$	$W_{i,t}^{Survey} - W_{i,t}^{Actual}$	$ W_{i,t}^{Survey} - W_{i,t}^{Actual} $	$W_{i,t}^{Survey} - W_{i,t}^{Actual} = 0$	
Sender: researcher							
Info: job opportunities	-0.0173 (0.125)	0.00867 (0.0612)	-0.0590** (0.0282)	-0.490 (0.601)	0.198 (0.276)	-0.0138 (0.0191)	
Sender: institute							
Info: job opportunities	0.284** (0.140)	0.166** (0.0713)	-0.0415* (0.0244)	0.167 (0.519)	0.336 (0.260)	-0.0238 (0.0200)	
Sender: researcher							
Info: job opp. & wages	0.269* (0.157)	0.0508 (0.0760)	-0.0530* (0.0317)	1.189* (0.604)	0.353 (0.254)	-0.0344 (0.0235)	
Sender: institute							
Info: job opp. & wages	-0.207 (0.161)	-0.0380 (0.0668)	0.00785 (0.0322)	0.121 (0.543)	0.0692 (0.258)	0.0103 (0.0318)	
2019/2020	0.238 (0.208)	0.184* (0.0962)	-0.120*** (0.0454)	-0.464 (1.224)	0.643 (0.424)	-0.0231 (0.0254)	
Sender: researcher							
Info: job opportunities	× 2019/2020	-0.00801 (0.265)	-0.397** (0.153)	0.205*** (0.0709)	-2.186 (1.656)	0.0274 (0.885)	0.115* (0.0674)
Sender: institute							
Info: job opportunities	× 2019/2020	-0.181 (0.234)	-0.389*** (0.118)	0.0716 (0.0659)	0.336 (1.287)	-1.107** (0.485)	0.0777* (0.0396)
Sender: researcher							
Info: job opp. & wages	× 2019/2020	-0.136 (0.265)	-0.302* (0.165)	0.217** (0.0846)	-0.846 (1.368)	-0.765 (0.693)	0.0378 (0.0380)
Sender: institute							
Info: job opp. & wages	× 2019/2020	0.471 (0.293)	-0.286** (0.122)	0.152** (0.0616)	0.0694 (1.497)	-0.675 (0.760)	-0.00231 (0.0428)
Constant							
	0.216** (0.0859)	1.154*** (0.0422)	0.284*** (0.0171)	-0.741** (0.368)	3.857*** (0.151)	0.0887*** (0.0129)	
Observations	2079	2057	1928	1928	1909	1928	
F-Stat Non-interacted (Wage) Treatments	2.907	2.044	1.950	2.142	0.992	1.286	
P-value F-Stat Non-interacted (Wage) Treatments	0.025	0.093	0.107	0.122	0.374	0.281	
F-Stat (Wage) Treatments + (Wage) Treatments x 19/20	0.570	2.622	3.496	0.066	0.805	0.222	
P-value F-Stat (Wage) Treatments + (Wage) Treatments x 19/20	0.685	0.039	0.010	0.936	0.450	0.801	

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. F-Stat Non-interacted Treatments is compared to 2018/2019 control group. F-Stat Treatments + Treatments × 19/20 is compared to 2019/2020 control group.

Table 3.C.14: Heterogeneous treatment effect on favorite occupation's job opportunities

	(1) Pr(Fav. Change)	(2) ΔO_j^{Actual}	(3) ΔO_j^{Actual} (Changed)
Treated	0.0249*** (0.00939)	-0.00210 (0.0107)	-0.226 (0.183)
Age (demeaned)	-0.00159 (0.00257)	0.00862* (0.00441)	0.221** (0.110)
Treated × Age (demeaned)	0.00300 (0.00371)	-0.0111* (0.00567)	-0.270** (0.120)
3rd year	-0.00226 (0.00584)	-0.00539 (0.00781)	-0.199 (0.161)
4th year	-0.0264*** (0.00896)	-0.0103 (0.0117)	-0.254 (0.345)
Treated × 3rd year	-0.00773 (0.00776)	0.00867 (0.0103)	0.303 (0.186)
Treated × 4th year	0.00653 (0.0129)	0.0155 (0.0162)	0.471 (0.383)
Male	0.0160*** (0.00528)	-0.00129 (0.00739)	-0.0415 (0.129)
Treated × Male	-0.00538 (0.00675)	0.00515 (0.00933)	0.0406 (0.151)
QOL score (demeaned)	-0.00357 (0.00216)	-0.00153 (0.00255)	-0.0126 (0.0401)
Treated × QOL score (demeaned)	0.0000950 (0.00267)	0.00280 (0.00338)	0.0445 (0.0495)
No. of Profiles Available=3	0.0185* (0.00945)	-0.00702 (0.00967)	-0.255 (0.161)
No. of Profiles Available=4	0.0171** (0.00824)	-0.0185** (0.00827)	-0.447*** (0.136)
Treated × No. of Profiles Available=3	-0.00217 (0.0114)	0.00296 (0.0132)	0.128 (0.209)
Treated × No. of Profiles Available=4	-0.00255 (0.00971)	0.0289*** (0.0110)	0.528*** (0.176)
Constant	0.0328*** (0.00790)	0.0192** (0.00791)	0.530*** (0.146)
Observations	25174	25174	1634

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Treated = All Treatments. Regressions at individual level. 2nd year, female students in schools where only 1 profile is available are baseline.

Table 3.C.15: Heterogeneous treatment effect on favorite occupation's hourly wages

	(1) Pr(Fav. Change)	(2) ΔW_{Actual}	(3) ΔW_{Actual} (Changed)
Wage information	0.0207* (0.0118)	0.0101 (0.0405)	0.205 (0.764)
Age (demeaned)	-0.00159 (0.00257)	0.00141 (0.0120)	0.0468 (0.305)
Wage information \times Age (demeaned)	0.00254 (0.00487)	-0.00235 (0.0234)	-0.133 (0.410)
3rd year	-0.00226 (0.00585)	-0.0238 (0.0222)	-0.468 (0.441)
4th year	-0.0264*** (0.00897)	0.0352 (0.0327)	1.472 (1.118)
Wage information \times 3rd year	0.00239 (0.0101)	0.0657 (0.0441)	1.112* (0.661)
Wage information \times 4th year	0.00676 (0.0150)	0.0222 (0.0611)	-0.0980 (1.397)
Male	0.0160*** (0.00529)	-0.0161 (0.0204)	-0.375 (0.366)
Wage information \times Male	0.00181 (0.00827)	0.0552 (0.0363)	0.549 (0.517)
QOL score (demeaned)	-0.00357 (0.00216)	0.00636 (0.00694)	0.103 (0.121)
Wage information \times QOL score (demeaned)	0.000869 (0.00327)	-0.0222* (0.0117)	-0.251 (0.157)
No. of Profiles Available=3	0.0185* (0.00946)	0.0370 (0.0297)	0.514 (0.651)
No. of Profiles Available=4	0.0171** (0.00825)	0.00310 (0.0246)	-0.00864 (0.554)
Wage information \times No. of Profiles Available=3	0.00197 (0.0141)	-0.00110 (0.0431)	-0.336 (0.786)
Wage information \times No. of Profiles Available=4	-0.00126 (0.0126)	0.0472 (0.0418)	0.461 (0.739)
Constant	0.0328*** (0.00791)	0.00990 (0.0228)	0.346 (0.581)
Observations	16036	16036	1034

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Wage info = Treatments 3 & 4. Treatments 1 & 2 are excluded from this analysis. are excluded Regressions at individual level. 2nd year, female students in schools where only 1 profile is available are baseline.

Table 3.C.16: Effect of sender on likelihood favorite occupation changing and its prospects

	(1)	(2)	(3)	(4)	(5)
	Pr(Fav. Change)	ΔO_j^{Actual}	ΔO_j^{Actual} (Changed)	ΔW_j^{Actual}	ΔW_j^{Actual} (Changed)
Sender: Female - Low status	-0.00203 (0.0106)	-0.000304 (0.0163)	0.00430 (0.252)	0.00653 (0.0675)	0.126 (0.954)
Sender: Male - High status	0.00575 (0.0118)	-0.00312 (0.0169)	-0.0667 (0.242)	0.0896 (0.0722)	1.105 (0.840)
Sender: Male - Low status	-0.00370 (0.0128)	0.0151 (0.0146)	0.262 (0.227)	0.111 (0.0822)	2.233 (1.326)
Male	0.000274 (0.0102)	0.00166 (0.0150)	0.0241 (0.232)	0.0736 (0.0811)	0.934 (1.047)
Sender: Female - Low status × Male	0.00811 (0.0147)	-0.00507 (0.0194)	-0.105 (0.293)	-0.0000434 (0.115)	-0.0565 (1.493)
Sender: Male - High status × Male	-0.00110 (0.0157)	0.0146 (0.0238)	0.209 (0.334)	-0.0299 (0.118)	-0.482 (1.383)
Sender: Male - Low status × Male	0.0239 (0.0172)	0.00338 (0.0214)	-0.119 (0.288)	-0.0904 (0.114)	-2.264 (1.568)
Constant	0.0655*** (0.00807)	0.0186 (0.0121)	0.284 (0.182)	0.0170 (0.0480)	0.229 (0.642)
Observations	9016	9016	629	4099	315

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual level. Regressions contain occupation dummies. Female students with female high status sender are baseline. Regressions (1), (2) and (3) contain treatments 1 & 3. Regressions (4), (5) and (6) conly contain treatment 3.

Table 3.C.17: Long-term treatment effect on prospects favorite occupation

	(1)	(2)	(3)	(4)
	ΔO_j^{Actual} (Experiment)	ΔO_j^{Actual} (Survey)	ΔW_j^{Actual} (Experiment)	ΔW_j^{Actual} (Survey)
Sender: researcher				
Info: job opportunities	0.0416 (0.0676)	-0.0825 (0.128)	0.0525 (0.0865)	0.0150 (0.294)
Sender: institute				
Info: job opportunities	-0.0334 (0.0310)	0.0211 (0.126)	-0.0432 (0.0478)	-0.0688 (0.367)
Sender: researcher				
Info: job opp. & wages	0.00715 (0.0353)	0.136 (0.160)	0.00715 (0.0726)	-0.455 (0.478)
Sender: institute				
Info: job opp. & wages	-0.0236 (0.0213)	-0.0768 (0.132)	-0.0236 (0.0432)	0.157 (0.334)
Constant	0.0236 (0.0213)	0.126 (0.0867)	0.0236 (0.0432)	0.0394 (0.175)
Observations	447	447	447	447

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. F-Stat Non-interacted Treatments is compared to 2018/2019 control group. F-Stat Treatments + Treatments \times 19/20 is compared to 2019/2020 control group.

Table 3.C.18: Heterogeneity in change in profiles considered immediately after intervention

	(1)	(2)	(3)
	Pr(Same Profile Choice)	No. of Theoretical Profiles	No. of Other Profiles
Treated	0.0543 (0.0495)	0.0123 (0.0519)	-0.151** (0.0704)
3 theoretical profiles available	0.197*** (0.0405)	-0.0491 (0.0506)	-0.185** (0.0854)
4 theoretical profiles available	0.153*** (0.0376)	0.00139 (0.0494)	-0.211*** (0.0711)
3 theoretical profiles available × Treated	-0.0572 (0.0550)	-0.00201 (0.0578)	0.219** (0.0940)
4 theoretical profiles available × Treated	-0.0816 (0.0550)	-0.0278 (0.0569)	0.242*** (0.0921)
Wage information	-0.0202 (0.0438)	-0.0614 (0.0551)	-0.0250 (0.0458)
3 theoretical profiles available × Wage information	-0.0400 (0.0581)	0.0950 (0.0665)	0.0163 (0.0625)
4 theoretical profiles available × Wage information	0.0516 (0.0506)	0.0567 (0.0600)	0.00929 (0.0724)
No. of theoretical profiles a priori		0.553*** (0.0190)	
No. of other profiles a priori			0.443*** (0.0329)
Constant	0.538*** (0.0356)	0.183*** (0.0475)	0.451*** (0.0649)
Observations	10671	5901	4772

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Table 3.C.19: Impact of occupational information on likelihood choosing related profile

	(1)	(2)
	Pr(Chose Profile of Occupation)	Pr(Chose Profile of Occupation)
Treated	-0.00154 (0.00870)	0.0150 (0.0188)
Job opportunities	-0.00240 (0.00184)	
Treated × Job opportunities	-0.0000956 (0.00250)	
Wage information	-0.00727 (0.00828)	-0.00130 (0.0174)
Wage information × Job opportunities	0.00111 (0.00244)	
Chose profile a priori	0.784*** (0.00657)	0.784*** (0.00657)
Hourly wage		0.000731 (0.000708)
Treated × Hourly wage		-0.000987 (0.00108)
Wage information × Hourly wage		-0.000146 (0.00103)
Constant	0.0609*** (0.00649)	0.0408*** (0.0129)
Observations	52785	52785

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Table 3.C.20: Treatment effect on prospects of chosen study program by gender

	(1)	(2)
	Job Opportunities Chosen Program	Average Hourly Wages Chosen Program
Treated	0.209 (0.177)	0.105 (0.269)
Male	0.000877 (0.226)	0.858** (0.387)
Treated × Male	-0.0310 (0.280)	0.146 (0.493)
Wage information	-0.346 (0.218)	0.409 (0.292)
Wage information × Male	0.293 (0.340)	-0.514 (0.504)
Constant	2.361*** (0.108)	17.15*** (0.191)
Observations	405	405

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual level.

Table 3.C.21: Treatment effect on prospects of chosen study program by QOL score

	(1) Job Opportunities Chosen Program	(2) Average Hourly Wages Chosen Program
Treated	0.147 (0.160)	0.151 (0.204)
QOL score	0.0300 (0.0634)	-0.0816 (0.0969)
Treated × QOL score	0.0363 (0.0883)	0.0145 (0.131)
Wage information	-0.220 (0.189)	-0.00441 (0.252)
Wage information × QOL score	-0.0242 (0.106)	0.313** (0.150)
Constant	2.351*** (0.0931)	17.49*** (0.134)
Observations	399	399

Note: standard errors clustered at the school level in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions at individual level.

4

Jobs Reports Affect Personal Job Loss Expectations

Abstract

Using data from the New York Federal Reserve's Survey of Consumer Expectations, we study how the United States Bureau of Labor Statistics' Employment Situation Reports (Jobs Reports) affect individuals' expectations about the likelihood of losing their own job. We do this in two steps. First, we estimate the information shocks of the Jobs Reports on expectations about the development of the national unemployment rate in the next twelve months. We do this by comparing survey responses shortly before and after publication of the reports. Second, we estimate how these shocks affect individuals' expectations about losing their own job in the same time frame. The results show that when a report is estimated to increase beliefs about the likelihood of the unemployment rate increasing by 1 percentage point, beliefs about the likelihood of personal job loss during that time increase by up to 0.22 percentage points. We further find that the information shock negatively affects individuals' beliefs about the likelihood of finding a new job if they were to lose their current one, but (surprisingly) positively affects individuals' beliefs about the likelihood of voluntarily leaving their job. Our results are robust to the use of different bandwidths around the reports' publication dates and placebo treatments provide reassurance that the information shock is indeed the mechanism driving the result.

This chapter is based on joint work with Didier Fouarge and Johannes Schuffels. We make use of the Survey of Consumer Expectations. Source: Survey of Consumer Expectations, © 2013-2020 Federal Reserve Bank of New York (FRBNY). The SCE data are available without charge at <http://www.newyorkfed.org/microeconomics/sce> and may be used subject to license terms posted there. FRBNY disclaims any responsibility for this analysis and interpretation of Survey of Consumer Expectations data.

4.1 Introduction

Providing individuals with information about the macroeconomy affects their personal economic expectations and behavior (Roth & Wohlfart, 2020). This evidence is based on an experimental study, however; it is unclear to what degree this is true for information that people acquire in their day-to-day lives. In this chapter, we study the effect of the publication of the United States Bureau of Labor Statistics' Employment Situation Reports on individuals' expectations about the likelihood of losing their own job.

Every month, the United States Bureau of Labor Statistics publishes its Employment Situation Report, often referred to as the 'Jobs Report'. The report contains, among other things, information about the unemployment rate in the United States. As the Jobs Reports receive considerable attention in the media, it is likely that these reports play an important role in shaping individuals' expectations about the employment situation at the national level. The question we ask is whether it affects expectations about their own job security as well.

Understanding how information about the development of the unemployment rate affects individuals' expectations about the likelihood of losing their job is important, as these expectations have a wide range of implications. At the individual level, an increase in the expected likelihood of job loss is related to a decrease in expected earnings (Stephens Jr, 2004; Campbell, Carruth, Dickerson, & Green, 2007) as well as (and perhaps therefore) consumption (Pettinicchi & Vellekoop, 2019; Hendren, 2017; Brown & Taylor, 2006), saving and borrowing (Kłopocka, 2017). Fears of unemployment tend to be warranted, as they are indeed related to actual job loss (Dickerson & Green, 2012) and lower wage

growth (Campbell et al., 2007). Economic expectations affect health outcomes too. The more likely individuals think it is that they will lose their job, the higher they score on depression scales (Mandal et al., 2011) and the more likely they are to develop a range of health issues (Caroli & Godard, 2016). At the macroeconomic level, the increase in precautionary savings because of unemployment fears may lead to a deflationary spiral under incomplete financial markets (Den Haan, Rendahl, & Riegler, 2018; Ravn & Sterk, 2017).

We employ an event study approach using data from the New York Federal Reserve's Survey of Consumer Expectations to study the extent to which the Jobs Reports affect individuals' expectations. Event study approaches have been commonly used in research on the effects of monetary policy announcements on macroeconomic expectations of various types of agents (see e.g., Bulligan, 2018; Bottone and Rosolia, 2019; De Fiore, Lombardi, and Schuffels, 2021; Lamla and Vinogradov, 2019; Mertens, Lewis, and Makridis, 2020). We restrict our sample to include working individuals only.

The two most relevant questions in the Survey of Consumer Expectations for our analyses are about (a) individuals' beliefs about the likelihood that the unemployment rate will increase in the twelve months following their response and (b) their beliefs about the likelihood that they will lose their own job in the same time frame. For every published Jobs Report, we first estimate its impact on individuals' beliefs about the development of the national unemployment rate by comparing respondents who answered the survey shortly before the report's publication to those who answered the survey shortly after the report's publication. Second, we estimate how the change in the expectations

about the national unemployment rate is related to individuals' expectations about losing their own job.

We cannot use the full sample for both steps, however. If there is a correlation between individuals' expectations about the aggregate unemployment and their expectations about personal job loss caused by their personal circumstances, this would bias the results. Especially as the number of observations surrounding each Jobs Report is limited. It seems likely that the two outcomes are indeed correlated as described. Earlier literature has shown that individuals who experience unemployment, become more pessimistic about aggregate unemployment as well (Kuchler & Zafar, 2019). There is further evidence that the way in which individuals perceive macroeconomic conditions depends on their lived experiences. For instance, individuals give disproportional weight to inflation experienced during their lifetime when forming inflation expectations (Malmendier & Nagel, 2016). This also impacts behavior. Having experienced low stock market returns decreases the willingness to take financial risks and own stocks (Malmendier & Nagel, 2011), and growing up in a recession increases the relative importance individuals assign to income compared to meaning in their jobs (Cotofan, Cassar, Dur, & Meier, 2020). We tackle the above mentioned issue by employing two ways of splitting the sample between the first and second step.

In the first strategy, we split the sample in two mutually exclusive groups. We use the first group to estimate the impact of the Jobs Report on the expectations about the national unemployment rate. We use this estimate as the treatment intensity variable when we estimate the impact of the Jobs Report on personal job loss expectations in the

second group. The advantage of this methodology is that we avoid the bias caused by individual-level correlations between expectations about the national unemployment rate and personal job loss expectations. The downside of this methodology is that our number of observations is effectively cut in half, significantly decreasing statistical power.

Our second strategy tackles the latter issue. It involves using a jack-knife procedure. For each individual, we estimate the information shock of the Jobs Report on expectations about the national unemployment rate among all other participants in the survey. We then use this leave-out estimate as the treatment intensity variable in the second step. While this strategy does not negatively impact our effective sample size, it does have a different downside. While the national unemployment rate expectations of the individual do not affect the estimate assigned to them, they do affect the estimate assigned to every other individual. This means that those with the most extreme opinions about the national unemployment rate, get assigned the most conservative coefficients. This may lead to a downward bias. We show, however, that the variance in interpretation predominantly comes from between-Jobs Report differences. This means the bias is unlikely to be large.

We find that a Jobs Report we estimate to increase expectations about the likelihood that the unemployment rate will increase by 1 percentage point, leads to an increase of up to 0.22 percentage points in the expected likelihood of personal job loss during the same period. While the exact size of the estimate varies by bandwidth and approach used, the qualitative impact is consistent across these dimensions. Using

the jackknife procedure, so as to conserve power, we find some evidence that expectations about personal job loss are more strongly affected for individuals below 40, and above 60 years of age, than for those between 40 and 60 years of age. We further show that Jobs Reports that lead to an increase in expectations about the likelihood of the unemployment rate increasing also negatively affect individuals' expected likelihood of finding a new job if they were to lose their current one. Somewhat surprisingly, the information shock positively affects individuals' expectations about leaving their jobs voluntarily. We argue that the most likely reason for this is increased search effort as a response to larger perceived job uncertainty. We find no evidence of an effect on expected earnings and expected spending. To ensure our result is not driven by idiosyncratic shocks affecting both unemployment expectations and personal job loss expectations, we conduct a placebo treatment analysis: we move the 'treatment' date forward by two weeks and show that the effects of the Jobs Reports disappear when we do so.

The rest of this chapter is structured as follows. Section 4.2 provides more context on the Jobs Reports. Section 4.3 describes the data and methodology in more detail. Section 4.4 presents our main results, as well as robustness checks of our estimates, our analyses of heterogeneous treatment effects and of alternative outcomes. Section 4.5 concludes.

4.2 The Jobs Reports

The Jobs Reports, formally known as the Employment Situation Summaries, are monthly-published reports by the United States Bureau

of Labor Statistics. The reports contain information on the change in nonfarm payroll employment as well as the unemployment rate in the United States. The report about a certain month is generally published in the first week of the subsequent month.

The reports garner considerable attention – both on financial markets and among the wider public. A number of event studies found significant movements in prices and trading activity after publications of Jobs Reports on exchange rate markets (e.g., Harris and Zabka, 1995), bond markets (e.g., Fleming and Remolona, 1997; Green, 2004) and stock markets (e.g., Graham, Nikkinen, and Sahlström, 2003; Rangel, 2011; Chan and Gray, 2018). This attention on financial markets is mirrored in economic journalism, the channel through which large parts of the general public receive information about the Jobs Reports. The New York Times referred to the monthly report as “the government’s most watched economic indicator” (2018). Indeed, Google shows that approximately 10,900 news articles have been published referencing the Jobs Reports between June of 2013 and October of 2019. At the time of writing this chapter, ahrefs.com’s backlink checker shows that the main employment situation summary page has 310,550 backlinks¹, which are links from other websites referring to the Jobs Report. The PDF version of the document² has 143,720.

Individuals also look for the Jobs Reports themselves. Panel (a) of Figure 4.1 shows the average relative search interest for the Jobs Reports from June 2013 to October 2019 in the days surrounding the report’s publication. The Panel clearly shows that interest in the reports peaks

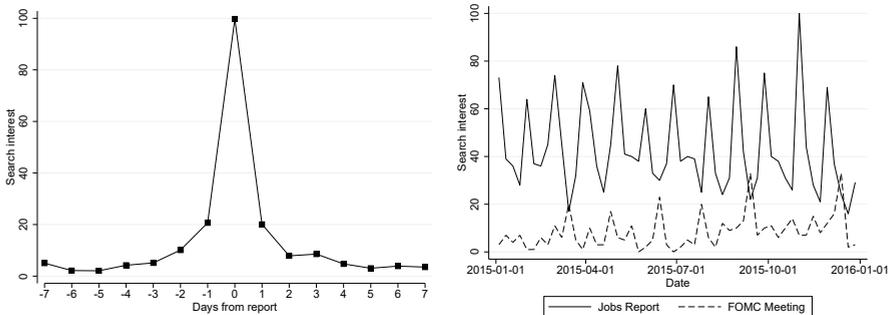
¹<https://ahrefs.com/backlink-checker>; Backlinks for URL <https://www.bls.gov/news.release/empisit.nr0.htm>; Retrieved 31 January, 2021.

²<https://www.bls.gov/news.release/pdf/empisit.pdf>

on the day of publication. There is some anticipation visible in the days leading up to the report’s publication, as well as some persisting interest afterwards. However, search volume is barely higher than 20% of the publication day peak on the days right before and after, declining to below 10% on the other days.

Unfortunately, Google Trends do not provide any insight into the actual search volume. However, Panel (b) shows how the search interest for the Jobs Reports compares to search interest for the Federal Open Market Committee (FOMC) meetings. These meetings are also one-day events that provide information about the macroeconomy. Similar to the Jobs Reports, there are clear search interest peaks in the periods where FOMC meetings took place. However, the search interest for the Jobs Reports is consistently higher than for the FOMC meetings. Given the findings of earlier literature that the FOMC meetings have an impact on peoples’ expectations (see e.g., De Fiore et al., 2021; Mertens et al., 2020), individuals are likely sufficiently aware of the Jobs Reports.

Figure 4.1: Search interest



(a) Interest surrounding publication (b) Interest compared to FOMC meeting
 Note: Panel (a) shows the average search interest over all reports. Panel (b) shows the interest in each report over time.

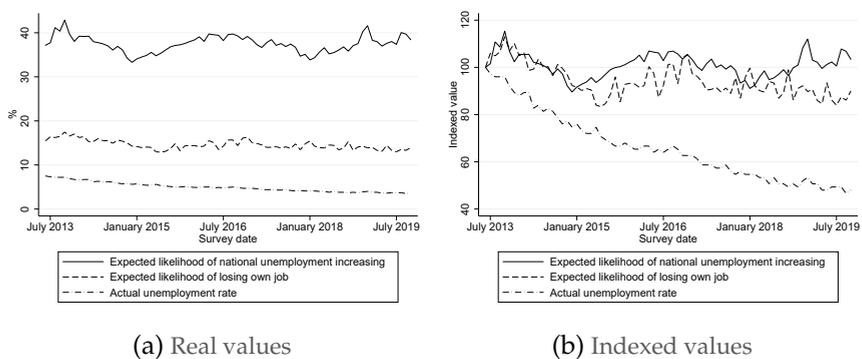
4.3 Data & Methodology

4.3.1 Data

The Survey of Consumer Expectations conducted by the Federal Reserve Bank of New York is a monthly survey with a rotating panel (see Armantier, Topa, Van der Klaauw, and Zafar (2017) for a detailed overview of the survey). We use data from the main module of the survey collected between June of 2013 and October of 2019 and restrict our sample to working individuals.

In the survey, respondents are asked about their expectations of a wide range of macroeconomic variables, such as the inflation rate, interest rate, stock prices, house prices and, of course, unemployment. In addition, people are asked about expectations related to their personal life, such as the likelihood that they will be financially better off in twelve months. Two questions are of particular importance to us. The first relates to the expectations about the development of unemployment in the United States as a whole. The question reads: *“What do you think is the percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now?”*. Individuals can answer on a scale from 0 to 100. Unfortunately, the survey does not include a question in which respondents are asked about their estimates of the national unemployment rate. This makes it impossible for us to control for their level expectations, or compare their prior to the contents of the Jobs Reports. The second question is about their expectations regarding their own job and reads as follows: *“What do you think is the percent chance that you will lose your main/current job during the next 12 months?”*. The scale of this question is the same.

Figure 4.2: Development of expectations and unemployment rate over time



Note: Panel (a) shows the average values for each variable per survey month. In Panel (b), values are indexed by their value in July 2013.

Panels (a) and (b) of Figure 4.2 respectively show the levels and indexed value of the answers to these questions and the actual unemployment rate over time. On average, individuals think that the likelihood that unemployment in the United States will increase in the next 12 months is slightly below 40%. This value hardly changed between 2013 and 2019. This is somewhat surprising, given the steady decrease of unemployment from 7.5 to 3.5% during this period. Individuals estimate that the likelihood that they will lose their job in the next 12 months is approximately 15%. Likewise, we do not find a strong indication that this figure is trending in any direction.

In addition to the questions mentioned above, the survey also contains questions about individuals' expectations of leaving their jobs voluntarily, the likelihood of finding a new job if they were to be displaced, expected changes in earnings conditional on remaining in their jobs and expected changes in spending. Appendix 4.A provides an

overview of the relevant questions. Lastly, the data contain broad information on peoples' age, numeracy skills, education level and household income.

Apart from the survey data, we have information on the exact date on which the Jobs Reports are published and their contents. We retrieve this information from the Bureau of Labor Statistics' website³. In total, 77 reports are included in our analyses.

4.3.2 Methodology

As the Jobs Reports provide individuals with information on the current unemployment rate, whereas the question in the Survey of Consumer Expectations is about the likelihood of it increasing in the next twelve months, it is not clear how individuals should interpret an individual Jobs Report. Analogous to monetary policy shocks identified using high-frequency data, we compare responses shortly before and after report publications to estimate the surprise component of Jobs Reports. We interpret the difference as the effect of the report's publication.

Figure 4.3 provides a visual representation of the treatment assignment if we were to choose a bandwidth of three days around the report. We exclude individuals who respond on the day of the report from the sample, as we do not have information on the exact time of day at which they responded. This means they could have responded either before or after the publication of the Jobs Report. Individuals who responded in the three days before the publication of the report are

³<https://www.bls.gov/bls/news-release/empsit.htm>

assigned to the control group. Individuals who responded in the three days after the publication of the report are assigned to the treatment group. Those who responded outside of this three-day bandwidth are excluded from the sample. In our estimations, we use bandwidths that vary from one to seven days. We refer to the group of individuals who responded within the bandwidth as a 'cohort'.

We analyze the impact of the Jobs Reports by estimating Jobs Report-specific coefficients of the impact on the national unemployment rate expectations. Using the full sample, it would amount to estimating the following Equation.

$$u_{i,t} = \alpha_t + \beta_t \times T_{i,t} + \zeta_i + \varepsilon_{i,t}. \quad (4.1)$$

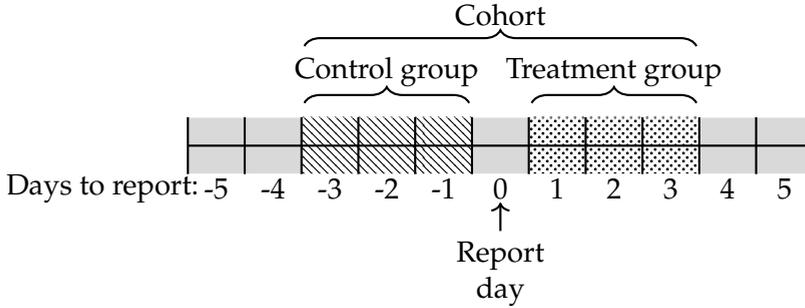
Here, $u_{i,t}$ denotes individual i 's expectations about the likelihood of the national unemployment rate increasing surrounding the Jobs Report published at time t . α_t is a period-specific constant (i.e., Jobs Report fixed effect). $T_{i,t}$ is the treatment indicator, which means β_t denotes the coefficient of the impact the Jobs Report published at time t had on the unemployment rate expectations. ζ_i denote the individual fixed effects and $\varepsilon_{i,t}$ the idiosyncratic error term.

The next step would be to include $\hat{\beta}_t$ in a similar Equation, with personal job loss expectations as the outcome variable:

$$J_{i,t} = \delta_t + \gamma \times \hat{\beta}_t \times T_{i,t} + \nu + \xi_i + \eta_{i,t}. \quad (4.2)$$

Here, $J_{i,t}$ denotes individual i 's belief at time t about the likelihood

Figure 4.3: Visual representation of treatment allocation



of losing their job in the next twelve months. δ_t is a period-specific constant. γ is the coefficient of interest. It indicates the impact of a Jobs Report that changed unemployment rate expectations by $\hat{\beta}_t$ on personal job loss expectations. v denotes the Jobs Report fixed effects, ξ_i denotes the individual fixed effects and $\eta_{i,t}$ the idiosyncratic error term.

Note that we do not use an instrumental variable strategy, but rather use $\hat{\beta}_t$ from Equation 4.1 as a proxy for the way in which the Jobs Report is interpreted. Using an instrumental variable approach would implicitly assume that the mechanism through which the Jobs Reports affect personal job loss expectations only goes through the expectations about the national unemployment rate. While it is likely that this mechanism plays a major role, it may not be the only factor. Specific information about, e.g., sectoral nonfarm payroll employment may also affect personal job loss expectations.

There is a problem with this methodology, however. Because we have

77 reports, the number of individuals in each cohort is limited and individuals thus have a large impact on $\hat{\beta}_t$. If there is a within-individual correlation between expectations about the development of the national unemployment rate and the likelihood of losing their own job caused by something other than the information shock, this would lead to biased results. This seems likely, given the results of prior research on the relationship between individuals' experiences and expectations (see e.g., Cotofan et al., 2020; Geishecker, Riedl, and Frijters 2012; Kuchler and Zafar, 2019 and Malmendier and Nagel, 2011, 2016). The reason that this leads to biased estimates is that any cohort where the treatment group is more optimistic (pessimistic) about the development of the national unemployment rate than the control group will have a negative (positive) value for $\hat{\beta}_t$. Because of the individual-level correlation between unemployment rate expectations and personal job loss expectations, it is likely that the treatment group in such a cohort will also be more optimistic (pessimistic) about the likelihood of personal job loss, creating an artificial positive correlation between our estimates of γ and β_t .

To solve this, we have to break the individual-level correlation. We propose two ways of doing this, each with its distinct advantages and disadvantages. The first strategy involves splitting the sample in two equally sized groups: the '50/50 sample split'. We use the first group to estimate Equation 4.1, and the other to estimate Equation 4.2. We repeat this process 500 times, to obtain estimates and standard errors for all of the variables in the model. The advantage of this methodology is that by using mutually exclusive samples to estimate the Equations, we avoid any bias caused by within-individual correlation between the two outcomes. The core identifying assumption is that the way in

which the first group interprets the Jobs Report is correlated with the way in which the second group does. We expand on this in Section 4.3.3. The downside of this methodology is that our number of observations is effectively cut in half, significantly decreasing statistical power.

The second strategy follows the same intuition: we use others' interpretation of the Jobs Reports as an explanatory variable when we estimate a Jobs Report's impact on personal job loss expectations. In this case, however, we use a 'Jackknife procedure'. For each individual, we estimate the information shock of the Jobs Report on expectations about the national unemployment rate among all other participants in the survey. We then use this 'leave-out estimate' as the treatment intensity variable in the second step. More specifically, we estimate the following two Equations.

$$\mathbf{u}_{-i,t} = \alpha + \beta_{i,t} \times \mathbf{T}_{-i,t} + \zeta_{-i} + \varepsilon_{-i,t}. \quad (4.3)$$

Reusing the notation from Equation 4.1, $\mathbf{u}_{-i,t}$ is a vector denoting the expected likelihood that the national unemployment rate will increase in the next twelve months according to all individuals except for i at time t . α is a constant. We do not use Jobs Report fixed effects. Vector $\mathbf{T}_{-i,t}$ indicates whether an individual was in the treatment group at time t . $\beta_{i,t}$ can be viewed as the Jackknife equivalent of β_t in Equation 4.1. It denotes the impact of the Jobs Report published at time t on unemployment expectations of every individuals except for i . ζ_{-i} denotes the individual fixed effects and $\varepsilon_{-i,t}$ is an idiosyncratic error term.

Taking $\hat{\beta}_{i,t}$ from Equation 4.3, we then estimate the following Equation.

$$J_{i,t} = \delta + \gamma \times \hat{\beta}_{i,t} \times T_{i,t} + \xi_i + \eta_{i,t}. \quad (4.4)$$

Again reusing notation, $J_{i,t}$ denotes individual i 's belief about the likelihood of losing their job in the next twelve months. δ is the constant. $T_{i,t}$ indicates whether individual i was in the treatment group at time t . γ is again the coefficient of interest. It indicates the impact of a Jobs Report that changed unemployment rate expectations by $\hat{\beta}_{i,t}$ on personal job loss expectations. ξ_i denotes the individual fixed effects, and $\eta_{i,t}$ is an idiosyncratic error term.

This strategy does not negatively impact our effective sample size, but it does have a different downside. While the national unemployment rate expectations of the individual do not affect the estimate assigned to them, they do affect the estimate assigned to every other individual. This means that those with the most extreme opinions about the national unemployment rate, get assigned the most conservative coefficients within the cohort. This may lead to a negative bias, especially if within cohort variance is large. Table 4.B.1 in the Appendix provides reassuring evidence that most of the variance is between cohorts, however. The second column of the first row shows that the overall variance of $\hat{\beta}_{i,t}$ is equal to $3.688^2 \approx 13.60$. The second row shows that the average within cohort variance is equal to 0.160; even for the cohort in which the variance is largest, it is only 1.40. This is reassuring, and means the bias is likely small. This potential bias is also the reason we do not use Jobs Report fixed effects in Equations 4.3 and 4.4.

4.3.3 Justification of identifying assumptions

The analysis we conduct relies on three main identifying assumptions:

1. The Jobs Reports actually move expectations.
2. Individuals' interpretations of the Jobs Reports are correlated with each other.
3. Expectations are not moved by anything other than the Jobs Reports.

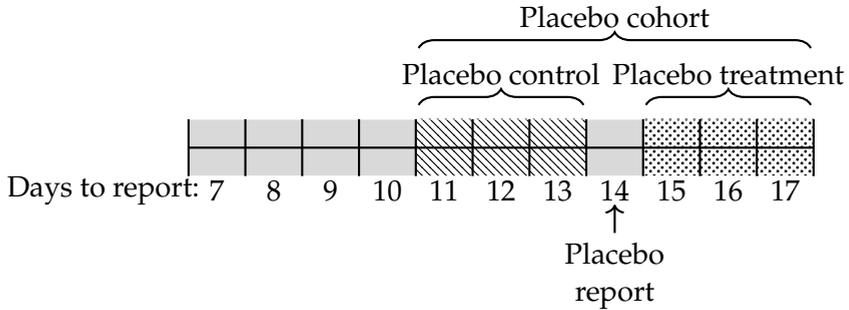
In this Section, we provide support for the validity of these assumptions. We do this based on the 50/50 sample split, as it is most cleanly identified; it does not suffer from the (small) downward bias that the Jackknife procedure does. The subsections below each describe an identifying assumption and provide supporting evidence.

To do so, we need a benchmark. For this, we use placebo treatment analyses. This entails moving the 'treatment' date forward by 14 days and re-estimating our Equations. Given that no reports were published on the days of the placebo treatments, there should not be any effect of the placebo treatment. Figure 4.4 provides a visual representation for the placebo treatment if we were to choose a three-day window around the placebo publication.

4.3.3.1 The Jobs Reports actually move expectations

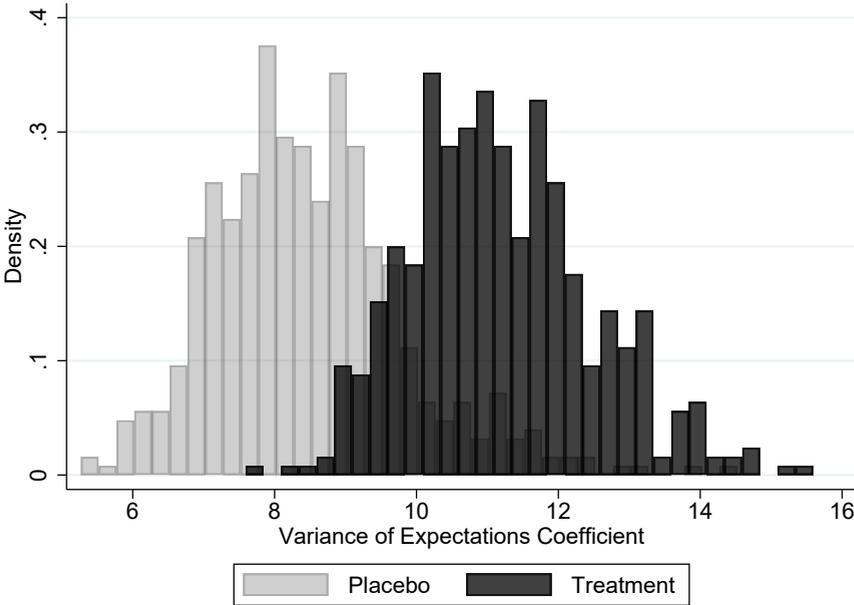
The first identifying assumption is that Jobs Reports actually move expectations. If this is indeed true, we would expect higher variance in expectations on the days surrounding the Jobs Reports than on other

Figure 4.4: Visual representation of placebo treatment allocation



days. We analyze this by studying the variance of the estimated coefficient $\hat{\beta}_t$ from Equation 4.1 for both the true publication of the Jobs Reports and the placebo treatment. Figure 4.5 shows the distribution of the variance of $\hat{\beta}_t$ from 500 replications for both the actual Jobs Report as well as the placebo reports. The Figure shows what we would expect if Jobs Reports indeed affect expectations: the variance of $\hat{\beta}_t$ looks to be much higher for the real treatment than for the placebo treatment. Intuitively, this means that expectations shift more strongly on report dates than on placebo dates. A Kolmogorov-Smirnov test confirms that the distributions are not the same. We thus conclude that the Jobs Reports actually move expectations about the unemployment rate.

Figure 4.5: Variance of $\hat{\beta}_t$ for real Jobs Report and placebo Jobs Reports



Note: light gray bars show density of estimated variance in expectations coefficient for the treatment impact of the placebo Jobs Reports. Dark gray bars show the same for the treatment impact of the actual Jobs Reports.

4.3.3.2 Individuals' interpretations of the Jobs Reports are correlated with each other

Our method further relies on the expectations of other individuals to get a measure of the way in which the Jobs Reports are interpreted. For this to make sense, it requires that peoples' expectations around the Jobs Report actually correlate with each other. To test this, we again employ the 50/50 sample split. We first estimate Equation 4.1, using the first subsample. Next, we estimate Equation 4.2 using the second subsample, but take the unemployment rate expectations as the outcome variable for both Equations instead. If the interpretations are indeed correlated, we would expect to see a positive value for $\hat{\beta}_t$. Table 4.1 shows the results from this exercise, using bandwidths of one to seven days around the publication dates of the Jobs Reports.

Table 4.1: Correlation of news shocks between groups

	$B = 1$	$B = 2$	$B = 3$	$B = 4$	$B = 5$	$B = 6$	$B = 7$
Treated	0.502 (0.671)	0.324 (0.372)	0.444 (0.28)	0.098 (0.236)	0.079 (0.211)	0.072 (0.203)	0.085 (0.177)
Treated $\times \hat{\beta}_t$	0.176* (0.095)	0.14 (0.086)	0.267*** (0.09)	0.34*** (0.08)	0.371*** (0.085)	0.424*** (0.087)	0.418*** (0.08)
Observations	1730	4819	8958	13208	16660	19929	22951
Average obs. per Group	865	2409	4479	6604	8330	9964	11475

Note: *Treated* is a dummy indicating whether the individual answered before (0) or after (1) the Jobs Report. $\hat{\beta}_t$ denotes the coefficient of the impact of the Jobs Report published at time t on national unemployment rate expectations. We include Jobs Report fixed effects in the regression, which means $\hat{\beta}_t$ drops out of the Equation. Results are based on 500 resamplings into the two groups. B indicates the bandwidth in days around the reports' publication dates. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

The estimates are positive for all bandwidths, although insignificant (or only marginally significant) for bandwidths of one and two days around the report. For a bandwidth of three and more days around the reports, the point estimate is between 0.267 and 0.424, and significant at the 1% level in all cases. It shows that interpretation of Jobs Reports are indeed consistent across individuals.

4.3.3.3 Expectations should not be moved by anything other than the Jobs Reports

The last identifying assumption is that the Jobs Reports are the only source of variation in expectations. If this is indeed true, the exercise from Section 4.3.3.2 should show mostly null results for the placebo treatment.

Table 4.2: Correlation of news shocks between groups - Placebo treatment

	$B = 1$	$B = 2$	$B = 3$	$B = 4$	$B = 5$	$B = 6$	$B = 7$
Placebo Treated	1.452 (0.787)	0.765 (0.564)	0.699 (0.52)	0.406 (0.21)	0.286 (0.153)	0.229 (0.171)	0.198 (0.169)
Placebo Treated $\times \hat{\beta}_t$	0.141 (0.117)	0.062 (0.099)	0.066 (0.097)	0.102 (0.093)	0.065 (0.092)	0.147* (0.081)	0.062 (0.083)
Observations	1984	5890	11225	16218	19685	22724	25776
Average obs. per Group	992	2945	5612	8109	9842	11362	12888

Note: *Placebo Treated* is a dummy indicating whether the individual answered before (0) or after (1) the placebo Jobs Report. $\hat{\beta}_t$ denotes the coefficient of the impact of the placebo Jobs Report on national unemployment rate expectations. We include placebo Jobs Report fixed effects in the regression, which means $\hat{\beta}_t$ drops out of the Equation. Results are based on 500 resamplings into the two groups. B indicates the bandwidth in days around the reports' publication dates. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

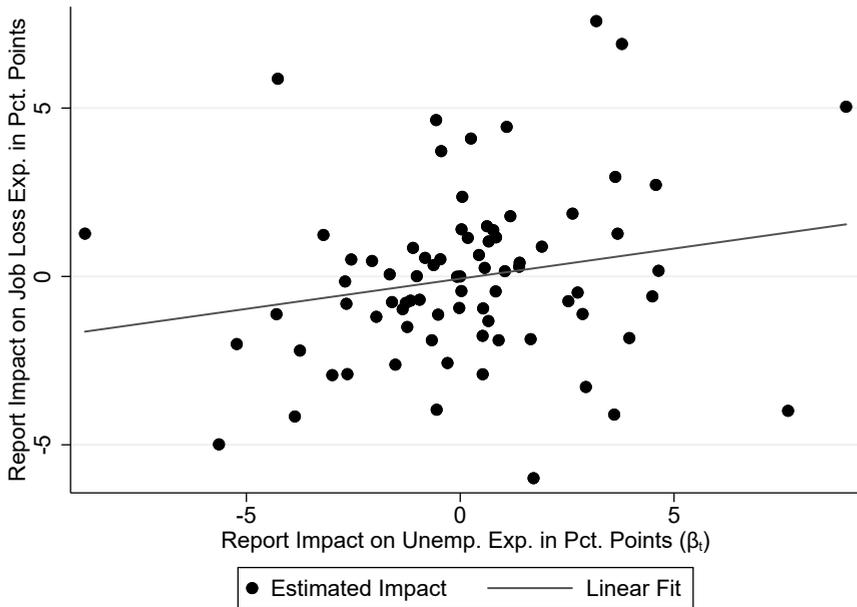
With the exception of Column (6), which is marginally significant, Table 4.2 indeed shows no significant correlation between the expectations surrounding the Jobs Report of the first and the second subsample. The point estimates are also much smaller, for all but the tightest bandwidths. We thus conclude that it is unlikely that any other systematic event close to the publication of the Jobs Reports has a major impact on expectations in the days surrounding it.

4.4 Results

4.4.1 An intuitive first look

Before we turn to our main results, we first conduct an intuitive exercise. Instead of estimating Equation 4.1 for just national unemployment rate expectations, we also do this for personal job loss expectations. This gives us two estimates for each Jobs Report: the impact on national unemployment rate expectations and the impact on personal job loss expectations. Figure 4.6 shows how these impacts correlate with each other. The x-axis shows the estimated impact of each Jobs Report on expectations about the development of the national unemployment rate. The y-axis shows the estimated impact of each Jobs Report on the likelihood of losing one's own job. The Figure shows what we would expect. The higher the impact on national unemployment rate expectations, the higher the impact on personal job loss expectations. In the next Section, we formalize this result.

Figure 4.6: Correlation between report impacts on expectations about national unemployment rate and personal job loss



Note: Figure shows how the estimated impact for both the national unemployment rate expectations and personal job loss expectations of each Jobs Report correlate with each other. Bandwidth is equal to 3.

4.4.2 Main results

Turning to our main results, Tables 4.3 and 4.4 show our estimates from Equation 4.2 (50/50 split) and 4.4 (Jackknife procedure), respectively. The value of $\text{Treated} \times \hat{\beta}$ can be interpreted as the percentage point change in personal job loss expectations if we estimate the impact of the Jobs Report on national unemployment rate expectations to be 1 percentage point. We display the results for bandwidths of one to seven days around the publication of the Jobs Reports. Estimates differ slightly between the two methods, but tell the same story: a Jobs Report that increases individuals' expectations about the likelihood of the national unemployment rate increasing also increases their personal job loss expectations. Looking at the Tables into more detail, Table 4.3 shows no significant results for bandwidths of up to three days. For bandwidths of four days and more around the report, the effects are statistically significant and vary between 0.14 to 0.223 percentage points. For the Jackknife procedure, the results are apparent for all bandwidths, but not significant for the largest bandwidths. Estimates hover between 0.11 and 0.192 percentage points. The fact that for small bandwidths, the results are more convincing using the Jackknife procedure is unsurprising. As stated before, the 50/50 split causes the sample to be effectively halved.

Table 4.3: Effect of cohort-specific news shock on personal job loss expectations - 50/50 split

	$B = 1$	$B = 2$	$B = 3$	$B = 4$	$B = 5$	$B = 6$	$B = 7$
Treated	-0.424 (0.466)	-0.193 (0.272)	0.137 (0.193)	-0.488 (0.152)	-0.04 (0.137)	0.017 (0.111)	0.052 (0.109)
Treated $\times \hat{\beta}_t$	0.02 (0.086)	0.085 (0.086)	0.131 (0.083)	0.169** (0.076)	0.223*** (0.083)	0.165** (0.068)	0.14** (0.07)
Observations	1730	4819	8958	13208	16660	19929	22951
Average obs. per Group	865	2409	4479	6604	8330	9964	11475

Note: *Treated* is a dummy indicating whether the individual answered before (0) or after (1) the Jobs Report. $\hat{\beta}_t$ denotes the coefficient of the impact of the Jobs Report published at time t on national unemployment rate expectations. We include Jobs Report fixed effects in the regression, which means $\hat{\beta}_t$ drops out of the Equation. Results are based on 500 resamplings into the two groups. B indicates the bandwidth in days around the reports' publication dates. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Table 4.4: Effect of cohort-specific news shock on personal job loss expectations - Jackknife

	$B = 1$	$B = 2$	$B = 3$	$B = 4$	$B = 5$	$B = 6$	$B = 7$
Treated	-1.157 (0.793)	-0.265 (0.483)	-0.167 (0.314)	-0.382 (0.278)	-0.191 (0.241)	-0.105 (0.213)	-0.148 (0.209)
$\hat{\beta}_{i,t}$	-0.0321 (0.0478)	-0.150** (0.0623)	-0.0382 (0.0652)	-0.0295 (0.0599)	-0.0413 (0.0585)	-0.0529 (0.0634)	-0.0614 (0.0617)
Treated $\times \hat{\beta}_{i,t}$	0.142** (0.0693)	0.192** (0.0828)	0.147* (0.0818)	0.178** (0.0887)	0.163** (0.0816)	0.110 (0.0838)	0.127 (0.0814)
Observations	1730	4819	8958	13208	16660	19929	22951

Note: *Treated* is a dummy indicating whether the individual answered before (0) or after (1) the Jobs Report. $\hat{\beta}_{i,t}$ denotes the coefficient of the impact of the Jobs Report published at time t on national unemployment rate expectations, when individual i is excluded from the regression. $\hat{\beta}_{i,t}$ is demeaned. B indicates the bandwidth in days around the reports' publication dates. Bootstrapped standard errors (500 replications) clustered at individual level between parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

It is worthwhile to confirm that placebo treatment analyses show no results with personal job loss expectations as the outcome as well. Tables 4.5 and 4.6 show the results of our placebo treatment analyses, again using the 50/50 split and Jackknife procedure, respectively. In contrast to the analysis conducted around publication days of Jobs Reports, the placebo treatment has no effects on personal job loss expectations that are significantly different from zero. This holds for both the 50/50 split and the Jackknife procedure. Additionally, all estimated coefficients are very close to zero. If idiosyncratic shocks that affect both national unemployment rate expectations and personal job loss expectations would frequently occur, the placebo analysis would likely reveal significant effects and the results presented in Tables 4.3 and 4.4 could not be causally linked to the publication of Jobs Reports. The absence of such effects therefore supports the causal interpretation of the results presented in this subsection.

Table 4.5: Effect of Placebo news shock on personal job loss expectations - 50/50 split

	$B = 1$	$B = 2$	$B = 3$	$B = 4$	$B = 5$	$B = 6$	$B = 7$
Placebo Treated	0.563 (0.716)	0.169 (0.541)	0.454 (0.499)	0.244 (0.158)	0.27 (0.123)	0.317 (0.127)	0.347 (0.155)
Placebo Treated $\times \hat{\beta}_t$	0.033 (0.116)	-0.067 (0.102)	0.028 (0.1)	0.064 (0.094)	-0.046 (0.089)	0.043 (0.082)	0.007 (0.086)
Observations	1984	5890	11225	16218	19685	22724	25776
Average obs. per Group	865	2409	4479	6604	8330	9964	11475

Note: *Placebo Treated* is a dummy indicating whether the individual answered before (0) or after (1) the placebo Jobs Report. $\hat{\beta}_t$ denotes the coefficient of the impact of the placebo Jobs Report on national unemployment rate expectations. Includes placebo Jobs Report fixed effects. Results are based on 500 resamplings into the two groups. B indicates the bandwidth in days around the reports' publication dates. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Table 4.6: Effect of Placebo news shock on personal job loss expectations - Jackknife

	$B = 1$	$B = 2$	$B = 3$	$B = 4$	$B = 5$	$B = 6$	$B = 7$
Placebo Treated	0.759 (0.923)	0.778 (0.477)	0.589* (0.327)	0.529** (0.270)	0.543** (0.227)	0.541** (0.224)	0.508** (0.210)
$\hat{\beta}_{i,t}$	-0.0308 (0.0507)	-0.0301 (0.0611)	-0.0713 (0.0639)	-0.0291 (0.0645)	-0.0582 (0.0688)	-0.0535 (0.0636)	-0.0248 (0.0664)
Placebo Treated $\times \hat{\beta}_{i,t}$	0.0753 (0.0723)	0.0173 (0.0861)	0.0198 (0.0832)	-0.0184 (0.0861)	0.0155 (0.0959)	0.0561 (0.0918)	-0.0132 (0.0875)
Observations	1984	5890	11225	16218	19685	22724	25776

Note: *Placebo Treated* is a dummy indicating whether the individual answered before (0) or after (1) the placebo Jobs Report. $\hat{\beta}_{i,t}$ denotes the coefficient of the impact of the placebo Jobs Report on national unemployment rate expectations, when individual i is excluded from the regression. $\hat{\beta}_{i,t}$ is demeaned. Bootstrapped standard errors (500 replications) clustered at individual level between parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

4.4.3 Heterogeneous treatment effects

The Jobs Reports may not have an equal impact on everyone. Differences in the attention individuals pay to macroeconomic news, their ability to interpret news and perceived business cycle sensitivity of job security may all affect the size of the impact. In this Section, we analyze how the impact of the Jobs Reports on personal job loss expectations differs between groups of individuals. For this analysis, we exclusively turn to the Jackknife procedure, as it provides more power and results look to be comparable. Table 4.7 shows how the treatment effect differs for different groups of individuals.⁴ We check for heterogeneous effects by age cohort, numeracy skills, education level and categories

⁴We do not re-estimate the first Equation, so that we can study how the impact of the information shock differs by group.

of household income. We find some evidence that the treatment effect is smaller for individuals aged between 40 and 60 years, and that it is higher for respondents with a college education and a household income of over \$50,000 if we include the characteristics separately. However, if we include all personal characteristics in a single regression, only the effect of age remains (marginally) significant. This analysis only tells us that for a given interpretation of the Jobs Report, individuals aged between 40 and 60 react less strongly than younger and older respondents. A possible mechanism for this result is that individuals in these age categories feel that their job security is less dependent on the general state of the economy than those younger and older. However, the Table does not allow us to explore this, as $\hat{\beta}_{i,t}$ is not age-cohort specific. We therefore do not know if the heterogeneity is driven by differing interpretations of the Jobs Reports or by differences in the translations of these shocks into personal job loss expectations.

Table 4.7: Heterogeneous treatment effects - Jackknife

	Age	Numeracy	Education	Household Inc.	All
Treated	-0.289 (0.487)	-0.418 (0.732)	0.246 (0.874)	-0.103 (0.624)	-0.0527 (1.076)
$\hat{\beta}_{i,t}$	-0.183* (0.110)	0.0688 (0.133)	0.140 (0.145)	0.185 (0.118)	0.0670 (0.222)
Treated $\times \hat{\beta}_{i,t}$	0.311** (0.143)	-0.0384 (0.165)	-0.0481 (0.182)	-0.160 (0.164)	-0.0239 (0.268)
40 to 60 $\times \hat{\beta}_{i,t}$	0.331** (0.141)				0.355** (0.143)
Over 60 $\times \hat{\beta}_{i,t}$	-0.0362 (0.239)				-0.0485 (0.244)
Treated \times 40 to 60 $\times \hat{\beta}_{i,t}$	-0.365** (0.183)				-0.342* (0.190)
Treated \times Over 60 $\times \hat{\beta}_{i,t}$	0.0278 (0.283)				0.102 (0.300)
High numeracy $\times \hat{\beta}_{i,t}$		-0.144 (0.148)			-0.0290 (0.162)
Treated \times High numeracy $\times \hat{\beta}_{i,t}$		0.252 (0.195)			0.0788 (0.218)
College $\times \hat{\beta}_{i,t}$			-0.239 (0.171)		0.00119 (0.198)
Some College $\times \hat{\beta}_{i,t}$			-0.113 (0.178)		-0.0187 (0.195)
Treated \times College $\times \hat{\beta}_{i,t}$			0.362* (0.218)		0.105 (0.258)
Treated \times Some College $\times \hat{\beta}_{i,t}$			-0.0787 (0.241)		-0.182 (0.244)
50k to 100k $\times \hat{\beta}_{i,t}$				-0.254* (0.148)	-0.264 (0.173)
Over 100k $\times \hat{\beta}_{i,t}$				-0.382** (0.173)	-0.404** (0.202)
Treated \times 50k to 100k $\times \hat{\beta}_{i,t}$				0.356* (0.197)	0.298 (0.228)
Treated \times Over 100k $\times \hat{\beta}_{i,t}$				0.545** (0.240)	0.432 (0.305)
Observations	8956	8954	8956	8886	8880
χ^2 Age	4.878				4.758
P-value χ^2 Age	0.087				0.093
χ^2 Edu			6.378		1.799
P-value χ^2 Edu			0.041		0.407
χ^2 HHI				5.585	2.267
P-value χ^2 HHI				0.061	0.322

Note: *Treated* is a dummy indicating whether the individual answered before (0) or after (1) the Jobs Report. Bandwidth is equal to 3. $\hat{\beta}_{i,t}$ denotes the coefficient of the impact of the Jobs Report published at time t on national unemployment rate expectations, when individual i is excluded from the regression. $\hat{\beta}_{i,t}$ is demeaned. Bootstrapped standard errors (500 replications) clustered at individual level between parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. χ^2 tests and P-values are for tests of joint significance for triple interactions of age, education and household income with 'Treated' and $\hat{\beta}_{i,t}$.

4.4.4 Alternative outcomes

Apart from different impacts across individuals, shocks to expectations about the national unemployment rate may not only affect individuals' expectations about losing their own job, but could impact a number of other expectations about their personal life as well. We study how it impacts individuals' expectations about how easy it will be to find a new job if they were to lose their current one, how likely they think it is they will voluntarily leave their job and how their earnings and spending will change. We again use the Jackknife procedure for this analysis.

The Column (1) of Table 4.8 shows the impact on individuals' expected likelihood of finding a new job within three months if they were to lose theirs now. In line with our prior finding, it decreases. Somewhat surprisingly, we find a positive effect on people's expectation about leaving their own jobs voluntarily in Column (2). One potential explanation is that individuals start spending more time looking for other jobs when they expect it to be more likely that they will lose their own job, potentially leading to a voluntary exit. Table 4.B.2 in the Appendix provides some support for this hypothesis. The more worried an individual is about losing their job, the more likely they are to spend time searching for different jobs.

Columns (3) and (4) of Table 4.8 show the effect on expected earnings. An expected increase in the unemployment rate should worsen individuals' bargaining position, potentially driving their wages down. We find no evidence that people expect their earnings in their current job to decrease, however. We also do not find any evidence of changes in planned spending either, as shown in columns (5) and (6).

Table 4.8: Alternative outcomes - Jackknife

	(1)	(2)	(3)	(4)	(5)	(6)
	Find new job	Leave voluntarily	Earnings increase (dummy)	Pct. earnings increase	Spending increase (dummy)	Pct. spending increase
Treated	-0.183 (0.441)	-0.208 (0.400)	0.00870* (0.00464)	0.231 (0.177)	0.00322 (0.00658)	-0.103 (0.245)
$\hat{\beta}_{i,t}$	0.139 (0.0918)	-0.0597 (0.0871)	-0.000560 (0.000901)	0.0314 (0.0367)	-0.000285 (0.00140)	-0.0596 (0.0474)
Treated \times $\hat{\beta}_{i,t}$	-0.214* (0.123)	0.219* (0.114)	-0.000795 (0.00123)	-0.0327 (0.0557)	-0.000927 (0.00190)	0.0685 (0.0600)
Observations	8954	8956	8956	8950	8955	8953

Note: *Treated* is a dummy indicating whether the individual answered before (0) or after (1) the Jobs Report. Bandwidth is equal to 3. $\hat{\beta}_{i,t}$ denotes the coefficient of the impact of the Jobs Report published at time t on national unemployment rate expectations, when individual i is excluded from the regression. Columns indicate outcome variable. $\hat{\beta}_{i,t}$ is demeaned. Bootstrapped standard errors (500 replications) clustered at individual level between parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. χ^2 tests and P-values are for tests of joint significance for triple interactions of age, education and household income with 'Treated' and $\hat{\beta}_{i,t}$.

4.5 Conclusion

The results from this chapter show that individuals indeed acquire information about macroeconomic conditions in their day-to-day lives and relate this to their personal situation. It not only affects individuals' expected likelihood of losing their own jobs, but their expectations about the likelihood of being able to find a new job conditional on losing theirs as well. News that people interpret as increasing the likelihood of the unemployment rate increasing thus makes individuals more pessimistic about their employment prospects through multiple channels: an increase in the expected likelihood of job loss, and a decrease in the expected likelihood of being able to find employment.

Our finding that individuals between 40 and 60 years of age are somewhat more sensitive to the information shocks requires further research.

One explanation is that these individuals feel their jobs are secure, even if the economy takes a turn for the worse. However, the analyses in this chapter do not allow us to answer this question.

Appendix 4.A Additional Survey Questions

Leave job voluntarily: *What do you think is the percent chance that you will leave your main/current job voluntarily during the next 12 months?*

Find new job: *Suppose you were to lose your (main) job this month. What do you think is the percent chance that within the following 3 months, you will find a job that you will accept, considering the pay and type of work?*

Earnings increase (dummy): *Please think ahead to 12 months from now. Suppose that you are working in the exact same job at the same place you currently work, and working the exact same number of hours. What do you expect to have happened to your earnings on this job, before taxes and deductions?*

Twelve months from now, I expect my earnings to have...

- *increased by 0% or more*
- *decreased by 0% or more*

Pct. earnings increase: *By about what percent do you expect your earnings to have [increased/decreased as in previous question]? Please give your best guess.*

Twelve months from now, I expect my earnings to have [increased/decreased] by -- %

Spending increase (dummy): *Now think about your total household spending, including groceries, clothing, personal care, housing (such as rent, mortgage payments, utilities, maintenance, home improvements), medical expenses (including health insurance), transportation, recreation and entertainment, education, and any large items (such as home appliances, electronics, furniture, or car payments). Over the next 12 months, what do you expect will happen to the total spending of all members of your household (including you)?*

Over the next 12 months, I expect my total household spending to...

- *increase by 0% or more*
- *decrease by 0% or more*

Pct. spending increase: *By about what percent do you expect your total household spending to [increase/decrease as in previous question]? Please give your best guess.*

Over the next 12 months, I expect my total household spending to [increase/decrease] by __%

Appendix 4.B Additional Tables

Table 4.B.1: Summary statistics of $\hat{\beta}_{i,t}$

	Mean	Std. Dev.	Median	Min	Max
$\hat{\beta}_{i,t}$	-.1780225	3.687733	-.0962545	-18.19821	17.56856
Within cohort variance of $\hat{\beta}_{i,t}$.160359	.1738912	.1020288	0	1.403147
Maximum value within cohort of $\hat{\beta}_{i,t}$	1.182855	3.687081	1.229935	-7.227622	17.56856
Minimum value within cohort of $\hat{\beta}_{i,t}$	-1.534244	3.663398	-1.421255	-18.19821	13.87807
Difference between maximum and minimum value of $\hat{\beta}_{i,t}$ w/i cohort	2.717099	1.496279	2.365716	0	13.39423
Observations	8958				

Note: first row displays general descriptive statistics of $\hat{\beta}_{i,t}$. The rows below show how $\hat{\beta}_{i,t}$ is distributed between cohorts.

Table 4.B.2: Expectations and Search

	Searched for Work	Hours Spent Searching
Expected Likelihood Losing Job	0.00360*** (0.000462)	0.0197*** (0.00392)
Constant	0.182*** (0.00636)	0.625*** (0.0540)
Observations	7614	7616

Note: results from a Fixed Effects regression. Standard errors clustered at individual level between parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

5

Stimulating Occupational Mobility among Unemployed Job Seekers

Abstract

We study the impact of online information provision to job seekers who are looking for work in occupations with relatively poor labor market prospects. The information is provided through a personalized email that contains suggestions about suitable alternative occupations and information about how the prospects of these alternatives compare to the job seekers' current occupation of interest. A second treatment adds a motivational video aimed at addressing the psychological hurdles of switching to a different occupation. We evaluate the interventions using a randomized field experiment with 30,129 unemployed job seekers, and we acquire additional descriptive information on beliefs and job search. We find no impact on received benefits and earnings in the first eight months after the treatment. The findings do show that treated individuals are 1.79 percentage points more likely to have found a job seven months after the intervention, although this difference decreases to 1.19 percentage points four months later. Moreover, treated individuals are between 5 and 6 percentage points more likely to have done so in an occupation different from their initial occupation of interest. This may be promising for their longer-term prospects.

This chapter is based on joint work with Michèle Belot, Didier Fouarge, Philipp Kircher, Paul Muller and Sandra Phlippen. We are grateful to the many colleagues at the Public Employment Office in the Netherlands (UWV) for the collaboration that enabled this project, with specific thanks to Yvonne Engels and Mario Keer for their key contributions and Peter Berkhout for his help in preparing the various datasets. We gratefully acknowledge the generous support through ESRC grant ES/L009633/1 and the financial support from ROA (Research Centre for Education and the Labour Market; Maastricht University). This study is registered in the AEA RCT Registry and the digital object identifier (DOI) is: 10.1257/rct.7374-1.0.

5.1 Introduction

Occupational transitions play a significant role in labor market adjustments to changes in the economy. The Covid-19 pandemic (del Rio-Chanona, Mealy, Pichler, Lafond, & Farmer, 2020; Forsythe, Kahn, Lange, & Wiczer, 2020), technological development (Autor, Levy, & Murnane, 2003), and automation (Autor, 2015; Brynjolfsson, Mitchell, & Rock, 2018; Frey & Osborne, 2017) have been associated with profound changes in the demand for certain occupations. Adjusting to such a changing environment means that workers need to transit from occupations in decline to occupations with better prospects. A major challenge is that workers may not be well informed about occupations they could or should consider. Moreover, even if well informed, there may also be psychological barriers to consider a career switch. The lack of familiarity and uncertainty about the fit with other occupations may constitute significant hurdles to occupational transitions. Evidence indeed suggests that when searching for jobs, individuals tend to narrowly focus on occupations in which they have experience (Belot, Kircher, and Muller, 2019; Faberman and Kudlyak, 2019).

In this chapter, we design and evaluate two low-cost digital interventions aimed at job seekers who are looking for work in occupations that are in low demand. The first intervention aims at addressing informational deficits. The second aims at tackling psychological barriers to occupational transitions. We conduct these interventions in collaboration with the Public Employment Office in the Netherlands (UWV).

The experiment involves 30,129 job seekers who recently became unemployed and search in one of 21 occupations that the employment office identified to have poor employment prospects. We send an email

to 20,125 of these job seekers, in which we inform them of the poor job prospects in their primary occupation of interest and suggest alternative occupations with better prospects that are particularly well-suited to their background. Each suggested occupation includes information about the job finding prospects, the skills required to do well in the occupation, and a link to a webpage with more detailed information about the occupation. In addition, the email contains a link to an online job search engine that job seekers can use to find relevant vacancies. The occupational suggestions are based on the most common occupational transitions observed from (i) millions of resumes from former job seekers and (ii) a longitudinal survey that is representative of the Dutch labor force. This ensures that the occupational suggestions are realistic switches for the targeted job seekers. From these common and attainable transitions, we include those that currently offer sufficiently good job finding prospects.

In the second treatment, we add a motivational component to our intervention. We sourced videos from a diverse group of individuals who made a successful transition from one occupation to another. In cooperation with a professional video maker, we compiled their stories into a motivational video that addresses the main challenges, costs and benefits of occupational transitions. Half of the email recipients in the information treatment also receive a link to the motivational video in the email.

We measure the impact of the interventions on benefits receipts, earnings and job finding probability using administrative data.¹ On top of

¹Further administrative data (either from online search records, or from caseworkers records) on job search activities is expected to become available for analysis in the near future.

that, we assess how the intervention impacts job search activities, labor market beliefs and the occupation in which job seekers ultimately find work using survey data collected before and after the interventions.

Preliminary data analyses show that job seekers are, on average, willing to look for alternative occupations and are confident that they will be able to do well in these occupations. However, job seekers are generally not aware of how bad job prospects are in their primary occupation of interest, compared to suitable alternatives. While most job seekers do consider one (or a couple of) alternative occupation(s), their assessment of the job finding chances in these alternatives is also hardly correlated with true job finding prospects. These findings point towards fertile grounds for our intervention.

Take-up of the information emails is high: we find that almost 80% of the treatment group opened at least one of the two informational emails. A sizeable share also clicked on at least one occupational suggestion for more information. The motivational video, on the other hand, did not attract interest. After an explicit reminder, still fewer than 10% of the recipients watched the video.

Despite the high engagement with the information intervention, we do not find a significant impact of either treatment on benefit receipt or labor earnings over the first eight months after the intervention.² However, we do find that treated individuals seem to find a job sooner, and when they do so, are more likely to end up in a different occupation than their primary (pre-intervention) occupation of interest. Our survey data cannot confirm that beliefs or search behaviors are affected by the treatments, but our sample size is small.

²Future versions of this chapter (as a research paper) will include up to eighteen months of post-experimental measurements.

Our first treatment contributes to the literature on advice and counseling. So far, little is known about this topic, as these policies are often combined with other policies such as monitoring and sanctions, making it difficult to disentangle underlying mechanisms (see Card, Kluge, and Weber (2018) for a recent review). Earlier literature has shown that subjective expectations about job finding prospects determine individuals' search efforts (Caliendo, Cobb-Clark, & Uhlendorff, 2015). However, these expectations are not always in line with reality, which may explain why individuals tend to spend too much time looking for work in low-demand occupations. Individuals partly form their beliefs through lived experiences (see e.g., Jäger, Roth, Roussille, and Schoefer (2022), who show that individuals strongly anchor their wage beliefs to current earnings), but their beliefs are also affected by information they acquire in their day-to-day lives, as shown in Chapter 4 of this thesis.

Information acquisition is an endogenous process (Wiederholt et al., 2010). Individuals only acquire information when they deem it to be worthwhile. In the case of labor market prospects, individuals will want to acquire more information if they have high macroeconomic risk exposure (Roth, Settele, & Wohlfart, 2022). This leads to a Catch-22 situation. Individuals who are unaware of the poor labor market prospects in their occupation of interest see no reason to search for information about more promising alternatives, and as such remain uninformed. Job seekers indeed tend to overestimate their employment prospects (Spinnewijn, 2015), and do not properly adjust beliefs as the unemployment spell lasts (A. Mueller et al., 2021). Directly providing these individuals with information can be effective, however; individuals do update their beliefs and behavior based on relevant in-

formation they receive (Roth & Wohlfart, 2020).

Related to this chapter, Altmann, Falk, Jäger, and Zimmermann (2018) evaluate the effects of a generic information intervention in Germany. They sent a brochure aimed at providing generic information about beneficial job search activities and motivate job seekers to exert more search effort early on in their unemployment spell. They find that this intervention is effective for job seekers at risk of long-term unemployment. Our study builds on previous work by Belot et al. (2019), who test information interventions on a small sample of job seekers in the UK. They observed job seekers' search behavior over the course of 3 months and find that personalized suggestions of alternative occupations affects job search and increases the chances of getting an interview. Our study is of a much larger scale, focuses on job seekers who search in occupations with poor prospects, and aims at an evaluation of the effects on the chances of finding employment. In addition, we collect detailed information on beliefs regarding employment prospects, allowing us to investigate the mechanism underlying the impact of providing labor market information.

Our second (motivational) intervention draws on the literature on social norms and role models. The social norm to work is a strong motivator to find employment (Kondo & Shoji, 2019). Role models may convey such social norms, as well as motivate individuals and display that certain career paths are possible. Earlier studies have shown that role models can be very effective in shaping individuals' education choices (e.g., Del Carpio and Guadalupe, 2021; Porter and Serra, 2020; Riley, 2022). Our setting is unique in that it combines factual information (targeting individuals' beliefs about the labor market) with a more

intangible part focusing on more personal aspects through role models (targeting individuals' beliefs about their own ability and chances of finding different employment).

From a policy perspective, getting unemployed job seekers back into employment is an important policy objective and the specific objective of employment agencies. For job seekers transitioning out of occupations with poor labor market prospects, finding work can be particularly challenging, as it may require them to consider alternative occupations for which prospects and the match with own skills are not easy to identify. Our study contributes to this policy challenge by evaluating the extent to which broadening the search behavior can help job seekers out of unemployment.

The rest of the chapter is structured as follows. In Section 5.2, we describe the context of our experiment and its design. We provide descriptive results regarding job search behavior of our sample (based on a pre-intervention survey) in Section 5.3. In Section 5.4, we present our empirical evaluation of the impact of the intervention using both administrative and survey data. Section 5.5 concludes.

5.2 Experimental Design

5.2.1 Institutional context

We evaluate the effectiveness of the information and motivation treatments through a large-scale randomized controlled trial in collaboration with the Public Employment Office in the Netherlands (UWV).

One of the employment office's core responsibilities is the administration and payment of employee insurances, including unemployment benefits. In the Netherlands, individuals can apply for unemployment benefits if they meet the following criteria: they are insured for unemployment, their hours of work are decreased by more than five hours per week, they are available to start a different job immediately, they have worked at least 26 out of the last 36 weeks, and their transition to unemployment was not their own fault (Uitvoeringsinstituut Werknemersverzekeringen (UWV), 2014). If individuals meet all of these criteria, they can register for unemployment benefits. As part of the registration process, the employment office asks individuals to complete an online CV that contains up to three 'search occupations'; occupations that the individual would like to find employment in.

Another core task of the employment office is to assist job seekers in finding employment, particularly those with a large distance to the labor market. To this end, the employment office provides a number of services. While job seekers do get assigned to a caseworker, the employment office also states that they "are calling on Dutch citizens to assume their own responsibility and on their self-reliance; the services we provide will increasingly be based on online self-service" (Uitvoeringsinstituut Werknemersverzekeringen (UWV), 2015). An important part of these 'online services' is the employment office's provision of labor market information. Using data on the number of registered job seekers with a certain 'search occupation', as well as the number of available vacancies, the employment office assesses the job opportunities in different occupations. They publish their findings on one of their subsidiary websites: *werk.nl*. While job seekers can find a lot of labor market information there, it is scattered throughout the

website. Moreover, the website is not personalized, meaning job seekers have to be well aware of their wants and needs to find relevant information. In our experiment, we attempt to (i) consolidate the available – and add new – labor market information about occupations, and (ii) provide this information in a personalized manner through email.

5.2.2 Sample selection

The aim of our experiment is to help unemployed job seekers who search in occupations with low employment prospects to consider different, more promising, occupations. The first step in constructing the sample of job seekers is to select the occupations that offer poor job prospects. Job seekers who search in these occupations are most likely to benefit from information on alternative occupations with better prospects. For this, we use the job finding score. The job finding score is a metric used by the employment office based on the ratio of vacancies to job seekers in the employment office's database and outflow rates of unemployment insurance recipients that is updated several times per year. These scores are computed for over 600 narrowly defined occupations (5-digit classification).³ The score runs from 2 (very poor job prospects) to 10 (excellent job prospects). We select for the experiment individuals interested in occupations with a score of 2, 3 or 4 in the spring of 2021, leading to 21 'selection occupations'.⁴ These

³The occupational classification used is called 'BRC+' which resembles the ISCO classification, but more detailed and slightly modified to better reflect the Dutch labor market.

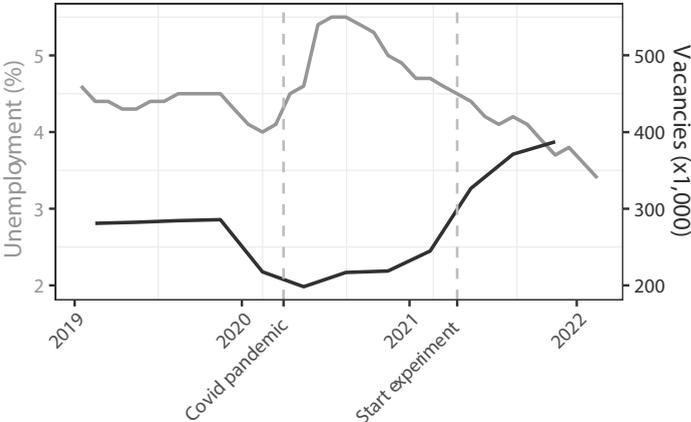
⁴Note that there were a couple of small amendments in the selection of these 21 occupations after consultation with labor market experts of the employment office. In particular, some occupations with a very small number of job seekers were removed,

21 occupations exhibit a substantial variety: they include low-skilled occupations such as waiters/bartenders, janitors and taxi drivers, but also skilled professions such as graphic designers, event organizers and social workers. The complete list can be found in Table 5.2 (including their relative share within the sample). Appendix Table 5.A.1 provides the original occupation names in Dutch.

Table 5.2 makes a distinction between ‘Covid occupations’ and ‘non-Covid occupations’). Due to the Covid pandemic, the state of the labor market fluctuated substantially, as illustrated by the fluctuations in unemployment and vacancy rates depicted in Figure 5.1. Until early 2020, unemployment was low and stable at around 4.5%, while it increased to 5.5% in the summer of 2020 and steadily decreased from there. Vacancies mirror this trend. Despite our selection occupations sharing low prospects in early 2021, they differ substantially in the longer-run trends. Most importantly, there is large variation in the degree to which occupations were affected by the varying social distancing measures that were imposed to minimize the number of Covid cases. We can in fact identify a subset of our selection occupations that offered poor prospects primarily because of the Covid measures, but offered substantially better prospects prior to the Covid pandemic *and* after many restrictions were lifted over the summer of 2021. We classify all selection occupations as ‘Covid-occupations’ if the job finding score decreased with at least two points at the onset of the Covid pandemic *and* increased at least two points in the summer of 2021. There are 7 ‘Covid occupations’. In Figure 5.2, we show how the job finding score evolves for the two groups. As expected, the Covid occupations (right panel) offer decent prospects before the pandemic and almost fully re-

and some occupations that were very similar were combined.

Figure 5.1: Unemployment and vacancies in the Netherlands



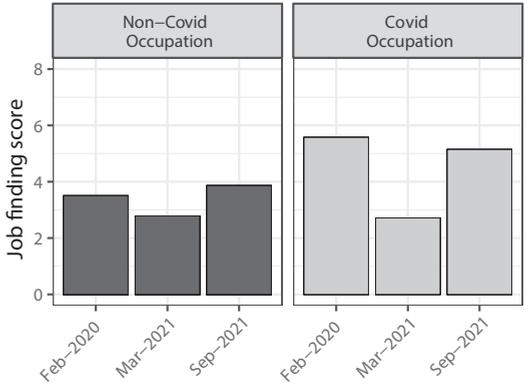
Note: left y-axis (light gray line) displays unemployment rate in the Netherlands. Right y-axis (dark gray line) displays the number of vacancies. Source: StatLine. StatLine is the electronic database of Statistics Netherlands.

cover in late 2021. For the non-Covid occupations (left panel) this is not the case, and job prospects have been structurally bad during the past years.

This distinction is essential for our analyses. Many job seekers may have anticipated that the Covid restrictions were temporary and these individuals may therefore have been less willing to consider switching occupations. Especially since the process of transiting to a new occupation may well take several months, which is precisely the time horizon over which the labor market prospects would be expected to improve. Providing an intervention to encourage occupational switches is less likely to be effective for this group.

We have access to all registered job seekers' records in the Netherlands and select all who have indicated on their CV that they are looking for a job in one of the 21 occupations with a very low job finding score. This implies that we also restrict our sample to job seekers who have completed their online CV, which automatically ensures a minimum level of computer skills. Given that we send our labor market information by email, this was considered desirable as we exclude those who may be less likely to read emails. Finally, we impose the restriction that, at the time of sample selection, job seekers should have at least one month of unemployment insurance benefits eligibility left, to ensure they would not automatically exit the sample before receiving the first intervention email.

Figure 5.2: Job prospects Covid and non-Covid occupations



Note: left-hand Graph shows the (unweighted) average job finding score for the ‘non-Covid’ occupations. Right-hand Graph shows the (unweighted) average job finding score for the ‘Covid’ occupations.

5.2.3 Interventions

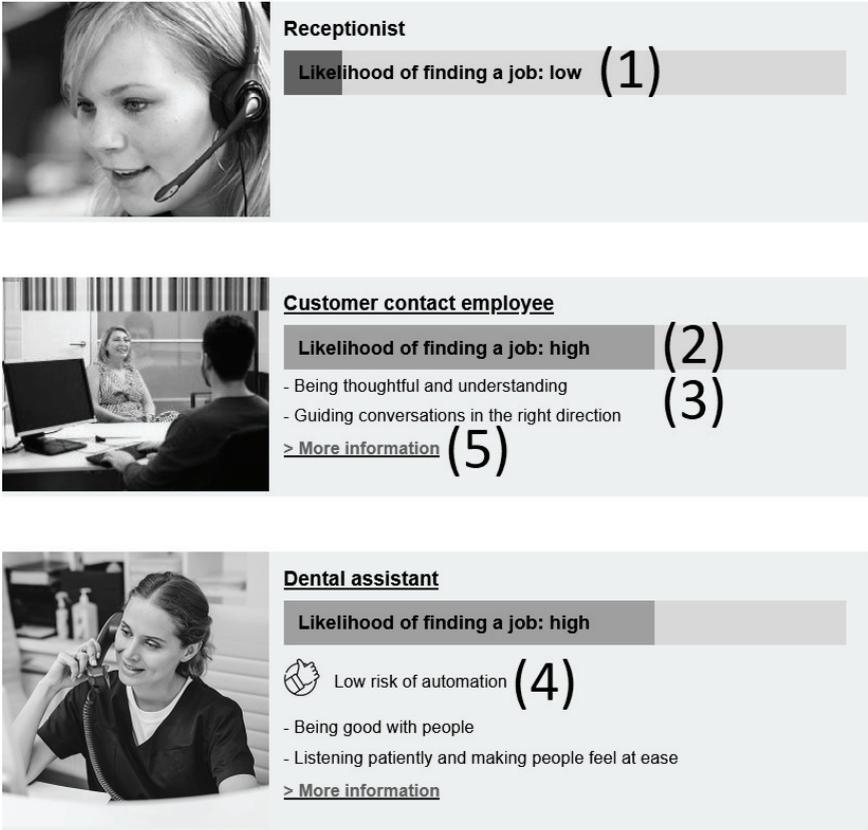
5.2.3.1 Information treatment

Our first treatment objective is to ensure job seekers (i) are aware of the poor labor market prospects in their occupation of interest and (ii) learn about suitable alternative occupations. We determine these suitable alternative occupations based on two metrics. First, we use historical occupational switches based on resume data that the employment office collects for all registered job seekers. This allows us to identify the occupations that other job seekers with skills, experience and educational backgrounds similar to the job seekers in our sample most often switch to. There is one caveat, which is that the resulting list of occupation is based on historical data only. While a high rate of switches is a clear indication that the skill requirements in the suggested occu-

pation are such that transitions are possible and that it had good labor market prospects at some point in time, it does not guarantee that prospects are still good at the time of the intervention. Therefore, as a second step, we select occupations with a high job finding score (see Section 5.2.2 for details). We only include occupations in our list of suitable alternatives if they have a job finding score of at least 6. The combination of these two criteria ensures that we send job seekers a list of occupations that (i) they are likely to be (or can easily become) qualified for and (ii) have good job finding prospects. Depending on a job seeker's preferred occupation, we selected 7 to 9 alternative occupations. While we generally chose the occupations with the largest number of historical switches of those that had good enough job opportunities, the staff at the employment office made some minor changes to the included occupations and the order in which we display them.

We present the information through an information visualization that we send to job seekers by email. In the email's introductory text, we stress a number of key points. First, we provide information about market tightness in the main occupation of interest. Specifically, we inform job seekers that the occupation in which they are currently looking for work has few vacancies available, but that a lot of people are looking for work in that occupation. This implies bad prospects of finding employment. Second, we mention that with their skills and experience, there are alternative occupations they would qualify for (or could relatively easily qualify for) that provide much better job prospects. In this way, we aim to inform job seekers about the urgency of considering alternatives, as well as reassure them that their skills and experience will fit in the new occupation.

Figure 5.3: Example of information email visualization



Note: example visualization for occupation of receptionist. Numbers indicate different aspects of the email, as described in the text.

Figure 5.3 shows an example of the visualization we use. We first list job seekers' primary occupation of interest, together with a bar of which the length and color represent the likelihood of finding a job (1).⁵

⁵For the length, we divided the full length of the bar (in grey) up into tenths. Depend-

Next, we show each of the alternative occupations that we matched to the job seeker's primary occupation of interest. The order in which we show these alternative occupations is largely based on the number of historical transitions we observed and, to a lesser degree, on the job opportunities associated with the alternative occupation. For each of the alternative occupations, we first show the job finding score in the same way as we did for their occupation of interest (2). Next, we show the two main skills associated with the occupation (3). While the use of historical switches between occupations ensures that all presented suggestions are relevant, individuals may have idiosyncratic skills that fit well with one occupation in particular. We want to ensure that job seekers realize that their existing skills and experience can be valuable in another occupation. Many of the occupations with poor prospects we select are at risk of being automated. The set of alternative occupations we propose to them have much better short-term job prospects (job finding score of at least 6). However, the longer-term prospects of these occupations vary. As job seekers may want to avoid occupations with poor long term prospects due to automation risks, we include this information in the treatment as well. If an occupation is at low risk of automation (25th percentile of automation risk or lower), we mention this to the job seeker (4).⁶ Lastly, there is a link for more information about the occupation (extended description, required certifications, various job titles, etc.) (5).

ing on the occupation's job finding score, it fills up the corresponding share of the bar. For the colors, we use the following categorization: job finding scores 2 to 4 are red, job finding score 5 and are yellow, and job finding scores 7 to 10 are green.

⁶The automation risk is measured with the indicator proposed by Nedelkoska and Quintini (2018).

5.2.3.2 Motivation treatment

The second intervention targets psychological barriers to consider an occupational transition. A professional short film video was assembled, with former job seekers sharing their personal occupational transition success stories. The aim of this video is to provide job seekers with relatable stories about motivational challenges associated with occupational transitions and how to overcome them. While job seekers might find our alternative occupational suggestions interesting, they may still wonder if they would really be able to make the switch. Listening to the personal stories of others who have experienced such occupational transitions may be a source of motivation, as evidenced by the role models literature discussed in the introduction. We recruited role models through an op-ed in a Dutch national newspaper. In this op-ed, we explained that a lot of people find occupational transitions to be difficult and perhaps even scary, and that individuals considering such a transition may benefit from learning about the experience of others. We asked individuals to submit a short, personal video. We selected nine recordings and asked a professional video maker to compile these clips into a 5-minute video. The video covers three main topics. First, the individuals introduce themselves and describe the transition they made (occupation they had before and new occupation). Second, they talk about how they experienced the transition. Third, they provide general advice and encouragement.

5.2.4 Randomization, data collection and timeline

We selected the sample on March 15, 2021, and ended up with 30,129 individuals who remained unemployed until the first email (April 12,

2021). These individuals constitute our experimental sample. Job seekers were randomly assigned to three equally sized groups: (1) the information group, (2) the information + motivation group and (3) the control group. Randomization was stratified by gender, unemployment duration and selection occupation. A random third was selected to receive pre- and post-intervention surveys (equally-sized across treatment groups). After selecting the baseline sample, we administered the pre-intervention survey followed by the intervention emails and the post-intervention survey. Subsequently, we sent out ‘outflow surveys’ to those who found jobs. Table 5.1 provides a precise timeline with corresponding sample sizes.

The pre- and post-survey contained questions about job search behavior (primary search occupations, alternative search occupation, applications and interviews), beliefs (job findings prospects in the primary and alternative occupations, beliefs about wages) and willingness to explore and search for occupations other than the primary occupation of interest. Further details can be found in Section 5.3 where we present descriptive statistics.

We sent the first intervention email on April 12. It contained the information visualization for both treatment groups and the video link for the motivational treatment group. In Section 5.4.1, we provide statistics on the engagement with the email. We find that a substantial share opened the email, but few clicked on the link to the video. As a result, we sent an extra email with only the video link to the corresponding treatment group on May 10. Finally, a general reminder email was sent containing a modified version of the information visualization on May 28. The modification was based on clicking statistics from the first email, the details of which can also be found in Section 5.4.1.

Table 5.1: Timeline experimental set-up and sample sizes

Date	Event	Treatment 1 (Information) N = 10,050	Treatment 2 (Info + video) N = 10,075	Control N = 10,004	Total N = 30,129
March 23	Pre-survey sent	3308	3310	3292	9910
	Respondents	899	959	931	2789
April 12	First email	<i>Information</i> 10,050	<i>Info + video</i> 10,075	<i>No email</i> 10,004	30,129
May 10	Second email		<i>Only video</i> 9022		
May 28	Third email	<i>Information</i> 8388	<i>Information</i> 8450	<i>No email</i> 8399	25,237
June 7	Post-survey	2766	2781	2752	8299
	Respondents	400	457	421	1278
June 24	Outflow survey 1	1848	1813	1799	5460
	Respondents	612	550	588	1750
Sept 9	Outflow survey 2	1427	1402	1411	4240
	Respondents	473	491	439	1403
Dec 1	Outflow survey 3	1057	1037	1004	3098
	Respondents	377	353	327	1057
April 5 (2022)	Outflow survey 4	411	402	443	1256
	Respondents	106	107	136	349

Note: all dates are in 2021, unless otherwise noted. Minor sample selection steps were applied prior to each intervention email: only those who (1) did not yet exit unemployment insurance, (2) had valid email addresses and (3) did not change their ‘unemployment-indication’ were included. Prior to the post-survey an additional subset was removed that either denied the consent statement in the pre-survey or that clicked the ‘unsubscribe’ button in the pre-survey. Each survey was followed by an email reminder after one week.

The administrative data that we use contains start of employment spells, earnings from employment and benefits receipts. However, it does not contain information about the occupation people work in. To collect more information on the occupations the unemployed exit to, we also administered outflow surveys. Every two to three months, we selected all job seekers in our sample for whom we observed in the administrative data a labor earnings increase of more than €300.- in the preceding months. For example, for the first outflow survey (in June) we selected recipients for whom monthly earnings in April and/or May were at least 300 euro higher than their highest monthly earnings in February and March. Such a substantial increase in earnings should reflect a new job. Since many job seekers hold part-time and temporary jobs during their unemployment spell, they may not have left the unemployment insurance system yet and therefore this is a preferred selection criteria. In addition, we also added everyone who left the unemployment insurance system with registered indication 'employed' to the outflow-survey sample. The outflow survey contains a number of questions about the new job (starting date, occupation, and a comparison of tasks relative to the pre-unemployment job). It is important to note that these outflow surveys are intended only for those who found a job. For that reason, we specify in the invitation that the survey is only relevant if individuals indeed found a job. Once individuals open the survey, they are asked once again if they indeed did find a job and only then do they continue on to the survey.

5.2.5 Hypotheses

The aim of the intervention is to make job seekers aware of suitable alternatives to the occupations they are currently looking for work in, and motivate them to look for work in these occupations. If effective, the likely impact on job finding is not straightforward, however. In the short term, the expected effect on the likelihood of finding a job is ambiguous. On the one hand, when individuals start looking for work in more promising occupations, they will likely have more vacancies to apply to, with fewer competing job seekers per vacancy. On the other hand, despite the relevancy of the suggested alternatives, job seekers will likely have less experience in these new occupations, which decreases their comparative advantage. Moreover, they might need some time to adjust their search efforts.

Once individuals have had time to adjust, a successful intervention would likely lead to treated job seekers ending up in different occupations. Since the alternative occupations offer better job opportunities, one would expect that these job seekers will more often be employed, and remain with the same employer for longer. While the differences in the demand for and supply of labor between these occupations may lead to higher wages in the alternative occupations, it is important to note that we do not take this into account in the intervention. We therefore make no predictions on changes in earnings conditional on having a job. However, total earnings are likely to be different between the control and treatment groups, because of differences in rates of employment.

5.3 Descriptive Results

Before turning to the analysis of the impact of the interventions in Section 5.4.2, we first provide descriptive statistics for our sample and document a range of descriptive findings regarding job search behavior and beliefs in our data. In Table 5.2 we show that the job seekers in our sample have a fairly long unemployment duration at the time of selection, with a mean of 32 weeks. This is not surprising, given that we selected job seekers from those occupations with the worst job finding prospects. Most job seekers still are entitled to substantial benefits (51 weeks on average). Our selection of occupations also resulted in a skewed gender distribution, with only 25% males and 75% females. The distribution across selection occupations shows that bartenders/waiters, office support staff and receptionists are by far the largest groups and in all of these women are over-represented. As stated in Section 5.2.2, seven of these occupations can be classified as ‘Covid occupations’, which are occupations which were hit particularly hard by the Covid pandemic.

Table 5.2: Sample descriptives: administrative data

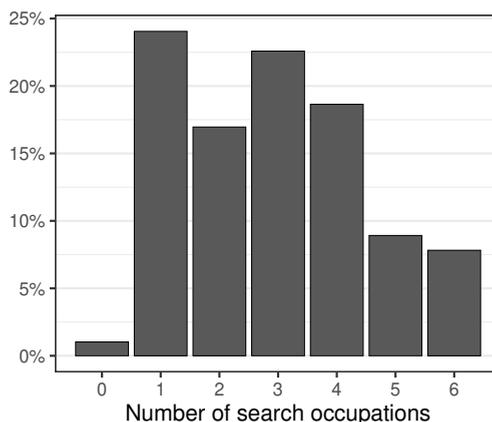
	Mean	Std.Dev.	Min	Max
Male	0.25	0.43		
Unemployment duration (wks.)	32.17	28.07	0.00	463.00
Remaining benefits (wks.)	50.98	29.70	4.14	188.71
Covid selection occ.	0.49	0.50		
Selection Occupation:				
<u>Non-covid occupations</u>				
Activity counsellor	0.03	0.18		
Archivist	0.01	0.10		
Video and sound technician	0.01	0.10		
Janitor/Concierge	0.03	0.17		
Animal caretaker	0.01	0.12		
Printer	0.01	0.10		
Graphic designer	0.03	0.16		
Office support staff	0.21	0.41		
Primary school teaching assistant	0.02	0.15		
Event/conference organizer	0.02	0.15		
Producer (television/film)	0.01	0.09		
Social worker	0.08	0.27		
Steward/stewardess	0.01	0.10		
Shop attendant household/leisure	0.01	0.12		
<u>Covid occupations</u>				
Hotel receptionist	0.02	0.13		
Hairdresser	0.02	0.14		
Bartender/waiter	0.16	0.37		
Canteen/Buffer employee	0.07	0.25		
Receptionist	0.17	0.37		
Travel agent	0.02	0.13		
Taxi driver	0.05	0.21		
Observations	30,129			

Note: based on administrative data of the full sample of experiment participants. Remaining benefits and unemployment duration are measured in March 2021.

5.3.1 How do job seekers search?

For the subsample that completed the pre-intervention survey (N = 2,789) we obtain a rich set of responses regarding job search beliefs and activities. While those invited to the survey were randomly selected, those who responded may not be. In Table 5.A.2 in the Appendix, we compare the survey respondents to the rest of the sample and conclude they are fairly similar. There are no significant differences in gender composition or unemployment duration. Only the remaining benefit rights are higher for survey respondents and there is a slight difference in the distribution across selection occupations. Based on observable characteristics, we conclude that we can interpret the survey responses as fairly representative of the full experimental population.

Figure 5.4: Number of search occupations



Note: y-axis displays the share of recipients that search for the number of occupations displayed on the x-axis. A value of 0 indicates the individual indicated they were not currently searching for work in any occupation. Job seekers could fill in at most 6 occupations, so the value of 6 may be interpreted as 6 or more.

Survey respondents first indicate what their primary search occupation is (typically the selection occupation) and which alternative occupations they consider. In Figure 5.4 we show how many occupations respondents list as their search occupations (their primary occupation, as well as alternatives). Almost 25% searches for work in only one occupation, while 40% searches in two or three occupations. Around 35% searches in more than three occupations. In Appendix Table 5.A.3, we show that most respondents (i) spend at least some hours per week exploring alternative occupations, (ii) are fairly willing to consider new occupations, (iii) have quite some confidence in their ability to work in an occupation in which they have no experience, and (iv) believe that their skills are transferable. Over 50% of respondents expects to widen their search in terms of occupations if they are still unemployed in two months.

For the primary and first alternative search occupation, we collect various measures of job search activities and elicit beliefs about the returns to job search (see Table 5.3). As the primary occupation is for most individuals the selection occupation, it has a low job finding score (3.2, row 1).⁷ The first alternative occupation that they search in offers better prospects with an average job finding score of 4.3. In the previous two weeks the average number of applications for jobs in the primary occupation is 3.1, while it is 2.5 for the first alternative occupation (row 2). The resulting number of job interviews follows a similar pattern: 0.43 for the primary occupation and 0.37 for the first alternative. The number of interviews per application is slightly higher for the alternative occupation (row 4), which is consistent with the higher job finding

⁷They can indicate at the start of the pre- and post-surveys that the selection occupation is not their primary occupation of search and provide a different primary occupation.

score. Expect for interviews, all of these differences are statistically significant.

Table 5.3: Comparison primary and alternative occupation (survey data)

	Primary (N=2789)		Alternative (N=2789)		Diff.	p
	Mean	Std. Dev.	Mean	Std. Dev.		
Job search activities:						
Job finding score	3.20	1.02	4.26	1.67	1.06	0.00
Applications sent (past 2 weeks)	3.14	6.16	2.51	4.78	-0.64	0.00
Job interviews (past 2 weeks)	0.43	1.32	0.37	1.04	-0.06	0.12
Interviews per application	0.16	0.43	0.20	0.45	0.04	0.01
Expectations:						
Expected job offer rate	0.10	0.11	0.10	0.11	0.00	0.67
Expected wage	2638.27	866.48	2698.29	1003.90	60.02	0.03
Reservation wage	2563.07	878.22	2596.49	933.67	33.42	0.21
Job stability	0.68	0.30	0.71	0.27	0.03	0.00
Exp. appl. if equal job offer rate	4.20	7.81	4.47	8.39	0.28	0.29
Exp. appl. if equal wage	4.28	7.59	4.41	7.99	0.13	0.62
Exp. job offer rate in 2 months	0.09	0.10	0.08	0.10	-0.01	0.11

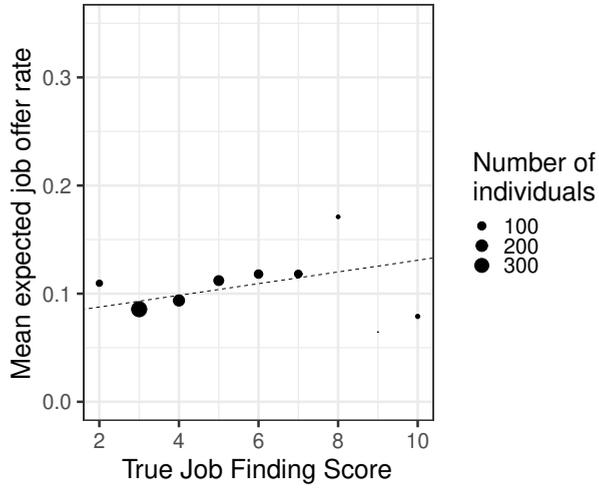
Note: *primary occupation* is the occupation that the respondent searches primarily in. *alternative occupation* is the occupation that the respondent considers the most important alternative occupation of search. The number of observations varies slightly across variables due to item non-response. *job stability* is defined as the expected probability of being able to keep a new job for at least two years. *exp. appl. if equal job offer rate* is the expected number of applications per week in case the job offer rate would be equal in the primary and alternative occupation. *exp. appl. if equal wage* is the expected number of applications per week in case the job offer rate *and* the expected wage would be equal in the primary and alternative occupation. *exp. job offer rate in 2 months* is the expected job offer rate in case the respondent is still unemployed in two months time.

5.3.2 How well are job seekers informed?

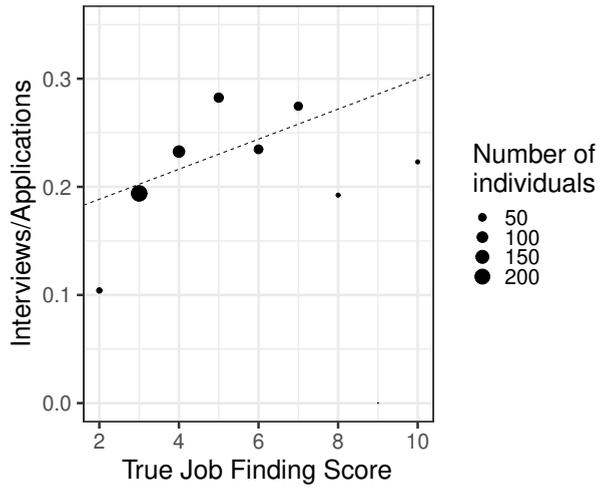
We elicit a range of beliefs about the returns to job search activities and labor market prospects. The key question of interest is whether expectations regarding job prospects in various occupations align with actual prospects. In addition, we explore whether these expectations drive job search activities. First, respondents indicate their belief about the number of applications it requires on average to obtain one acceptable job offer, both for their primary occupation and their first alternative. By inverting this number we obtain the expected job offer rate (per application), which is fairly small on average (0.10, row 5) and strikingly similar between the primary and alternative occupation. We do not have a direct measure of the actual job offer rate, but the job finding score shows a large difference between the primary and alternative occupation (row 1: 3.20 versus 4.26, respectively). We conclude that, on average, job seekers are not aware that job finding prospects are significantly better in their alternative occupations. Job seekers also expect to earn higher wages in the first alternative occupation and have a slightly higher reservation wage for the alternative, although the difference is not significant. Expectations about job stability (the probability of keeping a new job for at least two years), are slightly more optimistic for the alternative occupation with a small but significant difference. Finally, the last row shows that job seekers expect to update their expectations about job offer rates, but only slightly. If they are still unemployed in two months time, they expect the job offer rate to be 0.09 for the primary occupation (compared to 0.10 now) and 0.08 for the alternative occupation (compared to 0.10 now).

To assess how well job seekers are informed about job prospects, we

Figure 5.5: Expected and actual job finding prospects



(a)



(b)

Note: both Panels refer to the individuals main alternative occupation of interest. Panel (a) relates the expected job offer rate to the job finding score of the occupation. Panel (b) relates it to the ratio of interviews to applications.

link their beliefs to the actual job finding prospects. We exploit variation across individuals in their selection of alternative occupations. Specifically, we examine the relationship between job finding score and expected job offer rate in Figure 5.5. In Panel (a) we find that this relationship is fairly flat: regardless of the true job finding score, the expected job offer rate of an application is always close to 0.1. A linear regression produces a positive and significant slope coefficient ($\hat{\beta} = 0.003$, p-value = 0.03), but the magnitude is very small. A further observation is that most of the alternative occupations have only low to medium job prospects (3-5), with a small share having good prospects (6-7) and a tiny fraction having excellent prospects (8-10). These two facts suggest that job seekers do not select their alternative search occupations on the basis of better job prospects. First, most job seekers select alternatives with only marginally better job prospects. Second, even those who select high-prospect alternatives do not seem to be aware of these better job finding chances.

In Panel (b) of Figure 5.5 we investigate whether the better job prospects translate into better returns to job search based on the reported number of applications and interviews. We see some indication that indeed the occupations with a higher job finding score lead to a higher interview per application rate. The linear regression coefficient is much larger, but not statistically significant ($\hat{\beta} = 0.012$, p-value = 0.23).

Summing up, we draw the following two key conclusions regarding job search strategies of the job seekers in our sample.

1. While most job seekers indicate that they are willing (and confident) to search in alternative occupations, the majority searches only in 1-3 occupations.

2. Job seekers do not appear to be well-informed about the stark difference in actual job finding prospects between their primary search occupation and potential alternatives.

These two findings are both encouraging for the potential of information interventions that bring the variation in job prospects to job seekers' attention. We now proceed by analyzing our intervention's impact.

5.4 Empirical Analysis

5.4.1 Take-up: email opening and clicking statistics

Job seekers in the treatment groups received their first email with occupational information on April 12 (see Section 5.2). We first compare the suggested occupations to the occupations in which job seekers report they search, to assess to what degree we provide 'new' information. Then we present statistics on engagement: whether they opened the email and clicked on the links. These statistics provide an indication of 'treatment take-up'.

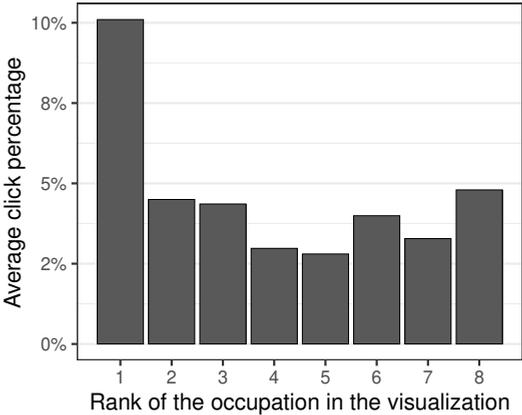
If most job seekers already search in a couple of occupations that we offer as 'high prospect alternatives', we are unlikely to provide novel information to them. In Figure 5.B.1 in the Appendix, we show the number of suggestions that an individual received in the emails that was already present in their search set as measured in the pre-intervention survey.⁸ It turns out that the vast majority (78%) searches in none of the suggested occupations before receiving the emails. A small group

⁸Both the search occupations and the suggestions are defined at the 5-digit BRC+ level.

was already searching in one of our suggested occupations (18%) and a negligible share already searches in more than one suggested occupation.

A total of 19,960 job seekers received the first email (both treatment groups). From these, 12,804 opened the email (64%). Each occupation is clickable for more information about the occupation (description, tasks, skills, related occupations, educational level). The share of recipients that clicks on each occupation provides a measure of how interesting each occupation is to job seekers. In total, we observe 4975 clicks on occupations. These are not evenly distributed across the total of 165 presented suggestions (21 selection occupations with each between 7 and 9 occupational suggestions).

Figure 5.6: Clicks on occupation by rank



Note: data is at the ‘suggestion-selection occupation’ level (165 observations). Click percentage is conditional on opening the email.

Table 5.4: Clicks on occupation by rank, job finding score, automation risk and no. of transitions

	<i>Dependent variable:</i>		
	Percentage of recipients that clicks		
	(1)	(2)	(3)
Rank 2	-0.06 *** (0.01)	-0.05*** (0.01)	-0.05*** (0.01)
Rank 3	-0.06 *** (0.01)	-0.04*** (0.01)	-0.05*** (0.01)
Rank 4	-0.07 *** (0.01)	-0.05*** (0.01)	-0.05*** (0.01)
Rank 5	-0.07 *** (0.01)	-0.06*** (0.01)	-0.06*** (0.01)
Rank 6	-0.06 *** (0.01)	-0.05*** (0.01)	-0.05*** (0.01)
Rank 7	-0.07 *** (0.01)	-0.05*** (0.01)	-0.05*** (0.01)
Rank 8	-0.05 *** (0.01)	-0.04*** (0.01)	-0.04*** (0.01)
Job finding score (tightness)			0.01*** (0.003)
Low automation risk			0.03*** (0.01)
Relative no. of transitions		0.09*** (0.02)	0.09*** (0.02)
Constant (Rank 1)	0.10 *** (0.01)	0.08*** (0.01)	0.01 (0.02)
Observations	165	165	165
R ²	0.27	0.33	0.45

Note: Table displays OLS regression at the ‘suggestion-selection occupation’ level. Ranks are dummies. Baseline category is rank 1. *Job finding score (tightness)* is a continuous variable. Low automation risk is a dummy. Relative no. of transitions is a continuous variable, indicating the fraction of total transitions from the selection occupation to the suggestion occupation (ranges from 0 to 1). * p<0.1; ** p<0.05; *** p<0.01

In Figure 5.6 we show how the number of clicks depends on the ranking of the suggestions within the visualization. The top one is by far most popular, while also the last ones are slightly more popular. The popularity of the first suggestion reflects both (i) that job seekers start at the top of the visualization when reading the email and (ii) that the suggestions are (primarily) ranked based on the number of historically observed transitions, suggesting that the higher placed suggestions are the most suitable ones. However, disentangling the two is possible because the ranking was not perfectly aligned with the number of transitions. Because of the fact that we let the employment office alter the occupational order to some degree, some occupations with a lot of transitions ended up at lower ranks. A simple regression at the occupation-level (165 observations) uncovers the importance of both rank and transitions. In Table 5.4 (Column 1) we include seven rank-indicators. In Column (2), we add the relative number of transitions (the share of all transitions within the 7-9 suggested occupations). Transitions, while not observable to job seekers, turn out to be highly statistically significant, and rank becomes less important. This is encouraging, as it suggests that our method of selecting suitable ('fitting') occupations seems effective. In Column (3) we add the job finding score and the indicator for the suggestion displaying 'low-automation risk', both of which also have a statistically significant and positive impact on the number of clicks. Again, this is encouraging, as it suggests that we are providing relevant information.

We sent a reminder email with a similar visualization on May 28th. In coordination with the communication experts from the employment office we decided to change the content slightly, to maximize the relevance of the message. Using the regression model from Column (3)

we generated predicted interest, controlling for the rank in the first email. Thus, we predict interest based on the job finding probability, the automation indicator and the number of transitions. Using these predictions we created a new ordering which was implemented in the second email. In addition only the new top-5 suggestions were included to make the message slightly shorter. The email was sent out to 16,838 individuals, of which 11,475 opened it (68.1%). Of those who opened it, 2,442 clicked on a link (21.3%). Over both emails, 15,867 individuals opened at least one (78.8%), of which 4,874 clicked on at least one link (30.7%).

The motivational treatment group received a version of the first email that contained an extra paragraph with a link to the motivational video. In contrast to suggestion links, very few people (0.5%) clicked on the video link. A likely explanation is that the video was only provided *after* the information visualization, and many readers may not have reached this part of the email. Of course, it might also be that job seekers are simply not interested in the video. We sent an additional email to this treatment group that *only* provided the video link (not the occupation information). This email led to a slightly higher click rate (7.5%), but still the overall share of the motivational treatment group that has seen the video remains low. Given the low 'take-up' of the video, our analysis in the next Section will combine the two treatment groups and only measure the effect of the informational content that both groups received.

5.4.2 Experimental analysis

Given the randomized assignment, the empirical strategy is straightforward and we can simply compare outcomes across the treatment and the control group. Following our pre-analysis plan, we first consider the primary outcomes, which are employment (earnings, hours and occupation) and benefit receipt. Subsequently we turn to job search behavior as measured in the post-intervention survey.

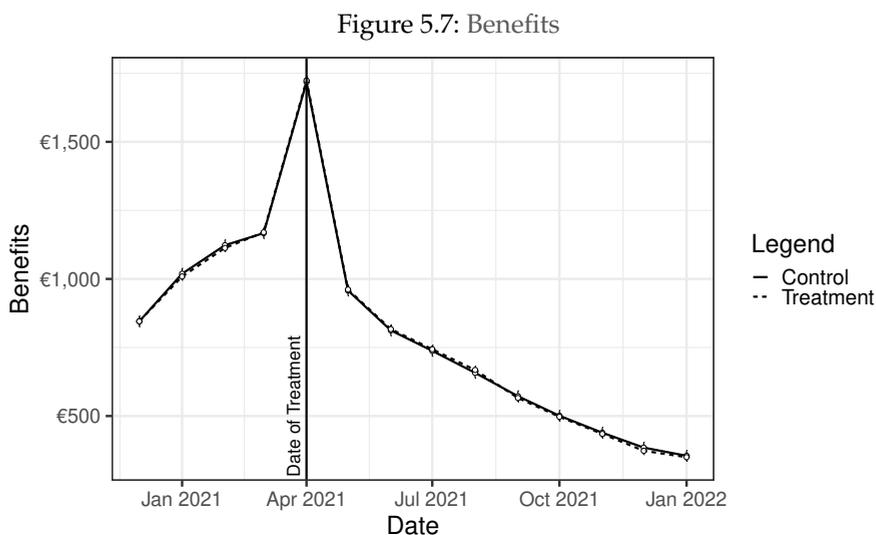
5.4.2.1 Balancing checks and further sample selection

Before turning to the analyses, it is worth checking whether our data is balanced on the most important dimensions. As randomization was stratified by gender, unemployment duration (in three bins) and selection occupation, we obtain near-perfect balance on these variables, as we show in Appendix Table 5.A.4. In Table 5.A.5 in the Appendix, we show that the samples are also balanced in terms of responses to the pre-intervention survey as well.

As stated before, however; a number of our selection occupations recovered swiftly after most Covid-restrictions were lifted. As such, demand for these occupations strongly increased again. Since individuals who were looking for work in these occupations are likely to be able to find a job in that occupation again, the treatment is likely not as effective for them and does not align with our initial question of interest. As such, we focus our analyses only on individuals looking for work in non-Covid occupations. Tables 5.A.6 and 5.A.7 show that we our sample is balanced on all relevant variables for this subsample as well.

5.4.2.2 Treatment effect on benefits receipt and labor earnings

Labor earnings and benefits receipt are measured using the administrative data provided by the public employment office. The data covers all experimental participants and we simply look at the month-by-month development of the outcome variable for the control and treatment groups.



Note: based on administrative data of the full experimental sample (of individuals with a non-Covid selection occupation). Dots indicate estimates, with vertical bars displaying the 95% confidence interval.

We consider benefit receipts the most accurate measure of benefit dependence as it is typically complex to define a specific binary point of outflow from unemployment insurance benefits. Many job seekers find temporary and part-time jobs while continuously receiving (fluctuating) unemployment insurance benefits. The amount of benefits received thus provides an aggregated and complete measure of

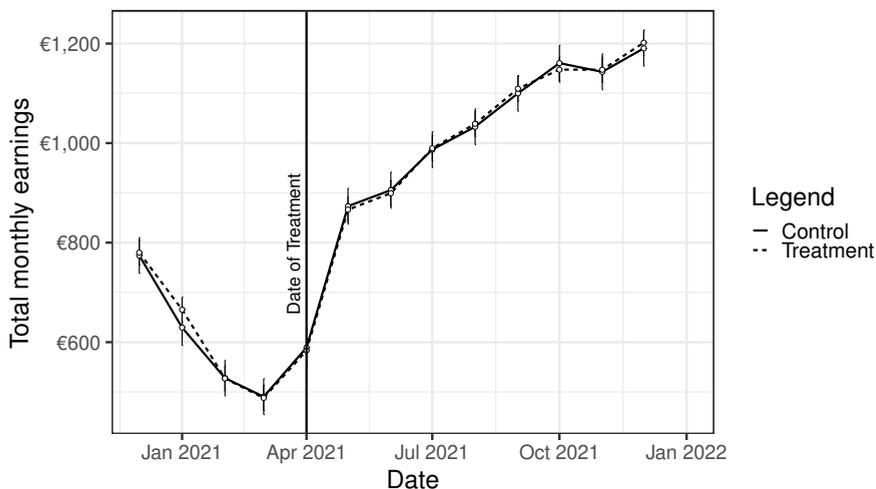
job search success. Figure 5.7 shows the mean benefit receipts for both groups, as well as the 95% confidence intervals, from December 2020 to January 2022. Before our intervention, from December 2020 to April 2021, benefit receipts increase equally for the control and treatment groups. Since we selected our sample on the 15th of March, 2021, benefits peak in April 2021. From April to May 2021, we observe a large drop in the received benefits for both the control and treatment group. Total received benefits drop from well over €1,600.- to below €1,000.-. From May 2021 onwards, total received benefits steadily declines for both the control and treatment group, ending up just above €300.- in January of 2022. There is essentially no difference between the control and treatment groups in any of the observed months. Given the tight confidence intervals, we are able to rule out that the treatment impact is larger than €30.-. Figure 5.B.2 in the Appendix paints a similar picture for the share of individuals receiving positive amounts of benefits.

Figure 5.8 displays the development of labor earnings over approximately the same time frame. The pattern is essentially the inverse of that of Figure 5.7. Earnings are lowest in March 2021, at the point where we selected our sample. From April to May of 2021, we observe a large increase, followed by a smaller, but steady increase in the following months. We again find no impact of the treatment, and can confidently rule out large effects.

5.4.2.3 Type of work found

While we find no difference in terms of monthly benefits or labor earnings, the means may mask more subtle differences in job finding be-

Figure 5.8: Earnings



Note: based on administrative data of the full experimental sample (of individuals with a non-Covid selection occupation). Dots indicate estimates, with vertical bars displaying the 95% confidence interval.

tween the groups. In addition, we would like to compare the occupations of the new jobs between the control and treatment group. To do so, we analyze the outflow survey. As described in more detail in Section 5.2.4, the outflow survey was sent at three-month intervals to all experiment participants for whom we observed a substantial increase in monthly earnings over the preceding months. We take such an increase in earnings as a strong indicator of job finding. As a result, the *invitation* to the survey is based purely on administrative data and can be used as a measure of job finding.⁹

⁹Note that practical challenges in terms of data access made it impossible to use actual data on job finding on a rolling basis for selecting survey recipients.

Table 5.5 provides a summary of the response to the survey. Of the 14,812 individuals in our non-Covid subsample, 6,730 received an invitation to fill out the outflow survey. Row 1 shows that while the share of treated individuals who received an invitation is slightly higher (0.458) than that of the control group (0.446), the difference is not statistically significant. Row 2 shows what share of individuals opened the survey. For both the full sample, as well as the sample who received an invitation we see that a larger share of treated individuals opened the survey. Looking at the full sample, the difference is 1.7 percentage points (or 12%). For the sample who received the invitation, the difference is 2.9 percentage points (or 10%). These differences are statistically significant and since the invitation specifically states that the survey is only intended for those who have found a job, this suggests a positive treatment effect on the likelihood of finding employment. The last row shows a similar picture. From the full sample, the share of treated individuals who have (i) received an invitation, (ii) opened the survey and (iii) answered they indeed found a job is 1.2 percentage points higher than in the control group; this represents a 12% difference. This difference is again statistically significant. When we look at the share of respondents as a fraction of the individuals who were invited to the survey, the effect is 2.0 percentage points, or 9%. A slightly smaller difference, and marginally significant.

The results point in the direction of slightly higher job finding rates in the treatment group. We further investigate potential dynamics in the job finding rates. In Figure 5.9 we plot the number of survey invitees (the assumed job finders) for the treatment and control group over time. We find that the gap in job finding grows over time until December 2021 (about 7 months after the intervention). At that point

Table 5.5: Outflow survey invitations and response rates

	Share of Sample				Share of Invited			
	Control (N=4930)	Treatment (N=9882)	Diff.	p	Control (N=2201)	Treatment (N=4529)	Diff.	p
Invited to survey	0.446	0.458	0.012	0.178				
Opened survey	0.134	0.151	0.017	0.007	0.301	0.33	0.029	0.019
Responded job found	0.104	0.116	0.012	0.034	0.234	0.253	0.02	0.084

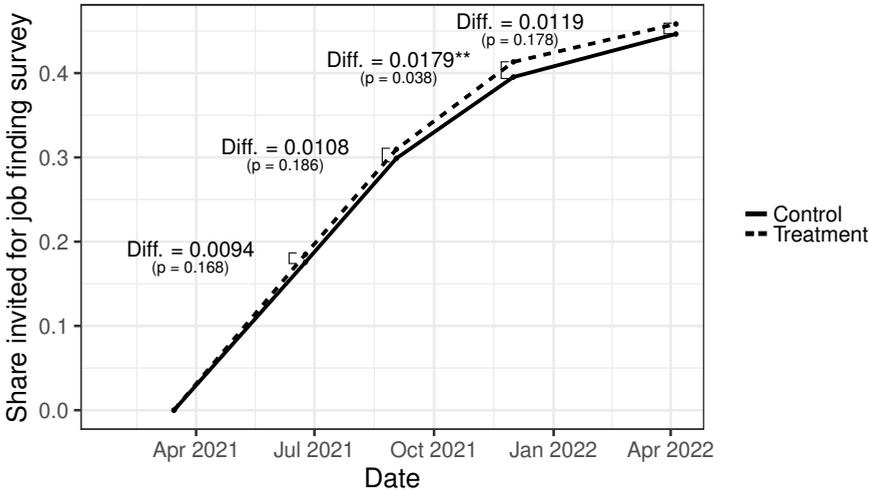
Note: the numbers for the control and treatment group in the rows below ‘Share of Sample’ are the number of observations relevant to the row (e.g. number of individuals invited to the survey) divided by the number of individuals in the full sample. The same holds for ‘Share of Invited’.

the difference is statistically significant. By April 2022 (about one year after the intervention) the difference remains but has become slightly smaller.

Table 5.6 provides more detailed insights into what type of jobs people found, for those who completed the outflow survey. The table compares the occupation of the new job with the ‘selection occupation’.¹⁰ Each row shows the share of respondents that found employment in the same occupation as the occupation they were selected for. The difference between the rows is the occupational classification used, going from very fine grained (5-digit) in Row 1, to very broad (2-digit) in Row 4. We find that for 15-19% the new job is the exact same occupation as the occupation that we selected them for. Using broader classifications this share grows to more than 50% (when comparing only 2 digits). Regardless of the occupational classification, a larger share of treated individuals indicate that they found a job that is different from the one they were selected for. This difference is approximately 5 percentage

¹⁰The survey asked for a free text job title, which were blindly coded into a 5-digit occupational code.

Figure 5.9: Outflow survey invitations



Note: invitations to outflow survey are based on administrative data. Individuals are invited if they either indicate to the employment office that they have found employment, or if they show a month-to-month earnings increase of over € 300.- in the two to three months before the invitations are sent out. Based on the full experimental sample (of individuals with a non-Covid selection occupation). Dots indicate estimates. Text to the left of the lines displays difference in percentage points and p-value of difference between brackets.

points. Note that errors in the classification are likely to occur, which would lead us to underestimate the numbers in Table 5.6. Since the classification was performed blindly with respect to the treatment status, there is however no reason to believe that this affects the difference between treatment and control group.

Table 5.6: Outflow to different occupations

	Control (N=514)		Treatment (N=1147)		Diff. in Means	p
	Mean	Std. Dev.	Mean	Std. Dev.		
Same 5-digit occupation	0.19	0.40	0.15	0.35	-0.05	0.06
Same 4-digit occupation	0.27	0.45	0.21	0.41	-0.06	0.04
Same 3-digit occupation	0.41	0.49	0.36	0.48	-0.05	0.10
Same 2-digit occupation	0.55	0.50	0.50	0.50	-0.05	0.13

Note: The mean for same 'x'-digit occupation denotes the share of individuals who found work in the same occupation as their selection occupation according to the 'x'-digit classification.

In summary, we find no difference in the mean labor earnings of mean benefit receipt, but a clear indication of higher job finding rates and more diverse occupations in the treatment group. We now turn to secondary outcomes of interest, which are job search behavior and beliefs about the labor market. These are only measured through the post-intervention survey and therefore only available for the small subset of participants that completed the survey.

5.4.2.4 Treatment effect on survey responses

For the outcome variables that we collected through the survey, we have precise pre-intervention measurements and we opt for a difference-in-differences model that controls for baseline differences to increase statistical power. The baseline specification is

$$Y_{it} = \beta_0 + \beta_1 P_{it} + \beta_2 T_i + \beta_3 P_{it} T_i + \varepsilon_{it}, \quad (5.1)$$

with T_i a treatment indicator and P_i a time period indicator (equal to 1 for the post-intervention period). Using the survey data, we first consider measurements of job search activities and beliefs. In Table 5.7 we show regression estimates. The number of observations varies across Columns, as we only include individuals who answered the respective questions in both the pre- and post-intervention surveys. We consider the outcomes weekly time spent on exploring alternative occupations (Column 1), total number of weekly applications (Column 2), total number of weekly interviews (Column 3), number of occupations included in the search (Column 4), the mean job finding score of the set of search occupations (Column 5) and the number of suggestions from the email that are included in the set of search occupations (Column 6). We find that the treatment effect estimate (Treatment \times Post, $\hat{\beta}_3$), is never statistically significantly different from zero. Thus we cannot reject that the treatment has no observable impact on job search activities as measured along these six dimensions.

Table 5.7: Difference-in-differences analysis survey outcomes: Job search activities

	<i>Dependent variable:</i>					
	Time exploring	Applications	Interviews	Number of search occupations	Mean jobfinding score	Suggestions used in search set
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment group	0.252 (0.532)	-1.993 (1.455)	0.232 (0.257)	-0.106 (0.151)	0.041 (0.100)	0.072 (0.049)
Post-period	-0.770 (0.614)	-0.877 (1.705)	0.012 (0.301)	-0.342* (0.175)	0.222* (0.115)	0.032 (0.056)
Treatment × Post	-0.577 (0.752)	0.002 (2.058)	0.061 (0.364)	0.015 (0.214)	-0.162 (0.142)	-0.015 (0.070)
Constant	5.354*** (0.434)	7.781*** (1.206)	0.554*** (0.213)	3.037*** (0.124)	3.765*** (0.081)	0.186*** (0.040)
Observations	964	466	522	964	910	910

Note: the dependent variables are weekly time spent on exploring alternative occupations (Column 1), total number of weekly applications (Column 2), total number of weekly interviews (Column 3), the mean job finding score of the set of search occupations (Column 4) and the number of suggestions from the email that are included in the set of search occupations (Column 5). Treatment (group) is a dummy indicating whether the individual is in the treatment group. Post(-period) is a dummy indicating whether the observation is from the post-experimental survey. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Next, we perform a similar analysis for the beliefs. Table 5.8 shows the results. We consider the expected job offer rate per application in the primary occupation (Column 1) and first alternative occupation (Column 2), the expected job stability of a job in the primary occupation (Column 3) and the alternative occupation (Column 4) and the probability of finding employment in the next two months (Column 5). Again, we find no statistically significant effects of the treatment on any of the belief measures.

Table 5.8: Difference-in-differences analysis survey outcomes: labor market beliefs

	<i>Dependent variable:</i>				
	Job offer rate per application primary	Job offer rate per application alternative	Expected stability primary	Expected stability alternative	Job finding probability
	(1)	(2)	(3)	(4)	(5)
Treatment group	-0.017 (0.064)	-0.015 (0.096)	-0.010 (0.030)	0.028 (0.035)	0.039 (0.035)
Post-period	0.097 (0.078)	-0.107 (0.113)	-0.020 (0.035)	-0.007 (0.040)	0.081** (0.041)
Treatment × Post	-0.018 (0.094)	0.048 (0.136)	0.016 (0.043)	0.013 (0.049)	-0.008 (0.050)
Constant	0.147*** (0.053)	0.282*** (0.081)	0.658*** (0.025)	0.669*** (0.028)	0.297*** (0.029)
Observations	334	203	964	614	688

Note: The dependent variables are: weekly time spent on exploring alternative occupations (Column 1), total number of weekly applications (Column 2), total number of weekly interviews (Column 3), the mean job finding score of the set of search occupations (Column 4) and the number of suggestions from the email that are included in the set of search occupations (Column 5).

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

These results are difficult to square with our finding that treated job seekers seem to have found employment in occupations different from their initial occupation of interest sooner. There are a number of possible explanations for the null effects we find on search behavior and beliefs. First, sample size becomes fairly small at this stage, with only around 300-600 observations for some outcomes (implying 150-300 individuals per treatment/control). Starting from an experimental sample of 30,000, this limits statistical precision. Indeed, wide confidence

intervals cannot reject substantial positive (or negative) impacts. Second, the small sample size also hints at the possibility of selective response: while those invited to answer the survey were randomly drawn, the sample that completed both the pre- and post-survey may not be representative of the full sample in terms of unobservables. Third, search activities and beliefs may be difficult concepts to measure in a survey, resulting in measurement error (in both the pre- and post-survey) and attenuation bias in our estimates. Obtaining administrative data on job search activities as registered by case workers and through logged activities on the national job search website is ongoing. These data are arguably more precise and will be available for the entire sample.

5.4.3 Remaining analyses

Various extensions of the analyses remain to be performed, as outlined in the pre-analysis plan that can be found in the AEA RCT registration. Firstly, there are a number of heterogeneity analyses to be done. We expect that job seekers' prior search strategy is an important determinant for treatment impact. For instance, we expect a more pronounced impact of providing information if job seekers initially search narrowly. In addition, we plan to explore heterogeneity by unemployment duration, expecting more willingness to consider alternatives among those that have been unemployed for a longer time. Both hypotheses are based on findings from Belot et al. (2019). Second, we will investigate other job search activities that are collected administratively by the employment office. These include both measures of job search collected by case workers (applications and interviews) and records of online ac-

tivities on the employment office's job search platform. Lastly, we will keep following our sample for at least ten more months to measure potential long-term impacts.

5.5 Conclusion

We provide unemployed job seekers looking for work in occupations with poor labor market prospects with information about suitable alternative occupations that offer better prospects. In addition we offer a motivational video aimed at overcoming behavioral hurdles associated with occupational transitions. Combining administrative data with pre- and post-intervention surveys to collect labor market beliefs, we measure how these interventions may contribute to opening up job seekers' job search horizon and stimulate them towards occupational mobility to jobs with better prospects.

Our descriptive statistics show that our sample of job seekers is likely to respond to the information treatment. While many report to be willing to explore new alternatives and are confident about their ability to work in a new occupation that matches their skillset, actual job search is fairly narrow in terms of occupations. Moreover, beliefs about job offer rates show that awareness of the large variation in labor market prospects across occupations is very limited.

We do not find that the interventions had any impact on benefits receipt or labor earnings up to at least eight months after the treatment. We do find an indication that those in the treatment group found employment sooner, and more often found employment in an occupation different from their previous occupation than those in the control

group. This means the effects we find are not confined to closely related occupations. Whether the difference in occupations of new jobs leads to more job stability, higher earnings and lower benefit dependence in the long-run remains to be studied in the future. Our survey data, collected two months after the experiment, does not show any impact on search behavior or beliefs, but the sample size is limited.

Appendix 5.A Additional Tables

Table 5.A.1: Selection occupations with low job prospects

Occupation	Occupation (Dutch name)
Activity counsellor	Activiteitenbegeleider
Animal caretaker	Dierenverzorger
Archivist	Archiefmedewerker
Bartender/waiter	Medewerker bediening/bar
Canteen/Buffer employee	Medewerker bedrijfsrestaurant of buffet
Event/conference organizer	Organisator van conferenties en/of evenementen
Graphic designer	Grafisch vormgever
Hairdresser	Kapper
Hotel receptionist	Hotelreceptionist
Janitor/Concierge	Conciërge/huismeester
Office support staff	Ondersteunend medewerker op een kantoor/secretariaat
Primary school teaching assistant	Onderwijsassistent basisonderwijs
Printer	Drukkerijmedewerker
Producer (television/film)	Productieleider/producent
Receptionist	Receptionist/telefonist
Shop attendant household/leisure goods	Verkoopmedewerker huishoudelijke en vrijetijd-sartikelen
Social worker	Sociaal werker
Steward/stewardess	Steward/stewardess
Taxi driver	Taxi- of particulier chauffeur
Travel agent	Reisadviseur/reisbureau-medewerker
Video and sound technician	Beeld- en geluidtechnicus

Table 5.A.2: Comparison of composition of survey-respondents and rest of sample

	Non-survey (N=27340)		Survey (N=2789)		Diff.	p
	Mean	Std. Dev.	Mean	Std. Dev.		
Male	0.25	0.43	0.24	0.42	-0.01	0.17
Unemployment duration	32.25	28.00	31.43	28.71	-0.83	0.15
Remaining benefits (wks.)	49.96	29.53	61.03	29.48	11.07	0.00
Covid selection occ.	0.49	0.50	0.51	0.50	0.02	0.05
Selection Occupation:						
Activity counsellor	0.03	0.18	0.04	0.19	0.00	0.79
Archivist	0.01	0.10	0.01	0.09	0.00	0.26
Video and sound technician	0.01	0.10	0.00	0.07	-0.01	0.00
Janitor/Concierge	0.03	0.17	0.04	0.19	0.01	0.04
Animal caretaker	0.01	0.12	0.01	0.11	0.00	0.33
Printer	0.01	0.10	0.01	0.10	0.00	0.55
Graphic designer	0.03	0.16	0.03	0.16	0.00	0.74
Hotel receptionist	0.02	0.13	0.02	0.12	0.00	0.57
Hairdresser	0.02	0.14	0.01	0.12	-0.01	0.01
Bartender/waiter	0.16	0.37	0.14	0.35	-0.02	0.00
Canteen/Buffer employee	0.06	0.25	0.08	0.27	0.02	0.00
Office support staff	0.21	0.41	0.23	0.42	0.02	0.06
Primary school teaching assistant	0.02	0.15	0.02	0.12	-0.01	0.00
Event/conference organizer	0.03	0.16	0.02	0.13	-0.01	0.01
Producer (television/film)	0.01	0.09	0.01	0.09	0.00	0.80
Receptionist	0.16	0.37	0.19	0.40	0.03	0.00
Travel agent	0.02	0.12	0.02	0.13	0.00	0.70
Social worker	0.08	0.27	0.07	0.25	-0.01	0.02
Steward/stewardess	0.01	0.10	0.01	0.08	0.00	0.00
Taxi driver	0.04	0.21	0.05	0.22	0.01	0.15
Shop attendant household/leisure	0.01	0.12	0.01	0.11	0.00	0.73

Note: based on administrative data. Remaining benefits and unemployment duration are measured in March 2021. *Difference* = Survey - Non-survey. Column *p* indicates p-value of t-test.

Table 5.A.3: Survey responses about broader job search

	Mean	Std.Dev.	Min	Max
Search occupations suggested	0.23	0.49	0.00	3.00
Weekly hours exploring alternatives	5.61	5.95	0.00	20.00
Willingness to consider other occupations (1-5)	3.39	0.87	0.00	5.00
Confidence in working without experience (1-5)	3.76	0.80	0.00	5.00
Believes that skills are transferable (1-5)	3.76	0.80	0.00	5.00
Probability to expand search in two months	0.54	0.29	0.00	1.00
Observations	2,789			

Note: Based on survey data collected before intervention.

Table 5.A.4: Balance table: administrative records

	Control (N=10004)		Treatment (N=20126)		Diff.	p
	Mean	Std. Dev.	Mean	Std. Dev.		
Unemployment duration	32.27	28.02	32.13	28.09	-0.14	0.68
Male	0.25	0.43	0.25	0.43	0.00	0.98
Covid selection occ.	0.49	0.50	0.49	0.50	0.00	0.77
Remaining benefit (wks.)	50.92	29.61	51.01	29.74	0.09	0.81
Selection Occupation:						
Activity counsellor	0.03	0.18	0.03	0.18	0.00	0.77
Archivist	0.01	0.10	0.01	0.10	0.00	0.71
Video and sound technician	0.01	0.10	0.01	0.10	0.00	0.80
Janitor/Concierge	0.03	0.17	0.03	0.17	0.00	0.74
Animal caretaker	0.01	0.12	0.01	0.12	0.00	0.95
Printer	0.01	0.10	0.01	0.10	0.00	0.96
Graphic designer	0.03	0.16	0.03	0.16	0.00	0.72
Hotel receptionist	0.02	0.13	0.02	0.13	0.00	1.00
Hairdresser	0.02	0.13	0.02	0.14	0.00	0.35
Bartender/waiter	0.16	0.37	0.16	0.37	0.00	0.96
Canteen/ Buffet employee	0.07	0.25	0.07	0.25	0.00	1.00
Office support staff	0.21	0.41	0.21	0.41	0.00	0.88
Primary school teaching assistant	0.02	0.15	0.02	0.15	0.00	0.55
Event/conference organizer	0.02	0.16	0.02	0.15	0.00	0.88
Producer (television/film)	0.01	0.09	0.01	0.09	0.00	0.86
Receptionist	0.17	0.37	0.17	0.37	0.00	0.69
Travel agent	0.02	0.12	0.02	0.13	0.00	0.64
Social worker	0.08	0.27	0.08	0.27	0.00	0.97
Steward/stewardess	0.01	0.10	0.01	0.10	0.00	0.84
Taxi driver	0.05	0.21	0.04	0.21	0.00	0.44
Shop attendant household/leisure	0.01	0.12	0.01	0.12	0.00	0.85

Note: based on administrative data. remaining benefits and unemployment duration are measured in March 2021. *Difference* = Treatment - Control. Column *p* indicates *p*-value of t-test.

Table 5.A.5: Balance table: survey responses

	Control (N=931)		Treatment (N=1858)		Diff.	p
	Mean	Std. Dev.	Mean	Std. Dev.		
Unempl. duration	31.50	30.02	31.39	28.04	-0.11	0.93
Male	0.24	0.43	0.24	0.42	0.00	0.85
Job finding score sel. occ.	3.02	0.64	3.03	0.63	0.01	0.64
Covid selection occ.	0.52	0.50	0.50	0.50	-0.02	0.44
Time exploring alternatives	5.76	6.07	5.54	5.88	-0.22	0.35
Willingness work in new occ.	3.43	0.87	3.37	0.86	-0.06	0.11
My skills are transferable	3.77	0.81	3.76	0.79	-0.02	0.61
Prob. job in 2 months	0.42	0.30	0.42	0.30	0.00	0.86
Appl. needed (primary)	44.24	56.80	41.00	54.48	-3.24	0.24
Appl. needed (alt.)	43.98	56.36	40.82	54.45	-3.16	0.31
Salary previous job	2692.06	1168.23	2698.25	1197.17	6.19	0.90
Hours previous job	28.38	8.64	28.40	8.56	0.02	0.94
Expected wage (main occ.)	2617.08	827.09	2648.98	885.78	31.90	0.36
Reservation wage (main occ.)	2544.51	850.98	2572.46	891.78	27.94	0.43
Expected wage (alt. occ.)	2659.37	927.10	2717.81	1040.07	58.43	0.20
Reservation wage (alt. occ.)	2567.39	870.77	2611.17	963.83	43.78	0.30
Applications (main occ.)	3.17	6.90	3.13	5.76	-0.04	0.88
Job interviews (main occ.)	0.43	1.33	0.43	1.32	0.00	0.93
Applications (alt. occ.)	2.62	5.42	2.45	4.41	-0.18	0.47
Job interviews (alt. occ.)	0.37	1.03	0.37	1.04	0.00	0.92
Applications (other occ.)	2.83	6.43	2.46	4.71	-0.37	0.20
Job interviews (other occ.)	0.37	1.11	0.36	0.99	-0.01	0.91

Note: Based on survey data collected before intervention. *Difference* = Treatment - Control. Column *p* indicates p-value of t-test.

Table 5.A.6: Balance table: administrative records non-Covid occupations only

	Control (N=5074)		Treatment (N=10243)		Diff.	p
	Mean	Std. Dev.	Mean	Std. Dev.		
Male	0.27	0.44	0.27	0.44		0.90
Unemployment duration (wks.)	34.41	29.07	34.45	29.63	0.05	0.92
Remaining benefits (wks.)	49.78	28.94	49.83	29.00	0.05	0.92
Covid selection occ.						
Selection Occupation:						
Activity counsellor	0.07	0.25	0.07	0.25		0.81
Archivist	0.02	0.14	0.02	0.14		0.69
Video and sound technician	0.02	0.14	0.02	0.14		0.78
Janitor/Concierge	0.06	0.23	0.06	0.23		0.77
Animal caretaker	0.03	0.16	0.03	0.16		0.93
Printer	0.02	0.14	0.02	0.14		0.94
Graphic designer	0.05	0.23	0.05	0.22		0.68
Office support staff	0.42	0.49	0.42	0.49		0.99
Primary school teaching assistant	0.04	0.20	0.04	0.21		0.58
Event/conference organizer	0.05	0.22	0.05	0.21		0.84
Producer (television/film)	0.02	0.13	0.02	0.13		0.89
Social worker	0.16	0.37	0.16	0.36		0.90
Steward/stewardess	0.02	0.14	0.02	0.14		0.86
Shop attendant household/leisure	0.03	0.16	0.03	0.16		0.88

Note: based on administrative data from individuals with a non-Covid primary occupation. Remaining benefits and unemployment duration are measured in March 2021. *Difference* = Treatment - Control. Column *p* indicates p-value of t-test.

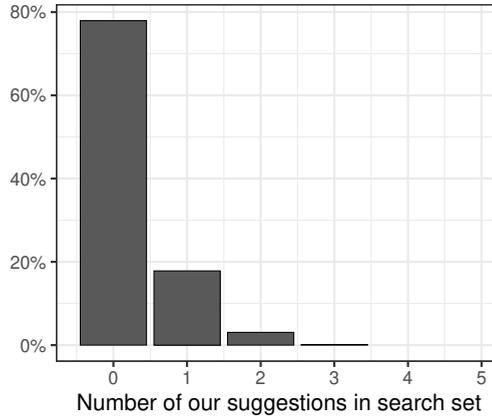
Table 5.A.7: Balance table: survey responses non-Covid occupations only

	Control (N=447)		Treatment (N=921)		Diff.	p
	Mean	Std. Dev.	Mean	Std. Dev.		
Unempl. duration	34.84	32.72	33.27	29.55	-1.57	0.39
Male	0.28	0.45	0.27	0.44	-0.01	0.57
Job finding score sel. occ.	3.01	0.59	2.99	0.54	-0.01	0.69
Time exploring alternatives	5.58	5.97	5.61	5.93	0.02	0.95
Willingness work in new occ.	3.45	0.88	3.41	0.86	-0.04	0.44
My skills are transferable	3.88	0.81	3.83	0.80	-0.05	0.27
Prob. job in 2 months	0.40	0.31	0.39	0.29	-0.01	0.65
Appl. needed (primary)	51.69	61.44	42.73	54.77	-8.95	0.03
Appl. needed (alt.)	50.60	64.09	39.39	51.36	-11.21	0.02
Salary previous job	2988.66	1212.55	2991.31	1199.65	2.65	0.97
Hours previous job	29.37	8.49	29.32	8.16	-0.06	0.90
Expected wage (main occ.)	2903.57	900.01	2918.46	929.40	14.89	0.78
Reservation wage (main occ.)	2772.45	848.93	2811.36	907.74	38.91	0.45
Expected wage (alt. occ.)	2866.69	899.44	2923.22	1005.29	56.53	0.37
Reservation wage (alt. occ.)	2716.34	828.04	2832.96	969.60	116.62	0.05
Applications (main occ.)	3.40	8.35	3.15	6.63	-0.25	0.60
Job interviews (main occ.)	0.36	1.07	0.44	1.58	0.08	0.29
Applications (alt. occ.)	2.86	6.51	2.23	3.68	-0.64	0.11
Job interviews (alt. occ.)	0.34	0.96	0.39	1.11	0.05	0.48
Applications (other occ.)	2.45	5.13	2.36	4.75	-0.09	0.80
Job interviews (other occ.)	0.37	1.16	0.34	1.05	-0.03	0.74

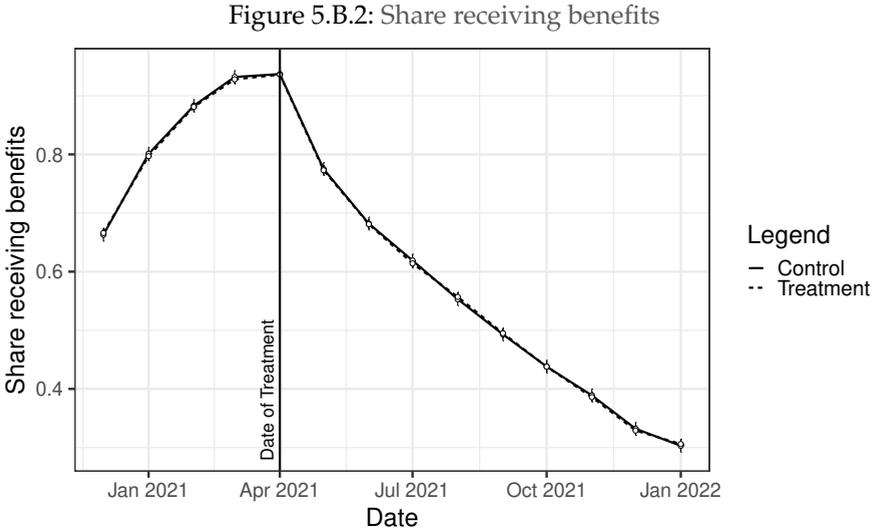
Note: based on survey data collected before intervention. Only includes survey respondents from non-Covid occupations. *Difference* = Treatment - Control. Column *p* indicates p-value of t-test.

Appendix 5.B Additional Figures

Figure 5.B.1: Number of suggestions initially in search set



Note: based on survey data collected before the interventions. Occupations mentioned in the survey were coded by hand and compared to occupational codes of provided suggestions.



Note: based on administrative data of the full experimental sample (of individuals with a non-Covid selection occupation). Dots indicate estimates, with vertical bars displaying the 95% confidence interval.

6

Conclusion

The studies in this thesis focus on the role of information in shaping individuals' beliefs and choices. I show that information matters. In **Chapters 2 and 4**, I provide evidence that individuals acquire and use the information that public institutions provide. **Chapter 2** shows that student satisfaction scores in one year matter for enrollment of first-year students in the next. Not only the satisfaction score of the program itself matters, those of its substitutes as well. Analyses exploiting rounding discontinuities show that a satisfaction score being rounded up to the next tenth increases first-year enrollment in the subsequent year by 1.70% to 3.52%, although the estimates are imprecise. How a program's score compares to that of its closest substitutes is of particular importance: conditional on being rounded up, a program that has a (slightly) higher published satisfaction score than at least one of

its top substitutes will see an increase in first-year enrollment in the subsequent year of up to 4.37% on average.

Chapter 4 shows that individuals acquire information about macroeconomic conditions in their day-to-day lives and relate this to their personal situation. When a Jobs Report is estimated to increase beliefs about the likelihood of the unemployment rate increasing by 1 percentage point, beliefs about the likelihood of personal job loss during that time increase by up to 0.22 percentage points. It not only affects individuals' expected likelihood of losing their own jobs, but also their expectations about the likelihood of being able to find a new job conditional on losing theirs. News that people interpret as increasing the likelihood of the unemployment rate increasing thus makes individuals more pessimistic about their employment prospects through multiple channels.

In **Chapters 3** and **5**, I present the results of two field experiments, and show that information can also be used as a tool to help students and job seekers make decisions. **Chapter 3** shows that students' prior beliefs about the labor market prospects of occupations they are interested in are highly inaccurate. Students overestimate both job opportunities and hourly wages of occupations that they like. The Chapter further shows that providing information about these prospects is effective in correcting belief errors in the short term. Survey data shows that these beliefs stick for at least a couple of months, but only for the job opportunities. Students who receive information are (between 0.88 and 2.16 percentage points) more likely to change their favorite occupation and, if they do so, switch towards occupations within their preference set that have better labor market prospects. I am unable

to confirm whether this change in preferences holds in the long term, however. Even though I do not see very strong effects on stated beliefs and preferences in the long term, I do see that students who received information enroll in profiles associated with occupations that have better labor market prospects (1.5% and 0.3% – €0.05 an hour – higher than the control group mean for job opportunities and wages, respectively). Similarly, students who received information about the hourly wages of occupations do enroll in post-secondary education programs that have better earnings prospects (2.5% – approximately €0.40 an hour – higher than the control group mean).

Chapter 5 shows that while many job seekers report to be willing to explore alternative occupations and are confident about their ability to work in a new occupation that matches their skillset, their actual job search is fairly narrow. Job seekers tend to look for work in occupations they have experience in. Moreover, beliefs about job offer rates show that awareness of the large variation in labor market prospects across occupations is very limited. One might thus expect that the job seekers are receptive to the information I provide them with. I do not find that the interventions had any impact on benefits receipt or labor earnings up to at least eight months after the treatment, however. I do find that those in the treatment group are 1.79 percentage points more likely to have found a job seven months after the intervention, and are between 5 and 6 percentage points more likely to have done so in an occupation different from their initial occupation of interest. Whether the difference in occupations of new jobs leads to more job stability, higher earnings and lower benefit dependence in the long-run remains to be studied in the future, when additional data become available. Survey data collected two months after the experiment does

not show any impact of the treatments on beliefs or search behavior, but the sample size is limited.

Although the contexts are different, it is worthwhile to consider how the results of the different studies in this thesis relate to each other. **Chapters 3 and 4** show a clear impact of information on beliefs in the short term. Looking at the longer term impacts, **Chapter 3** shows that while beliefs about job opportunities persist for up to seven months after the intervention, this is not true for hourly wages. In **Chapter 5**, I find no evidence of changes in beliefs. The lack of a strong impact on beliefs is hard to square with the finding that both Chapters do show an impact on behavior: the study program students enroll in, and occupations job seekers find work in, respectively. However, the samples of the survey data on beliefs collected for both field experiments are small.

Chapters 2, 3 and 5 all show that information affects choices. While the different effect sizes are hard to compare, the effect of student satisfaction scores seems large; particularly considering the fact that I mostly study the impact of minor changes in these scores. Contrarily, I provide individuals with information about sizeable differences in labor market prospects of occupations in the the field experiments. A potential explanation for the particularly large effects in **Chapter 2** is that prospective students can access Studiekeuze123.nl whenever they want. Students and job seekers that were part of the field experiments could theoretically login to the Qompas platform, or search their inbox for the information we provided. However, any prospective student that uses Google to search for a Dutch study program will most likely end up on Studiekeuze123.nl; at that moment, they are actively looking

for information, and may thus be much more receptive to it. Relatedly, one of the most striking findings from **Chapter 3** is that conditional on changing their favorite occupation, treated students move to occupations with much better labor market prospects. For the control group, the prospects of the occupations they switch to are not at all better than those of the occupations they switch from. It demonstrates that when individuals who are receptive to information receive it, it can have a large impact on their beliefs and choices. However, many individuals are not. A potentially valuable question for future research is therefore: what determines whether an individual is receptive to information? Of particular interest are the importance of targeting, timing, and presentation.

Another promising avenue for future (experimental) research is to go beyond information provision, and instead investigate how exactly individuals process information. How much attention do individuals pay to information they receive? Are they able to assess what figures about job finding rates and earnings tell them about their day-to-day lives in the (distant) future? Not just in terms of the likelihood of unemployment or earnings, but more broadly: what they will be able to consume, where they will live, how happy they will be, and so on. While the studies in this thesis show that information matters for beliefs and choices, these questions could shed further light on the mechanisms driving these results, and potentially guide future information provision endeavours.

Bibliography

- Alstadsæter, A. (2011). Measuring the consumption value of higher education. *CESifo Economic Studies*, 57(3), 458–479.
- Altmann, S., Falk, A., Jäger, S., & Zimmermann, F. (2018). Learning about job search: A field experiment with job seekers in Germany. *Journal of Public Economics*, 164, 33–49.
- Altonji, J. G., Arcidiacono, P., & Maurel, A. (2016). The analysis of field choice in college and graduate school: Determinants and wage effects. In *Handbook of the economics of education* (Vol. 5, pp. 305–396). Elsevier.
- Altonji, J. G., Blom, E., & Meghir, C. (2012). Heterogeneity in human capital investments: High school curriculum, college major, and careers. *Annu. Rev. Econ.*, 4(1), 185–223.
- Arcidiacono, P., Hotz, V. J., & Kang, S. (2012). Modeling college major choices using elicited measures of expectations and counterfactuals. *Journal of Econometrics*, 166(1), 3–16.
- Armantier, O., Topa, G., Van der Klaauw, W., & Zafar, B. (2017). An overview of the survey of consumer expectations. *Economic Policy Review*, 23(2), 51–72.
- Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of economic perspectives*, 29(3), 3–30.
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279–1333.
- Baker, R., Bettinger, E., Jacob, B., & Marinescu, I. (2018). The effect of labor market information on community college students' major

- choice. *Economics of Education Review*, 65, 18–30.
- Bean, J. P., & Bradley, R. K. (1986). Untangling the satisfaction-performance relationship for college students. *The Journal of Higher Education*, 57(4), 393–412.
- Belot, M., Kircher, P., & Muller, P. (2019). Providing advice to job-seekers at low cost: An experimental study on online advice. *The review of economic studies*, 86(4), 1411–1447.
- Betts, J. R. (1996). What do students know about wages? evidence from a survey of undergraduates. *Journal of human resources*, 31(1), 27–56.
- Bleemer, Z., & Mehta, A. (2022). Will studying economics make you rich? a regression discontinuity analysis of the returns to college major. *American Economic Journal: Applied Economics*, 14(2), 1–22.
- Bleemer, Z., & Zafar, B. (2018). Intended college attendance: Evidence from an experiment on college returns and costs. *Journal of Public Economics*, 157, 184–211.
- Bonilla-Mejía, L., Bottan, N. L., & Ham, A. (2019). Information policies and higher education choices experimental evidence from colombia. *Journal of Behavioral and Experimental Economics*, 83, 101468.
- Bottone, M., & Rosolia, A. (2019). Monetary policy, firms' inflation expectations and prices: causal evidence from firm-level data. *Bank of Italy Temi di Discussione (Working Paper) No, 1218*.
- Brewer, D. J., Eide, E. R., & Ehrenberg, R. (1999). Does it pay to attend an elite private college? cross-cohort evidence on the effects of college type on earnings. *Journal of Human Resources*, 34(1), 104–123.
- Britton, J., van der Erve, L., Belfield, C., Vignoles, A., Dickson, M., Zhu, Y., . . . Buscha, F. (2021). *How much does degree choice matter?* (Working Paper No. 21/24). IFS.

- Brown, S., & Taylor, K. (2006). Financial expectations, consumption and saving: a microeconomic analysis. *Fiscal Studies*, 27(3), 313–338.
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? In *Aea papers and proceedings* (Vol. 108, pp. 43–47).
- Bulligan, G. (2018). The effect of the eurosystem's expanded asset purchase programme on inflation expectations: evidence from the ecb survey of professional forecasters. *Bank of Italy Occasional Paper*(455).
- Caliendo, M., Cobb-Clark, D. A., & Uhlendorff, A. (2015, 03). Locus of Control and Job Search Strategies. *The Review of Economics and Statistics*, 97(1), 88–103.
- Campbell, D., Carruth, A., Dickerson, A., & Green, F. (2007). Job insecurity and wages. *The Economic Journal*, 117(518), 544–566.
- Card, D., Kluve, J., & Weber, A. (2018, 10). What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations. *Journal of the European Economic Association*, 16(3), 894–931.
- Caroli, E., & Godard, M. (2016). Does job insecurity deteriorate health? *Health economics*, 25(2), 131–147.
- Casselmann, B. (2018). Making sense of the jobs report: It's not always easy. *The New York Times*. Retrieved 2022-04-17, from <https://www.nytimes.com/2018/02/01/insider/insider-jobs-report.html>
- Chan, K. F., & Gray, P. (2018). Volatility jumps and macroeconomic news announcements. *Journal of Futures Markets*, 38(8), 881–897.
- Chu, A., & Westerheijden, D. F. (2018). Between quality and control: what can we learn from higher education quality assurance policy in the netherlands. *Quality in higher education*, 24(3), 260–270.
- Clark, A. E., Cotofan, M., & Layard, R. (2022). *The full returns to*

- the choice of occupation and education* (IZA Discussion Papers No. 15279). Institute for the Study of Labor (IZA). Retrieved from <https://docs.iza.org/dp15279.pdf>
- Conlon, J. J. (2019). Major malfunction: A field experiment correcting undergraduates' beliefs about salaries. *Journal of Human Resources*, 56(3), 922-939.
- Cotofan, M., Cassar, L., Dur, R., & Meier, S. (2020). Macroeconomic conditions when young shape job preferences for life. *The Review of Economics and Statistics*, 1-20.
- De Fiore, F., Lombardi, M. J., & Schuffels, J. (2021). *Are households indifferent to monetary policy announcements?* (Tech. Rep.). Bank for International Settlements.
- Del Carpio, L., & Guadalupe, M. (2021). More women in tech? evidence from a field experiment addressing social identity. *Management Science*, 68(5), 3196-3218.
- del Rio-Chanona, R. M., Mealy, P., Pichler, A., Lafond, F., & Farmer, J. D. (2020). Supply and demand shocks in the covid-19 pandemic: An industry and occupation perspective. *Oxford Review of Economic Policy*, 36(Supplement_1), S94-S137.
- Den Haan, W. J., Rendahl, P., & Riegler, M. (2018). Unemployment (fears) and deflationary spirals. *Journal of the European Economic Association*, 16(5), 1281-1349.
- Dickerson, A., & Green, F. (2012). Fears and realisations of employment insecurity. *Labour Economics*, 19(2), 198-210.
- Douglas, J. A., Douglas, A., McClelland, R. J., & Davies, J. (2015). Understanding student satisfaction and dissatisfaction: an interpretive study in the uk higher education context. *Studies in Higher Education*, 40(2), 329-349.
- Dur, R., & Van Lent, M. (2019). Socially useless jobs. *Industrial Relations: A Journal of Economy and Society*, 58(1), 3-16.

- Dutch Inspectorate of Education. (2020). *De staat van het onderwijs 2020* (Tech. Rep.).
- Faberman, R. J., & Kudlyak, M. (2019). The intensity of job search and search duration. *American Economic Journal: Macroeconomics*, 11(3), 327–57.
- Fleming, M. J., & Remolona, E. M. (1997). What moves the bond market? *Economic policy review*, 3(4).
- Forsythe, E., Kahn, L. B., Lange, F., & Wiczer, D. (2020). Labor demand in the time of covid-19: Evidence from vacancy postings and ui claims. *Journal of public economics*, 189, 104238.
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological forecasting and social change*, 114, 254–280.
- Geishecker, I., Riedl, M., & Frijters, P. (2012). Offshoring and job loss fears: An econometric analysis of individual perceptions. *Labour Economics*, 19(5), 738–747.
- Gibson, A. (2010). Measuring business student satisfaction: a review and summary of the major predictors. *Journal of Higher Education Policy and Management*, 32(3), 251–259.
- Government of the Netherlands. (n.d.-a). *Pre-vocational secondary education (vmbo)*. Retrieved 2019-08-06, from <https://www.government.nl/topics/secondary-education/pre-vocational-secondary-education-vmbo>
- Government of the Netherlands. (n.d.-b). *Wat zijn de vooropleidingseisen voor het middelbaar beroepsonderwijs (mbo)?* Retrieved 2019-08-11, from <https://www.rijksoverheid.nl/onderwerpen/middelbaar-beroepsonderwijs/vraag-en-antwoord/wat-zijn-de-vooropleidingseisen-voor-het-middelbaar-beroepsonderwijs-mbo>
- Graham, M., Nikkinen, J., & Sahlström, P. (2003). Relative importance

- of scheduled macroeconomic news for stock market investors. *Journal of Economics and Finance*, 27(2), 153–165.
- Green, T. C. (2004). Economic news and the impact of trading on bond prices. *The Journal of Finance*, 59(3), 1201–1233.
- Griffith, A., & Rask, K. (2007). The influence of the us news and world report collegiate rankings on the matriculation decision of high-ability students: 1995-2004. *Economics of Education Review*, 26, 244–255.
- Harris, E. S., & Zabka, N. M. (1995). The employment report and the dollar. *Current Issues in Economics and Finance*, 1(8).
- Hastings, J., Neilson, C. A., Ramirez, A., & Zimmerman, S. D. (2016). (un) informed college and major choice: Evidence from linked survey and administrative data. *Economics of Education Review*, 51, 136–151.
- Hastings, J., Neilson, C. A., & Zimmerman, S. D. (2013). *Are some degrees worth more than others? evidence from college admission cutoffs in chile* (Tech. Rep.). National Bureau of Economic Research.
- Hastings, J., Neilson, C. A., & Zimmerman, S. D. (2015). *The effects of earnings disclosure on college enrollment decisions* (Working Paper No. 21300). National Bureau of Economic Research. Retrieved from <http://www.nber.org/papers/w21300>
- Hazelkorn, E., Loukkola, T., & Zhang, T. (2014). *Rankings in institutional strategies and processes: Impact of illusion?* (Tech. Rep.). European University Association.
- Hendren, N. (2017). Knowledge of future job loss and implications for unemployment insurance. *American Economic Review*, 107(7), 1778–1823.
- Horstschräer, J. (2012). University rankings in action? the importance of rankings and an excellence competition for university choice of high-ability students. *Economics of Education Review*, 31, 1162–

1176.

- Inspectorate of Education. (2022). *Staat van het onderwijs 2022*.
- Jäger, S., Roth, C., Roussille, N., & Schoefer, B. (2022). *Worker beliefs about outside options* (Tech. Rep.). National Bureau of Economic Research.
- Jensen, R. (2010). The (perceived) returns to education and the demand for schooling. *The Quarterly Journal of Economics*, 125(2), 515–548.
- Kaplan, G., & Schulhofer-Wohl, S. (2018). The changing (dis-) utility of work. *Journal of Economic Perspectives*, 32(3), 239–58.
- Kerr, S. P., Pekkarinen, T., Sarvimäki, M., & RoopeUusitalo. (2020). Post-secondary education and information on labor market prospects: A randomized field experiment. *Labour Economics*, 33.
- Ketel, N., Leuven, E., Oosterbeek, H., & van der Klaauw, B. (2016). The returns to medical school: Evidence from admission lotteries. *American Economic Journal: Applied Economics*, 8(2), 225–54.
- Kirkeboen, L. J., Leuven, E., & Mogstad, M. (2016). Field of study, earnings, and self-selection. *The Quarterly Journal of Economics*, 131(3), 1057–1111.
- Kłopocka, A. M. (2017). Does consumer confidence forecast household saving and borrowing behavior? evidence for poland. *Social Indicators Research*, 133(2), 693–717.
- Kondo, A., & Shoji, M. (2019). Peer effects in employment status: Evidence from housing lotteries. *Journal of Urban Economics*, 113, 103195.
- Kuchler, T., & Zafar, B. (2019). Personal experiences and expectations about aggregate outcomes. *The Journal of Finance*, 74(5), 2491–2542.
- Lafortune, J. (2013). Making yourself attractive: Pre-marital investments and the returns to education in the marriage market. *Amer-*

- ican Economic Journal: Applied Economics*, 5(2), 151–78.
- Lamla, M. J., & Vinogradov, D. V. (2019). Central bank announcements: Big news for little people? *Journal of Monetary Economics*, 108, 21–38.
- Lergetporer, P., Werner, K., & Woessmann, L. (2018). *Does ignorance of economic returns and costs explain the educational aspiration gap? evidence from representative survey experiments* (IZA Discussion Papers No. 11453). Institute for the Study of Labor (IZA). Retrieved from <https://EconPapers.repec.org/RePEc:iza:izadps:dp11453>
- Machin, S., Marie, O., & Vujić, S. (2011). The crime reducing effect of education. *The Economic Journal*, 121(552), 463–484.
- Malmendier, U., & Nagel, S. (2011). Depression babies: do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics*, 126(1), 373–416.
- Malmendier, U., & Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1), 53–87.
- Mandal, B., Ayyagari, P., & Gallo, W. T. (2011). Job loss and depression: The role of subjective expectations. *Social Science & Medicine*, 72(4), 576–583.
- McGuigan, M., McNally, S., & Wyness, G. (2016). Student awareness of costs and benefits of educational decisions: effects of an information campaign. *Journal of Human Capital*, 10(4), 482–519.
- Mertens, K., Lewis, D. J., & Makridis, C. (2020). *Do monetary policy announcements shift household expectations?* (CEPR Discussion Paper No. 14360). Centre for Economic Policy Research. Retrieved from <https://repec.cepr.org/repec/cpr/ceprdp/DP14360.pdf>
- Ministry of Education, Culture and Science. (2010). *Staatsblad 2010*, 314.

- Ministry of Education, Culture and Science. (2018). *Staatsblad 2018*, 226.
- Mueller, A., Spinnewijn, J., & Topa, G. (2021). Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias. *American Economic Review*, 111(1), 324–63.
- Mueller, R. E., & Rockerbie, D. (2005). Determining demand for university education in ontario by type of student. *Economics of Education Review*, 24, 469–483.
- Nedelkoska, L., & Quintini, G. (2018). Automation, skills use and training. (202). Retrieved from <https://www.oecd-ilibrary.org/content/paper/2e2f4eea-en> doi: <https://doi.org/10.1787/2e2f4eea-en>
- Nuffic. (2019). *The education system of the netherlands*. Retrieved 2019-08-04, from <https://www.nuffic.nl/documents/459/education-system-the-netherlands.pdf>
- Oosterbeek, H., & Van Ophem, H. (2000). Schooling choices: Preferences, discount rates, and rates of return. *Empirical Economics*, 25(1), 15–34.
- Oreopoulos, P., & Dunn, R. (2013). Information and college access: Evidence from a randomized field experiment. *The Scandinavian Journal of Economics*, 115(1), 3–26.
- Oreopoulos, P., & Salvanes, K. G. (2011). Priceless: The nonpecuniary benefits of schooling. *Journal of Economic perspectives*, 25(1), 159–84.
- Pekkala Kerr, S., Pekkarinen, T., Sarvimäki, M., & Uusitalo, R. (2015). *Post-secondary education and information on labor market prospects: A randomized field experiment* (IZA Discussion Papers No. 9372). Institute for the Study of Labor (IZA). Retrieved from <https://EconPapers.repec.org/RePEc:iza:izadps:dp9372>
- Peter, F. H., & Zambre, V. (2017). Intended college enrollment and

- educational inequality: Do students lack information? *Economics of Education Review*, 60, 125–141.
- Pettinicchi, Y., & Vellekoop, N. (2019). Job loss expectations, durable consumption and household finances: Evidence from linked survey data.
- Porter, C., & Serra, D. (2020). Gender differences in the choice of major: The importance of female role models. *American Economic Journal: Applied Economics*, 12(3), 226–54.
- Rangel, J. G. (2011). Macroeconomic news, announcements, and stock market jump intensity dynamics. *Journal of Banking & Finance*, 35(5), 1263–1276.
- Rauhvargers, A. (2013). *Global university rankings and their impact: Report ii* (Tech. Rep.). European University Association.
- Ravn, M. O., & Sterk, V. (2017). Job uncertainty and deep recessions. *Journal of Monetary Economics*, 90, 125–141.
- Reuben, E., Wiswall, M., & Zafar, B. (2017). Preferences and biases in educational choices and labour market expectations: Shrinking the black box of gender. *The Economic Journal*, 127(604), 2153–2186.
- Riley, E. (2022, 01). Role Models in Movies: The Impact of Queen of Katwe on Students' Educational Attainment. *The Review of Economics and Statistics*, 1–48.
- Roth, C., Settele, S., & Wohlfart, J. (2022). Risk exposure and acquisition of macroeconomic information. *American Economic Review: Insights*, 4(1), 34–53.
- Roth, C., & Wohlfart, J. (2020). How do expectations about the macroeconomy affect personal expectations and behavior? *Review of Economics and Statistics*, 102(4), 731–748.
- Sa, C., Florax, R. J., & Rietveld, P. (2004). Determinants of the regional demand for higher education in the netherlands: A grav-

- ity model approach. *Regional Studies*, 38(4), 375–392.
- Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association*, 13(1), 130–167.
- Stephens Jr, M. (2004). Job loss expectations, realizations, and household consumption behavior. *Review of Economics and statistics*, 86(1), 253–269.
- Times Higher Education. (2022). *World university rankings*. Retrieved from <https://www.timeshighereducation.com/world-university-rankings/2022/world-ranking>
- Uitvoeringsinstituut Werknemersverzekeringen (UWV). (2014, Aug). *Ik word werkloos - uwv - particulieren*. Retrieved 2022-05-06, from <https://www.uwv.nl/particulieren/werkloos/ik-word-werkloos/detail/kan-ik-een-ww-uitkering-krijgen>
- Uitvoeringsinstituut Werknemersverzekeringen (UWV). (2015, Feb). *About us - executive board - organization: Uwv: Over uwv*. Retrieved 2022-05-06, from <https://www.uwv.nl/overuwv/english/about-us-executive-board-organization/detail/organization/public-employment-service>
- U.S. Department of Labor. (2022). *Employment and earnings by occupation*. Retrieved 2022-05-24, from <https://www.dol.gov/agencies/wb/data/occupations>
- Veerman, C., Berdahl, R., Bormans, M., Geven, K., Hazelkorn, E., Kan, A. R., ... Vossensteyn, J. (2010). *Threefold differentiation: recommendations of the committee on the future sustainability of the dutch higher education system*. Koninklijke Broese en Peereboom.
- Wiederholt, M., et al. (2010). Rational inattention. *The New Palgrave Dictionary of Economics (Online Edition ed.)*.
- Wiswall, M., & Zafar, B. (2017). Preference for the workplace, invest-

- ment in human capital, and gender. *The Quarterly Journal of Economics*, 133(1), 457–507.
- Zafar, B. (2011). How do college students form expectations? *Journal of Labor Economics*, 29(2), 301–348.
- Zafar, B. (2013). College major choice and the gender gap. *Journal of Human Resources*, 48(3), 545–595.

Impact

The main objective of the research described in this thesis is to better understand how information affects beliefs and choices related to education and work. This is an important topic, as the choices of what degree to obtain, and what job to apply to have potentially life-altering consequences. The studies in this thesis show that information indeed matters.

Chapter 2 shows that student satisfaction scores matter for enrollment decisions of first-year students. When a program's satisfaction score increases, so does first-year enrollment in that program. First-year enrollment decreases when the satisfaction score of a close substitute of the program goes up. These findings underline the importance of providing readily accessible information to students; something universities and public authorities invest a lot of resources in.

Chapter 3 shows that students generally overestimate the job opportunities and hourly wages of their favorite occupations. Providing information is effective in correcting these overestimations, particularly in the short term. The information also affects students' preferences and choices. Students who receive information are more likely to change their favorite occupation and, if they do so, switch towards occupations with better labor market prospects. However, I cannot confirm that this preference change persists in the long term. Students in schools that received information do more often enroll in profiles associated with occupations that have better labor market prospects and post-secondary education programs that have better earnings prospects.

The direct implication of these findings is that there is scope for making relevant information more accessible and salient to students; for instance, by integrating it in career guidance counseling efforts at schools.

Chapter 4 shows that individuals acquire information about macroeconomic conditions in their day-to-day lives and relate this to their personal situation. Whenever a Jobs Report is published that makes people think it is more likely that the unemployment rate will increase, people also think it is more likely that they will lose their own job. Moreover, it makes them more pessimistic about the likelihood of being able to find a new job conditional on losing theirs. News that people interpret as increasing the likelihood of the unemployment rate increasing thus makes individuals more pessimistic about their employment prospects through multiple channels. It is important to understand, as expectations about the likelihood of job loss affect individuals' behavior, and (mental) health. Moreover, it could potentially lead to a *deflationary spiral*; a situation where in turn consumption decreases, prices decrease, production decreases, wages decrease and therefore consumption decreases further. However, I do not find evidence that people expect their spending to change as the result of a Jobs Report.

Chapter 5 shows that while many job seekers looking for work in occupations that provide poor job prospects report to be willing to explore new alternatives and are confident about their ability to work in a new occupation that matches their skillset, actual job search is fairly narrow. Moreover, beliefs about job offer rates show that awareness of the large variation in labor market prospects across occupations is very limited. Despite that, providing job seekers with information on suitable alternatives that offer better prospects has no impact on benefits

receipt or labor earnings up to at least eight months after the treatment. It does seem to help people find work sooner, and in different types of occupations. This may be beneficial to the job seekers in the future. Once again, it shows that while a lot of information is out there; making an effort to provide personalized information through commonly used channels has an impact on decisions people make.

Aside from adding to our understanding of the role of information in beliefs and choices about education and work, the work described in this thesis has also led to output that can be put into practice easily. Qompas, the company we collaborated with for the study in **Chapter 3**, intends to implement our information intervention on their updated platform. Similarly, the employment office has expressed interest in using the findings from the study presented in **Chapter 5** to improve their service to job seekers. The methods of collecting and providing information developed in these Chapters are low-cost and easy to replicate. This allows similar institutions to employ them as well.

Summary

This thesis sheds light on how information affects beliefs and choices about education and work. It consists of four chapters written by the author of this thesis, in collaboration with several co-authors. The first two chapters cover educational choices. The two chapters thereafter are about work. On both topics, one chapter is based on observational data, and another on a field experiment. The two types of studies complement each other; the field experiments allow for cleaner identification of the causal impact of information provision, while the observational studies provide higher external validity. Together, they provide a comprehensive look into the impact information has on individuals' beliefs and choices about education and work.

Chapter 2 studies the impact of published student satisfaction scores (ranging from 1 to 5) on enrollment of first-year students for the near universe of existing higher education programs in the Netherlands between 2011 and 2019. To determine each programs' closest substitutes, the author uses pageview data from the largest Dutch educational information website. This allows for an analysis of the impact of changes in a program's own published student satisfaction score, but also the impact of changes in the student satisfaction scores of its substitutes. The author analyzes the impact of these satisfaction scores using fixed effects Poisson regressions and exploits rounding discontinuities to identify causal effects. On the whole, the findings show that student satisfaction scores matter for enrollment. An increase in a program's student satisfaction score leads to higher levels of enrollment, whereas

an increase in the student satisfaction scores of substitutes leads to lower levels of enrollment. Point estimates of the impact of a program's student satisfaction score being rounded up to the next tenth on first-year enrollment range between 1.70% and 3.52%, depending on the bandwidth around the threshold we consider. Conditional on being above the rounding threshold, a program being rounded up over at least one of its closest substitutes increases first-year enrollment by up to 4.37%.

Chapter 3 presents the results of a large-scale field experiment in which the author provides students at randomly selected schools with information about the job opportunities and hourly wages of a set of occupations they are interested in. The experiment takes place on an online career guidance counseling platform that is widely used in the Netherlands, and involves 28,267 pre-vocational secondary education students in 243 schools over a period of 2 years. The information improves the accuracy of students' beliefs, both in the short run (for job opportunities and hourly wages) and in the long run (for job opportunities only). Students who receive the information also change their favorite occupation 0.88 to 2.16 percentage points more often, and switch towards an occupation with better labor market prospects when they do so. Last, and most importantly, they select secondary school specializations related to occupations with better labor market prospects (1.5% and 0.3% – €0.05 an hour – higher than the control group mean for job opportunities and wages, respectively) and choose post-secondary education programs with higher expected wages (2.5% – approximately €0.40 an hour – higher than the control group mean).

In **Chapter 4**, the author uses data from the New York Federal Re-

serve's Survey of Consumer Expectations to study how the United States Bureau of Labor Statistics' Employment Situation Reports (Jobs Reports) affect individuals' expectations about the likelihood of losing their own job. This happens in two steps. First, the author estimates the information shocks of the jobs reports on expectations about the development of the national unemployment rate in the next twelve months by comparing survey responses shortly before and after publication of the reports. Second, the author estimates how these shocks affect individuals' expectations about losing their own job in the same time frame. The results show that when a report is estimated to increase beliefs about the likelihood of the unemployment rate increasing by 1 percentage point, beliefs about the likelihood of personal job loss during that time increase by up to 0.22 percentage points. The information shock further negatively affects individuals' beliefs about the likelihood of finding a new job if they were to lose their current one, but surprisingly has a positive effect on their beliefs about the likelihood of voluntarily leaving their job. The results are robust to the use of different bandwidths around the reports' publication dates and placebo treatments provide reassurance that the information shock is indeed the mechanism driving the result.

Chapter 5 studies the impact of online information provision to job seekers who are looking for work in occupations with relatively poor labor market prospects. The author provides the information through a personalized email containing suggestions about suitable alternative occupations and how the prospects of these alternatives compare to the job seekers' current occupation of interest. A second treatment adds a motivational video aimed at addressing the psychological hurdles of switching to a different occupation. The author evaluates the

interventions using a randomized field experiment with 30,129 unemployed job seekers, and acquires additional descriptive information on beliefs and job search. The results show no impact on received benefits and earnings in the first eight months after the treatment. The findings do show that treated individuals are 1.79 percentage points more likely to have found a job seven months after the intervention, although this difference decreases to 1.19 percentage points four months later. Moreover, treated individuals are between 5 and 6 percentage points more likely to have done so in an occupation different from their initial occupation of interest. This may be promising for their longer-term prospects.

To conclude, this thesis shows that information matters in important ways for beliefs and choices about education and work.

About the author

Bart de Koning was born on September 6, 1994 in Leiderdorp, the Netherlands. He received his bachelor's (2016) and master's degree (2017; cum laude) in Economics from Erasmus University Rotterdam. He subsequently started his doctoral studies at Maastricht University's school of business and economics under the supervision of Didier Fouarge and Robert Dur.



At Maastricht University, Bart taught a number of undergraduate courses on microeconomics and international economics. Between 2018 and 2021, he organized the school of business and economics' Ph.D. Brown Bag seminars. He presented his work at several international conferences and conducted two short-term research visits; to the University of Melbourne, Australia and IZA Bonn, Germany.

Bart will join Cornell University as a postdoctoral research associate in October of 2022.

ROA Dissertation Series

1. Lex Borghans (1993), *Educational Choice and Labour Market Information*, Maastricht, Research Centre for Education and the Labour Market.
2. Frank Cörvers (1999), *The Impact of Human Capital on International Competitiveness and Trade Performance of Manufacturing Sectors*, Maastricht, Research Centre for Education and the Labour Market.
3. Ben Kriechel (2003), *Heterogeneity Among Displaced Workers*, Maastricht, Research Centre for Education and the Labour Market.
4. Arnaud Dupuy (2004), *Assignment and Substitution in the Labour Market*, Maastricht, Research Centre for Education and the Labour Market.
5. Wendy Smits (2005), *The Quality of Apprenticeship Training, Conflicting Interests of Firms and Apprentices*, Maastricht, Research Centre for Education and the Labour Market.
6. Judith Semeijn (2005), *Academic Competences and Labour Market Entry: Studies Among Dutch Graduates*, Maastricht, Research Centre for Education and the Labour Market.
7. Jasper van Loo (2005), *Training, Labor Market Outcomes and Self-Management*, Maastricht, Research Centre for Education and the Labour Market.

8. Christoph Meng (2005), *Discipline-Specific or Academic? Acquisition, Role and Value of Higher Education Competencies*, Maastricht, Research Centre for Education and the Labour Market.
9. Andreas Ammermüller (2007), *Institutional Effects in the Production of Education: Evidence from European Schooling Systems*, Maastricht, Research Centre for Education and the Labour Market.
10. Bart Golsteyn (2007), *The Ability to Invest in Human Capital*, Maastricht, Research Centre for Education and the Labour Market.
11. Raymond Montizaan (2010), *Pension Rights, human capital development and well-being*, Maastricht, Research Centre for Education and the Labour Market.
12. Annemarie Nelen (2012), *Part-Time Employment and Human Capital Development*, Maastricht, Research Centre for Education and the Labour Market.
13. Jan Sauermann (2013), *Human Capital, Incentives, and Performance Outcomes*, Maastricht, Research Centre for Education and the Labour Market.
14. Harald Ulrich Pfeifer (2013), *Empirical Investigations of Costs and Benefits of Vocational Education and Training*, Maastricht, Research Centre for Education and the Labour Market.
15. Charlotte Büchner (2013), *Social Background, Educational Attainment and Labor Market Integration: An Exploration of Underlying Processes and Dynamics*, Maastricht, Research Centre for Education and the Labour Market.

16. Martin Humburg (2014), *Skills and the Employability of University Graduates*, Maastricht, Research Centre for Education and the Labour Market.
17. Jan Feld (2014), *Making the Invisible Visible, Essays on Overconfidence, Discrimination and Peer Effects*, Maastricht, Research Centre for Education and the Labour Market.
18. Olga Skriabikova (2014), *Preferences, Institutions, and Economic Outcomes: an Empirical Investigation*, Maastricht, Research Centre for Education and the Labour Market.
19. Gabriele Marconi (2015), *Higher Education in the National and Global Economy*, Maastricht, Research Centre for Education and the Labour Market.
20. Nicolas Salamanca Acosta (2015), *Economic Preferences and Financial Risk-Taking*, Maastricht, Research Centre for Education and the Labour Market.
21. Ahmed Elsayed Mohamed (2015), *Essays on Working Hours*, Maastricht, Research Centre for Education and the Labour Market.
22. Roxanne Amanda Korthals (2015), *Tracking Students in Secondary Education, Consequences for Student Performance and Inequality*, Maastricht, Research Centre for Education and the Labour Market.
23. Maria Zumbuehl (2015), *Economic Preferences and Attitudes: Origins, Behavioral Impact, Stability and Measurement*, Maastricht, Research Centre for Education and the Labour Market.

24. Anika Jansen (2016), *Firms' incentives to provide apprenticeships – Studies on expected short- and long-term benefits*, Maastricht, Research Centre for Education and the Labour Market.
25. Jos Maarten Arnold Frank Sanders (2016), *Sustaining the employability of the low skilled worker: Development, mobility and work redesign*, Maastricht, Research Centre for Education and the Labour Market.
26. Marion Collewet (2017), *Working hours: preferences, well-being and productivity*, Maastricht, Research Centre for Education and the Labour Market.
27. Tom Stolp (2018), *Sorting in the Labor Market: The Role of Risk Preference and Stress*, Maastricht, Research Centre for Education and the Labour Market.
28. Frauke Meyer (2019), *Individual motives for (re-)distribution*, Maastricht, Research Centre for Education and the Labour Market.
29. Maria Ferreira Sequeda (2019), *Human Capital Development at School and Work*, Maastricht, Research Centre for Education and the Labour Market.
30. Marie-Christine Martha Fregin (2019), *Skill Matching and Outcomes: New Cross-Country Evidence*, Maastricht, Research Centre for Education and the Labour Market.
31. Sanne Johanna Leontien van Wetten (2020), *Human capital and employee entrepreneurship: The role of skills, personality characteristics and the work context*, Maastricht, Research Centre for Education and the Labour Market.

32. Cécile Alice Jeanne Magnée (2020), *Playing the hand you're dealt, The effects of family structure on children's personality and the effects of educational policy on educational outcomes of migrant children*, Maastricht, Research Centre for Education and the Labour Market.
33. Merve Nezihe Özer (2020), *Essays on drivers and long-term impact of migration*, Maastricht, Research Centre for Education and the Labour Market.
34. Inge Ingeborg Henrica Maria Hooijen (2021), *Place attractiveness, A study of the determinants playing a role in residential settlement behaviour*, Maastricht, Research Centre for Education and the Labour Market.
35. Alexandra Marie Catherine de Gendre (2021), *Behavioral Barriers to Success in Education*, Maastricht, Research Centre for Education and the Labour Market.
36. Kim van Broekhoven (2021), *From creativity to innovation: Understanding and improving the evaluation and selection of ideas in educational settings*, Maastricht, Research Centre for Education and the Labour Market.
37. François M.B.M. Molin (2022), *Using digital formative assessments to improve learning in physics education*, Maastricht, Research Centre for Education and the Labour Market.
38. Bart Kasper de Koning (2022), *Empirical Studies on Information, Beliefs, and Choices in Education and Work*, Maastricht, Research Centre for Education and the Labour Market.