Information and Behaviour in Labour Markets

Author:
Eduard Krkoska

Supervisors:
Prof. Barbara Petrongolo
Prof. Johannes Abeler



University of Oxford

A thesis submitted in fulfillment of the requirements
for the degree of *Doctor of Philosophy in Economics*

January 2025

Declaration of Authorship

I, Eduard Krkoska, declare the existence of joint work in this thesis. The first chapter of this thesis is co-authored with Dr. Christian Johannes Meyer, Dr. Ashley Pople and Ms. Krittika Ray. The second and third chapters of this thesis are written by me alone. The third chapter is based partly on work completed during my MPhil in Economics.

Eduard Krkoska, January 2025

Information and Behaviour in Labour Markets

by Eduard Krkoska

Abstract

This thesis examines a number of questions in the field of labour economics. Each chapter addresses a separate economic question, with the theme bringing the papers together being the role of information in guiding the behaviour of labour market participants.

Chapter one experimentally tests the effects of introducing screening test requirements in the hiring process on the self-selection behaviour of jobseekers. Screening tests are commonly thought of as a mechanism for firms to obtain verifiable information on the quality of job candidates. We model these tests as a (cognitive and time) cost which jobseekers bear when applying for vacancies and show that the introduction of screening tests decreases the size of the applicant pool while increasing the quality of those who self-select into applying. In doing so, screening tests provide greater value than the information generated in the test results alone and can reduce wasted effort on both sides of the market by inducing poor quality matches to self-select out of the applicant pool.

Chapter two adds to the body of literature on income tax evasion under systems of third-party (employer) reporting. In this setting, tax evasion requires collusion between the employee and employer to hide the true salary information, which is difficult to verify, from the tax authorities. I provide evidence from major reforms in Romania which strongly indicates that there is widespread income under-reporting in the construction sector. The data also suggests that men are substantially more likely than women to under report their income.

Chapter three studies the evolution of group differences in labour markets over time, in particular addressing the role of information in driving the dynamics of the differences. Previous performance in the labour market acts as a signal of worker ability, but interpreting the successes or failures of those who might have been discriminated against is not straightforward. Depending on the information that employers have on the discriminatory preferences of *other* employers, group differences may be exacerbated over time, converge, or even reverse.

Word Count: Approximately 80,884

Acknowledgements

I would not have been able to complete my PhD without the generosity of so many people.

Firstly, I thank my supervisors. Prof. Barbara Petrongolo was my primary supervisor and Prof. Johannes Abeler was my secondary supervisor. Both have been extremely generous with their time and I appreciate that their different supervision styles were complementary to one another and jointly instrumental to completing this thesis. I am particularly grateful for their patience with my many questions and requests, and for their reassurances whenever I have felt doubtful. My third chapter builds on work completed during my MPhil, and I thank Prof. Willemien Kets for her supervision and support as my MPhil thesis supervisor.

For providing me funding, both through my MPhil in Economics as well as the DPhil, I thank the generosity of the Department of Economics and University College. I have benefitted from the Oxford-Chellgren graduate scholarship, as well as the Annual Bursaries provided by the Department of Economics.

Alongside my supervisors, a number of other academics at Oxford have been generous with their time and provided me with advice and feedback on my work. These include: Abi Adams, Alison Andrew, Sarah Clifford, Ian Crawford, Hamish Low, Sophocles Mavroeidis, Romuald Meango, Christian Meyer, Margaret Meyer, Niklas Moneke and Simon Quinn. I completed my BA in Economics at the University of Manchester, and I thank the lecturers who inspired me and gave me the foundation from which to build my knowledge at Oxford. In particular, I thank Dr. Peter Backus and Prof. Klaus Reiner Schenk-Hoppe. For emphasizing the importance of education, I thank my parents.

I thank my fellow PhD students who have supported me and provided valuable feedback on my work – Abelardo, Max, Stan, Alex, Juliette, Hannah, Luke, Nick and Shihang. Mihai and Tzvetan, who left Oxford following their MPhil in Economics, have been a constant source of support. For feedback on my second paper in particular, Mihai was very generous with his time.

Support from Ioana Gavril and Madalina Avram, who helped me to access the necessary data and understand Romanian tax policy, was instrumental in writing the second paper.

For ensuring I had a life outside of economics, I thank Dylan, Cherry, Cameron, Stefan, Prithwiraj, Abdullah, Mike, Meg, Rachel and Malitha. Going on walks with my dog was at times a welcome distraction from work, although I would often end up thinking about research, and sometimes generate ideas which were material to my thesis – thank you Brownie.

Finally, I thank my wife, Amy, for everything. I dedicate this thesis to you.

Contents

Declaration of Authorship iii

Abstract v

Acknowledgements vii

Introduction to the Thesis 1

1 Search Frictions and Self-Selection in Job Matching 5

1.1 Introduction . . . 6
1.2 Theoretical Framework . . . 9
1.3 Experimental Design . . . 12
1.3.1 The Online Labour Market . . . 12
1.3.2 External Validity . . . 13
1.3.3 Experimental set-up . . . 13
1.4 Empirical Strategy . . . 15
1.5 Results . . . 16
1.5.1 Introducing a Screening Task Improves Match Quality . . . 16
1.5.2 The Effects of Screening Task Length . . . 18
1.5.3 The Effects of the Content of the Screening Task . . . 18
1.5.4 Heterogeneity . . . 20
1.5.5 Applicant Welfare . . . 20
1.5.6 Selection on Observables . . . 21
1.6 Conclusion . . . 22

Appendix to Chapter 1 29

A.1 Experimental Details . . . 29
A.1.1 Measurement and Construction of Outcome Variables . . . 29
A.1.2 Experimental Design - Additional Information . . . 34
The Online Labour Market - Additional Information . . . 34
Recruitment of Participants . . . 34
Experiment Instructions . . . 36
A.1.3 Video Scripts . . . 37
A.2 Additional Results . . . 40
A.3 Additional Figures . . . 47

2 Tax Evasion, Employee-Employer Collusion, and Gender Differences **49**
2.1 Introduction . . . 50
2.2 Institutional Setting, Data and Identification . . . 53
2.2.1 Tax Evasion in Romania . . . 53
2.2.2 The Tax System and Policy Reform . . . 54
2.2.3 Data . . . 56
The Informal Sector . . . 57
2.3 Theory . . . 58
2.3.1 The Competitive Model . . . 58
2.3.2 The Monopsonistic Model . . . 59
Monopsony Model - Graphical Intuition . . . 60
2.3.3 Model Takeaways . . . 62
2.4 Empirical Analysis - Evidence of Tax Evasion . . . 62
2.4.1 Competition or Monopsony? . . . 62
2.4.2 Policy Effects on Employment . . . 64
2.4.3 Policy Take-Up . . . 65
2.4.4 Direct Evidence of Tax Evasion . . . 66
2.4.5 Individuals Report Above the Minimum Wage . . . 67
2.4.6 Joint Collusion . . . 69
2.4.7 Other Corruption in Romania . . . 72
2.4.8 Productivity Growth . . . 73
2.5 The Gender Gap in Tax Evasion . . . 73
2.5.1 Bargaining Power . . . 75
One (Wo)man bands . . . 75
2.5.2 Occupational Segregation . . . 77
Occupational Segregation: Implied Occupation by Hours Worked . . . 77
Office-Based Work: Implied by Industry Sector . . . 79
2.5.3 Firm Composition . . . 80
2.6 Conclusion . . . 82

Appendix to Chapter 2 **85**
B.1 Other Graphics . . . 85
B.1.1 Local-Industry Labour Market Flow Rates . . . 85
B.1.2 One-Person Firm Descriptive Graphs . . . 86
B.1.3 One-Person Firm Take-up . . . 87
B.1.4 Percent of Construction Employees Reporting Exactly the Minimum Wage . . . 88
B.2 Sample Details . . . 89
B.2.1 Treated NACE Code List: Female Share and Share of Construction Employment . . . 89
B.2.2 Wage Distributions by Sector and by Year . . . 90
B.2.3 Pre and Post-Reform Descriptive Statistics: by Gender and Sector . . . 91

B.2.4 Sample Details by Firm Composition . . . 93
B.2.5 Firm Size Distributions . . . 93
B.3 Other results . . . 94
B.3.1 Firm Industrial Code Switching . . . 94
B.3.2 Joint Collusion: Jumps to Reporting 3000 RON . . . 94
B.3.3 Joint Collusion: Discontinuity in One-Person Firms . . . 95
B.3.4 Joint Collusion: Jumps in Firms' Reporting . . . 96
B.3.5 Productivity Growth . . . 97
B.3.6 Manual Workers: Alternative Definition . . . 97
B.4 Standard Errors . . . 97

3 Group Differences in Labour Markets with Observational Learning and Information Imperfections 99
3.1 Introduction . . . 100
3.1.1 Related Literature . . . 101
3.2 The Simple Model . . . 104
3.2.1 Model setup . . . 104
3.3 Full Information Benchmark . . . 108
3.4 Naive Belief System . . . 109
3.4.1 Motivating the Belief System . . . 109
3.4.2 Ex-ante discrimination . . . 110
3.5 Partial Information . . . 112
3.5.1 Ex-ante discrimination . . . 113
3.6 Reduced-form Evidence . . . 116
3.6.1 Data . . . 116
3.6.2 Methodology . . . 116
3.6.3 Results . . . 117
3.7 The Full Model . . . 119
3.7.1 Model setup . . . 119
3.7.2 Full Information . . . 120
3.7.3 Period 1 . . . 120
Ex-ante Employment . . . 121
3.7.4 Period 2 . . . 121
Ex-ante Employment . . . 124
3.7.5 Naive Beliefs . . . 125
3.7.6 Ex-ante second period employment . . . 125
3.7.7 Partial Information . . . 127
3.7.8 Posterior Beliefs . . . 127
3.7.9 Period 2 . . . 128
Ex-ante Employment . . . 129
3.7.10 Implications . . . 130
3.8 Conclusion . . . 131

3.8.1 Policy Implications 132

Appendix to Chapter 3 **135**

C.1 Simple Model - Proofs 135

C.2 Full Model - Proofs 137

Conclusion to the Thesis **143**

Bibliography **147**

List of Figures

1.1 Cumulative Distribution of Job Performance of Control vs. Pooled Treatment 23
A1 Screening Task PDF Scan . . . 37
A2 Sequence Data Collection and Experiment . . . 47

2.1 Take Home Pay vs Reported Gross Salary: Pre-Reform . . . 55
2.2 Take Home Pay vs Reported Gross Salary: Post-Reform . . . 55
2.3 Competitive Model . . . 59
2.4 Monopsony Model: 'High Tax' case . . . 61
2.5 Monopsony Model: 'Low Tax' case . . . 61
2.6 Monopsony Model: Minimum Wages above Equilibrium Wage . . . 62
2.7 HHI: By Sector . . . 63
2.8 Share of Employment in the Construction Sector Over Time . . . 64
2.9 Take-Up of Policy . . . 65
2.10 Proportion of employees reporting a salary of exactly the minimum wage . . . 66
2.11 Event-Study - Reporting Above the Minimum Wage . . . 68
2.12 Event-Study - Reporting 3000 RON . . . 68
2.13 Event Study - Reporting 3000 RON by Pre-Reform Reported Salary . . . 69
2.14 2x2 DiD Estimates of Reporting Above the Minimum Wage by Firm Size . . . 70
2.15 Event Study - One-Person Firms . . . 71
2.16 One-Person Firms: Median Salaries and Firm Revenues . . . 71
2.17 Tax Evasion vs Hospital Bribery Rates . . . 73
2.18 Event Study - Reporting Above the Minimum Wage . . . 74
2.19 Event-Study - Reporting 3000 RON . . . 74
2.20 Event Study-Effects on One (Wo)man bands . . . 75
2.21 Event Study-Effects on Women by Boss's Gender . . . 76
2.22 Event Study-Effects on Men by Boss's Gender . . . 77
2.23 Event Study-Effects on Office-Based Workers . . . 78
2.24 Event Study-Effects on Manual-Based Workers . . . 79
2.25 Event Study - Project Development Industry Code . . . 80
2.26 Event Study - Female Dominated Firms . . . 81
2.27 Event Study - Mixed Gender Firms . . . 81
2.28 Event Study - Segregated Firms . . . 82
B1 Labour Market Flow Rates: By Sector . . . 85
B2 One-Person Firms: Mean Salaries and Firm Revenues . . . 86
B3 One-Person Firms: Median Salaries by Gender . . . 86

B4 One-Person Firms: Mean Salaries by Gender 86
B5 One-Person Firms: Median Firm Income by Gender 87
B6 One-Person Firms: Mean Firm Log Income by Gender 87
B7 One-Person Firms: Take-Up of Policy 87
B8 Proportion of Construction Employees Reporting Exactly the Minimum Wage 88
B9 Wage Distributions of the Construction Sector 90
B10 Wage Distributions of the Non-Construction Sector 90
B11 Firm Size Distributions by Sector 93
B12 Non-Construction to Constriction Classification Switches and Vice-Versa . . . 94
B13 2x2 DiD Estimates of Reporting 3000 RON by Firm Size 94
B14 Small Firms:2x2 DiD Estimates 95
B15 Small Firms: 2x2 DiD Estimates of Reporting Above Regular NMW by Gender 95
B16 Small Firms: 2x2 DiD Estimates of Reporting 3000 RON by Gender 96
B17 2x2 DiD Estimates of Log Firm Income by Firm Size 96
B18 Productivity Growth in Construction and Services 97
B19 Event Study-Effects in Competitive Industries 97

3.1 Employment rates across the life-cycle of working age labour force participants.100
3.2 Discrimination ratios in a two-period naive beliefs model 112
3.3 Discrimination ratios in the partial information model 116

List of Tables

1.1 Pooled Treatment Effect Estimates on Primary Outcomes 24
1.2 Pooled Treatment Effect Estimates on Job Tenure and Ability of Best Workers 24
1.3 Stages of Selection 25
1.4 Treatment Effects of Screening Length on Primary Outcomes 26
1.5 Treatment Effects of Screening Type on Primary Outcomes 27
1.6 Effects of introducing a friction on exploratory variables 27
1.7 Self-Selection Mechanism vs Observable Candidate Quality Predictors 28
A1 Notes: Data shared in correspondence from Prolific 34
A2 Treatment Arms 34
A3 Learning Mechanism 40
A4 Learning Mechanism: By Effort 41
A5 Heterogeneity: Application Rates 42
A6 Heterogeneity: Ravens Score 43
A7 Heterogeneity: Job Performance 44
A8 Alternative measures of task performance 44
A9 Balance Test 45
A10 Balance Test: Disaggregated 45
A11 Post Double Selection Results: Pre-Specified Primary Outcomes 46

2.1 Descriptive Statistics: Selection into Construction Sector by Gender 57
B1 NACE Sector Codes Treated With Income Tax Exemption 89
B2 Descriptive Statistics: Selection into Construction Sector by Gender: Pre-Reform 91
B3 Descriptive Statistics: Selection into Construction Sector by Gender: Post-Reform 92
B4 Descriptive Statistics: By Firm Composition 93
B5 Different choices for the calculation of standard errors 98

3.1 Summary - model setup 107
3.2 State-contingent employment probabilities for $g = B$ at $t = 2$ under naive beliefs 111
3.3 State-contingent employment probabilities for $g = B$ at $t = 2$ under partial information 114
3.4 Summary Statistics for 1979 Cohort 117
3.5 Empirical Results 118

Introduction to the Thesis

This thesis is comprised of three standalone papers which address strategic behaviour in labour markets, with particular focus on the role of information in shaping choices. I examine self-selection behaviour, tax-evasion decisions, as well as the way in which jobseekers are evaluated in a setting with labour market discrimination. The motivation underpinning each of these chapters is to improve understanding of how decisions made strategically in the labour market are shaped by a combination of common information, private information and social norms. As I will demonstrate in Chapter 1, mechanisms designed to transfer private information can also have the effect of shaping behaviours, and as I show in Chapter 2, income under-reporting can exist as an 'open-secret', where it is known that the behaviour is common practice, but is difficult for authorities to address. In Chapter 3, I show how the evaluation of candidates from a discriminated-against group depends crucially on the information environment.

Another motivation which binds the three papers together is to improve knowledge of inequalities, particularly between members of different groups. In the first paper, I show that screening procedures being a requirement in the hiring process does not result in women or minorities selecting out of the applicant pool disproportionately. In the second paper, after providing evidence suggesting that income under-reporting is prevalent in the Romanian construction sector, I show that men seem disproportionately more likely to under report their income, and that this has consequences for the measurement of the gender wage gap. In the final paper, I study employment gaps between members of observably different groups using a theoretical model of employer learning and taste-discrimination, where the gap between groups may increase or decrease over time.

My first paper develops a simple theoretical model of screening procedures as a cost in the application process, thus affecting the decision of whether a jobseeker self-selects into the applicant pool. In the model, mandatory screening procedures impose an opportunity cost on applicants' time, which is larger for better quality candidates. On the other hand, the effort cost associated with the procedure is smaller for better quality candidates. Screening procedures also allow jobseekers to learn more about the vacant job, and through a learning mechanism the decision of whether or not to take-up a job offer is affected. We bring our theoretical predictions to the data with a hiring experiment in an online labour market. Our results show that the introduction of a screening procedure results in substantial increases to the quality of the applicant pool and increases in job offer take-up rates, while also reducing the number of applicants. Decomposing the effects of opportunity and effort costs in the experiment lends some qualitative support to our theory, although the decomposed empirical results lack statistical precision. The induced self-selection improves the quality

of inexperienced applicants more than those with experience, indicating that screening tasks help jobseekers learn about the job and their personal match quality with the vacancy. The effects of the self-selection induced by the screening procedure improve the quality of the applicant pool so much that it is quantitatively comparable to (if not even better than) using only standard measures of ability, such as education status or relevant experience.

The second paper studies a major set of tax reforms in Romania's construction sector and uses the resulting changes in behaviour to infer whether workers in the sector are likely to be evading tax. Due to the system of third-party (employer) reporting, evasion occurs through employee-employer collusion whereby it is jointly agreed to falsely report the minimum wage to the tax authorities, on which tax is paid, while the remainder of the salary is taken cash-in-hand as an 'envelope wage'. The reform increased the minimum wage in the construction sector by approximately 50 percent, but reduced the rates of taxes and social contributions so that the total tax obligation was lower at this new construction-specific minimum wage compared to the regular one. Pre-reform, a substantial share of employees working in the sector were reporting a salary at exactly the minimum wage. Were these genuine, the policy would be expected to cause unemployment. However, the policy induces no unemployment, and due to the re-classification of firms into the construction sector and the emergence of new firms, the sector actually grows. The policies also had a condition which firms could potentially manipulate to be exempt from the reforms, but firms did not choose to do so, indicating that the true pre-reform wages were at least at the value of this new construction-specific minimum wage, but under-reported. Adding to the evidence that the changes in reported wages reflect the revelation of prior wages being under-reported, I examine one-person firms. Because the employee is their own employer, there is no incentive to mark-down wages relative to the marginal product of labour, and jumps of income reports (as well as firm revenue reports) from the regular minimum wage to the new construction-specific minimum wage indicate that the prior wage is almost certainly a false declaration for the purpose of evading tax. After providing the evidence which indicates that income under-reporting is prevalent, I highlight that men are much more likely than women to exhibit changes in reporting behaviour that is consistent with tax evasion. The gender composition of the firm appears to be a significant determinant of gender differences in reporting behaviour, with men and women in mixed-gender firms having very similar responses to the policy while diverging substantially in gender-segregated (including one-person) firms. This supports a risk and social preference based explanation, whereby the possibility for the joint agreement to collude breaking down deters men from attempting to participate in such a collusion.

In my final paper, I develop a dynamic model of employer learning to study group differences in employment rates over time. This contributes to the recent developments in the theoretical literature on the dynamics of discrimination, and highlights that the information environment is a key determinant of the evolution of group differences. In the model, ability is imperfectly observed, so previous employment outcomes are informative in updating beliefs about candidates. However, some jobseekers are discriminated against, which makes it more difficult, or impossible, to gain employment. I show that if employers are naive to the existence of discrimination by other employers, then discrimination will be exacerbated over

time, as unemployment caused by discrimination is falsely interpreted as a negative signal of ability, even by non-discriminatory employers. On the other hand, if employers are rational Bayesians and are aware of the existence of discrimination in the economy, then extra lenience is given to members of the discriminated-against group, and the negative signal associated with unemployment is rationally discounted. I show that in this setting, group differences can converge over time and that there is even the possibility for a reversal of group differences. One interpretation of behaviour in this information setting is that employers naturally develop an affirmative action policy simply by updating their beliefs about jobseeker ability correctly and profit maximising. The paper begins with a simple model, which generates sharp predictions and clear intuition, and then complements this with a richer model that relaxes a number of restrictive assumptions which had been made for mathematical convenience. Although the model predictions are less sharp, the economic intuitions developed in the simple model carry over to the less restrictive model.

Chapter 1

Search Frictions and Self-Selection in Job Matching

Abstract: Screening tests are often used by employers to learn about job applicants. In addition to generating information, they act as a costly friction in job search, which can affect self-selection of jobseekers into the applicant pool, potentially increasing or decreasing the match quality between workers and firms. We develop and experimentally test a theory of employer screening procedures as an application cost and the resulting self-selection behaviour of jobseekers on an online job matching platform. We find that introducing screening procedures for a short-term, task-based job opening reduces application rates, but the self-selection behaviour of applicants increases average match quality, thereby reducing wasted effort on both sides of the market. When we decompose the costs of completing screening procedures we find some evidence, although statistically weak, that increasing the opportunity cost of applicants' time drives a decrease in the quality of the applicant pool and that increasing the cost of effort incurred in undertaking screening procedures increases the quality. We also highlight that job screening procedures not only alter the costs of a job application but also applicant learning about the work, which affects take up of job offers.

We thank Rose Page and staff at the Centre for the Study of African Economies (CSAE) for research management. We gratefully acknowledge funding from the UK Foreign, Commonwealth & Development Office (FCDO) through the Private Enterprise Development in Low-Income Countries (PEDL) Programme. Fieldwork received ethical approval from the University of Oxford Economics Department Research Ethics Committee (#ECONCIA22-23-16). The project is registered with the AEA RCT Registry (#AEARCTR-0011711 , https://doi.org/10.1257/rct.11711-1.0). We gratefully acknowledge comments from Johannes Abeler, Abi Adams, Sam Altmann, Alison Andrew, Sarah Clifford, Stefan Dercon, Hamish Low, Barbara Petrongolo, Simon Quinn and seminar audiences at M-BEES 2024. Abelardo De Anda Casas kindly provided scans of archived labour force statistics which we used for a real-effort data entry task in our experiment. All errors are our own.

1.1 Introduction

With the proliferation of online platforms and matching technology, labour markets are increasingly "wired", transforming the ways in which firms and workers find each other (Pallais, 2014). Popular job search websites allow jobseekers to review hundreds of vacancies within minutes and applications can often be submitted with just a few additional clicks. Standard search theory predicts that such a decrease in search costs will unequivocally raise match quality and, ultimately, productivity (Goldfarb and Tucker, 2019; Pissarides, 1990).

At the same time, low marginal application costs in thick online markets can result in high volumes of ill-targeted job applications, sometimes called "congestion" in the economics literature (Agrawal et al., 2015; Agrawal, Lacetera, and Lyons, 2016).[1] In the business literature, Cappelli (2023) describes efforts by recruiters to make it "faster and cheaper [to] shovel more applicants into the funnel" as misguided and instead proposes to focus on smaller and better-qualified applicant pools. When job search is costless, the average match quality of an applicant with a given role may be low.

Employers are thus increasingly exploring the use of costly application procedures, such as screening tests and gamified job "previews"[2] – not just to sift through large numbers of applications in a more cost-effective way, but also as ordeals mechanisms to induce better self-targeting of jobseekers to vacancies. If the costs associated with these screening technologies can indeed change the self-selection behaviour of applicants, what is the effect of introducing application costs on applicant behaviour and match quality in an online labour market?

Our paper investigates the relationship between application costs induced by job screening requirements and subsequent match quality. We focus on self-selection into applications as the main mechanism of interest. We first present a simple theoretical framework to show how different types of application costs may affect selection. This contrasts the role of "effort costs", which can be thought of as cognitive costs that decrease in ability, against the role of "opportunity costs", which can be thought of as the lost value of applicants' time spent on applying. In our framework, effort costs screen out low-ability applicants at a higher rate than high-ability applicants, while the opposite is true for opportunity costs. Since high-ability jobseekers tend to have more valuable outside options and may thus be less willing to undertake more costly applications, the marginal effect of additional screening procedures on match quality is theoretically ambiguous. Moreover, the content, and not only the size, of the application costs may matter for self-selection. Application procedures that are relevant to the tasks of the job may present a learning opportunity to applicants about their fit for the job, particularly those inexperienced in necessary job-specific skills.

To measure the impact of application costs on match quality and empirically disentangle

[1] This is similarly observed in college applications and the academic job market for economists, where marginal application costs are close to zero (Roth, 2018; Coles et al., 2010).

[2] See, for example, Jaser and Petrakaki (2023) or Homegardner (2020) for popular media coverage. WEF (2023) reports now that 27 percent of employers use 'psychometric testing' as a screening device and 47 percent use 'proprietary skill assessments'. The HR website Recruiter.com, for example, recommends "five innovative ways to screen out the masses," including through attention checks in application forms, unique document requirements, or time-consuming video interviews (Ladimeji, 2023).

the role of effort and opportunity costs, we conduct an experiment on an online task matching platform, Prolific. This online labour market facilitates short-term gig (task-based) work (Stanton and Thomas, 2020), for example creating training datasets for large language models. We invite jobseekers on the platform to apply for a short-term data entry task in support of an unrelated university research project. We randomly expose applicants to screening tasks that vary in size and content. By cross-randomising both the cognitive difficulty of the test and the time taken to complete it, we separately identify the effects of effort costs (difficulty) and opportunity costs (time). In the control group, an application does not involve any screening test and therefore is virtually costless. All jobseekers who opt into applying are offered the data entry job.

We find that introducing application costs in the form of a job screening procedure significantly decreases application rates by 30 to 40 percentage points relative to the control group. At the same time, the screening requirements lead to a significantly higher-quality applicant pool: applicants have higher baseline cognitive ability and perform about 0.3 standard deviations better on the data entry job. Decomposing these application costs, we find that increasing the difficulty and the time required to undertake the screening procedure decreases application rates. Nevertheless, despite a smaller applicant pool, those in the higher-effort application group perform better in the data-entry task than those in the lower-effort group. Increasing the opportunity cost associated with the application decreases task performance on average. The marginal effects of increasing the difficulty of the screening procedure and increasing the time required to complete the task, while in line with our theoretical predictions, are small in magnitude relative to the size of imposing a screening requirement at all, and statistically weak.

We find evidence that learning about the job (both objective content and subjective fit) drives self-selection behaviour and subsequent improvements in match quality. Disaggregating results by workers' prior data experience, the effect of the screening procedure on shrinking the applicant pool and increasing its quality is strongest for inexperienced jobseekers, who we assume learn more about the job from the screening procedures. We also find that the screening requirement did not disproportionately deter female and non-white jobseekers from applying, but there were significantly stronger treatment effects of the procedure on the ability of the applicant pool among poor and white workers. Screening task requirements also lead to higher task satisfaction and lower perceived difficulty of the job among workers, which together with the lower application rates and higher abilities of treated workers suggests that both employers and employees benefit from screening frictions in an online labour market setting with congestion.

Our work contributes to several literatures. Costs of applying can reduce take-up of beneficial social programmes (Bertrand, Mullainathan, and Shafir, 2004) such as the Supplemental Nutrition Assistance Program (SNAP) (Finkelstein and Notowidigdo, 2019). However, there is an argument that imposing costs, such as work requirements for social welfare, acts as a screening device which deters those who are not in need of the programme (Besley and Coate, 1992). Nichols, Smolensky, and Tideman (1971) point out that waiting times, which imply an opportunity cost, also have the effect of rationing public goods. Complexity in

applications can play a similar role in screening (Kleven and Kopczuk, 2011).

There is also a small literature on thick markets with low application costs and resulting congestion externalities (Roth, 2018), with the economics job market used as an example. Theoretically, Arnosti, Johari, and Kanoria (2021) show that excess visibility of applicants to employers results in inefficient matches. Applicants who are a good match to an open position are likely to already be taken by another, lower quality, match. Reducing visibility would enhance welfare. Kanoria and Saban (2021) similarly finds that wasted effort can be reduced by limiting contact made by the larger side of a two-sided matching market. Empirically, Gürtzgen et al. (2021) find that online recruiting raises the number of applicants, as well as the share of unsuitable candidates per vacancy.

We build on a literature on online labour markets. Although they have the potential to reduce information frictions, difficulties remain in communication between employees and employers (Pelletier and Thomas, 2018). For example Pallais (2014) studies inefficient employment of inexperienced workers and Pallais and Sands (2016) examines the effects of worker referrals on the oDesk task-based platform. On the same platform, Stanton and Thomas (2016) study the signalling value of intermediaries in matching.

Online labour markets a growing part of the global labour market. Kassi, Lehdonvirta, and Stephany (2021) estimate that there are 163 million online labour market profiles globally and that there has been a substantial amount of growth in registered worker accounts in recent years. Other estimates have the figure of registered workers at 53 million (Codagnone, Biagi, and Abadie, 2016) and 48 million (Kuek et al., 2015). Using different methods of estimation Datta et al. (2023) put the figure at 154 or 133 million, highlighting the rapid growth over time. Our experiment takes place on an online matching market used for short term, freelance, 'gig' work. This kind of work is an increasingly large share of the economy, with 39 percent of the US workforce doing at least some freelance work in 2022 while 17 percent of the workforce combine traditional employment with freelance work (Upwork, 2022).

In a low-income country setting, small monetary payments to offset transport and opportunity costs (lost wages) for jobseekers can increase average applicant quality (Abebe, Caria, and Ortiz-Ospina, 2021). Our contribution is studying screening costs in a more general labour market setting and disentangling the main mechanisms.

Finally, we also contribute to a relatively small literature on employer screening technologies. Autor and Scarborough (2008) find that introducing employer testing procedures increases tenure while having no effect on minority hiring from a case study. Hoffman, Kahn, and Li (2018) study the potential trade-off between job testing technologies and managers' private information and find evidence that managers do not have superior private information and those who hire against the technology recommendation have worse employees.

At the Toulouse School of Economics, He and Magnac (2022) implement an experiment where an application cost for applying to different streams of master's programmes is introduced.[3] They find that introducing the small application cost reduces congestion without

[3]The cost is the writing of motivational letters, which induces both an effort cost and an opportunity cost of applicants' time

reducing match quality. Weinstein (2018) studies employer screening costs as a search friction on employers' recruitment behaviour, specifically, on how they target search in university campuses in the US, since search is costly. In contrast, we study the screening costs as a friction on job-seekers' self-selection into applying.

Policymakers and governments have supported the recent proliferation in online job platforms, for both gig and longer-term work. However, the extent to which employers engage with these technologies has largely been determined by the quality of the subsequent matches, measured by worker productivity and retention. This study demonstrates the importance of application costs, specifically from screening exercises implemented by employers, in shaping the behaviour of jobseekers on a task-matching platform. Our results suggest that introducing application costs through screening tests has an impact on the self-selection of jobseekers into the pool of applicants for a vacancy, in turn affecting the average ability of job applicants and the resulting match quality. Our experimental design unpicks the mechanisms by which screening costs affect the self-selection into the applicant pool. Our results are applicable not to just the growing market for virtual labour online, but for other labour markets in which there is little information, or it is difficult to credibly communicate information. Screening tests may therefore be especially valuable in developing economy contexts and in the markets for more inexperienced and less educated workers, since there is less information to communicate and it may be more difficult to signal ability reliably.

The paper proceeds as follows. Section 1.2 provides the conceptual framework from which our experiment follows, Section 1.3 presents the experiment context and describes our design. Section 1.4 describes the empirical strategy, Section 1.5 describes the results and Section 1.6 concludes.

1.2 Theoretical Framework

To structure our analysis, we provide a simple theoretical framework for job applications under different costly screening procedures. The model captures two different types of costs faced by jobseekers: an effort cost and an opportunity cost to applying. Effort cost can be thought of as cognitive cost, which decreases in ability or match quality for the specific position. Opportunity cost can be thought of as loss of income, loss of leisure time, or loss of the expected value derived from the next best job application that the applicant would otherwise have spent their time on. Our theoretical framework describes the effect that each of these costs has on the pool of applicants.

There are two sides to the job market, jobseekers ('he') and employers ('she'). Jobseekers are active decision makers who assess the expected value of applying for a vacancy and decide whether to apply or not. Employers are passive, to match our role as the experimenters who randomly impose different screening procedures on jobseekers and offer short-term task-based work automatically to all those whom apply.

There are two stages of selection. Firstly, the individual decides whether or not to apply for the job, and secondly they decide whether or not to take up the job offer if they had

applied for the job. This matches our experimental setting in which all job applicants are offered the job, but of course not everyone takes it up.

In the first stage of self-selection the individual compares the expected value of applying with not applying and makes the decision to apply according to the following simple rule:

$$\text{EV(apply)} > \text{EV(not apply)} \tag{1.1}$$

$$\text{EV(job)} - \text{cost of application} > \text{EV(not apply)} \tag{1.2}$$

Conditional on applying, the individual will receive a job offer and take-up the job according to the rule:

$$\text{EV(take-up job)} > \text{EV(not take-up job)} \tag{1.3}$$

To generate simple experimental predictions, we put a minimal amount of structure on these decision rules. We otherwise leave the framework as purposefully unstructured, so that it can guide the empirical analysis without imposing restrictive assumptions.

Jobseekers in our framework differ in their ability, a. In this context, ability can be thought of narrowly as task specific for a data entry job. It thus categorises applicants into those who have a higher or lower match quality with the employer. In our experiment, we use a number of proxies for ability. We assess general ability using performance on a Raven's Progressive Matrices test, a widely used measure of intelligence and abstract reasoning. We also look at the actual take-up of job offers made and real performance in the data entry task for those who take it up.

The job provides a lifetime value, V, comprising of the income generated by the work and the non-pecuniary value (the 'fit' between the worker and the job). When applying for the job, applicants do not know how good a match they are for the job and form an expectation of this lifetime value $\mathbb{E}(V)$. Individuals have a subjective belief, $\mathsf{Pr}(job)$, that they will receive a job offer. We impose no functional form assumptions on either $\mathbb{E}(V)$ or $\mathsf{Pr}(job)$.

When applying to a job, jobseekers face a screening. This can be thought of as a candidate assessment test, which may include batteries of questions to elicit cognitive and non-cognitive skills and fit with the job. This type of screening imposes costs: $C(e, a)$ denotes the composite cost of the application, and depends on the level of effort that needs to be exerted and the individual's ability. The effort requirement, e, could be thought of equivalently as the level of difficulty of the screening procedure.

To characterise the expected value of not applying to the job, we use $V(\mathcal{O}(t, a))$ as the individuals' value of the outside option, which varies depending on their ability and the length of time that the application takes. This yields the first selection model equation:

$$\mathsf{Pr}(job) \cdot \mathbb{E}(V) - C(e, a) > V(\mathcal{O}(t, a)) \tag{1.4}$$

We make the following assumptions:

- $\frac{\partial C(e,a)}{\partial e} > 0$
- $\frac{\partial^2 C(e,a)}{\partial e \partial a} < 0$
- $\frac{\partial V(\mathcal{O}(t,a))}{\partial t} > 0$
- $\frac{\partial^2 V(\mathcal{O}(t,a))}{\partial t \partial a} > 0$

These assumptions tell us that the cost of undergoing an application screening procedure is increasing in the effort required to complete the application, but that this marginal cost of effort is lower for higher ability individuals. Similarly, the value of the outside option is increasing in the time required to complete the screening procedure, and this marginal cost of time is higher for high ability candidates.

Putting together our simple framework and minimalist assumptions allows us to generate some predictions for the first-stage selection decision corresponding to our experimental setting.

An increase in t reduces the number of applicants by making the application relatively less attractive by increasing the opportunity cost of time. Because the value of time is higher for higher ability people, an increase in t will result in higher ability individuals self-selecting out of the application proportionally more than lower ability individuals, decreasing the quality of the applicant pool.

Increasing e, the effort required to complete the screening, also reduces the number of applicants by making the application less attractive. Completing a more difficult screening procedure is of course more costly. However, this cost is less burdensome to higher ability applicants, meaning that as a result of the self-selection decisions we would expect the composition of the applicant pool to become relatively higher quality.

At the second stage of selection, a job applicant (who has applied for the vacancy) has received an offer and has to make the choice of whether to accept it or not. The job lasts $\bar{t}$, which in our experimental setting is 20 minutes, and thus the applicant weighs the expected value of taking up the job with the value of the outside option. Having experienced the screening procedure, the applicant may update their expected value of the job depending on the difficulty of the procedure and their ability. The decision rule is to take up the offer under the following condition:

$$\mathbb{E}(V|e,a) > V(\mathcal{O}(\bar{t},a)) \tag{1.5}$$

We make no assumption on the sign of $\frac{\partial \mathbb{E}(V|e,a)}{\partial e}$, but assume that $\frac{\partial^2 \mathbb{E}(V|e,a)}{\partial e \partial a} > 0$. The logic is that the expected value placed on the job may be influenced by the effort required of the screening procedure and the individual's ability. To keep the model as general as possible, we do not assume that a more difficult screening makes the job more or less attractive. However, we assume that the effect of screening difficulty is complimentary with ability. That is, if a more difficult screening procedure has a positive effect on the value of the job, then it is

relatively more positive for high ability people and if a more difficult screening procedure has a negative effect on the value of the job, then it is relatively less negative for higher ability people. The predictions of the model for the job offer take-up decision arises immediately from the assumptions. It is unclear whether making the screening procedure more difficult increases or decreases take-up, but it does have the effect of improving the ability of the applicant pool.

1.3 Experimental Design

We take this theoretical framework to the data using a randomised control trial. Our experimental paradigm builds on previous work on labor market search frictions (Pallais, 2014; Pallais and Sands, 2016; Stanton and Thomas, 2016). Our design is a 'natural field experiment', as coined by Harrison and List (2004)[4], in which the jobseekers whom we study do not know that their job application, take-up of job offers, and resulting job performance is being observed for research purposes - they are in their 'natural' environment. We run the natural field experiment on an online matching platform usually used for short-term gig work, Prolific, where we invite jobseekers to apply for a short-term data entry task in support of a university research project. The experiment randomly varies both the time and effort costs associated with submitting an application, while including a control group of individuals who are able to costlessly apply for a short-term job.

1.3.1 The Online Labour Market

Prolific is an online task matching platform. Prolific advertises itself as a platform that primarily matches private sector and academic researchers with gig workers for research and data collection purposes, and has been called a "microtask" platform (ILO, 2022). Prolific can be considered as an MTurk alternative with greater regulation from the platform, such as ensuring that gig-workers are vetted upon entry and removed from the participant pool if they fail to meet quality standards.[5]

In our data, approximately 75% of participants report using Prolific as a primary or secondary source of income. 80% of respondents report being currently employed outside of the platform, either as an employee or self-employed. Of those employed, 38% of them report looking for full or part-time work in addition to their main job.[6] Additionally, there is strong demand to join Prolific from potential participants, and waiting lists to join can last several months. We thus interpret workers on Prolific as short-term gig-workers, similar to a university student doing food deliveries for part-time work, or a full-time employee doing gig-work on Uber when they are not working in their main job. The platform can be considered as something that strikes a balance between a traditional (online) labour market and a well-regulated pool of online research participants.

[4]Two examples of papes using a natural field experiment are (Kube, Maréchal, and Puppe, 2013) and (Leibbrandt and List, 2015)

[5]Eyal et al. (2021) confirm that data quality from Prolific is considerably higher than that from MTurk.

[6]23% of participants report being dissatisfied about their current or most recent job. In contrast, only 6% of participants report being dissatisfied with their most recent task in Prolific.

At the time of the experiment, there were approximately 120,000 active users on the platform. Approximately one-third of active[7] participants were located in the United States. Another third of participants are located in the UK. Approximately 60% were female, both those located in the US and worldwide.

The median age of our participants is 35, while the mean is 38. 58% of our sample holds a University degree, which includes 17% with a higher (Graduate or professional) degree. Of the total population of freelance workers, 49% hold a University degree (Upwork, 2022) and relatively younger workers are more likely to work in the freelance market (Upwork, 2023). Additional informaiton about the online platform and our sample are available in the Appendix to Chapter 1.

1.3.2 External Validity

Our natural field experiment has validity to environments outside of our experimental settings. Most directly, our setting is similar to other online labour markets for short-term gig work and should translate readily into those settings. More broadly, in settings where information is scarce, or hard to credibly communicate, such as some developing country settings and in the market for inexperienced workers, job screening procedures are likely to be an efficient way to generate information about jobseekers while also inducing positive self-selection into the applicant pool and reducing the size of the pool.

1.3.3 Experimental set-up

Our experiment tests whether introducing small application costs induces positive self-selection into the decision to enter the applicant pool for a vacancy. By cross randomising the time taken to complete the screening exercise and the difficulty of the exercise, we disentangle the effects of increasing the opportunity time costs and increasing effort required to complete the task, as proposed by our theoretical framework.

In line with the norms on the task matching platform, we first invite participants to undertake an initial survey hosted on Qualtrics, which includes questions about their demographics and work experience that will form part of our baseline survey (hereafter "baseline survey"). We also measure their ability (proxied by Raven's Progressive Matrices – a commonly used non-verbal test for fluid intelligence and abstract reasoning), risk and time preferences, and life satisfaction.

Upon completion of the baseline survey, we invite participants to apply for a second task (henceforth "data entry gig" or "data entry job"). The short-term job entails digitising PDF scans of archival data in support of a University of Oxford economic history research project. Workers are presented scans of a PDF document containing tables of labour force statistics in Spanish and are asked to transcribe the data in columns and rows using an online table in their screen. The statistics are from 1990 and record data on the size of the population and the pool of employed and unemployed individuals, by sex and by locality (municipalities

[7] Active workers need to have undertaken work within the last 90 days.

and parishes) in Venezuela. An example of these scanned data records is provided in the Appendix to Chapter 1.

We exogenously vary the application procedure for the data entry gig along effort and time dimensions. The willingness to bear the randomly assigned screening cost and thus apply for the data entry job is one of our primary outcomes of interest.

We randomise participants who successfully completed the baseline survey into one of the following treatment arms, which vary the size and type of application costs.

- **Control:** A control group is offered a 'no cost' option to express interest in being contacted for the data entry job. This group is simply required to click a button to submit an application.
- **Time cost treatments:** To assess the importance of opportunity costs of time in application decisions, we randomly vary the length of time required to complete the application screening procedure. The time required is either 1-minute, 3-minute, or 5-minutes.
- **Effort cost treatments:** To assess the importance of the effort cost of the applicant screening procedure, we vary the effort required to complete the screening. We implement a low-effort video procedure, in which the prospective applicant must watch a 1, 3, or 5 minute video explanation of the data entry task. To ensure that participants watch the video, they are required to answer questions that verify they paid attention to the video once completed and they are told about the attention checks in advance to ensure compliance. Additionally we implement a high-effort procedure. This is a 1, 3 or 5 minute long trial-run of the data digitisation gig for which we are hiring.

We target at least 320 participants in each of the seven experimental groups and the number of individuals in each treatment group is presented in Table A2.

Every applicant who applies for the data-entry gig then receives a job offer, and is given seven days to complete the work. We provide them with a reminder to complete the task at approximately the five day mark. The baseline survey and the experimental step (the invitation and screening tasks) together take up to 30 minutes to complete. The data entry task entails digitising PDF scans of archival data and lasts for 20 minutes.

Holding the time variable fixed, comparing the difficult screening task (trial-run data entry) with the easy task (watching the video explanation) identifies the effect of additional effort costs on our primary and secondary outcomes[8]. By holding the effort level (the difficulty) of the task fixed, comparing the lengths of time (1, 3 or 5 minutes) required to complete the screening exercise identifies the effect of increasing the opportunity cost of application on our outcomes variables. The vidoes provide a generic explanation of how to read rows and columns from a PDF and write the associated text and numbers into a given digital spreadsheet. Longer videos do not contain additional information of material substance, but instead generic statements (e.g. information about the University of Oxford's college system,

[8]Details on the construction and measurement of our outcome variables is available in the Appendix to Chapter 1.

its placement in various university rankings and recommendations for participants to place the place their mobile phone on silent and close other browsers).[9] Pooling the various treatments together allows us to make highly powered comparisons of modestly sized screening frictions against the no screening control. We also examine the effects on take-up of job offers and examine the mechanisms by which the screening tasks affect candidate quality and performance.

Once participants have completed the data entry job, we ask them to complete a very short endline survey. Here, we ask a short set of questions about their experience with the gig and their willingness to do similar work for us in the future.

Participation is remunerated at the hourly market wage (USD 9) on the online task matching platform. The flowchart in Figure A2 graphically represents the different stages of the experimental design.

To manage logistics, we implement the study over seven waves between July and August 2023 and stratified the sample on gender[10]. In each wave we ensured to exclude participants who previously completed the experiment, by adding their unique Prolific identification number to a blocklist.

We find that participants are well balanced in terms of baseline socio-economic characteristics across all groups (see Table A9 and Table A10).

1.4 Empirical Strategy

Our main empirical specification estimates pooled treatment effects as follows:

$$Y_i = \beta_0 + \beta \cdot T_i + \beta_f \mathbb{1}\{Female\} + \boldsymbol{\delta} \cdot X_i + \epsilon_i \tag{1.6}$$

where Y_i is the outcome of interest. T_i is a treatment dummy to indicate whether the participant faced any of the six screening treatments. X_i is the pre-specified vector of controls - a set of dummies for round of the survey while the female indicator is self-explanatory, and was also pre-specified. ε_i is a mean zero error term. β estimates the aggregate treatment effect of introducing a screening requirement relative to the no-screening control group.

In our second set of specifications, we disaggregate the overall effect into two effects: Equation 1.7 identifies only the effect of increasing time (opportunity cost) of the task by pooling together the high and low effort arms for each time arm. Equation 1.8 identifies only the effect of increasing effort (cognitive cost) of the task, pooling together the three time arms for each effort arm.

$$Y_i = \beta_0 + \sum_{t=1}^{3} \beta_t T_{t,i} + \beta_f \mathbb{1}\{Female\} + \boldsymbol{\delta} \cdot X_i + \epsilon_i \tag{1.7}$$

$$Y_i = \beta_0 + \sum_{e=1}^{2} \beta_e T_{e,i} + \beta_f \mathbb{1}\{Female\} + \boldsymbol{\delta} \cdot X_i + \epsilon_i \tag{1.8}$$

[9] The scripts of the low-effort video procedure and instructions for the trial-run are provided in the Appendix to Chapter 1.

[10] In our pilot study, the gender distribution of participants was unenven.

where $T_{t,i}$ and $T_{e,i}$ are indicators for individual i receiving treatments of different lengths of time and different difficulty respectively. The coefficients β_t estimate the effect of the time required to complete the application, which can be 1, 3 or 5 minutes. The coefficients β_e estimate the effect of either the high effort (trial data entry task) or low effort (video explanation) screening. All coefficients are interpreted relative the the pure control group.

In our final specification, we identify treatment effects for each of the six treatments compared to the control group:

$$Y_i = \beta_0 + \sum_{j=1}^{6} \beta_j \cdot T_{j,i} + \beta_f \mathbb{1}\{Female\} + \boldsymbol{\delta} \cdot X_i + \epsilon_i \tag{1.9}$$

The coefficients β_j estimate the treatment effects of each treatment arm relative to the control group.

In all our specifications, our main results control for strata (gender and survey wave) fixed effects. We check the robustness of our results to an array of other control variables using Post-Double Selection Lasso regressions (Belloni, Chernozhukov, and Hansen, 2014). In the main regressions we pre-specified, we provide both the usual robust standard errors as well as q-values which correct for multiple hypothesis testing. We specifically use the procedure of Benjamini, Krieger, and Yekutieli (2006), with the implementation discussed in detail by Anderson (2008).

1.5 Results

We first ask whether introducing a screening task in the application process impacts the application rate and the match quality relative to the no-screening control group using Equation 1.6. We then disentangle the effects of varying the content and size of the friction to unpack the self-selection behaviour of jobseekers using Equation 1.7 and Equation 1.8. Finally, we explore heterogeneous treatment effects and benchmark the various screening tasks in our experiment against alternative methods of selecting applicants.

1.5.1 Introducing a Screening Task Improves Match Quality

We find that introducing a screening task greatly reduces the application rate, but increases the quality of workers selecting into the applicant pool and their subsequent job performance. Table 1.1 reports the treatment effects on our three pre-specificed primary outcome variables, where we compare outcomes for jobseekers facing any screening task to the control group. All results below remain statistically significant at conventional levels after correcting for the false discovery rate in multiple hypothesis testing.

When invited to apply to the data entry job, 82 percent of jobseekers in the control group chose to apply. The screening requirement reduces the application rate by 37.2 percentage points (p=0.02), a fall of 45 percent compared to the control group.

We next investigate the average ability of the applicant pool by treatment status. Worker ability, as measured by our cognitive ability measure, the Raven's score, is 0.13 standard deviations (p=0.05) higher for the treatment group relative to the control group.

Not all who complete the screening task and apply for the job go on to take up the job offer, and this is further investigated in Table 1.3. However, conditional on taking up the job, the screening task requirement in the application process improves performance on the job by 0.3 standard deviations (p=0.07) relative to the control group. In raw numbers, treated applicants submit 9.2 more correct entries on average compared to the control mean of 33.6 correct entries – a 27.3 percent increase. This result is robust to alternative definitions of task performance (Appendix Table A8).

It is worth noting that employers may not care about the quality of the average applicant as much as they do about eliminating the worst candidates or selecting the best candidates. We examine this in two ways.

We first narrow down our group of interest to "top" candidates only. We find that introducing the screening task increases the proportion of high-ability applicants (those with cognitive ability scores above the 75th percentile of the control group score distribution) by 6.4 percentage points (p=0.02). However, we do not detect any significant changes in the share of applicants scoring above the 90th percentile of the control group (Table 1.2).

We next find that introducing the screening tasks leads to better on-the-job performance across the entire distribution of participants. Figure 1.1 compares the distributions of correct entries of untreated vs. treated participants. The performance distribution of treated applicants first-order stochastically dominates that of the control group (Somers' D test: $p < 0.01$). This means the screening procedure improves candidate quality across the entire performance spectrum, benefiting firms whether they aim to select top performers or eliminate low performers.

Worker ability and on-the-job performance are not the only measures of match quality employers may be interested in. Given the high rates of churn in labour markets[11], employers may care how many of them accept the job offer made to them, as well as how many of them actually follow through with the job, the latter being particularly relevant in short-term gig jobs such as the one we recruit for. In Table 1.3, we investigate these stages of selection - attrition from key stages of our study - as outcome variables of interest in their own right. Having made an application (and therefore having been offered the job), 55 percent of applicants in the control group start the job. Those who faced a screening task in the application process however were on average 14.7 percentage points (p=0.03) more likely to start the job, a 27 percent increase compared to the control group. Among those who start the job, 73 percent of those in the control group finish it, and here too we see that those who faced a screening task are on average 14.5 percentage points more likely to finish the job.

[11] Low offer acceptance rates are common across various labour markets. In a comparable online labour market (oDesk, now called Upwork), Pallais (2014) finds that 18 percent of those assigned to receive a job offer did not accept it, even though they initially applied. Acceptance rates of job offers are even lower in Pallais and Sands (2016) at 51 percent of their control group. In the offline labour market, Burks et al. (2015) find rejection rates of job offers between 20 to 25 percent, using personnel records from nine large firms across three industries.

From an employer's perspective, introducing the screening task unambiguously improves outcomes in hiring. It significantly reduces application rates, thereby decreasing the administrative burden of reviewing applications. Simultaneously, it improves the overall quality of the applicant pool, increasing both the average and top-end performance of candidates.

1.5.2 The Effects of Screening Task Length

We next examine whether the duration of screening tasks affects application rates and match quality. This analysis focuses on the role of the opportunity cost of time, a central mechanism in our conceptual framework outlined in Section 1.2. We vary the time required to complete the screening task (1, 3, or 5 minutes), pooling together both the difficult and easy screening tasks. Extensive piloting informs the choice of these durations, which represent 5 percent, 15 percent, and 25 percent of the 20-minute data entry task.

Our findings, presented in Table 1.4, indicate that changing the duration of the screening task has no statistically detectable effect on application rates or match quality. Quantitatively, the point estimates move in the expected direction, with increasing the time required to complete the screening decreasing both application rates and job performance, but these differences are small, particularly in the case of application rates, and are not statistically significant. Worker ability, as measured by Raven's score, does not moves monotonically with increases in the time, and is not statistically significant either.

1.5.3 The Effects of the Content of the Screening Task

The content of the screening task imposed on potential job applicants can alter both the size and quality of the applicant pool. We focus primarily on two potential mechanisms, the first of which is altering the effort required to complete the task, thereby affecting the pool of individuals willing to self-select into applying for the job. A more difficult screening task is a more costly friction in the application process, and is likely to be especially deterring to lower ability candidates. Secondly, the screening procedures reveal information to jobseekers about the job, which informs them of their particular match quality with the job as well as providing some training directly on how to do the job. Both the video and the trial-run provide more information on the nature of the job, while the trial-run provides some practice on how to complete the job.

We examine how the type of screening task affects jobseeker behaviour and match quality by varying exposure to a video about the job or a short data entry task in the application process. The video treatment entails watching a 1, 3 or 5 minute instructional video about the job.[12] The data entry task mirrors the content of the job requiring 1, 3 or 5 minutes of effort, depending on the timing treatment assignment. By engaging jobseekers in a shorter version of the work they are applying for we expect the task to be more difficult in terms of

[12]The instructions show one of the authors explaining and demonstrating how to read the PDFs and enter the data from the PDF into the table. The video also emphasises that it is important to enter the data correctly and explains some background information about the data contained in the PDF. No additional relevant information is contained in the longer videos.

requiring active engagement than the video, while also providing some degree of practice for the full-length job.

Table 1.5 reports effects of the video and task treatments relative to the control group, pooling together all duration arms.

We find that the data entry task has a stronger effect on job application decisions than the video treatment. Where the video treatment reduces application rates by 35 percentage points relative to the control group, the task treatment reduces application rates by 39.4 percentage points - a 4.4 percentage point difference (p=0.02). Despite reduced application rates, we find no statistically detectable differences in worker ability or job performance between the video and task trial-run treatments, though in line with our theoretical framework the task group has substantially larger point estimates for these effects.

To disentangle whether behaviour is being altered due to the differences in effort required to complete the screening procedures or due to the differences in jobseeker learning, we first note that intrinsic worker ability cannot reflect the effects of job practice, since the Raven's test is taken before any screening takes place.

We provide suggestive evidence that the screening task induces applicants to learn about the job, by disaggregating results by prior data entry experience. We find that the effects of the screening are strongest for inexperienced jobseekers, those who self-report as having no prior data entry experience. Table A3 examines heterogeneity in the effect of being assigned to the (pooled) screening treatment by prior data entry experience (any experience vs no experience), using Equation 1.10. We pool across all treatment arms, since the trial-run may act partly as a training while the video may act as a more indirect form of training and compare against the no screening control group. Inexperienced jobseekers are 8.5 percentage points (p=0.03) less likely to submit an application which requires a screening relative to experienced jobseekers. There are only very small effects of the treatment on worker ability among those with data entry experience, but significant positive effects among the inexperienced. The inexperienced also showcase much stronger treatment effects in job performance compared to the experienced, with the point estimate being 2.65 times larger for the former. Since the effects of the screening are concentrated on the inexperienced for both job performance but also worker ability as measured by the Raven's score, which is unaffected by the treatment as it is completed first, this indicates that workers who are inexperienced in data entry are learning more about their own fit for the job through the screening treatments, and self selecting according to their perceived job fit, rather than direct training explaining all of the difference between these groups of workers. Since the data entry trial-run screening provides a much more direct form of training compared to the video screening, we also look at the effects of undergoing a trial-run compared to the video, for the experienced vs inexperienced, in Table A4. This can disentangle whether inexperienced jobseekers have a larger treatment effect on their performance due to direct training from the task. The effect of undertaking the data digitisation trial-run compared to the video does not have a statistically detectable difference for experienced compared to inexperienced workers. This again suggests that there is likely to be little direct training effects of the screening procedures, but that undergoing one of these procedures teaches the jobseeker something about their own

match quality to the job.

1.5.4 Heterogeneity

In order to further unpack the drivers of the results, we conduct heterogeneity analysis by some theoretically relevant characteristics collected at baseline. These are: gender (pre-specified), ethnicity (white vs non-white), education, employment status, self-reported income, and poverty in the locality of residence (by zipcode). For each dimension of heterogeneity h, we estimate a saturated model of the form:

$$Y_i = \beta_1 \mathbb{1}\{h_i = 0\} + \beta_2 \mathbb{1}\{h_i = 0 \,\&\, T_i = 1) + \beta_3 \mathbb{1}\{h_i = 1\} + \beta_4 \mathbb{1}\{h_i = 1 \,\&\, T_i = 1) + \boldsymbol{\delta} \cdot X_i + \epsilon_i \tag{1.10}$$

In this specification, each dimension of heterogeneity h_i is a binary variable, and in the case of variables that are not binary (education, self-reported income, and zipcode poverty), a new binary variable is created that splits the sample into above or below median for that particular dimension of heterogeneity. For simplicity, as we are looking at heterogeneity in a pairwise manner taking one dimension at a time, we do not include gender in our vector of baseline controls X_i like in our main specifications, and instead only keep indicators for the round of the survey. The results are shown in Appendix Table A5 to Appendix Table A7. We find no heterogeneity in application rates among any dimensions. However, we do find that treatment effects on applicant ability are stronger among the poor and white subgroups. This pattern is also echoed somewhat in actual job performance, with stronger treatment effects among the white and those living in above-median poverty areas. Though our results are not dramatically different along demographic lines, the subset of significant results point towards poor, white workers driving our results. Even among the results that are statistically insignificant, we see that treatment effect point estimates are in almost all cases higher for those without education, unemployed, low income, and living in poorer zipcodes. Taken together, this provides some suggestive evidence on screening tasks in an online labour market setting having a greater effect on the self-selection behaviour of disadvantaged workers in particular.

To speculate on the underlying mechanisms generating these differential treatment effects, it may be that poorer individuals have higher opportunity costs of their time, and would spend their time on other forms of wage labour or non-wage household labour rather than undergo time and effort consuming screening procedures. White workers similarly may have more alternative options to generate income through wage labour than non-white workers and are thus more sensitive to having to undergo a screening procedure.

1.5.5 Applicant Welfare

The improvement in match quality is not only likely to be welfare-improving for employers, who can use screening tasks to decrease the size of the applicant pool while increasing its quality, but also for workers. After completing the data entry job, workers who incur the screening report higher task satisfaction of 0.22 standard deviations (p=0.01) relative to the

control group. Moreover, they are 0.16 standard deviations less likely (p=0.05) to report that the job was difficult, relative to the control group[13]. However, we do not observe any effects on our proxy measure for tenure, where we ask workers who complete the one-off data entry gig whether they would be interested in completing similar tasks in the future.[14] (see Table 1.2). Of course, while the workers who self-select into the job following a screening procedure seem to be a better 'fit' with the job, the costs of completing the procedure also enter into their utility functions, making the overall welfare effects on jobseekers ambiguous. In the case of a longer term job, however, the benefit of having an improved 'fit' with the job is likely to outweigh the costs of going through an initial screening procedure.

1.5.6 Selection on Observables

A mandated screening task acts as a beneficial friction to reduce congestion in job applications and increase the quality of the applicant pool. It is a low-cost and scalable intervention, and plausibly welfare enhancing for both sides of the matching market. Another way of reducing congestion is for firms to narrow the pile of candidates using observables, such as education, or simply make these a pre-requisite for applying. Allowing for the screening task to act as a self-selection mechanism is likely to be considerably more egalitarian than using observables such as higher education as a hiring criterion since such observables would favour historically advantaged groups. In Table 1.7 we compare the worker ability as measured by Raven's scores and realised job performance between those who selected into applying for the job with and without a screening task requirement. Next, we predict worker ability as measured by Raven's score and job performance by using a dummy for holding at least an undergraduate degree and a dummy for self-reporting some data entry experience, using the sample of those who applied for the job in the control group. The self-selection into applying which has been induced by the requirement to complete a screening task is at least as good at improving the quality of the applicant pool as would be using degree status as a pre-requisite for applying. The self-selection mechanism is also at least as good at improving the job performance as would be using data entry experience as a pre-requisite for applying, although this is not the case for predicting worker ability as measured by Raven's scores. Overall, the comparison of self-selection induced by a screening procedure with selection on observables tells us that the magnitudes of the effects we are able to induce are substantial.

[13] See Appendix Table 1.6 for results on these exploratory variables.

[14] We validate the hypothetical tenure question by offering those that completed the job the chance to apply for further work. Those that reported a hypothetical willingness to return for additional work were 9.5 percentage points more likely to apply for additional work. The regression of taking up additional work on the answer to the hypothetical question controls for data collection wave fixed effects and has a p-value<0.001, robust to heteroskedasticity. One concern is that those who took up the additional survey (660 individuals) which was used to apply for the additional task is a selected subsample of the 1056 individuals who completed the original data-entry task. Imputing a zero in the "application for additional work" dummy variable for those who chose not to complete the follow-up survey, which can be interpreted as dis-interest in additional data-entry work, reveals a qualitatively and quantitatively similar result. Those who expressed a hypothetical interest in additional work were 8 percentage points more likely to apply for the additional data-entry task (p=0.019).

1.6 Conclusion

We contribute to the literature on employer screening and search frictions by developing a theoretical framework on how the choice of tools used to screen applicants influences the self-selection of jobseekers in applying for a position. This is important to understand due to the proliferation of new job screening technologies in the labour market and the rising importance of online labour markets. Online testing tools allow firms to quickly learn something about the quality of applicants and can be especially useful in matching candidates who lack labour market experience or have experience in a separate labour market (such as in another country) with employers. The theoretical model disentangles the mechanisms by which screening tools can cause individuals to select into or out of the applicant pool. Increasing the difficulty, or effort required, of the technology can separate low and high ability jobseekers, since higher ability jobseekers find the difficulty of the screening less binding. Firms are able to acquire more information about applicants by increasing the number of different tasks, or screening procedures such as interviews, but by increasing the time the applicants need to spend on the application the opportunity cost of jobseekers' time increases. Under the assumption that high ability jobseekers have stronger outside options (a higher opportunity cost of their time), increasing the length of time required to complete screening procedures results in higher ability individuals self-selecting away from the applicant pool. In the model, individuals who complete a screening task can also use it to learn more about their fit with the job - which can improve subsequent match quality.

To bring our theoretical framework to the data, we run an experiment in an online labour market. We advertise a short-term data entry gig, for which we randomly vary the presence of a screening requirement in the application process. Randomly varying the difficulty and time required to complete the application, we are able to empirically disentangle the effort cost and opportunity cost mechanisms. We find qualitative evidence, though statistically weak, that increasing the difficulty of the screening task results in a better quality pool of jobseekers to self-select into applying for the job, but that increasing the length of time required to complete the screening task decreases the quality of the applicant pool. Although our empirical results are overall in line with our theory, the magnitudes of the effects of increasing the difficulty or length of the screening procedure are small compared to the presence of having a screening procedure at all - relative to no screening procedure. The message for employers choosing how to design applicant screening tools is thus that the presence of a costly screening procedure is highly effective at inducing positive selection into the applicant pool, while choosing the exact difficulty and time required to complete the procedures is a lower order concern.

Figures

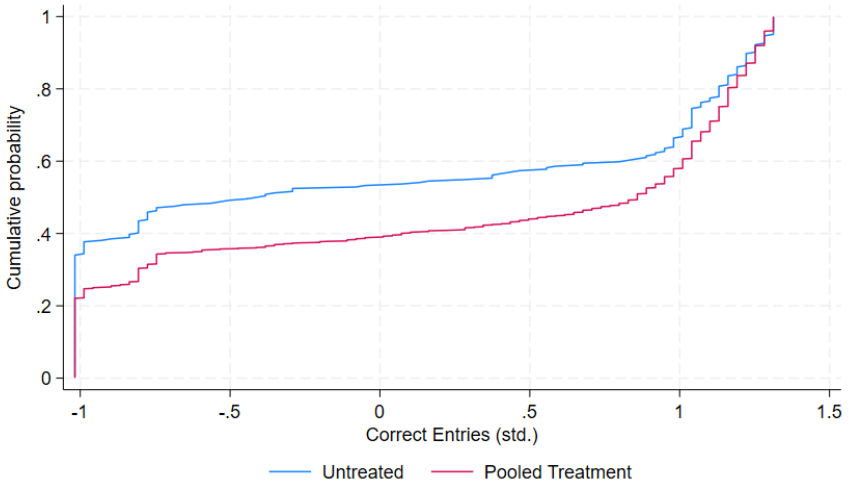


FIGURE 1.1. Cumulative Distribution of Job Performance of Control vs. Pooled Treatment

Notes: This figure shows the cumulative distribution functions (CDFs) of (stanardised) job performance for applicants who completed a screening procedure (pooled treatment group) and those who applied without screening (untreated control group,). Job performance is measured as the number of correct entries in the data entry task, normalized to mean zero and standard deviation one in the control group. The treatment group's CDF lies almost entirely to the right of the control group's CDF, indicating that screened applicants outperform unscreened applicants across the performance distribution. This implies that the screening procedure improves the quality of the applicant pool regardless of where employers set their hiring thresholds. A Somers' D test ($p < 0.01$) confirms first-order stochastic dominance.

Tables

Table 1.1. Pooled Treatment Effect Estimates on Primary Outcomes

	Application Rate	Worker Ability	Task Performance
	(1)	(2)	(3)
Pooled Treatment	-0.372***	0.131**	0.297***
	(0.02)	(0.05)	(0.07)
	[0.001]	[0.006]	[0.001]
DV Control Mean:	0.821	0.000	0.000
N	3470	1754	1164

Notes: Robust standard errors are reported in parentheses, sharpened q-values are reported in square brackets. Application rate is a dummy variable equal to one if an individual applied for the data entry task and zero otherwise. Worker ability is measured using the Raven's Progressive Matrices test score, normalised to have mean zero and standard deviation one among the control group of applicants. Task performance is the number of correct entries made in the data entry task conditional on accepting the offer, normalised to have mean zero and standard deviation one among the control group of applicants who completed the task. Control variables include gender and survey wave indicators. *, **, and *** indicate significance at the 10%, 5%, and 1% levels respectively.

Table 1.2. Pooled Treatment Effect Estimates on Job Tenure and Ability of Best Workers

	Tenure (Hypothetical)	Ravens > 75th pctile	Ravens > 90th pctile
	(1)	(2)	(3)
Pooled Treatment	0.041	0.064**	0.012
	(0.04)	(0.02)	(0.02)
	[0.404]	[0.032]	[0.535]
DV Control Mean:	0.682	0.268	0.132
N	997	1754	1754

Notes: Robust standard errors are reported in parentheses. Tenure (Hypothetical) is an indicator which takes on the value 1 if the response to "Would you do a task like this again in the future", is "Probably yes" or "Definitely yes" and 0 otherwise. Ravens > 75th (90th) pctile is a dummy taking on the value of one if the Raven's score is above the 75th (90th) percentile of the Raven's score distribution of the control group. Control variables include gender and survey wave indicators. *, **, and *** indicate significance at the 10%, 5%, and 1% levels respectively. FDR adjusted q-values are reported in the square brackets.

TABLE 1.3. Stages of Selection

	Applications	Started Job	Finish Job
	(1)	(2)	(3)
Pooled Treatment	-0.372***	0.147***	0.145***
	(0.02)	(0.03)	(0.03)
Sample:	All	Applied	Started Job
DV Control Mean:	0.82	0.55	0.73
N	3470	1754	1164

Notes: Robust standard errors are reported in parentheses. Application rate is a dummy variable equal to one if an individual applied for the data entry task and zero otherwise. Enter job is a dummy variable equal to one if an individual who was offered the data entry task began the data entry job and zero otherwise. Finish job is a dummy variable equal to one if an individual who was offered the data entry task and began the data entry job completed the job and the survey questions at the end of the job and zero otherwise. Control variables include gender and survey wave indicators. *, **, and *** indicate significance at the 10%, 5%, and 1% levels respectively. FDR adjusted q-values are reported in the square brackets.

TABLE 1.4. Treatment Effects of Screening Length on Primary Outcomes

	Application Rate	Worker Ability	Task Performance
	(1)	(2)	(3)
1 Minutes	-0.359***	0.153**	0.329***
	(0.02)	(0.07)	(0.08)
3 Minutes	-0.371***	0.063	0.314***
	(0.02)	(0.07)	(0.08)
5 Minutes	-0.386***	0.174***	0.251***
	(0.02)	(0.07)	(0.08)
p[q]-values:			
$\beta_{1min} = \beta_{3min}$:	0.592	0.169	0.838
	[1.000]	[1.000]	[1.000]
$\beta_{1min} = \beta_{5min}$:	0.223	0.746	0.307
	[1.000]	[1.000]	[1.000]
$\beta_{3min} = \beta_{5min}$:	0.505	0.083	0.413
	[1.000]	[1.000]	[1.000]
$\beta_{1min} = \beta_{3min} = \beta_{5min}$:	0.474	0.186	0.556
	[1.000]	[1.000]	[1.000]
DV Control Mean:	0.821	0.000	0.000
N	3470	1754	1164

Notes: The table presents the effects of varying the time required to complete the job application screening procedure (1 minute, 3 minutes or 5 minutes) relative to the no-screening control group on application rates, worker ability and job performance. The table also reports p-values from pairwise and joint tests of equality constraints between the estimated coefficients across the different time treatments. Robust standard errors are reported in parentheses. Application rate is a dummy variable equal to one if an individual applied for the data entry task and zero otherwise. Worker ability is measured using the Raven's Progressive Matrices test score, normalised to have mean zero and standard deviation one among the control group of applicants. Job performance is the number of correct entries made in the data entry task conditional on accepting the offer, normalised to have mean zero and standard deviation one among the control group of applicants who completed the task. Control variables include gender and survey wave indicators. *, **, and *** indicate significance at the 10%, 5%, and 1% levels respectively. FDR adjusted q-values are reported in the square brackets.

TABLE 1.5. Treatment Effects of Screening Type on Primary Outcomes

	Application Rate	Worker Ability	Task Performance
	(1)	(2)	(3)
Low Effort (video)	-0.350***	0.108*	0.270***
	(0.02)	(0.06)	(0.08)
High Effort (task)	-0.394***	0.157**	0.325***
	(0.02)	(0.06)	(0.08)
p[q]-values:			
$\beta_{task} = \beta_{video}$:	0.016	0.362	0.380
$\beta_{task} = \beta_{video}$:	[0.051]	[0.340]	[0.340]
DV Control Mean:	0.821	0.000	0.000
N	3470	1754	1164

Notes: Robust standard errors reported in parentheses. Application Rate is a dummy variable outcome taking on the value one if an individual self selects into applying for the task and zero otherwise. Worker Ability is the Raven's Score normalised around the distribution of scores for those in the control group. Because the comparison being made is against those who applied under different selection criteria (screening vs no screening) the sample used to make the comparison is of those who applied for the task, and this is also the sample used for the normalisation of the score around the control group. Job performance is the number of correct entries made in the data entry task. It has the smallest sample, since these scores are only available for those who applied for the task and took up the offer. The number of correct entries is normalised in the same way as Worker Ability/Raven's Scores. Control variables include gender and survey wave indicators. *,** and *** represent significance at the 10%, 5% and 1% level respectively. FDR adjusted q-values are reported in the square brackets.

TABLE 1.6. Effects of introducing a friction on exploratory variables

	Task Difficulty	Task Satisfaction	Risk Tolerance	Patience
	(1)	(2)	(3)	(4)
Pooled Treatment	-0.162*	0.224***	0.170***	0.070
	(0.08)	(0.08)	(0.05)	(0.05)
DV Control Mean:	0.000	0.000	0.040	0.031
N	997	997	1754	1754

Notes: Robust standard errors reported in parentheses. Task difficulty and task satisfaction are created using the self-reported subjective difficulty and satisfaction of the job, standardising the responses with respect to the control group to have a mean zero and a standard deviation of one. Risk tolerance and patience are created using self-reported willingness to take risks and willingness to give up something today in order to benefit in the future. Self-reported scores are given in values from zero to ten and then standardised with respect to the control group to have a mean zero and a standard deviation of one. Higher values of risk-tolerance represent more willingness to take risks and higher values of patience indicate more willingness to give something up today to gain something tomorrow. *,** and *** represent significance at the 10%, 5% and 1% level respectively.

TABLE 1.7. Self-Selection Mechanism vs Observable Candidate Quality Predictors

	Worker Ability		Task Performance	
	(1)	(2)	(3)	(4)
Pooled Treatment	0.131**		0.297***	
	(0.05)		(0.07)	
Degree Holder		0.064		-0.065
		(0.10)		(0.13)
Data Entry Experience		0.191*		0.206
		(0.10)		(0.14)
Sample:	Full	Control	Full	Control
N	1754	441	1164	244

Notes: Robust standard errors reported in parentheses. Worker ability is measured using the Raven's Progressive Matrices test score, normalised to have mean zero and standard deviation one among the control group of applicants. Job performance is the number of correct entries made in the data entry task conditional on accepting the offer, normalised to have mean zero and standard deviation one among the control group of applicants who completed the task. Degree holder is a dummy variable taking on the value one if the individual reported having at least an undergraduate university degree and zero otherwise. Data entry experience is a dummy variable taking on the value one if the individual self-reported having data entry experience and zero otherwise. Control variables include gender and survey wave indicators. *,** and *** represent significance at the 10%, 5% and 1% level respectively.

Appendix 1

A.1 Experimental Details

A.1.1 Measurement and Construction of Outcome Variables

This section details our key outcome measures capturing the effect of the screening tasks on self-selection behaviour, match quality, and welfare proxies. The construction of outcome variables follows our pre-analysis plan.[15]

Application behaviour

Application Rate
We measure the decision to select into the application process and then take up the job offer. Our first outcome measure of interest is thus the "application rate" for the advertised data entry job, measured as a dummy variable. "Applying to the job" in our experiment consists of choosing to progress into the next section of the survey to complete the allocated screening friction, and therefore be considered for the job we have described at the point of randomisation.

Starts Job
"Starts job" is a dummy variable taking on the value of one if the worker who had applied for the job begins the data entry job, which was sent out to job applicants typically the day after applications closed, and no more than one week after the applications had closed. The dummy variable is equal to zero if the job applicant did not begin the data entry job offered to them.

Finishes Job
"Finishes job" is a dummy variable taking on the value of one if the worker who had applied for the job completed the data entry job and the final survey questions, which were sent out to job applicants typically the day after applications closed, and no more than one week after the applications had closed. The dummy variable is equal to zero if the job applicant did not complete the data entry job and final survey questions offered to them.

Match quality

We measure match quality in two ways: worker ability and job performance.

[15] AEA RCT Registry #AEARCTR-0011711.

Worker ability

We measure general cognitive ability in the baseline survey using a subset of questions from Raven's progressive matrices, widely used as a measure of worker ability (Abebe, Caria, and Ortiz-Ospina, 2021; Fe, Gill, and Prowse, 2022). The test measures general fluid intelligence or the ability to make meaning out of complex information and to reproduce this information. We selected the moderately difficult middle module (the third our of five modules) of the test as our preferred subset, aligning with Prolific's norm for shorter tasks. We refer to this variable as "worker ability" throughout the paper.

Job performance

To measure performance in the data entry task, we count the number of correct entries. Each entry is a cell corresponding either to a local geographic region (a string entry) or one of the labour force statistics associated with that region (a numerical entry). The number of correct entries is then normalised to mean zero and standard deviation of one relative to the control group. Performance in the data entry task can be thought of as a measure of task-specific ability, or match quality. We henceforth call the measure "job performance".

We also create a more lenient measure of correct entries after allowing for minor stylistic mistakes in following instructions as a robustness check. These include omitting repeated rows and classifying words, such as "state" and "municipality". We will refer to this variable as "job performance (lenient)".

A related outcome concerning job performance is not just the number of correct entries, but the number of total entries made. This can be interpreted as a measure of brute effort, regardless of accuracy, and captures engagement with the task (rather than letting the timer run out, for instance). We will refer to this variable as "job performance (total entries)".

Worker welfare

To assess worker welfare, we examine several self-reported measures of job experience. Our primary pre-specified outcome is "hypothetical tenure," which gauges workers' willingness to undertake similar tasks in the future, serving as a proxy for actual job retention. Additionally, we explore two secondary outcomes: self-reported job satisfaction and perceived job difficulty, both measured upon task completion.

Hypothetical Tenure

We generate an indicator which takes on the value 1 if the response to "Would you do a task like this again in the future", with responses on a likert scale from 1 to 5 is "Probably yes" or "Definitely yes". The indicator takes on the value of zero otherwise.

Task Satisfaction

This variable is created using the question "How satisfied were you with your performance in this task?" with answers on a likert scale with 5 options from " Extremely dissatisfied"

to “Extremely satisfied”. A numeric variable is created by normalising to mean zero and standard deviation of one relative to the control group.

Task Difficulty
This variable is created using the question “On a scale from 1 to 10 with 1 being the most easy and 10 being the most difficult, how difficult did you find this task?” A numeric variable is created by normalising to mean zero and standard deviation of one relative to the control group.

Worker Characteristics

Finally, in our baseline survey, in addition to our measure of worker ability, we collect information on a number of worker characteristics. This includes zipcode of residence, highest level of education completed, occupation and industry of work, most recent wage in the labour market, past experience on the job platform, past experience in data entry, and risk and time preferences. We use these variables for our heterogeneity analysis, as well as to investigate how treated and control groups differ along some dimensions that employers may be interested in.

Gender
This variable is created simply using the pre-approved "Sex" question on Qualtrics, with available options being "Male", "Female", "Non-binary/ Third gender", and "Prefer not to say".

Ethnicity
Ethnicity data is recorded using Prolific’s preset ethnicity question, that asks participants to "Choose one or more races that you consider yourselves to be". For our heterogeneity analysis, we code our ethnicity variable to be "White" or "Non-White", where a participant is classified as being white if the first ethnicity they choose is the option "White or Caucasian".

Education
Education data is recorded by using Qualtrics’ suggested question on education. It asks "What is the highest level of education you have completed?", with step-wise answer options ranging from "Some Primary" at the lower end to "Graduate or professional degree" at the higher end. For the heterogeneity analysis, we split the sample into those who hold a University Bachelors’s degree (or more) and those who do not.

Employment
We use Qualtrics’ preset question on employment to record employment status. We code someone as being employed if they select "Working (paid employee)" and "Working (self-employed)", and unemployed otherwise.

Income

To record income levels of the participants, we ask the question "Approximately what is/was your hourly wage in US Dollars at your current/most recent job? (excluding your tasks on Prolific)". If they are paid a flat rate, we ask them to select the closest hourly equivalent depending on your hours worked. To encourage responses, we bin hourly wages according to Qualtrics' recommendations, with bin sizes ranging from 4 USD at the lower end of the wage distribution to 25 USD at higher end. We cap responses at "100 USD or more". We find that 51.73 percent of our sample earns less than 22 USD an hour, so we approximate our median at that bin mark. For our heterogeneity analysis, we create a dummy variable called "Income" that indicates whether the participant reports their latest hourly wage to be below 22 USD or 22 USD and above.

Local Poverty
We use Qualtrics' recommended zip code question and ask participants to report their 5-digit ZIP code[16]. For poverty data at the ZIP code level, we use 2017-2021 American Community Survey dataset, harmonised by IPUMS NHGIS. For each zip-code, we calculate a poverty rate by dividing the number of residents in the sample with "Income in the past 12 months below poverty level" by the total number of residents in the sample. We the merge this into our main dataset, and create a dummy variable indicating whether the participant resided in an area that was above or below the median ZIP code poverty rate among our participants (10 percent).

Data Entry Experience
We measure prior data entry experience with a simple yes or no question asking "Have you ever done any data entry tasks as part of your current job or in any of your past jobs?".

Risk Tolerance
We take responses to the question "Please tell me, in general, how willing or unwilling you are to take risks. Please use a scale from 0 to 10, where 0 means that you are "completely unwilling to take risks" and a 10 means you are "very willing to take risks". You can also use the values in-between to indicate where you fall on the scale." and normalise the responses to mean zero and standard deviation of one relative to the control group. Higher values indicate lower levels of risk aversion.

Patience
We take responses to the question "In comparison to others, are you a person who is generally willing to give up something today in order to benefit from that in the future or are you not willing to do so? Please use a scale from 0 to 10, where a 0 means you are "completely unwilling to give up something today" and a 10 means you are "very willing to give up something today". You can also use the values in-between to indicate where you fall on the

[16]Though US ZIP codes do not provide strict geographical boundaries, there are 41642 and are provided in several useful datasets to use as a unit of analysis.

scale." and normalise the responses to mean zero and standard deviation of one relative to the control group. Higher values indicate higher levels of patience.

A.1.2 Experimental Design - Additional Information

The Online Labour Market - Additional Information

All workers are anonymous and identified by a unique participant ID. This can be used to track workers over time and invite them for future task work. It can also be used to restrict workers from being invited for future work. Workers can be screened on pre-existing demographic information and characteristics, such as gender, ethnicity, nationality, languages spoken, religious and political beliefs, and education. Participants can also be screened on length of time they have been on the platform and number of submissions on the platform. Pay offers are made by the employers and tasks are taken up at the discretion of workers. All communication is done within Prolific, which provides a chat function. Because individually identifying information is hidden, all interactions, such as communication and payment, takes place on Prolific. Unlike other well-known job matching platforms, Prolific does not produce publicly visible ratings or reviews of workers.

Mean estimated completion time	14 minutes
Median completion time	14.21 minutes
Mean estimated pay per hour	£9.43
Median pay per hour	£13.13

TABLE A1. Notes: Data shared in correspondence from Prolific

Table A1 shows the platform-wide statistics that Prolific provided to us on request. The estimated completion times and rewards per hour are those set by those offering the tasks, while the median completion times and rewards per hour are the actual outcomes after participation. We use this information to decide on reasonable parameters and pay for our task, such that our task closely resembles what participants would expect to be offered on the platform on a regular basis.

Recruitment of Participants

We restrict participation in our study to English-speaking Prolific workers located in the US, but impose no other restrictions and then recruit them stratifying evenly by gender. The number of workers we recruit per treatment arm is given in Table A2.

TABLE A2. Treatment Arms

Effort	**Time**		
	1 minute	**3 minutes**	**5 minutes**
Low (video)	511	482	476
High (data entry)	491	461	512
Control (no effort or time)	537		

We also asked our participants their latest pay on Prolific to check if our subset of participants are meaningfully different to the platform-wide average (which includes all non-US participants). The mean hourly pay reported by our participants for their most recent task on Prolific was 11.95 GBP, and the median value was 10.71 GBP. On freelance websites such as Upwork, the median and mean hourly earnings are, respectively, USD 11.2 and USD 7.2.

Experiment Instructions

Control Group Job Application:

"You have now completed all questions and tasks. Thank you very much.

We would like to invite you to apply for a separate task digitising PDF scans of archived data. The task will take you no longer than 25 minutes and pay an amount of £3.50.

Would you like us to consider you for this task and contact you on Prolific if we choose to proceed in offering you the task?

To ensure that this study is correctly marked as completed, please select one of the options below." [Yes/No]

High-Effort (Trial Run) Group Job Application:

"You have now completed the study. Thank you for taking part. You will be paid out for this task in full sometime in the coming days.

We would like to invite you to apply for an additional task digitising PDF scans of archived data. The task will take you no longer than 25 minutes and pay an amount of £3.50.

This task requires attention to detail, accuracy, and speed. To test for this, we would like you to complete a screening procedure, which you can complete in the next step.

This will take [one/3/5] minute[s] and will be a simple trial run of the data entry task. You will see scans of archived data and be asked to enter this into a table.

If you would like to apply for this task, please press Continue below to do the screening procedure. Otherwise, please click on Exit Survey." [Continue/Exit Survey]

Low-Effort (Video) Group Job Application:

"You have now completed the study. Thank you for taking part. You will be paid out for this task in full sometime in the coming days.

We would additionally like to invite you to apply for an additional task digitising PDF scans of archived data. The task will take you no longer than 25 minutes and pay an amount of £3.50.

This task requires attention to detail, accuracy, and speed. To test for this, we would like you to complete a screening procedure, which you can complete in the next step.

This will take [one/3/5] minute[s] and will consist of a video explanation of how the data looks in its current format and how it should be digitised. There will then be a few short questions to verify that you have understood the instructions.

If you would like to apply for this task, please press Continue below to do the screening procedure. Otherwise, please click on Exit Survey." [Continue/Exit Survey]

High-Effort (Trial Run) Group Screening Instructions:

"We would like for you to complete a simple data digitisation task, converting PDF scans of historical data into a digital format.

Once you have completed one row of data, press the "add row" button to proceed with the next row of data. Do not click the "submit button".

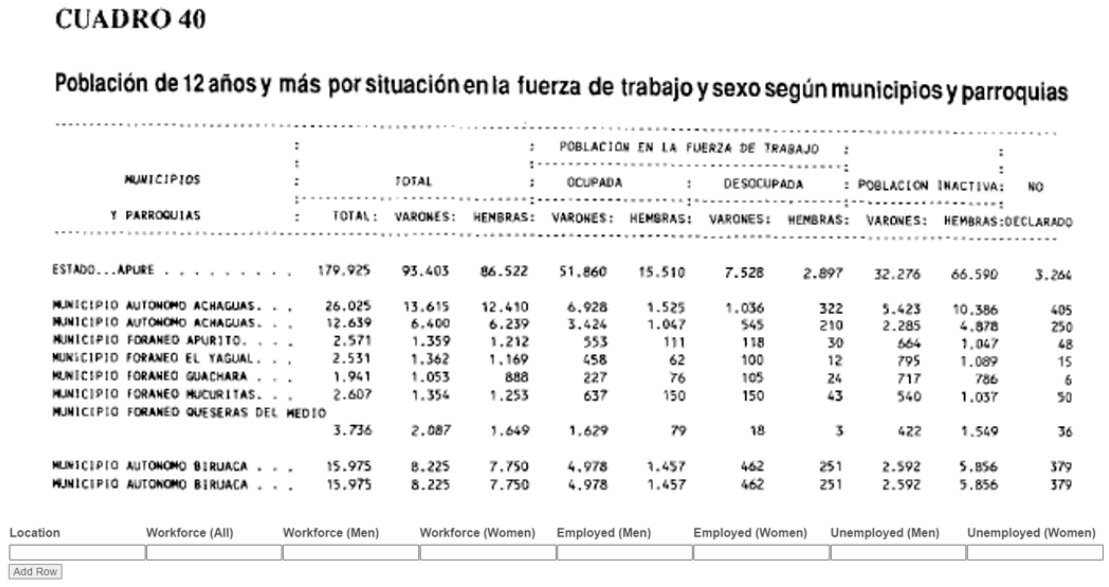
At the bottom of the next page is a timer for [1/3/5] minute[s]. Once the [minute/ 3 minutes / 5 minutes] is up, and only once it is up, you will be redirected to the next page.

Your response will be invalidated if you click on the "submit" button before the timer runs out.

Please type the data as it appears in the PDF scans. Because these are archived data records, the quality of the PDFs is not perfect. The first column gives the location, the second column gives the whole workforce, the third gives the male workforce, the fourth gives the female workforce, the fifth and sixth are the number of employed men and women and the seventh and eighth are the number of unemployed men and women.

Here is what the first three rows of scanned data looks like and how we would like it to be digitized:" [Image of scanned PDF and completed cells.]

Representative Example of PDF Scan:



CUADRO 40

Población de 12 años y más por situación en la fuerza de trabajo y sexo según municipios y parroquias

MUNICIPIOS Y PARROQUIAS	TOTAL			POBLACION EN LA FUERZA DE TRABAJO				POBLACION INACTIVA		NO DECLARADO
				OCUPADA		DESOCUPADA				
	TOTAL	VARONES	HEMBRAS	VARONES	HEMBRAS	VARONES	HEMBRAS	VARONES	HEMBRAS	
ESTADO...APURE	179.925	93.403	86.522	51.860	15.510	7.528	2.897	32.276	66.590	3.264
MUNICIPIO AUTONOMO ACHAGUAS. . .	26.025	13.615	12.410	6.928	1.525	1.036	322	5.423	10.386	405
MUNICIPIO AUTONOMO ACHAGUAS. . .	12.639	6.400	6.239	3.424	1.047	545	210	2.285	4.878	250
MUNICIPIO FORANEO APURITO. . . .	2.571	1.359	1.212	553	111	118	30	664	1.047	48
MUNICIPIO FORANEO EL YAGUAL. . .	2.531	1.362	1.169	458	62	100	12	795	1.089	15
MUNICIPIO FORANEO GUACHARA . . .	1.941	1.053	888	227	76	105	24	717	786	6
MUNICIPIO FORANEO MUCURITAS. . .	2.607	1.354	1.253	637	150	150	43	540	1.037	50
MUNICIPIO FORANEO QUESERAS DEL MEDIO	3.736	2.087	1.649	1.629	79	18	3	422	1.549	36
MUNICIPIO AUTONOMO BIRUACA . . .	15.975	8.225	7.750	4.978	1.457	462	251	2.592	5.856	379
MUNICIPIO AUTONOMO BIRUACA . . .	15.975	8.225	7.750	4.978	1.457	462	251	2.592	5.856	379

Location	Workforce (All)	Workforce (Men)	Workforce (Women)	Employed (Men)	Employed (Women)	Unemployed (Men)	Unemployed (Women)

Add Row

FIGURE A1. Screening Task PDF Scan

Notes: This is a representative image of the data entry task with cells provided to type in the data. This is part of a single page PDF scan. In the full length screening job, full length pages are shown along with buttons for the “previous” and “next” pages. In the screening trial-runs, only a single page of the PDF is shown since this is all which could be reasonably completed within the time limits.

A.1.3 Video Scripts

One Minute Video

Hello and welcome to the overview of the data digitisation task. In this task, what we’d like for you to do is to simply digitise some scanned archived records of Venezuelan labour force statistics that we can use in one of our research projects. It’s very important for our project that all this data is recorded accurately. So systematic errors when converting the

scanned archives into a digitial format can make it more difficult for us to understand the patterns in the data and at worst we could arrive at incorrect conclusions from our research. The way to digitise the archived records is simple. So, all that we want you to do is record each of these observations, and each observation is a location. So, for example, the first location is Estado Apure, you can see that the workforce for all people is A Hundred and Seventy Nine Thousand, Nine Hundred and Twenty Five and all the rows correspond to what you see here. Then all you want to do is keep adding rows, so for example the next row would be Municipio Autonomo Achaguas. In the event that you see the same location twice , please record both observations. Please record this every time you see the location. So in this example you would record Municipio Autonomo Achaguas twice. Thank you and we hope you choose to take up this task.

Three Minute Video

Hello and welcome to the data digitisation task. First, a little bit about the University. We're a group of researchers from the University of Oxford. We are currently ranked the best, or amongst the best, by various university rankings such as The Times Higher Education Guide. We have thirty nice constituent colleges and are the oldest university in the English Speaking World. According to the Research Excellence Framework, the work that we do is world leading research. What we'd like for you to do is to help us with one of our current projects. And, what we need in this project is the digitisation of some archived PDF scans of labour force records from Venezuela. It's crucial for out project that the data is transcribed accurately. Systematic errors in converting the scanned records into the digital format can make it much harder for us to understand what's going on in the data and at worst it could even result in mistaken conclusions from our research. I'll now explain briefly how to go about this task. All that we'd like for you to do is transcribe what you see in the scanned PDFs, which look like this, into a digital format. So, for example, the location for the first observation is Estado Apure with a total workforce of One Hundred and Seventy Nine Thousand, Nine Hundred and Twenty Five and so on. The next column would be the workforce for men and the following column would be the workforce for women and so on. All that you'd then do is add a row and proceed with the next observation which in this case is Municipio Autonomo Achaguas. In the event that you see the same location name twice, so for example here Municipio Autonomo Achaguas, what we'd like for you to do is to record every instance that you see the observation, even if you do see the same location name twice. Thanks very much and we hope that you choose to take on this data digitisation task.

Five Minute Video

Hello. Welcome to the data digitisation task video explainer. I'll first talk to you briefly about the university. We are researchers at the University of Oxford and are very proud that we consistently rank first or amongst the best by various university guides such as The Times Higher Education Guide. As a university we have 39 constituent colleges and we are the oldest university in the English speaking world. According to the Research Excellence

Framework, the work that we do is world leading. What we'd like your help with is digitising some archived data records for a new piece of research coming out of our university. These archived data records are labour force statistics from Venezuela and these data will ultimately be used for a project which aims to understand the evolution of these statistics during a period of political turmoil. In order for this project to work as efficiently as possible it is crucial that data is recorded accurately. Systematic errors in converting the scanned archived data records into the digitial format can make it much more difficult for us to understand the patterns in the underlying data. At worst, this could even result in mistaken conclusions from our research. I should note that the university takes health and safety seriously. We have undertaken a risk assessment which indicates that this task is of low risk, but we will just point out that you should think about staying hydrated and taking short breaks as necessary while you work. It can be hard to work productively on repetitive tasks for long periods of time. Some things you can do to help are try to turn off your phone or put it on silent mode. You could close other browsers on your computer, such as social media websites. And, taking a short break when you feel like you might need one can also help for you to work efficiently. I'll now go over how to actually digitise these archived records, which is a fairly straightforward procedure. What we are asking you to do is record for each observation, which is a location in Venezuela, statistics on the labour force. This is for the population as a whole and disaggregated by gender. What you'll see here is that the first location in the dataset is Estado Apure. The total workforce is A Hundred and Seventy Nine Thousand, Nine Hundred and Twenty Five individuals, so you just write that down in this cell here. You then proceed along the other columns, noting down the workforce for men, the workforce for women, the employed men, the employed women, the unemployed men and the unemployed women. When you are done with the first observation in this dataset, all that you do next is add a row and proceed with the next location. The next location here was Municipio Autonomo Achaguas and you'd just full in each column as you did with Estado Apure. In the event that you see a location twice, for example, Municipio Autonomo Achaguas, we ask that you record both, or all, observations. You would not simply record one observation, you would record both of these observations of Municipio Autonomo Achaguas and also down here Municipio Autonomo Biruaca. Thank you very much for listening to this video explained and we hope that you choose to take part in the data digitisation task.

A.2 Additional Results

TABLE A3. Learning Mechanism

Dependent Variable	Applied	Worker Ability	Job Performance
	(1)	(2)	(3)
Impacts for No Data Experience			
Treatment	-0.430	0.306	0.472
	(0.033)	(0.094)	(0.122)
Control mean	0.804	-0.178	-0.109
Impacts for Data Experience			
Treatment	-0.345	0.044	0.211
	(0.023)	(0.067)	(0.086)
Control mean	0.831	0.074	0.058
No heterogeneity (p)	0.034	0.024	0.080
Obs.	3470	1754	1164

Notes: Results from OLS regressions using a saturated model of testing for heterogeneity by data experience. For each outcome variable (column), the last panel reports the p-value of a test of the null hypothesis that the effect of treatment is not heterogeneous across those with and without prior data entry experience. Robust standard errors reported in parentheses.

TABLE A4. Learning Mechanism: By Effort

Dependent Variable	Applied	Worker Ability	Job Performance
	(1)	(2)	(3)
Impacts for No Data Experience			
High Effort Treatment	-0.057	0.122	0.129
	(0.031)	(0.091)	(0.117)
Low Effort Mean	0.401	0.104	0.303
Impacts for Data Experience			
High Effort Treatment	-0.035	0.023	0.022
	(0.023)	(0.064)	(0.074)
Low Effort Mean	0.502	0.109	0.287
No heterogeneity (p)	0.564	0.372	0.439
Obs.	2933	1313	920

Notes: Results from OLS regressions using a saturated model of testing heterogeneity in effects by effort level (excluding the no effort control), by data experience. For each outcome variable (column), the last panel reports the p-value of a test of the null hypothesis that the effect of treatment is not heterogeneous across those with and without prior data entry experience. Robust standard errors reported in parentheses.

TABLE A5. Heterogeneity: Application Rates

Heterogeneity by	Gender	Ethnicity	Education	Unemployment	Income	Poverty
	(1)	(2)	(3)	(4)	(5)	(6)
Impacts for	Female	Non-White	No degree	Unemployed	Low	High
Treatment	-0.367	-0.389	-0.411	-0.342	-0.361	-0.369
	(0.027)	(0.037)	(0.027)	(0.041)	(0.026)	(0.026)
Control mean	0.825	0.864	0.854	0.828	0.825	0.825
Impacts for	Male	White	Degree	Employed	High	Low
Treatment	-0.377	-0.369	-0.342	-0.378	-0.384	-0.376
	(0.027)	(0.022)	(0.026)	(0.022)	(0.028)	(0.028)
Control mean	0.817	0.809	0.795	0.819	0.816	0.816
No heterogeneity (p)	0.789	0.636	0.067	0.434	0.549	0.853
Obs.	3470	3470	3470	3435	3470	3470

Notes: Results from OLS regressions using a saturated model of testing heterogeneity. The heterogeneity variables are binary, with the continuous variables (Education, Income and Poverty) coded to be binary according to whether the observation is above or below median. Income is self-reported hourly wage, and poverty denotes the percentage of those below the poverty line in the respondent's zipcode of residence. Column headings indicate the dimension of heterogeneity studied. For each treatment and the last panel reports the p-value of a test of the null hypothesis that the effect of treatment is not heterogeneous across the dimension under study. Robust standard errors reported in parentheses.

TABLE A6. Heterogeneity: Ravens Score

Heterogeneity by	Gender	Ethnicity	Education	Unemployment	Income	Poverty
	(1)	(2)	(3)	(4)	(5)	(6)
Impacts for	Female	Non-White	No degree	Unemployed	Low	High
Treatment	0.180	-0.072	0.137	0.271	0.243	0.175
	(0.074)	(0.108)	(0.074)	(0.114)	(0.072)	(0.077)
Control mean	0.237	0.177	0.159	-0.098	0.224	0.196
Impacts for	Male	White	Degree	Employed	High	Low
Treatment	0.082	0.205	0.120	0.089	0.001	0.067
	(0.081)	(0.063)	(0.079)	(0.062)	(0.083)	(0.076)
Control mean	0.224	0.250	0.289	0.029	0.239	0.265
No heterogeneity (p)	0.368	0.028	0.874	0.164	0.028	0.319
Obs.	1754	1754	1754	1738	1754	1754

Notes: Results from OLS regressions using a saturated model of testing heterogeneity. The heterogeneity variables are binary, with the continuous variables (Education, Income and Poverty) coded to be binary according to whether the observation is above or below median. Income is self-reported hourly wage, and poverty denotes the percentage of those below the poverty line in the respondent's zipcode of residence. Column headings indicate the dimension of heterogeneity studied. For each treatment and the last panel reports the p-value of a test of the null hypothesis that the effect of treatment is not heterogeneous across the dimension under study. Robust standard errors reported in parentheses.

TABLE A7. Heterogeneity: Job Performance

Heterogeneity by	Gender	Ethnicity	Education	Unemployment	Income	Poverty
	(1)	(2)	(3)	(4)	(5)	(6)
Impacts for	Female	Non-White	No degree	Unemployed	Low	High
Treatment	0.300	0.046	0.342	0.511	0.408	0.472
	(0.097)	(0.132)	(0.105)	(0.140)	(0.095)	(0.096)
Control mean	-0.024	0.140	-0.008	-0.120	-0.031	-0.133
Impacts for	Male	White	Degree	Employed	High	Low
Treatment	0.294	0.386	0.267	0.244	0.168	0.107
	(0.103)	(0.083)	(0.096)	(0.081)	(0.105)	(0.102)
Control mean	0.026	-0.044	0.007	0.037	0.040	0.147
No heterogeneity (p)	0.964	0.030	0.600	0.099	0.089	0.009
Obs.	1164	1164	1164	1152	1164	1164

Notes: Results from OLS regressions using a saturated model of testing heterogeneity. The heterogeneity variables are binary, with the continuous variables (Education, Income and Poverty) coded to be binary according to whether the observation is above or below median. Income is self-reported hourly wage, and poverty denotes the percentage of those below the poverty line in the respondent's zipcode of residence. Column headings indicate the dimension of heterogeneity studied. For each treatment and the last panel reports the p-value of a test of the null hypothesis that the effect of treatment is not heterogeneous across the dimension under study. Robust standard errors reported in parentheses.

TABLE A8. Alternative measures of task performance

	Correct Answers	Correct Entries (Lenient)	Total Entries
	(1)	(2)	(3)
Pooled Treatment	0.297***	0.223***	0.261***
	(0.07)	(0.07)	(0.07)
DV Control Mean:	0.000	0.000	0.000
N	1164	1164	1164

Notes: Robust standard errors are reported in parentheses. Correct Answers is the number of correct entries made in the data entry task conditional on accepting the offer. Correct Entries with Lenience is the number of correct answers, adjusted for very minor discrepancies in following the exact rules. Total Entries is the total number of entries made, irrespective of accuracy. All outcomes standardised to mean zero standard deviation of one among the control group of applicants who completed the task. More detailed information on the construction of the outcome variables can be found in chapter 1.6. Control variables include gender and survey wave indicators. *, **, and *** indicate significance at the 10%, 5%, and 1% levels respectively.

TABLE A9. Balance Test

	Control	Pooled Treatment	p-val
Male	0.48	0.50	0.59
Employed	0.78	0.82	0.03
Non-White	0.22	0.25	0.11
Age	38.20	38.19	0.98
Degree Holder	0.55	0.58	0.18
Risk Tolerance	0.00	0.14	0.00
Patience	0.00	0.04	0.38

Notes: This table displays raw means for several descriptive variables grouped by the pooled treatment status. The number of observations in the pooled treatment group is 2,933 while it is 537 in the control. The p-value is associated with a test of equality of means between the two treatment groups.

TABLE A10. Balance Test: Disaggregated

	Control	Task			Video			
		1 Min.	3 Min.	5 Min.	1 Min.	3 Min.	5 Min.	p-val
Male	0.48	0.49	0.50	0.47	0.50	0.49	0.51	0.87
Employed	0.78	0.84	0.80	0.80	0.85	0.85	0.78	0.00
Non-White	0.22	0.24	0.29	0.24	0.26	0.25	0.22	0.13
Age	38.20	37.92	38.06	38.32	38.64	37.78	38.38	0.95
Degree Holder	0.55	0.61	0.56	0.59	0.63	0.55	0.56	0.05
Risk Tolerance	0.00	0.11	0.09	0.13	0.21	0.09	0.18	0.02
Patience	0.00	0.11	0.00	0.04	0.03	-0.01	0.08	0.42

Notes: This table displays raw means for several descriptive variables grouped by the treatment status. The p-value is associated with a joint test of equality of means between the all of the treatment groups.

TABLE A11. Post Double Selection Results: Pre-Specified Primary Outcomes

	Applications		Raven's Scord (std.)		Correct Answers (std.)	
1 minute task	-0.395***	-0.393***	0.206*	0.203*	0.286**	0.288**
	(0.03)	(0.03)	(0.09)	(0.08)	(0.10)	(0.10)
3 minute task	-0.395***	-0.396***	0.029	0.022	0.449***	0.421***
	(0.03)	(0.03)	(0.09)	(0.09)	(0.10)	(0.10)
5 minute task	-0.393***	-0.392***	0.228**	0.227**	0.254**	0.251**
	(0.03)	(0.03)	(0.08)	(0.08)	(0.10)	(0.10)
1 minute video	-0.325***	-0.326***	0.113	0.108	0.369***	0.365***
	(0.03)	(0.03)	(0.08)	(0.08)	(0.10)	(0.10)
3 minute video	-0.349***	-0.354***	0.094	0.100	0.190	0.201*
	(0.03)	(0.03)	(0.08)	(0.08)	(0.10)	(0.10)
5 minute video	-0.379***	-0.377***	0.119	0.122	0.246*	0.257**
	(0.03)	(0.03)	(0.08)	(0.08)	(0.10)	(0.10)
N	3470	3456	1754	1748	1164	1160
Lasso	No	Yes	No	Yes	No	Yes

Notes: The controls we allow the Lasso to choose from are: outside labour market wage, degree status, gender, age, risk aversion, time preference, and ethnicity.

A.3 Additional Figures

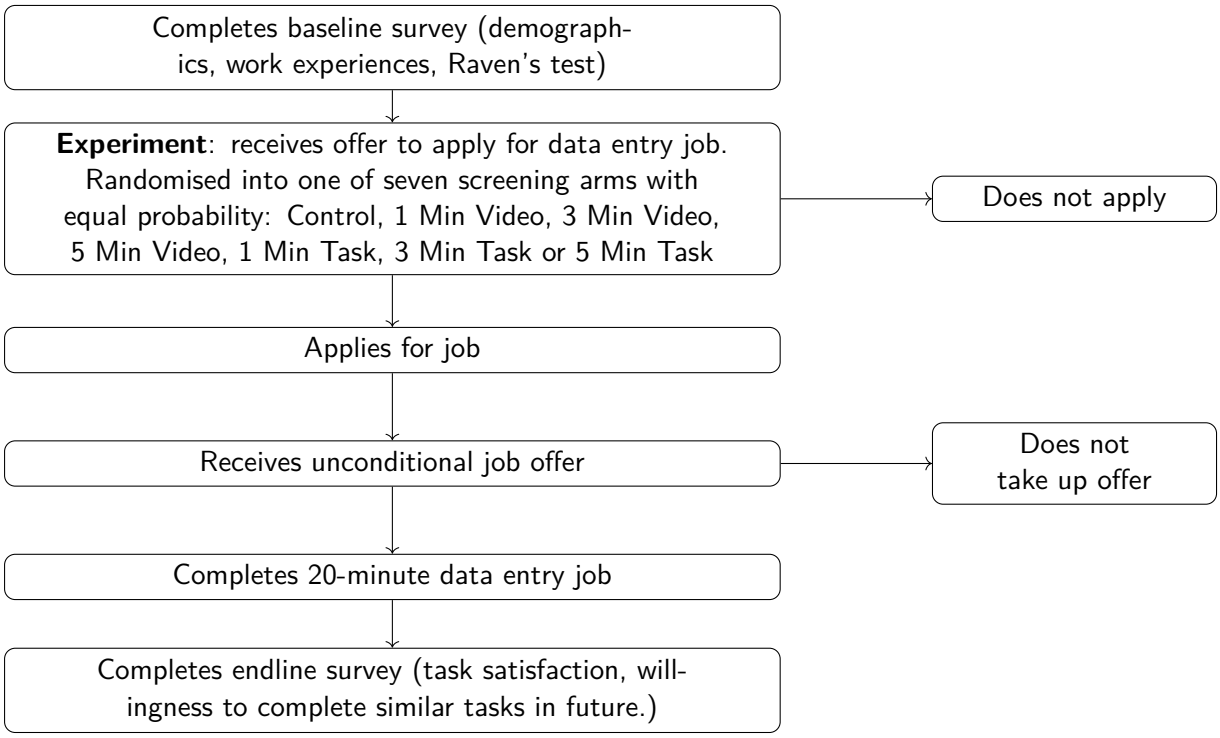


FIGURE A2. Sequence Data Collection and Experiment

Notes: There is an additional step after the endline survey. Those who indicated a willingness to take up additional work in the future were invited to apply for additional work. This provides a validation for a hypotherical question. Then, of those who indeed applied for additional work, one-third were randomly offerered to immediately take it up, while the remaining two-thirds were not given the offer due to budget constraints.

Chapter 2

Tax Evasion, Employee-Employer Collusion, and Gender Differences

Abstract: Third-party tax reporting systems are typically thought of as robust to tax evasion, since the evasion of income tax would require a joint collusion between employee and employer. Using administrative tax records from Romania and a major policy reform, however, I provide evidence suggesting that a large number of workers in the construction sector are evading income tax despite the third-party reporting system. The reforms change the tax and minimum wage rates in the sector substantially, and, ruling out alternative explanations, I use the behavioural responses to provide evidence of tax evasion. I add to the literature documenting labour market disparities between men and women by showing that men appear substantially more likely to be evading income tax than women. Investigating the mechanisms suggests that this is not being driven by gender differences in bargaining power or occupational segregation within the sector, but by segregation in the kinds of firms that men and women work for. Gender differences in preferences are a likely explanation of the observed difference in behaviour.

I gratefully acknowledge comments from Johannes Abeler, Abi Adams, Alison Andrew, Sarah Clifford, Mihai Codreanu, Hamish Low, Mathias Jensen, Niclas Moneke and Barbara Petrongolo and from conference participants at the Royal Economic Society 2024 Annual Conference, European Economic Association 2024 Annual Conference and European Association of Labour Economists 2024 Annual Conference. The Romanian Ministry of Public Finance kindly provided me with the data used in this paper and I received particularly generous support from Ioana Gavril and Madalina Avram. This paper makes use of confidential administrative data and received ethics approval from the University of Oxford with reference number ECONCIA22-23-04.

2.1 Introduction

Third-party reporting systems (where salaries are reported to the authorities by the employer, rather than the employee) are often considered to be a strong mechanism to induce truth-telling[1]. Leveraging a tax reform in Romania, I add to the body of evidence in the tax literature highlighting that employer reporting is not a silver bullet to tax evasion and that joint collusion between employee and employer can be sustained (Bjørneby, Alstadsæter, and Telle, 2021; Kumler, Verhoogen, and Frías, 2020; Bergolo and Cruces, 2014; Paulus, 2015)[2]. While gender differences in labour market outcomes, such as participation and wages, have received large amounts of attention, one potential difference in labour market behaviour, which has so far been unexplored, is the gender gap in tax evasion. Through the Romanian policy reform, I find that men appear to have a higher propensity to tax evade than women.

My setting is Romania, a country which is formal enough that the majority of workers and firms are registered with the tax authorities, for which there are strong incentives as a number of benefits such as health insurance, social insurance and pensions are linked to employment. However, it is commonly believed that a substantial amount of tax evasion takes place in a joint employee-employer collusion. This is often called 'grey work' and involves the employer reporting to the authorities that the salary of the employee is the minimum wage, while the remainder of the actual wage is paid cash-in-hand, and thus no income tax or other social contributions are paid on this additional payment[3]. Hidden from the authorities, the additional payment is often called an 'envelope wage'. Using administrative employee-employer tax records, I leverage a reform in Romania which eliminated income tax, substantially reduced the level of other social contributions and hiked the minimum wage by approximately 50% for those in the construction sectors. Eligible construction sector firms are determined by NACE code, with the policies affecting all of the firms' workers, regardless of the occupation within the firm. An important condition of the reform was that it only applied to firms which report that at least 80% of their revenue comes from construction activities. Otherwise, the set of policies do not apply. Importantly, the tax breaks mean that the tax minimising salary to report to the authorities is no longer the national minimum wage (NMW), but rather this policy-specific construction sector minimum wage (CMW).

I first show that following the introduction of the policy, workers in Romania's construction sector were almost 30 percentage points less likely to report a salary at exactly the NMW, from a pre-reform baseline of 34 percent. Workers in this sector were also approximately 30 percentage points more likely to report their salary as being exactly at the CMW, which is the new tax minimising level of income to report. Workers who were reporting a salary at exactly the NMW pre-reform were much more likely to jump to reporting the new,

[1] Kleven, Kreiner, and Saez (2016) describes third-party reporting as the reason modern tax enforcement is successful.

[2] Although this is not the case in Denmark, where only self-reported income is under-reported (Kleven and Kopczuk, 2011)

[3] Discussed, for example in this Romanian newspaper article https://www.bursa.ro/impactul-optimizarii-fiscale-a-muncii-in-romania-intre-legalitate-si-ilegalitate-21096646?fbclid=IwAR14gxCzEQpROHHcHWqd-9GCyyctYwu4V1TMi-zbHm6spQO2ktFt7g_tldM and this one https://www.europafm.ro/romania-in-direct-psd-si-pnl-maresc-si-salariul-minim-de-ce-o-treime-din-angajati-castiga-doar-salariul-minim-video/?fbclid=IwAR2dTjUZ7fq4KVJMCVb45vYyOP6u81-_c9LwobTKmzyD8BSRX3gKeeGsrho

50% higher, CMW, despite it being a larger income jump than for other employees earning above the NMW who were mechanically nearer to the threshold. As would be expected when income under-reporting requires joint employee-employer collusion, employees are less likely to change their income reporting behaviour in larger firms, and in the very largest firms the effects of the policy on income reporting behaviour are almost zero. An important special case is solo-firms. These are firms with just one employee who operates the firm, usually providing a relatively more specialist service, and where all of the firm revenue is paid out as a salary to the solo employee. In this case, the employee is not bargaining with an employer, since he or she is the firm operator. And, in this case, a hike in the reported wage (and corresponding declared firm revenue) of 50% from the NMW to the new CMW does not have any other plausible explanation other than revealing that income is under-reported for the purpose of avoiding tax.

Evidence of envelope wages, particularly in Central and Eastern European countries has previously been provided in the tax evasion literature. Williams and Padmore (2013) provide evidence from a representative survey conducted in the European Union in 2007, which found that envelope wages were ubiquitous in several countries and notably Romania, where 22 percent of employees reported that they receive these cash in hand payments.[4] Bíró, Prinz, and Sándor (2022) leverage a reform in Hungary which increased the probability of audit for workers reporting a sufficiently low gross salary[5]. Post-reform, the mass of employees reporting at exactly the minimum wage was decreased, and a new mass of employees appeared at the cutoff salary, where employees reporting under the cutoff faced the audit threat. With Latvian data, Gavoille and Zasova (2023a) present empirical facts strongly indicative of gross wage under-reporting. Firstly, it is shown that those reporting exactly the minimum wage are less likely to become unemployed following a minimum wage hike compared to those earning slightly above the minimum wage, indicating that many of those reporting exactly the minimum wage were in fact earning additional income off the books. Secondly, those reporting a salary at exactly the minimum wage who move from a small to a big firm report much larger increases in their salary compared to those who previously reported a salary slightly above the minimum wage, and also moved from a small to a large firm. Again equipped with Latvian data, using a machine-learning algorithm trained on a match of administrative data with survey data, and then applied to the complete Latvian administrative data, Gavoille and Zasova (2023b) estimates that 18 percent of employees are employed in a firm which engages in employee-employer collusion to under-report gross wages. Other empirical work studies how the gap between consumption and reported income changes in reaction to minimum wage or tax changes, in a way that is consistent with income under-reporting (Gorodnichenko, Martinez-Vazquez, and Sabirianova Peter, 2009; Tonin, 2011). This is based on the seminal methodology of Pissarides and Weber (1989) which relates individuals' reported income with food expenditure, assumes that wage employees report their income correctly to generate food

[4]There is also separate survey data focussing specifically on the three Baltic countries showing that envelope payments are commonplace (Meriküll and Staehr, 2010; Putniņš and Sauka, 2015). A survey in the Czech and Slovak Republics during economic transition showed that individuals' income under-reporting was related to their perception of government services (Hanousek and Palda, 2004).

[5]The threshold for the audit threat was at 2 times the minimum wage.

expenditure equations, and then applies these expenditure equations to the self-employed to back out the extent to which self-employees under-report their income.

In the theoretical literature, the level of risk-aversion is usually thought of as the primary driver of the tax evasion decision (Allingham and Sandmo, 1972), whereby the individual considers the probability of being caught and the penalty if caught under their own risk preferences. There are a number of other factors influencing the decision, such as reputational concerns, the technology available for avoiding or evading taxes (Mayshar, 1991) and the costs that might be incurred while avoiding taxes (Slemrod, 2001). Social norms, preferences for truth-telling and concerns of fairness have also been suggested as factors which can influence the decision (Andreoni, Erard, and Feinstein, 1998; Kirchler, Hoelzl, and Wahl, 2008; Traxler, 2010).

To date, there is minimal research on gender differences in tax evasion. Torgler and Valev (2010) use data from the World Values Survey and European Values Survey to show that women have greater aversion to corruption and tax evasion, while Kastlunger et al. (2010) and D'attoma, Volintiru, and Steinmo (2017) find that women were more tax compliant in separate lab experiments. Using an expenditure based approach, Bjørkheim and Nygård (2024) finds evidence that self-employed women in Norway evade more than men. Comparing administrative tax data with self-reported income from a household survey, and assuming that self-reports reflect the sum of both reported income and that which is hidden from the authorities, Robayo-Abril, Balaban, and Wronski (2024) find evidence that employees in Romania under-report their income to the authorities and that women do so to a lesser extent.[6]

There are a number of reasons why men and women may have different propensities to evade tax. A large literature has shown that there are gender differences in risk aversion (Borghans et al., 2009), lying aversion (Dreber and Johannesson, 2008; Childs, 2012; Abeler, Nosenzo, and Raymond, 2019), as well as social preferences and competitiveness (Croson and Gneezy, 2009; Backus et al., 2023; Fehr, Naef, and Schmidt, 2006; Niederle and Vesterlund, 2007; Niederle, 2017). Additionally, Roussille (2021) highlights the role of the gender 'ask-gap' on a job matching platform and Li and Zafar (2023) shows there is also a gender ask-gap for college regrades. Exley and Kessler (2022) shows that there is a, related, gender gap in willingness to self-promote. This may be particularly relevant in the Romanian setting because income reporting is done by the employer and tax evasion requires joint collusion with the firm, and this may require the employee to initiate the collusion. There may be gendered differences in collaboration between employee and employer[7], differentially affecting men and women's opportunity to evade tax. Biasi and Sarsons (2022) shows that when there is flexibility to bargain over wages, women lose out in pay compared to men. In the case where employees and employers jointly collude (or negotiate) over tax evasion, beliefs about others' willingness to engage in tax evasion can influence own willingness to engage.

[6]There is also evidence that construction is one of the sectors with the highest degrees of tax-evasion.

[7]Dreber and Johannesson (2008) showed that there were gendered differences in economics research collaborations, contributing to a gender output gap and Cullen and Perez-Truglia (2023) shows that social interactions at work contribute to a gender promotion gap. Beaman, Keleher, and Magruder (2018) shows that women may also be disadvantaged in employee referrals.

Given its sensitive nature, negotiations over a collusion to evade taxes may be particularly susceptible to gendered differences. Exley et al. (2022) shows that *beliefs* about gender differences in social preferences are large while Bursztyn, González, and Yanagizawa-Drott (2020) demonstrates that beliefs about gendered social norms may be incorrect. A lack of transparency (Baker et al., 2023) that tax evasion can be jointly agreed, or negotiated, may also contribute to a gender gap in tax evasion.

My data comprise of the entire matched employee-employer administrative tax records of Romanian firms and workers. The setting and data are particularly compelling to study tax evasion. It is widely believed that most tax evasion is at the intensive margin, where employees (and employers) are registered formally with the tax authority, but significantly under-report their true income to minimise the tax burden. So-called 'gray work', where they jointly report earning exactly the minimum wage, while the remainder of the true salary is paid cash in hand by the employer is said to be 'as difficult to stop as it is widespread' by one Romanian news outlet[8]. While I cannot observe tax evasion in the black market, I am able to use the tax reform to identify income under-reporting in the formal labour market using the tax records. Bearing in mind that individuals self-select into the formal labour market, the men and women in my sample are likely to be relatively more comparable than men and women in the labour market as a whole, including the black market. My estimates of tax evasion and the gender gap in tax evasion may in fact be under-estimates.

2.2 Institutional Setting, Data and Identification

2.2.1 Tax Evasion in Romania

Tax evasion in Romania is thought to be a pervasive issue. Anecdotally, institutions are strong enough that the majority of firms and employees must be formally registered with the tax authorities and report salaries, along with the corresponding tax payments. Operating formally as a firm and an employee also provides the benefits of eligibility for certain forms of government support, such as maternity/paternity benefits, unemployment insurance and pension benefits.

The anecdotal evidence is that tax evasion in Romania is often a 'gray' form of evasion, where the employee is formally registered with the tax authority, but falsely reports earning exactly the minimum wage, while the remainder of the true salary is paid cash in hand as an 'envelope wage'. One Romanian newspaper bringing this phenomenon to light has asked "Why do one third of Romanian employees earn only the minimum wage? Is it tax evasion or the power of the economy?"[9] In the construction sector, notorious for tax evasion, this figure was 50% before the tax reform. Qualitative work (Vâlsan, Druică, and Ianole-Călin, 2020) has shown that public attitudes towards tax evasion are only mildly intolerant of it.

[8] https://www.avocatnet.ro/articol_19270/Munca-la-gri-intre-evaziune-fiscala-si-avantajele-unor-taxe-mai-mici.html?fbclid=IwAR0ceBd1NTLp99ENYjPXUrs1bMsDOIyMsFAXF54Z6W61KNji_qREXPycFe8

[9] https://www.europafm.ro/romania-in-direct-psd-si-pnl-maresc-si-salariul-minim-de-ce-o-treime-din-angajati-castiga-doar-salariul-minim-video/?fbclid=IwAR2dTjUZ7fq4KVJMCVb45vYy0P6u81-_c9LwobTKmzyD8BSRX3gKeeGsrho

2.2.2 The Tax System and Policy Reform

The tax and social contributions system in Romania is based on three main components. Firstly, a 10% flat rate of health insurance contribution, secondly a 25% flat rate social insurance contribution and finally a flat 10% income tax on the remaining salary after health and social insurance contributions have been deducted and after any personal allowance. The personal allowance depends on the wage (and only applies to the income tax, not other contributions). At the 2018 minimum wage of 1900 it is 510 RON for employees without dependents[10], and decreases by 15 RON for every additional 50 RON earned above the minimum wage. The personal allowance system remains the same in the following years, but as the national minimum wage (NMW) increases, the size of the personal allowance becomes mechanically smaller[11].

In January 2019, workers in the construction sector, and some manufacturing, became subject to a number of policy changes so long as their firm met certain conditions. The NACE sector codes related to manufacturing are 'close to' construction, such as the manufacturing of construction inputs like bricks and concrete, rather than sectors such as the manufacturing of apparel, which is likely to be more formal and more gender neutral. The national minimum wage in 2019 was 2080 RON, increasing[12] from 1900 RON in 2018. At the same time, the government introduced a construction sector specific minimum wage (CMW) of 3000 RON for firms reporting that at least 80% of their revenues came from construction activities. This policy was unannounced and came into force as a 'Government Emergency Ordinance' at the end of December 2018.

In addition to exempting the employee from the 10% personal income tax, the policy also exempted employees from paying the usual health insurance contribution of 10%. Social insurance contributions are usually taxed at a rate of 25% of the income, but the policy reduced this figure to 21.25% for those in the construction sector. The total taxes and social contributions for construction workers (without dependents) reporting a salary at exactly the 2018 NMW would be 738 RON. At the slightly higher 2019 NMW it was 817 RON, but for those in construction at exactly the CMW, benefiting from the policy, it was a substantially lower burden of 638 RON. Even those with the largest personal allowance would pay a total amount of tax and social contributions of 737 RON at the 2019 NMW, still higher than the taxes plus social contributions owed at the higher CMW. Thus for the purpose of minimising the tax burden, a report of 3000 RON in the construction sector dominates a report of 2080 RON as the regular NMW. The tax schedules of a childless worker[13], pre and post-reform, are presented in Figure 2.1 and Figure 2.2 respectively.

[10] With dependents, the allowances are: 670 RON with one dependent, 830 RON with two dependents, 990 RON with three dependents and 1,310 RON with four or more dependents.

[11] At the 2019 minimum wage of 2080, the personal allowance was 465 and at the 2020 minimum wage of 2230 it was 420 (for those without dependents). To address this, a reform implemented in mid-2022 linked the personal allowances directly to the NMW rather than a fixed nominal value, ensuring that the system adjusts dynamically with annual changes to the minimum wage.

[12] In nominal terms and relative to the median, Romania has experienced substantial increases in the minimum wage over the past two decades (Robayo-Abril, Zamfir, and Wroński, 2024), without experiencing unemployment or inflarionary effects (Heemskerk, Voinea, and Cojocaru, 2018; Pantea, 2020).

[13] The median construction and non-construction worker reports zero dependents

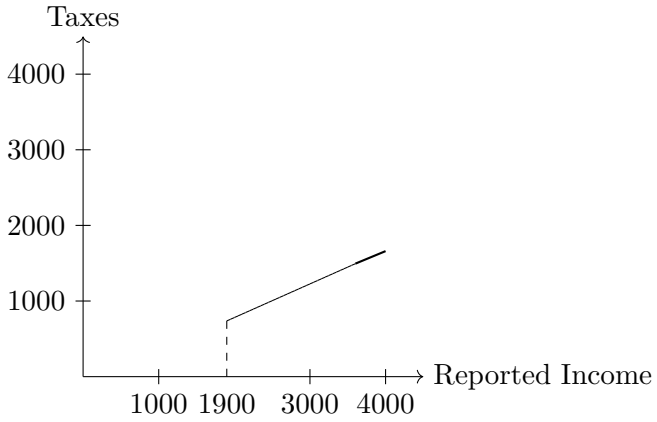


FIGURE 2.1. Take Home Pay vs Reported Gross Salary: Pre-Reform
Notes: The vertical dashed line intersects the x-axis at the 2018 minimum wage

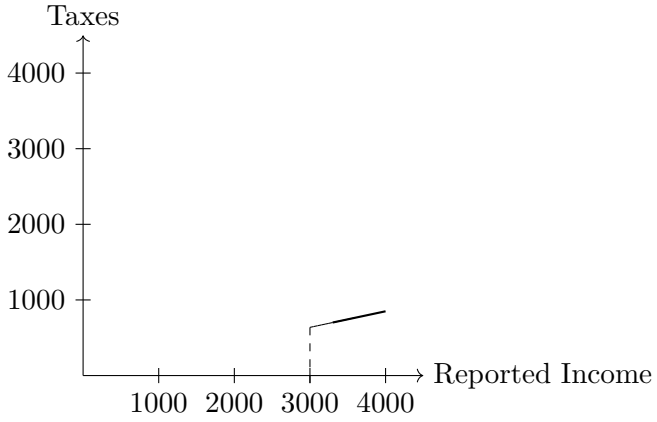


FIGURE 2.2. Take Home Pay vs Reported Gross Salary: Post-Reform
Notes: The vertical dashed line intersects the x-axis at the 2019 construction sector minimum wage

More minor details of the policy were a decrease in the labour insurance contribution owed by the firm from the regular 2.25% to 0.34%. Employees earning a salary of over 30,000 RON in the construction sector were also ineligible for the personal tax breaks. In both 2018 and 2019 (pre and post-reform respectively) only 0.1% of construction sector workers reported a salary of over 30,000 RON.

The reason for the policy reform was to promote the construction sector in Romania[14]. As a relatively poorer EU country, construction workers have strong incentives to work in richer countries in Western Europe. Even after the reform which made Romania relatively more attractive, there are more Romanians working in construction outside of Romania than inside of Romania[15]. Around the time of the reform, the Romanian government had ambitions for large infrastructure projects but a shortage of workers due to the construction

[14]It was conceived as a ten-year strategy to last until the end of 2028, however for fiscal reasons the policy incentives were eliminated in January 2025.

[15]According to the following Romanian news website warning that in the absence of the policy, more tax evasion would occur https://www.romania-insider.com/construction-companies-warning-fiscal-allowances-2023

workers going abroad for higher salaries[16]. Thus, the policy was intended to make work in Romania's construction sector relatively more attractive compared with the outside options while also addressing the issue of undeclared earnings. A full list of the treated NACE codes is in the Appendix to Chapter 2, along with the female share of employment in each of the NACE sector codes and the share of the total construction employment that the code makes up.

2.2.3 Data

I use administrative tax records provided by the Romanian Ministry of Finance. The matched employee-employer data are monthly and cover 2018 through 2020. I use the data up to and including January of 2020. In the February of 2020 an issue of missing data prevents me from extending the length of the panel. For my analysis, I create a balanced panel, keeping only individuals whom are active in the labour market throughout the whole period and do not move between the construction sector and the rest of the private sector labour market.

Income reporting in Romania is done by the firm[17] (otherwise known as 'third party reporting'), which is what I observe in the data. Because the data are tax records, demographic variables are limited, but crucially include gender. The number of dependents is also observed, but is likely to be under-reported. Dependents increase employees' tax-free allowance, but there is a burden of providing documentation to prove that the employee has the dependents claimed. Otherwise, for each individual I observe monthly income, monthly hours of actual work, contracted hours of work, the worker's county (the NUTS3 level and called *judet* in Romanian) and locality (*localitate* in Romanian), and the firm that the individual works for. At the firm level, industry NACE codes are provided, which are used to identify whether firms have been treated by the policy. In Table 2.1 I present descriptive statistics of the full sample by gender and construction vs non-construction sectors. Women make up only 14% of the construction sector, in spite of having a much more equal share of overall labour force participation, and may thus be differentially selected in the sector compared to men.

[16] https://www.romania-insider.com/romania-construction-workers

[17] The tax system has been modernised and digitised, through a partnership with the World Bank: https://www.worldbank.org/en/news/press-release/2013/04/26/romania-modernizes-revenue-administration

	(1) Construction Sector			(2) Non-Construction Sector		
	Median	Mean	S.D.	Median	Mean	S.D.
Age	42.00	42.45	11.21	42.00	42.15	11.98
Hours Worked	160.00	145.36	49.12	160.00	147.80	46.39
Number of Dependents	0.00	0.27	0.66	0.00	0.34	1.01
Firm Size	27.00	118.25	248.33	97.00	1252.53	3316.14
Female Share of Firm	0.23	0.31	0.24	0.70	0.66	0.23
Full-Time Contract	1.00	0.92	0.28	1.00	0.90	0.30
Bucharest	0.00	0.17	0.38	0.00	0.15	0.36
Observations	1576734			64263673		

Panel (a): Women

	(1) Construction Sector			(2) Non-Construction Sector		
	Median	Mean	S.D.	Median	Mean	S.D.
Age	43.00	42.88	11.74	43.00	42.84	12.79
Hours Worked	160.00	145.62	47.52	160.00	151.74	40.91
Number of Dependents	0.00	0.29	0.72	0.00	0.33	0.79
Firm Size	30.00	114.80	241.06	103.00	1416.81	3691.08
Female Share of Firm	0.09	0.11	0.11	0.30	0.34	0.24
Full-Time Contract	1.00	0.97	0.18	1.00	0.92	0.27
Bucharest	0.00	0.09	0.29	0.00	0.13	0.34
Observations	9809794			63315734		

Panel (b): Men

TABLE 2.1. Descriptive Statistics: Selection into Construction Sector by Gender
Notes: Hours worked are measured monthly. Bucharest is a dummy variable for working in the capital city of Bucharest. The data covers from Jan 2019 up to and including Jan 2020. The public sector has been removed and the panel has not been balanced. The number of unique men and women is 4,355,720 and 3,869,582 respectively.

Women in the construction sector are quite observationally similar to women in the non-construction sector, and to men in the construction sector in their age, working hours and number of dependents. Women in the construction sector, however, are substantially more likely to live in the capital city of Bucharest compared to men in the sector and work for firms with substantially (almost three times) higher shares of women in the firm, implying there is some segregation between the kinds of firms that men and women work at. In the Appendix to Chapter 2 I present corresponding descriptive tables pre and post-reform and include information on reported salaries which is more interpretable when split across the reform rather than averaged across the pre and post-reform periods.

The Informal Sector

Using the Ministry of Finance provided tax register data, I am unable to observe Romania's informal economic sector. This is a potential concern as the construction sector is likely to be especially informal and, as the policy reform makes working in the formal sector more attractive (even if still under-reporting the true wage) I am unable to observe transitions

between the formal and informal sector. I also cannot distinguish the people who enter the sector, as they might transition from unemployment, from employment abroad or from the black market. The World Bank[18] provides cross-country estimates on the size of the informal economy using model based estimates and survey data, the details of which are described by Elgin et al. (2021). The percentage of informal output relative to GDP was estimated, using the range between the two most preferred estimation strategies[19], to be 26.6% - 29.2% in Romania in 2018. To put this into context, the informal sector's share of GDP was estimated to be 11.9% - 12.3% in the UK. Romania is comparable with other Central and Eastern European countries, however. The corresponding figures are 27.8% - 31.8% in Bulgaria, 23.3% - 24.7% in Poland and 16.1% - 16.6% in Slovakia.

According to the same World Bank source, this time using World Bank Enterprise Survey data from 2013, 95.9% of Romanian firms were formally registered when they began their operations in the country. This figure is similar to nearby Central and Eastern European countries.

Although my data do not cover the informal economy, coverage of organised firms is very high. The selected sample of workers who are operating formally (whether or not they under-report their true income) may make the men and women relatively more comparable to each other, in terms of their preferences for honesty, risk and for the benefits offered by being employed formally.

2.3 Theory

I analyse the theoretical effects of the policy changes under both a competitive and monopsonistic labour market framework. After generating the theoretical predictions, I then compare these with the effects of the policy in the data.

2.3.1 The Competitive Model

In this model there is a downward sloping demand curve, $D(w)$, and upwards sloping supply curve. Workers make their labour supply choice based on the post-tax wage, so with a proportional tax on their wages their supply function is $S((1-t)w)$. For graphical simplicity, I assume that these are linear functions, where changes in the income tax rate cause the labour supply curve to rotate.

A decrease in the income tax rate rotates labour supply outwards, as workers are willing to supply their labour for a lower *gross* wage, since the net wage is mechanically higher. The resulting change in the free market equilibrium is thus one with a higher quantity of labour being employed in the market, and at a lower gross wage. The increase to the minimum wage, however, constrains the demand for labour and the resulting equilibrium post-reform would be one with a contraction of employed labour. Thus, the theoretical

[18] https://www.worldbank.org/en/research/brief/informal-economy-database

[19] A Dynamic General Equilibrium estimation strategy and a Multiple Indicators Multiple Causes Model Based estimation strategy

prediction under a competitive framework is a decrease in construction sector employment at the higher minimum wage.

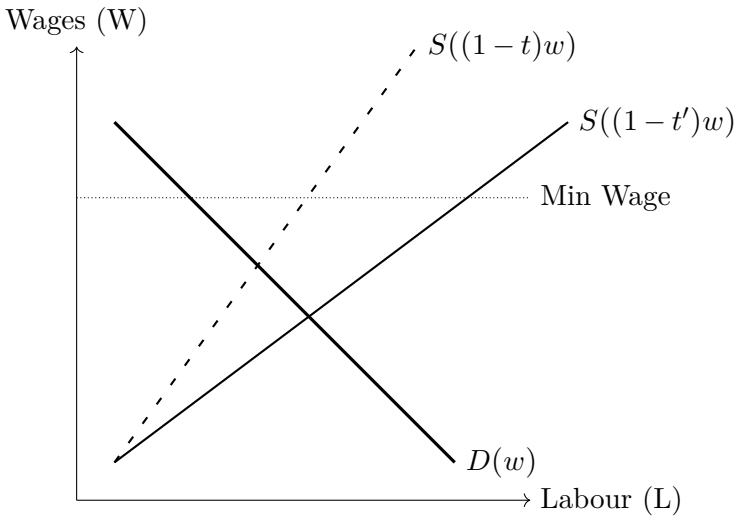


FIGURE 2.3. Competitive Model

2.3.2 The Monopsonistic Model

Since, as I will show later, the competitive model is inconsistent with the post-reform expansion of employed labour in the construction sector, I move onto a non-competitive model of a monopsonistic employer. The labour supply curve is $N((1-t)w)$, which can be re-written as $w(N,t)$.

The firm problem is:

$$\pi = Y(N) - w(N,t)N$$

Solving the problem by choosing the quantity of workers yields the FOC:

$$\underbrace{Y'(N)}_{MRPL} = \underbrace{w(N,t) + w_N(N,t)N}_{MCL}$$

This can be written equivalently with respect to the *net-of-tax rate*:

$$\underbrace{Y'(N)}_{MRPL} = \underbrace{w(N,1-t) + w_N(N,1-t)N}_{MCL}$$

I now analyse the effects of the policy under monopsonistic competition.

A decrease in the income tax rate (an increase in the net of tax rate) decreases the MCL. The gross wage which must be paid to the marginal worker is lower, since it is the net wage which matters to labour. By the same logic, the required increase in the wage which

applies to all of the firms' labour is lower. The firm does not need to increase the gross wage as much, since the net wage is mechanically increased. The RHS of the equilibrium condition has decreased, and since $Y'(N)$, is a decreasing function, equilibrium employment must increase.

Now suppose there is a binding minimum wage above the free market equilibrium. Denote the free market equilibrium wage as w^* with an associated MCL* and N^*.

Given the minimum wage is above w^*, there will be more labour willing to be supplied at w^{min} since $N((1-t)w^{min}) > N((1-t)w^*)$. The firm can thus expand its output to $Y(N^1)$, by paying all workers at the minimum wage. To expand output above this level, wages would need to rise above the minimum wage, which would jump the costs away from the minimum wage and back on to the MCL curve. Since MCL is increasing in N and MRPL is decreasing, and $N^1 > N^*$, this would be loss making, and would therefore not happen. The firm would operate at w^{min}, N^1 and $Y(N^1)$ so long as MRPL>= w^{min}. If at N^1, MRPL< w^{min} the firm would re-optimise and hire at N^2 such that $\underbrace{Y'(N^2)}_{MRPL} = \underbrace{w^{min}}_{MCL}$. The graphical intuition of this logic is presented in Figure 2.6.

Monopsony Model - Graphical Intuition

The logic of the model can also be made graphical.

A simple linear example starting with the labour supply function:

$$\begin{aligned} N &= \lambda w(1-t) \\ \iff w &= \frac{N}{\lambda(1-t)} \end{aligned}$$

This means the marginal cost of labour is:

$$MCL = \frac{N}{\lambda(1-t)} + \frac{N}{\lambda(1-t)} = \frac{2N}{\lambda(1-t)}$$

Revenue function:

$$\begin{aligned} Y(N) &= AN^{\alpha} \\ \implies Y'(N) &= A\alpha N^{\alpha-1} \end{aligned}$$

For graphical purposes I set $\lambda = 1$, $A = 20$ and $\alpha = 0.5$

I start with a 'high tax' case, setting $t = 0.4$

The model can be visualised in Figure 2.4.

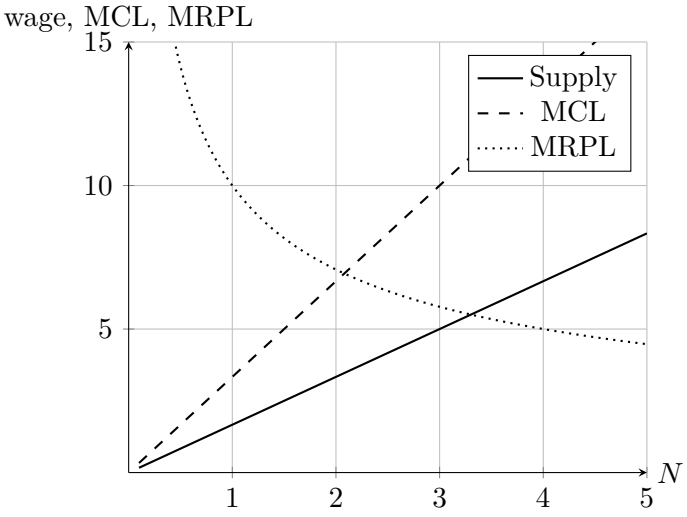


FIGURE 2.4. Monopsony Model: 'High Tax' case

The 'low tax' case, setting $t = 0.2$ is presented in Figure 2.5.

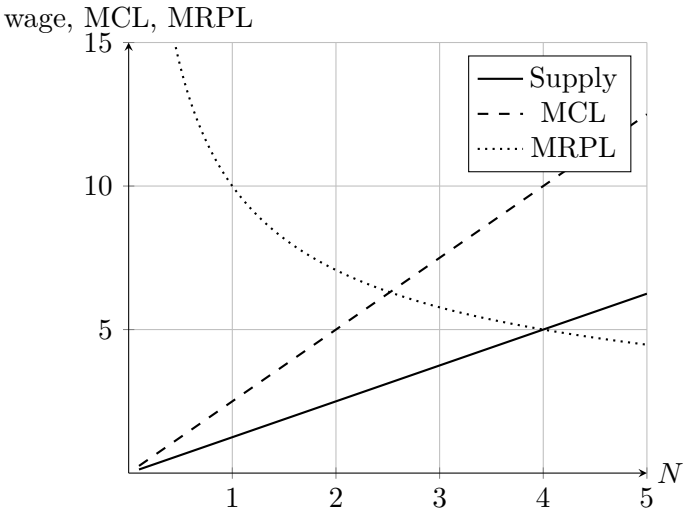


FIGURE 2.5. Monopsony Model: 'Low Tax' case

We can observe that lowering the tax rate rotates the labour supply and MCL curves rightwards, expanding equilibrium employment and decreasing the equilibrium gross wage.

Specifically, the previous equilibrium was $(N, w) = (3^{2/3}, 5 \cdot 3^{-1/3})$ while the new one is $(4^{2/3}, 5 \cdot 4^{-1/3})$

Finally, I present the employment effects of the minimum wage. Beginning with the free market equilibrium, the firm sets MRPL equal to MCL and employs labour at point A on Figure 2.6 with equilibrium wage w^a. A minimum wage of w^b, above the free-market equilibrium, would increase equilibrium employment, to point B. This is because the marginal worker is cheaper, only requiring to be paid that minimum wage, and the wage for all others in the firm is already set to that minimum wage and would not need to be increased. At a higher minimum wage, w^c, workers are still willing to supply their labour, but due to the

decreasing MRPL the firm would not wish to employ all those who are willing to supply their labour, and instead set employment to point C.

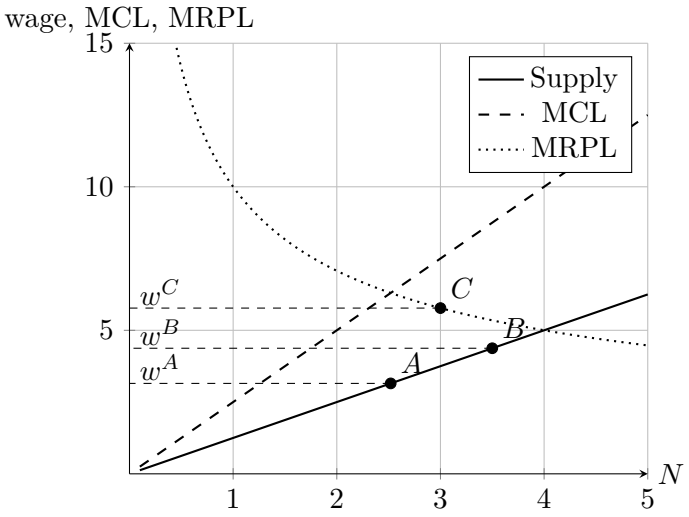


FIGURE 2.6. Monopsony Model: Minimum Wages above Equilibrium Wage

2.3.3 Model Takeaways

In a competitive labour market framework, a decline in the income tax rate would unambiguously decrease the equilibrium gross wage and expand employment. However, this effect would be nullified by a rise in the minimum wage above the previous equilibrium. The effect of the minimum wage would be an unambiguous decrease in employment, along with the mechanical rise in the wage.

Under a monopsonistic framework, a decline in the income tax rate would also unambiguously decrease gross wages and expand employment. The effect of the minimum wage hike, however, is more ambiguous. If the new minimum wage is not too large compared to the equilibrium wage, then it can have the effect of increasing employment, in contrast to the competitive framework. Once the minimum wage is sufficiently large compared to the free-market equilibrium, employment begins to decline.

2.4 Empirical Analysis - Evidence of Tax Evasion

2.4.1 Competition or Monopsony?

Because the predictions of the effects of the policy depend on the labour market setting, I first provide empirical context of the competitive landscape in the construction labour market. I use a Herfindahl–Hirschman Index[20] (HHI) to give local-industries a competitiveness score. To define a local industry I use 4-digit NACE sector codes within each of Romania's 42[21] Counties. While HHIs are often used in product markets to assess market concentration, I

[20]The index is described by the Federal Reserve in Rhoades (1993), which also details the history of the index.

[21]Including the capital city which is its own county

instead use market shares of employment rather than output to assess labour market concentration. Specifically, for each firm i in a local-industry j, I calculate their market share of employment:

$$MS_{i,j} = \frac{\#\ Employees_i}{\#\ Employees_j}$$

I then calculate the HHI for each local-industry j using the formula:

$$HHI_j = \sum_{\forall i \in j} (MS_{i,j})^2$$

For each local-industry, I thus have a measure of market concentration, where an index value close to zero implies a relatively competitive market with many small firms and a value close to one implies a more concentrated market dominated by one or a small number of very large firms.

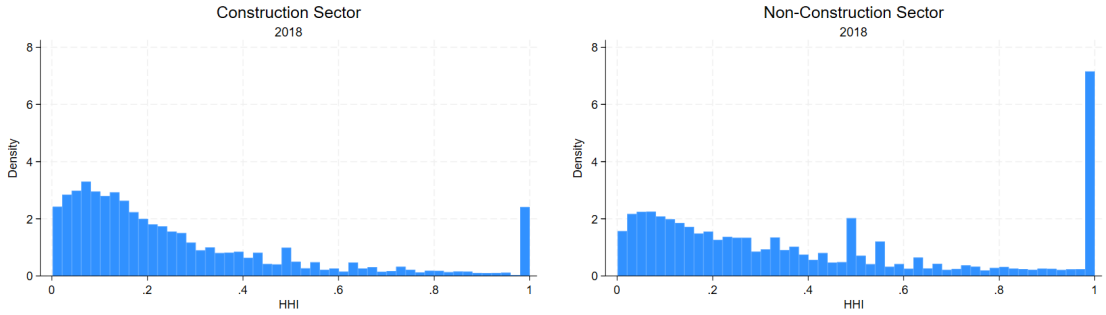


FIGURE 2.7. HHI: By Sector

Notes: There are 36 NACE codes in the construction sector and 569 in the non-construction sector. On average in 2018 each 4-digit NACE sector in construction had 11,882 employees and each NACE sector in non-construction had 9,031. On average in 2018, the number of firms in each 4-digit NACE sector was 1,304 in construction and 748 in non-construction.

As shown in Figure 2.7, 4-digit NACE sectors classified as construction compared to non-construction are much less concentrated. Empirically, this lends more support to a competitive labour market framework compared to a monopsonistic one. However, the expansion of employment in the sector runs in contrast to the competitive predictions.

An alternative way to measure the labour market competitiveness of an industry would be to look at mobility rates of workers within the local-industry. The logic being that there may be only a handful of relatively large firms which employ all of the workers in that local-industry, but if it is easy for workers to switch between firms then the market is effectively competitive, with workers switching to the firm which offers a higher wage and thus resulting in workers receiving close to their full marginal product, with only small wage markdowns. I show in the Appendix to Chapter 2 that the implication of the HHI scores, that the construction labour market appears empirically more competitive than the non-construction

labour market, is robust to using labour flow rates as an alternative measure of labour market competitiveness.

The logic underpinning the introduction of the policy reforms also supports a view that the labour market for construction workers is competitive. The intentions of the policy were to promote construction work and especially to incentivise Romanian workers who were working abroad to return to the country, by making work in the sector relatively more attractive. This of course implies that there are labour market shortages and thus that workers should have strong bargaining power with employers, taking home the bulk of their marginal product.

2.4.2 Policy Effects on Employment

The competitive labour market framework predicted a decline in construction sector employment, while the non-competitive model has an ambiguous prediction, whereby the size of the minimum wage hike determines the employment effects. Minimum wages which are not too far above the free-market wage lead to an increase in employment, while a sufficiently large minimum wage relative to the free-market wage will result in a decline in employment.

In Figure 2.8 I show that employment in the construction sector in fact grew following the reform in January 2019, from just below 8 percent of total employment to just under nine percent of total employment. This contradicts the predictions of the competitive model and could be consistent with monopsony, whereby firms expand following a new minimum wage. However, growth in the sector is not the result of firm expansion. Firms, in fact, do not expand at all. All of the growth in this sector comes from the switching of firms' industrial classifications from non-construction to construction as well as from the entry of new firms into the construction sector. It is unclear where the new firms come from. They could be made up of workers from abroad returning to Romania, as intended from the policy. Instead, they may be previously informal workers, either previously doing construction, or something else.

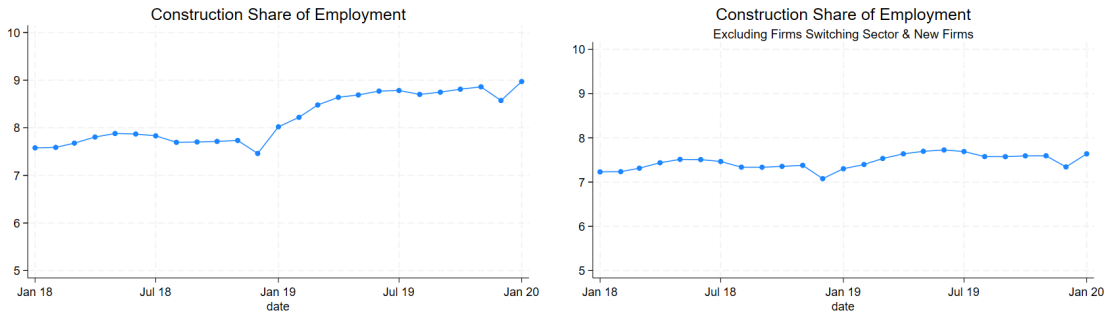


FIGURE 2.8. Share of Employment in the Construction Sector Over Time

Notes: The figure on the left shows the proportion of employees in the construction sector each month from January 2018 until January 2020, while on the right is the same figure but excluding firms which have ever changed their 4-digit industry classification or newly registered after 2018.

Overall, these empirical results are inconsistent with the competitive labour market model. Combined with the fact that the labour market for construction workers appears

to be competitive, with a large number of very small firms, this is consistent with income under-reporting, whereby the higher minimum wage for construction workers did not bite, since many of these workers take home a higher wage than is mandated by the new legal minimum.

The employment effects could be consistent with a non-competitive labour market model, but this is unlikely to be the case. Firstly, the market appears empirically to be highly competitive. Secondly, the CMW is substantially higher than the regular NMW - approximately 50 percent higher, which makes it difficult to believe that this is a sufficiently small minimum wage hike to increase employment or have no effect. Thirdly, if wages genuinely increased from the regular NMW to the CMW, this would imply that wages were marked down relative to the marginal product of labour by at least 1/3. This is a substantially higher mark down of wages than the one estimated in the US construction labour market[22], of about 1/5[23] (Kroft et al., 2020).

2.4.3 Policy Take-Up

One important feature of the policy was that is only applied to employees in firms which self-report[24] that at least 80 percent of their revenues are derived from construction activities. This means that the policy is manipulable and potentially non-binding. If it were the case that firms wanted to avoid the increasing wage bill, they could falsely report just under 80 percent and no longer be subject to the higher CMW. In Figure 2.9 I show that there is no evidence of manipulation at the cut-off.

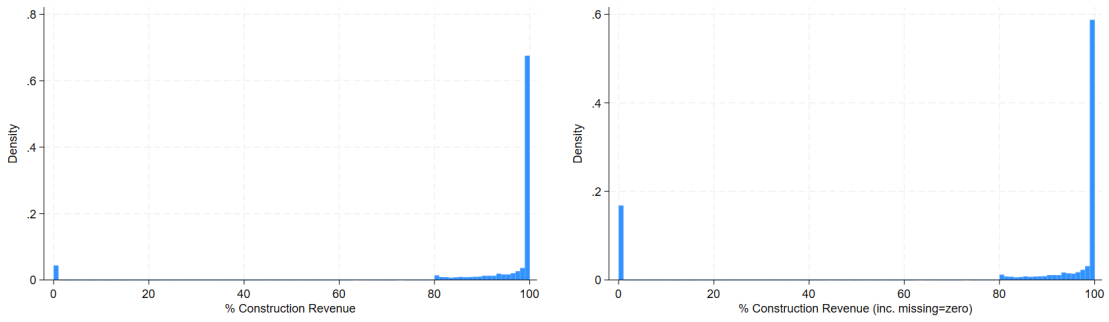


FIGURE 2.9. Take-Up of Policy

Notes: The figure on the left shows the distribution of the reported percentage of revenue which is derived from construction activities. On the right hand side is the corresponding graph with zeros filled in for missing reports for construction sector firms.

The majority of firms with a construction NACE code report that 100 percent of their revenues are in fact derived from construction activities. In contrast to what may be a priori expected, there is no bunching to the left of the cutoff value of 80, but rather a small amount

[22] Using exogenous demand shocks shifting treated firms along the labour supply curve.

[23] An almost identical markdown of wages is implied by the labour supply elastities estimated by Lamadon, Mogstad, and Setzler, 2022 for the whole U.S. labour market

[24] In January 2024 a digitisation process became implemented through the 'e-Factura' programme to improve invoicing transparency.

of bunching to the right of 80, indicating that taking up the policy is attractive to firms and workers. A small number of firms report a value of zero for their percentage of revenues derived from construction activities, which can be explained as mistakes in form-filling (the alternative explanation, that this is an overly confident method of avoiding the policy, is difficult to believe). There are also some firms who leave the reporting of this value as missing, which if these missing values are interpreted as zero, is presented on the right hand side of Figure 2.9. Combining the reported values of zero with the missings still makes up less than 20 percent of employees in the construction sector. Overall, it appears that the policy represents a net positive for firms and workers in the construction sector, and there is almost no evidence of manipulation to avoid the take-up of the policy.

2.4.4 Direct Evidence of Tax Evasion

I first present evidence of tax evasion in Romania, specifically the 'grey' form of evasion where employees and employers collude to falsely report the minimum wage on paper, on which taxes and social contributions are paid, while the remainder of the salary is paid cash in hand. In Figure 2.10 I present the proportion of employees reporting exactly the minimum wage in the construction sector and the non-construction private sector, by the size of the firm which they work for. It is immediately apparent that smaller firms are substantially more likely to report exactly the minimum wage, consistent with the logic that the employee-employer collusion is harder to sustain in larger firms. In the construction sector, reporting exactly the minimum wage is more prevalent than in the non-construction private sector in 2018, consistent with the sector being one with a higher prevalence of tax evasion. However in 2019 this behaviour is almost eliminated, as the tax minimising gross salary to report is no longer the minimum wage, but the higher value of 3000 RON. The construction sector employees reporting the regular NMW in 2019 are those working in firms which do not meet the criteria for the take-up of the policy, due to a lack of awareness or purposeful manipulation.

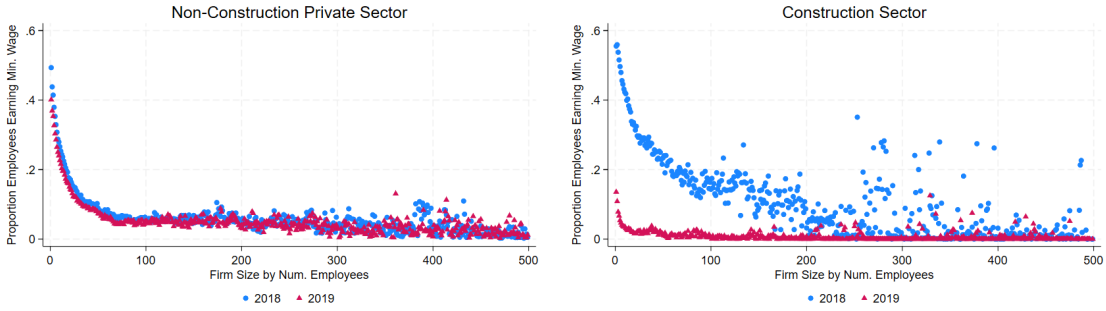


FIGURE 2.10. Proportion of employees reporting a salary of exactly the minimum wage

Notes: These graph presents the proportion of employees reporting a salary of exactly the minimum wage for firms of different sizes, up to a firm size of 500 employees. The data is split into two calendar years and the construction sector treated by a policy reform in 2019 and the non-construction private sector. There are no sample restrictions other than the exclusion of the public sector and firms with over 500 employees.

Prior to the reform[25] 50% of Romanian construction workers overall were reporting that they earn exactly the minimum wage, while this figure dropped to just 16% post reform. In the non-construction sector, the figure remained stable at around 25% pre and post-reform. In the Appendix to Chapter 2 I present the wage distributions for the construction and non-construction private sector, both pre and post-reform. The distribution of reported salaries shows a jump in excess mass from the regular NMW to the CMW of 3000 RON in the construction sector, but no such jump in the remainder of the private sector.

2.4.5 Individuals Report Above the Minimum Wage

Following the reform, the optimal way to minimise the tax burden changes, from reporting the regular NMW, to now reporting a salary of 3000 RON, at the CMW. Of course, this would require that employees are being paid a salary of 3000 RON, and previously they were making at least this salary, as a combination of the reported salary and the payment under the table. To first test whether individuals are more likely to report a salary above the NMW in the construction sector, after the reform, I run the following difference-in-difference regression:

$$y_{i,t} = \beta_0 + \sum_{\forall s/-1} \beta_s \mathbb{1}\{treated_{s=t}\} + \alpha_i + \alpha_t + \sum_{j=1}^{n} \delta_j x_{i,t}^j + \epsilon_{i.t}$$

In this regression, β_0 is an intercept, the series of β_s coefficients are the difference-in-difference treatment effects estimated relative to the base period, one month prior to the policy's introduction. Individuals are treated if they work in a construction sector (defined by having an eligible NACE code). α_i and α_t are the individual and time fixed effects and the series of coefficients δ_j are included for the control variables x^j.

As an outcome variable, I use an indicator for the individual earning a salary above the regular NMW. The results, in Figure 2.11, indicate that following the reform, construction workers are almost 30 percentage points more likely to report a salary above the minimum wage. Before the reform, in the balanced sample I use here, 66 percent of construction workers were reporting a salary above the NMW, meaning that approximately one-third of workers were reporting at exactly the minimum wage before the reform, but once the incentives to under-report at the minimum wage were removed, almost all of these workers began to report a higher salary.

This form of tax evasion should be less prevalent in large firms, where collusion between employees and employers is less likely to be sustained, so I run my event study specification using very large firms only, specifically those in the top 0.5% of the firm size distribution (which are firms with above 203 employees). The results in Figure 2.11 show that the reform has almost no effect on these firms in reporting above the NMW.

[25] In the full sample, not the balanced panel which removes individuals who move between the construction and non-construction private sector during the sample period.

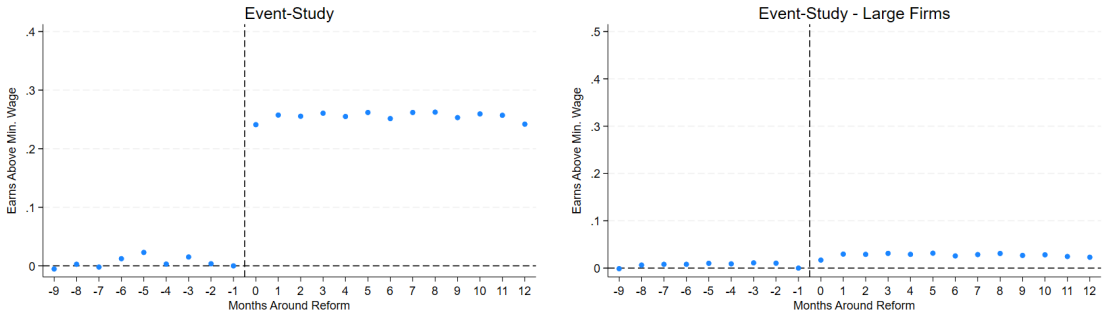


FIGURE 2.11. Event-Study - Reporting Above the Minimum Wage

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 3,761,657 individuals in each of the 22 months, of which 1,766,113 individuals are in very large firms. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

Since the reported salary which minimises the tax burden, following the reform, is 3000 RON in the construction sector, I run my difference in difference specification on an indicator of earning exactly 3000 RON. Following the reform, workers in the construction sector are much more likely to report at this CMW, which minimises the tax burden, but this effect is muted in large firms, where tax evasion as an employee-employer collusion is less likely to take place.

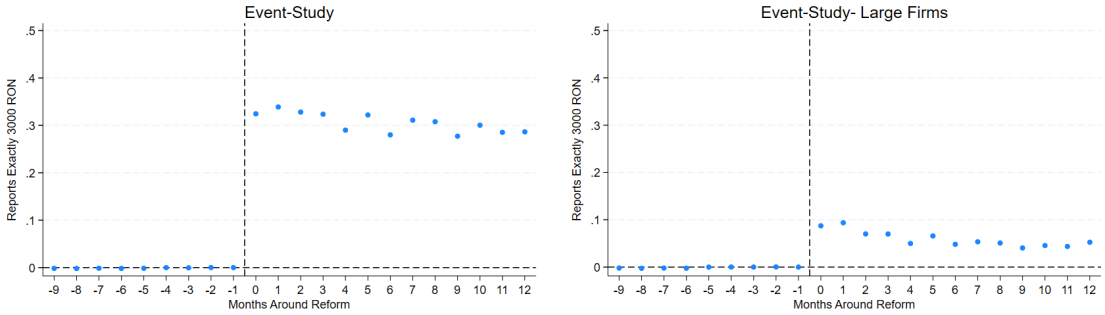


FIGURE 2.12. Event-Study - Reporting 3000 RON

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 3,761,657 individuals in each of the 22 months, of which 1,766,113 individuals are in large firms. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

As a further validation, I check the probability of jumping to a salary of 3000 RON following the reform for workers who were reporting exactly the minimum wage in the pre-reform period, and those who were on a wage close to, but slightly above (specifically, in the range 1906 to 2000 RON), the minimum wage. Presented in Figure 2.13, I show that those who were previously reporting exactly the minimum wage are about twice as likely compared to the average construction worker to jump to reporting a salary of 3000 RON, even though it is a larger jump in reported income than for those closer to 3000 RON, consistent with the

minimum wage being used as a tax minimising strategy rather than a true take home wage. Additionally, those previously reporting exactly the minimum wage are even more likely to jump to a report of 3000 RON than those earning marginally more than the minimum wage (labelled as "Almost Min. Wage Workers" in the figure), even though the individuals with a salary above the minimum wage have a smaller gap between their previous salary and the CMW of 3000 RON.

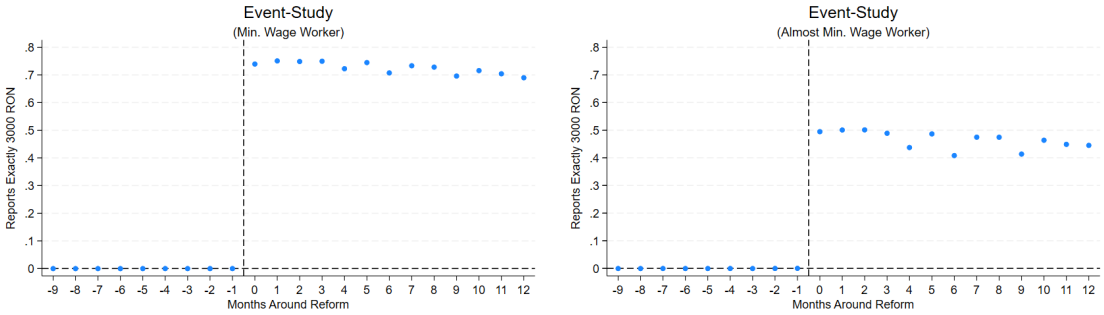


FIGURE 2.13. Event Study - Reporting 3000 RON by Pre-Reform Reported Salary

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 242,723 minimum wage workers (defined by earning the minimum wage in the pre-reform period) and 263,288 workers in the range of 1906 to 2000 RON (just above the minimum wage) in each of the 22 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

While the majority of movement to this new salary report comes from individuals who were previously reporting the minimum wage, and thus likely to be engaging in the 'grey' form of tax evasion, there are also some individuals between the regular NMW and the CMW who move to this reported salary due to an honest compliance with the policy, or for the tax advantages.[26] Jumps to 3000 RON are thus a combination of individuals falsely reporting the minimum wage before the policy introduction and switching to the new tax minimising salary report as well as some individuals who were not previously under-reporting their income but move to a new 'genuine' salary of 3000 RON due to compliance with the policies.

2.4.6 Joint Collusion

Given that under third-party reporting, evading income tax would require a joint employee-employer collusion, it is assumed that this is much easier to maintain in smaller firms rather than larger firms. To test this, I run the following 2x2 difference-in-difference regression across firms of different sizes:

$$y_{i,t} = \beta_0 + \beta_1 \mathbb{1}\{treated\} + \beta_2 \mathbb{1}\{post\} + \beta_{DiD} \mathbb{1}\{treated\} \cdot \mathbb{1}\{post\} + \sum_{j=1}^{n} \delta_j x_{i,t}^{j} + \epsilon_{i.t}$$

[26] There may also be salience or reference-dependence effects driving some employers to pay this salary, as Kleven (2016) suggests these effects may also explain excess bunching at kink points of income tax schedules.

The outcome variable is an indicator for the employee earning above the national minimum wage. The results in Figure 2.14 indicate that individuals in smaller firms are much more likely to jump up from reporting the NMW following the policy reform. The analogous graph for jumps to reporting a salary of exactly 3000 RON is in the Appendix to Chapter 2.

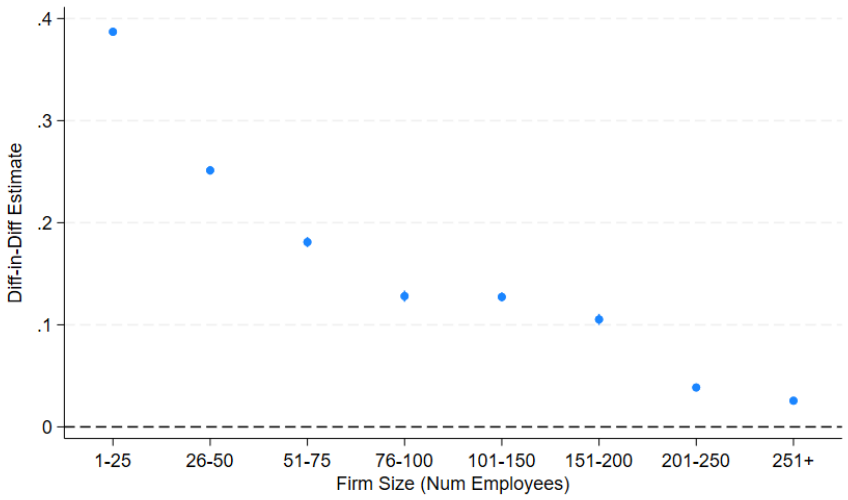


FIGURE 2.14. 2x2 DiD Estimates of Reporting Above the Minimum Wage by Firm Size
Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 3,670,969 individuals in each of the 24 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

This result is consistent with the assumption that workers collude with their employers to engage in 'grey' work, jointly reporting the minimum wage to the tax authorities and receiving the remainder of the salary cash in hand, but that this collusion is unsustainable in larger, more formal, firms.[27]

The issue of employee-employer collusion can be addressed directly by examining one-person firms. These are firms in which the employee is the owner and operator of the firm. In this case, all of the economic rents from economic activity are enjoyed by the single employee and the employee would have no incentive to mark-down his or her wages below the marginal revenue product of labour. We can see in Figure 2.15 that these one person firms react strongly to the policy, with sizeable jumps away from the regular NMW and towards the CMW.

[27]Larger firms may also be involved with larger infrastructure projects financed with government or EU funds which would require greater transparency and record keeping.

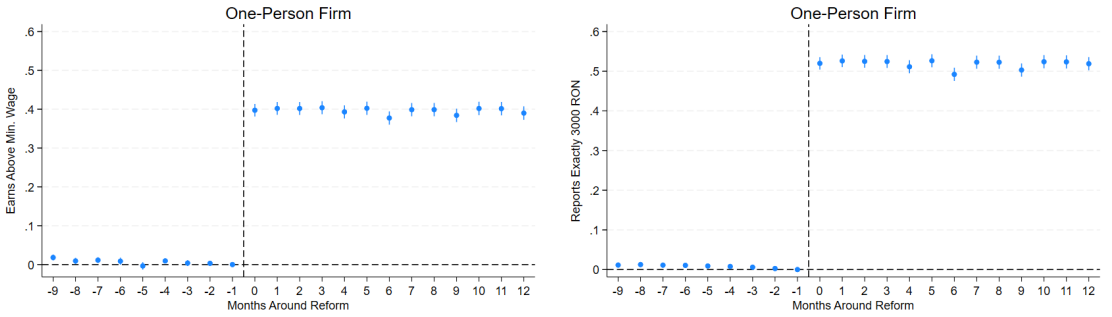


FIGURE 2.15. Event Study - One-Person Firms

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 124,505 individuals in each of the 22 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

In 2018, before the policy reform, 69 percent of employees in one-person firms (in the balanced sample) were reporting a salary at exactly the NMW. Given that an employee has no incentive to mark-down their own wages, it would be a puzzle to explain this behaviour outside the context of income under-reporting. One alternative would be that the one-person firm avoids tax, rather than evades, by paying out only the minimum wage as a salary and then paying out the remainder as dividends, at a lower dividend tax rate[28]. However, this does not seem to occur. In Figure 2.16 I show that as the median salary for those in these one person firms jumped from the old NMW to the new CMW, so too did the reported firms' incomes.

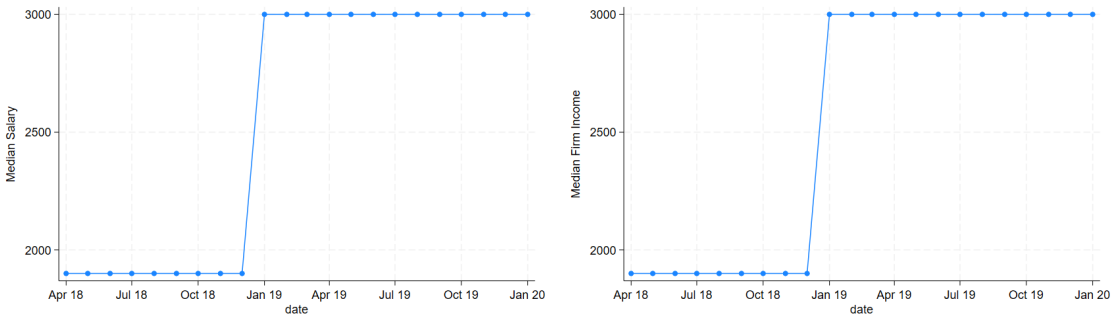


FIGURE 2.16. One-Person Firms: Median Salaries and Firm Revenues

Notes: The sample only includes individuals in one-person firms who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The graphs show median reported income and firm revenue of these one-person firms.

This is not an artefact which only appears at the median values of reported salaries and firm revenues, but is also apparent at the mean of the log of salaries and firm revenues, as presented in the Appendix to Chapter 2. The evidence is thus entirely consistent with, and difficult to rationalise in any other way than, these one-person firms reporting only

[28] Which is only 5%.

the minimum wage as the salary, and also as the firms' revenue, to cover the salary. The remainder of what the one person firm earns is kept as cash in hand and never observed by the tax authority. Following the reform, the reported salaries jumped by approximately 50 percent, as did the reported firm revenues.[29]

Interestingly, while one person firms could be considered a discontinuity in the relationship shown in Figure 2.14, whereby the jumps away from the regular NMW and towards the CMW is decreasing in the firm size, I identify a lower level of jumps in one person firms compared to marginally larger firms (shown in the Appendix to Chapter 2). This is unexpected, since with two employees and above the chances of collusion breaking down jump discontinuously from zero to some larger probability. This may be explained by selection into firm size, where relatively more honest workers select into working alone and avoiding dis-honest colleagues. It may also be an effect of working in a small tight-knit firm, where workers mutually reassure each other that they can collectively get away with under-reporting their true income. An alternative explanation is that there may be differences in strategic sophistication. Recognising that the total tax burden is lower at the substantially higher CMW than the regular NMW requires an understanding of new tax system post-reform. As shown in the Appendix to Chapter 2, one person firms are more likely to be reporting the percentage of their revenues which are derived from construction as zero and as missing, explaining how post-reform, they are less likely to to jump away from the minimum wage than marginally larger firms.

2.4.7 Other Corruption in Romania

To validate that my identification strategy is indeed picking up individuals who are tax evading, I compare, by Romanian counties, the tax evasion estimates to the hospital bribery rate. The logic to this exercise is that if in some counties attitudes are relatively more relaxed about tax evasion, the case may be similar for bribery, in which case these would be expected to correlate with one another. Bribery rates are calculated using Romanian government survey data from the 'Patients Feedback Mechanism'. I run my 2x2 difference-in-differences specification, estimating jumps in reported wage above the regular NMW separately for each of the Romanian counties, while the patients' feedback data provides me with the percentage of individuals who report that they have been asked for a bribe in a hospital in each county. The binscatter relating these is presented in Figure 2.17 and shows a positive, although weak, correlation between reported hospital briberies and my estimates of tax evasion.

[29] In the Appendix to Chapter 2 I show that in general small firms increase their reported revenue substantially more than larger firms in response to the policy.

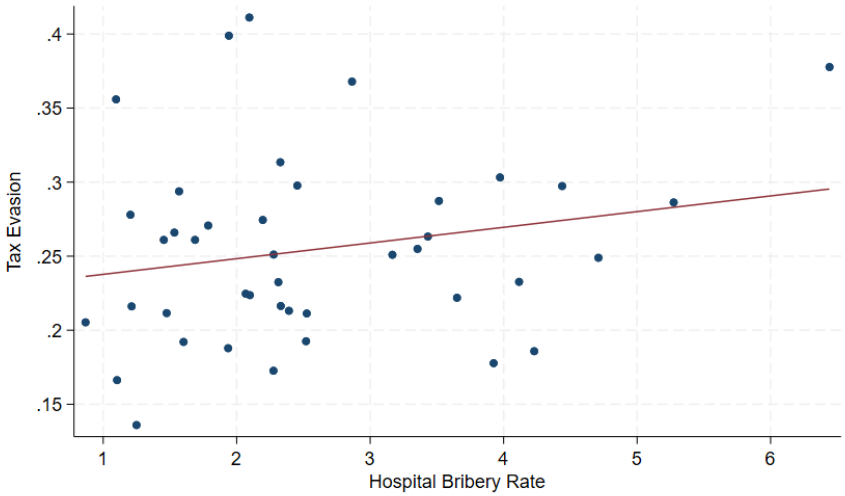


FIGURE 2.17. Tax Evasion vs Hospital Bribery Rates

Notes: On the x-axis I present the percent of surveyed patients who report that they have been asked for a bribe in a Romanian hospital, by Romanian county. The hospital bribery data uses the years 2021 through 2023. On the y-axis I present the post X treated DiD estimate for tax evasion, where the outcome in the regression is a dummy variable for reporting a salary above the minimum wage, by Romanian county. Each point in the scatter represents one of the 42 Romanian counties and a line of best fit for these points (linear regression) is overlayed. Hospital corruption data accessed from 'The patient's feedback mechanism', https://data.gov.ro/dataset/mecanismul-de-feed-back-al-pacientului-2023. The regression coefficient is 0.01 with a p-value of 0.18.

2.4.8 Productivity Growth

The policy reform could plausibly have positive effects on worker productivity in the sector, which would be reflected in wages. Due to the magnitude of the jump in construction workers' wages (the median wage jumped approximately 50 percent between 2018 and 2019), it seems implausible that this is reflecting gains in the marginal product of labour. In the Appendix to Chapter 2 I use Eurostat data to show that construction output in Romania only experienced mild growth in 2019 relative to 2018 and was comparable to the growth of the services sector. It is thus unlikely that construction workers jumping from a report of the minimum wage, onto a wage report at exactly 3000 RON, is reflecting jumps in productivity.

2.5 The Gender Gap in Tax Evasion

Under the interpretation that jumps in income away from the minimum wage and towards the threshold value of 3000 RON indicates that tax evasion is occurring, I run event-study difference-in-differences regressions separately for men and women to compare the rates of tax evasion. I use the following specification:

$$y_{i,t} = \beta_0 + \sum_{\forall s/-1} \beta_s \mathbb{1}\{treated_{s=t}\} + \alpha_i + \alpha_t + \sum_{j=1}^{n} \delta_j x_{i,t}^j + \epsilon_{i.t}$$

The results in Figure 2.18 show that men are just under 30 percentage points more likely to report a salary above the regular NMW following the reform, while women are just under 20 percentage points more likely to do so.

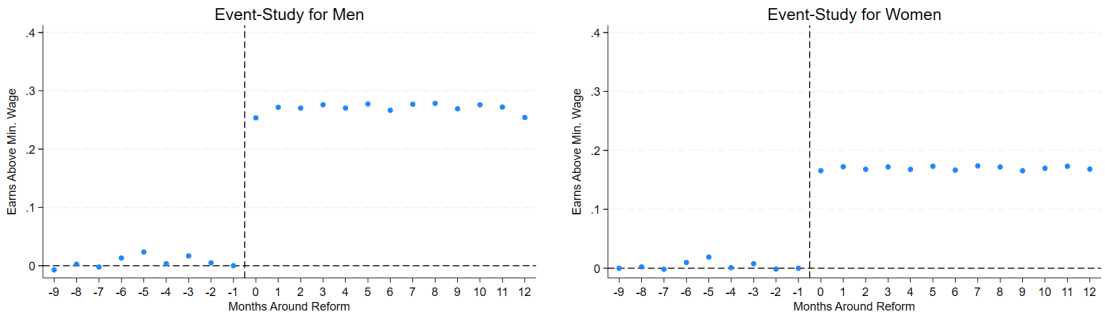


FIGURE 2.18. Event Study - Reporting Above the Minimum Wage

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 1,838,477 male individuals and 1,923,180 female individuals in each of the 22 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

The change in the tax minimising reported salary was moved by the policy, from the NMW to the CMW, 3000 RON. Figure 2.19 shows that men are substantially more likely than women to alter their reporting behaviour to this new level, although both men and women are slightly more likely to move to the CMW than to just increase their reported salary from the regular NMW.

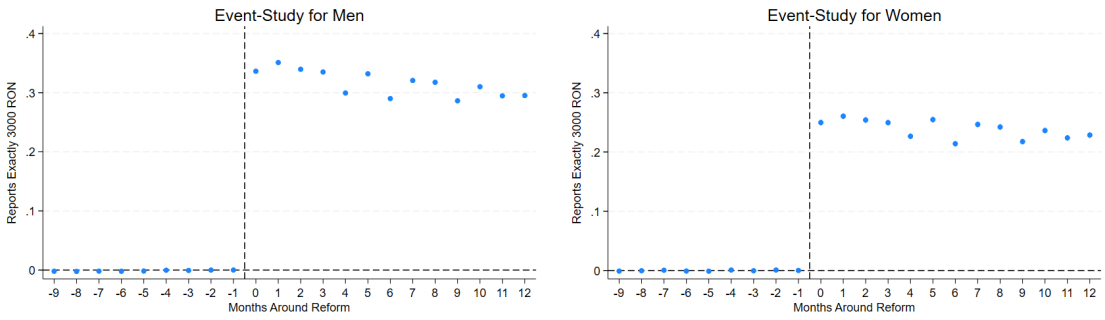


FIGURE 2.19. Event-Study - Reporting 3000 RON

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 1,838,477 male individuals and 1,923,180 female individuals in each of the 22 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

The evidence of a gender gap in tax evasion is clear, but the mechanisms generating this are difficult to establish. Although differences in preferences, such as preferences over risk and honesty, are one potential explanation, men and women may also have different abilities to under-report income for other reasons. Income under-reporting in Romania requires a

joint employee employer collusion, and women may be less willing or able to bargain with their employer to tax evade. A less sophisticated explanation than differences in bargaining over tax-evasion would be that men were just more likely to be genuine minimum wage earners pre-reform compared to women. It may also be the case that men work in less office-based, 'visible' occupations, such as on construction sites whereas women work in more visible and formal occupations, such as management of building projects or accounting for construction firms. There may also be segregation in the types of firms men and women work for which may generate differences in tax evasion. Empirically, women do work in firms with substantially larger shares of women than do men.

2.5.1 Bargaining Power

One (Wo)man bands

Firms with a single employee ('One Man/Woman bands') eliminate concerns about the bargaining between employee and employer over the decision to tax evade, since the single employee has full decision making power. In addition, a one (wo)man band would receive the entirety of the economic rent generated, rather than engaging in a bargaining process over this. As I showed earlier, one person firms tend to earn no profits, but rather report a firm income and a salary exactly equal to the NMW pre-reform, and exactly equal to the CMW post-reform[30]. The results, presented in Figure 2.20, again yield a gap between the implied under-reporting of income between men and women for the purpose of tax evasion. In fact, in these firms, the gender gap in tax evasion widens compared to the full sample. The widening of the gap comes mostly from men who are more likely to be tax evading in one man bands, while women are only marginally more likely to tax evade in these firms.

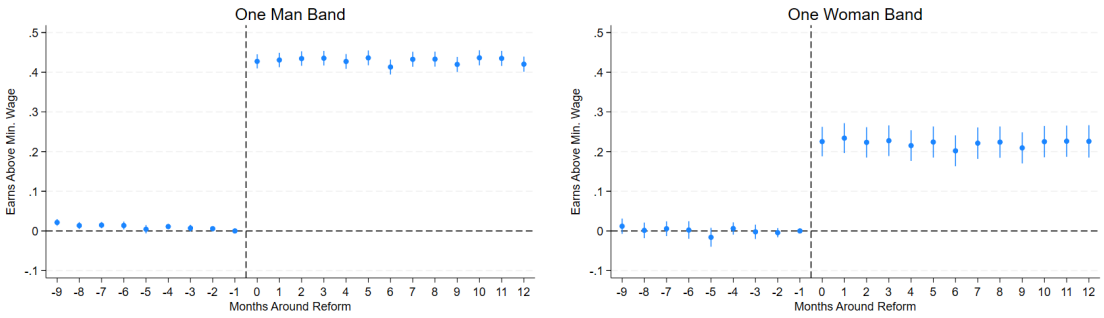


FIGURE 2.20. Event Study-Effects on One (Wo)man bands

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample contains 61,553 firms of one employee for men and 61,071 for women, of which, respectively, 5,623 and 1,208 are in the construction sector. The standard errors are clustered at the individual level and the controls are the contracted hours and the hours actually worked.

The change in behaviour for men may point to gender differences in lying aversion, whereby men have a higher tolerance for lying, but are less likely to do so in firms with

[30] In the Appendix to Chapter 2 I show that this is indeed the case for both men and for women, not just one-person firms overall.

other employees since this is more visible and has a higher chance of the collusion breaking down.

An approach to understanding the role of bargaining power in firms which have more than one employee would be to look at the gender of the boss. Since my data do not contain information on firm hierarchy, or even occupation, this is challenging to observe directly. This is made more difficult by the fact that classifying the highest earner within a firm as the boss is not possible in the cases when all workers in the firm report the same salary, which is very common when falsely reporting the minimum wage.[31] My approach to proxying for the boss of the firm is to assume that the oldest worker in the firm is the boss. While this may be unreasonable in large firms, for smaller firms it is more plausible that the oldest worker is the boss, or has some kind of authority in the firm due to seniority. I choose to look only at firms with five or fewer employees and classify the boss as the eldest worker in that firm. Then, using the gender of the boss, I can classify the firm as having either a male or female boss. Excluding the boss themselves, I then run my regressions on the subsample of employees in firms with five or fewer employees to examine the heterogeneity arising from the gender of the boss on the male and female employees in these firms.

The results of these regressions are presented in Figure 2.21 and Figure 2.22. It appears that the gender of the boss does matter, especially for the change in reporting behaviour of women. Compared to those with a female boss, women with a male boss are approximately three times more likely to have been previously reporting a salary at exactly the minimum wage and then jump to a higher salary post-reform.

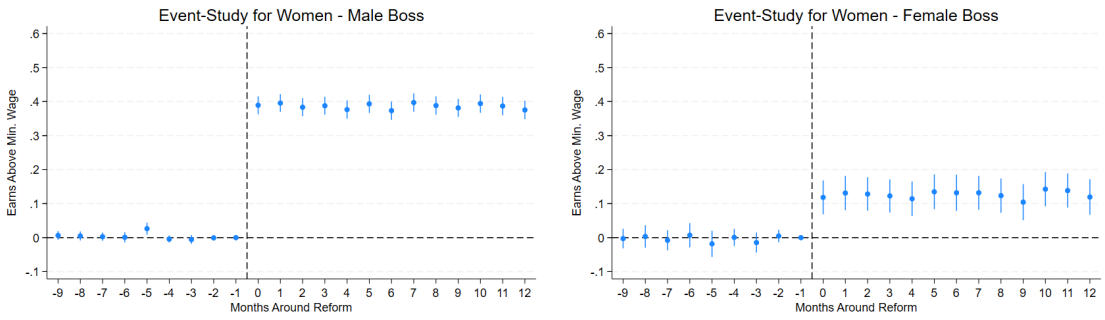


FIGURE 2.21. Event Study-Effects on Women by Boss's Gender

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. It only includes employees in firms with five or fewer employees, and excludes the boss of these firms, which is defined to be the individual with the highest age. The sample contains 50,380 men with female bosses and 95,097 women with male bosses. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

For male workers, the results are substantially more subdued. Those with a male boss,

[31] In the Appendix to Chapter 2 I show the proportion of employees within a firm whom are earning the minimum wage. This distribution is bi-modal, whereby in the majority of firms either all employees are reporting exactly the minimum wage or none of the employees are.

compared to a female boss, are more likely to have been previously reporting a salary at exactly the minimum wage and jump to a higher salary post-reform, but with a lower differential effect of the boss's gender than what is seen for female employees.

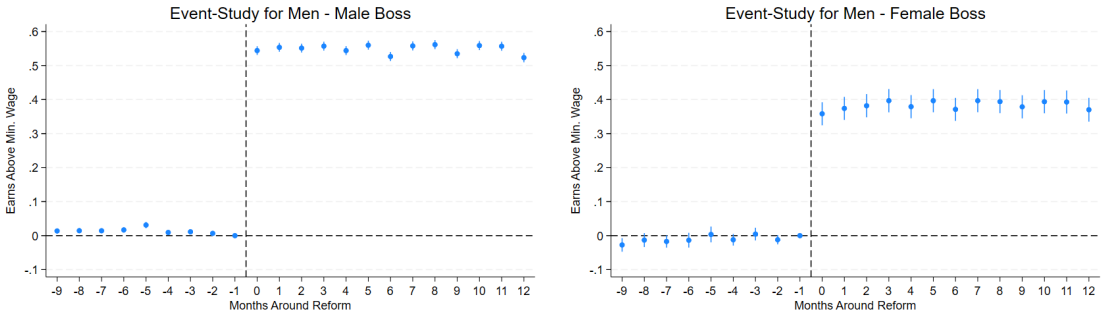


FIGURE 2.22. Event Study-Effects on Men by Boss's Gender

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. It only includes employees in firms with five or fewer employees, and excludes the boss of these firms, which is defined to be the individual with the highest age. The sample contains 91,150 women with female bosses and 112,000 men with male bosses. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

Overall, it seems that the bargaining power of an employee does matter and that male bosses are more likely to allow their employees to under-report their income than female bosses. At the same time, the results of the one man and one woman-bands suggest that bargaining power does not explain all of the difference in reporting behaviour between men and women and that preferences may play a role in the gender difference.

2.5.2 Occupational Segregation

The gender gap in propensity to tax evade may be driven by men and women working in different kinds of occupations, within the same industry of construction. Because the administrative tax data does not contain information on the occupation, I try to infer whether or not work is manual or office-based used information which is available in the dataset, as well as focus on one specific NACE industry code within the treated sector which involves a specific and gender neutral office-based type of work.

Occupational Segregation: Implied Occupation by Hours Worked

Manual work in a construction site is much more seasonal than office-jobs based in an indoors office space. I thus use individual-level variation in monthly hours worked to classify those with highly variable working hours as manual workers and those with less variable working hours as office based, addressing occupational segregation between men and women. If the strategy of using variation in working hours to separate office-based vs manual work is indeed

picking up these different types of occupations, it would be expected that the income under-reporting implied by jumps in reported income is higher for the manual work than for the office-based work.

We can see in Figure 2.23 that office-based workers, defined as having a variance in monthly working hours below the 10^{th} percentile of the distribution[32], have relatively lower implied tax evasion than my main specification presented in Figure 2.18, consistent with it being more difficult to evade taxes in a more formal office-based environment, but that the gender gap in tax evasion, although somewhat subdued, remains.

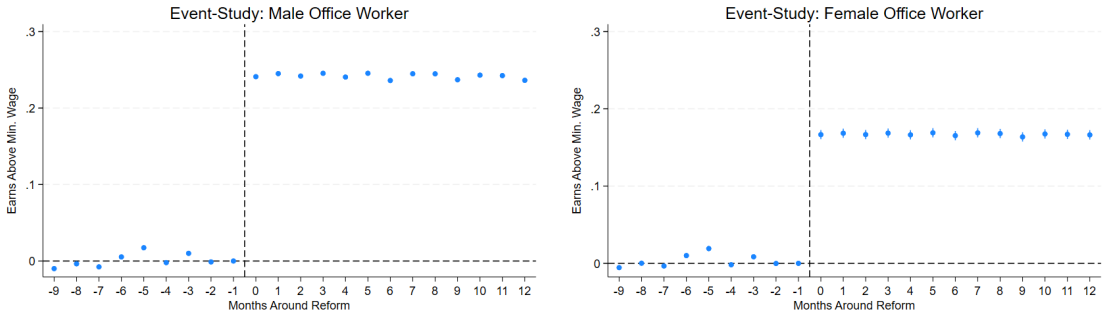


FIGURE 2.23. Event Study-Effects on Office-Based Workers

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 1,021,741 male individuals and 956,322 female individuals in each of the 22 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

We can conversely observe that for manual workers, defined by having monthly working hours variance above the 90^{th} percentile in the distribution, the tax evasion implied by my identification strategy is relatively high for both men and women. This is reassuring for my definition of manual workers picking up individuals who are indeed doing more work on construction sites and have greater ability to evade taxes. Indeed, within this category of occupation the gender gap in identified tax evasion remains. The results are presented in Figure 2.24.

[32]This value is 96 and is also the median variance in individual hours worked. This is because the variance simply comes from variation in the number of working days in the calendar month, with these individuals otherwise having a simple flat working schedule.

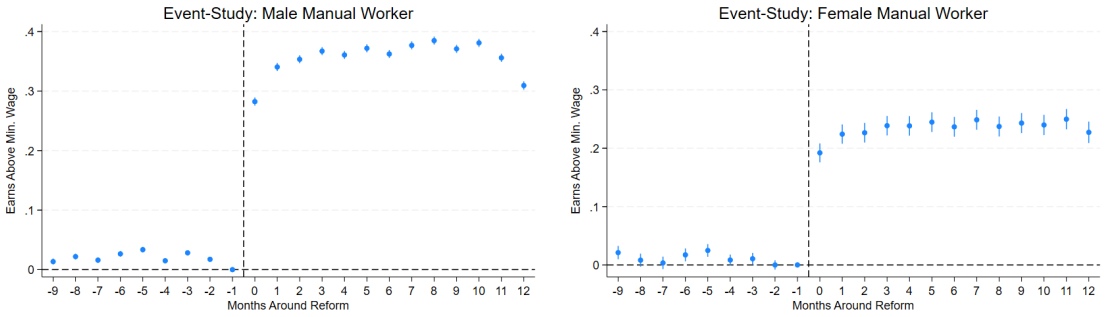


FIGURE 2.24. Event Study-Effects on Manual-Based Workers

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 186,418 male individuals and 192,606 female individuals in each of the 22 months. Within the treated construction sector figures are 22,268 and 2,975 respectively in each of the 22 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

An alternative way to implicitly separate manual workers from office-based workers, rather than using the simple within-individual variance in hours worked, is to make more explicit the type of seasonality that would be expected for manual workers in a country where cold winters makes working on a construction site more challenging. I create another definition of a manual worker by constructing, for each individual, the ratio of hours worked in the months of December, January and February to the hours worked during the rest of the year. A low ratio means that hours dip in the colder months relative to the rest of the year. Equipped with a distribution of values of this ratio, I define my second indicator of a manual worker as a worker with a ratio under the 10^{th} percentile of this distribution. The results are similar to what I present in this section and can be found in the Appendix to Chapter 2.

The results presented for office workers and manual workers so far suggest that men appearing to have a higher propensity to tax evade than women is not primarily driven by segregation in the kinds of work that men and women do.

Office-Based Work: Implied by Industry Sector

My previous definitions of office and manual work were based on reasonable but imperfect assumptions about differences in seasonality between the different kinds of work. As a further check of whether or not the gender gap is driven by differences in the kinds of work that men and women do, I run my preferred difference-in-difference specification using only "building project development"[33] as a treated industry, and comparing it with the non-construction private sector. This industry classification implies primarily office-based work. It is a treated profession which is relatively gender-neutral (34% female) and more comparable with other

[33]This is the 4 digit NACE code: 4110. The description for the code is the following: development of building projects for residential and non-residential buildings by bringing together financial, technical and physical means to realise the building projects for later sale.

private sector jobs than jobs in construction involving manual labour. The number of individuals I observe who work in this NACE code over my sample period is 3,948. The results of the difference-in-difference are presented in Figure 2.25.

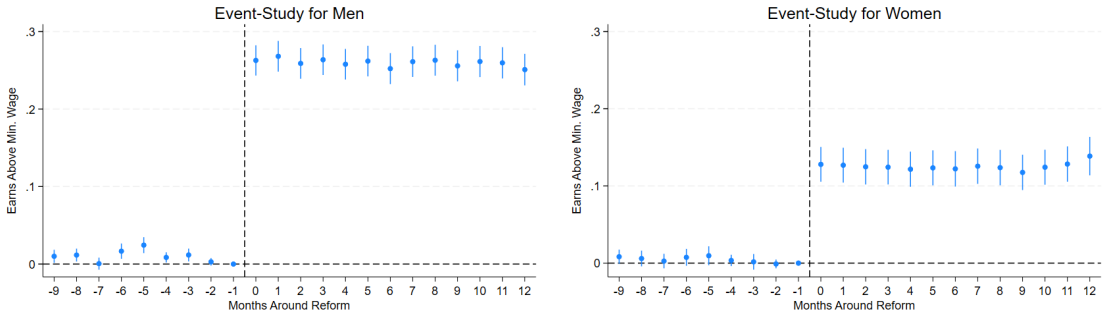


FIGURE 2.25. Event Study - Project Development Industry Code

Notes: The specification in the regressions includes controls for actual hours worked, the number of employees in the firm and the type of contract (full or part-time). The number of individuals I observe who work in this NACE code over my sample period is 3,948. Standard errors are clustered at the individual level and 95% confidence intervals are plotted around the point estimates.

The results for the "building project development" profession are, as expected, marginally weaker than for the construction and manufacturing sector as a whole, since the industry is more formal/office-based than traditional construction jobs. The gender gap in reporting behaviour remains, with men substantially more likely to jump away from reporting a minimum wage salary post reform compared to women.

2.5.3 Firm Composition

In the wider Romanian economy, and especially in the construction sector, men and women may be segregated. Besides entering different industries, men and women may enter different kinds of firms and have different kinds of occupations within those firms. Since the men in construction work in relatively more male dominated firms compared to the women, I run my preferred specification on female dominated firms, defined as having the female share of employment strictly above 50 percent. The results in Figure 2.26 suggest that segregation in the types of firms that men and women work at explain a substantial part of the gender gap in behaviour that is consistent with tax evasion, since the proportion of employees who jump away from minimum wage reporting is smaller in these firms for women, compared with the full sample of all firms, but substantially smaller for men in these firms, who are in fact marginally less likely to exhibit this behaviour than women in these female dominated firms.

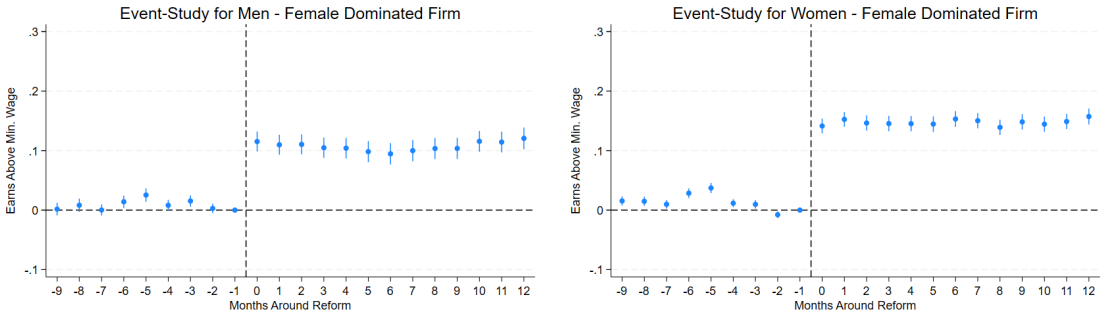


FIGURE 2.26. Event Study - Female Dominated Firms

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 520,639 male individuals and 1,433,371 female individuals in each of the 22 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

This result is not an artefact of female dominated firms. When looking at firms which are not necessarily dominated by women, but are a mixed gender firm, with at least one man and at least one woman, a similar picture emerges. As shown in Figure 2.27, the result is not a complete convergence in the change in reporting behaviour, but a near convergence, with men only modestly more likely than women to have previously been reporting a salary at exactly the minimum wage and then jumping up post-reform.

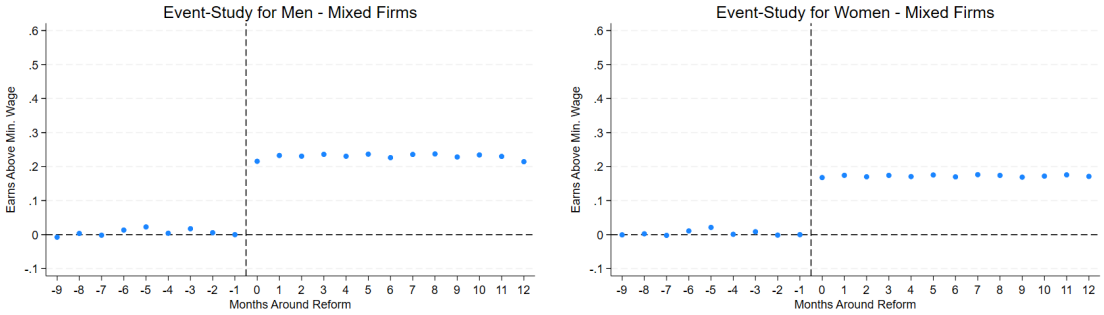


FIGURE 2.27. Event Study - Mixed Gender Firms

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 1,828,239 men in mixed firms and 1,735,504 women in mixed firms in each of the 22 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

On the other hand, we can observe a substantial gender gap when looking at firms which are completely segregated, in which there are only men or only women. The results are similar to when comparing one man and one woman-bands, where almost half of men were reporting a salary at exactly the minimum wage pre-reform and jump to a higher report post-reform, compared with a figure closer to 20 percent for women.

I present descriptive statistics on mixed gender and segregated firms in the Appendix to Chapter 2. Workers in mixed firms are only marginally older and have marginally more dependents. The key observable difference between these kinds of firms is the size, as measured by number of employees, with mixed firms being substantially larger. Although all of my regressions control for firm size, it is likely that the larger mixed gender firms are much more professional than gender segregated firms and thus feature minimal levels of tax evasion. Interestingly, the female-only segregated firms are smaller than the male only segregated ones, despite having smaller post-reform jumps away from the minimum wage which are consistent with tax evasion.

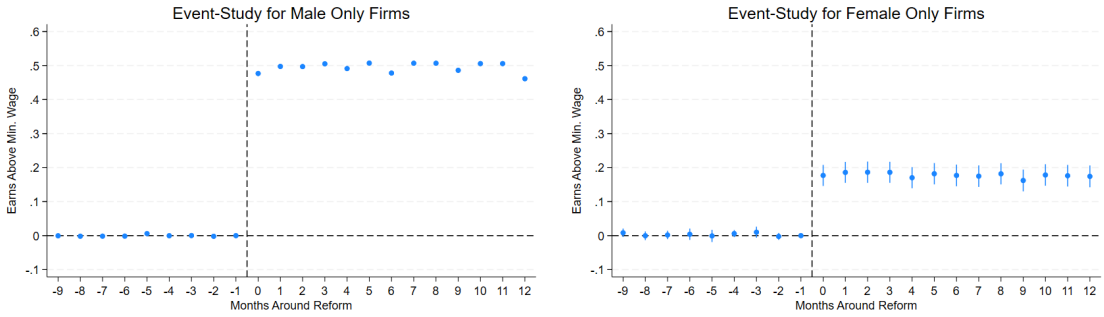


FIGURE 2.28. Event Study - Segregated Firms

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 142,925 male individuals and 150,299 female individuals in each of the 22 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

The evidence suggests that the composition of firms is a strong determinant of the propensity of tax evasion, but only for men. In one-man bands and in firms comprising only of men, men are substantially more likely to be changing their income reporting behaviour in a way that is consistent with tax evasion compared to those in the full sample. At the same time, when they are in female dominated firms, or even firms which are simply mixed-gender to any degree, men are less likely to be tax evading than they are in the full sample. On the other hand, the proportion of women under-reporting their income remains largely fixed, regardless of the gender composition of the firm. This suggests that it may not be the ability to evade tax which is binding for women, but rather that they have stronger preferences for honesty compared to men. Meanwhile, when men find it easier to evade tax, such as when they are in a single-person firm or a firm without any women, they are more likely to do it.

2.6 Conclusion

Equipped with a series of policy reforms and administrative tax data, I have provided evidence which points strongly to the existence of a large number of Romanian construction workers evading income taxes through the use of 'envelope payments', whereby the minimum wage is declared on paper and the remainder of the true salary is paid under the table. As well

as adding to the empirical literature on tax evasion, I contribute by highlighting the gender difference in under-reporting. It appears that men are substantially more likely to be under-reporting their income compared to women.

As the policy reforms pushed reported incomes closer to the true total incomes, the reported gender wage gap flipped from being 8% in favour of women, to women being out-earned by 7%[34]. Under the interpretation that the reporting changes post-reform reflect primarily previously hidden income now becoming revealed, this implies that tax evasion has significant implications for the empirical measurement of the wage gap.[35] Disentangling the reasons why men and women have different reporting behaviour is challenging. Risk and/or social preferences are likely to play key roles in the explanation. When men operate in solo firms, or firms in which all employees are male, they are substantially more likely to adjust their behaviour in a way that is indicative of tax evasion compared to when they are employed in mixed-gender firms. The reporting behaviour of women, however, is largely unchanged under different firm compositions. Men may thus perceive it to be more risky to engage in income under-reporting in the presence of female colleagues. This could also be explained by selection, where the men who have preferences against tax-evasion are more likely to work in the more honest firms, with a higher composition of women.

For policymakers, this paper highlights the prevalence of income tax evasion, even in a setting with third-party reporting. In addition to the objective of increasing tax revenues, improving tax compliance can fulfil an important secondary objective of boosting gender equality. When income tax evasion comes in the form of falsely declaring the minimum wage and collecting the remainder of the true salary as a cash in hand 'envelope wage', increasing the minimum wage can serve as a simple and pragmatic policy to boost compliance so long as it is done carefully.

[34] And from 6% in favour of women to zero in balanced sample

[35] Alstadsæter, Johannesen, and Zucman (2019) similarly find that accounting for unreported assets increases top wealth shares and thus tax evasion has an impact of the empirical measurement of wealth inequality.

Appendix 2

B.1 Other Graphics

B.1.1 Local-Industry Labour Market Flow Rates

As an alternative approach to using HHI scores to define the competitiveness of a labour market I use the rates of labour market flows between firms within a local-industry, where the local industry is defined as the four-digit NACE code within one of Romania's 42 counties. To generate the mobility scores, or flow rates, I use the following procedure. For each individual I generate an outflow indicator if in the next month they are observed in the same county and the same four-digit NACE code, but they are matched with a different firm. For each individual I have another indicator for inflows, which takes on the value of one if in the last month they were in the same county and the same four-digit NACE code, but they were in matched with a different firm. Then for each local-industry (four-digit NACE code and county) I count the aggregate number of employees who ever worked there. This is just adding up the number of employees in each month across all months, and is not using only unique employees. Then for each local-industry I create the labour flow rate by adding up the inflows and outflows and dividing by the aggregate employment. Thus for each local-industry I have a measure of how much employees transition between firms in the same local-industry, relative to the size of the industry.

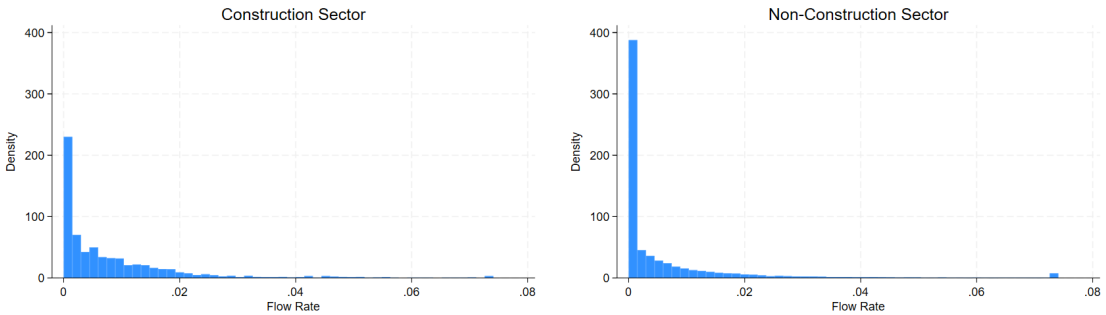


FIGURE B1. Labour Market Flow Rates: By Sector

Notes: The median flow rate in the non-construction sector is 0 and it is 0.004215 in the construction sector, with means of 0.0067 and 0.0084 respectively.

As can be seen in Figure B1, local-industries in the construction sector have somewhat higher rates of labour market flows.

B.1.2 One-Person Firm Descriptive Graphs

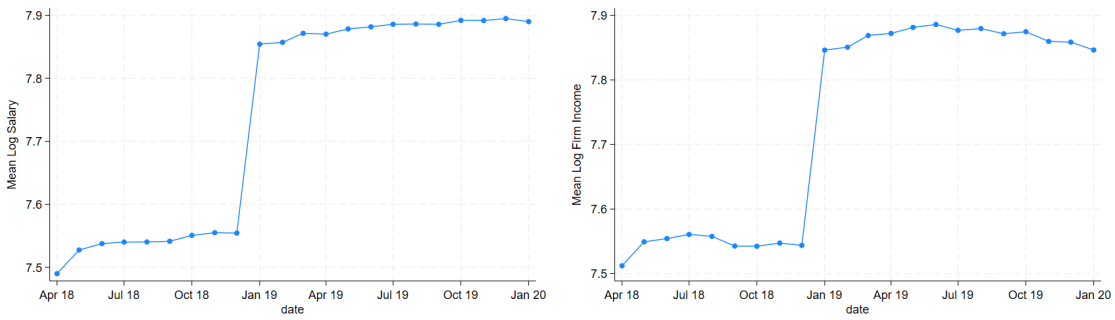


FIGURE B2. One-Person Firms: Mean Salaries and Firm Revenues

Notes: The sample only includes individuals in one-person firms who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The graphs show mean log reported income and firm revenue of these one-person firms.

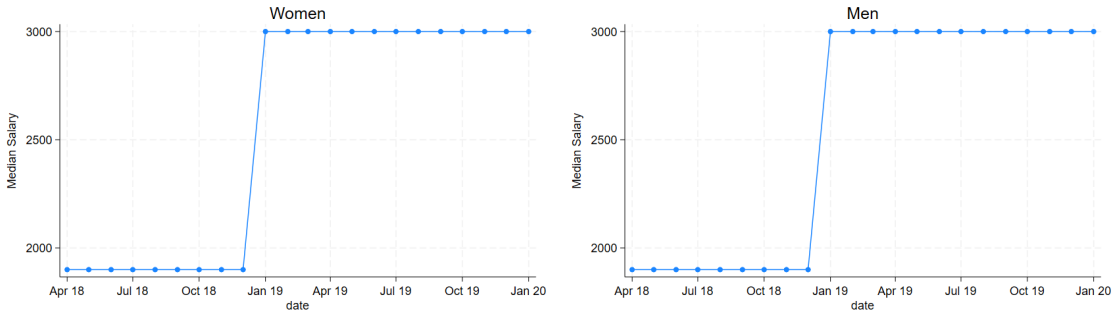


FIGURE B3. One-Person Firms: Median Salaries by Gender

Notes: The sample only includes individuals in one-person firms who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The graphs show median reported income of these one-person firms.

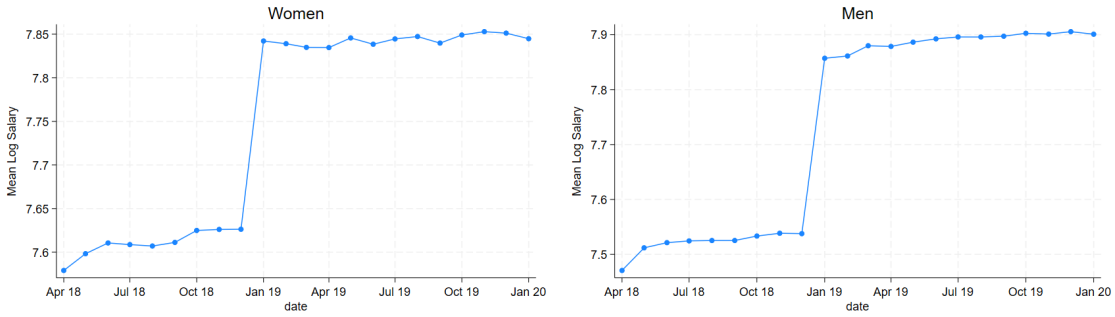


FIGURE B4. One-Person Firms: Mean Salaries by Gender

Notes: The sample only includes individuals in one-person firms who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The graphs show mean log reported income of these one-person firms.

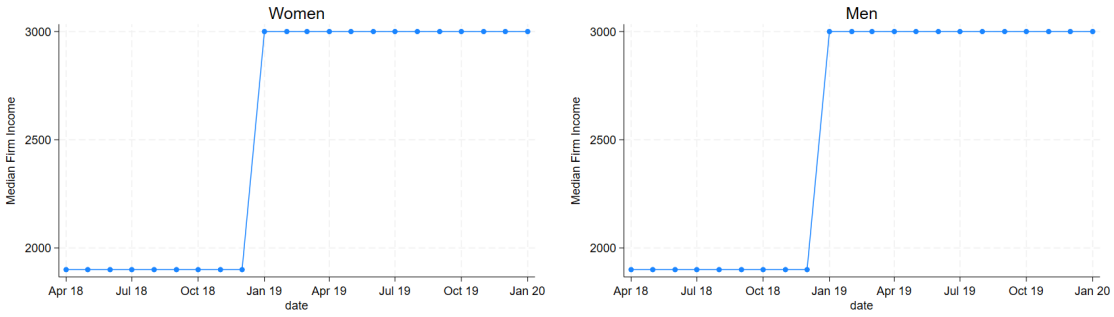


FIGURE B5. One-Person Firms: Median Firm Income by Gender

Notes: The sample only includes individuals in one-person firms who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The graphs show median reported firm income of these one-person firms.

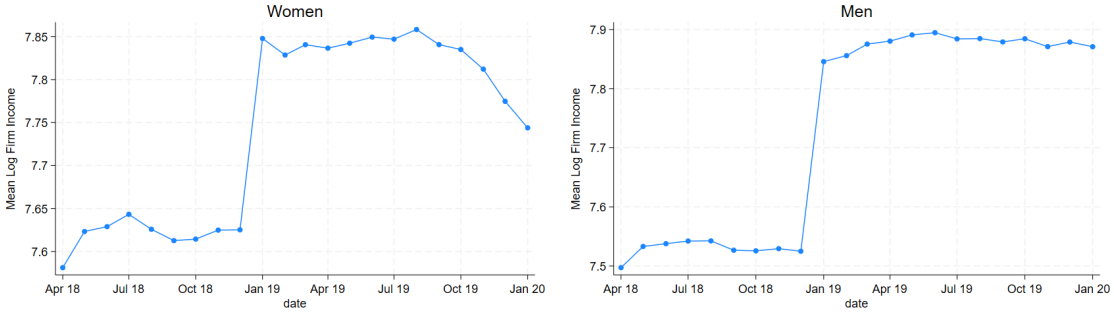


FIGURE B6. One-Person Firms: Mean Firm Log Income by Gender

Notes: The sample only includes individuals in one-person firms who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The graphs show mean reported firm log income of these one-person firms.

B.1.3 One-Person Firm Take-up

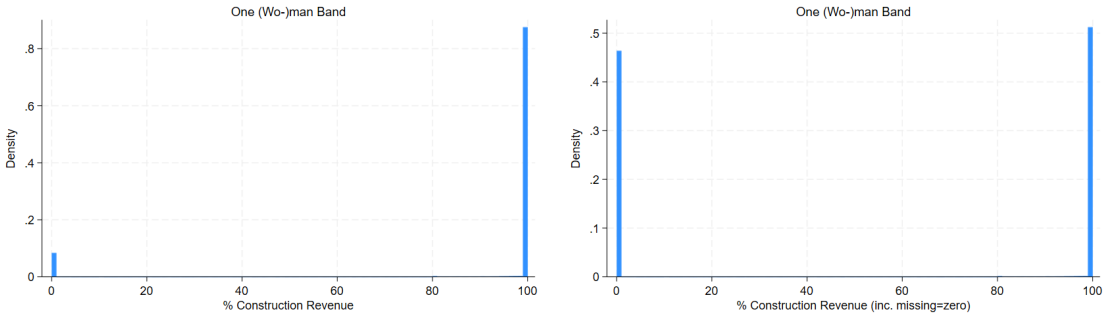


FIGURE B7. One-Person Firms: Take-Up of Policy

Notes: In both figures is the sample of one-person firms. The figure on the left shows the distribution of the reported percentage of revenue which is derived from construction activities. On the right hand side is the corresponding graph with zeros filled in for missing reports for construction sector firms.

B.1.4 Percent of Construction Employees Reporting Exactly the Minimum Wage

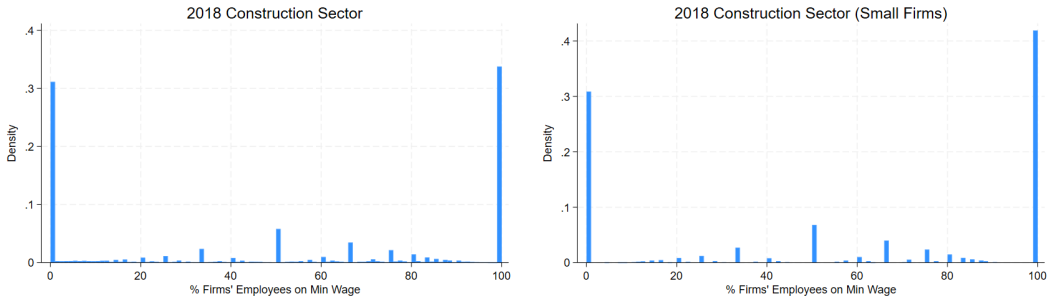


FIGURE B8. Proportion of Construction Employees Reporting Exactly the Minimum Wage
Notes: The graphs show the distribution across construction sector firms of the proportion of employees within a firm that reports a salary at exactly the minimum wage. The year used is 2018, which is before the reform, with the month arbitrarily chosen as April. The graph on the right hand side makes a restriction to small firms, with the threshold for a small firm being arbitrarily chosen as those with ten or fewer employees.

B.2 Sample Details

B.2.1 Treated NACE Code List: Female Share and Share of Construction Employment

Industry NACE Code	Female Share	Share of Construction Employment
2331	0.333	0.002
4110	0.324	0.017
2312	0.29	0.009
1623	0.238	0.02
2223	0.226	0.029
2362	0.224	0.002
4212	0.211	0.005
2332	0.208	0.004
2370	0.204	0.006
2512	0.202	0.011
2369	0.2	0.003
4213	0.195	0.003
4332	0.173	0.017
2361	0.171	0.02
2364	0.168	0.002
4291	0.154	0.01
4312	0.15	0.011
811	0.148	0.007
2363	0.146	0.01
4222	0.145	0.013
4313	0.145	0.002
4322	0.144	0.079
4321	0.143	0.076
4311	0.143	0.004
4334	0.135	0.008
4329	0.133	0.012
4221	0.132	0.013
4333	0.13	0.013
812	0.129	0.011
4211	0.128	0.086
4299	0.125	0.013
4339	0.124	0.005
4399	0.116	0.036
4391	0.112	0.007
4120	0.112	0.43
4331	0.103	0.003

TABLE B1. NACE Sector Codes Treated With Income Tax Exemption

Notes: I present the full list of industry codes which are eligible for the income tax exemption, ordered by the decreasing female share of employment within the industry. I also present for each sector code, the share of the construction employment it makes up.

B.2.2 Wage Distributions by Sector and by Year

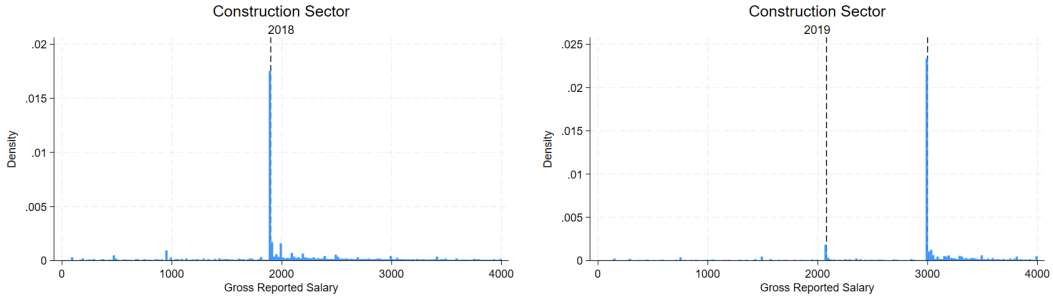


FIGURE B9. Wage Distributions of the Construction Sector

Notes: Histogram displaying, with a bin-width of 20 RON, the distribution of reported salaries in the construction sector in 2018 and 2019. For visualisation purposes only reported salaries up to 4000 RON are displayed. The dashed vertical line in 2018 is at the minimum wage of 1900, while in 2019 there is a dashed vertical line at the new minimum wage of 2080 and at 3000 RON.

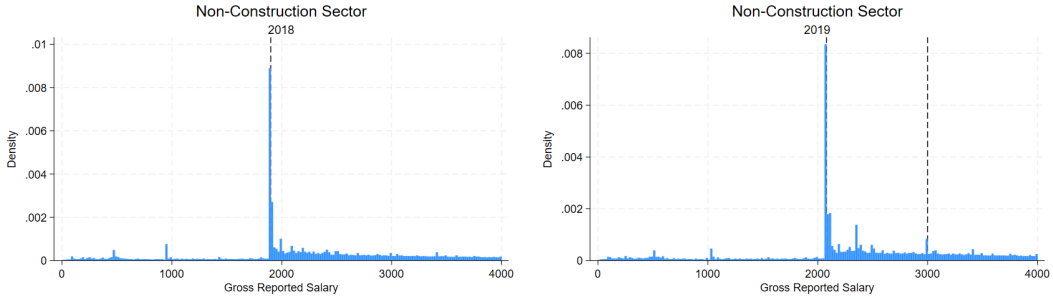


FIGURE B10. Wage Distributions of the Non-Construction Sector

Notes: Histogram displaying, with a bin-width of 20 RON, the distribution of reported salaries in the non-construction sector in 2018 and 2019. For visualisation purposes only reported salaries up to 4000 RON are displayed. The dashed vertical line in 2018 is at the minimum wage of 1900, while in 2019 there is a dashed vertical line at the new minimum wage of 2080 and at 3000 RON.

B.2.3 Pre and Post-Reform Descriptive Statistics: by Gender and Sector

	(1) Construction Sector			(2) Non-Construction Sector		
	Median	Mean	S.D.	Median	Mean	S.D.
Age	42.00	42.33	11.16	42.00	41.90	11.92
Hours Worked	160.00	144.10	49.58	160.00	147.00	46.71
Number of Dependents	0.00	0.28	0.65	0.00	0.34	1.10
Firm Size	28.00	119.69	247.18	96.00	1253.84	3304.47
Female Share of Firm	0.24	0.32	0.25	0.70	0.67	0.23
Full-Time Contract	1.00	0.91	0.29	1.00	0.90	0.30
Bucharest	0.00	0.17	0.38	0.00	0.15	0.36
Gross Reported Salary	2020.00	2977.09	5441.47	2548.00	3618.80	7831.29
Observations	740746			30785352		

Panel (a): Pre-Reform Women

	(1) Construction Sector			(2) Non-Construction Sector		
	Median	Mean	S.D.	Median	Mean	S.D.
Age	43.00	42.92	11.69	43.00	42.62	12.74
Hours Worked	160.00	143.26	48.68	160.00	150.76	41.61
Number of Dependents	0.00	0.30	0.73	0.00	0.33	0.74
Firm Size	31.00	115.68	240.24	102.00	1424.80	3695.99
Female Share of Firm	0.09	0.12	0.12	0.30	0.34	0.24
Full-Time Contract	1.00	0.96	0.20	1.00	0.92	0.27
Bucharest	0.00	0.09	0.29	0.00	0.13	0.34
Gross Reported Salary	1900.00	2614.99	4847.81	2646.00	4020.87	10274.67
Observations	4392411			30653062		

Panel (b): Pre-Reform Men

TABLE B2. Descriptive Statistics: Selection into Construction Sector by Gender: Pre-Reform
Notes: Hours worked are measured monthly. Bucharest is a dummy variable for working in the capital city of Bucharest. The data covers from Jan 2019 up to and including Dec 2019. The panel has not been balanced. The number of unique men and women is 3,874,789 and 3,485,161 respectively.

	(1) Construction Sector			(2) Non-Construction Sector		
	Median	Mean	S.D.	Median	Mean	S.D.
Age	42.00	42.55	11.25	42.00	42.37	12.03
Hours Worked	160.00	146.48	48.67	160.00	148.53	46.08
Number of Dependents	0.00	0.27	0.68	0.00	0.34	0.92
Firm Size	26.00	116.98	249.33	98.00	1251.32	3326.84
Female Share of Firm	0.22	0.30	0.24	0.70	0.66	0.23
Full-Time Contract	1.00	0.92	0.27	1.00	0.90	0.29
Bucharest	0.00	0.18	0.38	0.00	0.15	0.36
Gross Reported Salary	3000.00	3697.61	5526.12	2977.00	4105.54	6716.38
Observations	835988			33478321		

Panel (a): Post-Reform Women

	(1) Construction Sector			(2) Non-Construction Sector		
	Median	Mean	S.D.	Median	Mean	S.D.
Age	43.00	42.85	11.79	43.00	43.05	12.83
Hours Worked	160.00	147.53	46.48	160.00	152.67	40.22
Number of Dependents	0.00	0.28	0.71	0.00	0.34	0.84
Firm Size	29.00	114.08	241.72	104.00	1409.30	3686.45
Female Share of Firm	0.09	0.11	0.11	0.30	0.34	0.24
Full-Time Contract	1.00	0.97	0.17	1.00	0.92	0.27
Bucharest	0.00	0.09	0.29	0.00	0.13	0.34
Gross Reported Salary	3000.00	3405.57	3703.17	3078.00	4549.50	7691.33
Observations	5417383			32662672		

Panel (b): Post-Reform Men

TABLE B3. Descriptive Statistics: Selection into Construction Sector by Gender: Post-Reform
Notes: Hours worked are measured monthly. Bucharest is a dummy variable for working in the capital city of Bucharest. The data covers from Jan 2020 up to and including Jan 2021. The panel has not been balanced. The number of unique men and women is 3,921,9420 and 3,531,652 respectively.

B.2.4 Sample Details by Firm Composition

	(1) Men			(2) Women		
	Median	Mean	S.D.	Median	Mean	S.D.
Age	46.00	45.34	10.63	45.00	44.77	10.15
Hours Worked	168.00	160.26	27.09	160.00	158.22	30.82
Number of Dependents	0.00	0.42	0.82	0.00	0.33	0.72
Firm Size	55.00	162.76	286.41	43.00	155.91	292.62
Female Share of Firm	0.12	0.14	0.10	0.20	0.26	0.18
Full-Time Contract	1.00	0.98	0.15	1.00	0.94	0.23
Bucharest	0.00	0.11	0.31	0.00	0.19	0.39
Observations	3349724			658570		

Panel (a): Mixed-Gender Firms

	(1) Men			(2) Women		
	Median	Mean	S.D.	Median	Mean	S.D.
Age	43.00	43.32	10.20	43.00	43.78	9.70
Hours Worked	160.00	158.37	30.64	160.00	156.90	32.54
Number of Dependents	0.00	0.29	0.70	0.00	0.31	0.70
Firm Size	5.00	8.15	9.64	1.00	2.17	2.23
Female Share of Firm	0.00	0.00	0.00	1.00	1.00	0.00
Full-Time Contract	1.00	0.96	0.20	1.00	0.92	0.27
Bucharest	0.00	0.11	0.32	0.00	0.22	0.41
Observations	475746			20636		

Panel (b): Gender-Segregated Firms

TABLE B4. Descriptive Statistics: By Firm Composition

Notes: Hours worked are measured monthly. Bucharest is a dummy variable for working in the capital city of Bucharest. The data covers from April 2018 up to and including Dec 2019. The panel used is the balanced sample.

B.2.5 Firm Size Distributions

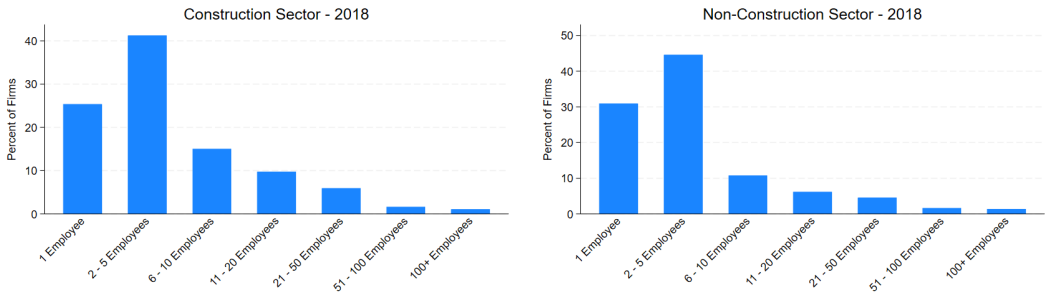


FIGURE B11. Firm Size Distributions by Sector

Notes: Bar graph presenting the distribution of firm sizes in the construction and non-construction industries.

B.3 Other results

B.3.1 Firm Industrial Code Switching

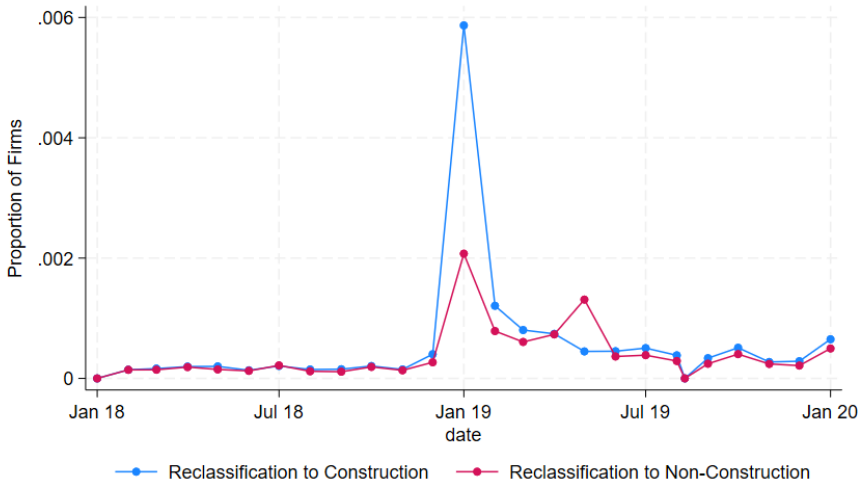


FIGURE B12. Non-Construction to Constriction Classification Switches and Vice-Versa
Notes: This graph plots the proportion of firms in each month which switch from a non-construction to a construction NACE sector code and the proportion of firms which which from a construction to a non-construction NACE code.

B.3.2 Joint Collusion: Jumps to Reporting 3000 RON

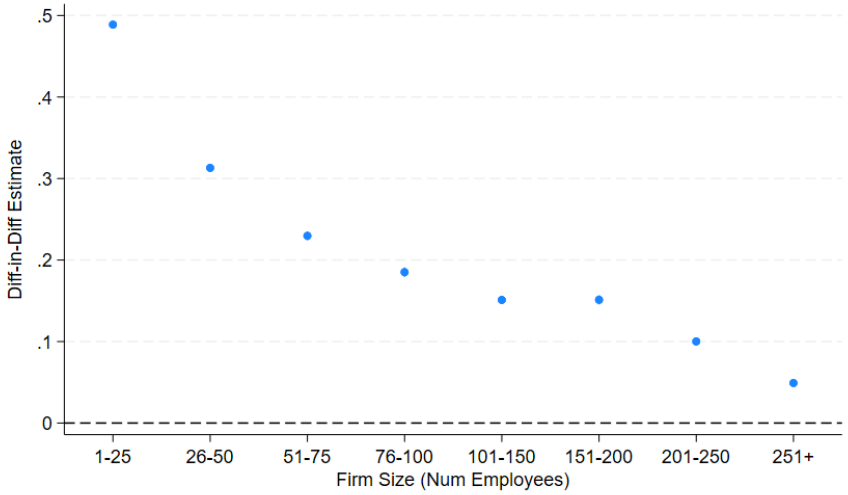


FIGURE B13. 2x2 DiD Estimates of Reporting 3000 RON by Firm Size
Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 3,670,969 individuals in each of the 24 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

B.3.3 Joint Collusion: Discontinuity in One-Person Firms

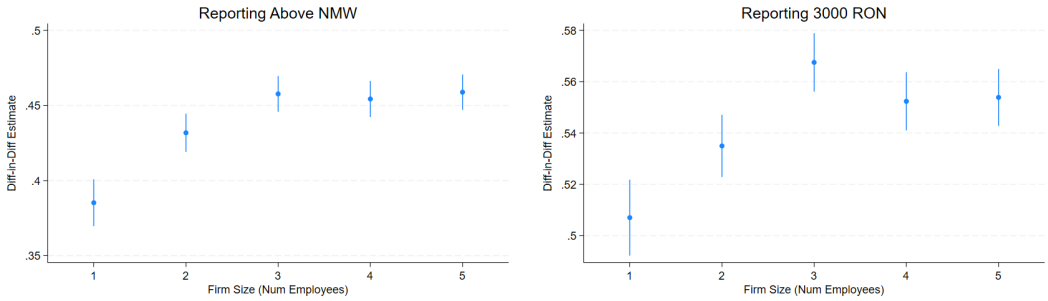


FIGURE B14. Small Firms:2x2 DiD Estimates

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 573,636 individuals in each of the 24 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

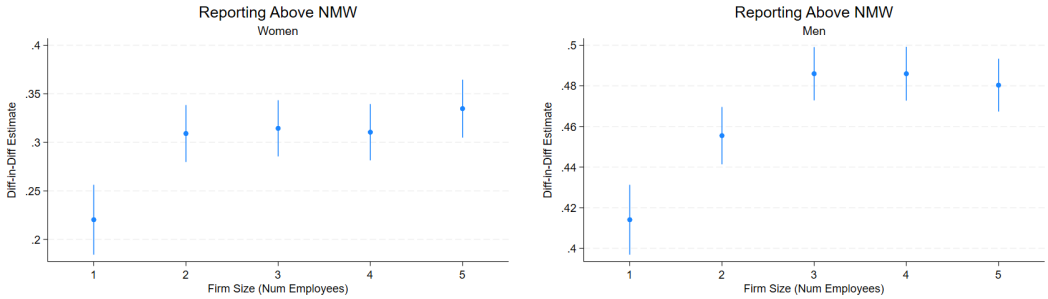


FIGURE B15. Small Firms: 2x2 DiD Estimates of Reporting Above Regular NMW by Gender

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 287,554 women and 286,082 men in each of the 24 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

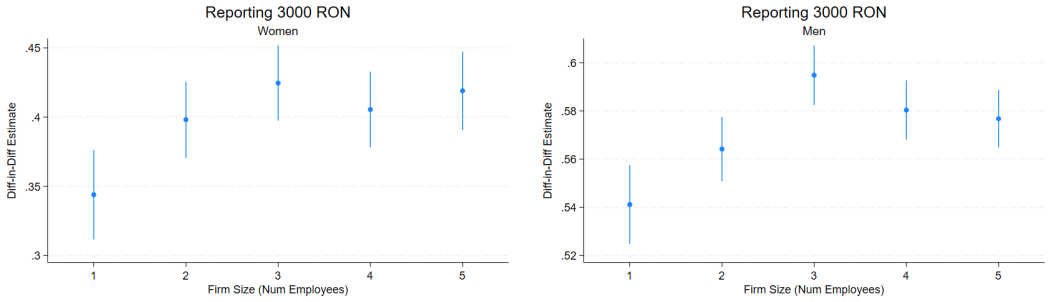


FIGURE B16. Small Firms: 2x2 DiD Estimates of Reporting 3000 RON by Gender

Notes: The sample only includes individuals who remained employed throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 287,554 women and 286,082 men in each of the 24 months. The standard errors are clustered at the individual level and the controls are the number of employees in the firm, the contracted hours and the hours actually worked.

B.3.4 Joint Collusion: Jumps in Firms' Reporting

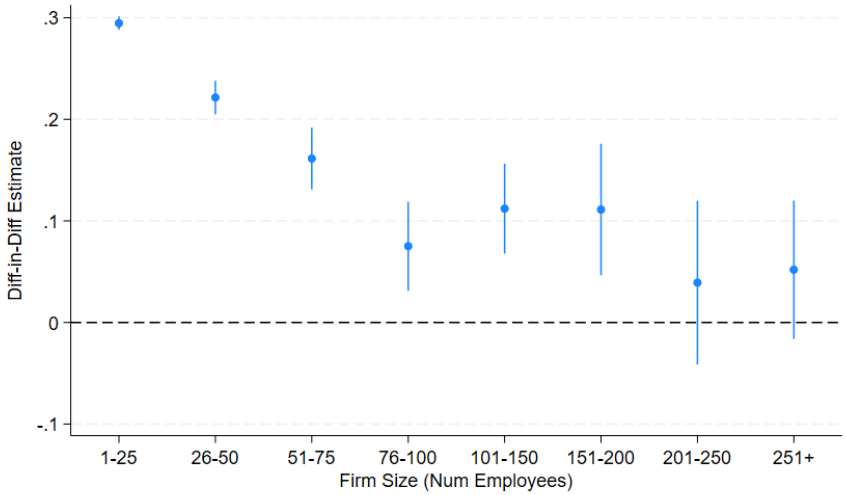


FIGURE B17. 2x2 DiD Estimates of Log Firm Income by Firm Size

Notes: The sample only includes firms who remained active throughout the entire period and excludes those who moved between the treated construction sector and other sectors. The sample size is 325,286 firms in each of the 24 months. The standard errors are clustered at the firm level.

B.3.5 Productivity Growth

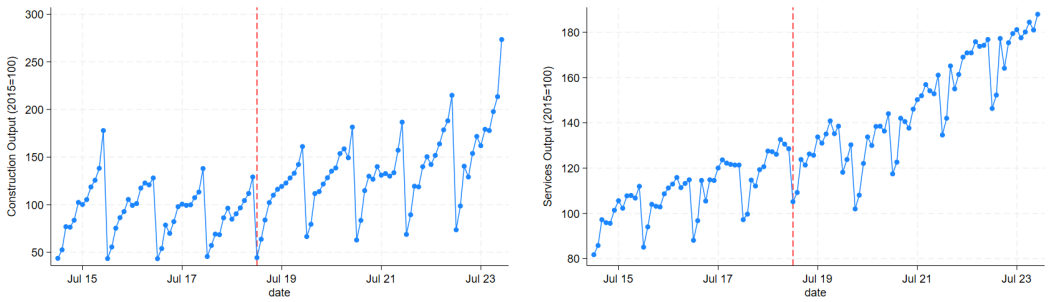


FIGURE B18. Productivity Growth in Construction and Services

Notes: The data are from Eurostat with online data code sts_copr_m for the construction data and sts_sepr_m for the services data. Output for construction and services are both normalised relative to the 2015 average level, which serves as the reference year. The data are not adjusted for seasonality or calendar year and measure production volume in the construction classification of output and services of the business economy (except financial and insurance activities).

B.3.6 Manual Workers: Alternative Definition

Here are the results for manual workers using an alternative definition of a manual worker. Rather than using the individual-level variance in hours worked, I construct, for each individual, the ratio of hours worked in the months of December, January and February to the hours worked during the rest of the year. A low ratio means that hours dip in the colder months relative to the rest of the year. Equipped with a distribution of values of this ratio, I define my second indicator of a manual worker as a worker with a ratio under the 10^{th} percentile of this distribution.

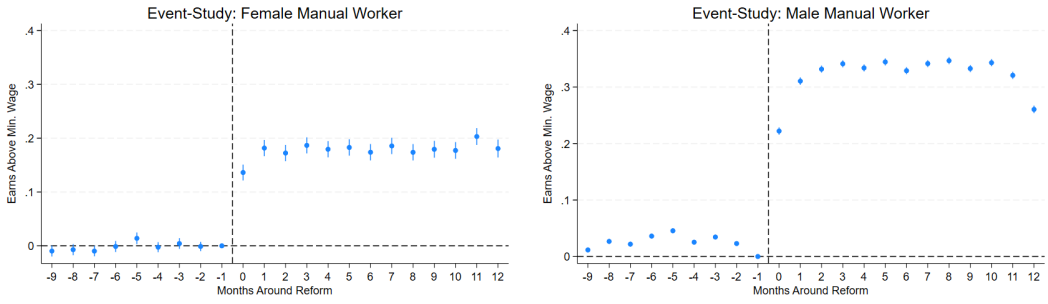


FIGURE B19. Event Study-Effects in Competitive Industries

Notes: The specification in the regressions includes controls for actual hours worked, the number of employees in the firm and type of contract (full or part-time). Standard errors are clustered at the individual level and 95% confidence intervals are plotted around the point estimates.

B.4 Standard Errors

I chose to cluster my standard errors at the individual level in my regression specifications. To verify that the significance of my results is not being driven by my choice of standard

errors, I run my preferred DiD event-study specification (across both genders) using a number of different choices for the calculation of standard errors. I present the standard errors for my difference-in-difference regressions on the outcome of a dummy for the individual earning above the minimum wage under: 1) individual-level clustering, 2) firm-level clustering, 3) 4-digit NACE industry code clustering, 4) clustering at the level of treatment (construction vs not construction industry) and 5) heteroskedasticity-consistent 'robust' standard errors. These are presented in Table B5 and the standard errors are similarly sized throughout the different specifications. Importantly, they remain remarkable small in under each standard error calculation, which can be attributed to the size of the data. Each point estimate is highly significant at conventional levels under and choice of standard error.

	Indiv	Firm	Industry	Constr.	HC Robust
Month -9	(0.001)	(0.002)	(0.002)	(0.000)	(0.001)
Month -8	(0.001)	(0.002)	(0.003)	(0.000)	(0.001)
Month -7	(0.001)	(0.002)	(0.003)	(0.000)	(0.001)
Month -6	(0.001)	(0.002)	(0.004)	(0.000)	(0.001)
Month -5	(0.001)	(0.002)	(0.005)	(0.000)	(0.001)
Month -4	(0.000)	(0.001)	(0.002)	(0.000)	(0.001)
Month -3	(0.001)	(0.002)	(0.004)	(0.000)	(0.001)
Month -2	(0.000)	(0.001)	(0.002)	(0.000)	(0.001)
Month 0	(0.001)	(0.006)	(0.038)	(0.000)	(0.001)
Month 1	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 2	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 3	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 4	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 5	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 6	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 7	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 8	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 9	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 10	(0.001)	(0.006)	(0.041)	(0.000)	(0.001)
Month 11	(0.001)	(0.006)	(0.039)	(0.000)	(0.001)
Month 12	(0.001)	(0.006)	(0.037)	(0.000)	(0.001)
N	82245572	82245572	82245572	82245572	82245572

TABLE B5. Different choices for the calculation of standard errors

Notes: I present the point-estimates and standard errors under 1) individual-level clustering, 2) firm-level clustering, 3) 4-digit NACE industry code clustering, 4) clustering at the level of treatment (construction vs not construction industry) and 5) heteroskedasticity-consistent 'robust' standard errors.

Chapter 3

Group Differences in Labour Markets with Observational Learning and Information Imperfections

Abstract: I use a rational model of observational learning, in which previous success or failure in the labour market acts as a signal of quality, to understand how group differences in employment can vary over time in a setting with discrimination. Under different information structures and (mis-)specifications of beliefs, group differences may be reduced or exacerbated over time, with the direction depending on the structure of information. A benefit-of-the-doubt effect occurs when some applicants are treated with extra lenience during evaluation, since they might have been discriminated against in the past. On the other hand, naive employers are subject to a signal-distortion effect, falsely interpreting unemployment caused by discrimination as a signal of low ability, exacerbating group differences.

This chapter draws from work completed during my MPhil in Economics, undertaken at the University of Oxford. I am grateful for the support of my MPhil supervisor, Willemien Kets, without whom it would not be possible to complete this DPhil chapter. I am also grateful to a number of individuals who have shared valuable comments and suggestions with me, at various stages of the project, including two anonymous MPhil thesis reviewers, Johannes Abeler, Mihai Codreanu, Olivia Flett, Christian Meyer, Tzvetan Moev, Barbara Petrongolo and Philip Trammell. As I worked on this project during my DPhil in Economics, Margaret Meyer was exceptionally generous with her time and supported me with the new theoretical analysis.

3.1 Introduction

There is a rich theoretical and empirical literature on discrimination in hiring[1], in which the dynamic effects of discriminatory behaviour are a relatively recent innovation.[2] I add to this literature by showing that the dynamic effects of discrimination in hiring depend on the information environment in the labour market, with gaps in employment between groups potentially converging or diverging depending on the scenario. To do this, I bridge a rational model of observational learning (often referred to as social or sequential learning) with the literature that aims to explain group differences in labour markets. The outcome variable of interest in my model is employment, and while other aspects of the job contract, particularly wages, are important outcomes, this paper focuses on how candidates are evaluated given their employment history. Figure 3.1 plots the working age employment rates for the black and white labour force subpopulations of the United States in 2019 and indicates that when first entering the labour force, there is a significant difference in how the two groups fare, but that this gap shrinks quickly over time. This pattern highlights the need for explanations of group differences in labour markets to focus not only on static differentials in pay or employment rates, but also on the dynamics, which has received less attention in the previous theoretical literature.

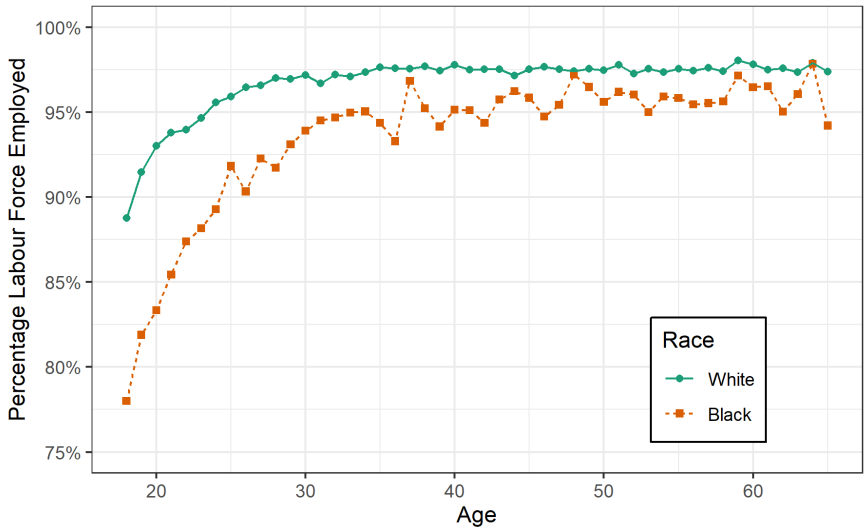


FIGURE 3.1. Employment rates across the life-cycle of working age labour force participants.
Notes: Current Population Survey (CPS) data for black and white labour force participants between the ages of 18 and 65 in the year 2019. The sub-sample sizes for the black and white groups are, respectively, 65057 and 528499.

I develop a model of job applicants with unknown ability, but an informative signal and an observable past performance in the labour market. This important feature captures the penalty to having been unemployed in the past. The model also features taste-discrimination, meaning that some employers have a dis-taste for a certain group of workers. I study the

[1] Bertrand and Duflo (2017)
[2] Bohren, Imas, and Rosenberg (2019)

employment dynamics under varying structures of employers' information about the discriminatory attitudes of the others, including a version with mis-specified beliefs.

The primary mechanism of interest in the model is how labour market outcomes in the past influence an employer's hiring choice in the present. A job applicant with a strong labour market history should be more likely to gain future employment than an applicant with a weak labour market history, since these outcomes signal something about a candidate's underlying ability. Taste-discrimination, however, results in job applicants from certain groups being unfairly rejected from jobs for which they are well qualified.

The main findings of the model are: (i) in a benchmark of full information about past discrimination, members of the discriminated group are less likely to find employment than those in the majority group, and they become no better or worse off over time; (ii) under mis-specified 'naive' beliefs, employers fail to account for the fact that members of the minority group may have been discriminated against in the past and treat unemployment as a negative signal, hence increasing the group differences over time; (iii) in a model of partial information, where employers are aware of the existence of taste-discriminators but cannot know for certain whether someone has been previously discriminated against, the difference in employment between groups is reduced over time (members of the discriminated-against group may even be more likely to find employment than those of the majority group), since the negative signal associated with unemployment is rationally discounted (because it might have been caused by discrimination). A panel data analysis of the US labour market suggests that the partial information model is the most in line with the stylised facts. The empirical analysis, however, should be thought of as primarily descriptive.

In the remainder of Section 3.1 I discuss related literature and in Section 3.2 I describe the preliminaries of the model. I study the model under full information and mis-specified beliefs in Section 3.3 and Section 3.4 respectively, and in Section 3.5 under partial information. In Section 3.6 I use a panel dataset to compare the stylised facts of the data against the different model results and in Section 3.7 I relax a number of assumptions made in the simple model. The economic logic developed in the simple model carries over to a more complex model, at the expense of some analytical tractability and the 'sharpness' of some predictions. Finally, in Section 3.8 I conclude.

3.1.1 Related Literature

An extensive field of literature has been devoted to explaining and quantifying group differences in a variety of economic markets, such as the markets for labour, housing and credit. Most commonly, the models employed have focused on either taste-based theories of discrimination (introduced by Becker, 1971), which formalise prejudice against a certain group of people, or statistical theories[3], which generally describe how group differences can arise without prejudice. The most influential work on the topic of discrimination has tended to study these theories in isolation, rather than combining the two.

[3]Fang and Moro (2011) provide a comprehensive overview of this literature.

Seminal models of statistical discrimination tend to rely on the intuition that employers[4] use group level characteristics to predict individual level attributes. For example, in the simple model proposed by Phelps (1972), workers in different groups have exogenous differences in average ability, resulting in discrimination between candidates who are equally qualified, as employers use group membership as a predictor of quality. As noted by Kim and Loury (2018), the theoretical literature is largely static and thus leaves several open questions about the evolution of group differences over time.

Alongside the theoretical explanations, there exists an influential field of empirical work attempting to quantify levels of discrimination. In one of the most well known articles in this sub-discipline, Bertrand, Mullainathan, and Shafir (2004) report evidence that a CV with a typically white name receives 50 percent more callbacks than an equivalent CV featuring a typically African-American name. Despite the size of the literature, there has been a gap in empirical work attempting to seriously distinguish between the different mechanisms which result in discrimination (Guryan and Charles, 2013).

Moreover, a specific feature of both the theoretical and empirical literature which has largely been missing is interdependence of hiring decisions between employers, particularly in the dynamic sense. In the long-term unemployment literature, however, this interdependence over time is a crucial aspect of the labour market dynamics.

The observation that those who have suffered from spells of joblessness are permanently disadvantaged in the labour market (Arulampalam, 2001) can be explained by a rational herding mechanism[5] (Oberholzer-Gee, 2008). The intuition is that when a potential employer meets someone who is long term unemployed, the employer will assume that the applicant has previously applied elsewhere and not been successful. Being unsuccessful with previous applications is a signal of low quality, since low quality applicants are more likely to be rejected and high quality applicants more likely to be hired. Thus, a poor past performance in the labour market, due to perhaps nothing more than misfortune, contributes to a further lack of success in job search. There is also empirical evidence (Barach and Horton, 2021) that a job applicant's previous compensation is used as a signal of unobserved ability, due to its reflection of marginal productivity.

Seminal rational herding models were developed by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), in which past decisions convey some information about the unknown quality[6] of an arbitrary good or service. A common example is an individual choosing which restaurant to eat at by first observing where others are eating, and using this information to help decide which restaurant to visit. Although an individual updating his or her beliefs conditional on the decisions of others is perfectly rational, it can quickly lead to information cascades, in which individuals discard their private information and follow the decisions of the preceding agents. Smith and Sørensen (2011) summarise this as a convergence in *beliefs* eventually resulting in a convergence in *actions*. Herd behaviour often

[4] In the case of housing markets, these are real estate agents, and in credit markets these are bankers.

[5] Another common explanation is that human capital depreciates over the unemployment spell (Acemoglu, 1995).

[6] Usually it is assumed that the object in question is objectively good or bad, such as a financial asset. Smith and Sørensen (2000) study a herding model where preferences might be heterogeneous.

leads to inefficiency, since the resulting actions are less reflective of private information and thus fail to provide meaningful information to future decision makers. This is called a 'herd externality' by Banerjee (1992).[7]

Although 'observational learning' and 'herding' are often used in reference to the same class of models, it is worth clarifying the terminology. Observational learning is the mechanism through which previous decisions are taken into account as a form of information, whereas herd behaviour is an outcome of many of these models, in which individuals forego their private information and follow the choices of previous decision makers. In my labour market discrimination setting, observational learning occurs when an employer uses an applicant's previous employment outcomes to learn something about the applicant's ability, while herd behaviour would occur if past discrimination results in permanent disadvantages.

In the related employer learning literature, Altonji and Pierret (2001) provide evidence that since employers learn about intrinsic ability over time, they progressively rely less on using easily observable variables such as ethnicity as a proxy. This provides evidence consistent with statistical theories of discrimination. I contribute to the employer learning literature by considering the more complex dynamic, whereby employers are not necessarily learning correctly about intrinsic ability over time, but learning through noisy signals which depend on the behaviour of other, potentially discriminatory, employers.

Although literature which models the dynamic aspects of discrimination is scarce compared to that which employs static models, there have been major contributions in recent years. Bohren, Imas, and Rosenberg (2019) show theoretically and experimentally that the dynamics of gender differences can reveal whether differential outcomes are driven by misspecified beliefs about the quality of men and women. If women are incorrectly thought to be of lower ability, but receive positive ability evaluations, those positive evaluations are especially impressive, resulting in a reversal of discrimination where women who were initially discriminated against are now favoured, since they have proved themselves as especially talented by overcoming a higher hurdle[8]. Much of the intuition in my model follows that which is in Bohren, Imas, and Rosenberg (2019), where positive or negative outcomes may be relatively more or less impressive, depending on the kind of hurdle which the individual initially faced. My contribution here is highlighting the role that the structure of information plays. In order for a positive evaluation to be especially impressive given an individual's circumstance, it must be known what their circumstance is, or what it might be. If an individual's circumstances are not taken into account, their history of evaluations may be interpreted very differently.

Extending the literature on discrimination dynamics Bohren, Hull, and Imas (2022) show how systemic discrimination can take place, whereby initial direct discrimination can become persistent, for example, by generating skill disparities. The persistence or cessation of initial

[7] Banerjee (1992) also shows that by constraining some people to only follow their private information, society's welfare can be ex-ante improved. To account for these externalities Smith, Norman Sørensen, and Tian (2021) solves herding models from the point of view of an infinitely-lived social planner, who takes into account that the actions of an agent today impact the choice of the decision maker tomorrow.

[8] A similar result is found in Fryer Jr (2007), which shows that an employer's beliefs about an employee they initially discriminated against may be 'flipped'.

discrimination can depend on the structure of information and the belief formation of other agents in the economy, which I show in my model.

My contribution is thus expanding the literature on the dynamics of discrimination by explicitly comparing the evolution of employment paths under different structures of information and showing that different information structures yield substantially different outcomes. I do this by borrowing from the theoretical literature on observational learning and applying it to my specific setting.

3.2 The Simple Model

In this section, I setup the preliminaries of a simple model. Following this, I study the model under different information structures and forms of beliefs, including those with misspecifications. I focus on the dynamic model in a two-period setting, since it is the most simple framework with which to study the mechanisms of interest and their economic intuition. The model is purposefully simple, with a number of assumptions made for analytical tractability and symmetry. I relax many of these assumptions in Section 3.7. The results of the full model, with relaxed assumptions, are less clear cut than those in the simple model, but the underlying mechanisms remain.

3.2.1 Model setup

This setup borrows from the seminal herding model of Bikhchandani, Hirshleifer, and Welch (1992). This herding model provides the most simple and tractable framework with which to examine how decisions in the past impact on choices in the present. I modify the framework and apply it to study how discrimination in the past impacts hiring decisions in the present.

Consider a single job applicant ("she") matched with a potential employer ("he") in two periods of discrete time. In each time period there is a random matching of the job applicant with a new employer.[9]

Applicant ability is unknown, either high or low with equal probability.

$$a \in \{L, H\} \quad \text{with} \quad \mathsf{Pr}(a = L) = \mathsf{Pr}(a = H) = \frac{1}{2}$$

Applicant output is a function of ability, where high ability applicants are productive and low ability applicants are completely unproductive. This is without loss of generality and can be thought of as a normalisation. With $V > 0$, the output function is:

$$F(a) = \begin{cases} V & \text{if } a = H \\ 0 & \text{if } a = L \end{cases}$$

Although ability is unknown, job applicants have a high or low signal in each period. The signals $x_t \in \{L, H\}$ are conditionally independent over time, and have the following form, with $p > \frac{1}{2}$:

[9] This is in contrast to search and matching models (Pissarides, 2000), in which a match typically forms via a matching function and dissolves usually with some exogenous probability.

$$\mathsf{Pr}(x_t = H|a = H) = \mathsf{Pr}(x_t = L|a = L) = p$$

The signal of quality can be interpreted as the applicant's performance in an interview, or performance in a standardised test implemented by employers.

Probabilities sum to one:

$$\sum_{x_t \in \{L,H\}} \mathsf{Pr}(x_t|a = H) = \sum_{x_t \in \{L,H\}} \mathsf{Pr}(x_t|a = L) = 1$$

To denote the information available to a potential employer at time t I use:

$$X_t = \begin{cases} x_{t=1} & \text{if } t = 1 \\ \{\mathcal{O}_{t=1}, x_{t=2}\} & \text{if } t = 2 \end{cases}$$

where $\mathcal{O}_{t-1} \in \{emp, \neg emp\}$ refers to the employment *outcome* observed from the previous period. As I will discuss when I solve the model, the observed outcome may be used to infer information about the first period signal.

Ex-ante identical job applicants belong to either the majority group (A) or the discriminated against group (B). I denote applicant type as $g \in \{A, B\}$. Group membership is publicly available knowledge.

There is an $\alpha \in (0, 1)$ probability that a firm operates with a taste for discrimination, imposing a penalty on applicants of type $g = B$. These employers act as if the cost of employing the applicant is $w + d$, where w is the cost of employing the worker and d is a 'psychic' cost of hiring a member of the $g = B$ group, often labeled as a coefficient of discrimination (Becker, 1971).

Assumption 3.2.1. (Fixed Employment Cost).

$$w = \tfrac{V}{2}$$

This employment cost can be thought of as an exogenously determined market wage that no firm can deviate from. While exogenously fixed wages may not be a very appealing assumption, it is commonly made in the labour market discrimination literature for simplicity (Coate and Loury, 1993; Kim and Loury, 2018; Fryer Jr, 2007; Gu and Norman, 2020). It could also be thought of as the costs of hiring and onboarding an employee. Should the employee then be hired an economic rent is generated, which is the gap between expected output and the cost of hiring (potentially including 'psychic' cost of hiring a member of the g group), and this rent will be shared between employee wages and firm profits according to a bargaining process. In the social learning literature it is also common for the cost of investment[10] (here the employment cost), as well as the value of the assets (here the worker output, V) to be fixed, hence Assumption 3.2.1 does follow the literature. The assumption is an augmented version of one made by Bikhchandani, Hirshleifer, and Welch

[10] Avery and Zemsky (1998) provide a discussion of how price adjustments can eliminate herd behaviour in social learning models.

(1992) which together with the tie-breaking assumption keeps the model symmetric and analytically tractable. In Section 3.7 I relax this assumption and allow w to be arbitrary.

While the information contained in X_t is common to all of the information structures and belief specifications I consider, knowledge of other employers' discriminatory attitudes varies, in turn affecting how an employer learns from observing past outcomes.

Since each employer is only matched with a single applicant per period, their preferences are to hire the applicant if their expected profits are positive and reject the applicant if expected profits are negative. I make the tie-breaking assumption that if the expected profit of hiring an applicant is zero, then the applicant is employed with probability $\frac{1}{2}$. Although the conventional economic assumption is that at indifference the employer would hire the applicant, this modeling choice follows Bikhchandani, Hirshleifer, and Welch (1992) and makes the model results more analytically tractable.

To formalise the strategies of employers in response to an observed information set and group membership (for now ignoring belief systems) I use $\sigma_t : \{X_t, g\} \to \{emp, \neg emp, randomise\}$. Because, however, I am primarily focused on employment outcomes for those in both groups, my results focus on ex-ante probabilities of employment (taking into account the randomisation) rather than strategies themselves.

The per-period expected profit function takes the form:

$$\mathbb{E}(\pi_t|X_t) = \mathsf{Pr}(a = H|X_t, g)V - (w + \mathbb{1}_{\{T=1\}}\mathbb{1}_{\{g=B\}} \cdot d)$$

where $\mathbb{1}_{\{T=1\}}$ is an indicator for discriminatory preferences and $\mathbb{1}_{\{g=B\}}$ is an indicator for the job applicant belonging to the discriminated against group. The posterior probability that the candidate is high ability depends on the information available to the firm.

I assume that job applicants have an outside option worth $\mathbb{O}_t$, which can be interpreted as unemployment benefits and the instantaneous utility function is such that $u_t(w) > u_t(\mathbb{O}_t)$, hence employment (with wage w) is preferred to non-employment and a job offer is never turned down.

A worker's Von Neuman-Morgenstern expected utility at time t is

$$\mathbb{E}(u_t) = \mathsf{Pr}(emp_t)(u_t(w) - u_t(\mathbb{O}_t)) + u_t(\mathbb{O}_t)$$

and with w and $\mathbb{O}_t$ fixed, it is strictly increasing in the probability of employment.

Job applicants can be thought of as essentially passive, being randomly matched with an employer in two periods of discrete time, always accepting a job offer if it is made. Employers on the other hand make consequential decisions, deciding whether or not to employ a given worker depending on the information available to them.

Assumption 3.2.2. (Strong dis-taste).

$$\mathsf{Pr}(a = H|\{x_{t=1} = H, x_{t=2} = H\})V = \tfrac{p^2}{p^2+(1-p)^2}V < w + d$$

This assumption ensures that two consecutive signals of high-ability are unable to compensate for the composite wage cost faced by taste-discriminating employers. Employers with

a taste for discrimination will *never* hire an applicant of type $g = B$, no matter how strong the ability relevant information in X_t is. The key components of the model are summarised in Table 3.1.

Players	Employers and job applicants
Timing	A single job applicant is repeatedly and randomly matched with a new employer in two periods of discrete time
Action space	*Employers*: Make job offer, reject the candidate, or randomise *Job applicants*: Accept or reject job offer
Preferences	*Employers*: Prefer higher expected profits *Job applicants*: Prefer employment to non-employment (a job offer is thus equivalent to being hired)
Information	The model is analysed under various information structures

TABLE 3.1. Summary - model setup

The timing of the model, for an arbitrary job applicant, begins with nature deciding whether or not the first employer has a taste for discrimination, which occurs with probability α. This employer then chooses to hire or reject the applicant (or randomise, which leads to one of these two outcomes as the realisation). Next, nature determines whether the second employer has a taste for discrimination, again occurring with probability α. The second employer will choose whether or not to hire the applicant and the model is concluded. At this point, I make no additional assumptions on the structure of information or belief formation in the labour market.

I consider what happens to members of both the majority and discriminated against group as they progress in their labour market matches. As a measure of the difference in outcomes between groups in the labour market I will use a discrimination ratio.

Definition 3.2.1. (Discrimination ratio).

The *discrimination ratio* is the ratio of hiring probabilities for those in the majority and the discriminated against group respectively.

$$Disc_t = \frac{\mathsf{Pr}(emp_t|g = A)}{\mathsf{Pr}(emp_t|g = B)}$$

Using this type of ratio (often called a callback ratio in correspondence studies, since the final employment outcome is not typically observed, but the response to a fictitious CV is) as a measure of group differences is common in the empirical literature (for example, Bertrand, Mullainathan, and Shafir, 2004; Banerjee et al., 2009; Carlsson and Rooth, 2007). Discrimination/callback ratios provide two convenient properties (Di Stasio and Heath, 2019). Firstly, they are an intuitive measure of group differences; a ratio of 2 would have the interpretation that members of the majority group are twice as likely to find employment than those of the discriminated against group. Secondly, this ratio has the property of

constant returns to scale[11], so that an economic shock which might double the employment rate for each group does not alter the ratio, but would alter the raw differential in employment rates.

Analysing the model in two-periods is sufficient for insights into how initial discrimination can evolve dynamically. Additionally, extending into further periods comes at the expense of extra assumptions, and additional dynamics which distract from the primary insights. A particularly important assumption which needs to be made is that of the observation window.[12] It is not realistic to assume that employers have a complete history of outcomes, especially for workers further along the life cycle who may have a large number of past jobs, or in industries with norms on the maximal CV length. Thus, neither in my simple model, nor in the full model with relaxed assumptions, do I study with model with more than two periods of time.

3.3 Full Information Benchmark

Although it is not the most realistic information structure I consider, I begin by analysing the model under full information. This structure of information is both a natural starting point for theoretical analysis and a useful point of comparison with the other structures. In this benchmark all of the employers know not only the proportion of taste discriminators in the economy, but also exactly which other employers were or were not discriminating.

It is intuitive that discriminating employers are unable to confound the true signal of employees whom have been rejected due to discrimination, since employers rationally discard any implied signal from the hiring choice of a discriminating employer, since he rejects applicants from the $g = B$ group with certainty. There is also no benefit-of-the-doubt effect, since employers know when a rejection is the result of being discriminated against and when it is not, and so have no reason to be especially lenient due to *potential* discrimination. In turn, this means that discrimination which is taste-based should not be exacerbated nor dampened in future periods via a signalling mechanism, and thus discrimination should be constant over time. Because the channels that arise under partial information and belief mis-specification do not arise in this setting, full information indeed provides a natural benchmark.

Lemma 3.3.1. *The ex-ante probability of employment for members of the majority group is constant over time*
$\mathsf{Pr}(emp_t|g = A) = \frac{1}{2} \forall t \in \{1, 2\}$

The proof is presented in the Appendix to Chapter 3.

This result follows intuitively from the fact that job applicants are high ability (and have a high ability signal) with probability one half, and the tie-breaking rule that at indifference an applicant is hired with probability one half. Those in the majority group do not suffer nor (directly) benefit from discrimination against the minority group.

[11] Formally, $Disc_t(\lambda \mathsf{Pr}(emp_t|g = A), \lambda \mathsf{Pr}(emp_t|g = B)) = \frac{\lambda \mathsf{Pr}(emp_t|g=A)}{\lambda \mathsf{Pr}(emp_t|g=B)} = Disc_t \ \forall \lambda > 0$

[12] Kabos and Meyer (2019) emphasise the importance of limited availability of observed decisions in the welfare analysis of observational learning models. (This is unpublished work which has been shared in private correspondence with one of the authors.)

Proposition 3.3.2. *Under full information, the discrimination ratio is constant, at* $Disc_t = \frac{1}{1-\alpha} \forall t \in \{1, 2\}$.

The proof is presented in the Appendix to Chapter 3.

Thus, Proposition 3.3.2 confirms the intuition that group differences remain constant over time when there is full information about the discriminatory attitudes of previous employers. Because it is known if a candidate has previously been unemployed due to discrimination, the unemployment is not interpreted as a negative signal. Rather, they are evaluated only through the second period signal, since no information about the candidate quality is transmitted from the unemployment caused by discrimination. Because unemployed candidates who have not been discriminated against are known to have not been discriminated against, they do not benefit from any lenience and are evaluated in the same way as a member of the A group would be. This result is not without loss of generality and relies on the specific assumptions of symmetry and the tie-breaking coin flip rule. Nevertheless, this version of the model is purposefully simple and serves as a point of reference for the more realistic assumptions in the setting of partial information and naive beliefs.

3.4 Naive Belief System

Under a naive belief system all employers, including those with a taste for discrimination, believe that the previous employers in the sequence are non-discriminating. I will first motivate this form of belief mis-specification and then present the main results in the two-preiod model.

Assumption 3.4.1. (Naive-beliefs).
Each employer, regardless of their taste for discrimination, has at time $t = 2$ *the belief that the matched job applicant has not been previously discriminated against. Mathematically:*

$$\hat{\mu}_{t=2}(T_{t=1} = 1|g, \mathcal{O}_{t=1}) = 0 \ \forall g \in \{A, B\}, \mathcal{O}_{t=1} \in \{emp, \neg emp\}$$

3.4.1 Motivating the Belief System

Although this form of belief mis-specification is a strong assumption, there are many avenues through which it can be motivated. Recall that in the model an α share of the employers have a taste for discrimination. It would be reasonable to assume that employers know that these employers exist, even if they do not know the exact proportion, or who the individual employers are. It would also be reasonable for employers to update their beliefs, for example noting that if a member of the discriminated against group is unemployed in the last period then it is likely to be the result of taste discrimination. This relies, however, on employers noting the correlation between membership of the group and previous employment outcomes.

The correlation neglect literature (Levy and Razin, 2015; Enke and Zimmermann, 2017) provides evidence that in many scenarios, individuals fail to note the correlation between observations which are related to one another and therefore do not update their beliefs correctly. Firms can be though of as mistakenly ignoring the fact that group membership and

the previous employment outcome are informative of potential discrimination. Indeed, the prevalence of 'bias awareness' campaigns as a strategy to reduce discriminatory behaviour may indicate that individuals do underestimate the existence of taste discrimination. Providing evidence that well-designed interventions can be effective Boring and Philippe (2021) indicate that an underestimation of discrimination is likely. Andre (2024) finds experimental evidence that meritocracy is 'shallow' with respect to individuals' past circumstances, which further supports a behavioural environment where (potential) past discrimination is not taken into account and employers behave as if there is no past discrimination.

Additionally, naive beliefs can be motivated by turning to the literature on implicit discrimination (Bertrand, Chugh, and Mullainathan, 2005), which suggests that sometimes discrimination may be unintentional and that those who discriminate may be unaware of their 'implicit attitude.' Since discrimination is more likely to be implicit rather than explicit in scenarios where decision makers are under time pressure, making judgments quickly and intuitively, it is especially likely to be applicable in hiring decisions where employers are busy with many tasks and evaluating numerous job candidates. The taste-discriminators in the economy may thus not realise that they impose a penalty against specific group members at the point of evaluation and the other employers may not explicitly consider the possibility that these job candidates have been discriminated against in the past. Thus, employers may fail to consider past discrimination, since it occurs implicitly, and act with an incorrect belief that there is no taste discrimination in the economy.

3.4.2 Ex-ante discrimination

Under the model featuring a naive belief system, since employers believe (with perceived certainty) that job applicants have never been discriminated against, they use $\mathcal{O}_{t=1}$ to directly infer $x_{t=1}$. This means that unemployment is interpreted as a low-ability signal, and employment is interpreted as a high-ability signal, no matter the circumstance. This interpretation is correct when applicants are evaluated without bias, but incorrect if unemployment arises from discrimination.

Definition 3.4.1. (Signal-distortion effect).

The *signal distortion effect* occurs when taste-based discrimination results in a job applicant's true signal x_{t-1} being distorted through $\mathcal{O}_{t-1}$ and falsely interpreted as $x'_{t-1}(\neq x_{t-1})$.

It is intuitive in this model that the signal distortion effect will hurt those in the $g = B$ group. Some high ability job candidates will be rejected on the basis of discrimination, but this observed employment outcome will incorrectly be perceived as a negative signal of ability.

Lemma 3.4.2. *In the first-period of the naive beliefs model the discrimination ratio is the same as in the full-information benchmark:*
$Disc_{t=1} = \frac{\mathsf{Pr}(emp_{t=1}|g=A)}{\mathsf{Pr}(emp_{t=1}|g=B)} = \frac{1}{1-\alpha}$

Proof:

At $t = 1$ the naive beliefs model is identical to the full information model. Apply the corresponding arguments from the proof to Proposition 3.3.2. □

This result follows intuitively from the fact that in the first period of the model there is no history to observe.

In the second period, however, the first period employment outcomes are observed and naive employers become subject to the signal-distortion effect, hurting members of the discriminated against group. This effect does not impact members of the majority group, since it is (correct) common knowledge that they are never discriminated against and thus are treated equivalently with the full information benchmark. Proposition 3.4.3 thus follows intuitively.

Proposition 3.4.3. *In the two period model with naive employer beliefs, the discrimination ratio increases between the first and the second period.*

The proof is presented in the Appendix to Chapter 3.

Again, this result is intuitive. Discrimination in the first period results in observed unemployment for those in the B group. Since employers in the second period are naive, the signal distortion effect exacerbates group differences, since some discriminated group applicants who had a high ability signal in the first period will have this signal incorrectly inferred as a low ability signal[13]. It is especially clear to note how this occurs by examining Table 3.2 and observing that in the history where a discriminated against group applicant has never been discriminated against then she is treated equivalently to a majority group member, but in the case where the applicant has been discriminated against at $t = 1$ but not at $t = 2$ then the first period discrimination has a lasting effect.

Employment matching history				
$t = 1$	Taste	No-taste	Taste	No-taste
$t = 2$	Taste	Taste	No-taste	No-taste
$\mathsf{Pr}(emp_{t=2} \mid g = B)$	0	0	$\frac{1}{4}$	$\frac{1}{2}$

TABLE 3.2. State-contingent employment probabilities for $g = B$ at $t = 2$ under naive beliefs

This result can be seen in Figure 3.2 below, which plots the discrimination ratios in periods 1 and 2 against the proportion of taste discriminating employers in the economy.

[13] Meanwhile, under the full information benchmark, this signal would simply be discarded when discrimination is known to have occurred and led to unemployment.

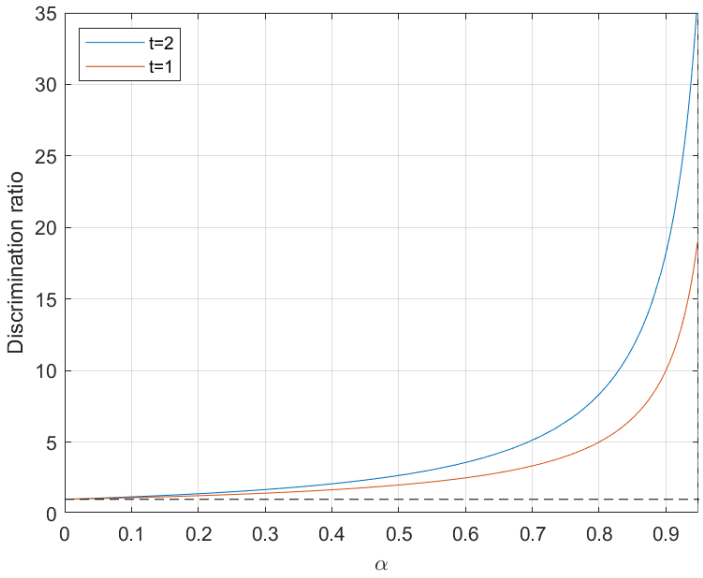


FIGURE 3.2. Discrimination ratios in a two-period naive beliefs model
Notes: The dashed line fixes a discrimination ratio of one, as a reference point.

3.5 Partial Information

Another framework which can be used to analyse the model is a partial information framework in which employers do not know exactly which other employers are taste-discriminators, but they have a correct prior on the proportion of discriminating employers, α.

One reasonable assumption which could achieve this information structure is to suppose an employer in the second period can observe the first period outcomes for both employee groups in the labour market as a whole. By observing the difference in the employment rates of those in the majority and discriminated against group, employers are able to deduce the proportion which taste-discriminate. Employers would not, however, know exactly which other employers taste-discriminate.

Lemma 3.5.1. *In the first-period of the partial information model the discrimination ratio is the same as in the full-information benchmark:*
$Disc_{t=1} = \frac{\Pr(emp_{t=1}|g=A)}{\Pr(emp_{t=1}|g=B)} = \frac{1}{1-\alpha}$

Proof:

At $t = 1$ the partial information model is identical to the full information model. Apply the corresponding arguments from the proof of Proposition 3.3.2. □

As with Lemma 3.4.2, this result follows intuitively from the fact that in the first period of the model there is no history to observe.

During the second round of matching, employers have a correct prior belief on the probability that an applicant's past employer was taste-discriminating. When an employer observes the first-period hiring outcome, which they can use to infer the signal, they update their beliefs on the type of the previous employer in the sequence.

Let the indicator function $T_t = \mathbb{1}\{\text{Taste-Discriminator}\}$ denote an employer's type in period t and $\mu_{t=2}(T_{t=1} = 1|\mathcal{O}_{t=1}, g)$ denote an employer's posterior belief that the matched employer at time $t = 1$ had a taste for discrimination, given the available information. In the partial information model beliefs are updated correctly, following Bayes' rule.

Lemma 3.5.2. *If a member of the B group is not hired in the first period, the updated belief that the previous employer was a taste-discriminator is higher than the unconditional probability that the previous employer was a taste-discriminator. Conversely, if the applicant was hired in the first period, then the updated belief that she was discriminated against is zero.*

The proof is presented in the Appendix to Chapter 3.

Of course, this result is intuitive - if a member of the discriminated group has been employed in the first period then she must have matched with an employer without taste for discrimination. If the applicant, on the other hand, is seen to be unemployed then it is more likely that she has been discriminated against since those who match with a discriminating employer always end up unemployed. It is also trivial to note that the outcomes of an A group member do not depend on the type of the matched employers, and thus the belief system is irrelevant to members of the A group.

Employers at time $t = 2$ thus use the first period employment outcome to infer the signal of quality taking into account that the employment decision may have been affected by taste-discrimination.

The main observation to be made in the partial information model is that when a member of the discriminated against group matches with a non-discriminating employer at $t = 2$ she gains a benefit of the doubt.

Definition 3.5.1. (Benefit-of-the-doubt effect).

The *benefit-of-the-doubt effect* occurs when unemployment in the first period for a discriminated against group member results in a strictly positive posterior belief that the applicant was discriminated against. Mathematically, $\mu_{t=2}(T_{t=1} = 1|g = B, \mathcal{O}_{t=1} = \neg emp) > 0$.

This effect can be noted, for example, in the case where the observed second period signal pair is $X_{t=2} = \{\neg emp, H\}$, since the rejection in the first period *might* have been the result of taste discrimination. Instead of inferring that the first signal was L, the employer takes into account that it *may* have been H, and that the candidate was simply rejected due to discrimination. This is in contrast to the naive beliefs model, since the signal-distortion effect makes the applicant worse off in that particular matching history, since the unemployment is believed with certainty to have resulted from an L signal.

3.5.1 Ex-ante discrimination

The evolution of discrimination in the partial information model follows from the behaviour of employers when matching with applicants from the discriminated against group. Note that the benefit-of-the-doubt effect only requires first-period unemployment for discriminated against applicants. It does not require that the unemployment was necessarily caused by

discrimination. Therefore, employers will be unable to distinguish between applicants who were previously discriminated against and those who were evaluated fairly but were rejected due to the realisation of a low ability signal. For those in the majority group, nothing has changed compared to the full information benchmark, since it is commonly known under both information structures that members of this group never suffer from taste discrimination. Proposition 3.5.3 thus follows intuitively.

Proposition 3.5.3. *In the two period partial information model, the discrimination ratio decreases between the first and the second period.*

The proof is presented in the Appendix to Chapter 3.

The intuition underlying this result is that there is an uncertainty about whether a rejection in the first period has been caused by a low ability signal, or by taste-discrimination. Since unemployment might have been the result of discrimination, which results in employers rationally discounting the implied first period signal, some members of the $g = B$ group gain from a benefit-of-the-doubt effect. Not only do candidates who have suffered from discrimination have this accounted for, but so do those who were not discriminated against but were rejected due to the realisation of a low ability signal.[14]

Employment matching history				
$t = 1$	Taste	No-taste	Taste	No-taste
$t = 2$	Taste	Taste	No-taste	No-taste
$\mathsf{Pr}(emp_{t=2} \mid g = B)$	0	0	$\frac{1}{2}$	$\frac{1+p-p^2}{2}$

TABLE 3.3. State-contingent employment probabilities for $g = B$ at $t = 2$ under partial information

Another interpretation of the mechanism driving this result is that an affirmative action policy arises naturally to profit maximising firms. There is a large literature in economics evaluating affirmative action policies which are imposed on a labour market (for example, Holzer and Neumark, 2000; Fryer and Loury, 2005; Moro and Norman, 2003). Such policies are usually motivated with an argument that they are correcting discrimination which has affected a given group in the past and/or the present. However, the benefit-of-the-doubt effect and the resulting decrease in the discrimination ratio can be viewed as an *internalisation* of such a policy, since rational Bayesian firms have an incentive to account for past discrimination when evaluating job candidates.

Given the result that under partial information the discrimination ratio decreases over time, a natural follow up question is whether this can be strong enough to reverse the differences in employment between groups (which are unambiguously in favour of the majority group in period 1).

[14]This effect is in line with a similar result found in Meyer (1991), in which an organisation evaluates two employees for a single promotion decision. Biasing a contest towards one employee in an early period may ultimately increase the other employee's chance of being awarded the promotion, since the past disadvantage is taken into account when forming beliefs.

Remark 3.5.4. *In the two period partial information model, there are some pairs of* (α, p) *for which the direction of employment differences between groups reverse.*

The proof is presented in the Appendix to Chapter 3.

Besides just identifying that this reversal is possible, the pairs of (α, p) which for which this occur exactly satisfy the inequality $(1-\alpha)(\alpha+(1-\alpha)(1+p-p^2)) > 1$.

This occurs when members of the discriminated against group are hindered very little by taste-discrimination whilst gaining a significant benefit-of-the-doubt due to the existence of the small proportion of taste discriminators. This can be seen in Panel (a) of Figure 3.3, which shows that the discrimination ratio lies below 1 for small values of α at $t = 2$. [15]

Since some members of the discriminated group gain a benefit that they otherwise would not, due to the partial information structure, it follows intuitively that if the signal strength, p, increased then the gain would be lower and the groups' employment probability would decrease. The reason for this is that, other than those who have genuinely been discriminated against in the first period, those who benefit from the extra lenience in evaluation are those who were rejected due to the realisation of a low ability signal. Since the signals are informative of ability, these individuals are more likely to have been from the low ability group. Thus, by increasing the strength of the signal, p, fewer low ability members of the discriminated against group 'slip through the cracks' of the benefit-of-the-doubt effect. In other words, when signals are more precise, the pool of those in first-period unemployment are more likely to be there because they are of genuine low ability, rather than an unlucky signal. This means that less of them gain in the second period from a benefit-of-the-doubt, since less of the people in this pool will go on to generate a high ability signal. In Panel (b) of Figure 3.3, the shaded area represents the combinations of small enough α and p for which the discrimination ratio lies below one, implying that those in the $g = B$ group are more likely to be employed than those in the $g = A$ group.

[15] The intuition for the requirement that α is small is simple. Because this result concerns not the change in group differences over time, but an *absolute* difference between groups, it must be the case that there is sufficiently little taste-discrimination in order for those in the B group to be more employable than those in the A group.

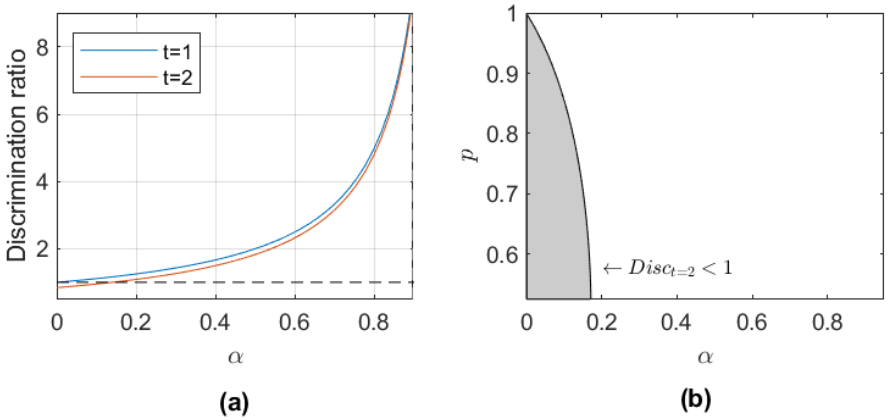


FIGURE 3.3. Discrimination ratios in the partial information model

Notes: Panel (a) plots discrimination ratios in the partial information model, which depend on both α and p; I have fixed $p = 0.75$. The dashed line fixes a discrimination ratio of one, as a reference point. Panel (b) plots the combinations of α and p which result in a discrimination ratio below 1, which is represented by the shaded area.

3.6 Reduced-form Evidence

My theoretical results make several predictions on the dynamics of group differences in employment over time, depending on the structure of information and belief systems in the economy. Since the predictions are empirically testable, I compare how my theoretical results match with the data. It is important to note that I am not making any claims about causality.

3.6.1 Data

The National Longitudinal Survey of Youth (NLSY79) follows a panel of 12,686 young people from 1979 who were aged 14 to 22 when first interviewed. From 1979 and until 1994 the survey was conducted on an annual basis and from then on has been conducted every other year. Included in the dataset are standard demographic variables, such as gender and race, as well as education and employment data. Labour force status is only available up to 2006 and because it is the primary outcome variable it is also the final year which I include in my analysis. Table 3.4 provides key descriptive statistics for the cohort in 1979.

3.6.2 Methodology

My theoretical results have shown that depending on the information structure in the economy, group differences in employment can be exacerbated or reduced over the life cycle of labour force participants. A precise name to describe the progression of the life cycle for individuals in the labour force is years of *potential experience*. I filter the dataset to only include observations of individuals who have completed their education, whether it finished at high school or at college level. This filtering firstly eliminates any time variation in education, meaning that education becomes an individual-level fixed effect which can be swept

Variable	N	Mean	Median	Pctl(25)	Pctl(75)
Panel (a): Continuous Variables					
Age	12,686	17.86	18	16	20
Panel (b): Dummy Variables					
$\mathbb{1}\{\text{Male}\}$	12,686	50.47%			
$\mathbb{1}\{\text{Black}\}$	12,686	25.02%			
$\mathbb{1}\{\text{Hispanic}\}$	12,686	15.78%			
$\mathbb{1}\{\text{In Education}\}$	12,686	58.54%			
$\mathbb{1}\{\text{Employed}\}$	12,685	40.69%			

TABLE 3.4. Summary Statistics for 1979 Cohort
Notes: In the NLSY79 dataset race is categorised as either black, hispanic, or non-black and non-hispanic.

out. It also allows for the calculation of individual potential experience, which I define as the difference between the current year and the last year in which the individual was in education.[16] For those in the dataset who were never in full time education during the years covered, I assume that their final year of education was 1979, the first year of the panel. I include only individuals who are members of the labour force.

I use fixed-effects models to estimate the returns to potential experience for both groups. Since I have filtered on those who have already completed their education, individual level education and ability are both time-invariant fixed effects which are swept away. The regressions which I run are:

$$\mathbb{1}\{\text{Employed}\}_{i,t} = \beta_1 \text{Potential Experience}_{i,t} + \beta_2 \mathbb{1}\{\text{Black}\}_i * \text{Potential Experience}_{i,t} + (\eta_i + \gamma_t + \epsilon_{i,t})$$

and,

$$\begin{aligned}\mathbb{1}\{\text{Employed}\}_{i,t} = {} & \beta_1 \text{Potential Experience}_{i,t} + \beta_2 \mathbb{1}\{\text{Black}\}_i * \text{Potential Experience}_{i,t} \\ & + \beta_3 (\text{Potential Experience}_{i,t})^2 + \beta_4 \mathbb{1}\{\text{Black}\}_i * (\text{Potential Experience}_{i,t})^2 + (\eta_i + \gamma_t + \epsilon_{i,t})\end{aligned}$$

I also run these regressions using time fixed-effects to account for any year-specific explanatory variation. Since the years are perfectly correlated with potential experience, these regressions are unable to estimate the β_1 parameter, however they do allow me to estimate β_2, which still has the interpretation as the returns to potential experience relative to the base group, as well as the coefficients on the quadratic variables.

3.6.3 Results

The results of these regressions, using individual only and individual and time fixed-effects, are reported in Table 3.5. The coefficients indicate a quadratic and concave relationship between potential experience and employment. The main coefficients of interest are those on

[16] Years of potential experience is often defined equivalently as: $age - education - 6$, when it is assumed that 6 is the age at which individuals begin their schooling (Card, 1999)

the Black ∗Potential Experience interaction term. The base group is all those in the dataset who are non-black and the coefficients on the interactions are interpreted as the returns to potential experience relative to the base group. Across all the specifications in Table 3.5 the Black ∗Potential Experience interaction term has a positive and highly statistically significant coefficient. Thus, this matches the graphical evidence and the theoretical predictions of the partial information model. The evidence thus aligns with a story in which members of a discriminated group suffer significantly when entering the labour market, but having a poor performance during the early years of life does not exacerbate group differences in the future. Rather, employers are aware that members of the discriminated against group are likely to have been discriminated against in the past and are thus given a benefit of the doubt by future employers. Of course, there are other interpretations of the data, outside the scope of my model.[17]

My results also indicate that the relationship between potential experience and employment is quadratic, with a higher degree of concavity for those in the black group relative to the base group. We can interpret this relationship as one in which employment for those in the black group quickly converges to the base group during the early years of their careers. However, this catch-up diminishes over time, and the extent of decreasing marginal returns to potential experience is greater for those in the black group.

	Dependent variable: Employment Dummy			
	(1)	(2)	(3)	(4)
Potential Experience	0.006***	0.015***		
	(0.0002)	(0.001)		
Black ∗Potential Experience	0.004***	0.008***	0.005***	0.009***
	(0.0005)	(0.001)	(0.0005)	(0.001)
Potential Experience Sq		−0.0004***		−0.0001***
		(0.00002)		(0.00003)
Black ∗Potential Experience Sq		−0.0002***		−0.0002***
		(0.00005)		(0.00005)
Indiv fixed-effects:	Yes	Yes	Yes	Yes
Time fixed-effects:	No	No	Yes	Yes
N	115530	115530	115530	115530

TABLE 3.5. Empirical Results

Notes: *p<0.1; **p<0.05; ***p<0.01
Standard errors in brackets. In all specifications standard errors are cluster-robust at the individual level.

Because the model presented is purposefully simple, and the empirical results are intended to only have a descriptive interpretation, I relax several of the assumptions made in the simple

[17] For example, a situation in which employers statistically discriminate at first, but rely less on group membership over time, as actual performance becomes observable.

model and use it to analyse the dynamics of discrimination. The intuitions developed by the simple model are robust to the full model, but without the purposeful symmetry imposed, the results deliver predictions which are less sharp. This emphasises the care which needs to be taken when interpreting the simple predictions and its link to the empirical analysis.

3.7 The Full Model

3.7.1 Model setup

Many aspects of the simple model remain, but the structure imposed on the signalling technology, is much less restrictive in the full model, with no coin-flipping rule to generate symmetry in the results. Additionally, strong dis-taste is relaxed, meaning that it is possible in some circumstances for group B job candidates to be hired by discriminating employers.

As before, job applicants come from two ex-ante identical groups, denoted by $g \in \{A, B\}$

And again, the ability of the job applicant is either high or low, denoted $a \in \{L, H\}$ and distributed according to $\mathsf{Pr}(a = H) = p_H$

Worker output is a function of ability:

$$Y(a) = \begin{cases} V \text{ if } a = H \\ 0 \text{ otherwise} \end{cases}$$

In each of two periods of discrete time a job candidate (she) is matched with an employer (he) who may or may not have a taste for discrimination. Taste for discrimination is denoted by the indicator function T_t, which takes on the value of 1 if the matched employer has a taste for discrimination and 0 otherwise.

The prior probability of matching with an employer who has a taste for discrimination is $\mathsf{Pr}(T = 1) = \alpha \in (0, 1)$

Job candidates have a signals of ability $\theta_t \in \Theta$ which are drawn from distributions f_H and f_L for high and low ability candidates respectively and are conditionally iid. The signal has the Monotone Likelihood Ratio Property (MLRP), meaning that:

$\frac{f_H(\theta_t)}{f_L(\theta_t)} > \frac{f_H(\theta'_t)}{f_L(\theta'_t)} \; \forall \theta_t > \theta'_t$

In words, this assumption means that a larger signal implies a relatively higher chance that the ability of the candidate is high compared to low.[18]

Denoting by $\mathcal{O}_t$ the *observed* employment outcome from period t, and at this point making no assumptions on firms' knowledge on the discriminatory attitudes of others, in each period firms observe:

$$X_t = \begin{cases} \theta_1 \text{ if } t = 1 \\ \{\mathcal{O}_1, \theta_2\} \text{ if } t = 2 \end{cases}$$

The cost of employment is exogenous and fixed at w $(< V)$, and firms of type $T = 1$ face a composite employment cost when employing a $g = B$ group member of $w + d$, where d represents the psychic cost of hiring a member of this group. Again, w can be interpreted as an exogenous determined wage, or more loosely as a fixed cost of hiring, with wages agreed according to some bargaining process.

[18] This is a commonly used assumption in the theoretical literature, for example, in Coate and Loury (1993).

Workers, again, always prefer employment to non-employment.

Firms' expected profits take the form:

$\mathbb{E}(\pi_t|X_t, g, T) = \mathsf{Pr}(a = H|X_t)V - (w + \mathbb{1}\{g = B\} \cdot \mathbb{1}\{T = 1\} \cdot d)$

3.7.2 Full Information

Under full information, firms observe X_t and are also aware of whether or not the previous firm in the sequence operated with a taste for discrimination. The information set in period 1 can be summarised by $\mathcal{I}_{1,full} = \{\theta_1, g\}$ and in period 2 by $\mathcal{I}_{2,full} = \{X_2, g, T_1\}$

3.7.3 Period 1

I begin by studying first period behaviour.

If $g = A$, or $g = B$ and $T = 0$, then expected profits are given by:

$\mathbb{E}(\pi_1|\theta_1, g, T) = \frac{f_H(\theta_1)p_H}{f_H(\theta_1)p_H + f_L(\theta_1)(1-p_H)}V - w.$

The hiring condition requires that the candidate must make at least zero economic profits:

$\mathbb{E}(\pi_1|\theta_1, g, T) = \frac{f_H(\theta_1)p_H}{f_H(\theta_1)p_H + f_L(\theta_1)(1-p_H)}V \geq w$

which can be rearranged to $\frac{f_L(\theta_1)}{f_H(\theta_1)} \leq \frac{V-w}{w}\frac{p_H}{1-p_H}$. I now define $h(\theta_t) \coloneqq \frac{f_L(\theta_t)}{f_H(\theta_t)}$. Since h is a decreasing function by MLRP, h^{-1} is also a decreasing function. The first period hiring condition can now be rewritten as:

$$\theta_1 \geq h^{-1}\left(\frac{V-w}{w}\frac{p_H}{1-p_H}\right) (\coloneqq \tilde{\theta}_1) \tag{3.1}$$

This equation tells us that for an individual to gain employent then her signal must be sufficiently big compared to a threshold value. The threshold value decreases with the value of hiring a high ability candidate and the prior probability that the candidate is high ability, and increases with the wage cost of hiring the candidate.

Note that in the equation above $h^{-1}(\cdot)$ is defined, since $V > w$. When there is taste-discrimination, there are two cases:

Strong Dis-taste. Under this special case, $w + d > V$, meaning that even if a job candidate were high ability with certainty, they would not be worth employing. In other words, members of the B group are always rejected when randomly matching with a taste-discriminator.

General case. In this more general case, $w + d \leq V$, meaning that when the ability relevant information is sufficiently strong, taste-discriminators are willing to hire B group members.

Applying the corresponding arguments when $g = B$ and $T = 1$, the hiring condition in the general case becomes:

$$\theta_1 \geq h^{-1}\left(\frac{V-(w+d)}{w+d}\frac{p_H}{1-p_H}\right) (\coloneqq \tilde{\theta}_1^d) \tag{3.2}$$

Becoming employed thus requires a stronger signal in this particular candidate-employer match, since $\tilde{\theta}_1^d > \tilde{\theta}_1$. Or in other words, when candidates are discriminated against, they must meet a higher standard in a interview or standardised test than a candidate who is not discriminated against.

Note that mathematically the function h is a ratio of density functions and thus has co-domain $\mathbb{R}^+$. In the Strong Dis-taste special case, $h^{-1}(\cdot)$ would not exist, which follows the economic intuition that there would not exist any sufficiently strong signal to make the taste-discriminator willing to hire the applicant.

Ex-ante Employment

For group A candidates the first-period employment probability is:

$$\mathsf{Pr}(emp_1|g = A) = (1 - F_H(\tilde{\theta}_1))p_H + (1 - F_L(\tilde{\theta}_1))(1 - p_H)$$

For group B candidates the first-period employment probability depends on the case. In the Strong Dis-taste case we have:

$$\mathsf{Pr}(emp_1|g = B) = (1 - \alpha)[(1 - F_H(\tilde{\theta}_1))p_H + (1 - F_L(\tilde{\theta}_1))(1 - p_H)]$$

Thus, so long as the group B candidate matches with a non-discriminating employer, they are as likely to be employed as an A group candidate. Otherwise, they are definitely unemployed.

While in the general case:

$$\begin{aligned}\mathsf{Pr}(emp_1|g = B) =&(1 - \alpha)[(1 - F_H(\tilde{\theta}_1))p_H + (1 - F_L(\tilde{\theta}_1))(1 - p_H)] \\ &+\alpha[(1 - F_H(\tilde{\theta}_1^d))p_H + (1 - F_L(\tilde{\theta}_1^d))(1 - p_H)] \\ &= \left(1 - \alpha F_H(\tilde{\theta}_1^d) - (1 - \alpha)F_H(\tilde{\theta}_1)\right)p_H + \left(1 - \alpha F_L(\tilde{\theta}_1^d) - (1 - \alpha)F_L(\tilde{\theta}_1))\right)(1 - p_H)\end{aligned}$$

In this case, when matched with a non-discriminating employer they are treated the same as a group A candidate would be, but when matched with a discriminating employer, they must meet a higher standard and are thus less likely to be employed.

Thus, overall, $\mathsf{Pr}(emp_1|g = A) > \mathsf{Pr}(emp_1|g = B)$

3.7.4 Period 2

In the second period there is now another ability signal, θ_2 as well as a realisation, $\mathcal{O}_1$, of being employed or not in the first period. Recall that the signals are conditionally independent, and since the employment outcome is a function of the random first period signal realisation, it is also independent of the second period signal.

The posterior belief that the candidate is high ability follows:

$\mathsf{Pr}(a = H|\mathcal{O}_1, \theta_2) = \frac{\mathsf{Pr}(\mathcal{O}_1|a=H)\mathsf{Pr}(\theta_2|a=H)\mathsf{Pr}(a=H)}{\mathsf{Pr}(\mathcal{O}_1|a=H)\mathsf{Pr}(\theta_2|a=H)\mathsf{Pr}(a=H)+\mathsf{Pr}(\mathcal{O}_1|a=L)\mathsf{Pr}(\theta_2|a=L)\mathsf{Pr}(a=L)}$

If $g = A$, or $g = B$, $T_1 = 0$ and $T_2 = 0$ then, following some rearranging of terms, expected profits when $\mathcal{O}_1 = emp$ are:

$$\mathbb{E}(\pi_2|emp, \theta_2, g, T) = \frac{1}{1 + \frac{(1-F_L(\tilde{\theta}_1))f_L(\theta_2)(1-p_H)}{(1-F_H(\tilde{\theta}_1))f_H(\theta_2)p_H}} V - w$$

Rearranging yields the employment condition:

$$\theta_2 \geq h^{-1}\left[\frac{V-w}{w}\frac{p_H}{1-p_H}\underbrace{\frac{1-F_H(\tilde{\theta}_1)}{1-F_L(\tilde{\theta}_1)}}_{>1 \text{ (MLRP implies FOSD)}}\right] (:= \tilde{\theta}_{2,emp})$$

Because of the extra information implied by the candidate being observed as employed in period one, the threshold second period signal for employment is less than that required in the first period, $\tilde{\theta}_{2,emp} < \tilde{\theta}_1$. The intuition is that an impressive outcome in period one makes the candidate a more convincing hire, so a less impressive signal in period two is required.

Conversely, if $\mathcal{O}_1 = \neg emp$ then:

$$\mathbb{E}(\pi_2|\neg emp, \theta_2, g, T) = \frac{1}{1 + \frac{F_L(\tilde{\theta}_1)f_L(\theta_2)(1-p_H)}{F_H(\tilde{\theta}_1)f_H(\theta_2)p_H}} V - w$$

Rearranging yields the employment condition:

$$\theta_2 \geq h^{-1}\left[\frac{V-w}{w}\frac{p_H}{1-p_H}\underbrace{\frac{F_H(\tilde{\theta}_1)}{F_L(\tilde{\theta}_1)}}_{<1 \text{ (FOSD)}}\right] (:= \tilde{\theta}_{2,\neg emp})$$

For similar reasons to first period employment lowering the second period signal threshold, first period unemployment raises the second period signal threshold $\tilde{\theta}_{2,\neg emp} > \tilde{\theta}_1$. Because first period unemployment is a negative indicator of the candidate's ability, an especially strong signal of ability is required in the second period to convince the employer to hire her.

I now turn to the general case where $g = B$, $T_1 = 1$ and $T_2 = 0$. Since it is known that $T_1 = 1$, the threshold for employment, $\tilde{\theta}_1^d$, is also known. Note that in the Strong Dis-Taste case, employment in the first period never occurs in this type of match.

Following similar rearranging as above, the employment condition when $\mathcal{O}_1 = emp$ is:

$$\theta_2 \geq h^{-1}\left[\frac{V-w}{w}\frac{p_H}{1-p_H}\underbrace{\frac{1-F_H(\tilde{\theta}_1^d)}{1-F_L(\tilde{\theta}_1^d)}}_{>1}\right] (:= \tilde{\theta}_{2,emp}^d)$$

Proposition 3.7.1. *A job candidate who is known to have been discriminated against in period 1 and employed regardless, but is not discriminated against in period 2, has a lower second period signal threshold than a candidate who was employed in period 1 but is known to have not been discriminated against. Mathematically:* $\tilde{\theta}_{2,emp}^d < \tilde{\theta}_{2,emp}$

The proof is presented in the Appendix to Chapter 3

I call this the *overcoming adversity* effect, where gaining employment while being discriminated against provides a stronger signal than gaining employment while being evaluated fairly. The logic is that exceeding a relatively more difficult hiring criteria is especially impressive.

The employment condition when $\mathcal{O}_1 = \neg emp$ is:

$$\theta_2 \geq h^{-1}\left[\frac{V-w}{w}\frac{p_H}{1-p_H}\underbrace{\frac{F_H(\tilde{\theta}_1^d)}{F_L(\tilde{\theta}_1^d)}}_{<1}\right] (:= \tilde{\theta}_{2,\neg emp}^d)$$

Proposition 3.7.2. *A job candidate who is known to have been discriminated against in period 1 and be unemployed, but is not discriminated against in period 2, has a lower second period signal threshold than a candidate who was also unemployed in period 1 but is known to have not been discriminated against. Mathematically:* $\tilde{\theta}_{2,\neg emp}^d < \tilde{\theta}_{2,\neg emp}$

The proof is presented in the Appendix to Chapter 3

The logic here is that while failing to meet the previous hiring criteria and being unemployed is a negative signal, it is less negative if the hiring criteria was relatively more difficult to overcome.

In the Strong Dis-taste case it is known that $T_1 = 1$, which implies rejection with certainly, and thus first period unemployment for the job candidate signals nothing about underlying ability[19].

The employment condition when not discriminated against is thus: $\theta_2 \geq \tilde{\theta}_1$. Or in other words, since no relevant information is provided by the first period unemployment outcome, since discrimination guaranteed unemployment, the candidate is evaluated according to the initial hiring criteria using only one signal of quality.

I now study the case in which $g = B$, $T_1 = 0$ and $T_2 = 1$. Of course, in the Strong Dis-Taste special case, this scenario would simply mean that the candidate would be unemployed with certainty. Recalling that the second employer knows for sure that the candidate was

[19] Applying the conditional version of Bayes' rule $\mathsf{Pr}(a = H|\theta_2, \mathcal{O}_1 = \neg emp, T_1 = 1, g = B)$
$= \frac{f_H(\theta_2)\mathsf{Pr}(\neg emp|a=H,T_1=1,g=B)p_H}{f_H(\theta_2)\mathsf{Pr}(\neg emp|a=H,T_1=1,g=B)p_H + f_L(\theta_2)\mathsf{Pr}(\neg emp|a=L,T_1=1,g=B)(1-p_H)} = \frac{f_H(\theta_2)p_H}{f_H(\theta_2)p_H + f_L(\theta_2)(1-p_H)}$

not previously discriminated against, the employment condition in the general case when $\mathcal{O}_1 = emp$ is:

$$\theta_2 \geq h^{-1}\left[\frac{V-(w+d)}{w+d}\frac{p_H}{1-p_H}\frac{1-F_H(\tilde{\theta}_1)}{1-F_L(\tilde{\theta}_1)}\right] (:= \tilde{\theta}^{\emptyset,d}_{2,emp})$$

When $\mathcal{O}_1 = \neg emp$ the employment condition becomes:

$$\theta_2 \geq h^{-1}\left[\frac{V-(w+d)}{w+d}\frac{p_H}{1-p_H}\frac{F_H(\tilde{\theta}_1)}{F_L(\tilde{\theta}_1)}\right] (:= \tilde{\theta}^{\emptyset,d}_{2,\neg emp})$$

Finally, I study the case in which $g = B$, $T_1 = 1$ and $T_2 = 1$. Recalling that now the second employer knows for sure that the candidate was previously discriminated against, in the general case, the employment condition when $\mathcal{O}_1 = emp$ is:

$$\theta_2 \geq h^{-1}\left[\frac{V-(w+d)}{w+d}\frac{p_H}{1-p_H}\frac{1-F_H(\tilde{\theta}^d_1)}{1-F_L(\tilde{\theta}^d_1)}\right] (:= \tilde{\theta}^{d,d}_{2,emp})$$

When $\mathcal{O}_1 = \neg emp$ the employment condition becomes:

$$\theta_2 \geq h^{-1}\left[\frac{V-(w+d)}{w+d}\frac{p_H}{1-p_H}\frac{F_H(\tilde{\theta}^d_1)}{F_L(\tilde{\theta}^d_1)}\right] (:= \tilde{\theta}^{d,d}_{2,\neg emp})$$

Note that under Strong Dis-Taste the employment probability will be zero in these last four match possibilities.

Ex-ante Employment

Unlike in the case with the simple model, the full version of the model has not been build using a purposefully symmetric structure, and thus, the discrimination ratio may increase or decrease between periods 1 and periods 2, depending on f_H and f_L and their corresponding CDFs. The employment probabilities are presented for the candidates in group A and group B, under both strong dis-taste and the more general case.

$$\begin{aligned}\Pr(emp_2|g=A) = &\Big((1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H + (1-F_L(\tilde{\theta}_1))(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\Big)\\ &+\Big(F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H + F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big)\end{aligned}$$

In the Strong Dis-taste special case:

$$\begin{aligned}\mathsf{Pr}(emp_2|g=B) =&(1-\alpha)^2\Big[(1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big]\\&+\alpha(1-\alpha)\Big[(1-F_H(\tilde{\theta}_1))p_H+(1-F_L(\tilde{\theta}_1))(1-p_H)\Big]\end{aligned}$$

In the general case:

$$\begin{aligned}\mathsf{Pr}(emp_2|g=B) =&(1-\alpha)^2\Big[(1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big]\\&+(1-\alpha)\alpha\Big[(1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}^{\emptyset,d}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1))(1-F_L(\tilde{\theta}^{\emptyset,d}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}^{\emptyset,d}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}^{\emptyset,d}_{2,\neg emp}))(1-p_H)\Big]\\&+\alpha(1-\alpha)\Big[(1-F_H(\tilde{\theta}^{d}_1))(1-F_H(\tilde{\theta}^{d}_{2,emp}))p_H+(1-F_L(\tilde{\theta}^{d}_1))(1-F_L(\tilde{\theta}^{d}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}^{d}_1)(1-F_H(\tilde{\theta}^{d}_{2,\neg emp}))p_H+F_L(\tilde{\theta}^{d}_1))(1-F_L(\tilde{\theta}^{d}_{2,\neg emp}))(1-p_H)\Big]\\&+\alpha^2\Big[(1-F_H(\tilde{\theta}^{d}_1))(1-F_H(\tilde{\theta}^{d,d}_{2,emp}))p_H+(1-F_L(\tilde{\theta}^{d}_1))(1-F_L(\tilde{\theta}^{d,d}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}^{d}_1)(1-F_H(\tilde{\theta}^{d,d}_{2,\neg emp}))p_H+(1-F_L(\tilde{\theta}^{d}_1))(1-F_L(\tilde{\theta}^{d,d}_{2,\neg emp}))(1-p_H)\Big]\end{aligned}$$

As expected, without the symmetry of the signalling technology imposed in the simple model, sharp predictions cannot be made about the direction of the discrimination ratio over time. The full information case is thus used as a benchmark against which to compare the partial information and naive beliefs cases.

3.7.5 Naive Beliefs

I again examine the information structure of naive beliefs in the full model, whereby employers are naive to the presence of discrimination and do not account for it when interpreting the previous employment status as a signal of ability.[20] As was the case in the simple model, everything in the first period of the model remains identical to the full information setting because there is no history to observe.

3.7.6 Ex-ante second period employment

Because group A candidates are never discriminated against, which aligns with employers' beliefs, all of the second period employment outcomes are identical with the full information benchmark. Meanwhile, group B candidates are evaluated according to the threshold rules already developed under the full information case, but only with the thresholds which apply to group A candidates (which are equivalent to the thresholds applied to group B candidates

[20] Mathematically, they have the belief $\hat{\mu}_2(\mathcal{O}_1, g=B, \theta_2)=0$

when discrimination in the first-period did not occur). Even when discriminated against in period one, employers behave as if that discrimination did not occur.

In the strong dis-taste special case the ex-ante employment probability for those in the B group is thus:

$$\begin{aligned}\mathsf{Pr}(emp_2|g=B) =&(1-\alpha)^2\Big[\Big((1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1))(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\Big)\\&+\Big(F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big)\Big]\\&+\alpha(1-\alpha)\Big[\Big((1-F_H(\tilde{\theta}_{2,\neg emp}))p_H+(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big)\Big]\end{aligned}$$

It can be seen that employment only occurs in this case when the candidate is not discriminated against. When they were discriminated against in the previous period, and were unemployed with certainty, they face the cutoff employment decision rule corresponding to having not been discriminated against.

Proposition 3.7.3. *The discrimination ratio decreases over time in the strong distaste special case under naive beliefs.*

The proof is presented in the Appendix to Chapter 3.

The intuition behind Proposition 3.7.3 follows that of the simple model, whereby unemployment due to discrimination occurs and is exacerbated since unemployment will be interpreted as a negative signal, when instead it should be rationally discounted[21].

In the general case, the ex-ante employment probability for those in the B group is:

$$\begin{aligned}\mathsf{Pr}(emp_2|g=B) =&(1-\alpha)^2\Big[\Big((1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1))(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\Big)\\&+\Big(F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big)\Big]\\&(1-\alpha)\alpha\Big[\Big((1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}^{\emptyset,d}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1))(1-F_L(\tilde{\theta}^{\emptyset,d}_{2,emp}))(1-p_H)\Big)\\&+\Big(F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}^{\emptyset,d}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}^{\emptyset,d}_{2,\neg emp}))(1-p_H)\Big)\Big]\\&\alpha(1-\alpha)\Big[\Big((1-F_H(\tilde{\theta}^d_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H+(1-F_L(\tilde{\theta}^d_1))(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\Big)\\&+\Big(F_H(\tilde{\theta}^d_1)(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H+F_L(\tilde{\theta}^d_1)(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big)\Big]\\&\alpha^2\Big[\Big((1-F_H(\tilde{\theta}^d_1))(1-F_H(\tilde{\theta}^{\emptyset,d}_{2,emp})p_H+(1-F_L(\tilde{\theta}^d_1))(1-F_L(\tilde{\theta}^{\emptyset,d}_{2,emp}))(1-p_H)\Big)\\&+\Big(F_H(\tilde{\theta}^d_1)(1-F_H(\tilde{\theta}^{\emptyset,d}_{2,\neg emp}))p_H+F_L(\tilde{\theta}^d_1)(1-F_L(\tilde{\theta}^{\emptyset,d}_{2,\neg emp}))(1-p_H)\Big)\Big]\end{aligned}$$

I can now make strict comparisons between the naive beliefs and full information settings to characterise the welfare of group B workers. Recall that due to the lack to symmetry imposed in the model, under full-information it cannot be said whether or not the discrimination ratio increases or decreases between periods one and two, or if it remains constant as

[21] It can also be seen under Strong Dis-Taste that group B candidates are worse off in the second period under naive beliefs relative to the full information benchmark.

it did under the simple model. I thus compare the two information settings to one another, since it cannot be said for sure whether or not the discrimination ratio increases or decreases in this general case.

Proposition 3.7.4. *Group B jobseekers are worse off under the naive beliefs setting relative to the full information setting. Mathematically* $\mathsf{Pr}(emp_2|g = B)$ *is smaller in the partial information setting than the full information setting.*

Proof:

Compare the expressions for $\mathsf{Pr}(emp_2|g = B)$ under the naive beliefs and full information setting. □

Proposition 3.7.5. *The discrimination ratio in period 2 is larger under naive beliefs than it is under full-information.*

Proof: $\mathsf{Pr}(emp_2|g = A)$ is the same under the full information and naive beliefs settings. Proposition 3.7.4 shows that $\mathsf{Pr}(emp_2|g = B)$ is smaller under naive beliefs, thus the discrimination ratio is higher in period 2 under naive beliefs than under full information. □

The intuition is simple. Under full information, being known to have been discriminated against generated lenience in the signals required to become employed in the second period. When it is believed by second period employers that no discrimination has occurred, no such lenience is granted. Discrimination thus worsens over time in the naive beliefs setting relative to the full information benchmark.

3.7.7 Partial Information

Under partial information, firms observe the X_t and are *un-aware* of whether or not the previous firm in the sequence operated with a taste for discrimination. The information set in period 1 can be summarised by $\mathcal{I}_{1,partial} = \{\theta_1, g\}$ and in period 2 by $\mathcal{I}_{2,partial} = \{X_2, g\}$. In the second period, firms use the information available to them to update their prior belief on the probability that the candidate was previously discriminated against. The posterior belief that the candidate was discriminated against is denoted by
$\mu_2(\mathcal{O}_1, g, \theta_2) \coloneqq \mathsf{Pr}(T_1 = 1|\mathcal{O}_1, g, \theta_2)$.

First period behaviour is completely equivalent under full and partial information, since knowledge of other firms' behaviour only matters in the second period.

3.7.8 Posterior Beliefs

In the partial information setting, beliefs on the probability that an individual has been discriminated against are formed rationally, with Bayes' rule applied to the given information.

In the Strong Dis-Taste special case, the posterior belief that a candidate has been discriminated against following employment is rationally zero, since under strong dis-taste a group B candidate would never be employed by a discriminating employer. Mathematically,
$\mu_2(\mathcal{O}_1 = emp, g = B, \theta_2) = 0$

The logic is that in the strong dis-taste special case, it is impossible for a member of the B group to be employed while matching with a discriminating-type employer, thus observing

that the candidate has been employed rules out the case that she has been discriminated against in period 1.

Lemma 3.7.6. *Following employment, in the general case, the posterior belief that a candidate has been discriminated against in the first period is increasing in the strength of the ability signal in the second period.*

$$\frac{\partial \mu_2(\mathcal{O}_1 = emp, g = B, \theta_2)}{\partial \theta_2} > 0$$

The proof is presented in the Appendix to Chapter 3.

When a job candidate faces a discriminating employer, she must generate an especially strong signal of ability to overcome the discrimination and receive a job offer. She is more likely to do so if she is indeed a high ability type. A stronger signal of ability in the current (second) period suggests that the candidate is more likely to be a high ability type, who also generated a strong signal in the last period, and thus it is relatively more likely that the candidate was previously employed by being able to overcome discrimination.

Lemma 3.7.7. *Following unemployment, the posterior belief that a candidate has been discriminated against in the first period is increasing in the strength of the ability signal in the second period.*

$$\frac{\partial \mu_2(\mathcal{O}_1 = \neg emp, g = B, \theta_2)}{\partial \theta_2} > 0$$

The proof is presented in the Appendix to Chapter 3.

The logic of Lemma 3.7.7 is that a candidate may either have been previously unemployed due to discrimination or due to a poor realisation of the first period signal of ability. Given that a strong signal of quality in the current (second) period indicates that the candidate is more likely to be of high ability, the candidate assessor therefore is more likely to conclude that the candidate has suffered in the past from discrimination.

3.7.9 Period 2

Because taste-discrimination does not affect group A members, and this is also known by second period employers, the results concerning their employment outcomes are identical to those in the full information model. The period one actions and employment outcomes are identical in the partial and full information settings.

Now begin by considering a $g = B$ group member who was employed in the first period, and is not discriminated against in the second period.

Expected profits are:

$$\mathbb{E}(\pi_2 | emp, \theta_2, g = B, T_2 = 0) = \left[\frac{\mu_2(\mathcal{O}_1 = emp, g = B, \theta_2)}{1 + \frac{(1 - F_L(\tilde{\theta}_1^d)) f_L(\theta_2)(1 - p_H)}{(1 - F_H(\tilde{\theta}_1^d)) f_H(\theta_2) p_H}} + \frac{1 - \mu_2(\mathcal{O}_1 = emp, g = B, \theta_2)}{1 + \frac{(1 - F_L(\tilde{\theta}_1)) f_L(\theta_2)(1 - p_H)}{(1 - F_H(\tilde{\theta}_1)) f_H(\theta_2) p_H}} \right] v - w$$

Thus the employment threshold, defined by $\tilde{\theta}_{2,emp}^{partial}$, is in the interval $(\tilde{\theta}_{2,emp}^d, \tilde{\theta}_{2,emp})$

If the candidate is discriminated against in the second period, then the expected profits simply take the form:

$$\mathbb{E}(\pi_2|emp,\theta_2,g=B,T_2=1)=\left[\frac{\mu_2(\mathcal{O}_1=emp,g=B,\theta_2)}{1+\frac{(1-F_L(\tilde{\theta}_1^d))f_L(\theta_2)(1-p_H)}{(1-F_H(\tilde{\theta}_1^d))f_H(\theta_2)p_H}}+\frac{(1-\mu_2(\mathcal{O}_1=emp,g=B,\theta_2))}{1+\frac{(1-F_L(\tilde{\theta}_1))f_L(\theta_2)(1-p_H)}{(1-F_H(\tilde{\theta}_1))f_H(\theta_2)p_H}}\right]v-(w+d)$$

Define the threshold in this case as $\tilde{\theta}_{2,emp}^{partial,d}\in(\tilde{\theta}_{2,emp}^{d,d},\tilde{\theta}_{2,emp}^{\emptyset,d})$

Now consider a $g=B$ group member who was not employed in the first period, and is not discriminated against in the second period.

Expected profits are:

$$\mathbb{E}(\pi_2|\neg emp,\theta_2,g=B,T_2=0)=\left[\frac{\mu_2(\mathcal{O}_1=\neg emp,g=B,\theta_2)}{1+\frac{F_L(\tilde{\theta}_1^d)f_L(\theta_2)(1-p_H)}{F_H(\tilde{\theta}_1^d)f_H(\theta_2)p_H}}+\frac{(1-\mu_2(\mathcal{O}_1=\neg emp,g=B,\theta_2))}{1+\frac{F_L(\tilde{\theta}_1)f_L(\theta_2)(1-p_H)}{F_H(\tilde{\theta}_1)f_H(\theta_2)p_H}}\right]v-w$$

The employment threshold condition $\tilde{\theta}_{2,\neg emp}^{partial}$ is in the interval $(\tilde{\theta}_{2,\neg emp}^{d},\tilde{\theta}_{2,\neg emp})$

Finally consider a $g=B$ group member who was not employed in the first period, and is discriminated against in the second period.

Expected profits are:

$$\mathbb{E}(\pi_2|\neg emp,\theta_2,g=B,T_2=1)=\left[\frac{\mu_2(\mathcal{O}_1=\neg emp,g=B,\theta_2)}{1+\frac{F_L(\tilde{\theta}_1^d)f_L(\theta_2)(1-p_H)}{F_H(\tilde{\theta}_1^d)f_H(\theta_2)p_H}}+\frac{(1-\mu_2(\mathcal{O}_1=\neg emp,g=B,\theta_2))}{1+\frac{F_L(\tilde{\theta}_1)f_L(\theta_2)(1-p_H)}{F_H(\tilde{\theta}_1)f_H(\theta_2)p_H}}\right]v-(w+d)$$

Define the cutoff threshold as $\tilde{\theta}_{2,\neg emp}^{partial,d}\in(\tilde{\theta}_{2,\neg emp}^{d,d},\tilde{\theta}_{2,\neg emp}^{\emptyset,d})$

Ex-ante Employment

Having defined the second period cutoff strategies which dictate the employment decision, ex-ante probabilities can be calculated based on the probabilities of meeting these cutoffs.

For those in the A group the probability remains the same:

$$\begin{aligned}\mathsf{Pr}(emp_2|g=A)=&\Big((1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1))(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\Big)\\&+\Big(F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big)\end{aligned}$$

For the B group it is, in the Strong Dis-Taste case:

$$\begin{aligned}\mathsf{Pr}(emp_2|g=B) =&(1-\alpha)^2\Big[(1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}^{partial}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}^{partial}_{2,\neg emp}))(1-p_H)\Big]\\&+\alpha(1-\alpha)\Big(1-F_H(\tilde{\theta}^{partial}_{2,\neg emp}))p_H+(1-F_L(\tilde{\theta}^{partial}_{2,\neg emp}))(1-p_H)\Big]\end{aligned}$$

In the general case:

$$\begin{aligned}\mathsf{Pr}(emp_2|g=B) =&(1-\alpha)^2\Big[(1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}^{partial}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}^{partial}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}^{partial}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}^{partial}_{2,\neg emp}))(1-p_H)\Big]\\&+(1-\alpha)\alpha\Big[(1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}^{partial,d}_{2,emp}))p_H+(1-F_L(\tilde{\theta}_1))(1-F_L(\tilde{\theta}^{partial,d}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}^{partial,d}_{2,\neg emp}))p_H+F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}^{partial,d}_{2,\neg emp}))(1-p_H)\Big]\\&+\alpha(1-\alpha)\Big[(1-F_H(\tilde{\theta}^{d}_{1}))(1-F_H(\tilde{\theta}^{partial}_{2,emp}))p_H+(1-F_L(\tilde{\theta}^{d}_{1}))(1-F_L(\tilde{\theta}^{partial}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}^{d}_{1})(1-F_H(\tilde{\theta}^{partial}_{2,\neg emp}))p_H+F_L(\tilde{\theta}^{d}_{1})(1-F_L(\tilde{\theta}^{partial}_{2,\neg emp}))(1-p_H)\Big]\\&+\alpha^2\Big[(1-F_H(\tilde{\theta}^{d}_{1}))(1-F_H(\tilde{\theta}^{partial,d}_{2,emp}))p_H+(1-F_L(\tilde{\theta}^{d}_{1}))(1-F_L(\tilde{\theta}^{partial,d}_{2,emp}))(1-p_H)\\&+F_H(\tilde{\theta}^{d}_{1})(1-F_H(\tilde{\theta}^{partial,d}_{2,\neg emp}))p_H+F_L(\tilde{\theta}^{d}_{1})(1-F_L(\tilde{\theta}^{partial,d}_{2,\neg emp}))(1-p_H)\Big]\end{aligned}$$

3.7.10 Implications

Proposition 3.7.8. *In the special case of strong dis-taste, the discrimination ratio decreases between period 1 and 2 in the partial information setting. Mathematically* $\frac{\mathsf{Pr}(emp_2|g=A)}{\mathsf{Pr}(emp_2|g=B)} < \frac{\mathsf{Pr}(emp_1|g=A)}{\mathsf{Pr}(emp_1|g=B)}$

The proof is presented in the Appendix to Chapter 3.

The logic is the same as in the simple model, where there is a benefit of the doubt effect benefitting those who were evaluated fairly and unemployed due to having a weak first period signal realisation, but were treated with extra lenience due to the possibility of having been discriminated against.

Recall that under the full information setting, it was ambiguous whether or not the difference between groups over time would increase or decrease. This is due to the signal structure no longer being built with purposeful symmetry. Using it now as a frame of reference, it is ambiguous whether or not members of the B group become relatively better off over time in the partial information structure. This applies to both the strong dis-taste special case and the general case. The reason is that those who were not discriminated against previously are better off under partial information. Even though they were evaluated fairly in period 1, they are given a benefit of the doubt when they were unemployed which means they are evaluated with more lenience in the second period. Likewise, if employed in the first period,

there is a chance that they had to overcome a more demanding hiring criteria, making them appear potentially more impressive. But the opposite case holds when they were previously discriminated against. It is more impressive to have overcome adversity when you are *known* to have been discriminated against compared to when you *might* have been discriminated against but also you might have been evaluated fairly, and analogously it is less damaging to be unemployed if you are known to have been discriminated against compared to it only being a possibility.

Thus, members of the B group in period 2 are not necessarily better off in the partial information case compared with the full information setting.

The decrease of the discrimination ratio over time was a key result in the simple model under the partial information setting, and that it does not survive under the more flexible full model demonstrates the value of the more general model with less stringent assumptions. In particular, the coarse nature of the signals and the imposed symmetry in the simple model limit the theoretical possibilities of the analysis.

3.8 Conclusion

The long-term effects of joblessness have been documented (Arulampalam, 2001; Gregg and Tominey, 2005) and strongly suggest that unemployment does have a 'scarring' effect. People who have been unsuccessful in the labour market end up with a significantly poorer performance in future years than their counterparts who have managed to avoid unemployment. I shed light on how observational learning can exacerbate discrimination that may have affected individuals early on in their careers through a signal distortion effect. It is also possible, however, that group differences in labour market outcomes may be reduced over time as members of discriminated groups gain from a benefit-of-the-doubt effect, if past outcomes are evaluated more leniently by accounting for discrimination which may have occurred. I have analysed how these effects arise under different assumptions on the structure of information and beliefs and shown that these are crucial in determining the evolution of group differences over time.

When belief formation is naive, and employers assume that the previous employer a candidate has interacted with did not taste-discriminate, group differences can become exacerbated over time through a signal distortion effect, where unemployment caused by taste discrimination is interpreted as a signal of low ability. The result is a compounding effect of initial discrimination. On the other hand, when information is imperfect but employers know that taste-discrimination exists and update their beliefs as rational Bayesians, group differences in the labour market can be dampened over time. This happens as members of the discriminated against group who have previously been rejected are given a benefit of the doubt and treated with extra lenience, even if their initial rejection has not been caused by discrimination. Under some circumstances, such as a sufficiently small level of taste-discrimination in the economy, the benefit-of-the-doubt effect causes a reversal in the employment outcome between the groups, where discriminated against applicants are eventually more likely to find employment than their counterparts in the majority group.

The outcome of interest in my model, employment rates between groups over time, is readily available in standard labour market data and I have compared my theoretical predictions with the stylised facts of the labour market. The data aligns with the partial information framework studied, in which employers are rational Bayesians and the group difference in employment rates decreases over time. However, the theoretical model, even with relaxed assumptions, is a simplification of reality and the empirical section is meant only for suggestive, rather than conclusive, evidence. There are also other theoretical mechanisms which could effect the evolution of group differences over time outside of my framework.

The simple model I present delivers sharp predictions with clear intuition, but when a more general model is presented, with minimal assumptions, the predictions become less sharp, although the general intuition remains. In particular, discriminated against groups are worst off under naive beliefs, since past discrimination, when unaccounted for, is exacerbated. However, comparing the cases where jobseekers are known to be discriminated against with the possibility that they might have been discriminated against, there are competing effects. Those who were not treated with prejudice benefit from the extra lenience that comes with potentially having been discriminated against. However, those who were treated with prejudice are better off if it is known that they are treated with prejudice and given lenience, than if it is uncertain how they have been treated and the level of lenience is lower.

3.8.1 Policy Implications

Because some people may suffer not just directly, but also indirectly and permanently from discrimination, this research has significant policy implications. But because the model predictions are sharply different under different structures of information and belief formation, it is not obvious what policy reforms to suggest. What is clear, however, it that the full information benchmark is not a very realistic depiction of how the labour market operates in practice. In addition, recall that one interpretation of the benefit-of-the-doubt effect was as an *internalisation* of an affirmative action policy which arises naturally in the model with partial information and rational Bayesian profit maximising firms. This suggests that if the labour market indeed operates in this way, then introducing additional distortions through policy might be unnecessary, since the market implements its own corrective action.

Since the naive beliefs model has been motivated using the realistic assumption that decision makers in charge of hiring ignore the possibility that a candidate has a weak labour market history due to discrimination[22], even though group membership is observable, effective policy should seek to increase awareness to those with mis-specified beliefs. One way in which awareness about other firms' behaviour might be achieved is through the public reporting of pay gaps which has, since 2017, been mandated for all firms in the UK with over 250 employees[23]

[22] Andre (2024) finds experimental evidence in favour of this worldview, in which meritocracy is 'shallow' with respect to the different circumstances of individuals.

[23] https://www.gov.uk/government/collections/gender-pay-gap-reporting. Blundell (2020) reports that this policy has been successful at reducing the gender pay gap at the affected firms.

Other potential policies which could help to mitigate signal-distortion effects which arise under naive beliefs are traditional bias awareness campaigns, for which Boring and Philippe (2021) found evidence of effectiveness. In general, the policies which are most likely to be of use in correcting the labour market distortions are those which counteract employer naivety.

Appendix 3

C.1 Simple Model - Proofs

Proof of Lemma 3.3.1:

The proof is direct and follows from checking cases where employers would hire or reject the job applicant, and cases of indifference.

At $t = 1$, if the signal is low $\mathbb{E}(\pi_{t=1}|x_{t=1})$ is strictly negative, leading to rejection, and if the signal is high the expected profit is positive, leading to employment. The ex-ante probability of employment is thus $\mathsf{Pr}(emp_{t=1}|g = A) = \mathsf{Pr}(x_{t=1} = H) = \frac{1}{2}$.

Since those in the majority group are never discriminated against, the signal $x_{t=1}$ is exactly inferred from $\mathcal{O}_{t=1} \in \{emp, \neg emp\}$.

At $t = 2$, check $\mathbb{E}(\pi_{t=1}|X_{t=2})$ for each of the four possible realisations of $X_{t=2} = \{\mathcal{O}_{t=1}, x_{t=2}\}$. By averaging across the possible ex-ante signal realisations, calculate the probability of employment as $\mathsf{Pr}(emp_{t=2}|g = A) = \frac{1}{2}$.

Thus, $\mathsf{Pr}(emp_t|g = A) = \frac{1}{2} \forall t \in \{1, 2\}$ □

Proof of Proposition 3.3.2:

Recall from Lemma 3.3.1 that $\mathsf{Pr}(emp_t|g = A) = \frac{1}{2} \forall t \in \{1, 2\}$.

Now consider a candidate in the $g = B$ group.

At $t = 1$, the applicant is rejected if matched with a taste-discriminator. Otherwise, by checking $\mathbb{E}(\pi_{t=1}|x_{t=1})$, she is employed so long as $x_{t=1} = H$, thus $\mathsf{Pr}(emp_{t=1}|g = B) = \frac{1-\alpha}{2}$.

Now consider an applicant who *was not* discriminated against at $t = 1$. At $t = 2$ the applicant is rejected if matched with a taste-discriminator. Otherwise, check $\mathbb{E}(\pi_{t=2}|X_{t=2})$ for each of the four possible signal realisations and note that they lead to employment with probability $\frac{1}{2}$.

Consider also an applicant who *was* discriminated against at $t = 1$. At $t = 2$ the applicant is rejected if matched with a taste-discriminator. Otherwise, because the employer knows that the applicant was previously rejected due to discrimination and $\mathcal{O}_{t=1} = \neg emp$ regardless of the true signal $x_{t=1}$, check $\mathbb{E}(\pi_{t=1}|x_{t=2})$ and note that this leads to employment with probability $\frac{1}{2}$.

This means that $\mathsf{Pr}(emp_{t=2}|g = B) = \frac{1-\alpha}{2}$.

Applying the definition of the discrimination ratio yields $Disc_t = \frac{1}{1-\alpha} \forall t \in \{1, 2\}$. □

Proof of Proposition 3.4.3:

Note from Lemma 3.3.1 members of the majority group are employed with probability $\frac{1}{2} \forall t \in \{1, 2\}$ and from Lemma 3.4.2 the discrimination ratio in period 1 is $Disc_{t=1} = \frac{1}{1-\alpha}$.

At $t = 2$, a discriminated against group member who was previously discriminated against is again rejected if she matches with another taste-discriminating employer. Otherwise, since the past rejection is naively interpreted as a signal of low ability, check the cases of $\mathbb{E}(\pi_{t=1}|\mathcal{O}_{t=1} = \neg emp, x_{t=2})$ and observe that employment occurs with probability $\frac{1}{2}\mathsf{Pr}(x_{t=2} = H) = \frac{1}{4}$.

At $t = 2$, a discriminated against group member who was *not* previously discriminated against is rejected if she matches with a taste-discriminating employer. Otherwise, check the four possible cases of $\mathbb{E}(\pi_{t=1}|\mathcal{O}_{t=1}, x_{t=2})$ and observe that employment occurs with probability $\frac{1}{2}$.

Now average over the ex-ante matching possibilities, which are outlined in Table 3.2, to find $\mathsf{Pr}(emp_{t=2}|g = B) = \frac{(1-\alpha)(2-\alpha)}{4}$.

Apply the definition of the discrimination ratio to observe that $Disc_{t=2} = \frac{2}{(1-\alpha)(2-\alpha)} \left(> Disc_{t=1} = \frac{1}{1-\alpha} \forall \alpha \in (0,1) \right)$ □

Proof of Lemma 3.5.2:

The proof follows the application of (the conditional version of) Bayes' rule.

For a discriminated group member previously unemployed the posterior belief follows:

$$\mu_{t=2}(T_{t=1} = 1|g = B, \mathcal{O}_{t=1} = \neg emp) = \frac{\mathsf{Pr}(\neg emp_{t=1}|T_{t=1} = 1, g = B)\mathsf{Pr}(T_{t=1} = 1|g = B)}{\mathsf{Pr}(\neg emp_{t=1}|g = B)}$$

$$= \frac{2\alpha}{1+\alpha} \quad (> \alpha \ \forall \alpha \in (0,1))$$

For a discriminated group member previously employed the posterior belief follows:

$$\mu_{t=2}(T_{t=1} = 1|g = B, \mathcal{O}_{t=1} = emp) = \frac{\mathsf{Pr}(emp_{t=1}|T_{t=1} = 1, g = B)\mathsf{Pr}(T_{t=1} = 1|g = B)}{\mathsf{Pr}(emp_{t=1}|g = B)} = 0 \quad \square$$

Proof of Proposition 3.5.3:

Recall from Lemma 3.3.1 that members of the majority group are employed with probability $\frac{1}{2} \forall t \in \{1, 2\}$ and recall from Lemma 3.5.1 the discrimination ratio in period 1 is $Disc_{t=1} = \frac{1}{1-\alpha}$.

At $t = 2$ a member of the discriminated against group who has been discriminated against in the first period is again unemployed if matched with a taste-discriminator. Otherwise, noting that the first period outcome is unemployment with certainty, check $\mathbb{E}(\pi_{t=1}|X_{t=2} = \{\neg emp, x_{t=2}\})$ for both possible signal realisations at $t = 2$ and observe that the employment probability is $\frac{1}{2}$.

Again, at $t = 2$ a member of the discriminated against group who has *not* been discriminated against in the first period is unemployed if matched with a taste-discriminator. Otherwise, check $\mathbb{E}(\pi_{t=1}|X_{t=2} = \{\mathcal{O}_{t=1}, x_{t=2}\})$ for each of the four possible realisations of $X_{t=2} = \{\mathcal{O}_{t=1}, x_{t=2}\}$ and observe that the employment probability is $\frac{1+p-p^2}{2}$.

Average across the possible ex-ante matching histories of employer types, which are summarised in Table 3.3, to observe that $\mathsf{Pr}(emp_{t=2}|g = B) = \frac{1-\alpha}{2}(\alpha + (1-\alpha)(1+p-p^2))$.

Thus, the ex-ante discrimination ratio in period 2 is:
$Disc_{t=2} = \frac{1}{(1-\alpha)(\alpha+(1-\alpha)(1+p-p^2))} \left(< Disc_{t=1} = \frac{1}{1-\alpha}\right)$ □

Proof of Remark 3.5.4:

In the first period, it is unambiguous that $\mathsf{Pr}(emp_{t=1}|g=B) < \mathsf{Pr}(emp_{t=1}|g=A)$. A reversal thus requires that $\mathsf{Pr}(emp_{t=2}|g=B) > \mathsf{Pr}(emp_{t=2}|g=A)$.

Firstly, note the following comparative statics, showing that the second period employment probability of those in the discriminated against group is decreasing (1) with the proportion of taste discriminators in the economy and (2) with the strength of the ability signal.

1. $\frac{\partial \mathsf{Pr}(emp_{t=2}|g=B)}{\partial \alpha} = \frac{1-2\alpha-2(1-\alpha)(1+p-p^2)}{2} < 0$
2. $\frac{\partial \mathsf{Pr}(emp_{t=2}|g=B)}{\partial p} = \frac{1}{2}(1-\alpha)^2(1-2p) < 0$

Now take these parameters to the limit:
$$\lim_{\alpha\to 0, p\to\frac{1}{2}} \underbrace{\frac{(1-\alpha)}{2}(\alpha+(1-\alpha)(1+p-p^2))}_{\mathsf{Pr}(emp_{t=2}|g=B)} = \frac{5}{8}\left(> \mathsf{Pr}(emp_{t=2}|g=A) = \frac{1}{2}\right)$$
Thus, for sufficiently small α and p, $\mathsf{Pr}(emp_{t=2}|g=B) > \mathsf{Pr}(emp_{t=2}|g=A)$ □

C.2 Full Model - Proofs

Proof of Proposition 3.7.1:

$\tilde{\theta}^d_{2,emp} < \tilde{\theta}_{2,emp}$ requires that $h^{-1}\left[\frac{V-w}{w}\frac{p_H}{1-p_H}\frac{1-F_H(\tilde{\theta}^d_1)}{1-F_L(\tilde{\theta}^d_1)}\right] < h^{-1}\left[\frac{V-w}{w}\frac{p_H}{1-p_H}\frac{1-F_H(\tilde{\theta}_1)}{1-F_L(\tilde{\theta}_1)}\right]$

This condition holds whenever $\frac{1-F_H(\tilde{\theta}^d_1)}{1-F_L(\tilde{\theta}^d_1)} > \frac{1-F_H(\tilde{\theta}_1)}{1-F_L(\tilde{\theta}_1)}$

Which can be re-written as: $\frac{\int_{\tilde{\theta}^d_1} f_H(\theta)\,d\theta}{\int_{\tilde{\theta}^d_1} f_L(\theta)\,d\theta} > \frac{\int_{\tilde{\theta}_1} f_H(\theta)\,d\theta}{\int_{\tilde{\theta}_1} f_L(\theta)\,d\theta}$

Recall that $\tilde{\theta}^d_1 > \tilde{\theta}_1$ and split up the integrals on the RHS

$$\frac{\int_{\tilde{\theta}^d_1} f_H(\theta)\,d\theta}{\int_{\tilde{\theta}^d_1} f_L(\theta)\,d\theta} > \frac{\int_{\tilde{\theta}_1}^{\tilde{\theta}^d_1} f_H(\theta)\,d\theta + \int_{\tilde{\theta}^d_1} f_H(\theta)\,d\theta}{\int_{\tilde{\theta}_1}^{\tilde{\theta}^d_1} f_L(\theta)\,d\theta + \int_{\tilde{\theta}^d_1} f_L(\theta)\,d\theta}$$

$$\iff \frac{\int_{\tilde{\theta}^d_1} f_H(\theta)\,d\theta}{\int_{\tilde{\theta}^d_1} f_L(\theta)\,d\theta} > \frac{\int_{\tilde{\theta}_1}^{\tilde{\theta}^d_1} f_H(\theta)\,d\theta}{\int_{\tilde{\theta}_1}^{\tilde{\theta}^d_1} f_L(\theta)\,d\theta}$$

MLRP verifies the inequality. □

Proof of Proposition 3.7.2:

$\tilde{\theta}^d_{2,\neg emp} < \tilde{\theta}_{2,\neg emp}$ requires that $h^{-1}\left[\frac{V-w}{w}\frac{p_H}{1-p_H}\frac{F_H(\tilde{\theta}^d_1)}{F_L(\tilde{\theta}^d_1)}\right] < h^{-1}\left[\frac{V-w}{w}\frac{p_H}{1-p_H}\frac{F_H(\tilde{\theta}_1)}{F_L(\tilde{\theta}_1)}\right]$

This inequality holds when $\frac{F_H(\tilde{\theta}^d_1)}{F_L(\tilde{\theta}^d_1)} > \frac{F_H(\tilde{\theta}_1)}{F_L(\tilde{\theta}_1)}$

Recall that $\tilde{\theta}^d_1 > \tilde{\theta}_1$ and write in integral form

$$\frac{\int^{\tilde{\theta}_1} f_H(\theta)\,d\theta + \int_{\tilde{\theta}_1}^{\tilde{\theta}^d_1} f_H(\theta)\,d\theta}{\int^{\tilde{\theta}_1} f_L(\theta)\,d\theta + \int_{\tilde{\theta}_1}^{\tilde{\theta}^d_1} f_L(\theta)\,d\theta} > \frac{\int^{\tilde{\theta}_1} f_H(\theta)\,d\theta}{\int^{\tilde{\theta}_1} f_L(\theta)\,d\theta}$$

$$\iff \frac{\int_{\tilde{\theta}_1}^{\tilde{\theta}_1^d} f_H(\theta)\,d\theta}{\int_{\tilde{\theta}_1}^{\tilde{\theta}_1^d} f_L(\theta)\,d\theta} > \frac{\int^{\tilde{\theta}_1} f_H(\theta)\,d\theta}{\int^{\tilde{\theta}_1} f_L(\theta)\,d\theta}$$

MLRP verifies the inequality. □

Proof of Proposition 3.7.3:

Verify that $\frac{\mathsf{Pr}(emp|g=A)}{\mathsf{Pr}(emp|g=B)} > \frac{1}{1-\alpha}$, which can be verified in the following simplified expression.

$$\begin{aligned}
&\Big[\Big((1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H + (1-F_L(\tilde{\theta}_1))(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\Big)\\
&+ \Big(F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H + F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big)\Big]\\
>\\
&(1-\alpha)\Big[\Big((1-F_H(\tilde{\theta}_1))(1-F_H(\tilde{\theta}_{2,emp}))p_H + (1-F_L(\tilde{\theta}_1))(1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)\Big)\\
&+ \Big(F_H(\tilde{\theta}_1)(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H + F_L(\tilde{\theta}_1)(1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big)\Big]\\
&\alpha\Big[\Big((1-F_H(\tilde{\theta}_{2,\neg emp}))p_H + (1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)\Big)\Big]
\end{aligned}$$

□

Proof of Lemma 3.7.6:

$$\mu_2(\mathcal{O}_1 = emp, g = B, \theta_2) = \frac{\left[(1-F_H(\tilde{\theta}_1^d))f_H(\theta_2)p_H + (1-F_L(\tilde{\theta}_1^d))f_L(\theta_2)(1-p_H)\right]\alpha}{\left[(1-F_H(\tilde{\theta}_1^d))f_H(\theta_2)p_H + (1-F_L(\tilde{\theta}_1^d))f_L(\theta_2)(1-p_H)\right]\alpha + \left[(1-F_H(\tilde{\theta}_1))f_H(\theta_2)p_H + (1-F_L(\tilde{\theta}_1))f_L(\theta_2)(1-p_H)\right](1-\alpha)}$$

Divide the top and bottom by $f_L(\theta_2)$

$$\mu_2(\mathcal{O}_1 = emp, g = B, \theta_2) = \frac{\left[(1-F_H(\tilde{\theta}_1^d))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1^d))(1-p_H)\right]\alpha}{\left[(1-F_H(\tilde{\theta}_1^d))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1^d))(1-p_H)\right]\alpha + \left[(1-F_H(\tilde{\theta}_1))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1))(1-p_H)\right](1-\alpha)}$$

$$\begin{aligned}
\frac{\partial \mu_2(\mathcal{O}_1 = emp, g = B, \theta_2)}{\partial \frac{f_H(\theta_2)}{f_L(\theta_2)}} =& \frac{(1-F_H(\tilde{\theta}_1^d))p_H\alpha\left(\left[(1-F_H(\tilde{\theta}_1^d))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1^d))(1-p_H)\right]\alpha + \left[(1-F_H(\tilde{\theta}_1))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1))(1-p_H)\right](1-\alpha)\right.}{\left(\left[(1-F_H(\tilde{\theta}_1^d))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1^d))(1-p_H)\right]\alpha + \left[(1-F_H(\tilde{\theta}_1))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1))(1-p_H)\right](1-\alpha)\right)^2}\\
&- \frac{\left[(1-F_H(\tilde{\theta}_1^d))\alpha p_H + (1-F_H(\tilde{\theta}_1))(1-\alpha)p_H\right]\left[(1-F_H(\tilde{\theta}_1^d))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1^d))(1-p_H)\right]\alpha}{\left(\left[(1-F_H(\tilde{\theta}_1^d))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1^d))(1-p_H)\right]\alpha + \left[(1-F_H(\tilde{\theta}_1))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1-F_L(\tilde{\theta}_1))(1-p_H)\right](1-\alpha)\right)^2}
\end{aligned}$$

$$\frac{\partial \mu_2(\mathcal{O}_1 = emp, g = B, \theta_2)}{\partial \frac{f_H(\theta_2)}{f_L(\theta_2)}} > 0$$

$$\iff$$

$$(1 - F_H(\tilde{\theta}_1^d))\alpha p_H(\left[(1 - F_H(\tilde{\theta}_1^d))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - F_L(\tilde{\theta}_1^d))(1 - p_H)\right]\alpha$$
$$+ \left[(1 - F_H(\tilde{\theta}_1))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - F_L(\tilde{\theta}_1))(1 - p_H)\right](1 - \alpha))$$
$$>$$
$$[(1 - F_H(\tilde{\theta}_1^d))\alpha p_H + (1 - F_H(\tilde{\theta}_1))(1 - \alpha)p_H]\left[(1 - F_H(\tilde{\theta}_1^d))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - F_L(\tilde{\theta}_1^d))(1 - p_H)\right]\alpha$$

$$\iff$$

$$(1 - F_H(\tilde{\theta}_1^d))p_H\left[(1 - F_H(\tilde{\theta}_1))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - F_L(\tilde{\theta}_1))(1 - p_H)\right](1 - \alpha)$$
$$>$$
$$(1 - F_H(\tilde{\theta}_1))(1 - \alpha)p_H\left[(1 - F_H(\tilde{\theta}_1^d))\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - F_L(\tilde{\theta}_1^d))(1 - p_H)\right]$$

$$\iff$$

$$(1 - F_H(\tilde{\theta}_1^d))p_H\left[(1 - F_L(\tilde{\theta}_1))(1 - p_H)\right](1 - \alpha) > (1 - F_H(\tilde{\theta}_1))(1 - \alpha)p_H\left[(1 - F_L(\tilde{\theta}_1^d))(1 - p_H)\right]$$

$$\iff$$

$$(1 - F_H(\tilde{\theta}_1^d))(1 - F_L(\tilde{\theta}_1)) > (1 - F_H(\tilde{\theta}_1))(1 - F_L(\tilde{\theta}_1^d)) \iff \frac{1 - F_H(\tilde{\theta}_1^d)}{1 - F_L(\tilde{\theta}_1^d)} > \frac{1 - F_H(\tilde{\theta}_1)}{1 - F_L(\tilde{\theta}_1)}$$

Which is verified by MLRP. □

Proof of Lemma 3.7.7:

To observe how the belief varies with the current signal θ_2, divide the top and bottom of the fraction by $f_L(\theta_2)$ and take the derivative wrt the likelihood ratio. Start with the special case of strong dis-taste:

$$\mu_2(\mathcal{O}_1 = \neg emp, g = B, \theta_2) = \frac{[\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - p_H)]\alpha}{\left[\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - p_H)\right]\alpha + \left[\frac{f_H(\theta_2)}{f_L(\theta_2)}F_H(\tilde{\theta}_1)p_H + F_L(\tilde{\theta}_1)(1 - p_H)\right](1 - \alpha)}$$

$$\frac{\partial \mu_2(\mathcal{O}_1 = \neg emp, g = B, \theta_2)}{\partial \frac{f_H(\theta_2)}{f_L(\theta_2)}} = \frac{p_H\alpha\left(\left[\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - p_H)\right]\alpha + \left[\frac{f_H(\theta_2)}{f_L(\theta_2)}F_H(\tilde{\theta}_1)p_H + F_L(\tilde{\theta}_1)(1 - p_H)\right](1 - \alpha)\right)}{\left(\left[\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - p_H)\right]\alpha + \left[\frac{f_H(\theta_2)}{f_L(\theta_2)}F_H(\tilde{\theta}_1)p_H + F_L(\tilde{\theta}_1)(1 - p_H)\right](1 - \alpha)\right)^2}$$
$$- \frac{[p_H\alpha + F_H(\tilde{\theta}_1)p_H(1 - \alpha)][\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - p_H)]\alpha}{\left(\left[\frac{f_H(\theta_2)}{f_L(\theta_2)}p_H + (1 - p_H)\right]\alpha + \left[\frac{f_H(\theta_2)}{f_L(\theta_2)}F_H(\tilde{\theta}_1)p_H + F_L(\tilde{\theta}_1)(1 - p_H)\right](1 - \alpha)\right)^2}$$

$$\frac{\partial \mu_2(\mathcal{O}_1 = \neg emp, g = B, \theta_2)}{\partial \frac{f_H(\theta_2)}{f_L(\theta_2)}} > 0$$

$$\iff$$

$$p_H \left(\left[\frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + (1 - p_H) \right] \alpha + \left[\frac{f_H(\theta_2)}{f_L(\theta_2)} F_H(\tilde{\theta}_1) p_H + F_L(\tilde{\theta}_1)(1 - p_H) \right] (1 - \alpha) \right)$$

$$>$$

$$[p_H \alpha + F_H(\tilde{\theta}_1) p_H (1 - \alpha)][\frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + (1 - p_H)]$$

$$\iff$$

$$p_H \alpha \left[\frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + (1 - p_H) \right] + p_H (1 - \alpha) \left[\frac{f_H(\theta_2)}{f_L(\theta_2)} F_H(\tilde{\theta}_1) p_H + F_L(\tilde{\theta}_1)(1 - p_H) \right]$$

$$>$$

$$p_H \alpha [\frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + (1 - p_H)] + F_H(\tilde{\theta}_1) p_H (1 - \alpha) [\frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + (1 - p_H)]$$

$$\iff$$

$$F_L(\tilde{\theta}_1) p_H (1 - \alpha)(1 - p_H) > F_H(\tilde{\theta}_1) p_H (1 - \alpha)(1 - p_H) \iff F_L(\tilde{\theta}_1) > F_H(\tilde{\theta}_1)$$

Which is verified by FOSD.

In the general case then:

$$\mu_2(\mathcal{O}_1 = \neg emp, g = B, \theta_2) = \frac{\left[F_H(\tilde{\theta}_1^d) f_H(\theta_2) p_H + F_L(\tilde{\theta}_1^d) f_L(\theta_2)(1 - p_H) \right] \alpha}{\left[F_H(\tilde{\theta}_1^d) f_H(\theta_2) p_H + F_L(\tilde{\theta}_1^d) f_L(\theta_2)(1 - p_H) \right] \alpha + \left[F_H(\tilde{\theta}_1) f_H(\theta_2) p_H + F_L(\tilde{\theta}_1) f_L(\theta_2)(1 - p_H) \right] (1 - \alpha)}$$

Dividing the top and bottom of the fraction by $f_L(\theta_2)$

$$\mu_2(\mathcal{O}_1 = \neg emp, g = B, \theta_2) = \frac{\left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1 - p_H) \right] \alpha}{\left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1 - p_H) \right] \alpha + \left[F_H(\tilde{\theta}_1) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1)(1 - p_H) \right] (1 - \alpha)}$$

$$\frac{\partial \mu_2(\mathcal{O}_1 = \neg emp, g = B, \theta_2)}{\partial \frac{f_H(\theta_2)}{f_L(\theta_2)}} =$$

$$\frac{F_H(\tilde{\theta}_1^d) p_H \alpha \left(\left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1 - p_H) \right] \alpha + \left[F_H(\tilde{\theta}_1) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1)(1 - p_H) \right] (1 - \alpha) \right)}{\left(\left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1 - p_H) \right] \alpha + \left[F_H(\tilde{\theta}_1) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1)(1 - p_H) \right] (1 - \alpha) \right)^2}$$

$$- \frac{[F_H(\tilde{\theta}_1^d) p_H \alpha + F_H(\tilde{\theta}_1) p_H (1 - \alpha)] \left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1 - p_H) \right] \alpha}{\left(\left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1 - p_H) \right] \alpha + \left[F_H(\tilde{\theta}_1) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1)(1 - p_H) \right] (1 - \alpha) \right)^2}$$

$$\frac{\partial \mu_2(\mathcal{O}_1 = \neg emp, g = B, \theta_2)}{\partial \frac{f_H(\theta_2)}{f_L(\theta_2)}} > 0$$

$$\iff$$

$$F_H(\tilde{\theta}_1^d) p_H \alpha \left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1-p_H) \right] + F_H(\tilde{\theta}_1^d) p_H (1-\alpha) \left[F_H(\tilde{\theta}_1) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1)(1-p_H) \right]$$

$$>$$

$$F_H(\tilde{\theta}_1^d) p_H \alpha \left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1-p_H) \right] + F_H(\tilde{\theta}_1) p_H (1-\alpha) \left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1-p_H) \right]$$

$$\iff$$

$$F_H(\tilde{\theta}_1^d) \left[F_H(\tilde{\theta}_1) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1)(1-p_H) \right] > F_H(\tilde{\theta}_1) \left[F_H(\tilde{\theta}_1^d) \frac{f_H(\theta_2)}{f_L(\theta_2)} p_H + F_L(\tilde{\theta}_1^d)(1-p_H) \right]$$

$$\iff$$

$$F_H(\tilde{\theta}_1^d) F_L(\tilde{\theta}_1) > F_H(\tilde{\theta}_1) F_L(\tilde{\theta}_1^d) \iff \frac{F_H(\tilde{\theta}_1^d)}{F_L(\tilde{\theta}_1^d)} > \frac{F_H(\tilde{\theta}_1)}{F_L(\tilde{\theta}_1)}$$

This expression is verified by MLRP. □

Proof of Proposition 3.7.8:

Evaluate the discrimination ratio in periods 1 and 2 first following the case of unemployment and then the case of employment, to simplify the algebra. The unconditional discrimination ratio is a weighted-average of these two expressions.

Following unemployment we have:

$$\frac{\mathsf{Pr}(emp_2|g=A,\neg emp)}{\mathsf{Pr}(emp_2|g=B,\neg emp)} < \frac{\mathsf{Pr}(emp_1|g=A)}{\mathsf{Pr}(emp_1|g=B)}$$

$$\frac{(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H + (1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H)}{(1-\alpha)[(1-F_H(\tilde{\theta}^{partial}_{2,\neg emp}))p_H + (1-F_L(\tilde{\theta}^{partial}_{2,\neg emp}))(1-p_H)]} < \frac{(1-F_H(\tilde{\theta}_1))p_H + (1-F_L(\tilde{\theta}_1))(1-p_H)}{(1-\alpha)[(1-F_H(\tilde{\theta}_1))p_H + (1-F_L(\tilde{\theta}_1))(1-p_H)}$$

$$\iff$$

$$(1-F_H(\tilde{\theta}_{2,\neg emp}))p_H + (1-F_L(\tilde{\theta}_{2,\neg emp}))(1-p_H) < (1-F_H(\tilde{\theta}^{partial}_{2,\neg emp}))p_H + (1-F_L(\tilde{\theta}^{partial}_{2,\neg emp}))(1-p_H)$$

The fact that $\tilde{\theta}^{partial}_{2,\neg emp} < \tilde{\theta}_{2,\neg emp}$ verifies that this inequality holds. In either period, those in the B group are unemployed when matching with a taste discriminator. But in the second period, unemployment is viewed with more lenience since members of the B group might have been taste-discriminated in the past.

Following employment:

$$\frac{\mathsf{Pr}(emp_2|g=A,emp)}{\mathsf{Pr}(emp_2|g=B,emp)} = \frac{\mathsf{Pr}(emp_1|g=A)}{\mathsf{Pr}(emp_1|g=B)}$$

$$\frac{(1-F_H(\tilde{\theta}_{2,emp}))p_H + (1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)}{(1-\alpha)[(1-F_H(\tilde{\theta}_{2,emp}))p_H + (1-F_L(\tilde{\theta}_{2,emp}))(1-p_H)]} = \frac{(1-F_H(\tilde{\theta}_1))p_H + (1-F_L(\tilde{\theta}_1))(1-p_H)}{(1-\alpha)[(1-F_H(\tilde{\theta}_1))p_H + (1-F_L(\tilde{\theta}_1))(1-p_H)]}$$

In the first period, members of the B group suffer from taste discrimination, but if they meet a non-discriminating employer, they are treated equivalently to an A group member. In the second period, following employment, it is known that being employed means being

not discriminated against in period one, and hence, conditional again on meeting a non-discriminating employer, members of both groups are treated according to the same standard.

Because the discrimination ratio decreases following unemployment, and remains constant following employment, ex-ante discrimination ratio decreases. □

Conclusion to the Thesis

This thesis has addressed a number of different topics and questions in the field of labour economics. While the questions addressed are varied, all are brought together by the common theme of information and strategic behaviour. To conclude my thesis, I will summarize each of my papers and discuss some of the policy implications of my results. Then, I will discuss the major limitations of my papers and make suggestions for future work.

In the first chapter, I developed a theoretical model of jobseeker self-selection decisions, when job applications require a screening procedure. Putting the theory to the data, I show that the introduction of screening requirements for vacancies leads to a substantial improvement to the quality of the applicant pool while reducing the number of applicants. Benchmarking the induced effects of self-selection into the applicant pool against standard predictors of ability, such as education and experience, I find that the screening procedure was comparable to, if not more effective than, using these predictors. Importantly, requiring screening procedures to be completed in the application process does not differentially deter women or minorities from applying. In practice, the screening procedures should be easy to implement, since details such as the level of difficulty of the screening task or the length of time required to complete the task matter much less for improving outcomes than simply having a task at all compared to not having one.

The implications of this paper are obvious. Employers can improve the quality of, and shrink the size of, the applicant pool for a given vacancy by introducing a screening test requirement in the application process. This is easy and practically costless to implement and could be as simple as a requirement to provide a written answer to a non-generic question, or watch a short video. Although policymakers cannot dictate how firms make their hiring decisions in a free-market, screening tests could be made a requirement for public sector jobs.

In the second chapter, I provide evidence which strongly suggests that income under-reporting is occurring in Romania's construction sector. I do so by examining the income reporting and other behavioural responses to the tax reforms, and concluding that there is no other plausible explanation for the behaviour I observe besides tax evasion. Having provided this evidence, I turn to address gender differences in reporting behaviour and find that men are much more likely than women to exhibit behaviour consistent with income under-reporting. While challenging to dis-entangle the mechanisms, the data points to risk and social preferences as a key driver. In mixed-gender firms, men and women have very similar income reporting responses following the tax reform, but in gender-segregated firms the men are much more likely to be under-reporting their income. This suggests that men may be averse to the risk of the joint collusion to tax evade breaking down and may be less trusting in mixed-gender firms. Of course, this heterogeneity may not be an effect of

being employed in different kinds of firms, but instead may be a selection effect, where more dis-honest men select into gender-segregated firms, including one-man bands. A result of the apparent gender gap in tax evasion is that it distorts the gender wage gap generated by the data compared to the gender wage gap of true take home salaries.

One policy implication of this paper is that extra care must be taken when calculating gender wage gaps using administrative data in settings where tax evasion is likely to be prevalent. An alternative would be to use surveys, with emphasis for respondents that the data is strictly confidential and could not be obtained by the authorities. To address income tax evasion in systems with employer reporting more generally, it can be difficult to induce truthful reporting and tax paying. Enforcement, through the use of audits, is costly and not foolproof when business transactions take place using cash payments. Reducing the prevalence of cash transactions in an economy is also no easy task. Where tax evasion through false declarations of the minimum wage is thought to be commonplace, authorities can increase the minimum wage primarily as a tool of tax enforcement rather than employee protection. As shown in the case of Romania, this can be done at a sectoral level, but could also be done at regional levels and for different age categories to target the groups most likely to be under-reporting their income.

The third chapter develops a theoretical model to understand the dynamic evolution of group differences in employment over time, focussing on the role of the information environment in shaping employers' hiring decisions. When employers have naive beliefs, incorrectly behaving as if there is no discrimination in the labour market, unemployment acts as a negative signal of ability, exacerbating differences in employment rates between discriminated-against and non-discriminated groups. On the other hand, if employers are aware of the presence of discrimination and behave as rational Bayesians, they treat members of the discriminated-against group with extra lenience, since unemployment may have been caused by discrimination. This could be viewed as a rationally developed affirmative action policy by employers, and can result in group differences converging over time, and possibly reversing.

This paper thus suggests that effective public policy can focus on improving awareness of discrimination, for example, through bias-awareness campaigns. If hiring managers are aware of the fact that previous labour market outcomes are directly a function of past discrimination, they should rationally take this into account when hiring for a vacancy.

Each chapter leaves questions unanswered and could be made more robust in future work.

In the first chapter, our hiring experiment takes place on an online labour market and advertises a vacancy for a short and un-skilled task. The conclusions reached may look different in a more realistic setting, when hiring for a longer-term or permanent job and in a more traditional labour market matching platform. Given that screening tasks are widespread in hiring for skilled jobs, if the observational data were available, the effects of these tasks could be analysed in a real-world setting. Because hiring data tend to be highly confidential, a more pragmatic approach for future research may be larger scale hiring experiments, in offline labour markets and for different types of work.

The second chapter used observed behavioural responses to a major set of policy changes to infer that tax evasion was prevalent in the affected sector. The paper thus relies on the

assumption that the observed behaviour can only be credibly explained by income under-reporting. This is, of course, an assumption. Directly observing tax evasion is, by its nature, very difficult. If it was easy, the tax authorities would be able to intervene with the evading individuals and firms. Unfortunately, due to its hidden nature, studying tax evasion will have to rely on non-direct methods for inference, which have been discussed in the literature review of the paper. Work on tax evasion in Romania, and employee-employer collusion and gender gaps in tax evasion specifically, would benefit from complementary studies using different methods and sources of data.

The final chapter relied on a theoretical model to study the evolution of group differences in labour markets over time. Although I began with a simple model and then provided a richer model with much less restrictive assumptions, even the fuller model was not assumption free. The model only featured two periods of discrete time and all matches in period 1 were forced to dissolve in period 2 so that employers and jobseekers would randomly match again. This eliminates the possibility of employers retaining their employees, and improving the learning within the firm about the employee quality. The model also does not feature any employee training or skill formation/dissolution with employment/unemployment. All of these features are likely to matter in real-life labour markets. Empirically, disentangling different mechanisms which drive differences in labour market outcomes between groups is notoriously challenging. While I have provided descriptive data in my paper, clever research design could help to better link the theory with data.

Bibliography

Abebe, Girum, A Stefano Caria, and Esteban Ortiz-Ospina (2021). "The selection of talent: Experimental and structural evidence from ethiopia". In: *American Economic Review* 111.6, pp. 1757–1806.

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond (2019). "Preferences for truth-telling". In: *Econometrica* 87.4, pp. 1115–1153.

Acemoglu, Daron (1995). "Public Policy in a Model of Long-term Unemployment". In: *Economica* 62.246, pp. 161–178.

Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons (2016). "Does standardized information in online markets disproportionately benefit job applicants from less developed countries?" In: *Journal of International Economics* 103, pp. 1–12. ISSN: 0022-1996. DOI: 10.1016/j.jinteco.2016.08.003.

Agrawal, Ajay et al. (2015). "Chapter 8 - Digitization and the Contract Labor Market". In: *Economic Analysis of the Digital Economy*. Chicago, IL: University of Chicago Press. DOI: 10.7208/chicago/9780226206981.003.0008.

Allingham, Michael G and Agnar Sandmo (1972). "Income tax evasion: A theoretical analysis". In: *Journal of public economics* 1.3-4, pp. 323–338.

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman (2019). "Tax evasion and inequality". In: *American Economic Review* 109.6, pp. 2073–2103.

Altonji, Joseph G and Charles R Pierret (2001). "Employer learning and statistical discrimination". In: *The Quarterly Journal of Economics* 116.1, pp. 313–350.

Anderson, Michael L (2008). "Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects". In: *Journal of the American Statistical Association* 103.484, pp. 1481–1495.

Andre, Peter (2024). "Shallow meritocracy". In: *Review of Economic Studies*, rdae040.

Andreoni, James, Brian Erard, and Jonathan Feinstein (1998). "Tax compliance". In: *Journal of economic literature* 36.2, pp. 818–860.

Arnosti, Nick, Ramesh Johari, and Yash Kanoria (2021). "Managing congestion in matching markets". In: *Manufacturing & Service Operations Management* 23.3, pp. 620–636.

Arulampalam, Wiji (2001). "Is Unemployment Really Scarring? Effects of Unemployment Experiences on Wages". In: *The Economic Journal* 111.475, F585–F606.

Autor, David H and David Scarborough (2008). "Does job testing harm minority workers? Evidence from retail establishments". In: *The Quarterly Journal of Economics* 123.1, pp. 219–277.

Avery, Christopher and Peter Zemsky (1998). "Multidimensional uncertainty and herd behavior in financial markets". In: *American economic review*, pp. 724–748.

Backus, Peter et al. (2023). "Gender, competition, and performance: Evidence from chess players". In: *Quantitative Economics* 14.1, pp. 349–380. DOI: https://doi.org/10.3982/QE1404. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/QE1404. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/QE1404.

Baker, Michael et al. (2023). "Pay transparency and the gender gap". In: *American Economic Journal: Applied Economics* 15.2, pp. 157–183.

Banerjee, A. et al. (2009). "Labor market discrimination in Delhi: Evidence from a field experiment." In: *Journal of Comparative Economics* 37.1, pp. 14–27.

Banerjee, Abhijit V (1992). "A simple model of herd behavior". In: *The Quarterly Journal of Economics* 107.3, pp. 797–817.

Barach, Moshe A and John J Horton (2021). "How do employers use compensation history? Evidence from a field experiment". In: *Journal of Labor Economics* 39.1, pp. 193–218.

Beaman, Lori, Niall Keleher, and Jeremy Magruder (2018). "Do job networks disadvantage women? Evidence from a recruitment experiment in Malawi". In: *Journal of Labor Economics* 36.1, pp. 121–157.

Becker, G. (1971). *The Economics of Discrimination*. 2nd ed. Chicago: University of Chicago Press.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen (2014). "Inference on treatment effects after selection among high-dimensional controls". In: *Review of Economic Studies* 81.2, pp. 608–650.

Benjamini, Yoav, Abba M Krieger, and Daniel Yekutieli (2006). "Adaptive linear step-up procedures that control the false discovery rate". In: *Biometrika* 93.3, pp. 491–507.

Bergolo, Marcelo and Guillermo Cruces (2014). "Work and tax evasion incentive effects of social insurance programs: Evidence from an employment-based benefit extension". In: *Journal of Public Economics* 117, pp. 211–228.

Bertrand, Marianne, Dolly Chugh, and Sendhil Mullainathan (2005). "Implicit discrimination". In: *American Economic Review* 95.2, pp. 94–98.

Bertrand, Marianne and Esther Duflo (2017). "Field experiments on discrimination". In: *Handbook of economic field experiments* 1, pp. 309–393.

Bertrand, Marianne, Sendhil Mullainathan, and Eldar Shafir (2004). "A behavioral-economics view of poverty". In: *American Economic Review* 94.2, pp. 419–423.

Besley, Timothy and Stephen Coate (1992). "Workfare versus welfare: Incentive arguments for work requirements in poverty-alleviation programs". In: *The American Economic Review* 82.1, pp. 249–261.

Biasi, Barbara and Heather Sarsons (2022). "Flexible wages, bargaining, and the gender gap". In: *The Quarterly Journal of Economics* 137.1, pp. 215–266.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). "A theory of fads, fashion, custom, and cultural change as informational cascades". In: *Journal of Political Economy* 100.5, pp. 992–1026.

Bíró, Anikó, Dániel Prinz, and László Sándor (2022). "The minimum wage, informal pay, and tax enforcement". In: *Journal of Public Economics* 215, p. 104728.

Bjørkheim, Julie Brun and Odd E Nygård (2024). "Gender Differences in Tax Evasion: Evidence from Norwegian Administrative Data". In.

Bjørneby, Marie, Annette Alstadsæter, and Kjetil Telle (2021). "Limits to third-party reporting: Evidence from a randomized field experiment in Norway". In: *Journal of Public Economics* 203, p. 104512.

Blundell, Jack (2020). "Wage responses to gender pay gap reporting requirements". In: *Available at SSRN 3584259*.

Bohren, J Aislinn, Peter Hull, and Alex Imas (2022). *Systemic discrimination: Theory and measurement*. Tech. rep. National Bureau of Economic Research.

Bohren, J Aislinn, Alex Imas, and Michael Rosenberg (2019). "The dynamics of discrimination: Theory and evidence". In: *American economic review* 109.10, pp. 3395–3436.

Borghans, Lex et al. (2009). "Gender differences in risk aversion and ambiguity aversion". In: *Journal of the European Economic Association* 7.2-3, pp. 649–658.

Boring, Anne and Arnaud Philippe (2021). "Reducing discrimination in the field: Evidence from an awareness raising intervention targeting gender biases in student evaluations of teaching". In: *Journal of Public Economics* 193, p. 104323.

Burks, Stephen V et al. (2015). "The value of hiring through employee referrals". In: *The Quarterly Journal of Economics* 130.2, pp. 805–839.

Bursztyn, Leonardo, Alessandra L González, and David Yanagizawa-Drott (2020). "Misperceived social norms: Women working outside the home in Saudi Arabia". In: *American economic review* 110.10, pp. 2997–3029.

Cappelli, Peter (2023). *Your Approach to Hiring Is All Wrong*. URL: https://hbr.org/2019/05/your-approach-to-hiring-is-all-wrong.

Card, David (1999). "Chapter 30 - The Causal Effect of Education on Earnings". In: ed. by Orley C. Ashenfelter and David Card. Vol. 3. Handbook of Labor Economics. Elsevier, pp. 1801–1863.

Carlsson, M. and D. Rooth (2007). "Evidence of ethnic discrimination in the Swedish labor market using experimental data." In: *Labour Economics* 14.4, pp. 716–729.

Childs, Jason (2012). "Gender differences in lying". In: *Economics Letters* 114.2, pp. 147–149.

Coate, Stephen and Glenn C Loury (1993). "Will affirmative-action policies eliminate negative stereotypes?" In: *The American Economic Review*, pp. 1220–1240.

Codagnone, Cristiano, Federico Biagi, and Fabienne Abadie (2016). "The passions and the interests: Unpacking the'sharing economy'". In: *Institute for Prospective Technological Studies, JRC Science for Policy Report*.

Coles, Peter et al. (2010). "The Job Market for New Economists: A Market Design Perspective". In: *Journal of Economic Perspectives* 24.4, pp. 187–206. DOI: 10.1257/jep.24.4.187.

Croson, Rachel and Uri Gneezy (2009). "Gender differences in preferences". In: *Journal of Economic literature* 47.2, pp. 448–474.

Cullen, Zoe and Ricardo Perez-Truglia (2023). "The old boys' club: Schmoozing and the gender gap". In: *American Economic Review* 113.7, pp. 1703–1740.

Datta, Namita et al. (2023). “Working Without Borders: The Promise and Peril of Online Gig Work”. In.

Di Stasio, V and A Heath (2019). *Are employers in Britain discriminating against ethnic minorities?. Oxford: Centre for Social Investigation.*

Dreber, Anna and Magnus Johannesson (2008). “Gender differences in deception”. In: *Economics Letters* 99.1, pp. 197–199.

D’attoma, John, Clara Volintiru, and Sven Steinmo (2017). “Willing to share? Tax compliance and gender in Europe and America”. In: *Research & Politics* 4.2, p. 2053168017707151.

Elgin, Ceyhun et al. (2021). “Understanding Informality”. In: *CEPR*.

Enke, Benjamin and Florian Zimmermann (Dec. 2017). “Correlation Neglect in Belief Formation”. In: *The Review of Economic Studies* 86.1, pp. 313–332.

Exley, Christine L and Judd B Kessler (2022). “The gender gap in self-promotion”. In: *The Quarterly Journal of Economics* 137.3, pp. 1345–1381.

Exley, Christine L et al. (2022). *Beliefs about gender differences in social preferences*. Tech. rep.

Eyal, Peer et al. (2021). “Data quality of platforms and panels for online behavioral research”. In: *Behavior Research Methods*, pp. 1–20.

Fang, Hanming and Andrea Moro (2011). “Chapter 5 - Theories of Statistical Discrimination and Affirmative Action: A Survey”. In: vol. 1. Handbook of Social Economics. North-Holland, pp. 133 –200.

Fe, Eduardo, David Gill, and Victoria Prowse (2022). “Cognitive skills, strategic sophistication, and life outcomes”. In: *Journal of Political Economy* 130.10, pp. 2643–2704.

Fehr, Ernst, Michael Naef, and Klaus M Schmidt (2006). “Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment”. In: *American Economic Review* 96.5, pp. 1912–1917.

Finkelstein, Amy and Matthew J Notowidigdo (2019). “Take-up and targeting: Experimental evidence from SNAP”. In: *The Quarterly Journal of Economics* 134.3, pp. 1505–1556.

Fryer, Roland G and Glenn C Loury (2005). “Affirmative action and its mythology”. In: *Journal of Economic Perspectives* 19.3, pp. 147–162.

Fryer Jr, Roland G (2007). “Belief flipping in a dynamic model of statistical discrimination”. In: *Journal of Public Economics* 91.5-6, pp. 1151–1166.

Gavoille, Nicolas and Anna Zasova (2023a). “Minimum wage spike and income underreporting: A back-of-the-envelope-wage analysis”. In: *Journal of Comparative Economics* 51.1, pp. 372–402.

— (2023b). “What we pay in the shadows: Labor tax evasion, minimum wage hike and employment”. In: *Journal of Public Economics* 228, p. 105027.

Goldfarb, Avi and Catherine Tucker (2019). “Digital economics”. In: *Journal of Economic Literature* 57.1, pp. 3–43.

Gorodnichenko, Yuriy, Jorge Martinez-Vazquez, and Klara Sabirianova Peter (2009). “Myth and reality of flat tax reform: Micro estimates of tax evasion response and welfare effects in Russia”. In: *Journal of Political economy* 117.3, pp. 504–554.

Gregg, Paul and Emma Tominey (2005). "The wage scar from male youth unemployment". In: *Labour Economics* 12.4, pp. 487 –509.

Gu, Jiadong and Peter Norman (2020). "A search model of statistical discrimination". In: *arXiv preprint arXiv:2004.06645*.

Guryan, Jonathan and Kerwin Kofi Charles (2013). "Taste-Based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots". In: *The Economic Journal* 123.572, F417–F432.

Gürtzgen, Nicole et al. (2021). "Does online search improve the match quality of new hires?" In: *Labour Economics* 70, p. 101981. ISSN: 0927-5371. DOI: https://doi.org/10.1016/j.labeco.2021.101981. URL: https://www.sciencedirect.com/science/article/pii/S0927537121000166.

Hanousek, Jan and Filip Palda (2004). "Quality of government services and the civic duty to pay taxes in the Czech and Slovak Republics, and other transition countries". In: *Kyklos* 57.2, pp. 237–252.

Harrison, Glenn W and John A List (2004). "Field experiments". In: *Journal of Economic literature* 42.4, pp. 1009–1055.

He, Yinghua and Thierry Magnac (2022). "Application costs and congestion in matching markets". In: *The Economic Journal* 132.648, pp. 2918–2950.

Heemskerk, Frank, Liviu Voinea, and Alexandra Cojocaru (2018). "Busting the Myth: The Impact of Increasing the Minimum Wage: The Experience of Romania". In: *World Bank Policy Research Working Paper* 8632.

Hoffman, Mitchell, Lisa B Kahn, and Danielle Li (2018). "Discretion in hiring". In: *The Quarterly Journal of Economics* 133.2, pp. 765–800.

Holzer, Harry and David Neumark (2000). "Assessing affirmative action". In: *Journal of Economic literature* 38.3, pp. 483–568.

Homegardner, Tammy (2020). *Pre-employment exams: What are they, and how can you prepare?* URL: https://www.forbes.com/sites/forbescoachescouncil/2020/03/26/pre-employment-exams-what-are-they-and-how-can-you-prepare/?sh=34324dcafc1f.

ILO (2022). *2021 World Employment and Social Outlook*. URL: https://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/---publ/documents/publication/wcms_771749.pdf.

Jaser, Zahira and Dimitra Petrakaki (2023). *Are you prepared to be interviewed by an ai?* URL: https://hbr.org/2023/02/are-you-prepared-to-be-interviewed-by-an-ai.

Kabos, Eszter and Margaret Meyer (2019). "A Welfare Analysis of a Steady-State Model of Observational Learning". In.

Kanoria, Yash and Daniela Saban (2021). "Facilitating the search for partners on matching platforms". In: *Management Science* 67.10, pp. 5990–6029.

Kassi, Otto, Vili Lehdonvirta, and Fabian Stephany (2021). "How many online workers are there in the world? A datadriven assessment". In: *Center for Open SCience*.

Kastlunger, Barbara et al. (2010). "Sex differences in tax compliance: Differentiating between demographic sex, gender-role orientation, and prenatal masculinization (2D: 4D)". In: *Journal of economic psychology* 31.4, pp. 542–552.

Kim, Young-Chul and Glenn C Loury (2018). "Collective reputation and the dynamics of statistical discrimination". In: *International Economic Review* 59.1, pp. 3–18.

Kirchler, Erich, Erik Hoelzl, and Ingrid Wahl (2008). "Enforced versus voluntary tax compliance: The "slippery slope" framework". In: *Journal of Economic psychology* 29.2, pp. 210–225.

Kleven, Henrik (2016). "Bunching". In: *Annual Review of Economics* 8, pp. 435–464.

Kleven, Henrik, Claus Thustrup Kreiner, and Emmanuel Saez (2016). "Why can modern governments tax so much? An agency model of firms as fiscal intermediaries". In: *Economica* 83.330, pp. 219–246.

Kleven, Henrik Jacobsen and Wojciech Kopczuk (2011). "Transfer program complexity and the take-up of social benefits". In: *American Economic Journal: Economic Policy* 3.1, pp. 54–90.

Kroft, Kory et al. (2020). *Imperfect competition and rents in labor and product markets: The case of the construction industry*. Tech. rep. National Bureau of Economic Research.

Kube, Sebastian, Michel André Maréchal, and Clemens Puppe (2013). "Do wage cuts damage work morale? Evidence from a natural field experiment". In: *Journal of the European Economic Association* 11.4, pp. 853–870.

Kuek, Siou Chew et al. (2015). "The global opportunity in online outsourcing". In.

Kumler, Todd, Eric Verhoogen, and Judith Frías (2020). "Enlisting employees in improving payroll tax compliance: Evidence from mexico". In: *Review of Economics and Statistics* 102.5, pp. 881–896.

Ladimeji, Kazim (2023). *Too many applicants? 5 innovative ways to screen out the masses*. URL: https://www.recruiter.com/recruiting/too-many-applicants-5-innovative-ways-to-screen-out-the-masses/.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler (2022). "Imperfect competition, compensating differentials, and rent sharing in the US labor market". In: *American Economic Review* 112.1, pp. 169–212.

Leibbrandt, Andreas and John A List (2015). "Do women avoid salary negotiations? Evidence from a large-scale natural field experiment". In: *Management Science* 61.9, pp. 2016–2024.

Levy, Gilat and Ronny Razin (2015). "Correlation neglect, voting behavior, and information aggregation". In: *American Economic Review* 105.4, pp. 1634–45.

Li, Cher Hsuehhsiang and Basit Zafar (2023). "Ask and you shall receive? gender differences in regrades in college". In: *American Economic Journal: Economic Policy* 15.2, pp. 359–394.

Mayshar, Joram (1991). "Taxation with costly administration". In: *The Scandinavian Journal of Economics*, pp. 75–88.

Meriküll, Jaanika and Karsten Staehr (2010). "Unreported employment and envelope wages in mid-transition: comparing developments and causes in the Baltic countries". In: *Comparative Economic Studies* 52, pp. 637–670.

Meyer, Margaret A (1991). "Learning from coarse information: Biased contests and career profiles". In: *The Review of Economic Studies* 58.1, pp. 15–41.

Moro, Andrea and Peter Norman (2003). "Affirmative action in a competitive economy". In: *Journal of Public Economics* 87.3-4, pp. 567–594.

Nichols, Donald, Eugene Smolensky, and T Nicolaus Tideman (1971). "Discrimination by waiting time in merit goods". In: *The American Economic Review* 61.3, pp. 312–323.

Niederle, Muriel (2017). "A gender agenda: a progress report on competitiveness". In: *American Economic Review* 107.5, pp. 115–119.

Niederle, Muriel and Lise Vesterlund (2007). "Do women shy away from competition? Do men compete too much?" In: *The quarterly journal of economics* 122.3, pp. 1067–1101.

Oberholzer-Gee, Felix (2008). "Nonemployment stigma as rational herding: A field experiment". In: *Journal of Economic Behavior & Organization* 65.1, pp. 30 –40.

Pallais, Amanda (2014). "Inefficient hiring in entry-level labor markets". In: *American Economic Review* 104.11, pp. 3565–3599.

Pallais, Amanda and Emily Glassberg Sands (2016). "Why the referential treatment? Evidence from field experiments on referrals". In: *Journal of Political Economy* 124.6, pp. 1793–1828.

Pantea, Smaranda (2020). "The effect of minimum wage hikes on employment: Evidence from regional panel data from Romania". In: *Economic Systems* 44.3, p. 100805.

Paulus, Alari (2015). *Tax evasion and measurement error: An econometric analysis of survey data linked with tax records*. Tech. rep. ISER Working Paper Series.

Pelletier, Adeline and Catherine Thomas (2018). "Information in online labour markets". In: *Oxford Review of Economic Policy* 34.3, pp. 376–392.

Phelps, E. (1972). "The Statistical Theory of Racism and Sexism". In: *American Economic Review* 62.4, pp. 659–661.

Pissarides, Christopher A. (1990). *Equilibrium Unemployment Theory*. English. Basil Blackwell, Oxford, UK, xiii, 173 p. : ISBN: 063115213.

Pissarides, Christopher A (2000). *Equilibrium unemployment theory*. MIT press.

Pissarides, Christopher A and Guglielmo Weber (1989). "An expenditure-based estimate of Britain's black economy". In: *Journal of public economics* 39.1, pp. 17–32.

Putniņš, Tālis J and Arnis Sauka (2015). "Measuring the shadow economy using company managers". In: *Journal of Comparative Economics* 43.2, pp. 471–490.

Rhoades, Stephen A (1993). "The herfindahl-hirschman index". In: *Fed. Res. Bull.* 79, p. 188.

Robayo-Abril, Monica, Georgiana Balaban, and Marcin Wronski (2024). "Tax compliance in Romania". In: *Policy Research working paper*.

Robayo-Abril, Monica, Mădălina Zamfir, and Marcin Wroński (2024). "Simulating Aggregate and Distributional Effects of Minimum Wage Increases in Romania". In.

Roth, Alvin E (2018). "Marketplaces, markets, and market design". In: *American Economic Review* 108.7, pp. 1609–1658.

Roussille, Nina (2021). "The central role of the ask gap in gender pay inequality". In.

Slemrod, Joel (2001). "A general model of the behavioral response to taxation". In: *International Tax and Public Finance* 8, pp. 119–128.

Smith, Lones, Peter Norman Sørensen, and Jianrong Tian (Feb. 2021). "Informational Herding, Optimal Experimentation, and Contrarianism". In: *The Review of Economic Studies* Accepted Manuscript.

Smith, Lones and Peter Sørensen (2000). "Pathological outcomes of observational learning". In: *Econometrica* 68.2, pp. 371–398.

Smith, Lones and Peter Norman Sørensen (2011). "Observational learning". In: *The New Palgrave Dictionary of Economics Online Edition*, pp. 29–52.

Stanton, Christopher T and Catherine Thomas (2016). "Landing the first job: The value of intermediaries in online hiring". In: *The Review of Economic Studies* 83.2, pp. 810–854.

— (2020). "The gig economy beyond local services and transportation". In: *CESifo Forum*. Vol. 21. 03. München: ifo Institut-Leibniz-Institut für Wirtschaftsforschung an der ..., pp. 21–26.

Tonin, Mirco (2011). "Minimum wage and tax evasion: Theory and evidence". In: *Journal of Public Economics* 95.11-12, pp. 1635–1651.

Torgler, Benno and Neven T Valev (2010). "Gender and public attitudes toward corruption and tax evasion". In: *Contemporary Economic Policy* 28.4, pp. 554–568.

Traxler, Christian (2010). "Social norms and conditional cooperative taxpayers". In: *European Journal of Political Economy* 26.1, pp. 89–103.

Upwork (2022). *Freelance forward 2022*. URL: https://www.upwork.com/research/freelance-forward-2022.

— (2023). *Gig economy statistics and key takeaways for 2024*. URL: https://www.upwork.com/resources/gig-economy-statistics.

Vâlsan, Călin, Elena Druică, and Rodica Ianole-Călin (2020). "State capacity and tolerance towards tax evasion: First evidence from Romania". In: *Administrative Sciences* 10.2, p. 33.

WEF (2023). URL: https://www.weforum.org/publications/the-future-of-jobs-report-2023/in-full/.

Weinstein, Russell (2018). "Employer screening costs, recruiting strategies, and labor market outcomes: An equilibrium analysis of on-campus recruiting". In: *Labour Economics* 55, pp. 282–299.

Williams, Colin C and Jo Padmore (2013). ""Envelope wages" in the European Union". In: *International Labour Review* 152.3-4.