ESSAYS ON EXPECTATIONS AND BELIEFS IN LABOR MARKETS

Gregor Bodo Jost, M.Sc.

Inaugural-Dissertation
zur Erlangung des akademischen Grades
Doctor rerum politicarum (Dr. rer. pol.)
der Technischen Universität Dortmund

Juli 2026

Veröffentlichung als Dissertation in der
Wirtschaftswissenschaftlichen Fakultät der
Technischen Universität Dortmund

Dortmund, 2026

Acknowledgements

Throughout the process of working on this dissertation, I benefited from the guidance and support of so many people. Looking back now, I realize even more clearly how important this has been. I would therefore like to use this opportunity to express my sincere gratitude to those who have accompanied me along the way.

First and foremost, I would like to thank my supervisor, Lukas Buchheim. I am very grateful that, when he took over the Chair of Microeconomics at TU Dortmund, he placed his trust in Marcel Vögele and me by hiring us as his first doctoral students in Dortmund. Since then, I have benefited greatly from his guidance, which he combined with the freedom to develop and pursue my own ideas, and from his feedback on my work, which always reflected a serious and thoughtful engagement with it. I am also particularly grateful for the way in which he supported my academic development, not least by encouraging and enabling my research stay at the University of Texas at Austin, even though this meant taking on additional work and teaching responsibilities during my absence.

I am also very grateful to my second supervisor, Almut Balleer, and to my third committee member, Michael Böhm, both of whom I could always turn to for helpful comments and advice.

I also want to thank my co-authors. I thank Lukas Buchheim, Sebastian Link, and Sascha Möhrle for involving me, at the very beginning of my doctoral studies, in their joint project, from which our separate joint project later emerged. I also thank my co-author and friend, Vincent Selz. Our joint project gave us many reasons to meet regularly, and while these meetings often extended well beyond work-related topics, working with him was not only fun but also very productive.

My thanks also go to the entire Department of Economics at TU Dortmund. Be it during our daily lunches, coffee breaks, or weekly seminars, our floor was always a lively, pleasant, and stimulating work environment. This made me genuinely enjoy coming to the office day after day. I would also like to thank Marcel Vögele, with whom I started my doctoral studies and shared an office throughout this time. I greatly appreciated working alongside him, preparing seminars together, and sharing everyday office life with him. I also owe thanks to Dominik Maxara: I could always approach Dominik when struggling with a proof, and he always took the time to think it through with me. I would also like

to thank the Department of Economics at the University of Texas at Austin for hosting me as a visiting student.

I am also grateful to the Ruhr Graduate School in Economics. Through the Jamboree, RGS Day, and the RGS Doctoral Conference, it provided a valuable community throughout my PhD. I would like to thank three people in particular: Marvin, Leon, and Leonie. Having met through the RGS, over time you became friends who enrich my life far beyond the academic context. Our weekly updates on Thursday evenings at the Amphütte have become a fixed point in my weeks, and I am very glad to have you all.

Finally, I am deeply grateful to my family: my parents, Katrin and Ralph, and my sisters, Caro and Camilla. Throughout my life, you have opened up a world full of opportunities to me and have always given me the absolute certainty that I can rely on your support, guidance, and love.

Beyond everything, I want to thank Becky, who means more to me than I could express in a few lines.

Contents

Introduction **1**

Main Chapters **6**

1 Downward Wage Rigidity and the Pass-Through from Expected Inflation to Wages **6**

1.1 Introduction . . . 7
1.2 Model . . . 9
1.2.1 Setup . . . 10
1.2.2 Firm-Level Wage Setting, Aggregation, and Pass-Through . . . 12
1.3 Model Results . . . 14
1.3.1 Pass-Through Without DNWR . . . 14
1.3.2 Pass-Through Under DNWR . . . 16
1.3.3 Pass-Through and Short-Run Inflation Uncertainty . . . 18
1.3.4 Pass-Through and Expected Long-Run Inflation . . . 19
1.4 Evidence . . . 20
1.4.1 Data and Empirical Specification . . . 20
1.4.2 Results . . . 23
1.5 Conclusion . . . 25

2 Overconfidence and the Gender Application Gap **26**

2.1 Introduction . . . 27
2.2 Model . . . 30
2.2.1 Setup . . . 31
2.2.2 The Direct and Indirect Effects of Overconfidence . . . 34
2.2.3 Implications of the Direct and Indirect Effects for the Application Gap . . . 37
2.3 Design of the Experiment . . . 39
2.3.1 Treatment . . . 41
2.3.2 Additional Measures . . . 43
2.3.3 Procedures . . . 44

2.3.4 Conditions and Hypotheses 44
2.4 Results 46
2.4.1 Gender Differences in Application Behavior and (Believed) Overconfidence 47
2.4.2 Causal Effects of (Believed) Overconfidence on Application Behavior 48
2.4.3 Overconfidence and the Gender Application Gap 54
2.5 Conclusion 56

3 Present Bias and On-the-Job Search **57**
3.1 Introduction 58
3.2 Three-Period Model 61
3.2.1 Setting and Timing 61
3.2.2 Preferences and Perceptions 62
3.2.3 Strategies and Solution Concept 63
3.2.4 Results 64
3.3 Infinite-Horizon Model 69
3.3.1 Strategies and Solution Concept 69
3.3.2 Results 72
3.4 Evidence 80
3.4.1 Data and Identification of Types 80
3.4.2 Results 84
3.5 Conclusion 88

Appendix **89**

A Appendix to Chapter 1 **89**
A.1 Preliminary Lemmas and Results 89
A.2 Proofs 100
A.2.1 Proof of Proposition 1.1. 100
A.2.2 Proof of Proposition 1.2. 100
A.2.3 Proof of Proposition 1.3. 102
A.2.4 Proof of Proposition 1.4. 104
A.3 Model Appendix 109
A.3.1 A Sufficient Parameter Condition for the Uncertainty Comparative Static 109
A.4 Empirical Appendix 111
A.4.1 Additional Figures and Tables 111

B Appendix to Chapter 2 **113**
B.1 Preliminary Lemmas and Results . . . 113
B.2 Proofs . . . 147
B.2.1 Proof of Lemma 2.2 . . . 147
B.2.2 Proof of Proposition 2.1 . . . 149
B.2.3 Proof of Proposition 2.2 . . . 149
B.2.4 Proof of Corollary 2.1 . . . 150
B.2.5 Proof of Corollary 2.2 . . . 150
B.3 Model Appendix . . . 151
B.3.1 Additional Figures . . . 151
B.4 Empirical Appendix . . . 152
B.4.1 Discussion on Using ORIV . . . 152
B.4.2 Additional Figures and Tables . . . 153
B.5 Experiment Instructions . . . 158

C Appendix to Chapter 3 **168**
C.1 Proofs of the Three-Period Model . . . 168
C.1.1 Proof of Lemma 3.1. . . . 168
C.1.2 Proof of Proposition 3.1. . . . 169
C.1.3 Proof of Proposition 3.2 . . . 171
C.2 Preliminary Lemmas and Results for the Infinite-Horizon Model . . . 177
C.3 Proofs of the Infinite-Horizon Model . . . 189
C.3.1 Proof of Lemma 3.2 . . . 189
C.3.2 Proof of Lemma 3.3 . . . 190
C.3.3 Proof of Corollary 3.1 . . . 194
C.3.4 Proof of Proposition 3.3 . . . 194
C.3.5 Proof of Corollary 3.2 . . . 197
C.4 Model Appendix . . . 198
C.4.1 Comparative Statics of Reservation Wages . . . 198
C.5 Empirical Appendix . . . 199
C.5.1 Details on Type Classification . . . 199
C.5.2 Additional Figures and Tables . . . 200

Bibliography **204**

List of Figures

1.1 Illustration of Wage Setting Without DNWR 15
1.2 Illustration of Wage Setting Under DNWR 17

2.1 Decision Screen 41
2.2 Gender Application Gap 48
2.3 Overconfidence and Believed Overconfidence 49
2.4 Heterogeneity of the Direct and Indirect Effect 52
2.5 Heterogeneity of the Indirect Effect by Own Belief 54

3.1 Continuation Payoffs from the Period-1 Perspective 65
3.2 Maximal Local Perception-Perfectness 73
3.3 Simulated Final Wage Distributions 78
3.4 Empirical Wage Distribution 87

A.1 Expected Inflation in the ifo Firm Survey Compared to Bundesbank Survey 112
A.2 Expected and Ex Post Realized Wage Growth 112

B.1 Equilibrium Branches 151
B.2 Vignettes 157

C.1 Comparative Statics of Reservation Wages 198
C.2 Distributions of Outcome Variables 200

List of Tables

1.1 Pass-Through and Uncertainty About Inflation 23
1.2 Pass-Through and Long-Run Inflation Expectations 24

2.1 Relative and Believed Overconfidence 49
2.2 Effects of (Believed) Overconfidence on Application Behavior 51
2.3 Gender Application Gap in INFO and CONTROL 55
2.4 Wage and Employment Outcomes 56

3.1 Worker Type, Search Beliefs, and Demanded Wage Increases 85

A.1 Pass-Through by Long-Run Inflation Expectations: Unconditional Median Split 111

B.1 Pearson Correlations Between Risk Measures and Overconfidence 153
B.2 Balance Table 153
B.3 Gender Application Gaps in CONTROL 154
B.4 Effects of (Believed) Overconfidence with Controls 155
B.5 Effects of (Believed) Overconfidence Excluding $q_i \notin [0.1, 0.9]$ 156
B.6 Anchoring Test 157

C.1 Separating Fully and Partially Sophisticated Workers 201
C.2 Excluding Self-Reported Patience as a Control Variable 202
C.3 Extensive Margin of Search Beliefs 203

Introduction

The outcomes that result from labor market choices made by workers and firms depend on a wide range of macroeconomic and individual factors. For example, the job search outcomes of workers may depend on overall labor market conditions, but also on individual characteristics such as their ability relative to other applicants. Similarly, the value of hiring another employee for a firm may depend on the current and future macroeconomic environment as well as on firm-specific business conditions. When making labor market choices, workers and firms therefore need to evaluate different options on the basis of their expectations and beliefs regarding such factors.[1] Understanding how expectations and beliefs influence labor market behavior, and how such expectations and beliefs vary systematically across labor market participants, can thus help explain patterns in outcomes at both the individual and aggregate levels.

This dissertation studies this broader question across three specific margins of labor market behavior. The first of these margins concerns firms' wage-setting decisions. A recent and growing literature studies how firms' inflation expectations affect wage-setting decisions (e.g., Abberger et al., 2025; Baumann et al., 2024; Buchheim et al., 2024; Coibion et al., 2018; Gautier et al., 2025; Savignac et al., 2024). The relationship between inflation expectations and wages is of particular interest for two reasons. First, the strength of the pass-through from expected inflation into wage growth can mute or amplify inflationary wage-price dynamics. Second, for a given path of nominal wages, inflation determines the real value of workers' compensation. Despite its relevance, however, there remains a gap between what theory and empirical estimates suggest about the size of this pass-through, with empirical estimates pointing to a substantially smaller pass-through than theory predicts. Chapter 1 addresses this gap by offering a theoretical explanation for the empirically low pass-through based on inflation uncertainty and downward nominal wage rigidity.

Turning to the supply side of the labor market, the second margin of labor market behavior studied in this dissertation concerns the job application behavior of workers. When workers choose which jobs to apply to, they may consider, among other factors, the

[1] Throughout this introduction, expectations typically refer to assessments of future macroeconomic conditions, while beliefs refer to assessments of one's own characteristics and behavior or those of other market participants. Each self-contained chapter defines the relevant concepts in its own setting.

wage a job offers and their perceived probability of being hired. This perceived probability may depend on how workers assess their own ability relative to other applicants. Such assessments may be particularly relevant for higher-wage vacancies, which tend to attract more applicants and are therefore more competitive (Banfi and Villena-Roldán, 2019; Belot et al., 2022). If beliefs about own relative ability differ systematically across groups of job seekers, these differences may thus generate systematic differences in application behavior. Chapter 2 takes up this idea by studying whether gender differences in beliefs about own relative ability can help explain recent evidence showing that men apply more often to higher-wage jobs than women, a pattern that accounts for a substantial part of the gender wage gap (Fluchtmann et al., 2024; Lochner and Merkl, 2026).

The third margin of labor market behavior studied in this dissertation concerns workers' on-the-job search and job-to-job transitions. Workers continue to make consequential labor market choices after entering employment. For example, by searching for other jobs while employed, they can switch to higher-paying jobs and thereby move up the wage ladder. Such job-to-job transitions are quantitatively important: especially early in the career, estimates for Germany and the U.S. suggest that they account for roughly one third of wage growth (e.g., Adda and Dustmann, 2023; Topel and Ward, 1992). Recent work studies how workers' beliefs about the wages they could earn in other jobs can influence on-the-job search behavior and job-switching decisions (e.g., Jäger et al., 2024; Pilossoph and Ryngaert, 2025). Chapter 3 takes a different perspective by shifting the focus from beliefs about outside opportunities to beliefs about own future search behavior. It studies how workers who underestimate their future tendency to avoid costly search may accept low-paying jobs as temporary steps on the wage ladder, but then do not continue searching on the job and become trapped in these jobs.

Together, the three self-contained chapters of this dissertation contribute to the literature on how expectations and beliefs shape labor market behavior by studying three specific choice margins: wage setting by firms, application decisions by job seekers, and on-the-job search and job-to-job transitions by workers. Methodologically, each chapter combines a theoretical model with an empirical analysis that assesses central predictions of the respective model. The remainder of this introduction provides a brief overview of each chapter.

Chapter 1. This chapter, which is joint work with Lukas Buchheim, Sebastian Link, and Sascha Möhrle, provides a theoretical rationale for the empirically small effect of inflation expectations on wage inflation. Our model combines downward nominal wage rigidity (DNWR), which prevents firms from cutting nominal wages, with uncertainty about future inflation. The interaction between uncertainty and DNWR creates an asymmetric risk in

wage setting. If realized inflation turns out higher than initially expected, the firm pays lower real wages than expected and can correct this by raising nominal wages. If realized inflation turns out lower than initially expected, however, the firm may end up paying real wages that are too high. DNWR then prevents the firm from correcting this through nominal wage cuts. Anticipating this risk, firms raise wages more cautiously ex ante, which lowers the pass-through from expected inflation into wage growth.

The model also links the size of the pass-through to the degree of uncertainty about future inflation and to the expected path of inflation. First, higher short-run inflation uncertainty lowers the pass-through because it raises the probability that realized inflation will be sufficiently low for DNWR to bind. Second, higher long-run inflation expectations raise the pass-through. When firms expect higher inflation in future wage spells, they anticipate higher wage increases at future adjustment dates. This lowers the risk that DNWR binds in the future and allows firms to set higher current wages.

We assess these theoretical predictions empirically using a large panel of German firms. The results are consistent with the uncertainty channel: in line with the model, the pass-through is lower among firms with greater uncertainty about future inflation. We also find some evidence that the pass-through increases with long-run inflation expectations, although this result is less robust across specifications.

Chapter 2. This chapter, which is joint work with Vincent Selz, studies whether gender differences in overconfidence about own relative ability can help explain the gender application gap. Existing work shows that women apply to lower-paying jobs than men and that this application behavior is quantitatively important for gender differences in wages (Fluchtmann et al., 2024; Lochner and Merkl, 2026). In particular, Fluchtmann et al. (2024) find that the gender application gap accounts for more than 70% of the gender wage gap among newly hired workers with similar observable characteristics. While this evidence highlights the importance of application behavior for gender differences in wages, the mechanisms behind the gender application gap remain less well understood. Chapter 2 studies one such potential mechanism: gender differences in overconfidence about own relative ability.

The chapter develops a directed search model in which job seekers compete for vacancies that differ in wages. The probability of receiving an offer increases in the ability of a job seeker relative to the other applicants. If job seekers overestimate their relative ability, that is, if they are overconfident, they assign a higher probability to receiving an offer from a higher-wage vacancy and apply to such vacancies more often. This is the direct effect of overconfidence. The model also identifies an indirect effect. If job seekers believe that others are overconfident, they expect stronger competition for higher-wage

vacancies. This discourages applications to these vacancies, especially among job seekers who are not themselves overconfident. Through these direct and indirect channels, gender differences in overconfidence can generate a gender application gap.

To test these channels empirically, we develop a laboratory experiment that closely follows the spirit of the model. Participants first work on a cognitive ability task. The experiment then elicits incentivized beliefs about own relative ability and about the relative ability beliefs of others, providing measures of overconfidence and of second-order beliefs about overconfidence. Participants subsequently enter a stylized labor market in pairs, where they compete for two vacancies by allocating application weights between them. As in the model, vacancies differ only in posted wages, and higher relative ability increases the probability of receiving an offer. In control pairs, neither participant receives information about own actual relative ability. In treated pairs, one participant receives precise information about own actual relative ability, which removes overconfidence. The other participant receives no information about own ability but knows that the competitor received this information, which removes the reason to expect that the competitor acts on overconfident beliefs. This design allows us to identify the direct and indirect effects of overconfidence on application behavior and relate them to the application gap.

In line with field evidence, we find that men allocate more application weight to the higher-wage vacancy than women. Men are also more overconfident about their relative ability. Among initially overconfident participants, receiving information about their own relative ability reduces applications to the higher-wage vacancy. This provides evidence for the direct effect of overconfidence. Comparing the gender application gap in the control group to the gap among informed participants suggests that this direct effect accounts for about one third of the application gap in the experiment. By contrast, we find no evidence for the indirect effect. Both women and men believe that others are overconfident on average, but these second-order beliefs about the overconfidence of others do not appear to affect application behavior.

Chapter 3. The last chapter of this dissertation studies how present bias and perceptions about future present bias shape job search, job offer acceptance behavior, and wage outcomes in a model with costly on-the-job search. Present bias, defined as the tendency to place disproportionate weight on immediate utility relative to future utility, appears widespread, with estimates suggesting that up to three quarters of individuals exhibit some degree of present bias (Cobb-Clark et al., 2024). This bias can matter whenever economically important choices require costly action today in exchange for delayed returns. The labor market is a natural setting for such choices, since many labor market decisions involve this type of trade-off and have persistent consequences for workers' wage

paths. For example, on-the-job search requires immediate and costly effort but may generate offers from higher-paying jobs and thereby lead to job-to-job transitions, which are an important source of wage growth over the career (e.g., Adda and Dustmann, 2023; Faberman et al., 2022; Topel and Ward, 1992). This central role of on-the-job search in wage growth, together with its intertemporal payoff structure, makes it a choice margin on which present bias may thus be particularly consequential.

The chapter develops an on-the-job search model with three types of workers. Time-consistent workers have no present bias. Naive and sophisticated workers are present-biased, but differ in their awareness of future present bias: naive workers do not anticipate it, while sophisticated workers do. Present bias reduces search because workers place extra weight on current search costs relative to future wage gains. Beyond this search margin, the model identifies a novel offer acceptance channel. Workers who do not anticipate their future present bias may accept low-wage jobs as temporary steps on the wage ladder. Once employed at these jobs, however, present bias may prevent further on-the-job search, leaving them trapped in low-wage jobs. Workers who anticipate their future present bias are more selective. They may demand larger wage increases to change jobs because some offers would end future search without providing a sufficiently high permanent wage. Through this greater selectivity, they avoid becoming trapped in low-wage jobs.

Chapter 3 concludes by examining whether descriptive evidence from the German Socio-Economic Panel Innovation Sample is in line with the model's predictions. The evidence is most consistent with the predicted differences in offer acceptance behavior: sophisticated workers demand the largest wage increases to be willing to change jobs, while naive and time-consistent workers do not differ significantly. Moreover, sophisticated workers tend to assign the lowest probabilities to engaging in future on-the-job search, which is consistent with their awareness of future present bias, although the evidence on this is statistically weaker.

CHAPTER 1

Downward Wage Rigidity and the Pass-Through from Expected Inflation to Wages

This chapter is co-authored with Lukas Buchheim, Sebastian Link and Sascha Möhrle.

Abstract

This paper provides a theoretical rationale for the empirically small effect of inflation expectations on wage inflation. We develop a partial equilibrium model of Taylor-type wage setting with inflation uncertainty and downward nominal wage rigidity (DNWR). When inflation turns out lower than expected, DNWR can leave firms with excessively high real wages that cannot be corrected through nominal wage cuts. Anticipating this risk, firms raise wages more cautiously, which lowers the pass-through from expected inflation into wage inflation. The model predicts that higher short-run inflation uncertainty decreases the pass-through by raising the risk of sufficiently low inflation realizations. It also predicts that higher long-run inflation expectations increase the pass-through because they raise anticipated future wages and thereby reduce the risk that DNWR will bind. Using survey data from German firms, we find evidence in line with the uncertainty prediction and less robust evidence that the pass-through increases with long-run inflation expectations.

1.1. Introduction

How large is the effect of inflation expectations on wage inflation? The answer to this question is central to our understanding of inflationary wage-price dynamics and the risk of wage-price spirals. If firms respond to higher expected inflation with substantial wage increases, inflation expectations themselves can raise firms' labor costs and thereby contribute to rising prices. A low pass-through from expected inflation into wage growth, by contrast, can dampen such wage-price dynamics. Motivated by the surge in inflation following the COVID-19 pandemic and Russia's invasion of Ukraine, a recent empirical literature has thus quantified this pass-through (e.g., Abberger et al., 2025; Baumann et al., 2024; Buchheim et al., 2024; Gautier et al., 2025; Savignac et al., 2024). Across different contexts, estimates of the pass-through are consistently low: a one percentage point increase in expected inflation over the next 12 months is typically associated with less than a 0.3 percentage point increase in average wage growth, with most estimates substantially below that. This robustly low pass-through is difficult to reconcile with benchmark time-dependent wage-setting models, which predict a pass-through of at least one half (Werning, 2022).[1]

This paper addresses the gap between low empirical estimates of the pass-through and larger predictions from benchmark time-dependent wage-setting models by introducing a theoretical mechanism that dampens wage growth responses to expected inflation. The model combines two well-documented empirical facts: firms are uncertain about future inflation (e.g., Candia et al., 2024; Weber et al., 2022), and downward nominal wage rigidity (DNWR) prevents firms from cutting nominal wages (e.g., Barattieri et al., 2014; Fehr and Goette, 2005; Schaefer and Singleton, 2023). The interaction between inflation uncertainty and DNWR creates an asymmetric risk in wage setting. If realized inflation turns out higher than initially expected, the firm ends up paying lower real wages than expected, which it can correct by raising nominal wages. If realized inflation turns out lower than initially expected, however, the firm may end up paying real wages that are too high. DNWR then prevents the firm from correcting this through nominal wage cuts. Anticipating this risk, firms raise wages more cautiously ex ante, which, in the aggregate, lowers the pass-through from expected inflation into wage inflation.

We formalize this mechanism in a partial equilibrium model of Taylor-type wage setting. Firms can adjust nominal wages only at regularly scheduled dates and hold probability distributions over future inflation rates. The model distinguishes between short-run

[1] Werning (2022) also shows that state-dependent menu-cost models can generate a very low, even negative, pass-through. The benchmark comparison here is to time-dependent wage-setting models, which provide a natural benchmark because wage adjustments often occur at fixed intervals (e.g., Druant et al., 2012; Gödl and Gödl-Hanisch, 2024; Le Bihan et al., 2012).

expectations over the current wage spell and long-run expectations over subsequent wage spells. In the benchmark case without DNWR, the model reproduces the Taylor-type pass-through of one half from Werning (2022): absent DNWR, firms center the expected real wage path around the profit-maximizing real wage. Since inflation gradually erodes real wages over the wage spell, firms do so by setting current wage growth equal to one half of expected inflation over the spell. Introducing DNWR lowers the pass-through below this benchmark because firms internalize the risk that low inflation realizations leave them with excessively high real wages that they cannot correct ex post through nominal wage cuts. In this way, the interaction of inflation uncertainty and DNWR moves the theoretical pass-through in the direction of the low estimates found in the empirical literature.

The model also links the size of the pass-through to the degree of uncertainty about short-run inflation and to the expected long-run path of inflation. First, higher short-run uncertainty lowers the pass-through, because it raises the probability that realized inflation will be sufficiently low for DNWR to bind. Second, higher long-run inflation expectations raise the pass-through. When firms expect higher inflation in future wage spells, they anticipate higher wage increases at future adjustment dates. This lowers the risk that DNWR binds in the future and allows firms to set higher current wages. The pass-through from expected inflation into wage inflation may thus be low at the beginning of a high-inflation period when uncertainty is high and long-run inflation expectations remain anchored at a low level. However, the pass-through may rise once high inflation is perceived as more persistent and uncertainty declines.

Anchoring long-run inflation expectations at a low level is widely viewed as an important tool for achieving price stability (e.g., Armantier et al., 2022; Bernanke, 2007; Gürkaynak et al., 2010; Mishkin, 2007). The long-run expectations channel identified in this paper provides a new theoretical rationale for this: low long-run inflation expectations make future DNWR constraints more relevant, so that firms raise wages more cautiously. This reduces the pass-through from short-run inflation expectations into wage inflation, which may mitigate self-reinforcing wage-price dynamics.

Finally, we assess the predictions of the model using quarterly survey data from a large panel of German firms. The survey elicits firms' short-run inflation expectations, expected wage growth, confidence in their inflation expectations, which we interpret as a measure of uncertainty, and longer-run inflation expectations. This allows us to estimate the pass-through from expected inflation into expected wage growth and to study whether this pass-through varies with uncertainty and long-run expectations as predicted by the model. In line with the empirical literature, we find that the average pass-through is substantially below the Taylor-type benchmark of one half. The results also support the uncertainty channel from the model: the pass-through is lower among firms that report

greater uncertainty about future inflation. We also find evidence that the pass-through tends to increase with long-run inflation expectations, although this evidence is less robust across specifications.

Our paper relates to two main strands of the literature. First and foremost, we contribute to the literature that studies the pass-through from inflation expectations into wages. A growing empirical literature uses firm and household survey data to estimate the size of this pass-through (e.g., Abberger et al., 2025; Baumann et al., 2024; Buchheim et al., 2024; Coibion et al., 2018; Gautier et al., 2025; Hajdini et al., 2023; Jain et al., 2024; Savignac et al., 2024). As discussed above, a robust finding in this literature is that the pass-through from expected inflation into wages is positive but low. This finding is difficult to reconcile with benchmark time-dependent wage-setting models (Werning, 2022). We contribute to this literature by providing a theoretical mechanism that lowers the pass-through and can thereby reconcile Taylor-type wage setting with the empirical estimates.

Second, we contribute to the literature on DNWR. Prior work documents that nominal wage cuts are rare and studies the macroeconomic implications of DNWR (e.g., Barattieri et al., 2014; Davis and Krolikowski, 2025; Fehr and Goette, 2005; Schaefer and Singleton, 2023; Schmitt-Grohé and Uribe, 2016). Related to our paper, Elsby (2008) and, building on this, Stüber and Beissinger (2012) show that DNWR can compress wage increases following realized productivity shocks. We add to this literature by studying how DNWR reduces the pass-through from expected inflation into wage inflation and how this pass-through varies with short-run inflation uncertainty and, through firms' anticipation of future wage-setting decisions, the expected path of future inflation.

1.2. Model

The model builds on three main assumptions. First, firms set nominal wages in a Taylor-type way: the periods in which firms can adjust their nominal wages are exogenously given, known in advance, and occur at fixed intervals. Second, firms are uncertain about future inflation: they hold probability distributions over future inflation rather than point forecasts held with certainty, with expected inflation corresponding to the mean of the distribution and uncertainty to its variance. Third, nominal wage adjustment is subject to downward rigidity: firms cannot reduce nominal wages. The interaction between inflation uncertainty and DNWR then leads firms to increase wages more cautiously, which constitutes the mechanism that dampens the pass-through from expected inflation into wage inflation.

These assumptions are empirically well motivated. Evidence suggests that firms often

adjust wages in a Taylor-type manner, with nominal wage changes occurring at regular adjustment dates. Druant et al. (2012) report that 54% of firms in the euro area typically change wages in specific months. For Germany, Gödl and Gödl-Hanisch (2024) show that 57% of firms change wages at an annual frequency and, in particular, do not plan to adjust this frequency in response to changes in inflation expectations. Using French micro data, Le Bihan et al. (2012) find no evidence in favor of state-dependent or Calvo-type wage setting and interpret their results as partly consistent with Taylor-type contracts. Moreover, firms report substantial uncertainty around their inflation forecasts; see Candia et al. (2024) for evidence from U.S. firms and Weber et al. (2022) for a broader survey documenting uncertainty as a systematic feature of the inflation expectations of households and firms. Finally, DNWR is a robust empirical feature of labor markets (e.g., Barattieri et al., 2014; Fehr and Goette, 2005; Schaefer and Singleton, 2023) for reasons such as adverse effort responses by workers to nominal wage cuts or legal restrictions (e.g., Babecký et al., 2010; Buchanan and Houser, 2022).

1.2.1. Setup

Wage and price variables enter the model in logs throughout the analysis. Since log changes closely approximate percentage changes for moderate changes in the underlying variables, we interpret log changes in the aggregate price level as inflation, and log changes in wages as wage growth at the firm level and as wage inflation at the aggregate level.

There is a continuum $[0, 1]$ of identical firms, each employing the same number of workers. Firms are infinitely lived. At any point in time, a firm's profits depend only on the log deviation of the real wage it pays, $w > 0$, from the profit-maximizing real wage, $w^* > 0$. Specifically, profits at any point in time are given by a function $\Pi(z)$ that is quadratic and strictly concave in the log real wage deviation, $z = \log(w) - \log(w^*)$, with $\Pi'(0) = 0$ and $\Pi''(z) = -c < 0$ for all $z \in \mathbb{R}$. Importantly, the profit-maximizing real wage w^* is constant over time and does not depend on inflation.

Firms can adjust nominal wages according to a Taylor-type schedule. In particular, firms can adjust nominal wages every $d \in \mathbb{N}$ periods, with $d \geq 1$. Thus, if a firm is eligible to set a new nominal wage in period t, it can adjust this wage again in period $t + d$. Firms discount payoffs between successive adjustment opportunities with discount factor $\delta \in (0, 1)$. For tractability, we abstract from discounting within a wage spell. Starting dates of wage spells are staggered across firms, with a share of $1/d$ of all firms being able to adjust nominal wages in each period.[2]

For an arbitrary period t, we refer to the time span from period t to period $t + d$ as

[2]This captures the idea that starting periods of wage spells are not synchronized across firms. For example, if wage spells last two periods, one half of firms adjusts in each period.

the short run. This time span corresponds to the wage spell that starts in period t. All subsequent time spans of length d, from $t+d$ to $t+2d$, from $t+2d$ to $t+3d$, and so on, constitute the long run from the perspective of period t.

Let P_τ denote the price level at time τ. In any period t, firms hold exogenous inflation expectations over different time horizons. Short-run expectations concern the cumulative inflation from period t to period $t+d$, given by $\log(P_{t+d}) - \log(P_t)$. Long-run expectations concern the cumulative inflation over all subsequent intervals of length d, from $t+d$ to $t+2d$, from $t+2d$ to $t+3d$, and so on. These expectations are described by independent normal distributions, $\mathcal{N}(\pi^e_{SR}, \sigma^2_{SR})$ for the wage spell from t to $t+d$ and $\mathcal{N}(\pi^e_{LR}, \sigma^2_{LR})$ for each subsequent wage spell, with cumulative distribution functions F_{SR} and F_{LR} and densities f_{SR} and f_{LR}, respectively. We assume $\pi^e_{SR}, \pi^e_{LR} > 0$ throughout.

Firms are therefore uncertain about the cumulative inflation over wage spells in the short and long run. Conditional on any possible realization of cumulative inflation over a wage spell, however, firms expect inflation to follow a deterministic within-spell path, accumulating evenly over the spell. Let π_{SR} denote an inflation realization from period t to period $t+d$. Conditional on π_{SR}, expected accumulated inflation from period t to any point in time $t+s$ with $s \in [0, d]$ thus equals

$$\pi^e_{t,t+s}(\pi_{SR}) = s\frac{\pi_{SR}}{d}.$$

Taking expectations over π_{SR} gives, for all $s \in [0, d]$,

$$\pi^e_{t,t+s} \coloneqq \mathbb{E}_{SR}\left[\pi^e_{t,t+s}(\pi_{SR})\right] = s\frac{\pi^e_{SR}}{d}.$$

For a firm i that sets a nominal wage $W_{i,t}$ in period t, let $w_{i,t} = W_{i,t}/P_t$ denote the real wage implied by this nominal wage choice. Define

$$x_{i,t} \coloneqq \log(w_{i,t}) - \log(w^*).$$

Thus, $x_{i,t}$ is the log deviation of the real wage implied by $W_{i,t}$ in period t from the profit-maximizing real wage w^*. Equivalently, $x_{i,t}$ is the log deviation of the nominal wage $W_{i,t}$ from the nominal wage $P_t w^*$ that implements w^* in period t. We refer to $x_{i,t}$ as the wage markup chosen by firm i in period t. Conditional on π_{SR}, the log real wage deviation the firm expects at any point in time $t+s$, with $s \in [0, d]$, is then

$$\log\left(w^e_{i,t+s}(\pi_{SR})\right) - \log(w^*) = x_{i,t} - s\frac{\pi_{SR}}{d}.$$

Profits accrue continuously over the wage spell. We adopt the interpretation that deviations of the real wage paid by firm i from the profit-maximizing real wage w^* constitute losses to the firm. For an initial wage markup $x_{i,t}$ and a short-run inflation realization

π_{SR}, the losses implied by the expected real-wage path over the current wage spell are therefore given by

$$L(x_{i,t}, \pi_{SR}) \coloneqq \int_0^d \left[\Pi(0) - \Pi\left(x_{i,t} - s\frac{\pi_{SR}}{d}\right)\right] ds.$$

Minimizing this loss function is equivalent to maximizing profits over the wage spell.

The optimal markup choice $x^*_{i,t}$ captures the forward-looking component of firm-level wage growth. If future conditions played no role in the current wage-setting decision, the firm would choose the nominal wage such that the real wage equals w^* and would therefore set $x^*_{i,t} = 0$. Thus, any nonzero optimal markup captures the component of wage growth driven by expected future inflation, taking into account future wage-setting constraints.[3]

Total firm-level wage growth may also include a catch-up component: if the real wage prior to adjustment differs from w^*, part of total wage growth reflects the correction of this initial gap. This component would arise even if forward-looking considerations played no role in wage-setting decisions. Since our focus is on the forward-looking pass-through from expected inflation into wage inflation, we abstract from this backward-looking component. Henceforth, when we refer to wage growth, we thus mean the forward-looking markup.[4]

1.2.2. Firm-Level Wage Setting, Aggregation, and Pass-Through

Solving the model consists of three steps. First, we solve for the optimal wage-setting decision of a firm i that can adjust its nominal wage in period t. This determines firm-level wage growth. Second, we aggregate the wage-setting decisions of all firms that adjust nominal wages in that period, yielding aggregate wage growth, i.e., wage inflation. Third, we relate wage inflation to expected inflation from period t to period $t+1$ to derive the pass-through from expected inflation into wage inflation.

Importantly, firms take inflation expectations, as described by the probability distributions, as given and do not update them through any model of expectation formation. Hence, neither individual wage-setting decisions nor aggregate wage growth feed back into expected inflation.[5] This partial equilibrium structure is a deliberate modeling choice: it shuts down any potential feedback channel from wage-setting decisions to inflation expec-

[3] Because the wage markup can be written either as the log deviation of the real wage from w^* or as the log deviation of the nominal wage from $P_t w^*$, this forward-looking component of wage growth can be interpreted in both real and nominal terms.

[4] The empirical pass-through literature disentangles the forward-looking pass-through from the catch-up component of wage growth by controlling for past inflation (Buchheim et al., 2024; Gautier et al., 2025) or by using randomized information provision that generates exogenous variation in inflation expectations (Abberger et al., 2025; Baumann et al., 2024; Savignac et al., 2024).

[5] The atomistic nature of firms already rules out a direct effect from any individual wage-setting decision to aggregate outcomes and, hence, inflation expectations. The assumption of exogenous inflation expectations additionally rules out feedback from aggregate wage-setting behavior to expected inflation.

tations and thereby allows us to isolate the direct causal effect of expected inflation on wage inflation.

We now characterize the wage-setting problem of a firm i that can adjust its nominal wage in period t. The firm chooses a wage markup $x_{i,t}$ by trading off expected losses during the current wage spell against the probability that DNWR binds at the next adjustment opportunity and the associated continuation losses. Let $x^*_{i,t+d}$ denote the wage markup the firm would choose in period $t+d$ if DNWR did not bind, based on its long-run inflation expectations.[6] Given short-run inflation π_{SR}, DNWR binds in period $t+d$ whenever

$$x_{i,t} - \pi_{SR} > x^*_{i,t+d},$$

or, equivalently, whenever $\pi_{SR} < x_{i,t} - x^*_{i,t+d}$. In that event, the wage set in period t, together with the realized short-run inflation, implies an end-of-spell markup above the markup the firm would like to choose in period $t+d$. Hence, the firm cannot optimally readjust its wage and thus carries forward an excessively high real wage into the next wage spell. If DNWR does not bind, the firm can choose the optimal markup $x^*_{i,t+d}$ in period $t+d$. For any markup $x_{i,t}$ chosen in period t, the firm's expected discounted loss is therefore

$$\begin{aligned} V_t(x_{i,t}) \coloneqq\ & \mathbb{E}_{SR}\big[L(x_{i,t}, \pi_{SR})\big] \\ & + \delta F_{SR}(x_{i,t} - x^*_{i,t+d}) \mathbb{E}_{SR}\big[V_{t+d}(x_{i,t} - \pi_{SR}) \mid \pi_{SR} < x_{i,t} - x^*_{i,t+d}\big] \\ & + \delta\big(1 - F_{SR}(x_{i,t} - x^*_{i,t+d})\big) V_{t+d}(x^*_{i,t+d}). \end{aligned} \tag{1.1}$$

The first term in the definition of V_t equals the expected losses during the current wage spell. The second term equals the discounted continuation losses when DNWR binds in period $t+d$, weighted by the probability of this event, $F_{SR}(x_{i,t} - x^*_{i,t+d})$. The third term equals the discounted continuation losses when DNWR does not bind in period $t+d$, weighted by the complementary probability $1 - F_{SR}(x_{i,t} - x^*_{i,t+d})$. The optimal period-$t$ markup then solves

$$x^*_{i,t} = \arg\min_{x_{i,t}} V_t(x_{i,t}).$$

As discussed above, $x^*_{i,t}$ describes the forward-looking component of wage growth of firm i. It remains to derive the implied aggregate nominal wage growth, i.e., wage inflation, and the corresponding pass-through from this individual wage growth. Since all firms are identical and only a share $1/d$ can adjust nominal wages in any period, the model-implied wage inflation driven by expected future inflation in period t equals

$$x^*_t = \frac{x^*_{i,t}}{d}.$$

[6] In the long-run problem, we restrict attention to stationary solutions: absent a binding DNWR constraint, the firm anticipates the same optimal markup at each future adjustment date. Corollary A.1 in Appendix A.1 establishes existence of this optimal markup and of the associated long-run value function.

We then define the pass-through as the ratio of wage inflation to expected inflation over the next period,

$$\phi^* \coloneqq \frac{x_t^*}{\pi_{t,t+1}^e}.$$

This measure provides the model analogue of the empirical pass-through from expected inflation to wage inflation. Empirically, this pass-through is typically estimated from a linear relationship between expected inflation over the next year and wage growth. It therefore provides a linear summary of how expected inflation relates to average wage growth, although this underlying relationship may in principle be nonlinear. Correspondingly, ϕ^* measures the wage inflation implied by the model per percentage point of expected inflation over the next period. The model-implied relationship between expected inflation and wage inflation may also be nonlinear, so ϕ^* likewise serves as a linear summary of this relationship for a given level of expected inflation. Since all results derived below hold for any positive level of expected inflation, this summary measure provides a robust theoretical counterpart to the empirical pass-through estimates.

1.3. Model Results

This section presents the model results on the pass-through from expected inflation into wage inflation. We first revisit the benchmark environment without DNWR, as discussed in Werning (2022), and show that the pass-through equals 1/2 in this setting. We then show that DNWR reduces the pass-through below this benchmark. We subsequently analyze two comparative statics under DNWR: greater uncertainty about short-run inflation decreases the pass-through, whereas higher long-run inflation expectations increase it.

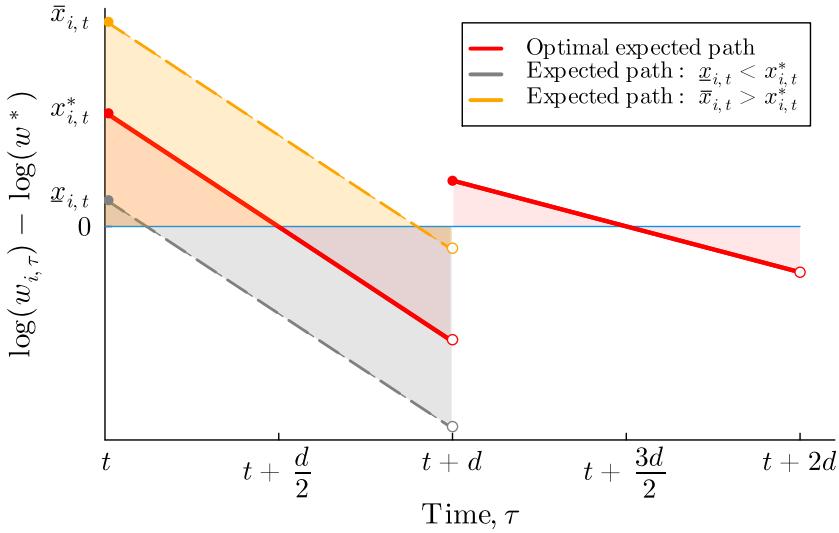
1.3.1. Pass-Through Without DNWR

In the benchmark environment without DNWR, firm i's wage choice in period t has no implications for future wage-setting decisions because the firm can always freely readjust nominal wages at the next adjustment opportunity. The firm's problem therefore reduces to choosing $x_{i,t}^*$ such that

$$x_{i,t}^* = \arg\min_{x_{i,t}} \mathbb{E}_{SR}\big[L(x_{i,t}, \pi_{SR})\big].$$

Proposition 1.1 states that, in this environment, firm-level wage growth equals one half of expected short-run inflation. Together with $x_t^* = x_{i,t}^*/d$ and $\pi_{t,t+1}^e = \pi_{SR}^e/d$, this implies that the pass-through from expected inflation into wage inflation equals 1/2.

Figure 1.1: Illustration of Wage Setting Without DNWR



Note: The figure illustrates expected paths of the log real wage deviation from the profit-maximizing real wage in the benchmark environment without DNWR. The horizontal axis shows time, τ, and the vertical axis shows $\log(w_{i,\tau}) - \log(w^*)$. From t to $t+d$, expected paths decline at a rate equal to π^e_{SR}/d; from $t+d$ to $t+2d$, the expected path declines at a rate equal to π^e_{LR}/d. From t to $t+d$, the red solid line shows the expected path under the optimal markup $x^*_{i,t} = \pi^e_{SR}/2$; the dashed lines show paths for markups below and above the optimal markup. From $t+d$ to $t+2d$, the red line shows the expected path under the optimal long-run markup. The shaded areas indicate the deviations of the real wage from the profit-maximizing level over the wage spell.

Proposition 1.1 *In the benchmark environment without DNWR, the individual wage growth of a firm i that can adjust nominal wages in period t equals*

$$x^*_{i,t} = \frac{1}{2}\pi^e_{SR}.$$

Aggregate wage growth in period t therefore equals

$$x^*_t = \frac{1}{2}\pi^e_{t,t+1}.$$

Hence, the pass-through from expected inflation into wage inflation equals

$$\phi^* = \frac{1}{2}.$$

Proof. See Appendix A.2.1. □

Following the argument in Werning (2022), Figure 1.1 illustrates the logic behind Proposition 1.1. Suppose that firm i chooses a markup $\bar{x}_{i,t} > x^*_{i,t} = \pi^e_{SR}/2$ in period t. Then, the real wage is initially relatively far above the profit-maximizing level. Only as expected inflation erodes the real value of the fixed nominal wage over the course of the spell, the expected real wage moves closer to the profit-maximizing level toward the end of the spell. Conversely, if the firm chooses $\bar{x}_{i,t} < x^*_{i,t} = \pi^e_{SR}/2$, the expected real wage

is initially closer to the profit-maximizing level but falls substantially below it over the course of the wage spell. In both cases, the firm incurs relatively large expected losses because the expected real wage is far from the profit-maximizing level for a sizable part of the wage spell.

The loss-minimizing choice centers the expected real wage path around w^*. The firm achieves this by setting $x^*_{i,t} = \pi^e_{SR}/2$, so that the expected real wage lies above w^* during the first half of the wage spell and below w^* during the second half. Aggregating individual wage-setting decisions then implies a pass-through from expected inflation into wage inflation of $\phi^* = 1/2$.[7]

1.3.2. Pass-Through Under DNWR

Proposition 1.2 shows that DNWR decreases the pass-through from expected inflation into wage inflation below the benchmark level of 1/2.

Proposition 1.2 *Under DNWR, the individual wage growth of a firm i that can adjust nominal wages in period t satisfies*

$$x^*_{i,t} < \frac{1}{2}\pi^e_{SR}.$$

Aggregate wage growth in period t therefore satisfies

$$x^*_t < \frac{1}{2}\pi^e_{t,t+1}.$$

Hence, the pass-through from expected inflation into wage inflation satisfies

$$\phi^* < \frac{1}{2}.$$

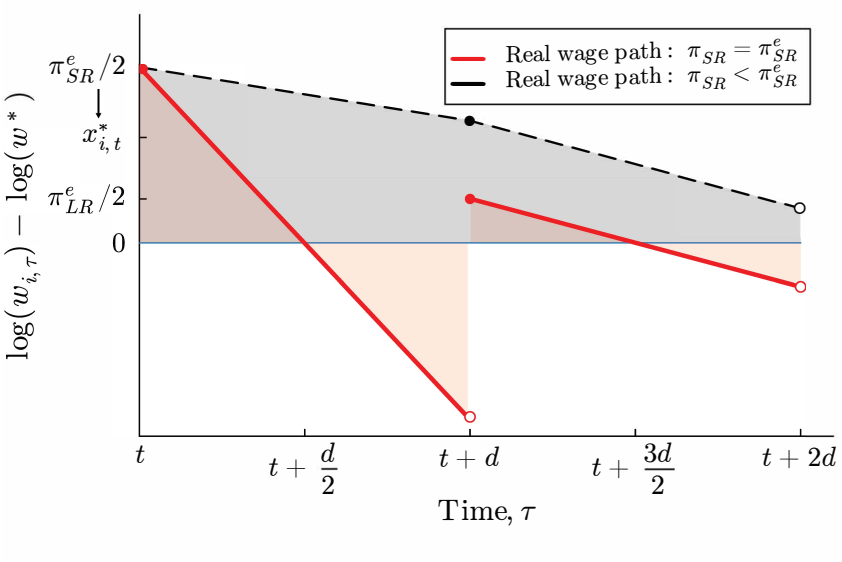
Proof. See Appendix A.2.2. □

Figure 1.2 illustrates the logic behind Proposition 1.2. Suppose firm i sets wages as in the benchmark environment without DNWR, choosing markups $x_{i,t} = \pi^e_{SR}/2$ in period t and $x_{i,t+d} = \pi^e_{LR}/2$ in period $t+d$. If realized short-run inflation equals its expected level, $\pi_{SR} = \pi^e_{SR}$, the end-of-spell markup lies below the markup that the firm wants to set in period $t+d$. The DNWR constraint therefore does not bind, and the firm can adjust its wage optimally in period $t+d$.

If realized short-run inflation is sufficiently below its expected level, this is no longer true. In this case, the end-of-spell markup exceeds the markup that the firm would like to set in period $t+d$. The DNWR constraint therefore binds and prevents the firm from

[7] The effect of inflation expectations on optimal wage setting generally depends on the firm's objective function. For example, an asymmetric profit function around the profit-maximizing real wage w^* could shift the optimal markup upward or downward relative to the benchmark.

Figure 1.2: Illustration of Wage Setting Under DNWR



Note: The figure illustrates expected paths of the log real wage deviation from the profit-maximizing real wage under DNWR. The horizontal axis shows time, τ, and the vertical axis shows $\log(w_{i,\tau}) - \log(w^*)$. The plotted paths use the benchmark markups from Proposition 1.1. From t to $t + d$, the red solid line represents the path if realized short-run inflation equals expected short-run inflation, $\pi_{SR} = \pi^e_{SR}$, while the dashed black line corresponds to a lower realization, $\pi_{SR} < \pi^e_{SR}$. At $t + d$, the lower inflation realization leaves the log real wage deviation above the desired markup, so DNWR binds. From $t + d$ to $t + 2d$, both paths decline at a rate equal to π^e_{LR}/d, reflecting expected long-run inflation. The shaded areas indicate the deviations of the real wage from the profit-maximizing level over the wage spell. The label $x^*_{i,t}$ is purely illustrative: it only indicates that the optimal markup under DNWR lies below the benchmark markup $\pi^e_{SR}/2$, not its quantitative magnitude.

optimally adjusting its wage. Consequently, the firm enters the next wage spell with an excessively high real wage, generating higher expected losses.

Anticipating this risk, the firm lowers the wage markup in period t, choosing $x^*_{i,t} < \pi^e_{SR}/2$.[8] This reduces the probability that low realized inflation leaves the firm with a binding DNWR constraint at the next adjustment opportunity. Aggregating these lower wage adjustments across firms yields a pass-through from expected inflation into wage inflation strictly below 1/2.

To make the mechanism behind Proposition 1.2 formally transparent, set $d = 1$, so that aggregate wage growth equals individual wage growth, and assume $\Pi(z) = -\frac{1}{2}z^2$. Then, aggregate equilibrium wage growth is characterized by the unique solution to

$$x^*_t = \frac{1}{2}\pi^e_{t,t+1} - \delta F_{SR}\left(x^*_t - x^*_{t+1}\right)\mathbb{E}_{SR}\left[V'_{t+1}(x^*_t - \pi_{SR}) \mid \pi_{SR} < x^*_t - x^*_{t+1}\right] < \frac{1}{2}\pi^e_{t,t+1}.$$

The aggregate wage growth thus equals the benchmark term $\pi^e_{t,t+1}/2$ from Proposition 1.1 minus a DNWR adjustment term. This adjustment term combines the probability that realized short-run inflation is low enough for DNWR to bind with the expected

[8] The same logic applies to the long-run problem: the optimal markup in period $t + d$ satisfies $x^*_{i,t+d} < \pi^e_{LR}/2$. For illustrative clarity, Figure 1.2 does not show this long-run DNWR adjustment.

marginal continuation loss conditional on that event. Hence, DNWR lowers the pass-through because of both the likelihood that DNWR binds at the next wage adjustment opportunity and the additional continuation losses implied when it binds.

Two distributional parameters are central to the probability that DNWR binds and thus to the size of the DNWR adjustment term. First, the variance of the short-run inflation distribution, which we interpret as uncertainty about short-run inflation, directly affects the probability that the short-run inflation realization will be sufficiently low for DNWR to bind. Second, the expected level of long-run inflation determines the future optimal markup, $x^*_{i,t+d}$, and thus the threshold for the short-run inflation realization, $x^*_{i,t} - x^*_{i,t+d}$, below which DNWR binds. Conditional on the short-run inflation distribution, this threshold determines the probability that DNWR binds. We next study the comparative statics with respect to these parameters in more detail.

1.3.3. Pass-Through and Short-Run Inflation Uncertainty

To derive the comparative statics result with respect to short-run inflation uncertainty, we impose one technical restriction on the equilibrium long-run wage markup:

$$x^*_{i,t+d} \geq -\frac{\pi^e_{SR}}{2}. \tag{1.2}$$

Together with $x^*_{i,t} < \pi^e_{SR}/2$ from Proposition 1.2, this restriction provides a conservative sufficient condition for the equilibrium DNWR threshold, $x^*_{i,t} - x^*_{i,t+d}$, to lie below the mean of the short-run inflation distribution. This threshold condition is used in the proof of Proposition 1.3. Appendix A.3.1 derives a sufficient condition for inequality (1.2) in terms of the distributional parameters.

Proposition 1.3 shows that higher short-run inflation uncertainty decreases the pass-through from expected inflation into wage inflation.

Proposition 1.3 *The pass-through from expected inflation into wage inflation is strictly decreasing in short-run inflation uncertainty. Formally, for any $\bar{\sigma}_{SR} > \sigma_{SR}$,*

$$\phi^*(\bar{\sigma}_{SR}) < \phi^*(\sigma_{SR}).$$

Proof. See Appendix A.2.3. □

Consider an increase in short-run inflation uncertainty in period t, meaning that the variance of the short-run inflation distribution rises. For a given wage choice $x_{i,t} < \pi^e_{SR}/2$ of a firm i that can adjust its wage in period t, the higher uncertainty raises the probability that realized short-run inflation will be sufficiently low for DNWR to bind in period $t+d$. In response, the firm lowers its wage growth in period t to counteract the increased risk

of being unable to readjust wages optimally in period $t + d$. In the aggregate, this lowers the pass-through from expected inflation into wage inflation.

In addition, higher short-run inflation uncertainty raises expected continuation losses conditional on the DNWR constraint being binding. Specifically, greater dispersion of the inflation distribution shifts the inflation realizations for which DNWR binds toward lower values. Among these realizations, the inherited real wage therefore tends to lie farther above the wage the firm would ideally like to implement in period $t + d$, raising continuation losses. Hence, DNWR becomes not only more likely to bind when uncertainty increases, but also more costly when it binds, which further contributes to the decrease in equilibrium pass-through.

1.3.4. Pass-Through and Expected Long-Run Inflation

Proposition 1.4 shows that a higher level of expected long-run inflation increases the pass-through from expected inflation into wage inflation.

Proposition 1.4 *The pass-through from expected inflation into wage inflation is strictly increasing in the level of expected long-run inflation. Formally, for any* $\bar{\pi}^e_{LR} > \pi^e_{LR}$,

$$\phi^*(\bar{\pi}^e_{LR}) > \phi^*(\pi^e_{LR}).$$

Proof. See Appendix A.2.4. □

Consider an increase in the level of expected long-run inflation. A firm i that can adjust its nominal wage in period t then anticipates a higher optimal markup in period $t + d$. For a given wage-setting decision $x_{i,t}$ in period t, the threshold $x_{i,t} - x^*_{i,t+d}$ below which short-run inflation must fall for DNWR to bind therefore decreases. Thus, the probability of a binding DNWR constraint falls for any period-t wage-setting decision.

Anticipating this lower risk of binding DNWR, the firm sets a higher wage markup in period t. This moves the firm closer to the benchmark wage-setting decision from Proposition 1.1, $x_{i,t} = \pi^e_{SR}/2$, which reduces expected losses during the wage spell from period t to period $t + d$. In the aggregate, this higher firm-level wage growth increases the equilibrium pass-through from expected inflation into wage inflation.

The dynamic structure that DNWR imposes on the wage-setting problem thus makes long-run inflation expectations relevant for current wage setting decisions. The particular mechanism derived in our framework provides a new theoretical rationale for why low, well-anchored long-run inflation expectations — that is, long-run expectations that are largely insensitive to variations in short-run inflation expectations — can contribute to achieving price stability (e.g., Armantier et al., 2022; Bernanke, 2007; Gürkaynak et al., 2010; Mishkin, 2007). If expectations become de-anchored so that high short-run inflation

expectations are accompanied by high long-run inflation expectations, this weakens the relevance of future DNWR constraints: firms anticipate higher desired markups at future adjustment opportunities, which lowers the probability that DNWR will bind. As a result, firms increase current wages more strongly. Since wages are an important determinant of firms' production costs and thus prices, this higher pass-through may contribute to the emergence of inflationary wage-price dynamics. Conversely, if long-run expectations remain anchored at a low level, DNWR remains a relevant constraint for future wage setting. Firms therefore increase current wages more cautiously, which dampens the pass-through from short-run inflation expectations into current wage growth and may thus reduce the risk of the emergence of self-reinforcing wage and price dynamics.

1.4. Evidence

This section provides a descriptive empirical assessment of the model's predictions. Specifically, we use firm-level survey data to examine whether the empirical pass-through from expected inflation into wage growth varies in the way predicted by the model: whether it is lower at higher levels of uncertainty about short-run inflation, and whether it is higher at higher levels of long-run inflation expectations.

1.4.1. Data and Empirical Specification

This subsection, which describes the data and empirical specification, closely follows Buchheim et al. (2024) and, in parts, uses their exact wording. See their paper for a more detailed description of the survey design, sample, data quality, and empirical specifications.

Survey Setting and Sample. Our empirical analysis uses quarterly surveys of a German firm panel. The quarterly panel was initiated by Link et al. (2026) in December 2020, and data on expected inflation have been available since then. Starting in December 2021, Buchheim et al. (2024) integrated wage-related questions into each quarterly survey wave until March 2026, which are specifically designed to measure the pass-through from expected inflation into wages.

The sample period covers a phase of unusually large variation in inflation. German CPI inflation exceeded the ECB's inflation target of 2 percent in the summer of 2021 and stood at about 5 percent at the start of our sample period. Inflation accelerated further after Russia's invasion of Ukraine in early 2022, reaching levels around 10 percent by the end of 2022, before gradually returning closer to pre-pandemic levels toward the end of the sample period.

The quarterly firm survey is conducted among firms participating in the ifo Business Survey (IBS), a long-standing monthly survey of a large and representative panel of Ger-

man firms.[9] On average, approximately 2,900 respondents participated in each survey wave. Given the focus of our analysis on wage setting, we exclude firms with fewer than 5 employees. We restrict the analysis to respondents who participated in at least two waves and provided information on all variables used in the regressions. This leaves an average of about 2,400 observations per wave.

The median firm has 49 employees. Respondents typically hold senior positions within their firms. In particular, 80 percent report having "very high influence" on their firms' decisions, which is consistent with evidence from Sauer et al. (2023) that IBS respondents usually hold upper management positions such as owner, CEO, or department head.[10]

Survey Design and Main Variables The empirical analysis requires survey measures that correspond to the main objects in the model: one-period-ahead inflation expectations, uncertainty about inflation over the same horizon, wage growth, and long-run inflation expectations. The survey provides empirical counterparts to each of these variables.

In each survey wave, firms report their expected CPI inflation over the next 12 months. This measure provides the empirical counterpart to one-period-ahead inflation expectations in the model. Firms appear to provide meaningful information about their expectations: average expected CPI inflation in our sample closely tracks the developments reported in the Bundesbank's representative firm survey; see Figure A.1 in Appendix A.4. This pattern suggests that the inflation expectation data are of high quality.

The survey also elicits firms' expected gross wage growth over the next 12 months. To ensure that firms have a specific employee in mind when answering the wage question, firms are asked to consider the gross wage growth of a "typical" employee with average qualifications, tenure, and job profile. To avoid measuring only wage changes that have already been agreed upon at the time of the survey, the survey proceeds in two steps. First, firms are asked about the part of future gross wage growth that has already been agreed upon, for example through currently valid collective bargaining agreements. Second, they are asked for total expected wage growth over the next 12 months, including both pre-agreed changes and anticipated outcomes from potential additional wage negotiations or new collective agreements. This measure appears to be accurate and informative about firms' actual wage-setting behavior: expected wage growth elicited between December 2021 and September 2024 is strongly correlated with realized wage growth reported by

[9] The IBS provides the basis for the ifo Business Climate Index, the most recognized leading indicator of the German business cycle. See Sauer et al. (2023) for details on the IBS and its representativeness. The IBS microdata have been used extensively in previous research in economics (e.g., Bachmann et al., 2013, 2019, 2021; Buchheim et al., 2022; Link et al., 2023).

[10] In the June 2021 wave, Link et al. (2026) elicited the respondent's influence on the firm's decisions regarding investment, production, personnel, and pricing on a 5-point scale. 65 percent of managers report having "very high influence" (highest category) on personnel decisions in the firm and 77 percent report to have "very high influence" on decisions in at least one of these areas.

the same firms 12 months later; see Figure A.2 in Appendix A.4.

The survey also contains measures that allow us to examine the model's comparative statics predictions. Uncertainty about short-run inflation expectations is measured through respondents' self-reported confidence in their 12-month inflation forecast. Respondents state their confidence on a scale from 1 to 5, where higher values indicate greater confidence and therefore lower subjective uncertainty about future inflation. The survey also elicits longer-run inflation expectations. Up to and including the first quarter of 2023, respondents reported their five-year-ahead inflation expectations. Starting in the second quarter of 2023, they reported their three-year-ahead inflation expectations instead.

The survey contains several additional variables that allow us to account for other possible determinants of wage growth expectations. These include expectations about the unemployment rate, firms' current and expected future business conditions, and firm characteristics such as the number of employees and the export share.

Empirical Specification. To estimate the pass-through from expected inflation into expected wage growth, we estimate the following specification:

$$w^e_{i,t} = \sum_g \phi_g \left(\pi^e_{i,t} \times \mathbb{1}_{\{G_{i,t}=g\}}\right) [+\phi^B \pi_t] + X'_{i,t}\gamma \; [+\delta_t + \alpha_i] + \epsilon_{i,t}. \tag{1.3}$$

Here, $w^e_{i,t}$ denotes firm i's expected gross wage growth over the next 12 months in survey wave t, and $\pi^e_{i,t}$ denotes its expected CPI inflation over the same horizon. The coefficients ϕ_g measure the pass-through from expected inflation into expected wage growth for group g.[11] In the uncertainty analysis, $G_{i,t}$ denotes respondents' confidence in their inflation expectations, so that g indexes the five confidence levels. In the long-run expectations analysis, $G_{i,t}$ indicates whether the firm has low or high long-run inflation expectations. We classify firms' long-run expectations relative to the survey-wave-specific median, so that the comparison is made within the same aggregate inflation environment. Appendix A.4 reports a robustness check using an unconditional median split across all observations.

We include realized CPI inflation over the past 24 months, π_t, as a control for backward-looking wage-setting considerations. The vector $X_{i,t}$ captures additional determinants of wage growth expectations, including the expected unemployment rate, firms' current and expected future business conditions, number of employees, export share, and respondent influence on firm decision making. Some specifications replace the time-invariant components of $X_{i,t}$ with firm fixed effects, α_i, while others add survey-wave fixed effects, δ_t, to absorb aggregate time variation. Since realized inflation varies only at the survey-wave level, it is absorbed by survey-wave fixed effects. We cluster standard errors at the firm level.

[11] We refer to the relationship between expected inflation and expected wage growth as the pass-through to maintain consistency with the model, without claiming a causal interpretation to the estimates.

1.4.2. Results

Pass-Through and Uncertainty about Inflation. Table 1.1 reports pass-through estimates by respondents' confidence in their inflation expectations. In line with the existing literature, the estimated pass-through is substantially below the benchmark value of 1/2 across all specifications and confidence levels: a 1 percentage point increase in expected inflation is associated with an increase in expected wage growth of only 0.06 to 0.22 percentage points.

Table 1.1: Pass-Through and Uncertainty About Inflation

	Expected Wage Growth next 12 Months			
	(1)	(2)	(3)	(4)
Expected Inflation next 12m × Conf=1	0.079***	0.063***	0.125***	0.067***
	(0.014)	(0.012)	(0.019)	(0.015)
Expected Inflation next 12m × Conf=2	0.108***	0.084***	0.155***	0.086***
	(0.012)	(0.010)	(0.018)	(0.014)
Expected Inflation next 12m × Conf=3	0.121***	0.089***	0.164***	0.089***
	(0.012)	(0.010)	(0.018)	(0.014)
Expected Inflation next 12m × Conf=4	0.161***	0.122***	0.204***	0.122***
	(0.013)	(0.011)	(0.018)	(0.015)
Expected Inflation next 12m × Conf=5	0.182***	0.135***	0.221***	0.131***
	(0.027)	(0.024)	(0.031)	(0.026)
Observations	44337	44094	44337	44094
R2	0.101	0.454	0.107	0.458
Firm-specifics	yes	no	yes	no
Firm FE	no	yes	no	yes
Survey Wave FE	no	no	yes	yes
Realized inflation	yes	yes	no	no

Note: This table reports results from regressions of expected wage growth over the next 12 months on expected CPI inflation over the same horizon interacted with respondents' confidence in their inflation expectations. Conf = 1 corresponds to the lowest confidence level and Conf = 5 to the highest confidence level. All specifications control for the expected unemployment rate over the next 12 months and firms' current and expected future business conditions. Columns (1) and (3) additionally include time-invariant firm controls. Columns (2) and (4) replace these controls with firm fixed effects. Columns (3) and (4) include survey-wave fixed effects. Realized CPI inflation over the previous 24 months is included in columns (1) and (2) and absorbed by survey-wave fixed effects in columns (3) and (4). Standard errors are clustered at the firm level. Levels of significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

In all specifications, the pass-through increases monotonically with firms' confidence in their inflation expectations. The magnitude is substantial: across specifications, the pass-through is about twice as high for firms with the highest confidence as for firms

Table 1.2: Pass-Through and Long-Run Inflation Expectations

	Expected Wage Growth next 12 Months			
	(1)	(2)	(3)	(4)
Expected Inflation next 12m×Low LR Expectations	0.078***	0.049***	0.153***	0.072***
	(0.014)	(0.013)	(0.023)	(0.018)
Expected Inflation next 12m×High LR Expectations	0.101***	0.047***	0.162***	0.071***
	(0.014)	(0.012)	(0.021)	(0.016)
Observations	35816	35298	35816	35298
R2	0.101	0.478	0.107	0.482
Firm-specifics	yes	no	yes	no
Firm FE	no	yes	no	yes
Survey Wave FE	no	no	yes	yes
Realized inflation	yes	yes	no	no

Note: This table reports results from regressions of expected wage growth over the next 12 months on expected CPI inflation over the same horizon interacted with indicators for low and high long-run inflation expectations. Firms are classified as having high long-run inflation expectations if their long-run inflation expectation lies above the survey-wave-specific median. All specifications control for the expected unemployment rate over the next 12 months and firms' current and expected future business conditions. Columns (1) and (3) additionally include time-invariant firm controls. Columns (2) and (4) replace these controls with firm fixed effects. Columns (3) and (4) include survey-wave fixed effects. Realized CPI inflation over the previous 24 months is included in columns (1) and (2) and absorbed by survey-wave fixed effects in columns (3) and (4). Standard errors are clustered at the firm level. Levels of significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

with the lowest confidence. Interpreting higher confidence as lower subjective uncertainty about future inflation, this finding supports the model prediction that higher inflation uncertainty reduces the pass-through from expected inflation into wage growth.

Pass-Through and Long-Run Inflation Expectations. Table 1.2 reports pass-through estimates separately for firms with low and high long-run inflation expectations, defined relative to the survey-wave-specific median. The estimates provide some evidence in the direction predicted by the model. In specifications without firm fixed effects, the pass-through is higher for firms with high long-run inflation expectations, with pass-through differences of 0.024 and 0.009 percentage points in columns (1) and (3), respectively. While small in magnitude, the difference is statistically significant at the 5% level in column (1). Once firm fixed effects are included, however, the difference becomes essentially zero. Table A.1 in Appendix A.4 shows similar patterns when classifying firms by the unconditional median of long-run inflation expectations across all observations. Under this split, the estimates are slightly more favorable to the model prediction, especially in the specification with survey-wave fixed effects but without firm fixed effects, where the difference is larger than in Table 1.2 and statistically significant at the 10%

level. However, the differences remain small and again become essentially zero once firm fixed effects are included. Thus, we interpret the overall results as providing some support for the model prediction that higher long-run inflation expectations raise the pass-through from expected inflation into current wage growth, though this interpretation should be treated with caution because the differences are small and not robust across specifications.

1.5. Conclusion

This paper provides a theoretical rationale for the empirically small effect of inflation expectations on wage inflation. We develop a Taylor-type wage-setting model with inflation uncertainty and DNWR. When inflation turns out lower than expected, DNWR can leave firms with excessively high real wages that cannot be corrected through nominal wage cuts. Anticipating this risk, firms raise wages more cautiously, which lowers the pass-through from expected inflation into wage inflation. The model predicts that higher short-run inflation uncertainty and lower long-run inflation expectations reduce the pass-through because both increase the risk that DNWR binds at future adjustment dates. Descriptive firm-level evidence provides support for the uncertainty channel and more limited support for the long-run expectations channel.

Two avenues for future research seem promising. First, the model is deliberately stylized and partial equilibrium in order to isolate the mechanism of interest. Embedding the mechanism into a general equilibrium framework could quantify its relevance for the pass-through and its implications for interdependent wage-price dynamics. Second, future empirical work could test the mechanism more directly by exploiting heterogeneity in firms' exposure to DNWR. For example, firms that rely more heavily on flexible pay components, such as bonuses, may face weaker constraints on downward adjustments of effective compensation. Linking such heterogeneity to firms' subjective inflation distributions and wage-setting decisions could provide a more direct test of the channel identified in this paper.

CHAPTER 2

Overconfidence and the Gender Application Gap

This chapter is co-authored with Vincent Selz.

Abstract

Women apply to lower-paying jobs than men (Fluchtmann et al., 2024; Lochner and Merkl, 2026). The mechanisms behind this gender application gap remain underexplored. This paper studies whether gender differences in overconfidence contribute to this gap. We develop a directed search model and identify two channels through which overconfidence affects application behavior. First, overconfident job seekers overestimate their hiring chances at high-wage vacancies and apply to them more often. Second, when job seekers believe that overconfidence is more prevalent among others, they expect higher competition for high-wage vacancies. This discourages applications there, especially among non-overconfident job seekers. Higher overconfidence among men therefore generates an application gap through both a direct and an indirect channel. We test these predictions in a laboratory experiment. The results support the direct channel, but not the indirect channel.

2.1. Introduction

Women apply to lower-paying jobs than men (Fluchtmann et al., 2024; Lochner and Merkl, 2026). This gender application gap accounts for more than 70% of the gender wage gap among newly hired workers with similar observable characteristics (Fluchtmann et al., 2024). Hence, substantial gender differences in wages arise already at the application stage, preceding differences in hiring (Goldin and Rouse, 2000), bargaining (Kiessling et al., 2024), or promotions (Hospido et al., 2022). Understanding the gender application gap is therefore central to understanding the gender wage gap. Despite its importance, however, the mechanisms underlying this gap remain largely unexplored.

In this paper, we study whether higher overconfidence among men about their relative ability contributes to the gender application gap. Prior work documents such gender differences in overconfidence (e.g. Bandiera et al., 2022; Barber and Odean, 2001; Cortés et al., 2023; Niederle and Vesterlund, 2007). We develop a directed search model that links the differences in overconfidence to application behavior and test its predictions in a laboratory experiment.

The model describes competition among job seekers for vacancies that differ in wages. The probability of receiving an offer increases in the ability of a job seeker relative to the other applicants. If job seekers overestimate their relative ability, that is, if they are overconfident, they therefore assign a higher probability to receiving an offer from a high-wage vacancy. Hence, they apply more frequently to high-wage jobs (*direct effect*), increasing competition for these vacancies. If job seekers believe that others are overconfident, the associated expected increase in competition discourages them from applying to high-wage vacancies (*indirect effect*).[1] However, overconfident job seekers are less sensitive to this indirect effect because they still expect that they can compete successfully for high-wage vacancies. Through these direct and indirect effects, higher overconfidence among men generates an application gap.

To test the two channels empirically, we design a laboratory experiment that closely follows the spirit of the model. Participants first work on a cognitive ability task, after which we elicit incentivized beliefs about own relative ability and about the relative ability beliefs of others, providing measures of overconfidence and believed overconfidence.[2] Pairs of participants subsequently enter a stylized labor market in which they compete against each other for two different vacancies by allocating application weights between the vacancies. As in the model, vacancies differ only in their posted wages, and higher

[1] In the following, we use the term *believed overconfidence* to refer to these second-order beliefs about the overconfidence among others.

[2] Our notion of overconfidence corresponds to overplacement in standard psychological terminology; see Moore and Healy (2008).

relative ability increases the probability of receiving an offer when applying. In control pairs, neither participant receives any information about their own actual relative ability. In treated pairs, one participant receives precise information about her own actual relative ability, which plausibly removes overconfidence.[3] The other participant receives no information about his own ability but knows that the competitor received this information, which plausibly eliminates beliefs that the competitor is overconfident. This design allows us to identify the direct and indirect effects of overconfidence and their implications for the application gap.

Through the lens of the model, two empirical conditions must hold for the mechanisms of interest to contribute to the application gap. First, for the direct channel, men must be more overconfident than women. Second, for the indirect channel, both men and women must additionally believe that others are overconfident, with no significant gender differences in the level of this believed overconfidence.[4] If these conditions hold, we can examine whether overconfidence increases the application weight allocated to the high-wage vacancy, whether believed overconfidence among others decreases this weight, and how much these mechanisms contribute to the application gap.

Consistent with field evidence (Fluchtmann et al., 2024; Lochner and Merkl, 2026), we find an application gap in the control group: men allocate about 15% more application weight to the high-wage vacancy than women. The data also support both conditions for our mechanisms of interest to contribute to this gap. First, in line with existing evidence, men are more overconfident than women, placing themselves about five percentiles higher in the ability distribution. Second, both women and men believe that others are overconfident on average, and these beliefs do not differ significantly by gender.

Next, we turn to the causal effects of overconfidence and believed overconfidence on application behavior. To test the direct effect, we compare the application decisions of participants who receive information about own relative ability to the decisions of participants in control pairs. This comparison provides clear evidence in favor of the direct effect of overconfidence: overconfidence increases the application weight allocated to the high-wage vacancy. The effect is sizable. On average, overconfidence raises the application weight by about 14%. To test the indirect effect, we compare participants whose competitor receives information about own ability to participants in control pairs. This

[3]Participants may also underestimate their relative ability and thus exhibit underconfidence. For simplicity and because this case plays a minor role empirically, we refer to underconfidence as negative overconfidence.

[4]Following the model logic, believed overconfidence can contribute to the application gap as long as women are less overconfident than men and their level of believed overconfidence is at least as high as that of men. Equality in the level of believed overconfidence therefore suffices. To our knowledge, there is no evidence on gender differences in believed overconfidence. We therefore have no ex ante reason to expect gender differences in believed overconfidence and focus on the equal-beliefs case.

comparison yields no evidence in favor of the indirect effect.

Finally, we quantify by how much gender differences in overconfidence contribute to the gender application gap through the direct effect. Comparing the application gap in the control group to the gap among informed participants shows that gender differences in overconfidence, via the direct effect, substantially contribute to the application gap in our stylized labor market, accounting for at least one third of the gap.

Our work relates to three main strands of the literature. First, we contribute to the literature on gender differences in job search behavior. This literature documents that women have lower reservation wages (Faberman et al., 2025; Le Barbanchon et al., 2021; McGee and McGee, 2025), start searching earlier (Cortés et al., 2023), and exert more search effort (Faberman et al., 2025). These patterns relate to differences in preferences for non-wage amenities, commuting, and mobility (Faberman et al., 2025; Le Barbanchon et al., 2021; Maestas et al., 2023), as well as to lower risk aversion and higher overconfidence among men (Cortés et al., 2023). To date, only Fluchtmann et al. (2024) and Lochner and Merkl (2026) examine gender differences in job application behavior and show that women target lower-paying jobs than men. Both studies provide some evidence on underlying mechanisms, including differences in targeted non-wage amenities and flexibility requirements. Nevertheless, a substantial share of the gender application gap remains unexplained. We contribute to this recent literature by examining whether gender differences in overconfidence can explain part of the application gap.

Second, we contribute to the literature on directed job search (Wright et al., 2021, provide an extensive overview about this type of models). We set up a directed job search model that incorporates overconfidence, second-order beliefs about overconfidence, and gender differences in overconfidence to study how these features affect individual and aggregate application outcomes from a theoretical perspective. Beyond this, we also add to the directed search literature empirically. Empirical tests of directed search models remain scarce due to data limitations, most notably the difficulty of observing vacancies that are comparable except for posted wages (Faberman and Menzio, 2018). Experimental settings are therefore particularly well suited to test predictions of directed job search models. To our knowledge, only Kloosterman (2016) in the laboratory and Belot et al. (2022) and He et al. (2022) in the field provide such tests. These studies validate the key prediction of directed job search models that vacancies with higher posted wages attract more applicants.[5] Our experimental design closely mirrors our theoretical setup, which allows us to directly test the implications of gender differences in overconfidence for the

[5]Related experimental work tests directed search models in other settings, such as buyer-seller markets with capacity constraints; see, for example, Cason and Noussair (2007) and Helland et al. (2017). These studies also provide support for core predictions of directed search models.

application gap as predicted by our theoretical framework.

Third, we contribute to the behavioral job search literature that studies how psychological traits and biases shape search behavior and outcomes, mainly through their effects on reservation wages. Prior work studies a range of psychological traits and biases in job search, including, for example, time preferences (DellaVigna and Paserman, 2005; Paserman, 2008), reference-dependent preferences (DellaVigna et al., 2017; Koenig et al., 2024), locus of control (Caliendo et al., 2015; McGee and McGee, 2016), overoptimism about future job prospects (Cortés et al., 2023; Mueller et al., 2021; Spinnewijn, 2015), and anchoring of beliefs about outside options to current wages (Jäger et al., 2024). We add to this literature by showing that gender specific differences in one trait — overconfidence — may substantially contribute to gender differences in application behavior.

2.2. Model

We develop a directed job search model to study the role of overconfidence for application behavior. We first abstract from gender and characterize the implications of overconfidence and believed overconfidence for application behavior. We then impose assumptions on how overconfidence is distributed across men and women, which allows us to derive implications for the gender application gap.

The model features two types of vacancies that differ only in their posted wages (low or high), and two types of job seekers that differ only in their abilities, (low or high). A subset of job seekers with low ability is overconfident and, hence, mistakenly believes to have high ability, while all other job seekers hold correct beliefs about own ability and are referred to as unbiased.[6] Each job seeker submits a single application to one vacancy. An application is a noisy signal of the true ability of the applicant. Then, for each vacancy the applicant with the best signal is hired. Because the signal is positively correlated with ability, high-ability applicants have a higher probability of being hired when both types apply to the same vacancy. Beliefs about own ability therefore shape beliefs about hiring probabilities, henceforth referred to as subjective hiring probabilities. In addition, believed overconfidence affects the competition that a job seeker expects to face at each vacancy and thus also affects subjective hiring probabilities. As a result, beliefs about own ability and believed overconfidence jointly determine optimal application decisions, and systematic differences in these beliefs between men and women imply systematic differences in application behavior.

[6]There is no underconfidence in our model. This is motivated by evidence from the literature showing that overconfidence is more prevalent than underconfidence for both men and women (Barber and Odean, 2001; Cortés et al., 2023; Niederle and Vesterlund, 2007).

2.2.1. Setup

Job Seekers and Vacancies. There is a continuum of job seekers and vacancies, each of measure two. All job seekers are risk neutral. A job seeker i is either of low ability, $\alpha(i) = \alpha_l$, or high ability, $\alpha(i) = \alpha_h > \alpha_l$. Half of all job seekers have high ability and half have low ability. Each vacancy j offers either a low wage, $\omega(j) = w_l$, or a high wage, $\omega(j) = w_h > w_l$ with half of all vacancies offering the high wage and half offering the low wage. Wages are not conditioned on the ability of the hired job seeker. Hence, the wage that is offered by a vacancy always corresponds to the wage that the job pays to the job seeker who is hired.

Each job seeker i holds a belief about its own ability, $\hat{\alpha}(i)$. We refer to this belief of a job seeker as its subjective ability. We denote a job seeker as unbiased if $\hat{\alpha}(i) = \alpha(i)$. All high-ability job seekers are unbiased. A low-ability job seeker is overconfident if $\hat{\alpha}(i) = \alpha_h > \alpha_l = \alpha(i)$. A share $x \in (0, 1)$ of all low-ability job seekers is overconfident. Each job seeker i holds a belief $\hat{x}(i)$ about this share, henceforth the subjective share. This subjective share captures believed overconfidence in the model.[7] We make the following assumption on the subjective share:

Assumption 2.1 *All job seekers hold the same belief $\hat{x}(i) = \hat{x} \in (0, 1)$ about the share of overconfident job seekers among low-ability job seekers. It is common knowledge among all job seekers that everyone has this belief.*

Assumption 2.1 makes the analysis tractable in two ways. First, it implies that all job seekers with the same subjective ability follow the same application strategy, which simplifies the derivation of the aggregate equilibrium. Second, it eliminates the need to specify beliefs beyond the second-order belief $\hat{x}$. Optimal application behavior of a job seeker i depends on her belief $\hat{x}(i)$ and on the expected application behavior of others, which in turn depends on beliefs about the beliefs of others regarding the share of overconfident job seekers, and so forth. Without Assumption 2.1, the analysis would therefore require specifying beliefs at all orders, making the model intractable.[8]

Applications and Hiring. Each job seeker sends out one application to a single vacancy. Following the directed search literature, we assume that job seekers' strategies are anonymous: they treat all vacancies of the same type identically (Galenianos and Kircher,

[7] We keep x and $\hat{x}(i)$ distinct to isolate the effect of believed overconfidence on application behavior. If the two instead co-moved, e.g., by imposing $\hat{x}(i) = x$, changes in believed overconfidence would be tied to changes in the actual prevalence of overconfidence, thereby mixing the implications of the belief channel for the gender application gap with those of changes in the actual composition of job seekers.

[8] For the experiment, we use an empirically relevant and testable counterpart of this rather technical model assumption, namely that men and women believe that others are overconfident on average, without significant gender differences in the average level of this belief.

2009; Kircher, 2009). This implies that a job seeker chooses a wage and then randomly applies to a vacancy that offers this wage. The anonymity assumption constitutes the key friction of directed search models and rules out trivial equilibria in which every job seeker applies to a different vacancy and is hired with a probability of one. Instead, some vacancies receive multiple applications, leading to competition for these vacancies, whereas others receive none.

To keep our labor market environment general, we represent the hiring preferences of firms through a simple rule: each firm strictly prefers hiring a high-ability applicant over a low-ability one. This assumption is natural since both types receive the posted wage of the respective vacancy and therefore impose identical costs on the firm. Yet, firms cannot perfectly observe the true ability of an applicant. Instead, each firm only receives a signal

$$s_i = \alpha(i) + \epsilon_i, \text{ where } \epsilon_i \sim \text{Gumbel}(0, \beta).$$

For this setup, Lemma 2.1 characterizes the optimal hiring behavior of firms and the resulting hiring probabilities. See Weibull et al. (2007) for details on the derivation.

Lemma 2.1 (Weibull, Mattsson and Voorneveld, 2007) *The optimal strategy of any firm j is to hire the applicant in its applicant pool P_j with the highest signal, i.e., the applicant $i \in P_j$ such that $s_i > s_k$ for all $k \in P_j$, $k \neq i$. The resulting hiring probability of applicant $i \in P_j$ is*

$$s(P_j, i) = \frac{\exp\left(\frac{\alpha(i)}{\beta}\right)}{\sum_{\ell \in P_j} \exp\left(\frac{\alpha(\ell)}{\beta}\right)}. \tag{2.1}$$

This formulation yields three key comparative statics. First, a high-ability applicant has a higher hiring probability than a low-ability applicant. Second, the hiring probability of an applicant declines as the number of competitors for the same vacancy increases. Third, the probability decreases as the abilities of competing applicants rise.

The parameter β governs how noisy the signals are. As $\beta \to 0$, a signal perfectly reveals the underlying true ability of an applicant. Then, a high-ability applicant i is hired with probability $s(P_j, i) = \frac{1}{n_H}$, where n_H denotes the number of high-ability applicants in P_j, and low-ability applicants are never hired in pools that also contain high-ability applicants. As $\beta \to \infty$, signals carry no information about underlying abilities, so the firm effectively randomizes uniformly over all applicants and $s(P_j, i) = \frac{1}{n_H + n_L}$ irrespective of the ability of applicant i.

Equilibrium. We follow the directed search literature by restricting attention to symmetric equilibria. In these equilibria, all job seekers of a given subjective ability are equally likely to apply to a vacancy of a given wage. The intuitive justification for restricting to

symmetric strategies is that asymmetric strategies would require extensive coordination among job seekers, since each one would need to know its own individual strategy, making such an equilibrium a less reasonable outcome in large markets (Galenianos and Kircher, 2009; Wright et al., 2021).

Denote by $\lambda_h \in [0,1]$ the probability with which a high-ability job seeker applies for a high-wage vacancy. Because overconfident low-ability job seekers believe that they have high ability, they also adopt this application strategy. Denote by $\lambda_l \in [0,1]$ the corresponding probability for an unbiased low-ability job seeker. In a continuum economy with infinitely many job seekers and vacancies, these probabilities together with the subjective share of overconfident job seekers, $\hat{x}$, determine what we refer to as the subjective mass of job seekers that applies for a given vacancy type. Specifically, denote by $\hat{\delta}_a^w$ the subjective mass of job seekers of type $a \in \{h,l\}$ applying for a vacancy of type $w \in \{h,l\}$. Then,

$$\begin{aligned} \hat{\delta}_h^h(\lambda_h) &= \lambda_h, \\ \hat{\delta}_h^l(\lambda_h) &= (1-\lambda_h), \\ \hat{\delta}_l^h(\lambda_h,\lambda_l) &= \hat{x}\lambda_h + (1-\hat{x})\lambda_l, \\ \hat{\delta}_l^l(\lambda_h,\lambda_l) &= \hat{x}(1-\lambda_h) + (1-\hat{x})(1-\lambda_l). \end{aligned} \tag{2.2}$$

The subjective applicant masses of high-ability job seekers that apply for the high- and low-wage vacancies correspond to the associated application probabilities, λ_h and $(1-\lambda_h)$ respectively. For low-ability job seekers, a subjective share $\hat{x}$ is overconfident and therefore adopts the application strategy of high-ability job seekers, resulting in a subjective applicant mass of $\hat{x}\lambda_h$ for high-wage vacancies and $\hat{x}(1-\lambda_h)$ for low-wage vacancies. A subjective share $(1-\hat{x})$ of low-ability job seekers follows the strategy of unbiased low-ability job seekers, resulting in a subjective applicant mass of $(1-\hat{x})\lambda_l$ for high-wage vacancies and $(1-\hat{x})(1-\lambda_l)$ for low-wage vacancies.[9]

Denote by $p_h(\hat{\delta}_h^w, \hat{\delta}_l^w)$ the subjective hiring probability of a job seeker who believes to have high ability, and by $p_l(\hat{\delta}_h^w, \hat{\delta}_l^w)$ the subjective hiring probability of an unbiased low-ability job seeker, when applying for a vacancy of wage type w, conditional on $(\hat{\delta}_h^w, \hat{\delta}_l^w)$ being the subjective masses of high- and low-ability applicants that also apply for a vacancy of the same wage.[10] A strategy pair $(\lambda_h^*, \lambda_l^*) \in [0,1]^2$ constitutes a symmetric equilibrium if λ_h^* is a best response for job seekers who believe they have high ability and

[9] We only need the subjective masses from (2.2) to derive the equilibrium application strategies of the different types of job seekers. To derive actual aggregate application behavior, however, we also require the actual masses of job seekers of each type that apply for the different types of vacancies. These are obtained by replacing the subjective share $\hat{x}$ of overconfident job seekers in (2.2) with the actual share x.

[10] The subjective hiring probabilities combine the selection probability $s(P_j, i)$ from Lemma 2.1 with the distribution over all possible applicant pools implied by anonymous strategies. Specifically, the number of competing applicants of each type follows a Poisson distribution with parameters given by the subjective masses $(\hat{\delta}_h^w, \hat{\delta}_l^w)$. See Lemma B.2 in Appendix B.1 for closed-form expressions of the probabilities.

λ_l^* is a best response for unbiased low-ability job seekers, given the subjective applicant masses implied by $(\lambda_h^*, \lambda_l^*)$ and $\hat{x}$. In a fully mixed equilibrium, $(\lambda_h^*, \lambda_l^*) \in (0,1)^2$, all job seekers must be indifferent between applying for high- and low-wage vacancies, so that

$$\begin{aligned} p_h\big(\hat{\delta}_h^h(\lambda_h^*),\ \hat{\delta}_l^h(\lambda_h^*, \lambda_l^*)\big) w_h &= p_h\big(\hat{\delta}_h^l(\lambda_h^*),\ \hat{\delta}_l^l(\lambda_h^*, \lambda_l^*)\big) w_l, \\ p_l\big(\hat{\delta}_h^h(\lambda_h^*),\ \hat{\delta}_l^h(\lambda_h^*, \lambda_l^*)\big) w_h &= p_l\big(\hat{\delta}_h^l(\lambda_h^*),\ \hat{\delta}_l^l(\lambda_h^*, \lambda_l^*)\big) w_l. \end{aligned} \tag{2.3}$$

If λ_h^* or λ_l^* is at the boundary, the corresponding indifference condition is replaced by the appropriate weak inequality.

2.2.2. The Direct and Indirect Effects of Overconfidence

Before characterizing the effects of overconfidence on application behavior, we first note that a symmetric equilibrium always exists.

Lemma 2.2 *There always exists a symmetric equilibrium.*

Proof. See Appendix B.2.1.

We first isolate how overconfidence affects application behavior independently of gender. We distinguish between the *direct effect*, which operates through beliefs about own ability, and the *indirect effect*, which operates through beliefs about the overconfidence of others. In the next step, we show how a gender application gap emerges in our model if overconfident job seekers are predominantly male.

We define the direct effect as the change in the probability that a low-ability job seeker applies for a high-wage vacancy when the job seeker is overconfident rather than unbiased:

$$D \coloneqq \lambda_h^* - \lambda_l^*.$$

Proposition 2.1 shows that job seekers who are overconfident have a weakly higher probability of applying to high-wage vacancies than unbiased low-ability job seekers. The difference is strictly positive unless all types of job seekers apply with probability one to high-wage vacancies.

Proposition 2.1 *In any symmetric equilibrium* $(\lambda_h^*, \lambda_l^*) \in [0,1]^2$,

$$D = \lambda_h^* - \lambda_l^* \geq 0,$$

with equality if and only if $\lambda_h^* = \lambda_l^* = 1$.

Proof. See Appendix B.2.2. □

Two main interdependent mechanisms give rise to this result. First, job seekers who perceive themselves as high ability assign themselves higher hiring probabilities than job

seekers who perceive themselves as low ability. This relative hiring advantage increases with competition, since additional competitors reduce the hiring prospects of low-ability applicants more strongly than those of high-ability applicants. Hence, because the high-wage market must be associated with lower hiring probabilities in equilibrium due to its wage premium, this competition makes the high-wage market relatively more attractive to job seekers who perceive themselves as high ability than to unbiased low-ability job seekers. Second, in the case of fully mixed equilibria, the requirement that both types are indifferent between the two markets dictates the composition of the applicant pools. While the hiring probabilities must be lower in the high-wage market, the relative hiring advantage of high-ability applicants must be the same across markets. Along a fixed relative hiring advantage, lower hiring probabilities can arise only from more high-ability and fewer low-ability applicants: replacing high-ability competitors with low-ability competitors would reduce effective competition and raise hiring probabilities. Thus, relative to the low-wage market, the high-wage market must contain more high-ability applicants and fewer low-ability applicants.[11]

We define the indirect effect as the change in the probability of applying for a high-wage vacancy when a job seeker believes that the share of overconfident job seekers increases. To study this effect, we restrict attention to continuous equilibrium paths along which application probabilities adjust smoothly as the subjective share $\hat{x}$ increases. We refer to such a path as an equilibrium branch.

Definition 2.1 *Fix $\hat{x}_0 \in (0,1)$ and an equilibrium $(\lambda_h^0, \lambda_l^0)$ at $\hat{x}_0$. An equilibrium branch is a continuous mapping*

$$\hat{x} \mapsto \big(\lambda_h^*(\hat{x}), \lambda_l^*(\hat{x})\big)$$

defined on an interval $I = [\hat{x}_0, \bar{x}) \subseteq (0,1)$ such that

$$(\lambda_h^*(\hat{x}_0), \lambda_l^*(\hat{x}_0)) = (\lambda_h^0, \lambda_l^0),$$

and $(\lambda_h^*(\hat{x}), \lambda_l^*(\hat{x}))$ is an equilibrium for all $\hat{x} \in I$.

We restrict attention to such equilibrium branches for two reasons. First, equilibrium uniqueness is not established: qualitatively different types of equilibria, such as fully mixed and (partially) pure strategy equilibria, may in principle coexist. We therefore avoid arbitrary switching between such potentially coexisting equilibria when $\hat{x}$ increases. Second,

[11] The relevance of these mechanisms differs across equilibrium configurations. Some boundary cases with $D < 0$ are ruled out primarily because job seekers who perceive themselves as high ability would have a sufficiently high probability of obtaining the high-wage job if they applied. The argument that $D > 0$ in any fully mixed equilibrium, by contrast, mainly depends on the restrictions imposed on the composition of applicant masses across the two markets.

equilibrium branches provide a natural equilibrium selection rule when a given type of equilibrium ceases to exist.

Since the indirect effect can only be relevant in equilibria in which overconfidence induces a change in the chosen strategy, that is, in equilibria with $D > 0$, we restrict attention to such initial equilibria.[12] We obtain the following result for the indirect effect along equilibrium branches.

Proposition 2.2 *Consider an equilibrium* $(\lambda_h^*(\hat{x}_0), \lambda_l^*(\hat{x}_0))$ *with* $\lambda_l^*(\hat{x}_0) \in (0,1)$ *for some* $\hat{x}_0 \in (0,1)$*. There exists an equilibrium branch through this equilibrium, which is unique up to restriction of its domain, and a unique* $\hat{x}^* > \hat{x}_0$ *such that the branch is defined at least on* $[\hat{x}_0, \hat{x}^*)$ *and*

$$\frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) < 0, \qquad \frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) = 0 \qquad \text{for every } \hat{x} \in [\hat{x}_0, \hat{x}^*).$$

If the branch extends beyond $\hat{x}^*$*, then for every* $\hat{x}$ *in its domain with* $\hat{x} \geq \hat{x}^*$*, it holds that* $\lambda_l^*(\hat{x}) = 0$ *and* $\lambda_h^*(\hat{x})$ *is weakly decreasing in* $\hat{x}$*. This latter result also applies to an equilibrium branch through an initial equilibrium with* $\lambda_l^*(\hat{x}_0) = 0$*, if such a branch exists.*

Proof. See Appendix B.2.3. □

As a remark, we state weak monotonicity rather than a derivative condition in the latter case because $\lambda_h^*(\hat{x})$ may fail to be differentiable at the point where it moves from the boundary $\{1\}$ to the interior $(0,1)$. In the interior case, the implicit function theorem establishes differentiability and yields the sign of the derivative. At the transition from the boundary to the interior, however, this argument does not apply, so we use a monotonicity statement instead.

Generally, higher believed overconfidence reduces the probability of job seekers to apply for a high-wage vacancy. Through the direct effect, a higher believed share of overconfident job seekers increases expected competition for high-wage vacancies and reduces expected competition for low-wage vacancies. Consequently, job seekers reduce their probability of applying for high-wage vacancies, giving rise to the indirect effect.

Importantly, this adjustment is driven entirely by unbiased low-ability job seekers as long as they apply with positive probability to high-wage vacancies. Along the unique equilibrium branch, their application probability decreases in a way that offsets the increase in expected competition implied by a higher subjective share $\hat{x}$. The total subjective mass of low-ability applicants therefore remains unchanged in both the high- and low-wage markets. Consequently, the environment faced by high-ability and overconfident job seekers

[12] The direct and indirect effects are interdependent. The indirect effect arises from the expected change in application probabilities of overconfident job seekers as implied by the direct effect. Changes in application probabilities induced by the indirect effect feed back into the direct effect itself.

does not change, and neither do their application probabilities, while unbiased low-ability applicants gradually exit the high-wage market. Once they have fully exited that market, further increases in believed overconfidence can no longer be offset through reductions in $\lambda_l^*(\hat{x})$. Only then may $\lambda_h^*(\hat{x})$ also fall, although it does not have to fall because job seekers who perceive themselves as high ability may still assign a sufficiently large probability to competing successfully for high-wage vacancies despite the higher expected competition.[13] We interpret this overall adjustment pattern as reflecting a lower responsiveness of job seekers who perceive themselves as high ability to believed overconfidence.

We thus identify two mechanisms through which overconfidence affects applications to high-wage vacancies. First, through the direct effect, overconfident job seekers are more likely to apply to a high-wage vacancy than unbiased low-ability job seekers. Second, through the indirect effect, higher believed overconfidence reduces the probability of applying to high-wage vacancies, with unbiased low-ability job seekers responding more strongly than job seekers who perceive themselves as high ability. We now introduce gender into the model and link these effects to the emergence of the application gap.

2.2.3. Implications of the Direct and Indirect Effects for the Application Gap

We assume that 50% of all job seekers with low and high ability are female and that a strictly positive share of job seekers with low ability is overconfident, $x > 0$.[14] Moreover, we assume that a majority of these overconfident job seekers is male:

Assumption 2.2 *Let m denote the share of men among all overconfident job seekers. We assume that $m > 1/2$ with $mx \leq 1/2$.*

The restriction $mx \leq 1/2$ is required to ensure that the total mass of male low-ability job seekers is $1/2$.

Denote by W_h and W_l the average probabilities with which women apply for high-wage and low-wage vacancies, respectively, and let M_h and M_l denote the corresponding probabilities for men. Define

$$G_h \coloneqq M_h - W_h,$$
$$G_l \coloneqq M_l - W_l.$$

Empirically, the gender application gap refers to the observation that women allocate a lower share of applications to high-wage vacancies and a higher share to low-wage

[13] In particular, when $\lambda_h^*(\hat{x}) = 1$, job seekers who perceive themselves as high ability need not respond to an increase in $\hat{x}$. See Figure B.1 in Appendix B.3 for a numerical illustration of exemplary equilibrium branches.

[14] Up to this point, no assumption about x was required, as only the subjective share $\hat{x}$ was necessary to derive equilibrium strategies in the previous section.

vacancies than men. In the model, we map this definition into gender differences in average application probabilities. We say that there is a gender application gap if $G_h > 0$ and $G_l < 0$. Because $G_h = -G_l$, we restrict attention to G_h; all results for G_l follow with the opposite sign.

The gender difference in average high-wage application probabilities is proportional to the direct effect: $G_h = x(2m-1)D$. For fixed $x > 0$ and $m > 0.5$, the sign of G_h is thus governed by the sign of D, and its magnitude is increasing in D. Since Proposition 2.1 establishes that the direct effect is strictly positive except for the case $\lambda_h^* = \lambda_l^* = 1$, gender differences in overconfidence generate an application gap through the direct effect.

Corollary 2.1 *The direct effect implies a gender application gap in equilibrium. In particular, since $D \geq 0$, it holds that $G_h \geq 0$, with equality if and only if $\lambda_h^* = \lambda_l^* = 1$. Moreover, the size of the gender application gap is strictly increasing in the direct effect.*

PROOF. See Appendix B.2.4. □

To characterize the implications of the indirect effect for the gender application gap, we distinguish two cases. First, if $\lambda_l^*(\hat{x})$ is decreasing in $\hat{x}$ while $\lambda_h^*(\hat{x})$ remains constant, an increase in believed overconfidence reduces only the probability with which unbiased low-ability job seekers apply to a high-wage vacancy. Since unbiased low-ability job seekers are predominantly female, this widens the gender application gap. Conversely, if $\lambda_l^*(\hat{x})$ remains constant while $\lambda_h^*(\hat{x})$ is decreasing in $\hat{x}$, only high-ability job seekers and overconfident low-ability job seekers reduce the probability with which they apply to high-wage vacancies. Since overconfident job seekers are predominantly male, this narrows the gender application gap.

Together with Proposition 2.2, this implies that as long as unbiased low-ability job seekers apply with positive probability to high-wage vacancies, any increase in believed overconfidence widens the gender application gap along the unique equilibrium branch. Only once all unbiased low-ability job seekers have exited the high-wage market may the gap narrow.

Corollary 2.2 *Consider an equilibrium $(\lambda_h^*(\hat{x}_0), \lambda_l^*(\hat{x}_0))$ with $\lambda_l^*(\hat{x}_0) \in (0,1)$ for some $\hat{x}_0 \in (0,1)$. Along the unique equilibrium branch through this equilibrium, as defined in Proposition 2.2, an increase in the level of believed overconfidence strictly increases the gender application gap through the indirect effect for every $\hat{x} \in [\hat{x}_0, \hat{x}^*)$. If the branch continues beyond $\hat{x}^*$, then the gender application gap is weakly decreasing for $\hat{x} \geq \hat{x}^*$.*

PROOF. See Appendix B.2.5. □

2.3. Design of the Experiment

The design of our laboratory experiment closely follows the spirit of our theoretical framework.[15] To study whether gender differences in (believed) overconfidence can explain gender differences in job-application behavior, we construct a controlled environment that captures the two main elements of the model. First, participants need to differ in relative abilities and must form beliefs about their own relative ability, as well as second-order beliefs about the beliefs of others. Second, participants compete in a stylized labor market for vacancies that differ in wages and their chances of success depend on their relative abilities, as well as on the quantity of competition they face. The controlled environment then enables us to exogenously provide information to plausibly remove overconfidence and beliefs that others are overconfident. This allows us to identify their causal effects on application behavior, and hence test the direct and indirect effects of overconfidence and their implications for the gender application gap.

Each laboratory session consists of two sequential stages. In the first stage, participants work on a cognitive ability task, namely Raven Progressive Matrices. Raven matrices are a widely used measure of abstract reasoning and fluid intelligence (Raven, 2000). For our purposes, they are particularly well suited because they do not require prior knowledge and are sufficiently familiar in structure for participants to form meaningful beliefs about their relative performance.[16] Participants have 15 minutes to solve up to 11 matrices. Participants are ranked according to their performance in this task: the primary determinant of the performance of a participant is the number of matrices that she solved correctly. The secondary determinant is completion time: if participants solve the same number of matrices, we rank the participant better that solved the matrices faster. The 15-minute window is chosen such that most participants can complete all matrices, ensuring that we have continuous variation in completion time, which prevents ties. Participants are informed about how their performance is determined but are not directly incentivized for performance. Instead, they are told that their performance will influence their payoffs in the next stage. This mirrors real-world settings such as education or skill development, where effort and ability are not directly rewarded but influence subsequent labor market opportunities.

After completing the cognitive ability task, participants are randomly matched into pairs, which later serve as the competitive units in the stylized labor-market stage. Before entering that stage, we elicit two beliefs of the participants using an incentivized belief-elicitation procedure. In the first question, participants state their belief about the prob-

[15]The pre-registration can be found under: socialscienceregistry.org under ID: AEARCTR-0014968

[16]Henceforth, we use performance in the Raven matrices task and ability interchangeably.

ability that they performed better in the Raven matrices than their randomly matched pair member. This is equivalent to eliciting their belief about their relative standing in the performance distribution of participants in the same session. The difference between the elicited belief and their actual relative performance serves as our measure of relative overconfidence. In the second belief question, participants must report their belief of how, on average, the other participants within the same session responded to the first question. This belief yields our measure of believed overconfidence.

Following Zimmermann (2020), both belief elicitation questions are incentivized using quadratic scoring rules that reward accuracy. For the second belief elicitation, payoffs are steeper to account for the fact that the true group averages typically lie within a tighter range.[17] While we disclose the specific formula and payoff scheme to the participants, we emphasize, similar to Zimmermann (2020) and as suggested by Haaland et al. (2023), that even though they might appear complicated, the core idea is simple and it is always in their best interest to provide truthful answers. We do not disclose their accrued payment until the end of the experiment, so the participants cannot deduce the accuracy of their beliefs before the labor-market stage takes place.

Within each pair, members are ranked according to their relative performance in the Raven task. The better-performing participant i is classified as the high-ability type ($\alpha(i) = \alpha_h$), and the worse-performing participant j as the low-ability type ($\alpha(j) = \alpha_l$).

Then, each pair of participants enters a stylized labor market in which they compete for two vacancies that differ only in wages: a high-wage vacancy that pays 12€ and a low-wage vacancy that pays 7€. If a pair member obtains an offer from one vacancy, the wage of this vacancy is realized at the end of this stage without any further tasks. If a pair member receives offers from both vacancies, the high-wage offer materializes for this pair member. Importantly, when one participant receives both offers, the low-wage offer is not assigned to the other participant but is forfeited; the other participant therefore receives 0€ in this case.

Each participant can freely allocate 100 application points between the two vacancies. Allowing for quasi-continuous application decision rather than a binary apply/not-apply choice provides three advantages. First, allowing participants to allocate application points between the two vacancies corresponds more closely to the strategy space of the model, in which job seekers allocate probability weights between the vacancies. Second, it captures the intensity of the application choice and, hence, allows participants to express otherwise inframarginal application behavior. Third, it generates smoother variation in observed application behavior, which increases statistical power.

A vacancy sends out one offer if it receives a positive number of application points

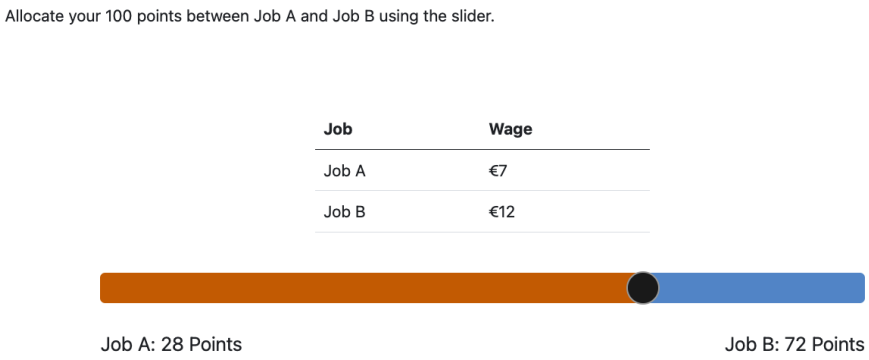
[17] See instructions in Appendix B.5 for the precise payoffs schemes.

and no offer otherwise. Conditional on a vacancy receiving positive application points, the probability that participant i receives the high-wage or low-wage offer depends on the application points of both participants in the pair and on their ability ranks:

$$p_i(w_h) = \frac{\lambda_i \alpha(i)}{\lambda_i \alpha(i) + \lambda_j \alpha(j)};\; p_i(w_l) = \frac{(100 - \lambda_i)\alpha(i)}{(100 - \lambda_i)\alpha(i) + (100 - \lambda_j)\alpha(j)}, \tag{2.4}$$

where $\alpha_h = 2$, $\alpha_l = 1$, and $\lambda_i \in [0, 100]$ denotes the number of application points that participant i assigns to the high-wage vacancy. This assignment rule is designed to mirror the comparative statics of the hiring probability function (2.1) from our model while keeping it tractable for the participants. Accordingly, the probability of obtaining an offer from a given vacancy increases with own ability rank and with the number of application points allocated to it, while it decreases with the intensity of competition, that is, the number of application points submitted to the same vacancy by the other pair member. The environment thus captures the central trade-off in directed search between potential earnings of a job and the expected competition for it. Because it is crucial that this assignment rule is well understood, participants need to answer three comprehension questions correctly before they could make their application decision. As depicted in Figure 2.1, participants have to move an uninitialized slider between both jobs to allocate their application points.

Figure 2.1: Decision Screen



Note: The figure shows the decision screen in the application stage of the experiment. Participants allocate their application points between the high-wage and low-wage vacancies by moving an initially uninitialized slider.

2.3.1. Treatment

As of now, we have obtained incentivized measures of individual (believed) overconfidence and can observe application behavior in a stylized labor market. To identify the causal effects of the two overconfidence channels, we use a between-groups information treatment that plausibly removes overconfidence and beliefs that others are overconfident.

Pairs are randomly assigned to either CONTROL or TREATMENT. In CONTROL, neither member of the pair receives any feedback about their own performance nor about the performance of the other pair member. In TREATMENT, one randomly selected participant within each pair—the INFO PLAYER—receives accurate information about her relative performance in the Raven task relative to all other participants in the same session. Specifically, participants are told how many individuals in the session had performed worse and how many had performed better, the total number of participants, and their implied ex ante probability of being classified as the high-ability type when being randomly matched with one other participant of the same session.[18] This information treatment is designed to remove any misperceptions about the own relative performance and thus overconfidence. It allows us to estimate the causal effect of overconfidence on application behavior, i.e., the direct effect of overconfidence.

The other member of the pair—the OTHER-INFO PLAYER—is informed that his pair member has received this feedback but does not see the information himself. Hence, the OTHER-INFO PLAYER learns that his competitor has all the relevant information to form accurate beliefs about her own relative performance. This information treatment is designed to remove any believed overconfidence, which allows us to estimate the causal effect of believed overconfidence on application behavior, i.e., the indirect effect of overconfidence.

Whereas creating a labor market with only two job seekers is a stark departure from our model, it enables us to create environments where we can plausibly remove believed overconfidence for a single job seeker by providing information about relative performance to only one other job seeker, i.e., to the INFO PLAYER. In any alternative design with $N > 2$ job seekers in the labor market, removing believed overconfidence for a single job seeker would require revealing relative performance information to the remaining $N - 1$ job seekers, which would substantially reduce statistical power. By combining these informational roles within pairs, the design enables a clean identification of the channels through which (believed) overconfidence shapes application behavior. To maintain a balanced sample across informational roles CONTROL, INFO, and OTHER-INFO we oversampled TREATMENT with two-to-one odds relative to CONTROL pairs.

One possible concern for this identification is that the information treatment affects behavior not only through the targeted belief channels, but also through additional feedback responses. For the INFO PLAYER, such feedback would require third-level reasoning: she would need to anticipate that the OTHER-INFO PLAYER changes his behavior because he knows that she received information. For the OTHER-INFO PLAYER, an analogous concern

[18] The feedback was presented in an explicit, transparent format to make the information easily interpretable and salient, see instructions in Appendix B.5.

would require an additional layer of reasoning. Since the literature on level-k reasoning generally finds that the most frequent types involve only one or two levels of reasoning, with higher-level types being rare (Arad and Rubinstein, 2012; Crawford et al., 2013), we do not view these additional feedback channels as serious threats to identification.

2.3.2. Additional Measures

Directly after the Raven task, participants report how many matrices they believe they solved correctly. This unincentivized question provides a measure of perceived absolute performance.

To conclude the labor-market stage, we implement an unincentivized vignette exercise as a supplementary test of the indirect effect. Participants allocate their application points in two hypothetical scenarios. In one scenario, the competitor is described as unbiased: the competitor is better than 50% of participants and correctly believes to be better than 50%. In the other scenario, the competitor is described as overconfident: the competitor is better than 50% of participants but believes to be better than 80%. The order of the two scenarios is randomized. Additionally, participants are also randomly assigned to one of two own-ability conditions: either they are told to imagine that they are better than 70% or that they are better than 30% of participants.[19]

After the labor-market stage, we elicit risk aversion via an multiple price list (MPL) where participants make 17 binary choices between a lottery and a certain payment, with the safe amount increasing in 50-cent increments from 4€ to 12€. To simplify the decision process, we instructed participants to identify the point at which they would switch from preferring the lottery to preferring the certain payment. All remaining choices were then automatically filled according to this cross-over-point, thus imposing consistency on their choices.[20] Before they had to make their choices, participants were told that if the MPL is selected for payment one of the 17 choices will be randomly chosen with equal probability and implemented. We use these choices to derive constant relative risk aversion (CRRA) parameters, which enter the empirical analysis as a control variable.

Finally, a short post-experimental survey was administered to collect information about the sociodemographic and academic backgrounds of the participants. The survey also elicited self-reported willingness to take risks (Dohmen et al., 2011). As participants are instructed before the experiment starts, they are then payed out a participant fee and

[19]To maintain consistency with the labor-market task, we deliberately frame the vignette instructions in terms of percentiles in the ability distribution and remind participants how these percentiles map into the probability of being the high-ability type. We note, however, that ranks are deterministic from the perspective of participants in the vignette scenarios, unlike in the actual labor-market task.

[20]As briefly discussed by Bronchetti et al. (2023), the use of single cross-over-points is established in the experimental literature and can make the administration easier and faster for the participants.

the payment from the belief elicitation with certainty and either the outcome of the labor market stage or the outcome of the MPL is selected randomly for payment.

2.3.3. Procedures

The experiment was programmed in oTree (Chen et al., 2016) and conducted in the Frankfurt Laboratory for Experimental Economic Research (FLEX) in Frankfurt and the Düsseldorf Institute for Competition Economics (DICELab) in Düsseldorf from December 2024 to June 2025. Participants were invited via ORSEE (Greiner, 2015) and hroot (Bock et al., 2014). In total, we had 430 participants distributed across 24 sessions (16 in Frankfurt and eight in Düsseldorf). Session sizes ranged from eight to 24 participants, with a median of 20. As pre-registered, from all 430 participants, we removed participants who only answered one Raven task correctly as this corresponds to performing worse than random chance in the Raven matrices. Because we want to examine differences between men and women, we restrict our sample to men and women yielding a total sample of 421 participants. Table B.2 in Appendix B.4 provides descriptive statistics and balancing checks. It shows that our randomization generally resulted in well balanced treatment groups. Sessions lasted approximately 50 minutes and participants earned 11.55 € on average.

2.3.4. Conditions and Hypotheses

While the setup of the experiment aims to capture the main elements of the theoretical model, it also departs from the model in some aspects. We therefore need to translate the main assumptions and predictions of the model into testable counterparts for the experiment.

Before turning to the specific assumptions and predictions, we note one important difference between the model and the experiment that applies throughout. In the model, overconfidence is a binary attribute: a job seeker is either overconfident or unbiased. In the experiment, by contrast, we measure overconfidence as a continuous variable. Specifically, we define the level of overconfidence of participant i as

$$\tilde{q}_i = b_i^{own} - q_i, \tag{2.5}$$

where b_i^{own} denotes the elicited belief about own relative ability and q_i denotes actual relative ability. The difference between the two captures the level of overconfidence.

Analogously, we define believed overconfidence as

$$\phi_i = b_i^{other} - \frac{\sum_{j=1}^{N_s} q_j - b_i^{own}}{N_s - 1}, \tag{2.6}$$

where b_i^{other} is participant i's belief about the average reported relative ability of all other participants in the session, N_s denotes session size, and q_j denotes participant j's realized

relative ability. This measure captures the difference between participant i's belief about others' average reported relative ability and the counterfactual average belief that would arise if all other participants reported their true percentiles, conditional on participant i's own self-assessment being correct.[21]

Using the level-measure of overconfidence instead of a binary classification has three advantages. First, it assigns each participant a deterministic level of overconfidence that only depends on her own ability and beliefs. In contrast, strictly imposing the model's binary notion of overconfidence would introduce randomness, because the classification would then also depend on the actual relative ability of the matched competitor. Second, the level-measure better reflects real labor markets, where individuals compare themselves to many potential competitors. Third, the continuous variation in the levels of overconfidence increases statistical power, just as using application intensity rather than a binary application decision. Measuring believed overconfidence in terms of levels then follows immediately from defining overconfidence in levels.

Using levels of overconfidence instead of shares yields the testable counterpart of Assumption 2.2 that men are more overconfident than women.

Condition 2.1 *On average, men exhibit higher levels of overconfidence than women for a given level of actual relative ability.*

Because our measure of overconfidence directly depends on actual relative ability, overconfidence may systematically vary with it. We therefore formulate Condition 2.1 for given levels of actual relative ability.

The model also assumes that job seekers hold a common belief about the share of overconfident job seekers, as stated in Assumption 2.1. In the experiment, we require that both men and women believe that others are overconfident on average, with no significant gender differences in the level of this believed overconfidence.

Condition 2.2 *On average, both men and women believe that others are overconfident, with no significant gender differences in the level of believed overconfidence.*

Condition 2.1 and Condition 2.2 are necessary for the direct and indirect effects to potentially contribute to the gender application gap in the experiment. The next step is to test whether we can empirically validate the existence of the direct and indirect effects. We therefore translate Proposition 2.1 and Proposition 2.2 into the following hypotheses for the experiment.

[21] This counterfactual average lies within the interval of $[45, 53.57]$ for all participants.

Hypothesis 2.1 *The application weight that a participant allocates to the high-wage vacancy increases in the level of overconfidence of that participant.*

Hypothesis 2.2 *The application weight that a participant allocates to the high-wage vacancy decreases in the level of believed overconfidence. The magnitude of this decrease is smaller for more overconfident participants.*

If we can empirically validate that men exhibit higher levels of overconfidence than women (Condition 2.1) and if the application weight assigned to the high-wage vacancy increases in the level of overconfidence (Hypothesis 2.1), gender differences in overconfidence through the direct effect should contribute to the gender application gap. Removing overconfidence should therefore decrease this gap. This leads to the following hypothesis, which is the testable counterpart of Corollary 2.1.

Hypothesis 2.3 *Under Condition 2.1 and Hypothesis 2.1, the gender gap in application weight assigned to the high-wage vacancy is smaller in* Info-Treatment *than in* Control.

In addition, if men exhibit higher levels of overconfidence than women (Condition 2.1) but similar positive levels of believed overconfidence (Condition 2.2), and if the application weight allocated to the high-wage vacancy decreases in the level of believed overconfidence but less so for more overconfident participants (Hypothesis 2.2), then gender differences in overconfidence should also contribute to the gender application gap through the indirect effect. In this case, removing believed overconfidence should decrease the gap. This leads to the following hypothesis, which is the testable counterpart of Corollary 2.2.

Hypothesis 2.4 *Under Condition 2.1, Condition 2.2 and Hypothesis 2.2, the gender gap in application weight assigned to the high-wage vacancy is smaller in* Other-Info-Treatment *than in* Control.

2.4. Results

Our empirical analysis proceeds as follows. First, we present descriptive evidence on gender differences in application behavior and the level of (believed) overconfidence to examine whether Condition 2.1 and Condition 2.2 are fulfilled. Next, we assess whether the direct and indirect effects of overconfidence exist, that is, we test Hypothesis 2.1 and Hypothesis 2.2. Finally, we examine how these effects relate to the gender application gap, thus evaluating Hypothesis 2.3 and Hypothesis 2.4, and briefly discuss implications for subsequent labor-market outcomes.

2.4.1. Gender Differences in Application Behavior and (Believed) Overconfidence

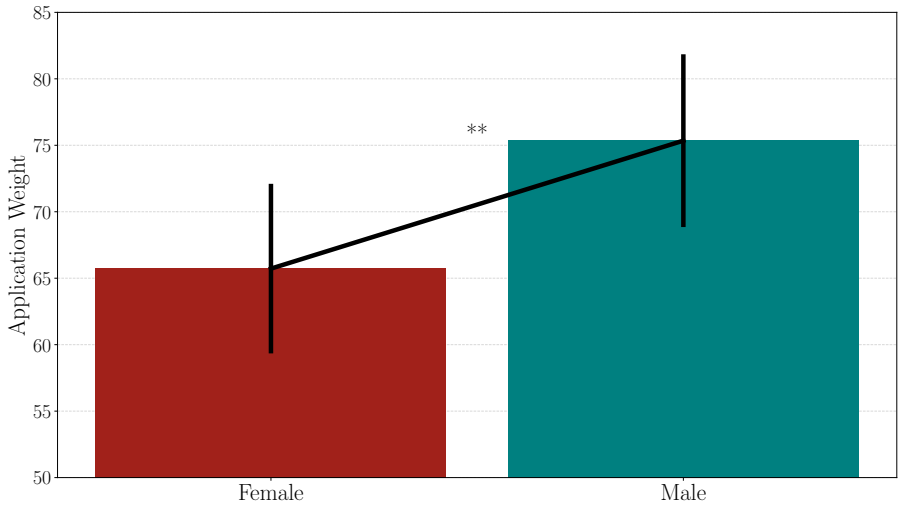
In a first step, we focus on the application weight allocated to the high-wage job in CONTROL. In CONTROL, participants receive no information about their actual relative performance and, hence, this provides us with a benchmark for behavior in an environment where beliefs are uninformed. Figure 2.2 displays the average application weight that men and women in CONTROL allocate to the high-wage vacancy. In line with the findings in the field (Fluchtmann et al., 2024; Lochner and Merkl, 2026), we document a gender application gap: men allocate about 10 application points more to the high-wage vacancy than women ($p = 0.035$), which corresponds to a gender gap of about 15%. Controlling for actual relative ability and other covariates, this gender difference shrinks to about 7.7-8.0 application points but remains significant on the 10% level (Table 2.3). In addition, Table B.3 in Appendix B.4 shows that once we account for overconfidence, the gender difference shrinks by 24% to 31% across specifications and becomes insignificant at any conventional level, already pointing to the central role of overconfidence in the gender application gap.[22]

As outlined above, two conditions must hold for the direct and indirect effects of overconfidence to contribute to the application gap. First, for the direct effect, men must, on average, display higher levels of overconfidence than women for a given level of actual relative ability (Condition 2.1). Second, for the indirect effect, both men and women must (i) believe that others are overconfident on average, without (ii) significant gender differences in the level of believed overconfidence (Condition 2.2).

Panel (a) of Figure 2.3 compares overconfidence among men and women. Specifically, it plots average beliefs of men and women about relative ability within ten equally-spaced actual ability bins. Most bin-specific means lie above the 45-degree line, indicating that participants are, on average, overconfident which is in line with the literature (e.g. Barber and Odean, 2001; Cortés et al., 2023; Niederle and Vesterlund, 2007). For both women and men, the fitted relationships are fairly flat: individuals up to approximately the 70th percentile of the actual ability distribution tend to overestimate their relative ability, while those above this threshold tend to underestimate it. In almost every ability bin, men report higher relative ability beliefs than women. As reported in Table 2.1, controlling for relative ability, women's belief are about 4.7 percentiles lower than those of men which is significant on the 1%-level and remains significant on the 5%-level if we control for other

[22]We considered applying the Obviously Related Instrumental Variables (ORIV) correction of Gillen et al. (2019) to address measurement error in the risk-preference control. However, our self-reported risk measure is correlated with overconfidence, making it unsuitable as an instrument for the incentivized risk measure. See Appendix B.4.1 for details.

Figure 2.2: Gender Application Gap



Note: The figure shows the average number of application points that men and women in the control group allocate to the high-wage vacancy. Error bars represent 95% confidence intervals based on robust standard errors. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

covariates.[23] We thus conclude that Condition 2.1 is fulfilled.

Panel (b) of Figure 2.3 shows that both men and women expect others to be overconfident. On average, men believe that others overestimate their position in the ability distribution by about 9.5 percentiles, while women expect an overestimation of about 7.7 percentiles. Relative to the actual average overconfidence of 11.2 percentiles, believed overconfidence thus amounts to about 85% of actual overconfidence for men and 69% for women. Raw gender differences in believed overconfidence are not significant ($p = 0.153$) and remain insignificant throughout all specifications in Table 2.1 (e.g. $p = 0.303$ in our preferred specification). Thus, we find no evidence of systematic gender differences in believed overconfidence, indicating that Condition 2.2 is also satisfied.

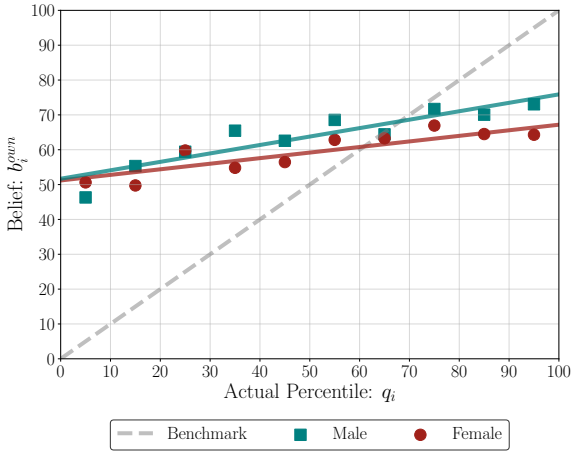
2.4.2. Causal Effects of (Believed) Overconfidence on Application Behavior

Having established higher overconfidence among men and similar positive levels of believed overconfidence across genders, we now examine whether overconfidence and believed overconfidence affect application behavior. In the terminology of our model, this corresponds to testing for the direct (Hypothesis 2.1) and indirect (Hypothesis 2.2) effects of overconfidence. To test for the direct effect, we compare the application behavior of participants in Info to participants in Control. Column (1) of Table 2.2 shows that participants who receive information about their relative ability (Info) allocate, on average, about 5 fewer

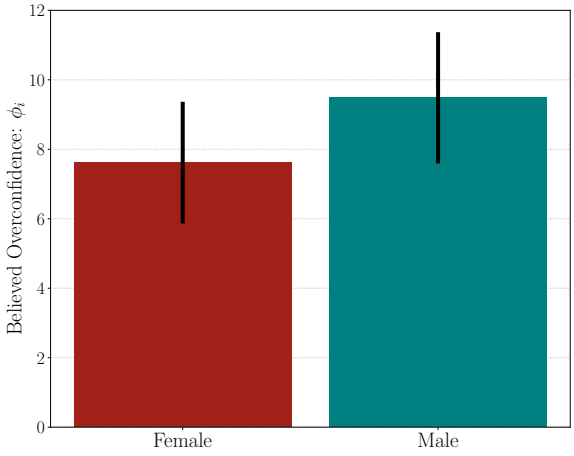
[23] Without controlling for actual relative ability, the gender difference in overconfidence is not statistically significant. This arises since women perform slightly worse in the Raven matrices in combination with beliefs not being perfectly correlated with actual relative ability as depicted in Panel (a) of Figure 2.3.

Figure 2.3: Overconfidence and Believed Overconfidence

(a) Ability Belief By Actual Ability



(b) Believed Overconfidence



Note: Panel (a) shows average beliefs about own relative performance by gender, separately for ten equally spaced bins of actual relative performance. The fitted lines show the corresponding gender-specific OLS relationships between beliefs about own relative performance and actual relative performance. The gray dashed 45-degree line indicates all points at which believed and actual relative performance coincide. Panel (b) shows the average level of believed overconfidence by gender as defined by Equation 2.6. Error bars represent 95% confidence intervals based on robust standard errors.

Table 2.1: Relative and Believed Overconfidence

	Relative Overconfidence: $\tilde{q}_i$			Believed Overconfidence: ϕ_i		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	11.480***	54.116***	64.762***	9.483***	3.871**	6.858**
	(2.007)	(2.061)	(3.114)	(0.957)	(1.602)	(2.781)
Female	-0.572	-4.665***	-3.332**	-1.867	-1.328	-0.492
	(2.874)	(1.654)	(1.631)	(1.307)	(1.289)	(1.276)
Actual Percentile		-0.804***	-0.832***		0.106***	0.095***
		(0.028)	(0.027)		(0.022)	(0.023)
Additional Controls			✓			✓
Wald-test: Female = 0	0.842	0.005	0.041	0.153	0.303	0.700
Observations	421	421	421	421	421	421
R^2	0.000	0.680	0.700	0.005	0.062	0.089

Note: The table reports OLS regression results. In columns (1)–(3), the dependent variable is overconfidence, $\tilde{q}_i$, as defined in Equation 2.5. In columns (4)–(6), the dependent variable is believed overconfidence, ϕ_i, as defined in Equation 2.6. Female is a dummy variable equal to one for women. Actual Percentile denotes the participant's percentile rank in the ability task. Additional controls include age, lab fixed effects, last high school math grade, and the CRRA parameter. * p $<$ 0.1, ** p $<$ 0.05, *** p $<$ 0.01. Robust standard errors in parentheses.

application intensity points to the high-wage vacancy than participants in CONTROL. The difference between groups is not statistically significant on any conventional level ($p = 0.150$). This is unsurprising since this treatment effect pools 32% of the participants in INFO who are ex ante underconfident and therefore receive positive feedback about their own relative ability with participants who are ex ante overconfident and therefore receive negative feedback.

Hence, a test of Hypothesis 2.1 must allow the treatment effect to vary with the sign and magnitude of overconfidence. Thus, our main regression specification is

$$\lambda_i = \beta_0 + \beta_1 \mathrm{T}_i + \beta_2 \tilde{q}_i + \beta_3 (\mathrm{T}_i \times \tilde{q}_i) + \epsilon_i \tag{2.7}$$

where $T_i = 1$ if participant i is in INFO and $T_i = 0$ if participant i is in CONTROL, and $\tilde{q}_i$ is the measure of overconfidence as defined in Equation 2.5. In this specification, the interaction coefficients estimates how the difference between beliefs about relative ability and actual relative ability affects application behavior, representing the direct effect of overconfidence.

Column (2) of Table 2.2 presents the estimation results of regression specification (2.7). The coefficient is estimated at -0.308 and is significant on the 1%-level ($p = 0.009$). Participants who are ex ante more overconfident reduce their allocated high-wage application points by more when provided with information about their actual relative ability. Specifically, a one-percentile increase in overconfidence raises high-wage application points by about 0.3 points. Hence, higher overconfidence causally increases high-wage application points, providing evidence for the presence of the direct effect of overconfidence (Hypothesis 2.1).[24]

This effect is not driven by participants at the bottom or top end of the ability distribution (see Table B.5 in Appendix B.4). Moreover, we find no indication that participants simply anchor their application points to the provided information about relative ability. Participants in INFO are not more likely to choose a high-wage application-point allocation equal or close to the provided information than participants in CONTROL are to choose an allocation equal or close to their own ability belief, irrespective of the width of the interval around the information or belief (see Table B.6 in Appendix B.4).

Panel (a) of Figure 2.4 shows how the treatment effect varies by the sign of overconfidence. It shows that the effect is driven by ex ante overconfident participants, who reduce their application points allocated to the high-wage vacancy. In particular, they

[24]The baseline coefficient on overconfidence, β_2, captures the association between overconfidence and high-wage application points in CONTROL. It is close to zero in Table 2.2 because overconfidence is negatively correlated with actual relative ability. Once actual relative ability is controlled for, this association becomes positive and highly significant (see Table B.4 in Appendix B.4). The coefficient of interest, β_3, remains essentially unchanged, as expected given random assignment to treatment.

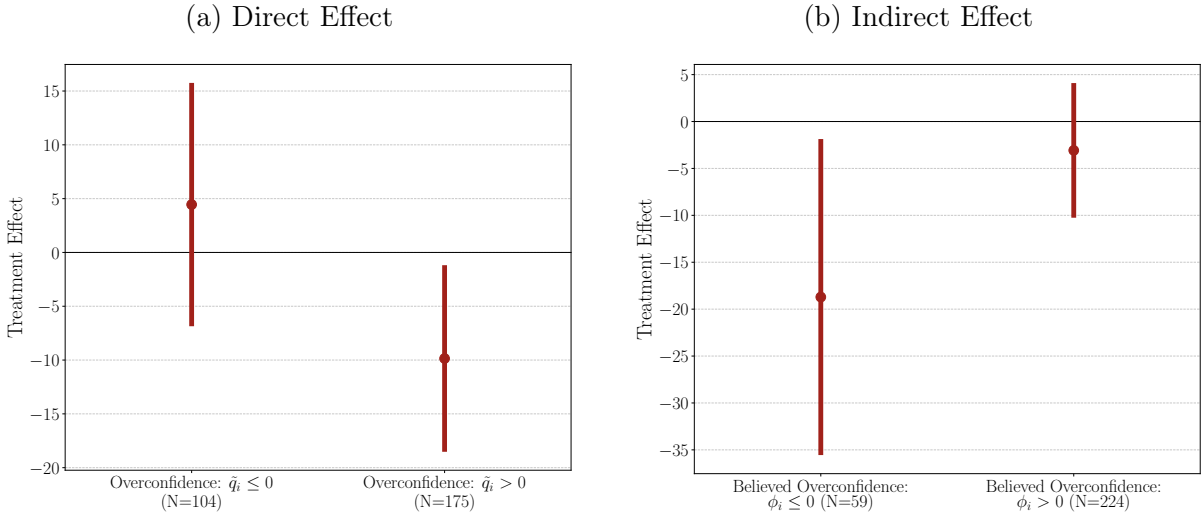
Table 2.2: Effects of (Believed) Overconfidence on Application Behavior

	Dependent variable: Application Weight			
	(1)	(2)	(3)	(4)
INFO	-5.114	-1.818		
	(3.548)	(3.516)		
Overconfidence		-0.015		
		(0.078)		
INFO × Overconfidence		-0.308***		
		(0.119)		
OTHER-INFO			-5.764*	-8.355*
			(3.439)	(4.341)
Believed Overconfidence				0.259
				(0.220)
OTHER-INFO ×				
Believed Overconfidence				0.214
				(0.294)
Constant	70.433***	70.567***	70.433***	68.482***
	(2.310)	(2.419)	(2.310)	(2.943)
Observations	279	279	283	283
R^2	0.007	0.058	0.010	0.039

Note: The table reports OLS regression results. The dependent variable is the application weight allocated to the high-wage vacancy. INFO is a dummy variable equal to one for participants who receive information about their own actual relative ability. OTHER-INFO is a dummy variable equal to one for participants whose competitor receives information about the competitor's own actual relative ability. Columns (1) and (2) compare INFO participants to participants in CONTROL. Columns (3) and (4) compare OTHER-INFO participants to participants in CONTROL. The measure of overconfidence is defined in Equation 2.5. The measure of believed overconfidence is defined in Equation 2.6. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

allocate about 10 application points less to the high-wage vacancy than participants in CONTROL who are also overconfident ($p = 0.026$). This effect is sizeable and amounts to about 14% of the application weight allocated to the high-wage vacancy in CONTROL. Analogously, participants who are ex ante underconfident react positively to information regarding their relative performance. However, this is a smaller effect at about half the size of the effect of overconfident participants and not significant at conventional levels. It shows that the direct effect is primarily driven by overconfident participants who receive information about their relative performance. This test can also be seen as a different specification of Hypothesis 2.1 where the treatment effect only depends on the sign and

Figure 2.4: Heterogeneity of the Direct and Indirect Effect



Note: Panel (a) reports the coefficient on INFO from regressions corresponding to column (1) of Table 2.2, estimated separately for participants with positive and nonpositive overconfidence. Panel (b) reports the coefficient on OTHER-INFO from regressions corresponding to column (3) of Table 2.2, estimated separately for participants with positive and nonpositive believed overconfidence. Error bars represent 95% confidence intervals based on robust standard errors.

not on the level of overconfidence.

Next, we test the indirect effect, i.e., Hypothesis 2.2. For this, first we test whether the indirect effect generally exists. We compare participants in OTHER-INFO where the other pair member receives information on their relative performance to participants in CONTROL. Analogously, to the analysis of the direct effect, we expect the treatment effect to depend on the magnitude and direction of believed overconfidence. Accordingly, we estimate the following regression specification:

$$\lambda_i = \beta_0 + \beta_1 \mathrm{T}_i + \beta_2 \phi_i + \beta_3 (\mathrm{T}_i \times \phi_i) + \epsilon_i \tag{2.8}$$

where $T_i = 1$ if participant i is in OTHER-INFO and $T_i = 0$ if participant i is in CONTROL, and ϕ_i is the measure of believed overconfidence as defined in Equation 2.6. We are interested in the coefficient of the interaction term which measures how the level of believed overconfidence causally affects application behavior. Then, we test the second part of Hypothesis 2.2 by analyzing whether participants with higher beliefs about their own relative ability are less responsive to believed overconfidence. To do so, we estimate regression specification (2.8) separately for different parts of the ability beliefs distribution and compare the resulting coefficients of the interaction term.

Column (4) of Table 2.2 presents the estimation results of our regression specification (2.8). Although the coefficient of the interaction term has the expected sign, it is not statistically significant on any conventional level ($p = 0.468$). Hence, we find no evidence

that participants adjust application behavior depending on their initial level of believed overconfidence when being informed that a competitor received information on their relative ability. This suggests, that believed overconfidence does not influence application behavior, on average. Panel (b) of Figure 2.4 shows how the effect differs by the sign of the believed overconfidence. Interestingly, participants in OTHER-INFO with negative levels of believed overconfidence as well as those with positive levels decrease their high-wage application points compared to respective participants in CONTROL. Whereas the group of participants with negative levels of believed overconfidence (21% of participants) move into the expected direction, i.e., they revise their high-wage application points downwards ($p = 0.029$), the group with positive levels of believed overconfidence also revises their high-wage application points slightly downwards, though the effect is not significant on any conventional level. In particular, for the group with positive believed overconfidence where we would expect that they revise their application decision upwards, we can rule out that the true effect is larger than 4.1 application points on the 97.5% confidence level.

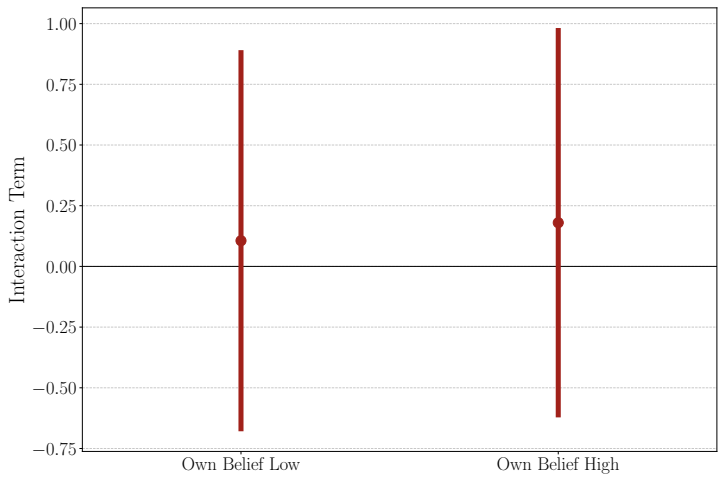
So far, our estimates suggests that the indirect effect does not play a role for application behavior. However, our findings may still be consistent with the predictions from the model and Hypothesis 2.2 if it reflects that participants with high ability beliefs are insensitive to believed overconfidence. Figure 2.5 shows the estimates of β_3 from regression specification (2.8) for participants with low and high beliefs about their own relative ability. Estimates for both groups are very similar (0.106 to 0.180, $p = 0.897$) which suggests that the indirect effect does not depend on the beliefs of participants about their relative ability.

As a supplementary exercise, we examine the unincentivized vignette task, with results reported in Figure B.2 in Appendix B.4. The average within-participant difference in high-wage application weight between the overconfident- and unbiased-competitor scenarios is directionally misaligned with the predicted effect (Panel (a) of Figure B.2). Participants assigned to the lower own-ability condition, however, are 8 percentage points more likely to exhibit a negative difference ($p = 0.080$), suggesting that they are more likely to reallocate application weight away from the high-wage vacancy when the competitor is overconfident (Panel (b) of Figure B.2). While this is consistent with the logic of the indirect effect, we still conclude that the overall evidence provides no support for Hypothesis 2.2.[25]

Summing up, we have established the following results: (i) there is a gender application gap, with men allocating more application points to the high-wage vacancy; (ii) men are more overconfident than women conditional on actual relative ability; and (iii) overconfidence causally increases the application points allocated to the high-wage vacancy. In addition, (iv) both men and women believe that others are overconfident on

[25] All the discussed results concerning believed overconfidence and the indirect effect are fully robust to using a simpler measure of believed overconfidence, $b_i^{other} - 50$, instead of the measure defined in (2.6).

Figure 2.5: Heterogeneity of the Indirect Effect by Own Belief



Note: The figure reports estimates of the coefficient of the interaction term, β_3, from specification (2.8), estimated separately for participants with below-median and above-median beliefs about their own relative ability. Error bars represent 95% confidence intervals based on robust standard errors.

average, with no significant gender differences in the level of this belief, but (v) participants, independent of their belief about own relative ability, do not react strategically to believed overconfidence, providing no evidence for the indirect effect. Since these results rule out the indirect effect as a contributor to the gender application gap, we now focus on the direct channel. Specifically, we test whether the gender application gap is smaller in INFO-TREATMENT than in CONTROL, as predicted by Hypothesis 2.3, and quantify how much of the application gap can be attributed to higher overconfidence among men.

2.4.3. Overconfidence and the Gender Application Gap

Table 2.3 reports the gender application gap in CONTROL and INFO respectively with and without additional controls. Whereas in CONTROL women allocate 7.7–8 application points fewer to the high-wage vacancy, which is significant on the 10%-level in these regression specifications ($p = 0.089$ and $p = 0.085$), in INFO the gender gap shrinks to 4.4–5.4 application points and becomes insignificant on any conventional level ($p = 0.361$ and $p = 0.287$). While the differences between the two gender gaps are not significant and should thus be interpreted with caution, the reductions themselves are sizeable: depending on the chosen regression specification, the gender application gap decreases by about 42% or 33% with controls. Hence, providing information about relative ability and thereby plausibly removing overconfidence reduces the gender application gap substantially, supporting Hypothesis 2.3.[26]

[26]This evidence is also consistent with the results in Table B.3 in Appendix B.4, which show that controlling directly for overconfidence substantially reduces the application gap in CONTROL.

Table 2.3: Gender Application Gap in Info and Control

	Dependent variable: Application Weight			
	Control	Info	Control	Info
	(1)	(2)	(3)	(4)
Female	-7.753*	-4.445	-8.000*	-5.383
	(4.554)	(4.868)	(4.649)	(5.054)
Actual Percentile	0.175**	0.487***	0.205**	0.491***
	(0.069)	(0.079)	(0.081)	(0.085)
Constant	65.280***	42.176***	62.748***	45.604***
	(5.515)	(6.365)	(10.846)	(11.625)
Additional Controls			✓	✓
Wald-test: Female = 0	0.089	0.361	0.085	0.287
Observations	141	138	141	138
R^2	0.069	0.226	0.101	0.265

Note: The table reports OLS regression results. The dependent variable is the application weight allocated to the high-wage vacancy. Regressions in columns (1) and (3) use participants in Control; regressions in columns (2) and (4) use participants in Info. *Female* is a dummy variable equal to one for women. Actual Percentile denotes the participant's percentile rank in the ability task. Additional controls include age, lab fixed effects, last high school math grade, and the CRRA parameter. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Our design is not particularly suited to study labor market outcomes, as they propagate through our specific assignment rule (2.4). This assignment rule introduces randomness into wage and employment outcomes, and combined with our limited sample, this prevents significant gender differences in application behavior from translating into significant gender differences in realized labor market outcomes. Nonetheless, we report in Table 2.4 the results of linear probability models where obtaining the high-wage vacancy (conditional on being employed) and not getting an offer from any of the two vacancies, i.e., ending up unemployed, constitute the dependent variables. Results for Info and Control are reported side-by-side. In Control females are 14 percentage points less likely to obtain a high-wage vacancy whereas in Info this difference shrinks to only 4.5 percentage points. While this shrinkage in the gender wage gap is in line with what we should expect, none of the coefficients are significant on conventional levels. For the probability of ending up unemployed, we find virtually no difference (1.8 and 0.3 percentage points) between men and women both in Control and Info.

Table 2.4: Wage and Employment Outcomes

	LPM High-Wage		LPM Unemployed	
	CONTROL	INFO	CONTROL	INFO
Female	-0.136	-0.045	-0.018	-0.003
	(0.096)	(0.085)	(0.069)	(0.081)
Actual Percentile	0.004**	0.005***	-0.003***	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.468**	0.531***	0.228	0.264**
	(0.191)	(0.180)	(0.161)	(0.127)
Additional Controls	✓	✓	✓	✓
Wald-test: Female = 0	0.155	0.594	0.800	0.968
Observations	112	102	141	138
R^2	0.122	0.193	0.107	0.048

Note: The table reports linear probability model estimates. In columns (1) and (2), the sample is restricted to participants who receive an offer, and the dependent variable is a dummy variable equal to one if the participant is hired by the high-wage vacancy. In columns (3) and (4), the dependent variable is a dummy variable equal to one if the participant receives no offer. Columns (1) and (3) use participants in CONTROL; columns (2) and (4) use participants in INFO. Female is a dummy variable equal to one for women. Actual Percentile denotes the participant's percentile rank in the ability task. Additional controls include age, lab fixed effects, last high school math grade, and the CRRA parameter.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

2.5. Conclusion

This paper studies whether gender differences in overconfidence contribute to the gender application gap. In a directed search model, we show that such a gap can emerge through a direct effect, which operates through overestimation of own relative ability, and an indirect effect, which operates through beliefs about overconfidence among other job seekers. We test both effects in a laboratory experiment. We find support for the direct effect: men are more overconfident than women, overconfidence increases applications to high-wage vacancies, and gender differences in overconfidence account for a substantial share of the application gap. By contrast, although both men and women believe that others are overconfident on average, we find no evidence for the indirect effect.

Our findings contribute to our understanding of the mechanisms behind the gender application gap. While the experiment provides a clean test of the direct and indirect effects in a stylized environment, the results should be interpreted in light of this setting. A natural avenue for future research is to test whether the direct effect also operates in actual labor markets.

CHAPTER 3

Present Bias and On-the-Job Search

Abstract

This paper studies how present bias shapes job search, job offer acceptance behavior, and wage outcomes in a model with costly on-the-job search. Present-biased workers search less than time-consistent workers because they overweight current search costs relative to future wage gains. Awareness of future present bias affects job offer acceptance behavior. Sophisticated workers anticipate that a new job may become permanent because future on-the-job search will be distorted by present bias, and therefore demand larger wage increases to switch jobs. Naive workers, by contrast, overestimate future on-the-job search and may accept relatively low offers under the mistaken belief that these jobs are temporary steps on the wage ladder, even when they ultimately become permanent. The same mechanism raises the reservation wage of unemployed sophisticated workers. Simulations show that these channels generate systematic wage differences: naive workers are overrepresented in lower-wage jobs, sophisticated workers in intermediate-wage jobs, and time-consistent workers in higher-wage jobs. Descriptive survey evidence is consistent with the predicted differences in job acceptance behavior across types.

3.1. Introduction

Present bias, defined as the tendency to place disproportionate weight on immediate utility relative to future utility, appears widespread, with estimates suggesting that up to three quarters of individuals exhibit some degree of present bias (Cobb-Clark et al., 2024). This bias can matter whenever economically important choices require costly action today in exchange for delayed returns. The labor market is a natural setting for such choices, since many labor market decisions involve this type of trade-off and have persistent consequences for lifetime earnings. One such example is the decision of workers whether to engage in on-the-job search. On-the-job search demands immediate and costly effort but may generate offers from higher-paying jobs and thereby lead to job-to-job transitions, which are a main source of wage growth over the career (e.g., Adda and Dustmann, 2023; Faberman et al., 2022; Topel and Ward, 1992). This central role of on-the-job search in wage growth, together with its intertemporal payoff structure, makes it a choice margin on which present bias may be particularly consequential. Yet research on the relationship between present bias and on-the-job search is surprisingly limited.

This paper addresses this gap by studying how present bias and perceptions about future present bias affect job search, acceptance behavior, and wage outcomes in a model with costly on-the-job search. Present bias reduces search because workers place extra weight on current search costs relative to future wage gains. Beyond this search margin, the model identifies a novel offer acceptance channel. Workers who do not anticipate their future present bias may accept low-wage jobs as temporary steps on the wage ladder. Once employed at these jobs, however, present bias may prevent further on-the-job search, leaving workers trapped in low-wage jobs. Workers who anticipate their future present bias are more selective: they may demand larger wage increases to change jobs because some offers would end future search without providing a sufficiently high permanent wage. Through this greater selectivity, they avoid being trapped in low-wage jobs. I provide descriptive survey evidence that is consistent with these predicted differences in acceptance behavior.

To make these mechanisms transparent, I first develop a three-period model. A single employed worker chooses whether to engage in costly on-the-job search in periods 1 and 2. The worker can be one of three types. A time-consistent worker has no present bias. Naive and sophisticated workers are present-biased; naive workers do not anticipate their future present bias, whereas sophisticated workers do. Search in each period generates a job offer characterized only by a wage drawn from an exogenous distribution. If the worker accepts the offer, she earns the new wage in the following period; if she rejects it, she remains at her current wage. Following O'Donoghue and Rabin (1999), I characterize behavior using the solution concept of perception-perfect strategies: current selves choose

optimally given their perceptions about the strategies of future selves.

The model yields three results. First, in both search periods, a time-consistent worker searches up to higher wage levels than a present-biased worker because she places greater weight on future wage gains relative to current search costs. Second, among present-biased workers, perceptions about future present bias affect current search decisions. A naive worker expects her future self to search at wage levels at which a sophisticated worker knows her future self will not search. As a result, a naive worker may postpone search in period 1 because she mistakenly expects her period-2 self to search later. A sophisticated worker anticipates that postponing search today means stopping search altogether; therefore, the wage level at which she stops searching in period 1 is weakly higher than for a naive worker. Third, perceptions about future present bias also affect offer acceptance behavior. A naive worker accepts any offer that raises her wage because she expects to continue searching up to higher wage levels in the future. A sophisticated worker anticipates whether accepting such offers would end future search and may therefore reject them if the wage gain is too small to justify accepting the offer as her final job.

I next extend the analysis to an infinite-horizon model in which workers start unemployed. This extension serves two purposes. First, it allows me to study reservation wages during unemployment. With on-the-job search, perceptions about future present bias affect the value of entering employment at a given wage because they affect whether the worker expects to continue searching after accepting the job. This contrasts with the model of unemployed search without on-the-job search in DellaVigna and Paserman (2005), in which present bias no longer affects behavior after employment is entered. Second, the infinite horizon allows me to simulate wage paths over the career and therefore generates predictions for wage outcomes across types.

For sophisticated workers, the infinite-horizon model does not admit a pure perception-perfect strategy because present bias, when correctly anticipated, creates a region of wages for which current behavior and beliefs about future behavior cannot be made mutually consistent. Behavior in this region is therefore ambiguous. I address this issue by introducing a new refinement of perception-perfect strategies that generates self-consistent search and no-search regions while making the ambiguous region as small as possible.

In the infinite-horizon model, the on-the-job search and offer acceptance results from the three-period model carry over. Moreover, the same acceptance mechanism that makes sophisticated employed workers demand larger wage increases also makes sophisticated unemployed workers more selective and raises their reservation wages. Simulations show how these search and acceptance decisions translate into different wage outcomes across types. Because naive workers search less than time-consistent workers but accept low wages under the mistaken belief that they will continue searching as if they were time-consistent,

they tend to become trapped in low-wage jobs and are therefore overrepresented at the lower end of the wage distribution. The greater selectivity of sophisticated workers, which stems from their anticipation of future present bias, moves them away from low-wage jobs and toward intermediate wage levels. Time-consistent workers search up to the highest wage levels and are thus overrepresented at the upper end of the distribution.

Finally, I examine whether there is empirical support for three predictions of the model using data from the German Socio-Economic Panel Innovation Sample (SOEP-IS). The SOEP-IS includes a module that classifies workers as time-consistent, naive, or sophisticated (Cobb-Clark et al., 2024), as well as information on expected future on-the-job search, the wage increase demanded to change jobs (Jäger et al., 2024), and actual wages. The strongest evidence concerns offer acceptance: in line with the model, sophisticated workers demand significantly larger wage increases to be willing to change jobs, while naive and time-consistent workers do not differ. The evidence for anticipated future search is also directionally consistent with the model: sophisticated workers report lower expected future on-the-job search than naive and time-consistent workers, although these differences are statistically insignificant in most specifications. Observed wage distributions by type are partially consistent with the model, especially for naive workers, who are most concentrated at the bottom of the distribution.

This paper contributes to the small literature on present bias and job search behavior. Most notably, DellaVigna and Paserman (2005) and, in a follow-up paper, Paserman (2008) study unemployed job search under present bias in a setting without on-the-job search. Their model predicts that present bias reduces search effort and, for sophisticated workers, lowers reservation wages because they anticipate reduced future search while unemployed and therefore want to leave unemployment sooner. Belot et al. (2025) develop a method for eliciting time preferences and use it to test the predictions of DellaVigna and Paserman (2005); they find that less present bias is associated with more job interviews and more search effort, while reservation wages are essentially unaffected.[1] Relative to this work, I study a setting in which present bias remains relevant after workers enter employment by introducing on-the-job search. This yields an opposing mechanism relative to DellaVigna and Paserman (2005): sophisticated workers anticipate whether an accepted job will be long-lasting rather than merely serving as a first step on the wage ladder, and may therefore demand higher wages to leave unemployment.

Studies of present bias in settings with on-the-job search are scarce. Meyer (2018) studies present bias and on-the-job search among low-wage workers in Ethiopia and shows that

[1] This null average effect does not necessarily speak against reservation-wage effects of present bias. In both DellaVigna and Paserman (2005) and my model, these effects operate only through sophistication. Not distinguishing between naive and sophisticated present-biased workers but pooling them instead, as is done in Belot et al. (2025), may therefore attenuate the effects towards zero.

present-biased workers exert lower search effort. I go beyond this search-effort margin by studying a dynamic environment in which workers can continue to search after changing jobs and differ in their awareness of future present bias. These features make offer acceptance a central choice margin in addition to search, with implications for wage outcomes. Also related but conceptually different, Fahn and Seibel (2022) study present bias and on-the-job search in an optimal contracting model with limited commitment. In contrast, I study present bias in a standard job search framework without contracting frictions.

More broadly, this paper contributes to the behavioral job search literature, which studies how psychological traits and biases shape search behavior and labor market outcomes. Prior work studies a range of such psychological traits and biases in job search, including, for example, reference dependence (e.g. DellaVigna et al., 2017; Koenig et al., 2024), locus of control (e.g. Caliendo et al., 2015; McGee and McGee, 2016), overoptimism (e.g. Cortés et al., 2023; Mueller et al., 2021; Spinnewijn, 2015), and risk preferences (e.g. Pannenberg, 2010). I add to this literature by showing how present bias and perceptions of future present bias affect job search effort, offer acceptance behavior, and wage outcomes in a dynamic job search environment.

3.2. Three-Period Model

3.2.1. Setting and Timing

Time is discrete, $t \in \{1, 2, 3\}$. There is a single agent (henceforth worker) who is employed in period $t = 1$ and receives a wage $w_1 \in [w_l, w_h]$. In the same period, the worker decides whether or not to search for another job at cost $c > 0$. Throughout, I focus on the economically relevant case in which search costs are small enough that all worker types strictly prefer to search if their current wage is close enough to, but above, w_l.[2] Following canonical search models (e.g., Burdett and Mortensen, 1998; McCall, 1970), a worker who searches receives a job offer characterized solely by its wage w', drawn from an exogenous offer distribution. I take this distribution to be uniform on $[w_l, w_h]$, with cumulative distribution function F and density f. After observing the wage offer w', she decides whether to accept or reject the offer, and period $t = 1$ ends.

In period $t = 2$, the worker earns wage $w_2 = w'$ if she searched and accepted the job offer in period $t = 1$ and $w_2 = w_1$ otherwise. The decision environment in period $t = 2$ is identical to the one in period $t = 1$. Hence, the worker may again decide to search for another job at cost $c > 0$. If she searches, she again draws a job offer $w'' \sim \mathcal{U}[w_l, w_h]$, and

[2]This condition can be guaranteed in both the finite and infinite-horizon models by explicit upper bounds on c as functions of the remaining parameters. Conversely, $c > 0$ implies that all worker types strictly prefer not to search at wages sufficiently close to, but below, w_h.

then decides whether to accept this offer.

In period $t = 3$, the worker earns the wage $w_3 = w''$ if she searched and accepted the job offer in period $t = 2$ and $w_3 = w_2$ otherwise. Then, the model ends.

3.2.2. Preferences and Perceptions

The worker is risk neutral. Hence, in any given period $t \in \{1, 2\}$, her instantaneous utility is

$$u(w_t) = \begin{cases} w_t - c, & \text{if she searches in period } t \\ w_t, & \text{otherwise} \end{cases}$$

and $u(w_3) = w_3$. The worker can have two kinds of time preferences. First, the worker may have *time-consistent* preferences. Under time-consistent preferences, the relative weight assigned to instantaneous utilities in two consecutive periods is independent of when the worker evaluates the payoff stream. Such preferences are represented by standard exponential discounting. Specifically, the present-value payoffs in periods $t \in \{1, 2\}$ of a deterministic flow of instantaneous utilities $\big(u(w_t)\big)_{t=1}^{3}$ are

$$V_1 = \sum_{t=1}^{3} \delta^{t-1} u(w_t) \quad \text{and} \quad V_2 = u(w_2) + \delta u(w_3), \text{ with } \delta \in (0, 1).$$

Second, the worker may have *present-biased* time preferences. Then, the relative utility weighting of two initially future periods shifts towards the earlier period once this period becomes the present one. In the model, such preferences are represented by quasi-hyperbolic discounting (e.g., Laibson, 1997). Specifically, the present-value payoffs in periods $t \in \{1, 2\}$ of a deterministic flow of instantaneous utilities $\big(u(w_t)\big)_{t=1}^{3}$ equals

$$V_1 = u(w_1) + \beta \sum_{t=2}^{3} \delta^{t-1} u(w_t) \quad \text{and} \quad V_2 = u(w_2) + \beta\delta u(w_3), \text{ with } \delta, \beta \in (0, 1).$$

When period 1 is the present, utility in period 3 is discounted relative to period 2 by δ. Once period 2 becomes the present, however, utility in period 3 is discounted by $\beta\delta$ rather than δ, increasing the relative weight on the instantaneous utility from period 2.

A present-biased worker can differ in how she perceives her own future present bias (e.g., O'Donoghue and Rabin, 1999). She is *sophisticated* if she correctly anticipates her future present bias, so that her belief $\hat{\beta}$ about her future present bias equals her actual present bias, $\hat{\beta} = \beta$. She is *naive* if she believes that her future selves will be time-consistent, $\hat{\beta} = 1 > \beta$. In total, the model thus considers three different worker types: time-consistent ($\phi = \text{TC}$), naive ($\phi = \text{N}$), and sophisticated ($\phi = \text{S}$).[3]

[3]The literature also considers partially naive agents with beliefs $\hat{\beta} \in (\beta, 1)$, which we omit for simplicity.

3.2.3. Strategies and Solution Concept

I restrict attention to pure strategies. Let $\mathcal{W} \coloneqq [w_l, w_h]$ denote the wage space. For each worker type $\phi \in \{TC, N, S\}$ and each period $t \in \{1, 2\}$, a period-t strategy is a pair

$$s_{t,\phi} \coloneqq (q_{t,\phi}, a_{t,\phi}),$$

where $q_{t,\phi} : \mathcal{W} \to \{0, 1\}$ is a search rule and $a_{t,\phi} : \mathcal{W} \times \mathcal{W} \to \{0, 1\}$ is an acceptance rule. I interpret $q_{t,\phi}(w) = 1$ as search and $q_{t,\phi}(w) = 0$ as no search at current wage w; likewise, I interpret $a_{t,\phi}(w, w') = 1$ as accepting offer w' and $a_{t,\phi}(w, w') = 0$ as rejecting offer w' after searching at current wage w. The acceptance rule is specified for all pairs (w, w'), but affects payoffs only if the worker searches at w. A complete strategy of type ϕ is

$$s_\phi \coloneqq (s_{1,\phi}, s_{2,\phi}),$$

which specifies search and acceptance behavior in periods 1 and 2.

Following O'Donoghue and Rabin (1999, 2001), I adopt the solution concept of perception-perfect strategies:

Definition 3.1 *Consider a worker of type $\phi \in \{TC, N, S\}$. Let $s^*_{1,\phi}(\hat{s}_2)$ denote the optimal period-1 strategy of the worker given her perception $\hat{s}_2$ about her period-2 strategy, and let $s^*_{2,\phi}$ denote her optimal period-2 strategy.*

- *For $\phi \in \{TC, S\}$, the perception-perfect strategy is*

$$s^*_\phi = \big(s^*_{1,\phi}(s^*_{2,\phi}),\ s^*_{2,\phi}\big).$$

- *For $\phi = N$, the perception-perfect strategy is*

$$s^*_N = \big(s^*_{1,N}(s^*_{2,TC}),\ s^*_{2,N}\big).$$

For each worker type, a perception-perfect strategy requires the period-1 search decision at every current wage and, conditional on search, the period-1 acceptance decision after every realized offer to maximize expected payoff given the worker's perception of her period-2 strategy. Time-consistent and sophisticated workers have correct perceptions about their future type, so their period-1 strategy is based on their actual period-2 strategy. In contrast, a naive worker believes that her future self will be time-consistent and therefore bases her period-1 strategy on $s^*_{2,TC}$, even though in period 2 she actually follows $s^*_{2,N}$.

Throughout the paper, I resolve indifference between searching and not searching in favor of search, and indifference between accepting and rejecting an offer in favor of

rejection. All uniqueness claims for strategies should therefore be understood under this tie-breaking convention.[4]

3.2.4. Results

Period 2. Period 2 is the last period in which the worker makes decisions. Hence, optimal behavior does not depend on perceptions about future strategies, and the continuation payoff simply equals the discounted wage earned in period 3. Conditional on searching, all worker types therefore accept a job offer if and only if $w'' > w_2$. Moreover, since the gain from searching rather than not searching is strictly decreasing in the current wage, the period-2 search strategies are characterized by type-specific search thresholds. Lemma 3.1 compares the optimal period-2 strategies across types.

Lemma 3.1 *For the optimal period-2 strategy, $s^*_{2,\phi}$, all worker types $\phi \in \{TC, N, S\}$ search if and only if $w_2 \leq w^*_{2,\phi}$, with*

$$w^*_{2,S} = w^*_{2,N} < w^*_{2,TC}.$$

Moreover, all worker types accept a job offer with wage w'' if and only if $w'' > w_2$.

Proof. The result follows directly from the full characterization of the period-2 strategies in Lemma C.1 in Appendix C.1.1. □

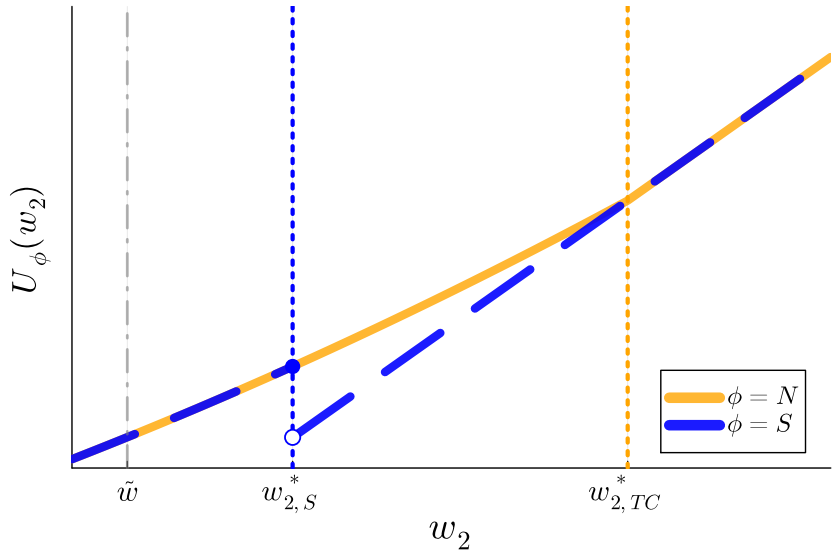
The time-consistent type of worker has a strictly higher search threshold than the present-biased types. The naive and sophisticated workers share the same threshold because perceptions about future behavior, and therefore about future present bias, are irrelevant in period 2. The difference relative to the time-consistent worker arises from the investment character of on-the-job search: search entails an immediate cost in exchange for a potential future wage gain. Present-biased workers discount this future gain more heavily relative to the current search cost, which lowers the wage level below which search is optimal for them.

Period 1. In period 1, the optimal strategy of a worker depends on her beliefs about her period-2 strategy. A time-consistent worker and a sophisticated worker correctly anticipate their period-2 strategies, whereas a naive worker mistakenly believes that her period-2 self will follow the time-consistent period-2 strategy.

From the perspective of period 1, these beliefs determine the continuation payoffs from entering period 2 at wage w_2. For wages between the period-2 search thresholds of

[4]Since this tie-breaking convention only selects actions at points of exact indifference, it does not affect payoffs. It only avoids having to list payoff-equivalent strategies that differ in their prescribed behavior at such points of indifference. Moreover, since offers are drawn from a continuous distribution, the acceptance convention concerns only probability-zero events.

Figure 3.1: Continuation Payoffs from the Period-1 Perspective



Note: The figure plots the continuation payoffs for types $\phi \in \{N, S\}$ from the perspective of period 1, $U_\phi(w_2)$, as a function of the period-2 wage w_2. The blue long-dashed line shows the continuation payoff of the sophisticated type, and the orange solid line shows the continuation payoff of the naive type. The finely dashed vertical lines indicate the type-specific expected period-2 search thresholds from the perspective of period 1. The gray dash-dotted vertical line indicates $\tilde{w}$, the lower bound of the period-1 wage interval on which the sophisticated worker's acceptance strategy is non-monotone. The continuation payoff of the time-consistent type is omitted for clarity, as it differs from that of the naive type only by multiplication with $1/\beta$.

time-consistent and present-biased workers, $w_2 \in (w^*_{2,S}, w^*_{2,TC})$, time-consistent and naive workers expect to search in period 2, while a sophisticated worker correctly anticipates that her future present bias will prevent her from searching. From the period-1 perspective, all types discount period-3 payoffs relative to period 2 by the exponential discount factor δ. Hence, evaluated from period 1, all types view search in period 2 as optimal for all $w_2 \leq w^*_{2,TC}$. A sophisticated worker therefore anticipates suboptimal non-search by her period-2 self whenever $w_2 \in (w^*_{2,S}, w^*_{2,TC})$. This lowers her continuation payoff for such wages relative to the continuation payoff from searching in period 2. By contrast, a naive worker expects to search in period 2 throughout this interval and therefore assigns a higher continuation payoff to these wages than the sophisticated worker.

This anticipation of suboptimal period-2 behavior generates a downward discontinuity in the continuation payoff of the sophisticated worker at her period-2 search threshold. This discontinuity, illustrated in Figure 3.1, is at the heart of the result that sophisticated workers may reject wage increases. Suppose the worker's period-1 wage lies below $w^*_{2,S}$ and she receives an offer just above $w^*_{2,S}$. Accepting the offer raises the period-2 wage but also eliminates future search. Since search at such wages is optimal from the period-1 perspective, entering period 2 with the new wage may imply a lower continuation pay-

off than entering period 2 with the current wage and continuing to search. Hence, the sophisticated worker may reject the higher offer.[5]

To characterize the set of period-1 wages for which this non-monotone acceptance behavior can occur, let $U_S(w)$ denote the continuation payoff of the sophisticated worker from the period-1 perspective of entering period 2 with wage w, under her anticipated period-2 behavior. Define

$$\tilde{w} \coloneqq \inf\{w \in [w_l, w^*_{2,S}] : U_S(w) \geq \lim_{z \downarrow w^*_{2,S}} U_S(z)\}. \tag{3.1}$$

Thus, $\tilde{w}$ is the lowest wage at which the sophisticated worker weakly prefers, from the period-1 perspective, entering period 2 at that wage and searching to entering period 2 just above $w^*_{2,S}$ and not searching.[6] For any $w_1 \in [\tilde{w}, w^*_{2,S}]$, define $\psi(w_1) \geq 0$ as the unique value such that

$$U_S(w_1) = \lim_{z \downarrow (w^*_{2,S} + \psi(w_1))} U_S(z). \tag{3.2}$$

Hence, $\psi(w_1) \geq 0$ is the wage increment above $w^*_{2,S}$ that makes the worker indifferent, from the period-1 perspective, between entering period 2 at wage w_1 and searching, and entering period 2 at wage $w^*_{2,S} + \psi(w_1)$ and not searching.[7] Then, for any $w_1 \in [\tilde{w}, w^*_{2,S}]$, offers $w' \in (w^*_{2,S}, w^*_{2,S} + \psi(w_1))$ raise the wage but still lower the continuation payoff because they induce the worker to suboptimally stop searching in period 2 without providing a sufficiently large wage gain. The sophisticated worker therefore rejects these offers.[8] By contrast, she accepts offers $w' \in (w_1, w^*_{2,S}]$, which raise the wage without changing future search behavior, and offers $w' > w^*_{2,S} + \psi(w_1)$, whose wage gain is large enough to compensate for giving up future search. Offers $w' \leq w_1$ are rejected trivially.

A time-consistent worker anticipates period-2 behavior that is also optimal from the period-1 perspective. Hence, any wage increase raises her continuation payoff, and she accepts all offers with $w' > w_1$. A naive worker behaves analogously because she believes

[5]This behavior is commitment-device-like in spirit, since the sophisticated worker uses a current choice to enforce future behavior. It is not a commitment device in the strict sense often used in the literature, however, since commitment devices typically restrict future choice sets either by removing tempting options or by imposing penalties for undesired behavior (e.g., Bryan et al., 2010; Gul and Pesendorfer, 2001; Laibson, 2015). In the mechanism here, rejecting the offer does not restrict future choices. Instead, the worker preserves search incentives by avoiding a state in which her future self would stop searching.

[6]The downward discontinuity, together with continuity and strict monotonicity of U_S on the search region, implies that $\tilde{w}$ is always well defined. If the sophisticated worker strictly prefers entering period 2 at w_l and searching to this right-limit no-search payoff, then $\tilde{w} = w_l$. Otherwise, $\tilde{w} \in (w_l, w^*_{2,S})$ and the worker is indifferent between the two continuation payoffs. See Figure 3.1 for an illustration of $\tilde{w}$.

[7]Strict inequality $\psi(w_1) > 0$ holds whenever $w_1 \in (\tilde{w}, w^*_{2,S}]$. Equation (3.2) can then be written simply as $U_S(w_1) = U_S(w^*_{2,S} + \psi(w_1))$.

[8]If $\psi(w_1) > 0$, then for an offer equal to the upper bound, $w' = w^*_{2,S} + \psi(w_1)$, she is indifferent and rejects by the tie-breaking convention.

that her period-2 self will follow the time-consistent period-2 strategy. Therefore, she does not anticipate that accepting a wage just above $w^*_{2,S} = w^*_{2,N}$ will deter future search. Thus, both types accept all offers with $w' > w_1$. Using $\tilde{w}$ and $\psi(w_1)$ as characterized in (3.1) and (3.2), Proposition 3.1 formalizes these insights.

Proposition 3.1 *Suppose the worker earns wage w_1 in period 1. The perception-perfect period-1 job offer acceptance strategies satisfy the following:*

(i) *The time-consistent and naive types accept a job offer with wage w' if and only if $w' > w_1$.*

(ii) *The sophisticated type may reject certain wage increases. In particular, for every $w_1 \in [\tilde{w}, w^*_{2,S}]$, she rejects all offers $w' \in (w^*_{2,S},\ w^*_{2,S} + \psi(w_1)]$.*

PROOF. The result follows directly from the full characterization of the period-1 acceptance strategies in Lemma C.2 in Appendix C.1.2. □

The full characterization, including cases outside the interval $[\tilde{w}, w^*_{2,S}]$, is given by Lemma C.2 in Appendix C.1.2. For wages below $\tilde{w}$, the sophisticated worker accepts any wage increase. For wages just above $w^*_{2,S}$, she may also accept wage decreases that restore future search. This case can arise because, as shown below, the non-stationarity of the three-period environment makes the period-1 search threshold exceed the period-2 search threshold. This, however, is a purely technical implication of the non-stationarity of the finite-horizon model so I do not interpret it as a meaningful prediction of the model.

The main insight of Proposition 3.1 is that, while naive and time-consistent workers accept any wage increase in period 1, sophisticated workers may be more selective and reject some wage increases because they anticipate when a new job would suboptimally distort future search and thereby make the job permanent. Naive workers, by contrast, may accept relatively small wage increases under the mistaken belief that they will continue searching on the job, but subsequently become trapped in jobs they initially perceived as temporary steps on the wage ladder. This difference in acceptance behavior yields the testable prediction that sophisticated workers, on average, demand larger wage increases to change jobs than naive and time-consistent workers.

Having characterized the period-2 strategies and period-1 acceptance decisions, it remains to derive period-1 search behavior. As in period 2, the gain from searching is strictly decreasing in the current wage, so each worker type follows a threshold rule: type ϕ searches in period 1 if and only if $w_1 \leq w^*_{1,\phi}$. Proposition 3.2 compares these thresholds across types.

Proposition 3.2 *The perception-perfect period-1 search strategies are characterized by thresholds. Each worker type $\phi \in \{TC, N, S\}$ searches if and only if $w_1 \leq w^*_{1,\phi}$. For each type, $w^*_{1,\phi} > w^*_{2,\phi}$. The thresholds satisfy the following comparisons across types:*

(i) *The time-consistent type has the highest period-1 search threshold:*

$$w^*_{1,TC} > w^*_{1,\phi} \qquad \text{for } \phi \in \{N, S\}.$$

(ii) *The sophisticated type has a weakly higher period-1 search threshold than the naive type:*

$$w^*_{1,S} \geq w^*_{1,N},$$

with strict inequality if and only if $\beta < 1/(1+\delta)$.

PROOF. The result follows directly from the full characterization of the period-1 search strategies in Lemma C.3 in Appendix C.1.3. □

Two insights emerge from Proposition 3.2. First, as in period 2, the period-1 search thresholds of present-biased workers lie strictly below the period-1 threshold of the time-consistent worker. Again, this is a direct implication from present-biased workers discounting future gains from search more heavily relative to the current search cost than the time-consistent type of worker.

Second, the search threshold of the sophisticated worker is weakly above that of the naive worker. This inequality is strict whenever the present bias is sufficiently strong, that is, if β is sufficiently small. In this case, the thresholds of both present-biased types lie in the interval $(w^*_{2,S}, w^*_{2,TC})$.[9] For wages in this interval, the two present-biased types differ in their beliefs about period-2 behavior: the naive worker mistakenly expects to search in period 2 and therefore believes that not searching in period 1 still leads to future search. This creates an incentive to postpone search to the next period. In contrast, the sophisticated worker correctly anticipates that, if she does not search in period 1 and therefore enters period 2 at the same wage, she will not search in period 2. This creates a stronger incentive for her to search in period 1. As a result, $w^*_{1,N} < w^*_{1,S}$.

If β is sufficiently large, the period-1 thresholds of the two present-biased types coincide. For sufficiently mild present bias, both present-biased workers find search worthwhile in period 1 up to wages strictly above the time-consistent period-2 threshold $w^*_{2,TC}$. At such wages, both workers are above their expected period-2 thresholds: the naive worker expects $w^*_{2,TC}$, while the sophisticated worker anticipates $w^*_{2,S} < w^*_{2,TC}$. Both types therefore anticipate not searching in period 2 at these wages, which is also optimal from the period-1 perspective. Hence, the anticipated future behavior coincides for the sophisticated and

[9] Recall that $w^*_{2,S} = w^*_{2,N}$.

naive workers despite the naive worker's misperception about her future search threshold. Their continuation payoffs therefore coincide, so the two types have the same period-1 search threshold. Importantly, this case can arise only because of the non-stationarity of the finite-horizon environment, which implies $w^*_{1,TC} > w^*_{2,TC}$. As a result, the period-1 search thresholds of present-biased workers can lie above the period-2 time-consistent threshold while still lying below the period-1 time-consistent threshold.

While the three-period model provides a tractable framework to analyze on-the-job search and acceptance behavior across worker types, it abstracts from unemployment and therefore cannot speak to job search by unemployed workers. Existing models study the search behavior of present-biased unemployed workers without the option of continued on-the-job search (DellaVigna and Paserman, 2005). Incorporating an unemployment stage into the framework allows to study unemployed job search when present bias continues to affect search and acceptance behavior after workers enter employment, a channel absent from this existing work. Moreover, the three-period model cannot deliver predictions for wage outcomes over the career: without an initial unemployment stage, any assumption about the worker's starting wage would be arbitrary yet important for final outcomes, and wages also evolve only over a very limited horizon.

I therefore next extend the framework to an infinite-horizon model in which workers start unemployed. This extension allows the analysis to characterize reservation wages and the wage distributions generated by type-specific search and acceptance behavior.

3.3. Infinite-Horizon Model

The worker lives for infinitely many periods. In the initial period $t = 1$, the worker is unemployed, receives unemployment benefits $b \in [0, w_h)$, and draws a job offer with wage $w' \sim \mathcal{U}[w_l, w_h]$.[10] If the worker accepts the offer, she becomes employed in period $t = 2$ at wage $w_2 = w'$ and from then onward the decision environment in each period mirrors that of the finite horizon model. If she rejects the offer, she remains unemployed and the period 1 environment repeats.

3.3.1. Strategies and Solution Concept

As is standard in dynamic games with stationary environments, I restrict attention to stationary Markov strategies (e.g., Harris and Laibson, 2001; Maskin and Tirole, 2001). For the on-the-job decisions studied first, a stationary Markov strategy of an employed

[10] One could alternatively assume a search cost $c_u > 0$ while unemployed. However, c_u would have to be sufficiently small to ensure that all worker types search when unemployed. Then, changes in c_u only affect the value of unemployment, equivalent to changing b. Search costs when unemployed are thus disregarded, implicitly setting $c_u = 0$.

worker of type $\phi \in \{TC, N, S\}$ is a pair

$$s_\phi := (q_\phi, a_\phi),$$

where $q_\phi : \mathcal{W} \to \{0, 1\}$ is a search rule and $a_\phi : \mathcal{W} \times \mathcal{W} \to \{0, 1\}$ is an acceptance rule, with the same interpretation as in the three-period model. Stationarity means that the same functions prescribe behavior in every period: search is a function only of the current wage (the state variable), and, conditional on search, offer acceptance is a function only of the current wage and the realized offer. All results below should be interpreted within this class of strategies.

Definition 3.2 *Let Σ denote the set of pairs $s = (q, a)$ as defined above. For $\hat{s} \in \Sigma$, let $B_\phi(\hat{s}) = (q_\phi^{\hat{s}}, a_\phi^{\hat{s}}) \in \Sigma$ denote the current best-response pair of an employed worker of type ϕ, given that all future employed selves are perceived to use $\hat{s}$ in every future period.*

- *For $\phi \in \{TC, S\}$, a perception-perfect stationary Markov strategy satisfies*

$$s_\phi^* = B_\phi(s_\phi^*).$$

- *For $\phi = N$, the perception-perfect stationary Markov strategy satisfies*

$$s_N^* = B_N(s_{TC}^*).$$

Henceforth, I refer to (perception-perfect) stationary Markov strategies simply as (perception-perfect) strategies.[11]

Under perception-perfectness the worker chooses an optimal strategy in each period given her perception of her future strategy. For time-consistent and sophisticated workers, this perception is correct. Hence, the strategy anticipated to be followed in the future must coincide with the strategy used by the current self. For naive workers, the current self optimizes under the perception that future selves will follow the time-consistent strategy.

For time-consistent and naive workers, the infinite-horizon model admits a unique pure perception-perfect on-the-job search strategy. Search behavior is characterized by thresholds, w_{TC}^* and w_N^*, respectively, and conditional on searching, any offer above the current wage is accepted. For the sophisticated worker, this is no longer the case.

Lemma 3.2 *There exists no pure perception-perfect on-the-job search strategy for the sophisticated worker in the infinite-horizon model.*

[11] When unemployed workers are studied below, I use the same stationarity convention. The unemployed worker's offer acceptance decision is characterized by a stationary reservation wage, given the perceived stationary continuation strategies after accepting or rejecting an offer.

PROOF. See Appendix C.3.1. The proof uses Lemma C.6 in Appendix C.2, which shows that, under existence, a perception-perfect strategy would need to have a threshold structure. It then rules out all such threshold strategies as perception-perfect strategies. □

The non-existence result arises because any candidate threshold must satisfy two consistency requirements at once: search must be optimal at all wages in the search region, given the belief that future selves will search at those wages, and no-search must be optimal at all wages in the no-search region, given the belief that future selves will not search at those wages. These requirements cannot be jointly satisfied.

To see why, suppose that search is consistent up to a candidate threshold w_S, so that the sophisticated worker searches if and only if $w \leq w_S$. Since search is consistent for the sophisticated worker, she must weakly prefer searching at the threshold to not searching, given the belief that all future selves follow this threshold strategy. When making this current search decision, she discounts the future gains from search by $\beta\delta$. In the continuation payoff generated by future behavior, however, the future self evaluates the same search decision with discount factor δ between future periods. Hence, if search is weakly optimal at the threshold for the current sophisticated self, search at the threshold must yield a strictly higher continuation payoff than keeping the wage forever. It follows that the continuation payoff drops discontinuously when crossing the threshold from below.

This drop generates a discontinuous increase in the gain from a one-time deviation to search just above the threshold. By searching in the current period, the worker has a positive probability of drawing an offer above her current wage and thereby leaving the low continuation value associated with permanent no-search at the current wage. This makes a deviation to search strictly profitable, violating perception-perfectness.[12]

This logic is specific to sophistication. For a time-consistent worker, the continuation value is continuous at the search threshold because current and future selves evaluate the cost of search relative to its gains in the same way and, at the threshold, the worker is indifferent between searching and not searching. For a naive worker, perceived future behavior differs from actual future behavior: she believes that future selves follow the time-consistent strategy. Her perceived continuation payoff is therefore proportional to the time-consistent continuation payoff and hence also continuous. Thus, the discontinuity-based inconsistency does not arise for these types.

The non-existence result is local and reflects an inconsistency between current optimal-

[12] A deviation to search may also generate offers below w_S, which can return the worker to the search region and restore the higher continuation payoff associated with future search. The proof does not rely on this channel; strict profitability of a deviation to search already follows from the positive probability of drawing an offer above the current wage. Also, the argument does not rely on assigning the threshold itself to the search region. If the candidate strategy prescribed no-search at w_S, the same discontinuity argument applies by comparing wages just below w_S with w_S itself.

ity and beliefs about the future strategy in a neighborhood of any candidate threshold. At the same time, the underlying threshold logic remains valid away from this region: there exist intervals of wages over which search or no-search is consistent with the corresponding beliefs. I therefore introduce a new refinement of perception-perfectness that identifies wage regions over which current behavior and beliefs about the future strategy are consistent and imposes a maximality requirement on the size of such regions.

Definition 3.3 *A pair $(\underline{w}; \bar{w})$ with $\underline{w} < \bar{w}$ is locally perception-perfect if, under the conjecture that future selves search if and only if $w \leq \underline{w}$, search (followed by optimal offer acceptance) is indeed optimal for all $w \leq \underline{w}$ and no-search is optimal for all $w \geq \bar{w}$. A locally perception-perfect pair $(\underline{w}; \bar{w})$ is maximal under the following two conditions:*

(i) **Maximal consistency of search:** *$\underline{w}$ is the largest threshold such that search is consistent for all $w \leq \underline{w}$.*

(ii) **Maximal consistency of no-search:** *$\bar{w}$ is the smallest wage such that no-search is consistent for all $w \geq \bar{w}$.*

3.3.2. Results

Characterization of the Locally Perception-Perfect Pair. The next lemma characterizes the unique maximal locally perception-perfect pair for the sophisticated worker.

Lemma 3.3 *There exists a unique maximal locally perception-perfect pair for the sophisticated worker. It is characterized by*

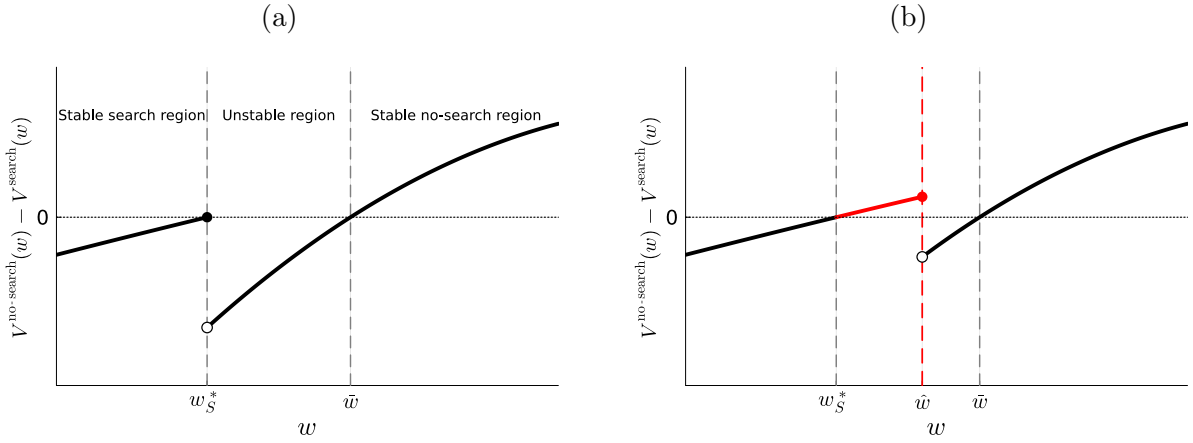
$$\underline{w} := w_S^* = w_h - \sqrt{\frac{2(w_h - w_l)(1-\delta)c}{\beta\delta}} - \frac{c(1-\beta)}{\beta}, \qquad \bar{w} := w_S^* + \frac{c(1-\beta)}{\beta}.$$

Proof. See Appendix C.3.2. □

The pair $(w_S^*; \bar{w})$ induces two regions within which current behavior and beliefs about the future strategy are consistent. For $w \in [w_l, w_S^*]$, search is optimal under the conjecture that future selves search if and only if their wage is at most w_S^*. For $w \in [\bar{w}, w_h]$, no-search is optimal under the same conjecture. Moreover, w_S^* is the largest threshold that sustains such consistent search behavior for all wages below w_S^*. Similarly, $\bar{w}$ is the smallest wage that sustains no-search as consistent behavior for all wages above $\bar{w}$. Hence, both the consistent search region and the consistent no-search region are maximal under $(w_S^*; \bar{w})$.

Figure 3.2 illustrates this characterization. Panel (a) plots the current-period payoff difference between not searching and searching for given wages. For wages $w \leq w_S^*$, the difference is evaluated under the conjecture that the wage itself constitutes the future search threshold. Since, for any fixed conjectured threshold, the payoff difference is strictly

Figure 3.2: Maximal Local Perception-Perfectness



Note: Panel (a) plots the current-period payoff difference between not searching and searching, $V^{\text{no-search}}(w) - V^{\text{search}}(w)$, as a function of the current wage w. The gray vertical dashed lines indicate the wages defining the maximal locally perception-perfect pair, $(w_S^*, \bar{w})$. For wages $w \leq w_S^*$, the difference is evaluated under the conjecture that the current wage w itself constitutes the future search threshold. For wages $w > w_S^*$, the difference is evaluated under the conjecture that the worker does not search in the future at wage w and, if she deviates to search today, she accepts all and only higher wage offers. Panel (b) plots the same payoff difference for candidate thresholds above w_S^*. The red vertical dashed line indicates one particular alternative candidate threshold $\hat{w} > w_S^*$.

increasing in the current wage, a threshold defines a consistent search region whenever the plotted payoff difference is weakly negative.[13] In Panel (a), the difference is negative for all $w < w_S^*$ and zero at w_S^*, so w_S^* induces a consistent, or stable, search region $[w_l, w_S^*]$.

Panel (b) shows that the threshold cannot be extended beyond w_S^*. If a higher wage $\hat{w}$ is conjectured to be the threshold, then no-search is strictly optimal at $\hat{w}$. By continuity, the same holds in a neighborhood below $\hat{w}$, contradicting consistency of search. Hence, w_S^* is maximal.

For wages $w > w_S^*$, Panel (a) plots, under the conjecture that future selves do not search at these wages, an upper bound on the current-period payoff difference between not searching and searching, obtained by assuming that, when deviating to search, the worker accepts only higher offers. This assumption weakly lowers the value of search because it rules out the option of accepting lower wages to re-enter the search region. Hence, no-search cannot be consistent at a given wage if this upper bound difference is negative. In Panel (a), the upper bound is negative on $(w_S^*, \bar{w})$, so no-search cannot be consistent on this interval. At $w = \bar{w}$, the upper bound equals zero, and for all $w > \bar{w}$ it is strictly positive, so that $\bar{w}$ is the smallest candidate to define a stable no-search region. In fact, for $w \geq \bar{w}$, accepting lower wages is never optimal when deviating to search under

[13] It follows from the proof of Lemma 3.3 in Appendix C.3.2 that, for any fixed conjectured threshold, the search gain $V^{\text{search}}(w) - V^{\text{no-search}}(w)$ is strictly decreasing and continuous in w for wages weakly below that threshold. Consequently, the negative of this difference, $V^{\text{no-search}}(w) - V^{\text{search}}(w)$, is strictly increasing and continuous on that region.

the conjectured search threshold w_S^*, so the upper bound coincides with the true payoff difference. Thus, no-search is indeed consistent for all $w \geq \bar{w}$ under search threshold w_S^* so that $\bar{w}$ is the minimal wage such that $[\bar{w}, w_h]$ constitutes a stable no-search region.

The preceding discussion implies that wages in the interval $(w_S^*, \bar{w})$ are inherently unstable for the sophisticated worker: if such a wage is assigned to the no-search region, a one-time deviation to search is strictly profitable under the conjecture that future selves do not search at these wages. Conversely, if a wage in this interval is made the conjectured search threshold, so that search is expected at that wage in future periods, a one-time deviation to no-search becomes optimal.

Corollary 3.1 *For the sophisticated worker, every wage $w \in (w_S^*, \bar{w})$ is inherently unstable in the following sense:*

(i) *If w is conjectured to lie in the no-search region, so that future selves are not expected to search at w, then search is optimal at w in the current period.*

(ii) *If w is conjectured to lie in the search region as its upper boundary, so that the conjectured search region is $[w_l, w]$, then no-search is optimal at w in the current period.*

PROOF. See Appendix C.3.3. □

Behavior of the sophisticated worker in the unstable interval $(w_S^*, \bar{w})$ is thus not pinned down by local perception-perfectness. In this region, neither search nor no-search can be sustained as consistent behavior. The following analysis therefore focuses on the stable regions $[w_l, w_S^*]$ and $[\bar{w}, w_h]$.

On-the-Job Search and Acceptance Behavior. The model yields three main results for on-the-job search and acceptance behavior across types. Because these results closely mirror those in the three-period model, formal statements and proofs are relegated to the appendix.

First, present-biased workers stop searching at lower wages than the time-consistent worker. For the naive type, this follows from $w_N^* < w_{TC}^*$ (see Lemma C.4 and Lemma C.5). For the sophisticated type, the comparison is based on the onset of the stable no-search region, which satisfies $\bar{w} < w_{TC}^*$ (see Lemma C.4 and Lemma 3.3). Thus, stable no-search begins at a lower wage for the sophisticated worker than for the time-consistent worker. As discussed above, this result arises because present-biased workers place greater weight on the immediate costs of search relative to its future benefits.

Second, the stable no-search region begins at a higher wage for the sophisticated worker than for the naive worker: while the naive worker stops searching above w_N^*, no-search is consistent for the sophisticated worker only from $\bar{w}$ onward, with $w_N^* <$

$\bar{w}$ (see Lemma C.7). As in the three-period model, this reflects that the naive worker has an incentive to postpone search, whereas the sophisticated worker anticipates that postponement implies stopping search altogether.

Third, offer acceptance behavior differs across types (see Lemma C.8). For time-consistent and naive workers, conditional on searching, any offer $w' > w$ is accepted and any $w' \leq w$ is rejected since both expect their future selves to behave optimally from the current period's perspective. By contrast, the sophisticated worker exhibits non-monotone acceptance behavior in the stable search region: for wages sufficiently close to w_S^* from below, there exists a nonempty interval of higher wage offers that are rejected. As in the three-period model, this reflects that the sophisticated worker anticipates that certain wage increases would eliminate future search without being sufficiently large to justify accepting it as a permanent wage. Remaining at a lower wage therefore commits the sophisticated worker to future search by keeping the future self in the region in which search remains optimal. This rejection of higher offers arises precisely for wage offers that would move the worker into the unstable region $(w_S^*, \bar{w})$, implying endogenous avoidance of that region.

As a remark, the relative position of w_N^* and w_S^* is not characterized. It may therefore be that $w_N^* > w_S^*$, in which case the naive threshold lies in the unstable region $(w_S^*, \bar{w})$ of the sophisticated worker. In this region, the behavior of the sophisticated worker is not pinned down by the model, which limits meaningful comparisons of search behavior across the two types. Hence, the robust implication of the infinite-horizon model is that the onset of consistent no-search occurs at higher wages for the sophisticated worker than for the naive worker, while comparisons of search behavior itself are not robust.

Reservation Wages. I next turn from the analysis of on-the-job search behavior to offer acceptance decisions out of unemployment. Define the reservation wage r_ϕ^* as the minimum wage that an unemployed worker of type ϕ is willing to accept.[14] Proposition 3.3 characterizes the reservation wages of all types, given their perceived future on-the-job search behavior. While perceived future on-the-job search behavior is only locally consistent for the sophisticated worker, the reservation wage is optimal given the perception induced by the locally perception-perfect pair. I therefore refer to the reservation wages as perception-perfect for all types. Recall that $b \in [0, w_h)$ denotes unemployment benefits and $c > 0$ denotes the cost of on-the-job search. Also recall that w_{TC}^* denotes the time-consistent on-the-job search threshold, and that $(w_S^*, \bar{w})$ denotes the unique maximal locally perception-perfect pair of the sophisticated worker.

[14]This is stated explicitly here because, for the sophisticated type, acceptance behavior of the unemployed may also be non-monotone. For acceptance behavior of the unemployed, I assume that offers are accepted in case of indifference.

Proposition 3.3 *The perception-perfect reservation wages are characterized as follows:*

(i) *For* $\phi \in \{TC, N\}$*:*

(a) *If* $b + c \leq w^*_{TC}$*, then* $r^*_\phi = b + c$*.*

(b) *If* $b + c > w^*_{TC}$*, then* r^*_ϕ *is the unique solution to*

$$r^*_\phi = b + \frac{\delta}{1-\delta}\mathbb{P}(w' \geq r^*_\phi)\big(\mathbb{E}[w' \mid w' \geq r^*_\phi] - r^*_\phi\big) > w^*_{TC}.$$

(ii) *For* $\phi = S$*:*

(a) *If* $b + c \leq w^*_S$*, then* $r^*_S = b + c$*.*

(b) *If* $b + c > w^*_S$*, then* r^*_S *is the unique solution to*

$$r^*_S = b + \frac{\delta}{1-\delta}\mathbb{P}(w' \geq r^*_S)\big(\mathbb{E}[w' \mid w' \geq r^*_S] - r^*_S\big) > \bar{w}.$$

PROOF. See Appendix C.3.4. □

The time-consistent and naive types have identical reservation wages. This arises because accepting or rejecting an offer does not affect the instantaneous utility, which always equals the unemployment benefit b, but only future payoffs. Since the naive worker expects future selves to behave time-consistently, she effectively compares the same continuation values as the time-consistent worker when deciding whether to accept an offer, so both types set the same reservation wage.

For all types, two cases arise in the determination of the reservation wage. Consider an unemployed worker who receives an offer below her type-specific expected on-the-job search threshold. Accepting such an offer does not change her expected future search behavior: the worker expects to continue searching in the next period regardless of whether she accepts or rejects the offer. Hence, the acceptance decision depends only on the per-period payoff difference between accepting and rejecting. Accepting an offer of wage w yields $w - c$, while rejecting yields b. Thus, if $b + c$ lies below the type-specific on-the-job search threshold, the reservation wage equals $r^*_\phi = b + c$. This yields parts (i)(a) and (ii)(a) of Proposition 3.3.

Next, suppose that $b+c$ exceeds the type-specific expected on-the-job search threshold. Then no wage below this threshold is accepted, because rejecting the offer yields a higher per-period payoff. Hence, any acceptable wage must lie above the threshold. Accepting such a wage implies that the first job is perceived as permanent, because the worker expects no further on-the-job search. By contrast, rejecting the offer and remaining unemployed preserves continued search. Acceptance must therefore compensate the worker

both for the forgone unemployment benefit and for the value of continuing to search in unemployment and thus possibly generating higher future offers. This yields the fixed-point equations that define the reservation wage in parts (i)(b) and (ii)(b) of Proposition 3.3.

Therefore, the comparison of reservation wages across types depends on how $b + c$ compares to the relevant on-the-job search thresholds and, thus, on perceived future search behavior after accepting a job. If $b + c$ lies below the on-the-job search thresholds of all types, all types expect to continue searching after accepting a wage of $b + c$ and therefore set reservation wage $b + c$. If $b + c$ lies above the time-consistent threshold w^*_{TC}, all types anticipate that accepting a job ends future search. The reservation wage must then compensate a worker of any type for the value of future search that she gives up by leaving unemployment, again yielding equal reservation wages across types.

Differences arise when $b + c \in (w^*_S, w^*_{TC})$. In this region, the time-consistent worker expects continued on-the-job search after accepting wage $b + c$, and the naive worker expects the same continuation behavior. For the naive worker, this perception is mistaken whenever $b + c > w^*_N$: she then accepts the job as a temporary step on the wage ladder but then does not search on the job, so the job becomes permanent.[15] By contrast, the sophisticated worker expects that accepting such a job ends future on-the-job search and makes the job permanent. She therefore demands a higher wage to leave unemployment.[16]

Corollary 3.2 *The perception-perfect reservation wage of the sophisticated worker is weakly higher than that of the time-consistent and naive workers. Specifically,*

- *If* $b + c \leq w^*_S$ *or* $b + c \geq w^*_{TC}$, *then* $r^*_S = r^*_{TC} = r^*_N$.
- *If* $b + c \in (w^*_S, w^*_{TC})$, *then* $r^*_S > r^*_{TC} = r^*_N$.

Proof. See Appendix C.3.5. □

Allowing for on-the-job search therefore introduces a mechanism that can reverse the reservation-wage ranking in DellaVigna and Paserman (2005), where sophisticated workers anticipate reduced search while unemployed and therefore accept lower wages to leave unemployment sooner. In more general frameworks with both unemployment search and on-the-job search margins, the overall ranking of reservation wages by type may therefore depend on the relative strength of these opposing mechanisms.

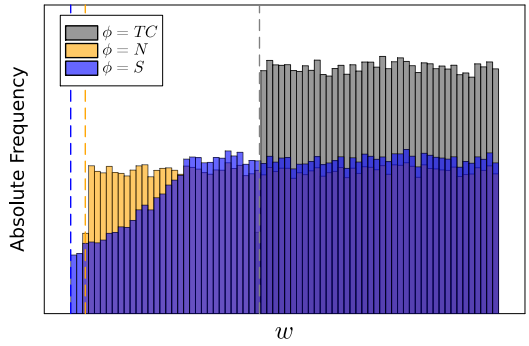
Wage Distributions. The model discussion concludes with simulations that illustrate how the mechanisms derived above affect wage distributions across types. The simulations

[15] The relative position of w^*_N and w^*_S is not characterized in general. However, $w^*_N < w^*_{TC}$ by Lemma C.5 in Appendix C.2. Hence, the interval (w^*_N, w^*_{TC}), on which the naive worker accepts jobs based on mistaken perceptions about future on-the-job search, is nonempty.

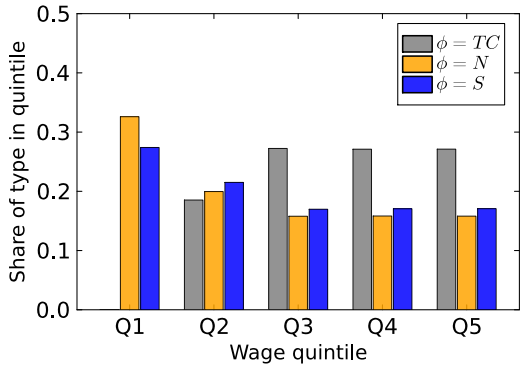
[16] Appendix C.4 illustrates the corresponding comparative statics of reservation wages with respect to unemployment benefits, on-the-job search costs, and present bias.

Figure 3.3: Simulated Final Wage Distributions

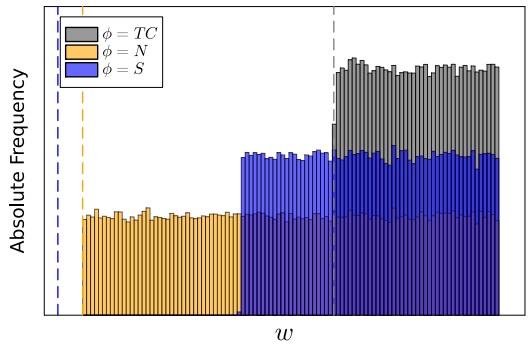
(a) $b + c < w_S^*$: Histogram



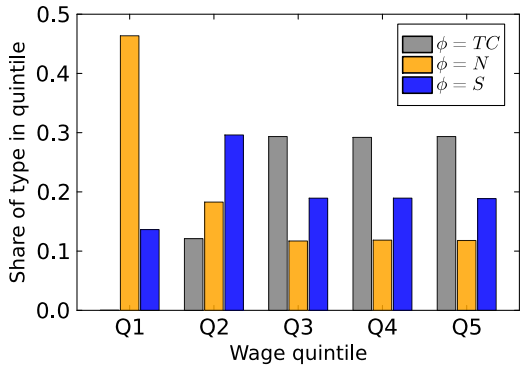
(b) $b + c < w_S^*$: Shares by Quintiles



(c) $b + c \in (w_S^*, w_{TC}^*)$: Histogram



(d) $b + c \in (w_S^*, w_{TC}^*)$: Shares by Quintiles



Note: The figure shows simulated final wage outcomes for 100,000 workers of each type. Panels (a) and (c) plot histograms of final wages by type. Panels (b) and (d) plot, for each type, the share of workers in each quintile of the overall final wage distribution. The first row corresponds to the case $b + c < w_S^*$, where all types have the same reservation wage, using the illustrative parameterization $(w_l, w_h, c, \delta, \beta, b) = (0, 1, 0.15, 0.88, 0.6, 0)$. The second row corresponds to the case $b + c \in (w_S^*, w_{TC}^*)$, where sophisticated workers set a higher reservation wage than time-consistent and naive workers, using the illustrative parameterization $(w_l, w_h, c, \delta, \beta, b) = (0, 1, 0.15, 0.88, 0.4, 0.31)$. Orange bars denote naive workers, blue bars denote sophisticated workers, and gray bars denote time-consistent workers. The vertical dashed lines in Panels (a) and (c) indicate the type-specific on-the-job search thresholds w_S^*, w_N^*, and w_{TC}^*.

are not intended to match observed moments of wage distributions quantitatively, but instead illustrate the qualitative patterns implied by the model.

Careers are simulated for 100,000 workers of each type. In each period, workers who search receive a wage offer drawn independently from the uniform distribution $\mathcal{U}[w_l, w_h]$. Each career is simulated for 100 periods, with each worker starting as unemployed, which is sufficient for virtually all workers to reach their final job. The wage in the final period is then used to construct the wage distributions. For the sophisticated type, the simulations assume no-search in the unstable region $(w_S^*, \bar{w})$, so that wages in this region are absorbing.[17]

The simulations focus on two cases in which behavioral differences across types arise. First, if $b + c < w_S^*$, all types share the same reservation wage. Differences in wage outcomes then arise from type-specific search thresholds and acceptance behavior. Second, if $b + c \in (w_S^*, w_{TC}^*)$, sophisticated workers set a higher reservation wage than time-consistent and naive workers.

Figure 3.3 presents the resulting wage distributions, both as absolute frequencies and as the share of each type located in each wage quintile. Across both cases, naive workers are concentrated at the lower end of the wage distribution, while time-consistent workers are concentrated at the upper end. Naive workers believe that they will continue to search in the future as time-consistent workers would, but then search only up to a lower threshold. Hence, naive workers accept certain relatively low-wage offers because they perceive them as temporary steps on the wage ladder, but then do not continue searching and therefore become trapped in these low-wage jobs. Time-consistent workers, by contrast, search up to the highest wage levels and therefore move furthest up the wage ladder.

Sophisticated workers are less concentrated at the bottom of the wage distribution than naive workers, but also less concentrated at the top than time-consistent workers. Relative to time-consistent workers, sophisticated workers search less and are therefore less likely to reach the highest wages. Relative to naive workers, however, they are more selective in their job acceptance behavior. Sophisticated workers anticipate when accepting a job would end future search. They therefore reject some small wage increases that would distort future search incentives and, in the second case, set a higher reservation wage. This shifts mass away from the lowest wage jobs and toward higher wages.

The distributional patterns in Figure 3.3 can thus be summarized as follows: sophisticated workers are most prevalent at intermediate wages (between $\bar{w}$ and w_{TC}^* in the first case, and between r_S^* and w_{TC}^* in the second case). Naive workers are most prevalent at

[17] This assumption is cautious relative to assuming search in the unstable region, because it weakens the distributional patterns discussed below. Footnote 18 explains the role of this assumption for the simulated distributions in more detail.

the lower end of the distribution (below $\bar{w}$ in the first case, and below r_S^* in the second case), while time-consistent workers are most prevalent at the upper end (above w_{TC}^*).[18]

3.4. Evidence

This section examines whether there is empirical support for three central claims of the model. The first concerns the anticipation of future on-the-job search: sophisticated workers should anticipate engaging less in future on-the-job search than naive and time-consistent workers.[19] The second concerns the implied acceptance behavior: sophisticated workers should demand larger wage increases to be willing to change jobs. The third concerns the resulting wage outcomes: differences in search and acceptance behavior should imply that naive workers are more concentrated in lower-wage jobs, sophisticated workers in intermediate-wage jobs, and time-consistent workers in higher-wage jobs.

The analysis uses data from the German Socio-Economic Panel Innovation Sample (SOEP-IS). For a subsample of respondents, the SOEP-IS contains a module that allows individuals to be classified as time-consistent, naive, or sophisticated. For a subset of these individuals, the survey further records beliefs about the likelihood of searching on the job during the next twelve months, the minimum wage increase demanded to be willing to change employers, and current wages.

The analysis is descriptive and does not claim to identify causal effects of being a particular type on the outcomes of interest. Instead, the goal of this exercise is to examine whether the patterns in the data align with the mechanisms and implications of the model.

3.4.1. Data and Identification of Types

The SOEP-IS is a representative panel study of the German population that surveys about 5,500 individuals annually (Richter and Schupp, 2015). In addition to a recurring core questionnaire, the SOEP-IS allows researchers to field own survey modules. The empirical analysis combines two of such modules: one that classifies individuals into their types, and one that elicits beliefs about future on-the-job search and demanded wage increases for job changes.

[18] In the parameterizations used here, $w_N^* > w_S^*$. Hence, there is an interval (w_S^*, w_N^*) that belongs to the unstable region of the sophisticated worker. Since behavior in this region is not sharply characterized, the fact that only sophisticated workers appear in this interval in Panel (a) is a consequence of the simulation convention that sophisticated workers do not search in the unstable region; I thus do not interpret it as a prediction of the model. Allowing sophisticated workers to search in the unstable region would move them over time toward wages above $\bar{w}$. Importantly, the higher prevalence of naive workers between w_N^* and $\bar{w}$ is not driven by this convention but arises despite it; under the alternative convention of search in the unstable region, it would be strengthened.

[19] Present-biased workers should also search less than time-consistent workers. This prediction cannot be tested due to data limitations.

The classification of individuals into their types follows Cobb-Clark et al. (2024), who fielded a module in the 2017 SOEP-IS to categorize individuals as sophisticated, naive, or time-consistent. Their approach uses a setting that, like job search, involves immediate effort costs and delayed benefits: reaching an ideal body weight. In 2017, the module elicited respondents' ideal and expected body weight one year ahead. These measures were combined with realized body weight from the 2018 core questionnaire.

The logic of the classification is as follows. Time-consistent individuals expect to reach their ideal weight and do so. Naive individuals also expect to reach their ideal weight but fail to do so because they overestimate their future costly effort. Sophisticated individuals anticipate future self-control problems and therefore expect not to reach their ideal weight. Cobb-Clark et al. (2024) distinguish between fully and partially sophisticated individuals according to whether the expected future weight coincides with realized future weight or lies between ideal and realized future weight:

$$\text{Ideal future weight} < \text{Exp. future weight} = \text{Actual future weight} \implies \text{Fully sophisticated}$$
$$\text{Ideal future weight} < \text{Exp. future weight} < \text{Actual future weight} \implies \text{Partially sophisticated}$$
$$\text{Ideal future weight} = \text{Exp. future weight} < \text{Actual future weight} \implies \text{Naive}$$
$$\text{Ideal future weight} = \text{Exp. future weight} = \text{Actual future weight} \implies \text{Time-consistent}$$

Applying the classification procedure by Cobb-Clark et al. (2024) classifies 1,155 individuals. In this sample, 24.9% are classified as time-consistent, 33.5% as naive, 15.8% as fully sophisticated, and 25.8% as partially sophisticated.[20]

The empirical analysis pools fully and partially sophisticated individuals. This pooling is natural from the perspective of the model because the relevant distinction is *whether* workers anticipate future present bias. Both types anticipate future present bias, differing only in the degree of anticipation: fully sophisticated workers anticipate it entirely, while partially sophisticated workers anticipate it to some extent. In both cases, anticipated future self-control problems lower expected future on-the-job search below the level considered optimal from the current perspective. Hence, fully and partially sophisticated individuals should exhibit qualitatively similar beliefs about job-search and acceptance behavior, which motivates pooling them in the empirical analysis.[21] Appendix C.5.2 reports robustness checks of the main regressions that distinguish fully and partially sophisticated individuals instead of pooling them.

[20] Appendix C.5.1 presents details on the precise classification procedure.

[21] Formally, the results for fully sophisticated workers extend to partially sophisticated workers by replacing the perceived future present-bias parameter under full sophistication, $\hat{\beta} = \beta$, with $\hat{\beta} \in (\beta, 1)$. Since $\hat{\beta} < 1$, partially sophisticated workers still anticipate suboptimal future search behavior. This preserves a discontinuity in expected continuation payoffs that drives higher selectivity in acceptance behavior, although the discontinuity occurs at a higher wage than under full sophistication.

The outcome variables of interest come from a survey module implemented by Jäger et al. (2024) in the 2019 and 2020 waves of the SOEP-IS. The module was administered to full-time and part-time employed individuals and contains two questions that provide empirical counterparts to the model's central predictions. First, respondents were asked:

> "Over the next 12 months, what is the probability that you will look for a new job at a different company?"

Respondents answered in percent. This question measures beliefs about future on-the-job search. Second, respondents were asked:

> "Imagine that you received a job offer from a different employer in your labor market region that would provide you with a comparable work environment. What wage would this other employer have to offer to ensure that you would leave your current employer?"

Respondents answered in percent gross wage increase. This question measures the wage increase demanded to change jobs.

A total of 1,271 employed individuals answered the search-belief question in 2019 and 1,051 in 2020. Because some respondents appear in both waves, this corresponds to 1,419 distinct individuals. For the demanded wage increase question, 1,006 individuals responded in 2019 and 779 in 2020, corresponding to 1,166 distinct individuals. Merging these samples with the type classification considerably reduces the sample size. For the search-belief question, the baseline sample contains 208 individuals in 2019 and 175 in 2020, with 239 distinct individuals overall. For the demanded wage increase question, the corresponding numbers are 171 individuals in 2019 and 134 in 2020, with 203 distinct individuals overall.

For the demanded wage increase outcome, the sample excludes 15 responses strictly above 50%, which are treated as outliers. Of these responses, 13 report demanded wage increases of at least 100%. Figure C.2a in Appendix C.5.2 shows the distribution of responses and illustrates that these observations are isolated from the rest of the distribution. Moreover, these responses are unlikely to reflect serious wage demands: in 11 cases, the implied demanded wage exceeds the respondent's stated belief about the typical occupational wage by more than 80%, and in 9 cases by more than 100%.

For the search-belief outcome, the full range of stated probabilities is kept. Neither high stated probabilities nor the large mass at zero constitutes a clear outlier, see Figure C.2b in Appendix C.5.2. Moreover, the large mass of responses at zero is a plausible feature of the outcome rather than a reporting problem.

To retain all remaining observations and avoid an arbitrary selection of one wave per individual, the analysis uses weighted least squares regressions in which individuals who responded in both years receive a weight of one half in each year. For the two main regression outcomes, the reported probability of engaging in on-the-job search during the next twelve months and the demanded wage increase required to be willing to change jobs, the outcome is regressed on worker type dummies, omitting time-consistent workers. Standard errors are clustered at the individual level.

The regressions include three sets of controls. The baseline specification controls for the respondent's self-reported gross wage in the previous month. Current wages may be related both to worker type and to the outcomes of interest, since they plausibly affect incentives to search and the wage increase demanded to be willing to change a job.[22]

The second set of controls adds demographic characteristics that may be related both to the type and to the outcomes of interest: years of education or training, gender, and age. Years of education or training are constructed within the SOEP-IS as the typical duration required to obtain the respondent's highest degree, yielding a numerical equivalent for each possible level of education or training.

The third set of controls adds individual traits. Cobb-Clark et al. (2024) show that intelligence and several Big Five personality traits are related to the type. Since such traits may plausibly also be related to intended job search and acceptance behavior, this set of controls includes measures for intelligence and the Big Five personality traits.[23] This set of controls also includes self-reported patience to account for differences in patience that may affect job-search behavior and beliefs independently of present bias and its anticipation.[24]

The analysis of wage outcomes uses the self-reported gross wages in the previous month. In the model, the outcome of interest is the final wage. In the data, such a terminal career wage may only be obtained by restricting attention to workers near the end of their careers, which would leave too few observations for a meaningful comparison across types. Wage distributions are therefore constructed using each individual's mean gross wage across available years; up to nine observations are available per individual. This measure

[22] In the 2019 wave, Jäger et al. (2024) provided half of the respondents with information on the monthly gross wage of a typical full-time employee in their occupation. Since treatment was randomly assigned, it is unrelated to worker type. Moreover, as Jäger et al. (2024) show, the treatment has weak and statistically insignificant effects on the outcomes of interest. This justifies pooling treated and untreated individuals. All results are fully robust to including a treatment dummy. They are also fully robust to including a dummy for the 2019 survey year.

[23] These traits are openness, conscientiousness, extraversion, agreeableness and emotional stability.

[24] In the context of the model, this control may help account for systematic differences in the exponential discount factor δ across types. Such differences could affect search intentions through a channel distinct from present bias and anticipation. At the same time, because it is self-reported, the patience variable may partly capture the degree of sophistication itself. Appendix C.5.2 therefore reports robustness checks that exclude self-reported patience from the set of individual-trait controls.

summarizes the individual's wage position over the panel and avoids an arbitrary choice of a single survey year. However, since this measure clearly differs from the final wage outcome, the corresponding evidence should be interpreted with caution.

To ensure comparability across workers, the sample is restricted to full-time workers. Observations with gross wages below the minimum wage in the respective year are excluded, since such reports may also indicate that the respondent was not in regular full-time employment at the time. These restrictions leave 553 individuals with at least one valid wage observation and type classification.

3.4.2. Results

Beliefs About the Probability to Search On the Job. The first prediction concerns beliefs about future on-the-job search. In the model, sophisticated workers anticipate that their future selves will search less than they currently consider optimal, with the time-consistent search rule providing this optimality benchmark. Naive workers, by contrast, do not anticipate this future deviation and expect to behave as time-consistent workers would. Sophisticated workers should therefore report lower probabilities to engage in future search than both naive and time-consistent workers, while naive and time-consistent workers should report similar beliefs.

Columns (1)–(4) of Table 3.1 report the regression results for the stated probability of searching on the job during the next twelve months. Across all specifications, sophisticated workers report lower probabilities of future on-the-job search than time-consistent workers. The estimated difference ranges from 4.0 to 8.2 percentage points. Sophisticated workers also report lower probabilities than naive workers. Naive workers themselves report lower expected search probabilities than time-consistent workers, but these estimates are smaller in magnitude and close to zero in specifications that include demographic controls.

The estimates therefore point in the direction predicted by the model, but the evidence is statistically weak. The coefficients on sophisticated workers are not statistically different from zero at conventional levels. The strongest support comes from the comparison between sophisticated and naive workers: in the full specification, sophisticated workers report search beliefs that are about 6.9 percentage points lower than those of naive workers, with a p-value of 0.062.

Given the large mass of responses at zero, Table C.3 in Appendix C.5.2 reports regression results using an indicator for reporting any positive probability of future on-the-job search. The results reveal similar patterns, although statistical significance is even more limited.

Taken together, while the estimates are consistent with the model in its sign, the limited statistical significance calls for a cautious interpretation of these results.

Table 3.1: Worker Type, Search Beliefs, and Demanded Wage Increases

	Belief to Search (%)				Demanded Wage Increase (%)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	16.042***	8.410	17.320***	12.856	20.136***	22.563***	21.460***	33.673***
	(5.569)	(10.121)	(6.399)	(14.733)	(2.342)	(5.304)	(2.628)	(7.022)
S	−4.472	−4.007	-8.206	−7.730	4.851**	4.765**	4.488*	4.153
	(4.796)	(4.723)	(5.271)	(5.205)	(2.195)	(2.195)	(2.522)	(2.512)
N	−1.921	−0.266	-2.311	−0.822	2.135	1.977	0.407	−0.502
	(5.042)	(4.964)	(5.599)	(5.646)	(2.158)	(2.119)	(2.445)	(2.430)
Demographic controls		✓		✓		✓		✓
Individual-trait controls			✓	✓			✓	✓
P-value: S = 0	0.352	0.397	0.121	0.140	0.028	0.031	0.077	0.100
P-value: N = 0	0.704	0.957	0.680	0.884	0.324	0.352	0.868	0.837
P-value: S = N	0.400	0.230	0.108	0.062	0.110	0.108	0.054	0.027
Observations	383	382	325	325	305	304	255	255
R^2	0.005	0.025	0.049	0.066	0.031	0.033	0.062	0.075

Note: Columns (1)–(4) report results from weighted least squares regressions of workers' reported probability of engaging in on-the-job search in the next year, measured in percent, on worker type dummies. Columns (5)–(8) report results from weighted least squares regressions of the wage increase demanded to be willing to change jobs, also measured in percent, on worker type dummies. The omitted category is time-consistent workers. The baseline specifications control for the gross wage in the previous month. Demographic and individual-trait controls are included as indicated. Standard errors are clustered at the individual level and reported in parentheses. The reported p-values are based on two-sided t tests. The p-values for $S = N$ are obtained from identical regressions in which naive workers are the omitted category. $^{*}: p < 0.1$; $^{**}: p < 0.05$; $^{***}: p < 0.01$.

Demanded Wage Increases for Job Changes. The second prediction concerns differences in acceptance behavior. In the model, sophisticated workers demand larger wage increases to change jobs than naive and time-consistent workers because they anticipate that job changes may suboptimally prevent future job search.

The estimates in columns (5)–(8) of Table 3.1 report results of the regression on the demanded wage increases to be willing to change jobs. These estimates are in line with the prediction from the model. On average, sophisticated workers demand larger wage increases to switch employers than time-consistent workers in all specifications, with estimated differences ranging from 4.2 to 4.9 percentage points. These differences are statistically significant at the 5% level in the first two specifications and at the 10% level with individual-trait controls; in the full specification, the estimate remains similar in magnitude and lies just outside conventional significance levels (p-value = 0.10). The comparison with naive workers points in the same direction: sophisticated workers demand larger wage increases than naive workers in all specifications. The difference lies just outside conventional significance levels in the first two specifications and becomes statistically significant at the 10% level with individual-trait controls and at the 5% level

in the full specification. Finally, as predicted by the model, naive workers do not differ significantly from time-consistent workers in any specification.

Taken together, the observed patterns support the model's prediction for type differences in job offer acceptance behavior. Table C.1 in Appendix C.5.2 shows that the patterns are robust to distinguishing fully from partially sophisticated workers: demanded wage increases are then ordered by the degree of sophistication. Finally, Table C.2 shows that the estimates remain very similar when self-reported patience is excluded from the set of individual-trait controls. In the full specification, the difference in demanded wage increases between sophisticated and time-consistent workers then also becomes statistically significant at the 10% level.

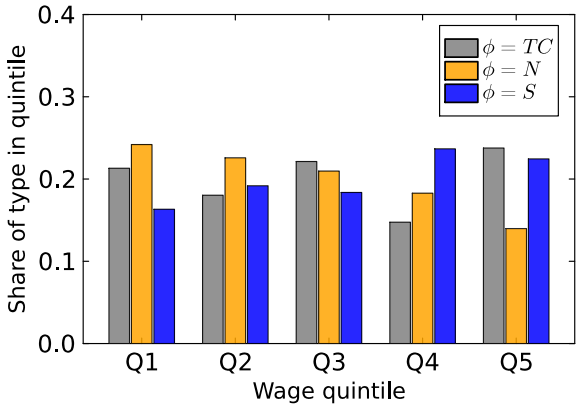
The demanded wage increase provides arguably the most direct evidence in favor of the model. Specifically, this outcome is less vulnerable to alternative explanations than the other outcomes studied. For example, time-consistent workers may have higher wages than present-biased workers because of general productivity differences, arising for instance from differences in procrastination. Such differences could explain wage differences without relying on the search mechanism. By contrast, the finding that sophisticated workers demand larger wage increases is harder to explain through alternative channels. A general productivity based explanation would more naturally predict that time-consistent workers, who may have higher productivity and thus better outside options, demand larger wage increases to change jobs. The estimates show the opposite: sophisticated workers demand larger wage increases. This pattern aligns closely with the acceptance channel in the model, in which sophisticated workers anticipate that accepting a job may end future search and therefore require a larger wage gain.

Wage Distributions. The third prediction concerns wage outcomes. In the model, type-specific differences in search and acceptance behavior translate into different positions in the wage distribution. Naive workers tend to accept relatively low wages under the expectation that they will continue searching in the future, but may subsequently stop searching and remain in relatively low-wage jobs. Sophisticated workers are more selective in job acceptance, which shifts them away from the lowest wages, but they still search less than time-consistent workers, who are more concentrated toward the upper end of the wage distribution.

Figure 3.4 examines this prediction using each individual's mean gross wage across available survey years. Individuals are assigned to quintiles of the overall wage distribution, and the figure reports, for each type, the share of individuals located in each quintile, analogous to the quintile shares in Figure 3.3.

The clearest pattern concerns naive workers. Their share declines monotonically across

Figure 3.4: Empirical Wage Distribution



Note: The figure assigns individuals to quintiles of the overall distribution of individual mean gross wages and reports, for each type, the share of individuals in each quintile. Individual mean gross wages are computed using all available yearly wage observations. The sample is restricted to full-time employed workers with an assigned type and excludes observations with gross wages below the statutory minimum wage in the respective year.

the wage distribution, from 24.2% in the bottom quintile to 14.0% in the top quintile. This is consistent with the model's prediction that naive workers become trapped in lower-wage jobs because they search less than time-consistent workers but accept jobs under overly optimistic beliefs about future on-the-job search.

The patterns for sophisticated and time-consistent workers are less clear. Sophisticated workers are less concentrated in the bottom quintiles than naive workers, consistent with the model. However, they are most represented in the fourth and fifth quintiles, so the data do not match the intermediate-wage prediction. Time-consistent workers have the largest share in the top quintile, consistent with the prediction that they move furthest up the wage distribution. However, their distribution does not show a clear monotonic shift toward higher wages.

Overall, the empirical wage distributions are thus only partially in line with the model and should be interpreted cautiously. Observed wages are a less direct empirical counterpart of the model than search beliefs or demanded wage increases. The empirical measure summarizes wage positions over the observed panel rather than the long-run wage positions predicted by the model. This distinction may be particularly relevant for time-consistent workers, who are predicted to search the most and whose observed wages may therefore be furthest from their eventual long-run positions. Moreover, wages are shaped by many factors outside the model. As discussed above, time-consistent workers may differ from present-biased workers in productivity or traits related to procrastination, which could affect wages independently of the search-based mechanisms identified in this model. This exercise should therefore be understood only as suggestive evidence on the model's

wage prediction.

3.5. Conclusion

This paper studies how present bias shapes job search, job offer acceptance behavior, and wage outcomes in a model with costly on-the-job search. Present-biased workers search less on the job than time-consistent workers, with naive workers especially prone to postponing search because they overestimate future search. Awareness of future present bias affects acceptance behavior: sophisticated workers anticipate that accepting a new job may eliminate future search and therefore demand larger wage increases to switch jobs, while naive workers may accept relatively low offers under the mistaken belief that these jobs are temporary steps on the wage ladder. The same channel also raises the reservation wage of unemployed sophisticated workers. These search and acceptance decisions generate systematic wage differences, with naive workers being overrepresented in lower-wage jobs, sophisticated workers in intermediate-wage jobs, and time-consistent workers in higher-wage jobs. In the infinite-horizon analysis, I introduce a new refinement of perception-perfect strategies that characterizes the stable regions of sophisticated search behavior.

Descriptive evidence from the SOEP-IS is broadly consistent with these predictions. Sophisticated workers report lower probabilities to engage in future on-the-job search than naive and time-consistent workers, although this evidence is statistically weak. The clearest and most direct empirical support concerns job acceptance: sophisticated workers demand larger wage increases to change jobs, while naive and time-consistent workers do not differ significantly. Wage distributions by type provide suggestive additional evidence for naive workers, who are most concentrated at the lower end of the distribution.

There are two natural directions for future work. Theoretically, the model studies unemployment search only through the way present bias distorts future on-the-job search; it abstracts from direct distortions in search effort while unemployed. A richer framework could combine the mechanism identified here with the lower unemployment-search effort emphasized by DellaVigna and Paserman (2005), allowing an assessment of their joint quantitative implications. Empirically, richer data linking time preferences, beliefs about future search, job-to-job transitions, and subsequent wage growth would allow for a more comprehensive test of the mechanisms and their quantitative importance.

APPENDIX A

Appendix to Chapter 1

A.1. Preliminary Lemmas and Results

Notation In the main text, we use firm- and time-specific notation such as $x_{i,t}$, $x_{i,t+d}$, V_t, and V_{t+d} to describe the wage-setting problem of a firm i across wage spells. In the proofs, we simplify notation by suppressing firm and time subscripts. In particular, we write x^*_{SR} for the optimal markup of an individual firm at the beginning of the current wage spell and x^*_{LR} for the optimal markup in the stationary continuation problem from subsequent wage spells onward. Thus, x^*_{SR} corresponds to the main-text object $x^*_{i,t}$, while x^*_{LR} corresponds to the stationary continuation markup represented in the main text by $x^*_{i,t+d}$, which is chosen again in periods $t + 2d$, $t + 3d$, and so on. We also write x_{SR} for a generic markup choice in the current wage spell when no confusion can arise. Likewise, V_{SR} and V_{LR} denote the associated short-run and long-run value functions. This change is purely notational.

Roadmap of the Mathematical Appendix

This section gives a short overview of how the model's results are derived. The appendix consists of two major blocks. First, Lemma A.1–Lemma A.3 and Corollary A.1–Corollary A.2 establish existence and properties of the long-run and short-run value functions and their derivatives. They also establish existence of the stationary long-run optimal wage-setting solution. Second, the proofs of Proposition 1.1, Proposition 1.2, Proposition 1.3, and Proposition 1.4 use these objects to derive the model predictions.

Step 1: Establishing existence and properties of the value functions and the stationary long-run optimal wage-setting solution. The first step establishes existence and properties of the value functions and derivative objects used throughout the rest of the appendix. It also establishes existence of the optimal wage-setting solution in the stationary long-run problem. For the long-run problem, the proofs proceed from the second derivative to the first derivative and finally to the value function itself. This order

is useful because the second derivative satisfies the simplest recursive equation and can be characterized by a standard contraction argument. Integrating this second derivative yields the first derivative. The first derivative is the key object for wage setting, because the optimal markup is characterized by the first-order condition. Further integration then yields the value function itself.

More specifically, Lemma A.1–Lemma A.3 solve a generic stationary version of the recursive problem. Lemma A.1 starts from a candidate stationary markup k and establishes existence, uniqueness, boundedness, and continuity of the corresponding second derivative V_k''. It also shows that V_k'' is bounded below by $cd > 0$, which gives the convexity needed for a well-defined optimum. Lemma A.2 then integrates this second-derivative characterization and shows that there exists a pair (k^*, V') such that V' satisfies the relevant recursive equation and $V'(k^*) = 0$. Lemma A.3 finally integrates once more and constructs the associated value function itself.

Two corollaries translate these technical objects into economic objects. Corollary A.1 applies the preceding lemmas to the stationary long-run problem and defines the long-run value function V_{LR}, its first and second derivatives, V'_{LR} and V''_{LR}, and the long-run optimal markup x^*_{LR}. Corollary A.2 characterizes the short-run value function V_{SR} and derives the expressions for V'_{SR} and V''_{SR}, taking the long-run characterizations from Corollary A.1 as the continuation objects.

Step 2: Proving the model predictions. The second step uses these objects to prove the model predictions. Proposition 1.1 first establishes the benchmark without DNWR. Proposition 1.2 then shows how DNWR reduces the pass-through relative to this benchmark. Proposition 1.3 shows that higher short-run inflation uncertainty lowers the pass-through. Proposition 1.4 shows the opposite effect for long-run inflation expectations.

Lemma A.1 *Let $k \in \mathbb{R}$, $c, d \in \mathbb{R}_{>0}$ and $\delta \in (0,1)$. Let $\pi \sim \mathcal{N}(\pi^e, \sigma^2)$ with cumulative distribution function $F(z) = \mathbb{P}(\pi \leq z)$. There exists a unique bounded and continuous function $V_k'' : \mathbb{R} \to \mathbb{R}$ such that, for all $x \in \mathbb{R}$,*

$$V_k''(x) = cd + \delta F(x-k)\mathbb{E}\left[V_k''(x-\pi) \mid \pi < x-k\right] \qquad \text{and} \qquad V_k''(x) \in \left[cd, \frac{cd}{1-\delta}\right].$$

Proof. Fix an arbitrary $k \in \mathbb{R}$. The functional equation in Lemma A.1 is equivalent to

$$V_k''(x) = cd + \delta \int_{-\infty}^{x-k} V_k''(x-\pi) f(\pi) d\pi,$$

where $f(\pi)$ denotes the probability density function of the normal distribution $\mathcal{N}(\pi^e, \sigma^2)$.

Consider the space of all continuous and bounded functions on $\mathbb{R}$, denoted by $\mathcal{C}_b(\mathbb{R})$. Equipped with the supremum norm $\|g\|_\infty := \sup_{x\in\mathbb{R}} |g(x)|$, this space is complete. Let $g \in \mathcal{C}_b(\mathbb{R})$ and define the operator T by

$$(Tg)(x) = cd + \delta \int_{-\infty}^{x-k} g(x-\pi) f(\pi) d\pi.$$

If (i) T maps $\mathcal{C}_b(\mathbb{R})$ into itself, that is, $T : \mathcal{C}_b(\mathbb{R}) \to \mathcal{C}_b(\mathbb{R})$, and if also (ii) T is a contraction mapping, then by the Banach fixed-point theorem, there exists a unique continuous and bounded function $V_k'' \in \mathcal{C}_b(\mathbb{R})$ that solves the functional equation in Lemma A.1. We first show property (i). Fix $g \in \mathcal{C}_b(\mathbb{R})$ and define

$$h(x) \coloneqq (Tg)(x) = cd + \delta \int_{-\infty}^{x-k} g(x-\pi) f(\pi) d\pi = cd + \delta \int_{-\infty}^{\infty} g(x-\pi) \mathbf{1}\{\pi < x-k\} f(\pi) d\pi.$$

We first show continuity of $h(x)$. Let $x_n \to x$. We need to show that $h(x_n) \to h(x)$. For each $\pi \neq x-k$, continuity of g implies $g(x_n - \pi) \to g(x-\pi)$, and $\mathbf{1}\{\pi < x_n - k\} \to \mathbf{1}\{\pi < x-k\}$. Since $\pi = x-k$ has measure zero, the integrand therefore converges pointwise almost everywhere.

Moreover, the sequence $\left|g(x_n-\pi)\,\mathbf{1}\{\pi < x_n - k\}\, f(\pi)\right|$ is dominated by an integrable function since

$$\left|g(x_n-\pi)\,\mathbf{1}\{\pi < x_n - k\}\, f(\pi)\right| \leq \|g\|_\infty f(\pi),$$

and

$$\int_{-\infty}^{\infty} \|g\|_\infty f(\pi)\, d\pi = \|g\|_\infty < \infty.$$

The dominated convergence theorem therefore implies

$$\lim_{n\to\infty} h(x_n) = cd + \delta \lim_{n\to\infty} \int_{-\infty}^{\infty} g(x_n - \pi) \mathbf{1}\{\pi < x_n - k\} f(\pi) d\pi$$

$$
\begin{aligned}
&= cd + \delta \int_{-\infty}^{\infty} \lim_{n\to\infty} g(x_n - \pi)\mathbf{1}\{\pi < x_n - k\} f(\pi) d\pi \\
&= cd + \delta \int_{-\infty}^{\infty} g(x - \pi)\mathbf{1}\{\pi < x - k\} f(\pi) d\pi \\
&= h(x).
\end{aligned}
$$

Hence, $h(x_n) \to h(x)$, so h is continuous. Moreover, h is bounded:

$$
|h(x)| \leq cd + \delta \int_{-\infty}^{x-k} |g(x-\pi)|\, f(\pi)\, d\pi \leq cd + \delta \|g\|_\infty \int_{-\infty}^{x-k} f(\pi)\, d\pi \leq cd + \delta \|g\|_\infty.
$$

Hence, $h \in \mathcal{C}_b(\mathbb{R})$ so that T maps $\mathcal{C}_b(\mathbb{R})$ into itself.

We next show property (ii), i.e., T is a contraction mapping. We verify Blackwell's sufficient conditions. For monotonicity, consider $g_1, g_2 \in \mathcal{C}_b(\mathbb{R})$ with $g_1(x) \leq g_2(x)$ for all x. Then for all $x \in \mathbb{R}$,

$$
(Tg_1)(x) = cd + \delta \int_{-\infty}^{x-k} g_1(x-\pi)\, f(\pi)\, d\pi \leq cd + \delta \int_{-\infty}^{x-k} g_2(x-\pi)\, f(\pi)\, d\pi = (Tg_2)(x).
$$

Hence, T is monotone. For discounting, let $g \in \mathcal{C}_b(\mathbb{R})$ and let $r \geq 0$. For all $x \in \mathbb{R}$,

$$
T(g+r)(x) = cd + \delta \int_{-\infty}^{x-k} (g(x-\pi) + r)\, f(\pi)\, d\pi = Tg(x) + \delta F(x-k)\, r \leq Tg(x) + \delta r.
$$

Since $0 < \delta < 1$, T satisfies Blackwell's discounting condition with discount factor δ. By Blackwell's theorem, T is therefore a contraction on $(\mathcal{C}_b(\mathbb{R}), \|\cdot\|_\infty)$.

Because T maps the complete space $(\mathcal{C}_b(\mathbb{R}), \|\cdot\|_\infty)$ into itself and it is a contraction mapping, T admits a unique fixed-point $V_k'' \in \mathcal{C}_b(\mathbb{R})$ by Banach fixed-point theorem.

It remains to show that $V_k''(x) \in [cd, cd/(1-\delta)]$ for every $x \in \mathbb{R}$. Since $V_k'' \in \mathcal{C}_b(\mathbb{R})$, we can define

$$
M := \sup_{x\in\mathbb{R}} V_k''(x), \qquad I := \inf_{x\in\mathbb{R}} V_k''(x).
$$

First, suppose toward a contradiction that $I < 0$. Then, for every $x \in \mathbb{R}$,

$$
\begin{aligned}
V_k''(x) &= cd + \delta \int_{-\infty}^{x-k} V_k''(x-\pi) f(\pi) d\pi \\
&\geq cd + \delta I F(x-k) \geq cd + \delta I,
\end{aligned}
$$

where the last inequality uses $I < 0$ and $F(x-k) \leq 1$. Hence,

$$
I \geq cd + \delta I,
$$

or equivalently $I \geq cd/(1-\delta) > 0$, a contradiction. Hence $I \geq 0$. Therefore,

$$
V_k''(x) = cd + \delta \int_{-\infty}^{x-k} V_k''(x-\pi) f(\pi) d\pi \geq cd,
$$

so $I \geq cd$. For the upper bound, using $V_k''(x) \leq M$ for all x, we obtain

$$\begin{aligned} V_k''(x) &= cd + \delta \int_{-\infty}^{x-k} V_k''(x-\pi) f(\pi) d\pi \\ &\leq cd + \delta M F(x-k) \leq cd + \delta M, \end{aligned}$$

where the last inequality uses $M \geq cd > 0$ and $F(x-k) \leq 1$. Hence,

$$M \leq cd + \delta M,$$

so that

$$M \leq \frac{cd}{1-\delta}.$$

Combining the lower and upper bounds yields

$$V_k''(x) \in \left[cd, \frac{cd}{1-\delta}\right] \qquad \text{for all } x \in \mathbb{R}.$$

□

Lemma A.2 *Let $c, d \in \mathbb{R}_{>0}$ and $\delta \in (0,1)$. There exists a pair (k^*, V'), where $k^* \in \mathbb{R}$ and $V' : \mathbb{R} \to \mathbb{R}$ is differentiable with a bounded and continuous derivative, such that $V'(k^*) = 0$ and, for all $x \in \mathbb{R}$,*

$$V'(x) = cdx - \frac{cd}{2}\pi^e + \delta \int_{-\infty}^{x-k^*} V'(x-\pi)\, f(\pi)\, d\pi,$$

where f denotes the density of $\pi \sim \mathcal{N}(\pi^e, \sigma^2)$. Moreover, the pair (k^*, V') is unique among all pairs (k, v) such that $v : \mathbb{R} \to \mathbb{R}$ is differentiable with a bounded and continuous derivative, $v(k) = 0$, and the functional equation holds for all $x \in \mathbb{R}$.

PROOF. Consider any arbitrary $k \in \mathbb{R}$ and define $V_k''(x)$ as in Lemma A.1. Since $V_k''(x)$ is continuous, it admits anti-derivatives. Define $\tilde{V}_k'$ as the unique anti-derivative of V_k'' such that $\tilde{V}_k'(k) = 0$.

Define an additional function $\hat{V}_k'(x)$ as

$$\hat{V}_k'(x) \coloneqq cdx - \frac{cd}{2}\pi^e + \delta \int_{-\infty}^{x-k} \tilde{V}_k'(x-\pi)\, f(\pi)\, d\pi.$$

We first verify that the improper integral in the definition of $\hat{V}_k'$ exists. By Lemma A.1, V_k'' is bounded, so there exists $M < \infty$ such that $\|V_k''\|_\infty \leq M$. Since $\tilde{V}_k'(k) = 0$ and $\tilde{V}_k'$ is an anti-derivative of V_k'', it follows that

$$|\tilde{V}_k'(y)| \leq M|y-k| \qquad \text{for all } y \in \mathbb{R}.$$

Hence, for any $x \in \mathbb{R}$,

$$\int_{-\infty}^{x-k} \left|\tilde{V}_k'(x-\pi)\right| f(\pi)\, d\pi \leq M \int_{-\infty}^{x-k} |x-\pi-k| f(\pi)\, d\pi \\ \leq M\big(|x-k| + \mathbb{E}[|\pi|]\big) < \infty,$$

where $\mathbb{E}[|\pi|] < \infty$ by normality. So, the improper integral $\int_{-\infty}^{x-k} \tilde{V}_k'(x-\pi) f(\pi)\, d\pi$ exists for all x.

We next derive the derivative of $\hat{V}_k'(x)$. To justify differentiation under the integral sign at $x \in \mathbb{R}$, choose some $a < x-k$. Then rewrite the integral term as

$$\int_{-\infty}^{x-k} \tilde{V}_k'(x-\pi) f(\pi)\, d\pi = \int_{-\infty}^{a} \tilde{V}_k'(x-\pi) f(\pi)\, d\pi + \int_{a}^{x-k} \tilde{V}_k'(x-\pi) f(\pi)\, d\pi.$$

For the second term on the right-hand side, the Leibniz rule applies directly since $\tilde{V}_k'$ is differentiable and V_k'' is continuous. For the first term, differentiation under the integral sign is justified by dominated convergence: for $h \neq 0$, the mean value theorem gives, for some $\theta \in (0,1)$,

$$\left| \frac{\tilde{V}_k'(x+h-\pi) - \tilde{V}_k'(x-\pi)}{h} \right| f(\pi) = |V_k''(x+\theta h-\pi)| f(\pi) \leq M f(\pi),$$

with $Mf(\pi)$ integrable on $(-\infty, a]$. Moreover, the difference quotient converges pointwise to $V_k''(x-\pi)f(\pi)$ as $h \to 0$. Hence dominated convergence justifies differentiation under the integral sign for the first term. Thus,

$$\frac{\partial}{\partial x} \int_{-\infty}^{x-k} \tilde{V}_k'(x-\pi) f(\pi)\, d\pi = \int_{-\infty}^{x-k} V_k''(x-\pi) f(\pi)\, d\pi + \tilde{V}_k'(k) f(x-k) = \int_{-\infty}^{x-k} V_k''(x-\pi) f(\pi)\, d\pi,$$

where the second equality follows from $\tilde{V}_k'(k) = 0$. Hence,

$$\frac{\partial}{\partial x} \hat{V}_k'(x) = cd + \delta \int_{-\infty}^{x-k} V_k''(x-\pi)\, f(\pi)\, d\pi = V_k''(x).$$

Thus, $\hat{V}_k'$ is also an anti-derivative of V_k''. Hence, since $\hat{V}_k'$ and $\tilde{V}_k'$ are anti-derivatives of the same function, their difference is constant in x, though it may depend on k. Thus, $\tilde{V}_k'(x) = \hat{V}_k'(x) - r(k)$ and

$$\tilde{V}_k'(x) = cdx - \frac{cd}{2}\pi^e + \delta \int_{-\infty}^{x-k} \tilde{V}_k'(x-\pi) f(\pi) d\pi - r(k). \tag{A.1}$$

For each $k \in \mathbb{R}$, define

$$W_k(y) := V_k''(k+y) \qquad \text{for all } y \in \mathbb{R}.$$

Substituting $x = k + y$ into the functional equation from Lemma A.1 yields

$$\begin{aligned} W_k(y) &= cd + \delta \int_{-\infty}^{y} V_k''(k + y - \pi) f(\pi)\, d\pi \\ &= cd + \delta \int_{-\infty}^{y} W_k(y - \pi) f(\pi)\, d\pi. \end{aligned}$$

By the same contraction mapping argument as in Lemma A.1, this functional equation admits a unique bounded and continuous solution. Hence all functions W_k coincide. Denoting the common function by W, we obtain

$$V_k''(x) = W(x - k) \qquad \text{for all } x, k \in \mathbb{R}.$$

Let A be the unique anti-derivative of W such that $A(0) = 0$. Since W is independent of k, its normalized anti-derivative A is also independent of k. We claim that

$$\tilde{V}_k'(x) = A(x - k) \qquad \text{for all } x, k \in \mathbb{R}.$$

Indeed, the function $x \mapsto A(x - k)$ is an anti-derivative of V_k'' because

$$\frac{d}{dx} A(x - k) = W(x - k) = V_k''(x).$$

Also, it satisfies

$$A(k - k) = A(0) = 0.$$

Since $\tilde{V}_k'$ is the unique anti-derivative of V_k'' with $\tilde{V}_k'(k) = 0$, it follows that

$$\tilde{V}_k'(x) = A(x - k) \qquad \text{for all } x, k \in \mathbb{R}.$$

Substituting this identity into (A.1), we obtain

$$A(x - k) = cdx - \frac{cd}{2}\pi^e + \delta \int_{-\infty}^{x-k} A(x - k - \pi) f(\pi)\, d\pi - r(k) \qquad \text{for all } x, k \in \mathbb{R}.$$

Rewriting in terms of $y \coloneqq x - k$,

$$A(y) = cd(k + y) - \frac{cd}{2}\pi^e + \delta \int_{-\infty}^{y} A(y - \pi) f(\pi)\, d\pi - r(k) \qquad \text{for all } y, k \in \mathbb{R}.$$

Rearranging gives

$$cdk - r(k) = A(y) + \frac{cd}{2}\pi^e - cdy - \delta \int_{-\infty}^{y} A(y - \pi) f(\pi)\, d\pi \qquad \text{for all } y, k \in \mathbb{R}.$$

Since this equality holds for all $y, k \in \mathbb{R}$, the left-hand side, which depends only on k, must equal the right-hand side, which depends only on y, at a common constant. Hence there exists $C \in \mathbb{R}$ such that

$$cdk - r(k) = C \qquad \text{for all } k \in \mathbb{R},$$

or equivalently,

$$r(k) = cdk - C.$$

Thus, r is continuous and strictly increasing in k. It follows that there exists a unique $k^* \in \mathbb{R}$ such that $r(k^*) = 0$.

Define

$$V'(x) := \tilde{V}'_{k^*}(x).$$

Because $\tilde{V}'_{k^*}$ is an anti-derivative of V''_{k^*} as defined in Lemma A.1, the function V' is differentiable on $\mathbb{R}$ with derivative V''_{k^*}, which is bounded and continuous. Moreover, by construction,

$$V'(k^*) = \tilde{V}'_{k^*}(k^*) = 0.$$

Finally, (A.1) with $k = k^*$ and $r(k^*) = 0$ yields, for all $x \in \mathbb{R}$,

$$V'(x) = cdx - \frac{cd}{2}\pi^e + \delta \int_{-\infty}^{x-k^*} V'(x-\pi) f(\pi)\, d\pi.$$

Hence, (k^*, V') satisfies all required properties.

To prove uniqueness, let $(\bar{k}, \bar{V}')$ be any solution such that $\bar{V}' : \mathbb{R} \to \mathbb{R}$ is differentiable with bounded and continuous derivative, $\bar{V}'(\bar{k}) = 0$, and the functional equation holds for all $x \in \mathbb{R}$. Differentiating the functional equation shows that $\bar{V}''$ satisfies the equation in Lemma A.1 with $k = \bar{k}$. By uniqueness in Lemma A.1, we obtain $\bar{V}'' = V''_{\bar{k}}$. Since both $\bar{V}'$ and $\tilde{V}'_{\bar{k}}$ are anti-derivatives of $V''_{\bar{k}}$ and satisfy $\bar{V}'(\bar{k}) = \tilde{V}'_{\bar{k}}(\bar{k}) = 0$, it follows that $\bar{V}' = \tilde{V}'_{\bar{k}}$. Hence (A.1) holds with $k = \bar{k}$ and $r(\bar{k}) = 0$. Since $r(k) = cdk - C$ is strictly increasing, there exists a unique $k^* \in \mathbb{R}$ such that $r(k^*) = 0$. Therefore $\bar{k} = k^*$ and thus $\bar{V}' = \tilde{V}'_{k^*} = V'$. □

Lemma A.3 *Let $c, d \in \mathbb{R}_{>0}$ and $\delta \in (0,1)$, and let k^* and V' be as defined in Lemma A.2. There exists a twice differentiable function $V : \mathbb{R} \to \mathbb{R}$ such that*

$$V(x) = \mathbb{E}\big[L(x,\pi)\big] + \delta F(x-k^*)\mathbb{E}\big[V(x-\pi) \mid \pi < x - k^*\big] + \delta\big(1 - F(x-k^*)\big)V(k^*)$$

for all $x \in \mathbb{R}$, where F denotes the cumulative distribution function of $\pi \sim \mathcal{N}(\pi^e, \sigma^2)$. Moreover, the first derivative of V is given by V' with $V'(k^*) = 0$, and the second derivative is $V'' = V''_{k^*} \in \big[cd, \frac{cd}{1-\delta}\big]$ as defined in Lemma A.1.

Proof. For any $b \in \mathbb{R}$, let $V^b(x)$ denote an anti-derivative of $V'(x)$ as defined in Lemma A.2 satisfying $V^b(k^*) = b$. Define an additional function $\tilde{V}^b$:

$$\tilde{V}^b(x) := \mathbb{E}\big[L(x,\pi)\big] + \delta F(x-k^*)\mathbb{E}\big[V^b(x-\pi) \mid \pi < x - k^*\big] + \delta\big(1 - F(x-k^*)\big)V^b(k^*)$$

$$= \mathbb{E}[L(x,\pi)] + \delta \int_{-\infty}^{x-k^*} V^b(x-\pi) f(\pi)\, d\pi + \delta\big(1 - F(x-k^*)\big) V^b(k^*). \tag{A.2}$$

Before proceeding, we verify that the improper integral term in (A.2) exists. By Lemma A.1, V'' is bounded, so there exists $M > 0$ such that

$$|V''(x)| \le M \qquad \text{for all } x \in \mathbb{R}.$$

Since $V'(k^*) = 0$, it follows that for any $y \in \mathbb{R}$,

$$|V'(y)| \le M|y - k^*|.$$

Thus V' has at most linear growth. Since V^b is an anti-derivative of V' with $V^b(k^*) = b$, it follows that V^b has at most quadratic growth. Hence, there exists a constant $C < \infty$ such that

$$|V^b(y)| \le |b| + C(y - k^*)^2 \qquad \text{for all } y \in \mathbb{R}.$$

Since the improper integral

$$\int_{-\infty}^{x-k^*} \big(|b| + C(x - \pi - k^*)^2\big) f(\pi)\, d\pi$$

converges, so does

$$\int_{-\infty}^{x-k^*} V^b(x-\pi)\, f(\pi)\, d\pi < \infty.$$

We next justify differentiation under the integral sign of the improper integral for any $x \in \mathbb{R}$. Choose $a < x - k^*$. Then

$$\int_{-\infty}^{x-k^*} V^b(x-\pi) f(\pi)\, d\pi = \int_{-\infty}^{a} V^b(x-\pi) f(\pi)\, d\pi + \int_{a}^{x-k^*} V^b(x-\pi) f(\pi)\, d\pi.$$

For the second term on the right-hand side, the Leibniz rule applies directly since V^b is differentiable with derivative V', and V' is continuous. For the first term, differentiation under the integral sign is justified by dominated convergence: for $h \neq 0$, the mean value theorem gives, for some $\theta \in (0,1)$,

$$\left|\frac{V^b(x+h-\pi) - V^b(x-\pi)}{h}\right| f(\pi) = |V'(x+\theta h - \pi)| f(\pi) \le M|x + \theta h - \pi - k^*| f(\pi).$$

Restricting to $|h| \le 1$, we obtain

$$|x + \theta h - \pi - k^*| \le |x - \pi - k^*| + 1,$$

and hence

$$\left|\frac{V^b(x+h-\pi) - V^b(x-\pi)}{h}\right| f(\pi) \le M\big(|x - \pi - k^*| + 1\big) f(\pi),$$

with this upper bound being integrable on $(-\infty, a]$. Moreover, the difference quotient converges pointwise to $V'(x-\pi)f(\pi)$ as $h \to 0$. Hence dominated convergence justifies differentiation under the integral sign for the first term. Combining yields

$$\frac{\partial}{\partial x}\delta \int_{-\infty}^{x-k^*} V^b(x-\pi)f(\pi)\,d\pi = \delta \int_{-\infty}^{x-k^*} V'(x-\pi)f(\pi)\,d\pi + \delta V^b(k^*)f(x-k^*).$$

Differentiating the last term of (A.2) gives

$$\frac{\partial}{\partial x}\Big[\delta\big(1-F(x-k^*)\big)V^b(k^*)\Big] = -\delta V^b(k^*)f(x-k^*).$$

The boundary terms therefore cancel, so

$$\frac{\partial}{\partial x}\tilde{V}^b(x) = cdx - \frac{cd}{2}\pi^e + \delta \int_{-\infty}^{x-k^*} V'(x-\pi)f(\pi)\,d\pi = V'(x),$$

where the last equality follows from Lemma A.2. Hence, $\tilde{V}^b(x)$ also constitutes an antiderivative of $V'(x)$. It follows that $\tilde{V}^b(x)$ and $V^b(x)$ can differ from each other only up to a constant that is independent of x:

$$\tilde{V}^b(x) = V^b(x) + \Psi.$$

We next show that there exists a unique constant b^* such that $\tilde{V}^{b^*}(x) = V^{b^*}(x)$ for all x.

For any $m \in \mathbb{R}$, consider $V^{b+m}(x) = V^b(x) + m$. Then

$$\begin{aligned}\tilde{V}^{b+m}(x) &= \mathbb{E}[L(x,\pi)] + \delta F(x-k^*)\,\mathbb{E}\big[V^b(x-\pi)+m \mid \pi < x-k^*\big] + \delta\big(1-F(x-k^*)\big)(V^b(k^*)+m)\\ &= \tilde{V}^b(x) + \delta m.\end{aligned}$$

Using $\tilde{V}^b(x) = V^b(x) + \Psi$, we therefore obtain

$$\tilde{V}^{b+m}(x) - V^{b+m}(x) = (V^b(x) + \Psi + \delta m) - (V^b(x) + m) = \Psi - m(1-\delta).$$

Setting $m^* := \Psi/(1-\delta)$ yields $\tilde{V}^{b+m^*}(x) = V^{b+m^*}(x)$ for all x. Define $b^* := b + m^*$ and $V(x) := V^{b^*}(x)$. Then V is a twice differentiable solution of the functional equation from Lemma A.3. Its first derivative is V' with $V'(k^*) = 0$ as given by Lemma A.2, and its second derivative is $V'' := V''_{k^*}$ as given by Lemma A.1. □

Corollary A.1 *Let $\pi \sim \mathcal{N}(\pi^e_{LR}, \sigma^2_{LR})$, and let f_{LR} and F_{LR} denote its density and cumulative distribution function. Apply Lemma A.2 to this distribution, let (k^*, V'_{LR}) denote the resulting unique pair, and set $x^*_{LR} := k^*$. Let V_{LR} be the function constructed from V'_{LR} as in Lemma A.3. Then $V_{LR} : \mathbb{R} \to \mathbb{R}$ is twice differentiable and, for all $x \in \mathbb{R}$,*

$$V_{LR}(x) = \mathbb{E}_{LR}\big[L(x,\pi)\big] + \delta F_{LR}(x - x^*_{LR})\,\mathbb{E}_{LR}[V_{LR}(x-\pi) \mid \pi < x - x^*_{LR}]$$

$$+ \delta\big(1 - F_{LR}(x - x^*_{LR})\big)V_{LR}(x^*_{LR}).$$

Moreover, its first and second derivatives satisfy, for all $x \in \mathbb{R}$,

$$V'_{LR}(x) = cdx - \frac{cd}{2}\pi^e_{LR} + \delta \int_{-\infty}^{x - x^*_{LR}} V'_{LR}(x - \pi)\, f_{LR}(\pi)\, d\pi,$$
$$V''_{LR}(x) = cd + \delta \int_{-\infty}^{x - x^*_{LR}} V''_{LR}(x - \pi)\, f_{LR}(\pi)\, d\pi, \qquad V''_{LR}(x) \in \left[cd, \frac{cd}{1-\delta}\right].$$

Finally, x^*_{LR} *is the optimal wage markup in any period, given the anticipation that the same markup will be chosen in all future wage setting decisions under long-run inflation expectations* $\mathcal{N}(\pi^e_{LR}, \sigma^2_{LR})$.

Proof. Applying Lemma A.1–Lemma A.3 for the distribution $\mathcal{N}(\pi^e_{LR}, \sigma^2_{LR})$ yields the function V_{LR} and the unique pair (x^*_{LR}, V'_{LR}) as defined in the corollary. In particular, V_{LR} solves the long-run functional equation, and its first and second derivatives have the stated representations and bounds.

Moreover, $V'_{LR}(x^*_{LR}) = 0$. Since $V''_{LR}(x) \geq cd > 0$ for all $x \in \mathbb{R}$, the function V_{LR} is strictly convex. Hence x^*_{LR} is the optimal wage markup in any period, given the anticipation that the same markup will be chosen in all future wage setting decisions under long-run inflation expectations $\mathcal{N}(\pi^e_{LR}, \sigma^2_{LR})$. □

Corollary A.2 *Let* $\pi \sim \mathcal{N}(\pi^e_{SR}, \sigma^2_{SR})$*, and let* f_{SR} *and* F_{SR} *denote its density and cumulative distribution function. Let* x^*_{LR} *and* V_{LR} *be as defined in Corollary A.1. There exists a twice differentiable function* $V_{SR} : \mathbb{R} \to \mathbb{R}$ *such that, for all* $x \in \mathbb{R}$,

$$V_{SR}(x) = \mathbb{E}_{SR}\big[L(x, \pi)\big] + \delta F_{SR}(x - x^*_{LR})\, \mathbb{E}_{SR}[V_{LR}(x - \pi) \mid \pi < x - x^*_{LR}]$$
$$+ \delta\big(1 - F_{SR}(x - x^*_{LR})\big)V_{LR}(x^*_{LR}).$$

Its first and second derivatives are given by

$$V'_{SR}(x) = cdx - \frac{cd}{2}\pi^e_{SR} + \delta \int_{-\infty}^{x - x^*_{LR}} V'_{LR}(x - \pi)\, f_{SR}(\pi)\, d\pi,$$
$$V''_{SR}(x) = cd + \delta \int_{-\infty}^{x - x^*_{LR}} V''_{LR}(x - \pi)\, f_{SR}(\pi)\, d\pi, \qquad V''_{SR}(x) \in \left[cd, \frac{cd}{1-\delta}\right].$$

Proof. The expression for V_{SR} is well defined by Corollary A.1. Differentiation under the integral sign follows by the same Leibniz and dominated convergence arguments as in Lemma A.2 and Lemma A.3, yielding the expressions for V'_{SR} and V''_{SR}. The bounds on V''_{SR} follow directly from $V''_{LR}(x) \in [cd, cd/(1-\delta)]$ for all x. □

A.2. Proofs

A.2.1. Proof of Proposition 1.1.

Since $\Pi'(0) = 0$ and $\Pi''(x) = -c < 0$ for all $x \in \mathbb{R}$, the profit function must be of the form

$$\Pi(x) = -\frac{c}{2}x^2 + a \quad \text{with} \quad a \in \mathbb{R}.$$

Hence,

$$\begin{aligned}
\mathbb{E}_{SR}\left[L(x_{SR}, \pi_{SR})\right] &= \mathbb{E}_{SR}\left[\int_0^d \Pi(0) - \Pi\left(x_{SR} - s\frac{\pi_{SR}}{d}\right)\, ds\right] \\
&= \mathbb{E}_{SR}\left[\int_0^d \frac{c}{2}\left(x_{SR} - s\frac{\pi_{SR}}{d}\right)^2\, ds\right] \\
&= \frac{c}{2}\left(x_{SR}^2 d + \mathbb{E}_{SR}\left[\frac{1}{3}d\pi_{SR}^2\right] - x_{SR}d\pi_{SR}^e\right).
\end{aligned}$$

Taking the first order condition with respect to x_{SR} yields

$$2x_{SR}^* d = d\pi_{SR}^e \iff x_{SR}^* = \frac{\pi_{SR}^e}{2}.$$

This solution constitutes the unique minimum since

$$\frac{\partial^2}{\partial x_{SR}^2}\mathbb{E}_{SR}\left[L(x_{SR}, \pi_{SR})\right] = cd > 0.$$

Aggregate wage growth and the corresponding pass-through ϕ^* follow from

$$x_t^* = \frac{x_{SR}^*}{d} = \frac{\pi_{SR}^e}{2d} = \frac{d\pi_{t,t+1}^e}{2d} = \frac{1}{2}\pi_{t,t+1}^e.$$

□

A.2.2. Proof of Proposition 1.2.

By Corollary A.2, the derivative of the short-run value function is given by

$$V_{SR}'(x) = cdx - \frac{cd}{2}\pi_{SR}^e + \delta\int_{-\infty}^{x - x_{LR}^*} V_{LR}'(x-\pi)\, f_{SR}(\pi)\, d\pi \qquad \text{for all } x \in \mathbb{R}.$$

Furthermore, again by Corollary A.2,

$$V_{SR}''(x) \in \left[cd, \frac{cd}{1-\delta}\right] \qquad \text{for all } x \in \mathbb{R}.$$

Therefore, V_{SR} is strictly convex. Moreover, since $V_{SR}''(x) \geq cd > 0$ for all $x \in \mathbb{R}$, it follows that $V_{SR}'(x) < 0$ for sufficiently small x and $V_{SR}'(x) > 0$ for sufficiently large x. Hence, there exists a unique minimizer $x_{SR}^* \in \mathbb{R}$ such that $V_{SR}'(x_{SR}^*) = 0$.

We next show that the integral term in the expression for $V'_{SR}(x)$ is strictly positive for every $x \in \mathbb{R}$. Fix any $x \in \mathbb{R}$. For every $\pi < x - x^*_{LR}$, we have $x - \pi > x^*_{LR}$. By Corollary A.1,

$$V'_{LR}(x^*_{LR}) = 0 \qquad \text{and} \qquad V''_{LR}(y) \in \left[cd, \frac{cd}{1-\delta}\right] \qquad \text{for all } y \in \mathbb{R}.$$

Hence V'_{LR} is strictly increasing on $\mathbb{R}$. Since $V'_{LR}(x^*_{LR}) = 0$, it follows that

$$V'_{LR}(y) > 0 \qquad \text{for all } y > x^*_{LR}.$$

Therefore,

$$V'_{LR}(x - \pi) > 0 \qquad \text{for all } \pi < x - x^*_{LR}.$$

Because the normal density satisfies $f_{SR}(\pi) > 0$ for all $\pi \in \mathbb{R}$, we obtain

$$\delta \int_{-\infty}^{x - x^*_{LR}} V'_{LR}(x - \pi)\, f_{SR}(\pi)\, d\pi > 0 \qquad \text{for all } x \in \mathbb{R}.$$

Now evaluate the first order condition at $x = x^*_{SR}$. This gives

$$0 = V'_{SR}(x^*_{SR}) = cdx^*_{SR} - \frac{cd}{2}\pi^e_{SR} + \delta \int_{-\infty}^{x^*_{SR} - x^*_{LR}} V'_{LR}(x^*_{SR} - \pi)\, f_{SR}(\pi)\, d\pi.$$

Since the integral term is strictly positive, it follows that

$$cdx^*_{SR} - \frac{cd}{2}\pi^e_{SR} < 0.$$

Because $cd > 0$, this is equivalent to

$$x^*_{SR} < \frac{1}{2}\pi^e_{SR}.$$

Hence, it holds for aggregate wage growth in period t that

$$x^*_t = \frac{x^*_{SR}}{d} < \frac{\pi^e_{SR}}{2d} = \frac{1}{2}\pi^e_{t,t+1}.$$

Hence, the pass-through from expected inflation into wage inflation satisfies

$$\phi^* < \frac{1}{2}.$$

□

A.2.3. Proof of Proposition 1.3.

Consider a mean-preserving increase in short-run inflation uncertainty from σ_{SR} to $\bar{\sigma}_{SR} > \sigma_{SR}$. Let x^*_{SR} denote the optimal short-run mark-up under σ_{SR} and let $\bar{x}^*_{SR}$ denote the optimal short-run mark-up under $\bar{\sigma}_{SR}$. Let V_{SR} and $\bar{V}_{SR}$ be the corresponding short-run value functions, and let f, F and $\bar{f}, \bar{F}$ denote the corresponding densities and cumulative distribution functions.

By Corollary A.2, for all $x \in \mathbb{R}$,

$$V'_{SR}(x) = cdx - \frac{cd}{2}\pi^e_{SR} + \delta \int_{-\infty}^{x - x^*_{LR}} V'_{LR}(x-\pi)\, f(\pi)\, d\pi,$$
$$\bar{V}'_{SR}(x) = cdx - \frac{cd}{2}\pi^e_{SR} + \delta \int_{-\infty}^{x - x^*_{LR}} V'_{LR}(x-\pi)\, \bar{f}(\pi)\, d\pi.$$

We claim that

$$\bar{V}'_{SR}(x^*_{SR}) > 0.$$

If this claim holds, then $\bar{V}''_{SR}(x) \geq cd > 0$ by Corollary A.2 implies that its unique zero must lie to the left of x^*_{SR}, that is, $\bar{x}^*_{SR} < x^*_{SR}$. Thus, it remains to prove the claim.

Since x^*_{SR} is the unique minimizer of V_{SR}, it holds that

$$V'_{SR}(x^*_{SR}) = 0.$$

Moreover,

$$\bar{V}'_{SR}(x^*_{SR}) - V'_{SR}(x^*_{SR}) = \delta \int_{-\infty}^{x^*_{SR} - x^*_{LR}} V'_{LR}(x^*_{SR} - \pi)\big(\bar{f}(\pi) - f(\pi)\big)\, d\pi.$$

Therefore, to prove $\bar{V}'_{SR}(x^*_{SR}) > 0$, it suffices to show that

$$\int_{-\infty}^{x^*_{SR} - x^*_{LR}} V'_{LR}(x^*_{SR} - \pi)\big(\bar{f}(\pi) - f(\pi)\big)\, d\pi > 0. \tag{A.3}$$

Define

$$G(\pi) := \bar{F}(\pi) - F(\pi) \qquad \text{and} \qquad \gamma(\pi) := V'_{LR}(x^*_{SR} - \pi).$$

Then

$$G'(\pi) = \bar{f}(\pi) - f(\pi),$$
$$\gamma(x^*_{SR} - x^*_{LR}) = V'_{LR}(x^*_{LR}) = 0,$$
$$\gamma'(\pi) = -V''_{LR}(x^*_{SR} - \pi).$$

By Corollary A.1,

$$V''_{LR}(y) \in \left[cd, \frac{cd}{1-\delta}\right] \qquad \text{for all } y \in \mathbb{R},$$

so

$$\gamma'(\pi) = -V''_{LR}(x^*_{SR} - \pi) < 0 \qquad \text{for all } \pi \in \mathbb{R}.$$

Using this notation, the sufficient condition (A.3) becomes

$$\int_{-\infty}^{x^*_{SR}-x^*_{LR}} \gamma(\pi)G'(\pi)\,d\pi > 0.$$

Applying integration by parts yields

$$\int_{-\infty}^{x^*_{SR}-x^*_{LR}} \gamma(\pi)G'(\pi)\,d\pi = \Big[G(\pi)\gamma(\pi)\Big]_{-\infty}^{x^*_{SR}-x^*_{LR}} - \int_{-\infty}^{x^*_{SR}-x^*_{LR}} G(\pi)\gamma'(\pi)\,d\pi.$$

We show that the first term on the right-hand side is equal to zero. Since

$$\gamma(x^*_{SR} - x^*_{LR}) = V'_{LR}(x^*_{LR}) = 0,$$

the upper boundary term is zero. For the lower boundary, note that V''_{LR} is bounded by Corollary A.1, so V'_{LR} has at most linear growth. Hence γ has at most linear growth. Moreover,

$$|G(\pi)| \le F(\pi) + \bar{F}(\pi) \to 0 \qquad \text{as } \pi \to -\infty,$$

and since F and $\bar{F}$ are normal cumulative distribution functions, their lower tails decay faster than any linear growth. Therefore,

$$G(\pi)\gamma(\pi) \to 0 \qquad \text{as } \pi \to -\infty.$$

Thus,

$$\Big[G(\pi)\gamma(\pi)\Big]_{-\infty}^{x^*_{SR}-x^*_{LR}} = 0,$$

and so

$$\int_{-\infty}^{x^*_{SR}-x^*_{LR}} \gamma(\pi)G'(\pi)\,d\pi = -\int_{-\infty}^{x^*_{SR}-x^*_{LR}} G(\pi)\gamma'(\pi)\,d\pi.$$

For normal distributions with identical means and different variances, the higher-variance distribution satisfies

$$G(\pi) = \bar{F}(\pi) - F(\pi) > 0 \qquad \text{for all } \pi < \pi^e_{SR}.$$

By Proposition 1.2,

$$x^*_{SR} < \frac{\pi^e_{SR}}{2}.$$

Moreover, by the maintained parameter restriction,

$$x^*_{LR} \geq -\frac{\pi^e_{SR}}{2}.$$

Hence $x^*_{SR} - x^*_{LR} < \pi^e_{SR}$. Therefore,

$$G(\pi) > 0 \qquad \text{for all } \pi \leq x^*_{SR} - x^*_{LR}.$$

Since also $\gamma'(\pi) < 0$ for all π, it follows that

$$G(\pi)\gamma'(\pi) < 0 \qquad \text{for all } \pi \leq x^*_{SR} - x^*_{LR}.$$

Thus,

$$\int_{-\infty}^{x^*_{SR}-x^*_{LR}} \gamma(\pi)G'(\pi)\,d\pi = -\int_{-\infty}^{x^*_{SR}-x^*_{LR}} G(\pi)\gamma'(\pi)\,d\pi > 0,$$

which proves (A.3). Therefore,

$$\bar{V}'_{SR}(x^*_{SR}) > 0.$$

As shown above, this implies

$$\bar{x}^*_{SR} < x^*_{SR}.$$

Therefore, the pass-through from expected inflation into wage inflation also decreases. □

A.2.4. Proof of Proposition 1.4.

Consider a variance-preserving increase in long-run inflation expectations to $\bar{\pi}^e_{LR} > \pi^e_{LR}$. Denote by (x^*_{LR}, V_{LR}, f) and $(\bar{x}^*_{LR}, \bar{V}_{LR}, \bar{f})$ the long-run optimal wage, value function, and density under π^e_{LR} and $\bar{\pi}^e_{LR}$, respectively.

We first show that for any $y > 0$,

$$\bar{V}'_{LR}(\bar{x}^*_{LR} + y) < V'_{LR}(x^*_{LR} + y).$$

By Corollary A.1, the second derivatives satisfy for any $y \geq 0$

$$\bar{V}''_{LR}(\bar{x}^*_{LR} + y) = cd + \delta \int_{-\infty}^{y} \bar{V}''_{LR}(\bar{x}^*_{LR} + y - \pi)\,\bar{f}(\pi)\,d\pi,$$
$$V''_{LR}(x^*_{LR} + y) = cd + \delta \int_{-\infty}^{y} V''_{LR}(x^*_{LR} + y - \pi)\,f(\pi)\,d\pi.$$

Iterating each equation n times yields

$$\bar{V}''_{LR}(\bar{x}^*_{LR} + y) = cd + \sum_{m=1}^{n} \delta^m \bar{s}_m + \delta^{n+1}\bar{R}_n,$$

$$V''_{LR}(x^*_{LR}+y) = cd + \sum_{m=1}^{n} \delta^m s_m + \delta^{n+1} R_n,$$

where,

$$\begin{aligned}
\bar{s}_1 &= cd \int_{-\infty}^{y} \bar{f}(\pi_1) d\pi_1 \\
\bar{s}_2 &= cd \int_{-\infty}^{y} \int_{-\infty}^{y-\pi_1} \bar{f}(\pi_2)\bar{f}(\pi_1) d\pi_2 d\pi_1 \\
\bar{R}_1 &= \int_{-\infty}^{y} \int_{-\infty}^{y-\pi_1} \bar{V}''_{LR}(\bar{x}^*_{LR}+y-\pi_1-\pi_2)\bar{f}(\pi_2)\bar{f}(\pi_1) d\pi_2 d\pi_1
\end{aligned}$$

and for $m \geq 3$,

$$\begin{aligned}
\bar{s}_m &:= cd \int_{-\infty}^{y} \cdots \int_{-\infty}^{y-(\pi_1+\cdots+\pi_{m-1})} \prod_{j=1}^{m} \bar{f}(\pi_j)\, d\pi_m \cdots d\pi_1, \\
\bar{R}_n &:= \int_{-\infty}^{y} \ldots \int_{-\infty}^{y-(\pi_1+\ldots+\pi_n)} \bar{V}''_{LR}(\bar{x}^*_{LR}+y-(\pi_1+\ldots+\pi_{n+1})) \prod_{j=1}^{n+1} \bar{f}(\pi_j)\, d\pi_{n+1} \cdots d\pi_1.
\end{aligned}$$

The terms s_m and R_n are defined analogously, replacing $\bar{f}$, $\bar{V}''_{LR}$, and $\bar{x}^*_{LR}$ by f, V''_{LR}, and x^*_{LR}. Let now for $m \geq 1$

$$A_m(y) := \Big\{(\pi_1, \ldots, \pi_m) \in \mathbb{R}^m : \sum_{j=1}^{k} \pi_j \leq y \text{ for all } k = 1, \ldots, m\Big\}.$$

Then,

$$\bar{s}_m = cd \int_{A_m(y)} \prod_{j=1}^{m} \bar{f}(\pi_j)\, d\pi_m \cdots d\pi_1, \qquad s_m = cd \int_{A_m(y)} \prod_{j=1}^{m} f(\pi_j)\, d\pi_m \cdots d\pi_1.$$

Note that the products of the density functions simply equal the joint density of m i.i.d. draws from the underlying normal distributions. Further note that $\bar{f}(\pi) = f(\pi - \Delta)$ with $\Delta := \bar{\pi}^e_{LR} - \pi^e_{LR} > 0$. Define $x_j := \pi_j - \Delta$. Then $d\pi_j = dx_j$ and $\bar{f}(\pi_j) = f(x_j)$. Moreover,

$$\sum_{j=1}^{k} \pi_j \leq y \quad \Longleftrightarrow \quad \sum_{j=1}^{k} x_j \leq y - k\Delta \qquad (k = 1, \ldots, m).$$

and hence

$$\bar{s}_m = cd \int_{B_m(y)} \prod_{j=1}^{m} f(x_j)\, dx_m \cdots dx_1,$$

where

$$B_m(y) := \Big\{x \in \mathbb{R}^m : \sum_{j=1}^{k} x_j \leq y - k\Delta \text{ for all } k = 1, \ldots, m\Big\}.$$

Because $y - k\Delta < y$ for all $k \geq 1$, we have the strict inclusion $B_m(y) \subsetneq A_m(y)$. Because integrating the joint density over a set is by definition the probability that a random vector lies in that set, and the set $B_m(y) \subsetneq A_m(y)$, it follows that

$$\bar{s}_m = cd \int_{B_m(y)} \prod_{j=1}^{m} f(x_j)\, dx_m \cdots dx_1 < cd \int_{A_m(y)} \prod_{j=1}^{m} f(\pi_j)\, d\pi_m \cdots d\pi_1 = s_m \qquad \forall m \geq 1.$$

Moreover, $\bar{s}_m, s_m \in (0, cd)$ for all $m \geq 1$. Hence,

$$0 < \sum_{m=1}^{\infty} \delta^m \bar{s}_m \leq cd \sum_{m=1}^{\infty} \delta^m < \infty, \qquad 0 < \sum_{m=1}^{\infty} \delta^m s_m \leq cd \sum_{m=1}^{\infty} \delta^m < \infty.$$

Thus, both sums converge. Since V''_{LR} and $\bar{V}''_{LR}$ are bounded by Corollary A.1, the remainder terms $\delta^{n+1} R_n$ and $\delta^{n+1} \bar{R}_n$ converge to zero. Letting $n \to \infty$ therefore yields

$$\bar{V}''_{LR}(\bar{x}^*_{LR} + y) = cd + \sum_{m=1}^{\infty} \delta^m \bar{s}_m < cd + \sum_{m=1}^{\infty} \delta^m s_m = V''_{LR}(x^*_{LR} + y), \tag{A.4}$$

which follows from $\bar{s}_m < s_m$ for all $m \in \mathbb{N}$. Hence,

$$0 < \bar{V}''_{LR}(\bar{x}^*_{LR} + y) < V''_{LR}(x^*_{LR} + y) \quad \forall y > 0. \tag{A.5}$$

Since $\bar{V}'_{LR}(\bar{x}^*_{LR}) = V'_{LR}(x^*_{LR}) = 0$, we can write

$$\bar{V}'_{LR}(\bar{x}^*_{LR} + y) = \int_0^y \bar{V}''_{LR}(\bar{x}^*_{LR} + t)\, dt, \qquad V'_{LR}(x^*_{LR} + y) = \int_0^y V''_{LR}(x^*_{LR} + t)\, dt.$$

Because the integrand is strictly smaller pointwise, it follows that

$$\bar{V}'_{LR}(\bar{x}^*_{LR} + y) < V'_{LR}(x^*_{LR} + y) \quad \forall y > 0. \tag{A.6}$$

We next show that $\bar{x}^*_{LR} > x^*_{LR}$. Assume by contradiction that $\bar{x}^*_{LR} \leq x^*_{LR}$. Since

$$\bar{V}'_{LR}(\bar{x}^*_{LR}) = cd\bar{x}^*_{LR} - \frac{cd}{2}\bar{\pi}^e_{LR} + \delta \int_{-\infty}^{0} \bar{V}'_{LR}(\bar{x}^*_{LR} - \pi)\bar{f}(\pi)d\pi = 0,$$
$$V'_{LR}(x^*_{LR}) = cdx^*_{LR} - \frac{cd}{2}\pi^e_{LR} + \delta \int_{-\infty}^{0} V'_{LR}(x^*_{LR} - \pi)f(\pi)d\pi = 0,$$

this implies that

$$\int_{-\infty}^{0} \bar{V}'_{LR}(\bar{x}^*_{LR} - \pi)\bar{f}(\pi)d\pi > \int_{-\infty}^{0} V'_{LR}(x^*_{LR} - \pi)f(\pi)d\pi. \tag{A.7}$$

Using that $\bar{f}(\pi) = f(\pi - \Delta)$ with $\Delta := \bar{\pi}^e_{LR} - \pi^e_{LR} > 0$, perform the change of variables $\pi = z + \Delta$:

$$\int_{-\infty}^{0} \bar{V}'_{LR}(\bar{x}^*_{LR} - \pi)\bar{f}(\pi)d\pi = \int_{-\infty}^{-\Delta} \bar{V}'_{LR}(\bar{x}^*_{LR} - \Delta - z)f(z)dz.$$

Fix any $z < -\Delta$ and define $y := -\Delta - z > 0$. Then

$$\bar{x}^*_{LR} - \Delta - z = \bar{x}^*_{LR} + y.$$

By (A.6),

$$\bar{V}'_{LR}(\bar{x}^*_{LR} + y) < V'_{LR}(x^*_{LR} + y).$$

Since $V''_{LR} > 0$, the function V'_{LR} is strictly increasing. Moreover,

$$x^*_{LR} + y = x^*_{LR} - \Delta - z < x^*_{LR} - z.$$

Thus, for all $z < -\Delta$,

$$\bar{V}'_{LR}(\bar{x}^*_{LR} - \Delta - z) < V'_{LR}(x^*_{LR} - z).$$

Multiplying by $f(z) > 0$ and integrating over $-\infty$ to $-\Delta$ yields

$$\int_{-\infty}^{0} \bar{V}'_{LR}(\bar{x}^*_{LR} - \pi)\bar{f}(\pi)d\pi < \int_{-\infty}^{-\Delta} V'_{LR}(x^*_{LR} - z)f(z)dz.$$

Since $V'_{LR}(x^*_{LR} - z) > 0$ for all $z < 0$, we have

$$\int_{-\infty}^{-\Delta} V'_{LR}(x^*_{LR} - z)f(z)dz < \int_{-\infty}^{0} V'_{LR}(x^*_{LR} - \pi)f(\pi)d\pi.$$

Therefore,

$$\int_{-\infty}^{0} \bar{V}'_{LR}(\bar{x}^*_{LR} - \pi)\bar{f}(\pi)d\pi < \int_{-\infty}^{0} V'_{LR}(x^*_{LR} - \pi)f(\pi)d\pi,$$

which contradicts inequality (A.7). Therefore,

$$\bar{x}^*_{LR} > x^*_{LR}. \tag{A.8}$$

We now turn to the short run. Denote by (x^*_{SR}, V_{SR}) and $(\bar{x}^*_{SR}, \bar{V}_{SR})$ the short-run optimal mark-up and value function under π^e_{LR} and $\bar{\pi}^e_{LR}$, respectively. Let f_{SR} now denote the short-run density.

Assume by contradiction that $\bar{x}^*_{SR} \leq x^*_{SR}$. Since $\bar{V}_{SR}$ is strictly convex, $\bar{V}'_{SR}(x^*_{SR}) \geq 0$. Using the derivative representation from Corollary A.2, we obtain

$$\bar{V}'_{SR}(x^*_{SR}) = cdx^*_{SR} - \frac{cd}{2}\pi^e_{SR} + \delta \int_{-\infty}^{x^*_{SR} - \bar{x}^*_{LR}} \bar{V}'_{LR}(x^*_{SR} - \pi)\, f_{SR}(\pi)\, d\pi \;\geq\; 0$$

$$V'_{SR}(x^*_{SR}) = cdx^*_{SR} - \frac{cd}{2}\pi^e_{SR} + \delta \int_{-\infty}^{x^*_{SR} - x^*_{LR}} V'_{LR}(x^*_{SR} - \pi)\, f_{SR}(\pi)\, d\pi = 0.$$

Hence,

$$\int_{-\infty}^{x^*_{SR}-\bar{x}^*_{LR}} \bar{V}'_{LR}(x^*_{SR}-\pi)\, f_{SR}(\pi)\, d\pi \;\geq\; \int_{-\infty}^{x^*_{SR}-x^*_{LR}} V'_{LR}(x^*_{SR}-\pi)\, f_{SR}(\pi)\, d\pi. \tag{A.9}$$

Because $\bar{x}^*_{LR} > x^*_{LR}$, we have $x^*_{SR}-\bar{x}^*_{LR} < x^*_{SR}-x^*_{LR}$, and since $V'_{LR}(x) > 0$ for all $x > x^*_{LR}$,

$$\int_{-\infty}^{x^*_{SR}-x^*_{LR}} V'_{LR}(x^*_{SR}-\pi)\, f_{SR}(\pi)\, d\pi \;>\; \int_{-\infty}^{x^*_{SR}-\bar{x}^*_{LR}} V'_{LR}(x^*_{SR}-\pi)\, f_{SR}(\pi)\, d\pi. \tag{A.10}$$

Combining (A.9) and (A.10) yields

$$\int_{-\infty}^{x^*_{SR}-\bar{x}^*_{LR}} \bar{V}'_{LR}(x^*_{SR}-\pi)\, f_{SR}(\pi)\, d\pi \;>\; \int_{-\infty}^{x^*_{SR}-\bar{x}^*_{LR}} V'_{LR}(x^*_{SR}-\pi)\, f_{SR}(\pi)\, d\pi. \tag{A.11}$$

We now show that the integrand on the left-hand side of (A.11) is pointwise smaller. Fix any $\pi \leq x^*_{SR}-\bar{x}^*_{LR}$ and set $z := x^*_{SR}-\pi$. Then $z \geq \bar{x}^*_{LR}$. Write $z = \bar{x}^*_{LR}+\bar{y}$ with $\bar{y} := z-\bar{x}^*_{LR} \geq 0$ and also $z = x^*_{LR}+y$ with $y := z-x^*_{LR} = \bar{y}+(\bar{x}^*_{LR}-x^*_{LR}) > \bar{y}$. Using $V'_{LR}(x^*_{LR}) = \bar{V}'_{LR}(\bar{x}^*_{LR}) = 0$

$$\bar{V}'_{LR}(z) = \int_0^{\bar{y}} \bar{V}''_{LR}(\bar{x}^*_{LR}+t)\, dt \;\leq\; \int_0^{\bar{y}} V''_{LR}(x^*_{LR}+t)\, dt \;<\; \int_0^{y} V''_{LR}(x^*_{LR}+t)\, dt = V'_{LR}(z),$$

where the first inequality uses the previously established result $V''_{LR}(x^*_{LR}+t) > \bar{V}''_{LR}(\bar{x}^*_{LR}+t) > 0$ for all $t > 0$, see (A.5), and the second strict inequality uses $y > \bar{y}$ and $V''_{LR} > 0$.

Thus, for all $\pi \leq x^*_{SR}-\bar{x}^*_{LR}$,

$$0 \leq \bar{V}'_{LR}(x^*_{SR}-\pi) < V'_{LR}(x^*_{SR}-\pi),$$

and multiplying by $f_{SR}(\pi) > 0$ and integrating implies

$$\int_{-\infty}^{x^*_{SR}-\bar{x}^*_{LR}} \bar{V}'_{LR}(x^*_{SR}-\pi)\, f_{SR}(\pi)\, d\pi \;<\; \int_{-\infty}^{x^*_{SR}-\bar{x}^*_{LR}} V'_{LR}(x^*_{SR}-\pi)\, f_{SR}(\pi)\, d\pi,$$

which contradicts (A.11). Hence, $\bar{x}^*_{SR} > x^*_{SR}$, completing the proof. □

A.3. Model Appendix

A.3.1. A Sufficient Parameter Condition for the Uncertainty Comparative Static

We now derive a conservative sufficient condition for the technical restriction

$$x^*_{LR} \geq -\frac{\pi^e_{SR}}{2} \tag{A.12}$$

in terms of the distributional parameters.

By Corollary A.1, the long-run optimal markup x^*_{LR} satisfies

$$x^*_{LR} = \frac{\pi^e_{LR}}{2} - \frac{\delta}{cd}\int_{-\infty}^{0} V'_{LR}(x^*_{LR} - \pi) f_{LR}(\pi)\, d\pi. \tag{A.13}$$

To bound the integral term, note that Corollary A.1 implies

$$V'_{LR}(x^*_{LR}) = 0 \qquad \text{and} \qquad V''_{LR}(x) \in \left[cd, \frac{cd}{1-\delta}\right] \quad \text{for all } x \in \mathbb{R}.$$

Fix any $\pi \leq 0$. Then,

$$V'_{LR}(x^*_{LR} - \pi) = \int_0^{-\pi} V''_{LR}(x^*_{LR} + s)\, ds.$$

Using the upper bound on V''_{LR}, we obtain

$$V'_{LR}(x^*_{LR} - \pi) \leq \int_0^{-\pi} \frac{cd}{1-\delta}\, ds = \frac{cd}{1-\delta}(-\pi).$$

Multiplying by $f_{LR}(\pi) \geq 0$ and integrating over $(-\infty, 0]$ yields

$$\int_{-\infty}^{0} V'_{LR}(x^*_{LR} - \pi) f_{LR}(\pi)\, d\pi \leq \frac{cd}{1-\delta}\int_{-\infty}^{0} (-\pi) f_{LR}(\pi)\, d\pi.$$

Substituting this bound into (A.13) gives

$$x^*_{LR} \geq \frac{\pi^e_{LR}}{2} - \frac{\delta}{1-\delta}\int_{-\infty}^{0} (-\pi) f_{LR}(\pi)\, d\pi.$$

Hence, a sufficient condition for

$$x^*_{LR} \geq -\frac{\pi^e_{SR}}{2}$$

is

$$\frac{\pi^e_{LR}}{2} - \frac{\delta}{1-\delta}\int_{-\infty}^{0} (-\pi) f_{LR}(\pi)\, d\pi \geq -\frac{\pi^e_{SR}}{2}. \tag{A.14}$$

It remains to evaluate the integral. For $\pi \sim \mathcal{N}(\pi^e_{LR}, \sigma^2_{LR})$, the conditional mean under left truncation at zero is

$$\mathbb{E}[\pi \mid \pi \leq 0] = \pi^e_{LR} - \sigma_{LR} \frac{\phi(-\pi^e_{LR}/\sigma_{LR})}{\Phi(-\pi^e_{LR}/\sigma_{LR})},$$

where ϕ and Φ denote the density and cumulative distribution function of the standard normal distribution. Since

$$\int_{-\infty}^{0} (-\pi) f_{LR}(\pi)\, d\pi = -\mathbb{E}[\pi \mid \pi \leq 0]\mathbb{P}(\pi \leq 0),$$

it follows that

$$\begin{aligned}\int_{-\infty}^{0} (-\pi) f_{LR}(\pi)\, d\pi &= -\left(\pi^e_{LR} - \sigma_{LR} \frac{\phi(-\pi^e_{LR}/\sigma_{LR})}{\Phi(-\pi^e_{LR}/\sigma_{LR})}\right) \Phi\left(-\frac{\pi^e_{LR}}{\sigma_{LR}}\right) \\ &= \sigma_{LR}\phi\left(-\frac{\pi^e_{LR}}{\sigma_{LR}}\right) - \pi^e_{LR}\Phi\left(-\frac{\pi^e_{LR}}{\sigma_{LR}}\right).\end{aligned}$$

Using symmetry of ϕ and substituting into (A.14) yields the conservative sufficient condition

$$\pi^e_{LR} \geq \frac{2\delta}{1-\delta}\left[\sigma_{LR}\phi\left(\frac{\pi^e_{LR}}{\sigma_{LR}}\right) - \pi^e_{LR}\Phi\left(-\frac{\pi^e_{LR}}{\sigma_{LR}}\right)\right] - \pi^e_{SR}.$$

Generally, the condition holds for a broad region of the parameter space. Importantly, the sufficient parameter condition is highly conservative. It is stronger than necessary to guarantee the technical restriction (A.12), and this technical restriction is itself stronger than needed to establish the threshold condition required for the proof of Proposition 1.3.

A.4. Empirical Appendix

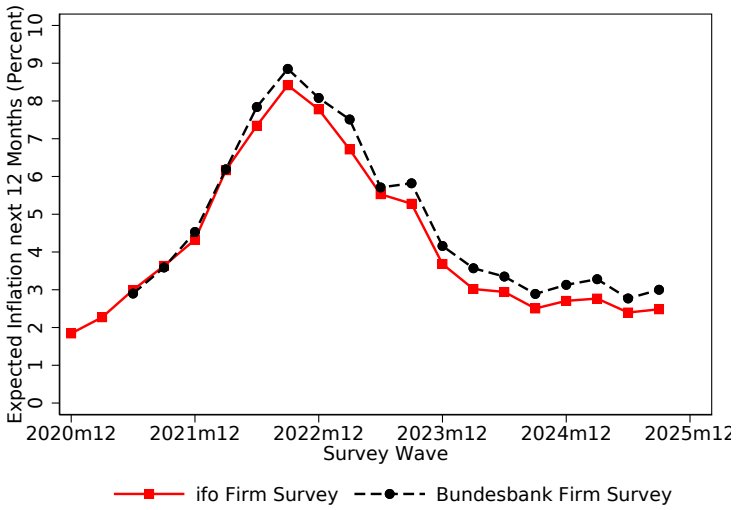
A.4.1. Additional Figures and Tables

Table A.1: Pass-Through by Long-Run Inflation Expectations: Unconditional Median Split

	Expected Wage Growth next 12 Months			
	(1)	(2)	(3)	(4)
Expected Inflation next 12m×Low LR Expectations	0.068***	0.045***	0.141***	0.072***
	(0.017)	(0.015)	(0.024)	(0.019)
Expected Inflation next 12m×High LR Expectations	0.096***	0.047***	0.160***	0.071***
	(0.014)	(0.011)	(0.021)	(0.016)
Observations	35816	35298	35816	35298
R2	0.101	0.478	0.107	0.482
Firm-specifics	yes	no	yes	no
Firm FE	no	yes	no	yes
Survey Wave FE	no	no	yes	yes
Realized Inflation	yes	yes	no	no

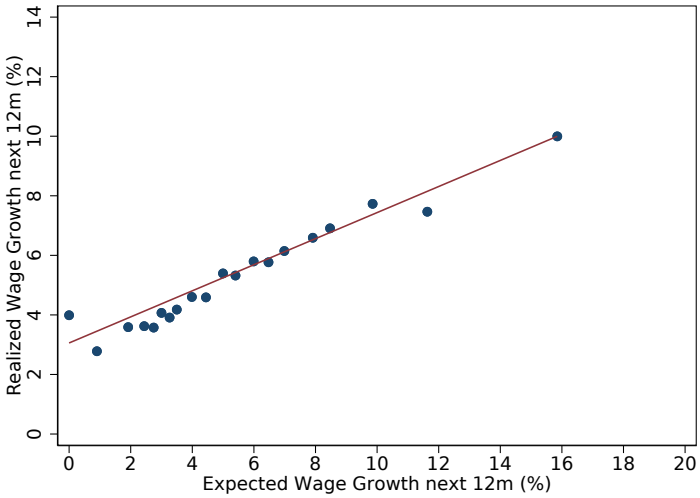
Note: This table reports results from regressions of expected wage growth over the next 12 months on expected CPI inflation over the same horizon interacted with indicators for low and high long-run inflation expectations. Firms are classified as having high long-run inflation expectations if their long-run inflation expectation lies above the unconditional median across all observations. All specifications control for the expected unemployment rate over the next 12 months and firms' current and expected future business conditions. Columns (1) and (3) additionally include time-invariant firm controls. Columns (2) and (4) replace these controls with firm fixed effects. Columns (3) and (4) include survey-wave fixed effects. Realized CPI inflation over the previous 24 months is included in columns (1) and (2) and absorbed by survey-wave fixed effects in columns (3) and (4). Standard errors are clustered at the firm level. Levels of significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure A.1: Expected Inflation in the ifo Firm Survey Compared to Bundesbank Survey



Note: This figure compares the development of average expected CPI inflation (12 months ahead) in Germany in the ifo firm surveys over time to the development of average expectations in the representative Bundesbank Online Panels of Firms.

Figure A.2: Expected and Ex Post Realized Wage Growth



Note: The figure shows, for bins of expected wage growth, the average ex post realized wage growth over the past 12 months reported in the survey 12 months later. By construction, the sample is restricted to observations referring to wage growth over the same period, i.e., expectations elicited between December 2021 and September 2024 and ex post realizations reported between December 2022 and September 2025.

APPENDIX B

Appendix to Chapter 2

B.1. Preliminary Lemmas and Results

Functional notation. Throughout this appendix, we use $(x, y) \in \mathbb{R}^2_+$ to denote generic arguments representing masses of high- and low-ability applicants. These variables serve as inputs to the hiring probability functions p_h and p_l, as well as to all functions derived from them below, including in particular the ratio R, and level-set mappings g_r and P_r.

Accordingly, ∂_x and ∂_y denote partial derivatives with respect to the first and second argument of these functions.

These variables are purely analytical objects. In particular, they should be distinguished from the model parameter x in the main text, which denotes the true share of overconfident job seekers, and from $\hat{x}$, which denotes the corresponding belief.

Market-specific applicant masses. To simplify notation, we denote by (x_H, y_H) and (x_L, y_L) the subjective masses of high- and low-ability applicants in the high-wage and low-wage markets, respectively. These objects correspond to $\hat{\delta}_a^w$ as defined in the main text.

We adopt this notation in place of $\hat{\delta}_a^w$ throughout the appendix in order to keep expressions tractable. In particular, the use of (x_H, y_H, x_L, y_L) avoids repeated double subscripts and superscripts, which would complicate the presentation of equilibrium conditions, derivatives, and comparative statics arguments.

When analyzing equilibrium behavior, we evaluate the functions p_h, p_l, and R at the corresponding market-specific arguments, for example $(x, y) = (x_H, y_H)$ in the high-wage market. For brevity, we will often suppress this substitution and directly write expressions such as $p_h(x_H, y_H)$.

Roadmap of the Mathematical Appendix

This section gives a short overview of how the model's results are derived. The appendix contains several preliminary lemmas. Some are purely technical, some establish properties of the key objects used in the proofs, and some derive or rule out equilibrium behavior.

One of the central objects in the proofs is the ratio

$$R(x, y) := \frac{p_h(x, y)}{p_l(x, y)},$$

where x and y denote generic masses of high- and low-ability applicants. This ratio measures the relative hiring advantage of high-ability applicants in a market with applicant composition (x, y). We first establish useful properties of R. These properties then allow us to compare the high-wage and low-wage markets in equilibrium.

Step 1: Making the hiring probabilities workable and establishing technical results. Lemma B.1 and Lemma B.2 derive a tractable representation of the hiring probabilities. This representation makes it possible to differentiate the hiring probabilities and compare how they react to changes in the applicant composition. Lemma B.3 provides a technical result on reweighting of probability measures that is used repeatedly in these comparisons. Lemma B.6 establishes a monotonicity property of an auxiliary function that is also used repeatedly.

Step 2: Showing that congestion raises the relative value of ability. Lemma B.4 shows that R is increasing in both applicant masses. The economic intuition is that more congestion hurts low-ability applicants relatively more than high-ability applicants. Hence, as the applicant pool becomes more crowded, high ability becomes more valuable in relative terms. This result is the first major building block: it gives the relative hiring advantage R a simple monotone structure.

Step 3: Studying level sets of the relative hiring advantage. The next step studies the level sets of R. This matters because, in a fully mixed equilibrium, both types are indifferent between the high-wage and low-wage markets. These two indifference conditions imply that the two market compositions must have the same value of R. Thus the high-wage and low-wage markets lie on the same level set of the relative hiring advantage.

Lemma B.5 shows that we can describe all market compositions with the same value of R in a simple way. Since the previous step shows that both applicant masses increase R, keeping R fixed requires a tradeoff between the two masses: if the mass of high-ability applicants is higher, the mass of low-ability applicants must be lower. Thus each level set of R can be represented by a unique decreasing function $g_r(x)$.

Lemma B.7 shows that $P_r(x) := p_l(x, g_r(x))$ is strictly decreasing in x. Thus, among market compositions with the same relative hiring advantage R, the market with more high-ability applicants gives low-ability applicants a lower hiring probability.[1] This result allows us to translate an equilibrium comparison of hiring probabilities into a comparison of applicant masses across the two markets.

Step 4: Proving the direct effect. The direct effect requires showing that high-ability job seekers apply to the high-wage market at least as strongly as low-ability job seekers, $D \geq 0$.

Lemma B.8 proves that $D > 0$ in every fully mixed equilibrium. Since both types mix, both types must be indifferent between the two markets. The two indifference conditions imply that the high-wage and low-wage market compositions lie on the same level set of R. The indifference conditions also imply that the hiring probabilities are lower in the high-wage market, because $w_h > w_l$. By Step 3, this can only happen if the high-wage market has the larger high-ability applicant mass. Since the relevant level set g_r is decreasing, the high-wage market must also have the smaller low-ability applicant mass. Together, these inequalities imply $D > 0$.

Lemma B.9 – Lemma B.11 rule out boundary cases that could otherwise undermine the direct-effect conclusion. Finally, the proof of Proposition 2.1 combines all these results to characterize all feasible equilibria and conclude $D \geq 0$.

Step 5: Preparing the indirect effect. The indirect effect asks how equilibrium behavior changes when the belief $\hat{x}$ changes. The key observation is that beliefs enter the hiring probabilities only through the perceived applicant composition. In particular, what matters for the high-wage market is the perceived mass of low-ability applicants,

$$z := \hat{x}\lambda_h + (1 - \hat{x})\lambda_l.$$

In a fully mixed equilibrium, the indifference conditions can therefore be characterized in terms of λ_h and z. Hence, if $\hat{x}$ changes, holding λ_h and z fixed preserves the indifference conditions. This requires an offsetting change in λ_l, as long as the implied value of λ_l remains interior.

Lemma B.12 shows that there is at most one fully mixed equilibrium for a given belief. This uniqueness result is important because it allows us to pin down the fully mixed equilibrium branch uniquely in Lemma B.13. Along this branch, λ_h remains constant, while λ_l adjusts to keep the perceived applicant composition unchanged.

[1] Conversely, along a fixed level set of R, increasing the mass of low-ability applicants requires lowering the mass of high-ability applicants. This replaces strong competitors with weak competitors. The reduction in high-ability competition dominates the increase in low-ability competition, so absolute hiring probabilities rise, while the relative hiring advantage of high ability remains fixed.

Lemma B.14 extends the same logic to equilibria in which high-ability job seekers apply to the high-wage market with probability one, while low-ability job seekers still mix. Lemma B.15 then considers equilibria in which low-ability job seekers do not apply to the high-wage market. Together, these results cover the relevant equilibrium cases needed to characterize the indirect-effect equilibrium branch.

Step 6: Proving the indirect effect. Lemma B.16 combines the insights from Step 5 and characterizes how an equilibrium branch can transition between the relevant equilibrium cases. The proof of Proposition 2.2 then follows directly from this characterization.

Step 7: Proving remaining results. The remaining results follow with little additional work. The proof of existence result, Lemma 2.2, uses a standard argument based on Kakutani's fixed point theorem. The proofs of Corollary 2.1 and Corollary 2.2 follow immediately from the direct- and indirect-effect results.

Lemma B.1 *Consider the hiring probability function in Lemma 2.1. Define*

$$a := \exp\left(\frac{\alpha_l - \alpha_h}{\beta}\right).$$

Then $a \in (0,1)$. A high-ability applicant facing i other high-ability applicants and j low-ability applicants in an applicant pool is hired with probability

$$\frac{1}{1+i+ja}.$$

A low-ability applicant facing i high-ability applicants and j other low-ability applicants in an applicant pool is hired with probability

$$\frac{a}{i+(1+j)a}.$$

PROOF. Since $\alpha_h > \alpha_l$ and $\beta > 0$, we have

$$\frac{\alpha_l - \alpha_h}{\beta} < 0.$$

Hence

$$a = \exp\left(\frac{\alpha_l - \alpha_h}{\beta}\right) \in (0,1).$$

Consider a high-ability applicant. If a high-ability applicant faces i other high-ability applicants and j low-ability applicants in an applicant pool, her hiring probability equals

$$\frac{\exp(\alpha_h/\beta)}{(1+i)\exp(\alpha_h/\beta) + j\exp(\alpha_l/\beta)} = \frac{1}{1+i+j\exp((\alpha_l-\alpha_h)/\beta)} = \frac{1}{1+i+ja}.$$

Consider a low-ability applicant. If a low-ability applicant faces i high-ability applicants and j other low-ability applicants in an applicant pool, her hiring probability equals

$$\frac{\exp(\alpha_l/\beta)}{i\exp(\alpha_h/\beta) + (1+j)\exp(\alpha_l/\beta)} = \frac{\exp((\alpha_l-\alpha_h)/\beta)}{i+(1+j)\exp((\alpha_l-\alpha_h)/\beta)} = \frac{a}{i+(1+j)a}.$$

□

Lemma B.2 *For every $(x,y) \in \mathbb{R}^2_+$,*

$$p_h(x,y) = \int_0^1 \exp\big(x(u-1) + y(u^a-1)\big)\,du,$$
$$p_l(x,y) = a\int_0^1 u^{a-1}\exp\big(x(u-1) + y(u^a-1)\big)\,du.$$

PROOF. Within our directed search model of a continuum economy with anonymous strategies, the number of high- and low-ability applications received by a firm follows independent Poisson distributions (e.g., Albrecht et al., 2004, 2006; Galenianos and Kircher,

2009; Kircher, 2009): let $x, y \in \mathbb{R}_+$ denote the mass of high- and low-ability applicants applying to a given type of vacancy. Then

$$f_x(i) = e^{-x}\frac{x^i}{i!} \qquad \text{and} \qquad f_y(j) = e^{-y}\frac{y^j}{j!}$$

denote the probabilities that a firm receives i high-ability applicants and j low-ability applicants, respectively.[2]

The functions $p_h(x, y)$ and $p_l(x, y)$ are obtained by taking expectations of the hiring probabilities from Lemma B.1 over all possible applicant pool compositions. Hence,

$$p_h(x,y) = \sum_{i=0}^{\infty}\sum_{j=0}^{\infty} f_x(i)f_y(j)\frac{1}{1+i+ja},$$
$$p_l(x,y) = \sum_{i=0}^{\infty}\sum_{j=0}^{\infty} f_x(i)f_y(j)\frac{a}{i+(1+j)a}.$$

We first derive the formula for $p_h(x, y)$. For any $c > 0$,

$$\frac{1}{c} = \int_0^1 u^{c-1}\,du.$$

Applying this with $c = 1 + i + ja$ yields

$$\frac{1}{1+i+ja} = \int_0^1 u^{i+ja}\,du = \int_0^1 u^i u^{aj}\,du.$$

Hence

$$p_h(x,y) = \sum_{i=0}^{\infty}\sum_{j=0}^{\infty} f_x(i)f_y(j)\int_0^1 u^i u^{aj}\,du.$$

Since all terms are nonnegative, Tonelli's theorem justifies interchanging summation and integration:

$$p_h(x,y) = \int_0^1 \left(\sum_{i=0}^{\infty} f_x(i)u^i\right)\left(\sum_{j=0}^{\infty} f_y(j)u^{aj}\right)\,du.$$

Using the probability mass function of the Poisson distribution, we obtain

$$\sum_{i=0}^{\infty} f_x(i)u^i = e^{-x}\sum_{i=0}^{\infty}\frac{(xu)^i}{i!} = e^{-x}e^{xu} = e^{x(u-1)},$$

[2]The Poisson distribution arises as the limit of a binomial distribution. In a finite economy, suppose n applications of a given type are sent independently and uniformly at random to m firms. Then the number of applications a given firm receives is binomially distributed with parameters n and $1/m$. As $n, m \to \infty$ with $n/m = k$, this converges to a Poisson distribution with parameter k. In the continuum model, the mass of firms and applicants of a given type is normalized to one, so that k corresponds to the fraction (equivalently, mass) of applicants of a given type applying to a vacancy of a given type (see Galenianos and Kircher, 2009).

$$\sum_{j=0}^{\infty} f_y(j)u^{aj} = e^{-y}\sum_{j=0}^{\infty}\frac{(yu^a)^j}{j!} = e^{-y}e^{yu^a} = e^{y(u^a-1)}.$$

Substituting gives

$$p_h(x,y) = \int_0^1 \exp\bigl(x(u-1)+y(u^a-1)\bigr)\,du.$$

We now turn to $p_l(x,y)$. Again using $\frac{1}{c} = \int_0^1 u^{c-1}\,du$, now with $c = i+(1+j)a$, we obtain

$$\frac{a}{i+(1+j)a} = a\int_0^1 u^{i+(1+j)a-1}\,du = a\int_0^1 u^i u^{aj}u^{a-1}\,du.$$

Hence,

$$p_l(x,y) = \sum_{i=0}^{\infty}\sum_{j=0}^{\infty} f_x(i)f_y(j)a\int_0^1 u^i u^{aj}u^{a-1}\,du.$$

Thus, again by Tonelli's theorem,

$$p_l(x,y) = a\int_0^1 u^{a-1}\left(\sum_{i=0}^{\infty} f_x(i)u^i\right)\left(\sum_{j=0}^{\infty} f_y(j)u^{aj}\right)\,du = a\int_0^1 u^{a-1}\exp\bigl(x(u-1)+y(u^a-1)\bigr)\,du.$$

Lastly, we note that both integrals exist. The integral in p_h exists because its integrand is bounded by 1. The integral in p_l exists because its integrand is bounded by u^{a-1}, which is integrable on $[0,1]$ because $1/a = \int_0^1 u^{a-1}du < \infty$. □

Lemma B.3 *Let μ be a probability measure on $[0,1]$. Assume that μ is not a point mass. Let $w : [0,1] \to [0,\infty]$ be measurable and satisfy*

$$0 < \int w\,d\mu < \infty.$$

Assume moreover that for all $u, v \in [0,1]$ with $u < v$ and $w(u), w(v) < \infty$,

$$w(u) > w(v).$$

Define ν on the subsets of $[0,1]$ by

$$\nu(A) := \frac{\int_A w\,d\mu}{\int w\,d\mu}.$$

Then ν is a probability measure on $[0,1]$. Moreover, for every bounded strictly increasing function $\varphi : [0,1] \to \mathbb{R}$,

$$\int \varphi\,d\nu < \int \varphi\,d\mu.$$

PROOF. Since $w \geq 0$ is measurable and $0 < \int w\, d\mu < \infty$, the function

$$\nu(A) := \frac{\int_A w\, d\mu}{\int w\, d\mu}$$

defines a probability measure on $[0, 1]$. We now prove the inequality

$$\int \varphi\, d\nu < \int \varphi\, d\mu.$$

By construction, $\nu(A) := \int_A \frac{w}{\int w\, d\mu}\, d\mu$, so $\frac{w}{\int w\, d\mu}$ is the Radon–Nikodym derivative of ν with respect to μ, that is, $d\nu/d\mu = \frac{w}{\int w\, d\mu}$.[3] Hence, for every bounded measurable function φ,

$$\int \varphi\, d\nu = \frac{\int \varphi w\, d\mu}{\int w\, d\mu}.$$

Hence, to prove

$$\int \varphi\, d\nu < \int \varphi\, d\mu,$$

it suffices to show that

$$\int \varphi w\, d\mu < \left(\int \varphi\, d\mu\right)\left(\int w\, d\mu\right).$$

Using covariance notation,

$$\int \varphi w\, d\mu = \left(\int \varphi\, d\mu\right)\left(\int w\, d\mu\right) + \mathrm{Cov}_\mu(\varphi, w).$$

So it remains to show that $\mathrm{Cov}_\mu(\varphi, w) < 0$.

Let $U, V \sim \mu$ be independent. Then for any functions g, h for which the expectations below are well defined,

$$\mathrm{Cov}_\mu(g, h) = \mathbb{E}[g(U)h(U)] - \mathbb{E}[g(U)]\mathbb{E}[h(U)].$$

Since U and V are independent and both have distribution μ,

$$\mathbb{E}[g(U)h(V)] = \mathbb{E}[g(U)]\mathbb{E}[h(V)] = \mathbb{E}[g(U)]\mathbb{E}[h(U)].$$

Thus,

$$\mathrm{Cov}_\mu(g, h) = \mathbb{E}[g(U)h(U)] - \mathbb{E}[g(U)h(V)].$$

By the analogous argument,

$$\mathrm{Cov}_\mu(g, h) = \mathbb{E}[g(V)h(V)] - \mathbb{E}[g(V)h(U)].$$

[3]Indeed, since $\nu(A) = \int_A \frac{w}{\int w\, d\mu}\, d\mu$ by construction, ν is absolutely continuous with respect to μ (that is, $\mu(A) = 0$ implies $\nu(A) = 0$). The Radon–Nikodym theorem then gives an almost everywhere unique function $f = d\nu/d\mu$ such that $\nu(A) = \int_A f\, d\mu$. By uniqueness, $f = w/\int w\, d\mu$.

Adding the two equalities and dividing by 2 gives

$$\mathrm{Cov}_\mu(g,h) = \frac{1}{2}\,\mathbb{E}[(g(U)-g(V))(h(U)-h(V))]\,.$$

Applying this identity with $g=\varphi$ and $h=w$, we obtain

$$\mathrm{Cov}_\mu(\varphi,w) = \frac{1}{2}\,\mathbb{E}[(\varphi(U)-\varphi(V))(w(U)-w(V))]\,.$$

Since $\int w\,d\mu < \infty$, the set

$$E := \{u \in [0,1] : w(u) = \infty\}$$

satisfies $\mu(E) = 0$. Hence

$$\mathbb{P}(U \in E) = \mathbb{P}(V \in E) = 0.$$

Therefore, with probability one, both $w(U)$ and $w(V)$ are finite.

Now let $u, v \in [0,1] \setminus E$. If $u < v$, then $\varphi(u) < \varphi(v)$ because φ is strictly increasing, while $w(u) > w(v)$ whenever they are finite by assumption in the lemma. Hence

$$(\varphi(u)-\varphi(v))(w(u)-w(v)) < 0.$$

If $u = v$, then

$$(\varphi(u)-\varphi(v))(w(u)-w(v)) = 0.$$

Thus, whenever $u, v \in [0,1] \setminus E$, we have

$$(\varphi(u)-\varphi(v))(w(u)-w(v)) \le 0,$$

with strict inequality whenever $u \neq v$. Since μ is not a point mass, $\mathbb{P}(U \neq V) = \mathbb{P}(U \neq V,\ U, V \notin E) > 0$. Therefore,

$$\mathbb{E}[(\varphi(U)-\varphi(V))(w(U)-w(V))] < 0,$$

and hence

$$\mathrm{Cov}_\mu(\varphi,w) < 0,$$

which implies

$$\int \varphi\,d\nu < \int \varphi\,d\mu.$$

□

Lemma B.4 *Define*

$$R(x,y) := \frac{p_h(x,y)}{p_l(x,y)}.$$

Then R has continuous partial derivatives on $\mathbb{R}^2_+$ and satisfies

$$R_x(x,y) > 0, \qquad R_y(x,y) > 0 \qquad \forall (x,y) \in \mathbb{R}^2_+.$$

PROOF. Fix $(x,y) \in \mathbb{R}^2_+$ and define

$$\Phi(u;x,y) := \exp\big(x(u-1) + y(u^a - 1)\big).$$

By Lemma B.2,

$$p_h(x,y) = \int_0^1 \Phi(u;x,y)\,du,$$
$$p_l(x,y) = a\int_0^1 u^{a-1}\Phi(u;x,y)\,du.$$

For $(x,y) \in \mathbb{R}^2_+$ and $u \in [0,1]$, we have $0 < \Phi(u;x,y) \le 1$ and

$$\partial_x \Phi(u;x,y) = (u-1)\Phi(u;x,y), \qquad \partial_y \Phi(u;x,y) = (u^a - 1)\Phi(u;x,y).$$

The relevant partial derivatives of the integrands thus exist pointwise for p_h for every $u \in [0,1]$ and for p_l for every $u \in (0,1]$. Moreover, the differentiated integrands for p_h are continuous in (x,y) for every $u \in [0,1]$ and dominated in absolute value by 1, while those for p_l are continuous in (x,y) for every $u \in (0,1]$ and dominated in absolute value by au^{a-1}, which is integrable on $[0,1]$. Dominated convergence therefore justifies differentiation under the integral sign and implies continuity of the partial derivatives. Since $p_l(x,y) > 0$, it follows that $R = p_h/p_l$ has continuous partial derivatives.

Next, define a probability measure μ on $[0,1]$ by

$$\mu(A) = \frac{\int_A \Phi(u;x,y)\,du}{\int_0^1 \Phi(v;x,y)\,dv} \qquad \text{for every } A \subseteq [0,1].$$

Since $\Phi(u;x,y) > 0$ for all $u \in [0,1]$, the measure μ assigns strictly positive mass to every interval of positive length and is therefore not a point mass. Also define

$$w(u) := \begin{cases} u^{a-1}, & \text{if } u \in (0,1], \\ \infty, & \text{if } u = 0. \end{cases}$$

Since $a \in (0,1)$, the exponent $a-1$ is negative, so

$$w(u) > w(v) \qquad \text{for all } 0 < u < v \le 1.$$

Moreover,

$$\int w\,d\mu = \frac{\int_0^1 u^{a-1}\Phi(u;x,y)\,du}{\int_0^1 \Phi(v;x,y)\,dv} \le \frac{\int_0^1 u^{a-1}\,du}{\int_0^1 \Phi(v;x,y)\,dv} = \frac{1}{a\int_0^1 \Phi(v;x,y)\,dv} < \infty.$$

Also, $\int w\,d\mu > 0$ since $w > 0$ on $(0,1]$. Define a further probability measure ν by

$$\nu(A) := \frac{\int_A w\,d\mu}{\int w\,d\mu} \qquad \text{for every } A \subseteq [0,1].$$

Thus, ν satisfies the assumptions of the reweighted measure from Lemma B.3.

We first study monotonicity of R with respect to x. Differentiating the probabilities yields

$$\partial_x p_h(x,y) = \int_0^1 (u-1)\Phi(u;x,y)\,du, \qquad \partial_x p_l(x,y) = a\int_0^1 (u-1)u^{a-1}\Phi(u;x,y)\,du.$$

Therefore,

$$\partial_x \log p_h(x,y) = \frac{\partial_x p_h(x,y)}{p_h(x,y)} = \frac{\int_0^1 (u-1)\Phi(u;x,y)\,du}{\int_0^1 \Phi(u;x,y)\,du} = \int (u-1)\,d\mu(u),$$
$$\partial_x \log p_l(x,y) = \frac{\partial_x p_l(x,y)}{p_l(x,y)} = \frac{a\int_0^1 (u-1)u^{a-1}\Phi(u;x,y)\,du}{a\int_0^1 u^{a-1}\Phi(u;x,y)\,du} = \int (u-1)\,d\nu(u).$$

The function $u \mapsto u-1$ is strictly increasing on $[0,1]$. Thus, Lemma B.3 implies

$$\int (u-1)\,d\nu(u) < \int (u-1)\,d\mu(u).$$

Thus

$$\partial_x \log R(x,y) = \partial_x \log p_h(x,y) - \partial_x \log p_l(x,y) > 0.$$

Since

$$\partial_x \log R(x,y) = \frac{R_x(x,y)}{R(x,y)} > 0$$

and $R(x,y) > 0$, it follows that

$$R_x(x,y) > 0.$$

The argument for y is analogous. Differentiating gives

$$\partial_y p_h(x,y) = \int_0^1 (u^a-1)\Phi(u;x,y)\,du,$$
$$\partial_y p_l(x,y) = a\int_0^1 (u^a-1)u^{a-1}\Phi(u;x,y)\,du,$$

hence

$$\partial_y \log p_h(x,y) = \int (u^a - 1)\, d\mu(u),$$
$$\partial_y \log p_l(x,y) = \int (u^a - 1)\, d\nu(u).$$

Because $u \mapsto u^a - 1$ is strictly increasing on $[0,1]$, Lemma B.3 again yields

$$\partial_y \log p_l(x,y) < \partial_y \log p_h(x,y),$$

so, following the same steps as above, we obtain

$$R_y(x,y) > 0.$$

□

Lemma B.5 *Fix $r > 0$ and define*

$$D_r := \{x \in [0,1] : \exists\, y \in [0,1] \text{ such that } R(x,y) = r\}.$$

Then D_r is an interval. Moreover, there exists a unique continuous function

$$g_r : D_r \to [0,1]$$

such that

$$R(x, g_r(x)) = r \qquad \forall x \in D_r.$$

On the interior of D_r, the function g_r is continuously differentiable and satisfies

$$g_r'(x) = -\frac{R_x(x, g_r(x))}{R_y(x, g_r(x))} < 0.$$

PROOF. Fix $x \in [0,1]$. By Lemma B.4, the map $y \mapsto R(x,y)$ is continuous and strictly increasing on $[0,1]$. Therefore the equation $R(x,y) = r$ has at most one solution $y \in [0,1]$. Moreover, because the map is increasing in y, such a solution exists if and only if r lies between its endpoint values:

$$x \in D_r \quad \Longleftrightarrow \quad R(x,0) \le r \le R(x,1).$$

Hence

$$D_r = \{x \in [0,1] : R(x,0) \le r\} \cap \{x \in [0,1] : r \le R(x,1)\}.$$

Now, again by Lemma B.4, both functions $x \mapsto R(x,0)$ and $x \mapsto R(x,1)$ are continuous and strictly increasing on $[0,1]$. Therefore the sets $\{x \in [0,1] : R(x,0) \le r\}$ and $\{x \in$

$[0,1] : r \leq R(x,1)\}$ are intervals (possibly empty or singletons). Their intersection D_r is therefore also an interval.

For every $x \in D_r$, there exists exactly one $y \in [0,1]$ satisfying $R(x,y) = r$. Define that unique value by

$$g_r(x) := y.$$

This defines a unique function on D_r. To show continuity of g_r, fix $x \in D_r$ and let $(x_n)_{n\geq 1}$ be a sequence in D_r with $x_n \to x$. We show that $g_r(x_n) \to g_r(x)$.

Since $g_r(x_n) \in [0,1]$ for all n, the sequence $(g_r(x_n))$ is bounded. Hence, by Bolzano–Weierstrass, every subsequence admits a further convergent subsequence. Let (x_{n_k}) be a subsequence such that

$$g_r(x_{n_k}) \to \bar{y} \in [0,1].$$

Since $R(x_{n_k}, g_r(x_{n_k})) = r$ for all k,

$$\lim_{k\to\infty} R(x_{n_k}, g_r(x_{n_k})) = r.$$

On the other hand, since $(x_{n_k}, g_r(x_{n_k})) \to (x, \bar{y})$ and R is continuous, we have

$$\lim_{k\to\infty} R(x_{n_k}, g_r(x_{n_k})) = R(x, \bar{y}).$$

Hence $R(x,\bar{y}) = r$. By uniqueness of the solution in y, it follows that $\bar{y} = g_r(x)$. Thus every convergent subsequence of $(g_r(x_n))$ converges to $g_r(x)$. Since every convergent subsequence of $(g_r(x_n))$ has the same limit $g_r(x)$, the whole sequence must converge to $g_r(x)$, i.e., $g_r(x_n) \to g_r(x)$. Hence g_r is continuous.

To derive differentiability of g_r on the interior of D_r, fix $x \in \mathrm{int}(D_r)$. We first show that $g_r(x) \in (0,1)$; since $x \in \mathrm{int}(D_r) \subseteq (0,1)$, this then justifies applying the implicit function theorem below. Suppose first that $g_r(x) = 0$. Then $R(x,0) = r$. Since $x \mapsto R(x,0)$ is strictly increasing, it follows that for every $x' > x$,

$$R(x',0) > r.$$

Because $y \mapsto R(x',y)$ is strictly increasing, we have

$$R(x',y) \geq R(x',0) > r \qquad \forall y \in [0,1],$$

so there exists no $y \in [0,1]$ such that $R(x',y) = r$. Hence $x' \notin D_r$ for all $x' > x$, which contradicts $x \in \mathrm{int}(D_r)$. Thus $g_r(x) \neq 0$. An analogous argument rules out $g_r(x) = 1$: if $g_r(x) = 1$, then $R(x,1) = r$, and the monotonicity properties of R imply that no $x' < x$

belongs to D_r, again contradicting $x \in \text{int}(D_r)$. Thus $g_r(x) \neq 1$. Therefore $g_r(x) \in (0,1)$ if $x \in \text{int}(D_r)$.

Consider now any $x \in \text{int}(D_r)$. Define $F(\tilde{x}, \tilde{y}) := R(\tilde{x}, \tilde{y}) - r$ on $(0,1)^2$. By Lemma B.4, R is continuously differentiable on $(0,1)^2$, hence F is continuously differentiable on $(0,1)^2$. Since F is continuously differentiable and

$$F(x, g_r(x)) = 0 \qquad \text{and} \qquad \partial_y F(x, g_r(x)) = R_y(x, g_r(x)) > 0.$$

we can apply the implicit function theorem. Specifically, there exists an open interval I_0 containing x and a continuously differentiable function $\tilde{g} : I_0 \to \mathbb{R}$ such that

$$F(\tilde{x}, \tilde{g}(\tilde{x})) = 0 \qquad \forall \tilde{x} \in I_0.$$

Because $x \in (0,1)$, $g_r(x) \in (0,1)$, and $\tilde{g}$ is continuous with $\tilde{g}(x) = g_r(x)$, we may shrink I_0 if necessary and obtain an open interval $I \subseteq I_0$ containing x such that $I \subseteq (0,1)$ and $\tilde{g}(I) \subseteq (0,1)$. We now show that $I \subseteq D_r$ and that $\tilde{g}$ coincides with g_r on I. Fix any $\tilde{x} \in I$. Since

$$F(\tilde{x}, \tilde{g}(\tilde{x})) = 0,$$

we have

$$R(\tilde{x}, \tilde{g}(\tilde{x})) = r.$$

Hence $\tilde{x} \in D_r$. By definition of g_r, $g_r(\tilde{x})$ is the unique $y \in [0,1]$ such that $R(\tilde{x}, y) = r$. Since $\tilde{g}(\tilde{x})$ is such a value, uniqueness implies $\tilde{g}(\tilde{x}) = g_r(\tilde{x})$. Therefore $I \subseteq D_r$ and $\tilde{g} = g_r$ on I.

Finally, it follows from the implicit function theorem that g_r is continuously differentiable on $\text{int}(D_r)$ with derivative

$$g_r'(x) = -\frac{\partial_x F(x, g_r(x))}{\partial_y F(x, g_r(x))} = -\frac{R_x(x, g_r(x))}{R_y(x, g_r(x))}.$$

Because both numerator and denominator are strictly positive by Lemma B.4, it follows that $g_r'(x) < 0$. □

Lemma B.6 *The function*

$$\psi(u) := \frac{1-u}{1-u^a}, \qquad u \in [0,1),$$

extends continuously to $u = 1$ by setting $\psi(1) := 1/a$, and the resulting function is strictly increasing on $[0,1]$.

PROOF. On $[0,1)$, the function is continuous because both numerator and denominator are continuous and $1-u^a > 0$ for $u < 1$. At $u = 1$, both numerator and denominator vanish. By l'Hôpital's rule,

$$\lim_{u\to 1} \frac{1-u}{1-u^a} = \lim_{u\to 1} \frac{-1}{-au^{a-1}} = \frac{1}{a}.$$

So defining $\psi(1) := 1/a$ makes ψ continuous on $[0,1]$.

It remains to prove strict monotonicity. For $u \in (0,1)$, differentiation yields

$$\psi'(u) = \frac{(-1)(1-u^a) - (1-u)(-au^{a-1})}{(1-u^a)^2} = \frac{-1 + au^{a-1} + (1-a)u^a}{(1-u^a)^2}.$$

Since $(1-u^a)^2 > 0$, the sign of $\psi'(u)$ is determined by the numerator

$$N(u) := au^{a-1} + (1-a)u^a - 1 = u^{a-1}\big(a + (1-a)u\big) - 1.$$

Define

$$f(u) := u^{a-1}\big(a + (1-a)u\big).$$

Then f is continuous on $(0,1]$ and $f(1) = 1$. Moreover,

$$f'(u) = a\,u^{a-2}(a-1)(1-u) < 0 \qquad \forall u \in (0,1),$$

since $a > 0$, $a-1 < 0$, $1-u > 0$, and $u^{a-2} > 0$ for all $u \in (0,1)$. So f is strictly decreasing on $(0,1)$. Since f is also continuous on $(0,1]$ and $f(1) = 1$, it follows that $f(u) > f(1) = 1$ for all $u \in (0,1)$. Therefore $N(u) > 0$ and $\psi'(u) > 0$ for all $u \in (0,1)$. Since ψ is continuous on $[0,1]$ and strictly increasing on $(0,1)$, ψ is strictly increasing on $[0,1]$. □

Lemma B.7 *Fix $r > 0$ and define*

$$P_r(x) := p_l(x, g_r(x)), \qquad x \in D_r.$$

Then P_r is strictly decreasing on D_r.

PROOF. Fix $x \in \operatorname{int}(D_r)$ and set $y := g_r(x)$, so that $R(x,y) = r$. As in the proof of Lemma B.4, define

$$\Phi(u; x, y) := \exp\big(x(u-1) + y(u^a - 1)\big).$$

As in the proof of Lemma B.4, define probability measures μ and ν on $[0,1]$ by

$$\mu(A) := \frac{\int_A \Phi(u;x,y)\,du}{\int_0^1 \Phi(v;x,y)\,dv}, \qquad \nu(A) := \frac{\int_A w\,d\mu}{\int w\,d\mu} \qquad \text{for every } A \subseteq [0,1],$$

where

$$w(u) := \begin{cases} u^{a-1}, & \text{if } u \in (0,1], \\ \infty, & \text{if } u = 0. \end{cases}$$

Now define

$$f(u) := 1 - u, \qquad h(u) := 1 - u^a.$$

Both functions are strictly decreasing on $[0,1]$. Set

$$A_\mu := \int f \, d\mu, \qquad B_\mu := \int h \, d\mu, \qquad A_\nu := \int f \, d\nu, \qquad B_\nu := \int h \, d\nu.$$

Taking derivatives of the probabilities as in the proof of Lemma B.4, we obtain

$$\begin{aligned} \partial_x \log p_h(x,y) &= \int (u-1) \, d\mu(u) = -A_\mu, \\ \partial_y \log p_h(x,y) &= \int (u^a - 1) \, d\mu(u) = -B_\mu, \\ \partial_x \log p_l(x,y) &= \int (u-1) \, d\nu(u) = -A_\nu, \\ \partial_y \log p_l(x,y) &= \int (u^a - 1) \, d\nu(u) = -B_\nu. \end{aligned}$$

Therefore

$$\begin{aligned} \partial_x \log R(x,y) &= \partial_x \log p_h(x,y) - \partial_x \log p_l(x,y) = A_\nu - A_\mu, \\ \partial_y \log R(x,y) &= \partial_y \log p_h(x,y) - \partial_y \log p_l(x,y) = B_\nu - B_\mu. \end{aligned}$$

By Lemma B.5, g_r is continuously differentiable on the interior of D_r and satisfies

$$g_r'(x) = -\frac{R_x(x,y)}{R_y(x,y)}, \qquad y = g_r(x).$$

Since

$$R_x(x,y) = R(x,y) \, \partial_x \log R(x,y), \qquad R_y(x,y) = R(x,y) \, \partial_y \log R(x,y),$$

it follows that

$$g_r'(x) = -\frac{\partial_x \log R(x,y)}{\partial_y \log R(x,y)}.$$

Hence

$$g_r'(x) = -\frac{A_\nu - A_\mu}{B_\nu - B_\mu}.$$

Since p_l and g_r are continuously differentiable on the interior of D_r, the composition $P_r(x) = p_l(x, g_r(x))$ is differentiable and the chain rule applies:

$$\frac{d}{dx}\log P_r(x) = \partial_x \log p_l(x,y) + g_r'(x)\partial_y \log p_l(x,y) = -A_\nu + \frac{A_\nu - A_\mu}{B_\nu - B_\mu}B_\nu = \frac{A_\nu B_\mu - A_\mu B_\nu}{B_\nu - B_\mu}.$$

To determine the sign, we study numerator and denominator separately.

The function $-h$ is bounded and strictly increasing on $[0,1]$. Hence Lemma B.3, applied to the reweighting from μ to ν together with the strictly increasing function $\varphi = -h$, yields

$$\int (-h)\, d\nu < \int (-h)\, d\mu.$$

Equivalently,

$$B_\nu > B_\mu.$$

So the denominator $B_\nu - B_\mu$ is strictly positive.

We next show that the numerator is strictly negative. Since $h(u) = 1-u^a$ is nonnegative on $[0,1]$, strictly positive on $[0,1)$, and neither μ nor ν is a point mass at $u = 1$, we have $B_\mu > 0$ and $B_\nu > 0$. Hence

$$A_\nu B_\mu - A_\mu B_\nu < 0 \iff \frac{A_\nu}{B_\nu} < \frac{A_\mu}{B_\mu}.$$

To compare these ratios, define new probability measures μ_h and ν_h on $[0,1]$ by reweighting μ and ν with the nonnegative function $h(u) = 1 - u^a$:

$$\mu_h(A) := \frac{\int_A h\, d\mu}{\int h\, d\mu}, \qquad \nu_h(A) := \frac{\int_A h\, d\nu}{\int h\, d\nu}.$$

Let ψ denote the continuous extension of $(1-u)/(1-u^a)$ from Lemma B.6. Since $h(1) = 0$, the measures μ_h and ν_h assign zero mass to $\{1\}$. Therefore

$$\int \psi\, d\mu_h = \int \frac{f}{h}\, d\mu_h = \int \frac{f(u)}{h(u)}\frac{h(u)}{\int h\, d\mu}\, d\mu(u) = \frac{\int f\, d\mu}{\int h\, d\mu} = \frac{A_\mu}{B_\mu},$$
$$\int \psi\, d\nu_h = \int \frac{f}{h}\, d\nu_h = \int \frac{f(u)}{h(u)}\frac{h(u)}{\int h\, d\nu}\, d\nu(u) = \frac{\int f\, d\nu}{\int h\, d\nu} = \frac{A_\nu}{B_\nu}.$$

Lemma B.6 implies that ψ is strictly increasing on $[0,1]$. We now compare μ_h and ν_h. Substituting the definition of ν into ν_h gives

$$\nu_h(A) = \frac{\int_A h\, w\, d\mu}{\int h\, w\, d\mu}.$$

Therefore ν_h is obtained from μ_h by reweighting with the function $w(u)$. Hence Lemma B.3 applied to the reweighting from μ_h to ν_h together with the strictly increasing function ψ yields

$$\int \psi\, d\nu_h < \int \psi\, d\mu_h.$$

That is,

$$\frac{A_\nu}{B_\nu} < \frac{A_\mu}{B_\mu}.$$

Hence the numerator is strictly negative. We can thus conclude that

$$\frac{d}{dx} \log P_r(x) < 0$$

for every interior point x of D_r. Since $P_r(x) = p_l(x, g_r(x)) > 0$, this implies

$$P_r'(x) < 0.$$

Hence P_r is strictly decreasing on the interior of D_r. Since P_r is moreover continuous on the entire interval D_r, strict monotonicity extends to the entire interval. □

Lemma B.8 *Suppose $(\lambda_h^*, \lambda_l^*)$ is a fully mixed equilibrium. Then*

$$D > 0.$$

PROOF. Define the subjective applicant masses in the high-wage market by

$$x_H := \lambda_h^*, \qquad y_H := \hat{x}\,\lambda_h^* + (1-\hat{x})\lambda_l^*,$$

and in the low-wage market by

$$x_L := 1 - \lambda_h^*, \qquad y_L := 1 - y_H.$$

We want to show $\lambda_h^* > \lambda_l^*$. Observe first that

$$x_H - y_H = \lambda_h^* - \hat{x}\,\lambda_h^* - (1-\hat{x})\lambda_l^* = (1-\hat{x})(\lambda_h^* - \lambda_l^*).$$

Since $\hat{x} \in (0,1)$ it is sufficient to prove $x_H > y_H$. Because the equilibrium is fully mixed, both types are indifferent between applying to the high-wage and low-wage markets. Hence

$$p_h(x_H, y_H)w_h = p_h(x_L, y_L)w_l, \qquad p_l(x_H, y_H)w_h = p_l(x_L, y_L)w_l.$$

Thus, for some $r > 0$,

$$R(x_H, y_H) = R(x_L, y_L) = r.$$

So both points (x_H, y_H) and (x_L, y_L) lie on the same level set $\{(x,y) \in [0,1]^2 : R(x,y) = r\}$. By Lemma B.5, this level set can be written as

$$y = g_r(x).$$

Now use the low type indifference condition. Since $w_h > w_l$, it follows that

$$p_l(x_H, y_H) = \frac{w_l}{w_h} p_l(x_L, y_L) < p_l(x_L, y_L).$$

We note that this implies $x_H \neq x_L$, since otherwise $y_H = g_r(x_H) = g_r(x_L) = y_L$, and therefore $p_l(x_H, y_H) = p_l(x_L, y_L)$, which is a contradiction. Since $x_H, x_L \in D_r$ and Lemma B.5 shows that D_r is an interval, it follows that D_r is a nondegenerate interval.

Since (x_H, y_H) and (x_L, y_L) lie on the same level set, we have

$$p_l(x_H, y_H) = P_r(x_H), \qquad p_l(x_L, y_L) = P_r(x_L).$$

Since P_r is strictly decreasing by Lemma B.7 and $P_r(x_H) < P_r(x_L)$, it follows that $x_H > x_L$. Using $x_L = 1 - x_H$, we obtain $x_H > 1 - x_H$ and hence

$$x_H > \frac{1}{2}.$$

By Lemma B.5, the function g_r is continuous on D_r and satisfies $g_r'(x) < 0$ on the interior of D_r. Hence g_r is strictly decreasing on the entire D_r. Because $x_H > x_L$,

$$g_r(x_H) = y_H < y_L = g_r(x_L).$$

Since $y_L = 1 - y_H$, we obtain $y_H < 1 - y_H$, so

$$y_H < \frac{1}{2}.$$

Therefore

$$x_H > \frac{1}{2} > y_H.$$

This implies $D = \lambda_h^* - \lambda_l^* > 0$. □

Lemma B.9 *There is no equilibrium with*

$$\lambda_l^* = 1 \qquad \textit{and} \qquad \lambda_h^* < 1.$$

PROOF. Assume throughout that $\lambda_l^* = 1$. We show that no value $\lambda_h = \lambda < 1$ can then be part of an equilibrium.

Define the subjective applicant masses in the high-wage market by

$$H(\lambda) := \big(x_H(\lambda), y_H(\lambda)\big) := \big(\lambda, \hat{x}\,\lambda + (1 - \hat{x})\big),$$

and in the low-wage market by

$$L(\lambda) := \big(x_L(\lambda), y_L(\lambda)\big) := \big(1 - \lambda, \hat{x}(1 - \lambda)\big).$$

We split the proof into two cases.

Case 1: $\lambda_h^* = 0$. In this case,

$$H(0) = (0, 1 - \hat{x}), \qquad L(0) = (1, \hat{x}).$$

We show that a high-ability job seeker then strictly prefers the high-wage market. Since $w_h > w_l$, it is sufficient to prove

$$p_h\big(H(0)\big) > p_h\big(L(0)\big).$$

By Lemma B.2,

$$p_h\big(H(0)\big) = \int_0^1 \exp\big((1-\hat{x})(u^a - 1)\big)\, du, \qquad p_h\big(L(0)\big) = \int_0^1 \exp\big((u-1) + \hat{x}(u^a - 1)\big)\, du.$$

It is enough to compare the exponents. For every $u \in (0, 1)$,

$$(1-\hat{x})(u^a - 1) - \big((u-1) + \hat{x}(u^a - 1)\big) = (1-u) - (1 - 2\hat{x})(1 - u^a).$$

Now $a \in (0, 1)$ implies $u^a > u$ for every $u \in (0, 1)$, hence $1 - u^a < 1 - u$. If $1 - 2\hat{x} \leq 0$, then

$$(1-u) - (1-2\hat{x})(1-u^a) \geq (1-u) > 0.$$

If $1 - 2\hat{x} > 0$, then $1 - 2\hat{x} < 1$ because $\hat{x} \in (0, 1)$, so

$$(1-2\hat{x})(1-u^a) < 1 - u^a < 1 - u,$$

which again implies

$$(1-u) - (1-2\hat{x})(1-u^a) > 0.$$

Thus for every $u \in (0, 1)$, the exponent in the integral for $p_h(H(0))$ is strictly larger than the one in $p_h(L(0))$. Therefore $p_h\big(H(0)\big) > p_h\big(L(0)\big)$. So a high-ability job seeker strictly prefers the high-wage market. Hence $(\lambda_h^*, \lambda_l^*) = (0, 1)$ cannot be an equilibrium.

Case 2: $0 < \lambda_h^* < 1$. If $(\lambda_h^*, 1)$ were an equilibrium, then high-ability job seekers would mix and therefore have to be indifferent between the two wage markets. Define the high type payoff difference

$$U_h(\lambda) := w_h\, p_h\big(H(\lambda)\big) - w_l\, p_h\big(L(\lambda)\big).$$

Thus, necessarily $U_h(\lambda_h^*) = 0$. We first show that U_h is strictly decreasing on $[0, 1]$. Differentiating $p_h(x, y)$ from Lemma B.2 with respect to x and y yields

$$\partial_x p_h(x, y) = \int_0^1 (u-1) \exp\big(x(u-1) + y(u^a - 1)\big)\, du < 0$$

$$\partial_y p_h(x,y) = \int_0^1 (u^a - 1)\exp\big(x(u-1) + y(u^a-1)\big)\,du < 0$$

since $u - 1 < 0$ and $u^a - 1 < 0$ for every $u \in (0,1)$. Now

$$\frac{d}{d\lambda}H(\lambda) = (1, \hat{x}), \qquad \frac{d}{d\lambda}L(\lambda) = (-1, -\hat{x}).$$

Therefore,

$$\frac{d}{d\lambda}U_h(\lambda) = w_h\Big(\partial_x p_h\big(H(\lambda)\big) + \hat{x}\,\partial_y p_h\big(H(\lambda)\big)\Big) + w_l\Big(\partial_x p_h\big(L(\lambda)\big) + \hat{x}\,\partial_y p_h\big(L(\lambda)\big)\Big) < 0.$$

So U_h is strictly decreasing. Next define

$$F(\lambda) := \log R\big(H(\lambda)\big) - \log R\big(L(\lambda)\big).$$

We claim that F is strictly increasing on $[0,1]$. We differentiate using the chain rule:

$$\frac{d}{d\lambda}F(\lambda) = (\log R)_x\big(H(\lambda)\big) + \hat{x}\,(\log R)_y\big(H(\lambda)\big) + (\log R)_x\big(L(\lambda)\big) + \hat{x}\,(\log R)_y\big(L(\lambda)\big).$$

By Lemma B.4, $R_x > 0$ and $R_y > 0$ on $[0,1]^2$. Since $R > 0$, this implies

$$(\log R)_x = \frac{R_x}{R} > 0, \qquad (\log R)_y = \frac{R_y}{R} > 0.$$

Therefore

$$\frac{d}{d\lambda}F(\lambda) > 0,$$

so F is strictly increasing.

Now evaluate F at $\lambda = \frac{1}{2}$. We have

$$H\left(\frac{1}{2}\right) = \left(\frac{1}{2}, 1 - \frac{\hat{x}}{2}\right), \qquad L\left(\frac{1}{2}\right) = \left(\frac{1}{2}, \frac{\hat{x}}{2}\right).$$

The two points have the same first coordinate, while

$$1 - \frac{\hat{x}}{2} > \frac{\hat{x}}{2}.$$

Since R is strictly increasing in its second argument by Lemma B.4, it follows that

$$R\left(H\left(\frac{1}{2}\right)\right) > R\left(L\left(\frac{1}{2}\right)\right),$$

hence

$$F\left(\frac{1}{2}\right) > 0.$$

Define

$$\lambda_0 := \inf\{\lambda \in [0,1] : F(\lambda) \geq 0\}.$$

Since $F(1/2) > 0$, the set is nonempty and thus λ_0 is well defined. Moreover, $1/2$ belongs to the set, so

$$\lambda_0 \leq \frac{1}{2}.$$

We next show that $U_h(\lambda_0) > 0$. If $\lambda_0 = 0$, this follows from Case 1, which showed that

$$U_h(0) = w_h\, p_h\big(H(0)\big) - w_l\, p_h\big(L(0)\big) > 0.$$

Now suppose $\lambda_0 > 0$. We first note that since R is continuous and strictly positive and since also H and L are continuous functions of λ, it follows that F is continuous on $[0,1]$. Since F is continuous and strictly increasing, it follows that if $\lambda_0 > 0$, then $F(\lambda_0) = 0$. Hence

$$R\big(H(\lambda_0)\big) = R\big(L(\lambda_0)\big).$$

Thus $H(\lambda_0)$ and $L(\lambda_0)$ lie on the same level set of R. Moreover, because $\lambda_0 \leq 1/2$,

$$x_H(\lambda_0) = \lambda_0 \leq 1 - \lambda_0 = x_L(\lambda_0).$$

If equality held, then $\lambda_0 = 1/2$, but this is impossible because $F(1/2) > 0$ while $F(\lambda_0) = 0$. Therefore

$$x_H(\lambda_0) < x_L(\lambda_0).$$

Since both points lie on the same level set of R, Lemma B.7 implies

$$p_l\big(H(\lambda_0)\big) > p_l\big(L(\lambda_0)\big).$$

Because the ratio $R = p_h/p_l$ is the same at both points, it follows immediately that

$$p_h\big(H(\lambda_0)\big) > p_h\big(L(\lambda_0)\big).$$

Then

$$U_h(\lambda_0) = w_h\, p_h\big(H(\lambda_0)\big) - w_l\, p_h\big(L(\lambda_0)\big) > 0,$$

because $w_h > w_l$ and $p_h(H(\lambda_0)) > p_h(L(\lambda_0))$. Thus in all cases $U_h(\lambda_0) > 0$.

Now suppose, toward a contradiction, that $(\lambda_h^*, 1)$ with $0 < \lambda_h^* < 1$ were an equilibrium. Then high-ability job seekers must be indifferent, so $U_h(\lambda_h^*) = 0$. Since U_h is strictly decreasing and $U_h(\lambda_0) > 0$, it follows that

$$\lambda_h^* > \lambda_0.$$

Because F is strictly increasing, this implies

$$F(\lambda_h^*) > F(\lambda_0) \geq 0.$$

Therefore

$$R\big(H(\lambda_h^*)\big) > R\big(L(\lambda_h^*)\big).$$

We now show that this contradicts $\lambda_l^* = 1$. Since high-ability job seekers are indifferent at λ_h^*,

$$\frac{w_l}{w_h} = \frac{p_h(H(\lambda_h^*))}{p_h(L(\lambda_h^*))}.$$

Consider the low type payoff difference

$$U_l(\lambda_h^*) := w_h\, p_l\big(H(\lambda_h^*)\big) - w_l\, p_l\big(L(\lambda_h^*)\big).$$

Dividing by w_h and substituting the expression for w_l/w_h gives

$$\frac{U_l(\lambda_h^*)}{w_h} = p_l\big(H(\lambda_h^*)\big) - \frac{p_h(H(\lambda_h^*))}{p_h(L(\lambda_h^*))} p_l\big(L(\lambda_h^*)\big).$$

Using $p_h = Rp_l$, we obtain

$$\frac{U_l(\lambda_h^*)}{w_h} = p_l\big(H(\lambda_h^*)\big) - \frac{R(H(\lambda_h^*))\, p_l(H(\lambda_h^*))}{R(L(\lambda_h^*))\, p_l(L(\lambda_h^*))}\, p_l\big(L(\lambda_h^*)\big) = p_l\big(H(\lambda_h^*)\big)\left(1 - \frac{R(H(\lambda_h^*))}{R(L(\lambda_h^*))}\right).$$

Since $R(H(\lambda_h^*)) > R(L(\lambda_h^*))$, it follows that

$$U_l(\lambda_h^*) < 0.$$

So low-ability job seekers strictly prefer the low-wage market. This contradicts $\lambda_l^* = 1$, which requires them to apply to the high-wage market with probability one.

We have therefore ruled out both possibilities $\lambda_h^* = 0$ and $0 < \lambda_h^* < 1$. Hence there is no equilibrium with $\lambda_l^* = 1$ and $\lambda_h^* < 1$. □

Lemma B.10 *There is no equilibrium with*

$$\lambda_l^* \in (0,1) \qquad \textit{and} \qquad \lambda_h^* = 0.$$

PROOF. Assume throughout that $\lambda_h^* = 0$ and $\lambda_l = \lambda \in (0,1)$. Given $\lambda_h^* = 0$, the subjective applicant masses in the high-wage market are

$$H(\lambda) := \big(x_H(\lambda), y_H(\lambda)\big) = \big(0, (1-\hat{x})\lambda\big),$$

and in the low-wage market they are

$$L(\lambda) := \big(x_L(\lambda), y_L(\lambda)\big) = \big(1, 1-(1-\hat{x})\lambda\big).$$

To rule out such an equilibrium, it is sufficient to show that a high-ability job seeker strictly prefers the high-wage market. Since $w_h > w_l$, it is sufficient to prove

$$p_h\big(H(\lambda)\big) > p_h\big(L(\lambda)\big).$$

By Lemma B.2,

$$p_h\big(H(\lambda)\big) = \int_0^1 \exp\big((1-\hat{x})\lambda(u^a-1)\big)\,du,$$
$$p_h\big(L(\lambda)\big) = \int_0^1 \exp\big((u-1) + \big(1-(1-\hat{x})\lambda\big)(u^a-1)\big)\,du.$$

It is enough to compare the exponents. For every $u \in (0,1)$,

$$(1-\hat{x})\lambda(u^a-1) - \Big((u-1) + \big(1-(1-\hat{x})\lambda\big)(u^a-1)\Big) = (1-u) + \big(1-2(1-\hat{x})\lambda\big)(1-u^a).$$

Set

$$c := (1-\hat{x})\lambda.$$

Since $\hat{x} \in (0,1)$ and $\lambda \in (0,1)$, we have $0 < c < 1$. Thus

$$(1-\hat{x})\lambda(u^a-1) - \Big((u-1) + \big(1-(1-\hat{x})\lambda\big)(u^a-1)\Big) = (1-u) + (1-2c)(1-u^a).$$

If $1-2c \geq 0$, then

$$(1-u) + (1-2c)(1-u^a) \geq (1-u) > 0.$$

If $1-2c < 0$, then $1-2c > -1$ because $c < 1$, and therefore

$$(1-u) + (1-2c)(1-u^a) > (1-u) - (1-u^a) > 0,$$

since $a \in (0,1)$ implies $u^a > u$ for every $u \in (0,1)$.

Hence for every $u \in (0,1)$, the exponent in the integral for $p_h(H(\lambda))$ is strictly larger than the exponent in the integral for $p_h(L(\lambda))$. Therefore

$$p_h\big(H(\lambda)\big) > p_h\big(L(\lambda)\big).$$

Hence, a high-ability job seeker strictly prefers the high-wage market so that $\lambda_h^* = 0$ cannot be part of an equilibrium when $\lambda_l^* \in (0,1)$. □

Lemma B.11 *There is no equilibrium with*

$$\lambda_l^* = 0 \qquad \textit{and} \qquad \lambda_h^* = 0.$$

PROOF. If $\lambda_l^* = \lambda_h^* = 0$, then no applicant applies to the high-wage market. Hence the subjective applicant masses are $(x_H, y_H) = (0, 0)$ and $(x_L, y_L) = (1, 1)$. A high-ability job seeker is hired with probability $p_h(0, 0) = 1$ in the high-wage market and with probability $p_h(1, 1) < 1$ in the low-wage market. Since also $w_h > w_l$, a high-ability job seeker strictly prefers the high-wage market. Therefore $(\lambda_h^*, \lambda_l^*) = (0, 0)$ cannot be an equilibrium. □

Lemma B.12 *Fix $\hat{x} \in (0, 1)$. Then there exists at most one fully mixed equilibrium with*

$$\lambda_h^* \in (0, 1) \qquad \textit{and} \qquad \lambda_l^* \in (0, 1).$$

PROOF. Fix $\hat{x} \in (0, 1)$. For any candidate pair $(\lambda_h, \lambda_l) \in (0, 1)^2$, define

$$x := \lambda_h, \qquad z := \hat{x}\, \lambda_h + (1 - \hat{x})\lambda_l.$$

Then the subjective applicant masses in the high-wage market are

$$(x_H, y_H) = (x, z),$$

while those in the low-wage market are

$$(x_L, y_L) = (1 - x, 1 - z).$$

It is sufficient to show that there is at most one equilibrium pair $(x, z) \in (0, 1)^2$. Define

$$\begin{aligned} F_h(x, z) &:= w_l\, p_h(1 - x, 1 - z) - w_h\, p_h(x, z), \\ F_l(x, z) &:= w_l\, p_l(1 - x, 1 - z) - w_h\, p_l(x, z). \end{aligned} \tag{B.1}$$

Then $(x, z) \in (0, 1)^2$ is an fully mixed equilibrium if and only if

$$F_h(x, z) = 0 \qquad \text{and} \qquad F_l(x, z) = 0.$$

Step 1: both functions are strictly increasing in both arguments.

By Lemma B.2,

$$\begin{aligned} \partial_x p_h(x, y) &= \int_0^1 (u - 1) \exp\big(x(u - 1) + y(u^a - 1)\big)\, du < 0, \\ \partial_y p_h(x, y) &= \int_0^1 (u^a - 1) \exp\big(x(u - 1) + y(u^a - 1)\big)\, du < 0, \\ \partial_x p_l(x, y) &= a \int_0^1 (u - 1) u^{a-1} \exp\big(x(u - 1) + y(u^a - 1)\big)\, du < 0, \end{aligned}$$

$$\partial_y p_l(x,y) = a\int_0^1 (u^a-1)u^{a-1}\exp\big(x(u-1)+y(u^a-1)\big)\,du < 0$$

for every $(x,y)\in\mathbb{R}^2_+$, because $u-1<0$ and $u^a-1<0$ for all $u\in(0,1)$. Therefore,

$$\begin{aligned}
\partial_x F_h(x,z) &= -w_l\,\partial_x p_h(1-x,1-z) - w_h\,\partial_x p_h(x,z) > 0,\\
\partial_z F_h(x,z) &= -w_l\,\partial_y p_h(1-x,1-z) - w_h\,\partial_y p_h(x,z) > 0,\\
\partial_x F_l(x,z) &= -w_l\,\partial_x p_l(1-x,1-z) - w_h\,\partial_x p_l(x,z) > 0,\\
\partial_z F_l(x,z) &= -w_l\,\partial_y p_l(1-x,1-z) - w_h\,\partial_y p_l(x,z) > 0
\end{aligned}$$

for every $(x,z)\in(0,1)^2$. Hence both F_h and F_l are strictly increasing in x and z.

Step 2: at any intersection point the $F_h=0$ curve is steeper than the $F_l=0$ curve.

Fix $(x,z)\in(0,1)^2$. Define

$$\begin{aligned}
\Phi_A(u) &:= \exp\big(x(u-1)+z(u^a-1)\big),\\
\Phi_B(u) &:= \exp\big((1-x)(u-1)+(1-z)(u^a-1)\big),\\
K(u) &:= w_h\,\Phi_A(u) + w_l\,\Phi_B(u).
\end{aligned}$$

Further define

$$f(u) := 1-u, \qquad h(u) := 1-u^a, \qquad w(u) := \begin{cases} u^{a-1}, & \text{if } u\in(0,1],\\ \infty, & \text{if } u=0. \end{cases}$$

Let $\psi(u) := f(u)/h(u) = (1-u)/(1-u^a)$ with $\psi(1) := 1/a$. By Lemma B.6, $\psi(u)$ is strictly increasing on $[0,1]$. Moreover, $w(u)=u^{a-1}$ is strictly decreasing on $(0,1]$.

Using the formulas for the partial derivatives from Step 1, we obtain

$$\begin{aligned}
\partial_x F_h(x,z) &= \int_0^1 f(u)\,K(u)\,du,\\
\partial_z F_h(x,z) &= \int_0^1 h(u)\,K(u)\,du,\\
\partial_x F_l(x,z) &= a\int_0^1 f(u)\,w(u)\,K(u)\,du,\\
\partial_z F_l(x,z) &= a\int_0^1 h(u)\,w(u)\,K(u)\,du.
\end{aligned}$$

Hence

$$\frac{\partial_x F_h(x,z)}{\partial_z F_h(x,z)} = \frac{\int_0^1 f(u)K(u)\,du}{\int_0^1 h(u)K(u)\,du}, \qquad \frac{\partial_x F_l(x,z)}{\partial_z F_l(x,z)} = \frac{\int_0^1 f(u)w(u)K(u)\,du}{\int_0^1 h(u)w(u)K(u)\,du}.$$

We now compare these two ratios. Define a probability measure μ on $[0,1]$ by

$$\mu(A) := \frac{\int_A K(u)\,du}{\int_0^1 K(v)\,dv} \qquad \text{for every } A\subseteq[0,1].$$

Next, reweight μ by the function h and define a probability measure μ_h by

$$\mu_h(A) := \frac{\int_A h \, d\mu}{\int_0^1 h \, d\mu}.$$

Also reweight the probability measure μ_h by the strictly decreasing function w, defining

$$\nu_h(A) := \frac{\int_A w \, d\mu_h}{\int w \, d\mu_h}.$$

Now, since $\psi(u)$ is bounded and strictly increasing on $[0,1]$ by Lemma B.6, and since w is strictly decreasing on $(0,1]$, Lemma B.3 yields

$$\int \psi \, d\nu_h < \int \psi \, d\mu_h.$$

Since $h(1) = 0$, both μ_h and ν_h assign zero mass to $\{1\}$. Therefore,

$$\int \psi \, d\mu_h = \int \frac{f}{h} \, d\mu_h = \frac{\int f \, d\mu}{\int h \, d\mu} = \frac{\int f(u)K(u) \, du}{\int h(u)K(u) \, du}.$$

Similarly,

$$\int \psi \, d\nu_h = \int \frac{f}{h} \, d\nu_h = \frac{\int \frac{f}{h} h \, w \, d\mu}{\int h \, w \, d\mu} = \frac{\int f(u)w(u)K(u) \, du}{\int h(u)w(u)K(u) \, du}.$$

Therefore

$$\frac{\int f(u)w(u)K(u) \, du}{\int h(u)w(u)K(u) \, du} < \frac{\int f(u)K(u) \, du}{\int h(u)K(u) \, du}.$$

Hence

$$\frac{\partial_x F_l(x,z)}{\partial_z F_l(x,z)} < \frac{\partial_x F_h(x,z)}{\partial_z F_h(x,z)}$$

and thus

$$-\frac{\partial_x F_h(x,z)}{\partial_z F_h(x,z)} < -\frac{\partial_x F_l(x,z)}{\partial_z F_l(x,z)}.$$

Let (x,z) be a point in $(0,1)^2$ with $F_h(x,z) = 0$. Since F_h is continuously differentiable on $(0,1)^2$ and $\partial_z F_h(x,z) > 0$ by Step 1, the implicit function theorem implies that in a neighborhood of x the zero set of F_h can be written as a continuously differentiable function $z_h(\cdot)$ satisfying

$$F_h(\tilde{x}, z_h(\tilde{x})) = 0$$

for all $\tilde{x}$ in some neighborhood of x, with derivative

$$z_h'(x) = -\frac{\partial_x F_h(x, z_h(x))}{\partial_z F_h(x, z_h(x))}.$$

Analogously, if $F_l(x, z) = 0$, then

$$z_l'(x) = -\frac{\partial_x F_l(x, z_l(x))}{\partial_z F_l(x, z_l(x))}.$$

Hence, whenever the two zero sets intersect at the same point (x, z), their slopes satisfy

$$z_h'(x) < z_l'(x).$$

Step 3: the two zero sets cannot intersect twice.

Assume, toward a contradiction, that there exist two distinct fully mixed equilibria

$$(x_1, z_1) \neq (x_2, z_2).$$

Without loss of generality, assume $x_1 < x_2$. Then, since F_h, F_l are strictly increasing in both its arguments, it must hold that $z_1 > z_2$.

Now fix any $x \in [x_1, x_2]$. Since F_h is strictly increasing in x and z, we have

$$F_h(x, z_2) \leq F_h(x_2, z_2) = 0, \qquad F_h(x, z_1) \geq F_h(x_1, z_1) = 0.$$

The inequalities are strict for $x \in (x_1, x_2)$. Because F_h is continuous and strictly increasing in z, for every $x \in [x_1, x_2]$ there exists a unique $z_h(x) \in [z_2, z_1]$ such that

$$F_h(x, z_h(x)) = 0.$$

Analogously, for every $x \in [x_1, x_2]$ there exists a unique $z_l(x) \in [z_2, z_1]$ such that

$$F_l(x, z_l(x)) = 0.$$

Moreover,

$$z_h(x_1) = z_l(x_1) = z_1, \qquad z_h(x_2) = z_l(x_2) = z_2.$$

By the implicit function theorem, both z_h and z_l are continuously differentiable in a neighborhood of every point at which they are defined. Moreover, if $z_h(x) = z_l(x)$, Step 2 implies

$$z_h'(x) < z_l'(x).$$

Define

$$H(x) := z_h(x) - z_l(x) \qquad \forall x \in [x_1, x_2].$$

Then H is continuous on $[x_1, x_2]$, and $H(x_1) = H(x_2) = 0$. Moreover, whenever $H(x) = 0$, we have

$$H'(x) = z_h'(x) - z_l'(x) < 0.$$

Since $H'(x_1) < 0$, there exists $\varepsilon > 0$ such that

$$H(x) < 0 \qquad \forall x \in (x_1, x_1 + \varepsilon).$$

Now define

$$x^* := \inf\{x \in (x_1, x_2] : H(x) = 0\}.$$

This set is nonempty because $H(x_2) = 0$, so x^* is well defined. By continuity of H, we have

$$H(x^*) = 0.$$

Since $H(x) < 0$ for x just to the right of x_1 and H is continuous, it follows that

$$H(x) < 0 \qquad \forall x \in (x_1, x^*).$$

But $H(x^*) = 0$ and thus $H'(x^*) < 0$. Hence there exists $\delta > 0$ such that

$$H(x) > 0 \qquad \forall x \in (x^* - \delta, x^*).$$

This contradicts $H(x) < 0$ for all $x \in (x_1, x^*)$. Therefore the system

$$F_h(x, z) = 0, \qquad F_l(x, z) = 0$$

has at most one solution in $(0, 1)^2$. Therefore there exists at most one equilibrium with

$$\lambda_h^* \in (0, 1) \qquad \text{and} \qquad \lambda_l^* \in (0, 1).$$

□

Lemma B.13 *Fix $\hat{x}_0 \in (0, 1)$ and suppose there exists a fully mixed equilibrium*

$$\lambda_h^*(\hat{x}_0) \in (0, 1), \qquad \lambda_l^*(\hat{x}_0) \in (0, 1).$$

Define

$$\bar{\lambda}_h := \lambda_h^*(\hat{x}_0), \qquad \bar{z} := \hat{x}_0\, \lambda_h^*(\hat{x}_0) + (1 - \hat{x}_0)\lambda_l^*(\hat{x}_0).$$

Then, for every belief $\hat{x} \in (0, 1)$ such that

$$\lambda_l^*(\hat{x}) := \frac{\bar{z} - \hat{x}\, \bar{\lambda}_h}{1 - \hat{x}} \in (0, 1),$$

the pair $\big(\bar{\lambda}_h, \lambda_l^*(\hat{x})\big)$ is the unique fully mixed equilibrium at belief $\hat{x}$. On this interval,

$$\frac{d}{d\hat{x}} \lambda_h^*(\hat{x}) = 0, \qquad \frac{d}{d\hat{x}} \lambda_l^*(\hat{x}) = \frac{\bar{z} - \bar{\lambda}_h}{(1 - \hat{x})^2} < 0.$$

PROOF. Define F_h, F_l as in (B.1) in the proof of Lemma B.12. We note that the system

$$F_h(x,z) = 0, \qquad F_l(x,z) = 0 \tag{B.2}$$

depends only on the induced applicant subjective masses (x,z). Consider a fully mixed equilibrium $\big(\lambda_h^*(\hat{x}_0), \lambda_l^*(\hat{x}_0)\big) \in (0,1)^2$. Define

$$\bar{\lambda}_h := \lambda_h^*(\hat{x}_0), \qquad \bar{z} := \hat{x}_0\,\bar{\lambda}_h + (1-\hat{x}_0)\lambda_l^*(\hat{x}_0).$$

Then $(\bar{\lambda}_h, \bar{z})$ solves (B.2). Now fix any $\hat{x} \in (0,1)$ such that

$$\lambda_l^*(\hat{x}) := \frac{\bar{z} - \hat{x}\,\bar{\lambda}_h}{1-\hat{x}} \in (0,1).$$

For this choice, $\hat{x}\,\bar{\lambda}_h + (1-\hat{x})\lambda_l^*(\hat{x}) = \bar{z}$. Hence, the subjective applicant masses coincide with those at belief $\hat{x}_0$. Hence, the pair $\big(\bar{\lambda}_h, \lambda_l^*(\hat{x})\big)$ is a fully mixed equilibrium at belief $\hat{x}$. By Lemma B.12, this fully mixed equilibrium is unique. Moreover,

$$\frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) = \frac{\bar{z} - \bar{\lambda}_h}{(1-\hat{x})^2} < 0$$

and since $\lambda_h^*(\hat{x}) = \bar{\lambda}_h$,

$$\frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) = 0.$$

□

Lemma B.14 *Fix $\hat{x}_0 \in (0,1)$ and suppose there exists a partially mixed equilibrium*

$$\lambda_h^*(\hat{x}_0) = 1, \qquad \lambda_l^*(\hat{x}_0) \in (0,1).$$

Define

$$\bar{z} := \hat{x}_0 + (1-\hat{x}_0)\lambda_l^*(\hat{x}_0).$$

Then for every belief $\hat{x} \in (0,1)$ such that

$$\lambda_l^*(\hat{x}) := \frac{\bar{z} - \hat{x}}{1-\hat{x}} \in (0,1),$$

the pair $(1, \lambda_l^*(\hat{x}))$ is the unique partially mixed equilibrium with $\lambda_h^*(\hat{x}) = 1$. Moreover,

$$\frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) = 0, \qquad \frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) = \frac{\bar{z} - 1}{(1-\hat{x})^2} < 0.$$

PROOF. Define

$$G(z) := w_l\, p_l(0, 1-z) - w_h\, p_l(1, z).$$

We note that

$$G(z) = 0 \tag{B.3}$$

does not depend on $\hat{x}$. Since also $\partial_y p_l(x, y) < 0$ everywhere, it follows that G is strictly increasing and thus any solution is unique. Let $\big(1, \lambda_l^*(\hat{x}_0)\big)$ be a partially mixed equilibrium at $\hat{x}_0$, and define

$$\bar{z} := \hat{x}_0 + (1 - \hat{x}_0)\lambda_l^*(\hat{x}_0).$$

Since low-ability job seekers mix at $\hat{x}_0$, they must be indifferent between the two markets. Hence $G(\bar{z}) = 0$. Now fix any $\hat{x} \in (0, 1)$ such that

$$\lambda_l^*(\hat{x}) := \frac{\bar{z} - \hat{x}}{1 - \hat{x}} \in (0, 1).$$

For this choice, $\hat{x} + (1 - \hat{x})\lambda_l^*(\hat{x}) = \bar{z}$. Hence the subjective applicant masses coincide with those at belief $\hat{x}_0$. At $\hat{x}_0$, high-ability job seekers weakly prefer the high-wage market and low-ability job seekers are indifferent. Therefore the same is true at belief $\hat{x}$, so $(1, \lambda_l^*(\hat{x}))$ is a partially mixed equilibrium. Also, the pair $\big(1, \lambda_l^*(\hat{x})\big)$ is the unique partially mixed equilibrium with $\lambda_h^*(\hat{x}) = 1$. Finally,

$$\frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) = \frac{\bar{z} - 1}{(1 - \hat{x})^2} < 0, \qquad \frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) = 0.$$

□

Lemma B.15 *For every $\hat{x} \in (0, 1)$, there exists at most one equilibrium with $\lambda_l^*(\hat{x}) = 0$. Moreover:*

(a) *On any interval of beliefs on which equilibria with $\lambda_l^*(\hat{x}) = 0$ and $\lambda_h^*(\hat{x}) \in (0, 1)$ exist,*

$$\frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) = 0, \qquad \frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) < 0.$$

(b) *On any open interval of beliefs on which equilibria with $\lambda_l^*(\hat{x}) = 0$ and $\lambda_h^*(\hat{x}) = 1$ exist,*

$$\frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) = 0, \qquad \frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) = 0.$$

PROOF. Fix $\hat{x} \in (0, 1)$ and suppose $\lambda_l^*(\hat{x}) = 0$. Then the subjective applicant masses in the high-wage market are $(x_H, y_H) = (\lambda_h, \hat{x}\,\lambda_h)$ while those in the low-wage market are $(x_L, y_L) = (1 - \lambda_h, 1 - \hat{x}\,\lambda_h)$.

Define the high type payoff difference between the high and low-wage markets by

$$U_h(\lambda_h, \hat{x}) := w_h\, p_h(\lambda_h, \hat{x}\lambda_h) - w_l\, p_h(1 - \lambda_h, 1 - \hat{x}\lambda_h). \tag{B.4}$$

An equilibrium with $\lambda_l^*(\hat{x}) = 0$ and $\lambda_h^*(\hat{x}) \in (0,1)$ therefore requires

$$U_h(\lambda_h^*(\hat{x}), \hat{x}) = 0,$$

while an equilibrium with $\lambda_l^*(\hat{x}) = 0$ and $\lambda_h^*(\hat{x}) = 1$ requires

$$U_h(1, \hat{x}) \geq 0.$$

We first show uniqueness. Since $\partial_x p_h(x,y) < 0$ and $\partial_y p_h(x,y) < 0$,

$$\partial_{\lambda_h} U_h(\lambda_h, \hat{x}) < 0$$

for every $(\lambda_h, \hat{x}) \in [0,1] \times (0,1)$. Therefore, there exists at most one interior value $\lambda_h \in (0,1)$ satisfying $U_h(\lambda_h, \hat{x}) = 0$. Moreover, if such an interior value exists, strict monotonicity implies $U_h(1, \hat{x}) < 0$ so $\lambda_h = 1$ cannot simultaneously be an equilibrium. Finally, $\lambda_h = 0$ together with $\lambda_l = 0$ is impossible by Lemma B.11. Hence there exists at most one equilibrium with $\lambda_l^*(\hat{x}) = 0$.

Next, consider any interval of beliefs on which equilibria with

$$\lambda_l^*(\hat{x}) = 0, \qquad \lambda_h^*(\hat{x}) \in (0,1)$$

exist. Trivially,

$$\frac{d}{d\hat{x}} \lambda_l^*(\hat{x}) = 0.$$

Also, whenever $\lambda_h > 0$, differentiation of (B.4) with respect to $\hat{x}$ gives

$$\partial_{\hat{x}} U_h(\lambda_h, \hat{x}) = w_h \, \lambda_h \, \partial_y p_h(\lambda_h, \hat{x}\lambda_h) + w_l \, \lambda_h \, \partial_y p_h(1 - \lambda_h, 1 - \hat{x}\lambda_h) < 0,$$

since $\partial_y p_h < 0$ everywhere.

Now fix any point $\hat{x}_0$ in such an interval. Since $\lambda_h^*(\hat{x}_0) \in (0,1)$, high-ability job seekers mix at $\hat{x}_0$ and are therefore indifferent. Hence, $U_h(\lambda_h^*(\hat{x}_0), \hat{x}_0) = 0$. Because also

$$(\lambda_h^*(\hat{x}_0), \hat{x}_0) \in (0,1) \times (0,1)$$

and

$$\partial_{\lambda_h} U_h(\lambda_h^*(\hat{x}_0), \hat{x}_0) < 0,$$

the implicit function theorem applies. Hence, there exists an open interval I containing $\hat{x}_0$ and a differentiable function $\tilde{\lambda}_h : I \to (0,1)$ such that

$$U_h(\tilde{\lambda}_h(\hat{x}), \hat{x}) = 0 \qquad \forall \hat{x} \in I.$$

By the uniqueness result proved above, $\tilde{\lambda}_h(\hat{x})$ must coincide with the equilibrium value $\lambda_h^*(\hat{x})$ for every $\hat{x} \in I$. Therefore

$$\frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) = -\frac{\partial_{\hat{x}} U_h(\lambda_h^*(\hat{x}), \hat{x})}{\partial_{\lambda_h} U_h(\lambda_h^*(\hat{x}), \hat{x})}.$$

Since both numerator and denominator are strictly negative, it follows that

$$\frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) < 0.$$

Because $\hat{x}_0$ was arbitrary, we conclude that on any interval of beliefs on which equilibria with $\lambda_l^*(\hat{x}) = 0$ and $\lambda_h^*(\hat{x}) \in (0,1)$ exist, one has

$$\frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) = 0, \qquad \frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) < 0.$$

Finally, consider any open interval of beliefs on which equilibria with $\lambda_l^*(\hat{x}) = 0$ and $\lambda_h^*(\hat{x}) = 1$ exist. On such an interval both equilibrium components are constant, hence

$$\frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) = 0, \qquad \frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) = 0.$$

□

Definition B.1 *We assign any equilibrium with $D > 0$ to one of the following classes:*

- Class (i): $\lambda_l^* \in (0,1)$ *and* $\lambda_h^* \in (0,1)$,
- Class (ii): $\lambda_l^* \in (0,1)$ *and* $\lambda_h^* = 1$,
- Class (iii): $\lambda_l^* = 0$ *and* $\lambda_h^* \in (0,1]$.

Lemma B.16 *Fix $\hat{x}_0 \in (0,1)$ and an equilibrium $(\lambda_h^0, \lambda_l^0)$ at $\hat{x}_0$ with $D > 0$.*

(a) *If $\lambda_l^0 > 0$, then there exists an equilibrium branch through $(\hat{x}_0, \lambda_h^0, \lambda_l^0)$, which is unique up to restriction of its domain. Moreover, there exists a unique $\hat{x}_1 \in (\hat{x}_0, 1)$ such that the branch is defined at least on $[\hat{x}_0, \hat{x}_1)$, belongs to Class* (i) *or Class* (ii) *for all $\hat{x} \in [\hat{x}_0, \hat{x}_1)$, and to Class* (iii) *if it continues beyond.*

(b) *If $\lambda_l^0 = 0$, then any equilibrium branch through $(\hat{x}_0, \lambda_h^0, \lambda_l^0)$ belongs to Class* (iii) *on its entire domain. Moreover, there is at most one such branch, up to restriction of its domain.*

(c) *Along any domain of a branch that belongs to Class* (iii)*, as $\hat{x}$ increases, switching between $\lambda_h^*(\hat{x}) \in (0,1)$ and $\lambda_h^*(\hat{x}) = 1$ can only occur from the boundary to the interior.*

PROOF. We consider the cases in which the initial equilibrium at $\hat{x}_0$ belongs to Class (i) and Class (iii). The argument for an initial Class (ii) equilibrium works analogously to the one for Class (i).

Assume the equilibrium at $\hat{x}_0$ belongs to Class (i). Then, by Lemma B.13,

$$\lambda_h^*(\hat{x}) = \bar{\lambda}_h \in (0,1), \qquad \lambda_l^*(\hat{x}) = \frac{\bar{z} - \hat{x}\,\bar{\lambda}_h}{1-\hat{x}}, \tag{B.5}$$

constitute the unique Class (i) equilibria for every $\hat{x} \in [\hat{x}_0, \hat{x}_1)$, where $\lambda_l^*(\hat{x}) \in (0,1)$. Hence,

$$\hat{x}_1 := \frac{\bar{z}}{\bar{\lambda}_h} \in (\hat{x}_0, 1).$$

Since both $\lambda_h^*(\hat{x})$, and $\lambda_l^*(\hat{x})$ are continuous, (B.5) constitutes an equilibrium branch on this interval.

We now show that any branch must remain in Class (i) on $[\hat{x}_0, \hat{x}_1)$ so that (B.5) constitutes the unique branch on this interval. Suppose, toward a contradiction, that this is false, and define

$$\tilde{x} := \inf\big\{\hat{x} \in [\hat{x}_0, \hat{x}_1) : \text{the branch does not belong to Class (i) at } \hat{x}\big\}.$$

Then $\tilde{x} \in (\hat{x}_0, \hat{x}_1)$. By definition of $\tilde{x}$, the branch belongs to Class (i) on $[\hat{x}_0, \tilde{x})$, so

$$\lambda_h^*(\hat{x}) = \bar{\lambda}_h \in (0,1), \qquad \lambda_l^*(\hat{x}) > 0 \qquad \forall \hat{x} \in [\hat{x}_0, \tilde{x}).$$

By continuity, this implies $\lambda_h^*(\tilde{x}) = \bar{\lambda}_h \in (0,1)$. If the branch belonged to Class (ii) at $\tilde{x}$, then $\lambda_h^*(\tilde{x}) = 1$, contradicting $\lambda_h^*(\tilde{x}) \in (0,1)$. If the branch belonged to Class (iii) at $\tilde{x}$, then $\lambda_l^*(\tilde{x}) = 0$. But by continuity it must hold that

$$\lambda_l^*(\tilde{x}) = \frac{\bar{z} - \tilde{x}\,\bar{\lambda}_h}{1-\tilde{x}},$$

which is strictly positive since $\tilde{x} < \hat{x}_1 = \bar{z}/\bar{\lambda}_h$, thus also yielding a contradiction. Therefore the branch cannot switch from Class (i) to Class (ii) or Class (iii) on $[\hat{x}_0, \hat{x}_1)$, proving uniqueness of the branch (B.5) on that interval.

Now suppose the branch is defined on $[\hat{x}_0, \bar{x})$ with $\bar{x} > \hat{x}_1$. By continuity, $\lambda_l^*(\hat{x}_1) = 0$. Hence, the equilibrium at $\hat{x}_1$ belongs to Class (iii).

We now show that the branch must remain in Class (iii) on $[\hat{x}_1, \bar{x})$. Suppose, toward a contradiction, that this is false. Then, analogously to above, there is a first point $\tilde{x} \in [\hat{x}_1, \bar{x})$ such that

$$\lambda_l^*(\tilde{x}) = 0 \qquad \text{and} \qquad \lambda_l^*(\hat{x}) > 0 \;\text{ for all } \hat{x} > \tilde{x} \text{ sufficiently close to } \tilde{x}.$$

Hence, on some right neighborhood of $\tilde{x}$, the branch belongs either to Class (i) or to Class (ii). If it belongs to Class (i) on such a neighborhood, then for some constants $\bar{\lambda}_h \in (0,1)$ and $\bar{z} \in (0,1)$

$$\lambda_l^*(\hat{x}) = \frac{\bar{z} - \hat{x}\,\bar{\lambda}_h}{1-\hat{x}}$$

for all $\hat{x}$ in that neighborhood. Since $\lambda_l^*(\tilde{x}) = 0$, it follows by continuity that

$$0 = \frac{\bar{z} - \tilde{x}\,\bar{\lambda}_h}{1-\tilde{x}}.$$

Since the right hand side is strictly decreasing, for every $\hat{x} > \tilde{x}$, $\lambda_l^*(\hat{x}) < 0$, a contradiction. The analogous argument rules out that the branch belongs to Class (ii) on some right neighborhood of $\tilde{x}$. Therefore the branch cannot leave Class (iii) after $\hat{x}_1$. Uniqueness of the branch after $\hat{x}_1$ then follows from uniqueness of any Class (iii) equilibrium by Lemma B.15. By the same argument, if the equilibrium at $\hat{x}_0$ already belongs to Class (iii), the unique branch is in Class (iii) on all of $[\hat{x}_0, \bar{x})$.

Finally, along a Class (iii) branch, moving from $\lambda_h^*(\hat{x}) \in (0,1)$ to $\lambda_h^*(\hat{x}) = 1$ as $\hat{x}$ increases is impossible, because on every interval on which $\lambda_h^*(\hat{x}) \in (0,1)$, Lemma B.15 implies that $\lambda_h^*(\hat{x})$ is strictly decreasing. Hence, it cannot attain the value 1 at a later point. Thus, if switching occurs along a Class (iii) branch, it can only be from $\lambda_h^*(\hat{x}) = 1$ to $\lambda_h^*(\hat{x}) \in (0,1)$. □

B.2. Proofs

B.2.1. Proof of Lemma 2.2

For any candidate strategy profile $(\lambda_h, \lambda_l) \in [0,1]^2$, define the subjective applicant masses in the high-wage market by

$$x_H(\lambda_h, \lambda_l) := \lambda_h, \qquad y_H(\lambda_h, \lambda_l) := \hat{x}\,\lambda_h + (1-\hat{x})\lambda_l,$$

and in the low-wage market by

$$x_L(\lambda_h, \lambda_l) := 1 - \lambda_h, \qquad y_L(\lambda_h, \lambda_l) := 1 - y_H(\lambda_h, \lambda_l).$$

These maps are continuous on $[0,1]^2$.

Because we consider a continuum economy, a unilateral deviation by a single job seeker does not affect these applicant masses. Therefore, each job seeker takes the induced applicant masses as fixed.

By the proof of Lemma B.2, the functions p_h and p_l are continuous on $\mathbb{R}^2_+$. Hence the payoff difference functions of a job seeker of type α,

$$U_\alpha(\lambda_h, \lambda_l) := w_h\, p_\alpha\big(x_H(\lambda_h, \lambda_l), y_H(\lambda_h, \lambda_l)\big) - w_l\, p_\alpha\big(x_L(\lambda_h, \lambda_l), y_L(\lambda_h, \lambda_l)\big),$$

are continuous on $[0,1]^2$.

Now define the best response correspondence of a job seeker of type α by

$$\Gamma_\alpha(\lambda_h,\lambda_l) := \begin{cases} \{1\}, & \text{if } U_\alpha(\lambda_h,\lambda_l) > 0 \\ [0,1], & \text{if } U_\alpha(\lambda_h,\lambda_l) = 0 \\ \{0\}, & \text{if } U_\alpha(\lambda_h,\lambda_l) < 0 \end{cases}$$

Thus $\Gamma_\alpha(\lambda_h,\lambda_l)$ is the set of optimal probabilities with which type α applies to the high-wage market given (λ_h,λ_l).

For every $(\lambda_h,\lambda_l) \in [0,1]^2$, the sets $\Gamma_h(\lambda_h,\lambda_l)$ and $\Gamma_l(\lambda_h,\lambda_l)$ are nonempty, compact, and convex subsets of $[0,1]$, since each of them is either $\{0\}$, $\{1\}$, or $[0,1]$. Define the joint best response correspondence

$$\Gamma(\lambda_h,\lambda_l) := \Gamma_h(\lambda_h,\lambda_l) \times \Gamma_l(\lambda_h,\lambda_l), \qquad (\lambda_h,\lambda_l) \in [0,1]^2.$$

Then $\Gamma(\lambda_h,\lambda_l) \subseteq [0,1]^2$ is nonempty, compact, and convex for every (λ_h,λ_l). We next show that Γ is an upper hemicontinuous correspondence. Let

$$(\lambda_h^n,\lambda_l^n) \to (\lambda_h,\lambda_l), \qquad (m_h^n,m_l^n) \to (m_h,m_l),$$

with $(m_h^n,m_l^n) \in \Gamma(\lambda_h^n,\lambda_l^n)$ for every n. We must show that $(m_h,m_l) \in \Gamma(\lambda_h,\lambda_l)$.

For high-ability job seekers, if $U_h(\lambda_h,\lambda_l) > 0$, continuity of U_h implies that $U_h(\lambda_h^n,\lambda_l^n) > 0$ for all sufficiently large n. Hence $m_h^n = 1$ for all sufficiently large n. Therefore (m_h^n) eventually becomes a sequence of constants, so $m_h = \lim_{n\to\infty} m_h^n = 1$. Since $U_h(\lambda_h,\lambda_l) > 0$, we have $\Gamma_h(\lambda_h,\lambda_l) = \{1\}$, and thus $m_h \in \Gamma_h(\lambda_h,\lambda_l)$.

If $U_h(\lambda_h,\lambda_l) < 0$, the same argument yields $U_h(\lambda_h^n,\lambda_l^n) < 0$ for all sufficiently large n, hence $m_h^n = 0$ for all sufficiently large n, and therefore $m_h = \lim_{n\to\infty} m_h^n = 0 \in \Gamma_h(\lambda_h,\lambda_l)$.

If $U_h(\lambda_h,\lambda_l) = 0$, then by definition $\Gamma_h(\lambda_h,\lambda_l) = [0,1]$, so $m_h \in \Gamma_h(\lambda_h,\lambda_l)$. Thus in all cases,

$$m_h \in \Gamma_h(\lambda_h,\lambda_l).$$

By an analogous argument, it follows for low-ability job seekers that

$$m_l \in \Gamma_l(\lambda_h,\lambda_l).$$

Hence,

$$(m_h,m_l) \in \Gamma_h(\lambda_h,\lambda_l) \times \Gamma_l(\lambda_h,\lambda_l) = \Gamma(\lambda_h,\lambda_l).$$

Therefore Γ is an upper hemicontinuous correspondence from the nonempty, compact, convex set $[0,1]^2$ into itself. By Kakutani's fixed point theorem, therefore there is $(\lambda_h^*, \lambda_l^*) \in [0,1]^2$ such that

$$(\lambda_h^*, \lambda_l^*) \in \Gamma(\lambda_h^*, \lambda_l^*).$$

Therefore $(\lambda_h^*, \lambda_l^*)$ is a symmetric equilibrium. □

B.2.2. Proof of Proposition 2.1

We use Lemma B.8, Lemma B.9, Lemma B.10, and Lemma B.11 to characterize all feasible equilibria.

First, assume that $\lambda_h^* = 1$ and $\lambda_l^* \in [0,1]$. Then, $D = \lambda_h^* - \lambda_l^* = 1 - \lambda_l^* \geq 0$, with equality if and only if $\lambda_l^* = 1$.

Second, assume that $\lambda_h^* \in (0,1)$. Then, by Lemma B.9, it cannot be that $\lambda_l^* = 1$. Thus either $\lambda_l^* = 0$ or $\lambda_l^* \in (0,1)$. In the former case, $D = \lambda_h^* - \lambda_l^* = \lambda_h^* > 0$. In the latter case, Lemma B.8 implies $\lambda_l^* < \lambda_h^*$, and therefore $D = \lambda_h^* - \lambda_l^* > 0$. So in either case $D > 0$.

Third, assume that $\lambda_h^* = 0$. By Lemma B.9, it cannot be that $\lambda_l^* = 1$. By Lemma B.10, it cannot be that $\lambda_l^* \in (0,1)$. By Lemma B.11, it cannot be that $\lambda_l^* = 0$. Hence there is no such equilibrium.

Therefore, the first two cases constitute all possible equilibria. Hence, in every equilibrium, $D \geq 0$ with equality if and only if $\lambda_h^* = \lambda_l^* = 1$. □

B.2.3. Proof of Proposition 2.2

We use the equilibrium classes as defined in Definition B.1.

By Lemma B.16(a), if $\lambda_l^*(\hat{x}_0) \in (0,1)$, there exists an equilibrium branch through $(\lambda_h^*(\hat{x}_0), \lambda_l^*(\hat{x}_0))$ that is unique up to restriction of its domain, and a unique threshold $\hat{x}^* > \hat{x}_0$ such that along this branch the equilibrium belongs to either Class (i) or Class (ii) for all $\hat{x} \in [\hat{x}_0, \hat{x}^*)$. Then, by Lemma B.13 and Lemma B.14,

$$\frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) < 0, \qquad \frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) = 0 \qquad \text{for all } \hat{x} \in [\hat{x}_0, \hat{x}^*).$$

Moreover, if the branch continues beyond $\hat{x}^*$, then by Lemma B.16(a) it belongs to Class (iii) on that continuation. Therefore, for every $\hat{x} \geq \hat{x}^*$ in the domain of the branch we have $\lambda_l^*(\hat{x}) = 0$, and Lemma B.15 together with Lemma B.16(c) implies that $\lambda_h^*(\hat{x})$ is weakly decreasing then, since it is strictly decreasing whenever $\lambda_h^*(\hat{x}) \in (0,1)$ and constant whenever $\lambda_h^*(\hat{x}) = 1$, and switching from the interior to the boundary cannot occur. The same holds if the initial equilibrium belongs to Class (iii), since then any unique branch belongs to Class (iii) on its entire domain by Lemma B.16(b). □

B.2.4. Proof of Corollary 2.1

Let x denote the share of overconfident job seekers among low-ability job seekers. The average probability of men to apply to high-wage vacancies equals $M_h = 0.5\,\lambda_h + xm\,\lambda_h + (0.5 - xm)\lambda_l$. The average probability of women to apply to high-wage vacancies equals $W_h = 0.5\,\lambda_h + x(1-m)\lambda_h + \big(0.5 - x(1-m)\big)\lambda_l$. Hence the gender application gap equals

$$G_h = M_h - W_h = x(2m-1)(\lambda_h - \lambda_l) = x(2m-1)D.$$

Since $D \geq 0$ and $x(2m-1) > 0$, we have $G_h \geq 0$ with equality if and only if $\lambda_l^* = \lambda_h^* = 1$. Moreover,

$$\frac{dG_h}{dD} = x(2m-1) > 0.$$

□

B.2.5. Proof of Corollary 2.2

Let x denote the share of overconfident job seekers among low-ability job seekers. Recall that the gender application gap equals $G_h(\hat{x}) = x(2m-1)(\lambda_h^*(\hat{x}) - \lambda_l^*(\hat{x}))$.

By Proposition 2.2, along the unique equilibrium branch,

$$\frac{d}{d\hat{x}}\lambda_h^*(\hat{x}) = 0, \qquad \frac{d}{d\hat{x}}\lambda_l^*(\hat{x}) < 0 \qquad \text{for every } \hat{x} \in [\hat{x}_0, \hat{x}^*).$$

Hence, $\lambda_h^*(\hat{x})$ is constant and $\lambda_l^*(\hat{x})$ is strictly decreasing, thus the gender gap is strictly increasing.

Next, consider $\hat{x} \geq \hat{x}^*$. By Proposition 2.2, if the equilibrium branch continues beyond $\hat{x}^*$, then $\lambda_l^*(\hat{x}) = 0$ and $\lambda_h^*(\hat{x})$ is weakly decreasing in $\hat{x}$. Hence, if the equilibrium branch continues beyond $\hat{x}^*$, the gender gap is weakly decreasing for $\hat{x} \geq \hat{x}^*$.

□

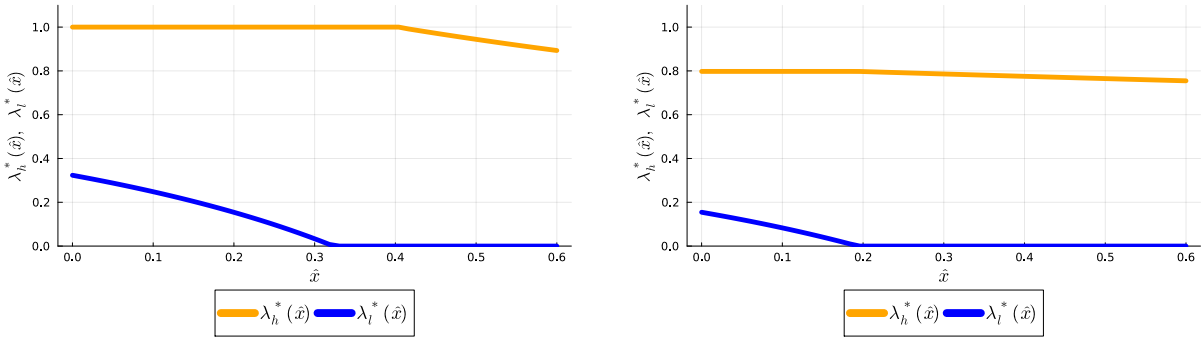
B.3. Model Appendix

B.3.1. Additional Figures

Figure B.1: Equilibrium Branches

(a) Initially Partially Pure Equilibrium

(b) Initially Fully Mixed Equilibrium



Note: The figure illustrates numerical results of the model and presents equilibrium branches for increasing levels of believed overconfidence, $\hat{x}$. The horizontal axis reports $\hat{x}$. The orange and blue lines show the equilibrium probability of applying to the high-wage vacancy for job seekers who perceive themselves as high ability, $\lambda_h^*(\hat{x})$, and for unbiased low-ability job seekers, $\lambda_l^*(\hat{x})$, respectively. Panel (a) uses $(\alpha_l/\alpha_h, w_l/w_h, \beta) = (0.7, 0.7, 0.8)$ and shows a case in which the branch starts from a partially pure strategy equilibrium, $\lambda_h^* = 1, \lambda_l^* \in (0,1)$. Panel (b) uses $(\alpha_l/\alpha_h, w_l/w_h, \beta) = (0.25, 0.8, 0.3)$ and shows a case in which the branch starts from a fully mixed strategy equilibrium, $\lambda_h^*, \lambda_l^* \in (0,1)$. In both panels, increases in $\hat{x}$ are first absorbed by reductions in $\lambda_l^*(\hat{x})$, while $\lambda_h^*(\hat{x})$ remains unchanged. Once unbiased low-ability job seekers have fully exited the high-wage market, further increases in $\hat{x}$ can no longer be offset through reductions in $\lambda_l^*(\hat{x})$ and may also reduce $\lambda_h^*(\hat{x})$.

B.4. Empirical Appendix

B.4.1. Discussion on Using ORIV

We initially considered applying the Obviously Related Instrumental Variables (ORIV) correction of Gillen et al. (2019) to address measurement error in the elicited risk-preference measure when estimating the gender application gap in CONTROL. ORIV is applicable in settings in which researchers observe duplicate elicitations of the same latent variable. The duplicate measures can then be used as instruments for one another if the measurement errors are orthogonal to each other and unrelated to the error term in the outcome equation. We do not view this latter requirement, i.e., the exclusion restriction in classical instrumental variable terminology, as credibly justified in our setting.

We observe the incentivized MPL-based estimate of the CRRA coefficient, $\hat{\rho}$, and a qualitative self-assessment of general willingness to take risks, following Dohmen et al. (2011). While both measures plausibly contain a common risk-preference component, the qualitative measure is significantly correlated with confidence measures, as shown in Table B.1. The first is overplacement, our overconfidence measure used throughout the paper, defined as the difference between the stated probability of being the better performing job seeker and the actual probability of being the better performing job seeker. The second is the supplementary measure of overestimation, defined as the difference between the believed and actual number of correctly solved Raven matrices. By contrast, the incentivized MPL-based measure is only weakly related to these confidence measures, if at all. This pattern suggests that the qualitative risk measure captures not only risk preferences, but also confidence-related variation. Since confidence is itself the central determinant of high-wage applications in our framework, the exclusion restriction required for using the qualitative risk measure as an instrument for $\hat{\rho}$ is difficult to justify.

A plausible explanation of the correlation between the qualitative risk question and the confidence measures is that respondents may infer their willingness to take risks from their willingness to enter risky situations. A respondent may report low willingness to take risks because she avoids such situations, even though this avoidance may reflect low confidence rather than risk aversion: if success probabilities in risky situations depend on performance and she expects to perform poorly, she also expects risky situations to have unfavorable outcomes. Hence, she avoids such situations. In that case, the qualitative risk measure captures both risk preferences and confidence-related beliefs.

Table B.1: Pearson Correlations Between Risk Measures and Overconfidence

	Overplacement	Overestimation
Incentivized CRRA, $\hat{\rho}$	-0.055	-0.089^{*}
Self-reported risk preference	-0.119^{**}	-0.128^{***}

Note: This table reports the Pearson correlation coefficients across $N = 421$ participants. The general risk question is reverse-coded so that higher values correspond to greater risk aversion, aligning the sign with the MPL-based $\hat{\rho}$. Overconfidence is constructed as the difference between the believed and actual number of correct Raven matrices. $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

B.4.2. Additional Figures and Tables

Table B.2: Balance Table

	Treatments			Kruskal-Wallis
	Control	Info	Other-Info	p-value
CRRA coefficient	0.47	0.50	0.59	0.12
	(0.05)	(0.05)	(0.05)	
Female	0.51	0.56	0.54	0.72
	(0.04)	(0.04)	(0.04)	
Last Math Grade in School	2.16	2.12	2.25	0.23
	(0.09)	(0.09)	(0.08)	
No. Raven Correct	7.92	7.88	7.44	0.08
	(0.19)	(0.18)	(0.18)	
Actual Percentile	52.01	52.61	46.37	0.16
	(2.62)	(2.57)	(2.50)	
Belief High Rank	61.07	63.24	60.17	0.67
	(1.51)	(1.49)	(1.55)	
Believed Overconfidence	7.52	8.32	9.60	0.37
	(1.07)	(1.23)	(1.09)	
Age				
<21	31.21%	31.88%	26.06%	0.50
21–24	50.35%	44.20%	40.85%	0.27
25–29	12.77%	15.94%	19.72%	0.28
30+	5.67%	7.97%	13.38%	0.07
N	141	138	142	

Note: The table reports descriptive statistics and p-values of a Kruskal-Wallis test of equality of all our three treatment groups.

Table B.3: Gender Application Gaps in CONTROL

	Dependent variable: Application Weight							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female	-9.626**	-7.753*	-7.688*	-5.313	-7.945*	-5.664	-8.000*	-6.046
	(4.565)	(4.554)	(4.631)	(4.329)	(4.607)	(4.397)	(4.649)	(4.382)
Actual Percentile		0.175**	0.198**	0.539***	0.178**	0.584***	0.205**	0.597***
		(0.069)	(0.080)	(0.129)	(0.070)	(0.130)	(0.081)	(0.130)
Overconfidence				0.474***		0.487***		0.492***
				(0.144)		(0.141)		(0.141)
CRRA Risk					1.303		2.090	2.679
					(4.206)		(4.338)	(3.833)
Constant	75.348***	65.280***	64.028***	40.832***	64.601***	33.911***	62.748***	32.002**
	(3.256)	(5.515)	(10.376)	(9.103)	(6.038)	(12.255)	(10.846)	(12.520)
Additional Controls			✓			✓	✓	✓
Wald-test: Female = 0	0.035	0.089	0.097	0.220	0.085	0.198	0.085	0.168
Observations	141	141	141	141	141	141	141	141
R^2	0.031	0.069	0.099	0.146	0.070	0.176	0.101	0.180

Note: The table reports OLS regression results using participants of CONTROL. The dependent variable is the application weight allocated to the high-wage vacancy. Column (1)-(3) depict the baseline gender differences without accounting for risk and overconfidence. Additional control variables include age, lab fixed effects and last high school math grade. Column (5) and Column (7) reports the addition of the CRRA risk parameter. Column (4) and Column (6) include overconfidence. Column (8) includes both. Hence, the comparison between Column (2) and (4), Column (3) and Column (6), and Column (7) and Column (8) respectively inform us on the extent to which the gender coefficient shrinks when accounting for overconfidence. The measure of overconfidence is defined in Equation 2.5. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table B.4: Effects of (Believed) Overconfidence with Controls

	Dependent variable: Application Weight			
	(1)	(2)	(3)	(4)
INFO	-4.605	-2.147		
	(3.310)	(3.234)		
Overconfidence		0.569***		
		(0.118)		
Actual Percentile	0.355***	0.673***	0.156***	0.131**
	(0.058)	(0.097)	(0.054)	(0.054)
INFO × Overconfidence		-0.334***		
		(0.111)		
OTHER-INFO			-3.656	-6.238
			(3.502)	(4.506)
Believed Overconfidence				0.164
				(0.224)
OTHER-INFO × Believed Overconfidence				0.211
				(0.300)
Constant	52.664***	28.712***	69.555***	68.230***
	(7.714)	(9.380)	(7.015)	(7.096)
Additional Controls	✓	✓	✓	✓
Observations	279	279	283	283
R^2	0.159	0.230	0.057	0.071

Note: The table reports OLS regression results. The dependent variable is the application weight allocated to the high-wage vacancy. INFO is a dummy variable equal to one for participants who receive information about their own actual relative ability. OTHER-INFO is a dummy variable equal to one for participants whose competitor receives information about the competitor's own actual relative ability. Columns (1) and (2) compare INFO participants to participants in CONTROL. Columns (3) and (4) compare OTHER-INFO participants to participants in CONTROL. The measure of overconfidence is defined in Equation 2.5. The measure of believed overconfidence is defined in Equation 2.6. The set of controls includes age, lab fixed effects, last high school math grade and the CRRA parameter. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

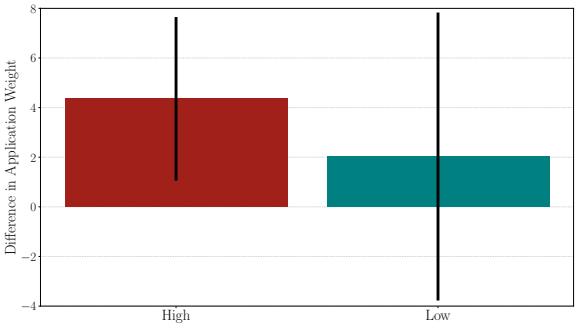
Table B.5: Effects of (Believed) Overconfidence Excluding $q_i \notin [0.1, 0.9]$

	Dependent variable: Application Weight			
	(1)	(2)	(3)	(4)
INFO	-6.300	-2.684		
	(3.930)	(4.430)		
Overconfidence		0.220**		
		(0.098)		
INFO × Overconfidence		-0.332**		
		(0.163)		
OTHER-INFO			-5.373	-11.605**
			(4.029)	(5.462)
Believed Overconfidence				-0.056
				(0.269)
OTHER-INFO ×				
Believed Overconfidence				0.629*
				(0.375)
Constant	70.538***	68.282***	70.538***	71.008***
	(2.654)	(3.058)	(2.654)	(3.551)
Observations	213	213	213	213
R^2	0.012	0.033	0.008	0.038

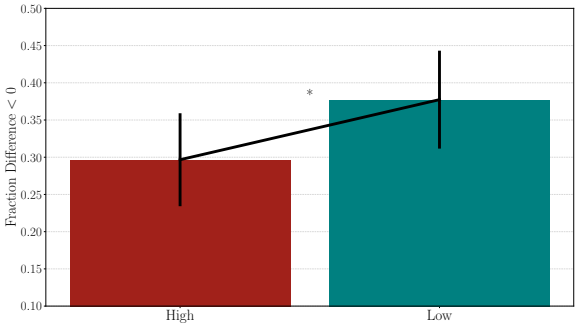
Note: The table reports OLS regression results. The dependent variable is the application weight allocated to the high-wage vacancy. The sample excludes participants with actual relative ability below the 10th percentile or above the 90th percentile. INFO is a dummy variable equal to one for participants who receive information about their own actual relative ability. OTHER-INFO is a dummy variable equal to one for participants whose competitor receives information about the competitor's own actual relative ability. Columns (1) and (2) compare INFO participants to participants in CONTROL. Columns (3) and (4) compare OTHER-INFO participants to participants in CONTROL. The measure of overconfidence is defined in Equation 2.5. The measure of believed overconfidence is defined in Equation 2.6. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Figure B.2: Vignettes

(a) Average of: Overconfidence - Unbiased



(b) Fraction Difference < 0



Note: The figure summarizes the results of the vignette study. Participants are split by whether they have higher ability than their competitor, *High*, or lower ability than their competitor, *Low*. Panel (a) shows the average within-participant difference in application weight allocated to the high-wage vacancy between the overconfident-competitor vignette and the unbiased-competitor vignette. Panel (b) shows the fraction of participants for whom this difference is negative. Error bars represent 95% confidence intervals based on robust standard errors. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.6: Anchoring Test

	Difference $= 0$	$\lvert$Difference$\rvert \leq 1$	$\lvert$Difference$\rvert \leq 2$	$\lvert$Difference$\rvert \leq 5$
	(1)	(2)	(3)	(4)
Info	0.002	0.031	0.060	0.025
	(0.032)	(0.035)	(0.037)	(0.046)
Constant	0.078***	0.078***	0.078***	0.163***
	(0.023)	(0.023)	(0.023)	(0.031)
Wald-test: Info $= 0$	0.958	0.380	0.108	0.580
Observations	279	279	279	279
R^2	0.000	0.003	0.009	0.001

Note: The table reports OLS regressions using participants from Control and Info. The dependent variables are dummy variables for choosing an application allocation equal or close to the potential anchor. Info is a dummy variable equal to one for participants who receive information about their own actual relative ability. The outcome equals one if the absolute difference between the high-wage application points and the relevant anchor lies within the window indicated in the column header. For participants in Info, the anchor is the provided information about actual relative ability. For participants in Control, the anchor is the participant's belief about own relative ability. * p < 0.1, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses.

B.5. Experiment Instructions

INTRODUCTION

Welcome to today's experiment!

This experiment consists of **three parts** and a final questionnaire. At the end of the experiment, you will certainly receive your participation fee of **€5.00**. In addition, you can earn money through your decisions during the experiment. Part 1 will be paid out **for sure**, and at the end of the experiment, **the computer will randomly select one of the remaining two parts (Part 2 or Part 3)**, and you will also be paid for that part.

Next

Information about Part 1

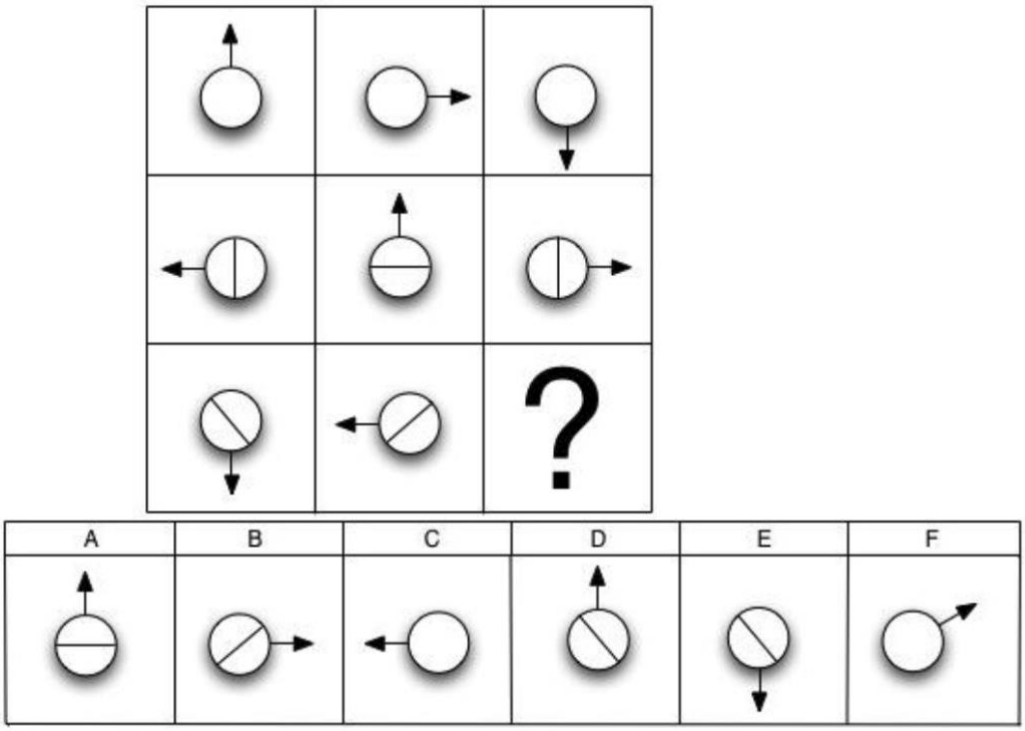
The first part of the experiment consists of **11 logic tasks**. On the following pages, you will see a figure with symbols at the top of the screen, in which one symbol is missing. Your task is to complete the empty field in such a way that a logical sequence of symbols is formed. To do so, you can choose from several symbols labeled Option A–F. You have **15 minutes** to complete Part 1 of the study. When you click "Next", your answer will be saved and you can proceed to the next task. You cannot go back and change your answer. Once you have completed all tasks, you can click "Finish". This will stop the measurement of your working time. However, the 15 minutes will run in full regardless of your working time. Even if you complete all tasks before the time has elapsed, you will still need to wait until the 15 minutes are over before continuing with the experiment.

Please note that your performance in Part 1 affects your payment in Part 2. Primarily, it is important that you solve as many tasks correctly as possible. Second, the time you take until you click "Finish" will be measured. If you do not complete all tasks and therefore do not click "Finish", your working time will be recorded as 15 minutes.

Next

RAVEN MATRICES

Task 1 of 11



Part 1

Before we begin with Part 2, we would like to ask you a few additional questions. Please answer the following questions as accurately as possible.

What do you think—how many of the 11 logic tasks did you answer correctly?

Next

Information

For Part 2 of the study, you were randomly assigned to a group of two together with another participant from the current session —that is, from all participants who are currently in the laboratory. Your expected payment in Part 2 depends on how well your performance in the logic task (Part 1) compares to that of the other member of your two-person group.

Your performance relative to the other member of your two-person group is evaluated as follows:

- Your performance is rated as **better than the performance of the other member of your two-person group** if you solved more logic tasks correctly than the other group member.
- If you solved the same number of logic tasks correctly as the other group member, your performance is rated as **better** if you required **less total time** than the other group member.

Based on your relative performance, you will be classified as **high** or **low**:

- **High**: If you performed **better** in the logic task than the other group member.
- **Low**: If you performed **worse** in the logic task than the other group member.

Next

BELIEF ELICITATION

Part 1: Self-assessment

Before Part 2 begins, we would like to know how you assess your own performance in the logic task from Part 1 compared to the other member of your two-person group. Specifically, we ask how likely you think it is that you were classified as **high**. **High** means that you performed **better than the other group member** in the logic task.

Please indicate your self-assessment by positioning the slider below. If you position the slider at (Low = 100%, High = 0%), this means that you believe you certainly performed worse than the other group member. If you position the slider at (Low = 0%, High = 100%), this means that you believe you certainly performed better than the other group member. For example, if you position the slider at (Low = 70%, High = 30%), this means that you believe there is a 70% probability that you performed worse, and a 30% probability that you performed better than the other group member.

The probability of being classified as **high** corresponds to the share of all participants in the current session who performed worse than you. If you think that you performed better than 70% of the other participants in the session, you should therefore position the slider at (Low = 30%, High = 70%).

The more accurate your self-assessment is relative to your actual performance, the higher your expected payment. You will be paid according to the following formula:

Your payment (in Euro) = $2 - 2((I(rank) - \frac{p}{100})^2$

where **I(rank)** is an indicator variable that takes the value 1 for your actual classification (**High** or **Low**) and **p** is your stated probability for that classification. Although this formula may appear complicated, the basic idea is very simple: **On average, your earnings are highest if you try to estimate your classification as accurately as possible.** Your maximum earnings from this estimate are 2 euros. Negative earnings are not possible.

Please provide your self-assessment by positioning the slider.



Low: % **High:** %

Next

Part 1: Assessment of others

We would also like to know what you think about how the **other participants in the current session**—that is, all other participants currently in the laboratory—**assessed themselves on average** in the previous question.

As in the previous question, the more accurate your answer is relative to the actual average self-assessment of the other participants in the session, the higher your payment will be. Specifically, your payment is calculated as follows:

Distance	< 2	< 4	< 6	< 8	< 10	< 12	< 14	< 16	< 18	< 20	≥ 20
Payment	€2.00	€1.90	€1.79	€1.67	€1.55	€1.41	€1.26	€1.10	€0.89	€0.63	€0.00

The distance is defined as the absolute difference between your answer and the actual average self-assessment of the other participants in your session. Your payment therefore decreases as the distance between your answer and the true average self-assessment increases. For example, if your answer deviates by 11 percentage points from the actual average, you will receive a payment of €1.41.

What do you think—how did the other participants in your session assess themselves **on average**?



Low: 39% **High:** 61%

Next

LABOR MARKET STAGE

Part 2: Application phase

Please read the following information carefully. You will have to answer comprehension questions correctly before you can proceed.

In this part, you apply for two jobs (Job A and Job B). The other person in your two-person group will also apply for these two jobs. You can distribute your application intensity across the two jobs. To do so, each of you receives **100 points that you can freely allocate between Job A and Job B**. The two jobs differ only in the wage that you receive if you are offered the job. The wages are as follows:

Job	**Wage**
Job A	€7.00
Job B	€12.00

Each of the two jobs sends exactly one offer. Your probability of receiving the offer from a given job depends on your application intensity for that job, the application intensity of the other group member, and your realized classification from the logic task (**High** or **Low**).

Please click "Next" to read the second part of the instructions. You can return to this part by clicking "Back".

Next

Part 2: Application phase

You and the other person in your two-person group apply simultaneously for **Job A** and **Job B**. Your probability of receiving an offer depends on your classification and on how many points you allocate to each job, relative to the other participant.

Your probability of receiving the offer from Job A is given by:

$$\text{Probability offer from Job A} = \frac{\text{Job-A-points}_i \times \text{R}_i}{(\text{Job-A-points}_i \times \text{R}_i) + (\text{Job-A-points}_j \times \text{R}_j)}$$

The probability of receiving the offer from **Job B** is calculated analogously. Note that if you allocate, for example, 100 points to Job A and 0 points to Job B, your probability of receiving Job B is **zero**, regardless of the other participant's allocation.

The factor by which your application intensity is multiplied depends on whether you were classified as High or Low:

- If you are classified as **High**, your points are weighted twice (**your factor = 2**), and the other participant's points are weighted once (**factor = 1**).
- If you are classified as **Low**, your points are weighted once (**your factor = 1**), and the other participant's points are weighted twice (**factor = 2**).

Reminder: Your classification was determined by your performance in the logic task relative to the other group member.

The basic idea behind the formula is simple:

- Your probability of receiving a job offer **increases** if you are classified as High.
- Your probability **increases** with your own application intensity.
- Your probability **decreases** with the other participant's application intensity.

Example

Suppose you are classified as **Low** and the other participant is classified as **High**. If you allocate 60 points to Job A and 40 points to Job B, and the other participant allocates 70 points to Job A and 30 points to Job B, the probabilities are:

- **Probability of receiving Job A (you):**

$$\frac{1 \times 60}{(1 \times 60) + (2 \times 70)} \times 100 = \frac{60}{60 + 140} \times 100 = \frac{60}{200} \times 100 = 30\%$$

- **Probability of receiving Job A (other participant):**

$$\frac{2 \times 70}{(1 \times 60) + (2 \times 70)} \times 100 = \frac{140}{60 + 140} \times 100 = \frac{140}{200} \times 100 = 70\%$$

Your payment in this part equals the wage of the job you receive. If you receive no offer, your payment is zero. If you receive offers from both jobs, you receive the higher wage (Job B).

Back

Next

Comprehension questions

To ensure that you have understood the instructions, please answer the following questions.

Click here to read the instructions again.

How does your probability of receiving an offer from Job A change if you increase your application intensity for Job A?

- Probability decreases
- Probability remains the same
- Probability increases

How does your probability of receiving an offer from Job B change if the other member of your two-person group increases their application intensity for Job B?

- Probability decreases
- Probability remains the same
- Probability increases

Suppose both members of a two-person group allocate the same application intensity to Job A. Which person has a higher probability of receiving an offer from Job A?

- Person classified as High
- Both are equally likely
- Person classified as Low

INFORMATION SCREENS

TREATMENT GROUP INFO SCREEN

As already explained, you and the other participants in your session were assigned to two-person groups. Based on your performance in the logic task relative to the other group member, you were classified as **high** or **low**.

Before the application phase begins, **one person in each two-person group** receives information about how they performed relative to **all other participants in the session**. Specifically, this person is informed about how many participants they performed better than and how many they performed worse than, as well as the resulting probability of being classified as **high**.

It is determined **at random** which person receives this information. If you receive the information, the other person does not. If you do not receive the information, the other person does.

Next

OTHER-INFO PLAYER SCREEN

Information

You do **not receive any information** about your performance in the logic task relative to the other participants in the session. However, the other person in your two-person group does receive information about how they performed relative to **all other participants in the session**.

Next

INFO PLAYER SCREEN

Information

You were randomly selected within your two-person group to receive information about your performance in the logic task relative to the other participants in your session. The other person in your two-person group does not receive any information.

You performed better than **7** of the **19** other participants in your session. This also means that you performed worse than **12** participants. Therefore, your probability of being classified as **high** in a two-person group is 37%.

Next

CONTROL GROUP INFO SCREEN

As already explained, you and the other participants in your session were assigned to two-person groups. Based on your performance in the logic task relative to the other group member, you were classified as **high** or **low**.

As explained earlier, this classification affects your probabilities of receiving a job offer. Please note that neither you nor the other member of your two-person group will receive any information about the performance in the logic task or the resulting classification.

Next

DECISION SCREEN

Allocate your 100 points between Job A and Job B using the slider.

Job	Wage
Job A	€7
Job B	€12



Job A: 28 Points

Job B: 72 Points

VIGNETTES

Information

Before you receive the results of the application phase, we ask you to consider **two stylized hypothetical scenarios**. For each scenario, please indicate **how you would have behaved in the experiment** if you had received the information from that scenario **before** making your application decision.

Note that the probability of being classified as **high** corresponds to the share of other participants whom you outperformed in the logic task. For example, if your performance was better than 70% of other participants, your probability of being classified as **high** is also 70%.

Next

Scenario 1

Imagine that before making your application decision, you receive the information that you performed better than **30% of all other participants** in the logic task. In addition, you receive the information that your competitor performed better than **50% of the other participants** and that he assesses his performance correctly. Your competitor therefore correctly believes that he performed better than **50% of the participants**. Your competitor receives no information about either your performance or their own performance in the logic task.

How would you have applied for the jobs in the experiment given this information? Allocate your 100 points between Job A and Job B using the slider.

Job	Wage
Job A	€7
Job B	€12



Job A: 33 Points Job B: 67 Points

Click here to read the instructions again.

Next

Scenario 2

Now imagine that, before making your application decision, you know that—just as in Scenario 1—you performed better than **30% of all other participants** in the logic task. In addition, you know that your competitor, as in Scenario 1, performed better than **50% of the participants**, but that **in contrast to Scenario 1**, he overestimates his performance. Your competitor incorrectly believes that he performed better than **80% of the participants**. As in Scenario 1, your competitor receives no information about either your performance or their own performance.

How would you have applied for the jobs in the experiment given this information? Allocate your 100 points between Job A and Job B using the slider.

Job	Wage
Job A	€7
Job B	€12



Job A: 38 Points Job B: 62 Points

Click here to read the instructions again.

Next

Results – Part 2

Of the jobs you applied for, you received **2 job offers**.

The highest-paying job you received pays **€12.00**. Therefore, your payment is **€12.00**.

Please note that, as mentioned earlier, you will only receive this payment if the computer randomly selects **Part 2** for payment at the end of the experiment.

Next

MULTIPLE PRICE LIST

Part 3

In Part 3 of the study, you will make **17 decisions**. Each decision consists of a choice between two options: A and B. If you choose Option A, you receive a fixed monetary amount for sure. If you choose Option B, your payment is determined by a random draw: **with a probability of 50% you receive €20, and with a probability of 50% you receive €0**. Once you have made all your decisions, the computer will randomly select one of the 17 decisions to determine your payment for Part 3.

To make the process easier for you, you can simply click on the decision at which you want to switch from **the certain payment (A)** to **the lottery (B)**. Your preferred options will then be highlighted in blue.

Next

Please choose either A or B in each decision. To make the process easier for you, you can simply click on the decision at which you want to switch from **the certain payment (A)** to **the lottery (B)**. Your preferred options will then be highlighted in blue.

	A		B
Decision 1:	12€ with certainty	OR	0€ with 50% probability und 20€ with 50% probability
Decision 2:	11.5€ with certainty	OR	0€ with 50% probability und 20€ with 50% probability
Decision 3:	11€ with certainty	OR	0€ with 50% probability und 20€ with 50% probability
Decision 4:	10.5€ with certainty	OR	0€ with 50% probability und 20€ with 50% probability
Decision 5:	10€ with certainty	OR	0€ with 50% probability und 20€ with 50% probability
Decision 6:	9.5€ with certainty	OR	0€ with 50% probability und 20€ with 50% probability
Decision 7:	9€ with certainty	OR	0€ with 50% probability und 20€ with 50% probability
Decision 8:	8.5€ with certainty	OR	0€ with 50% probability und

Part 3

Decision 6 was randomly selected for you.
In decision 6, you chose the lottery. The computer conducted the draw and **you receive €20**.

Please note that, as mentioned earlier, you will only receive this payment if the computer randomly selects **Part 3** for payment at the end of the experiment.

Next

APPENDIX C

Appendix to Chapter 3

C.1. Proofs of the Three-Period Model

This appendix proves the three-period model results stated in the main text. Each proof proceeds by defining and proving an auxiliary lemma that fully characterizes the corresponding period-specific behavior; the main-text result then immediately follows as the relevant implication of that full characterization.

C.1.1. Proof of Lemma 3.1.

Lemma C.1 *For the optimal period-2 strategy, $s^*_{2,\phi}$, it holds that*

(i) the time-consistent type searches if and only if

$$w_2 \leq w_h - \sqrt{\frac{2(w_h - w_l)c}{\delta}} \coloneqq w^*_{2,TC}$$

and accepts a job offer with wage w'' if and only if $w'' > w_2$.

(ii) The present-biased types $\phi \in \{S, N\}$ search if and only if

$$w_2 \leq w_h - \sqrt{\frac{2(w_h - w_l)c}{\beta\delta}} \coloneqq w^*_{2,\phi} < w^*_{2,TC}$$

and accept a job offer with wage w'' if and only if $w'' > w_2$.

Proof. Because period 3 is terminal, a worker who searches in period 2 and receives an offer w'' accepts the offer if and only if it exceeds the current wage, i.e., $w'' > w_2$.

We characterize the period-2 search decision for present-biased workers $\phi \in \{N, S\}$. The search rule for the time-consistent type follows immediately by setting $\beta = 1$.

Fix a period-2 wage $w_2 \in [w_l, w_h]$. Let $v^{\text{search}}_{2,\phi}(w_2)$ denote the period-2 payoff from searching and $v^{\text{no-search}}_{2,\phi}(w_2)$ the period-2 payoff from not searching. If the worker searches, she receives the current wage net of the search cost in period 2. In period 3, she earns either a higher wage if the offer she draws exceeds w_2, or otherwise keeps w_2. Hence,

$$v^{\text{search}}_{2,\phi}(w_2) = w_2 - c + \beta\delta\left((1 - F(w_2))\,\mathbb{E}[w'' \mid w'' > w_2] + F(w_2)w_2\right).$$

If the worker does not search, she certainly keeps wage w_2 in both periods 2 and 3, so

$$v_{2,\phi}^{\text{no-search}}(w_2) = w_2(1+\beta\delta).$$

The gain from search equals

$$H_{2,\phi}(w_2) \coloneqq v_{2,\phi}^{\text{search}}(w_2) - v_{2,\phi}^{\text{no-search}}(w_2) = \frac{\beta\delta}{2(w_h - w_l)}(w_h - w_2)^2 - c. \tag{C.1}$$

This function is strictly decreasing, continuous, and satisfies $H_{2,\phi}(w_h) = -c < 0$. Under the maintained parameter restriction, we also have $H_{2,\phi}(w_l) > 0$. Therefore, there exists a unique cutoff $w_{2,\phi}^* \in (w_l, w_h)$ such that $H_{2,\phi}(w_{2,\phi}^*) = 0$ and searching is optimal if and only if $w_2 \leq w_{2,\phi}^*$. Solving $H_{2,\phi}(w_{2,\phi}^*) = 0$ yields

$$w_{2,\phi}^* = w_h - \sqrt{\frac{2(w_h - w_l)c}{\beta\delta}} \quad \text{for} \quad \phi \in \{S, N\}.$$

Setting $\beta = 1$ gives $w_{2,TC}^*$ with $w_{2,TC}^* > w_{2,\phi}^*$ for $\phi \in \{S, N\}$. □

C.1.2. Proof of Proposition 3.1.

Lemma C.2 *Suppose the worker earns wage w_1 in period 1. The perception-perfect period-1 job offer acceptance strategies are as follows:*

(i) *The time-consistent and naive types accept an offer if and only if $w' > w_1$.*

(ii) *For the sophisticated type:*

(a) *If $w_1 < \tilde{w}$ or $w_1 \geq w_{2,S}^* + \psi(w_{2,S}^*)$, she accepts an offer if and only if $w' > w_1$.*

(b) *If $w_1 \in [\tilde{w},\, w_{2,S}^*]$, she accepts an offer if and only if $w' \in (w_1,\, w_{2,S}^*]$ or $w' > w_{2,S}^* + \psi(w_1)$.*

(c) *If $w_1 \in (w_{2,S}^*,\, w_{2,S}^* + \psi(w_{2,S}^*))$, define $w \in [w_l, w_{2,S}^*)$ such that $w_{2,S}^* + \psi(w) = w_1$ whenever such a wage exists. If such a wage exists, she accepts an offer if and only if $w' \in (w,\, w_{2,S}^*]$ or $w' > w_1$. Else, she accepts an offer if and only if $w' \in [w_l, w_{2,S}^*]$ or $w' > w_1$.*

Proof.

Time-consistent worker. Consider a wage $w \in [w_l, w_h]$ realized in period 2. By Lemma C.1, the worker searches if and only if $w \leq w_{2,TC}^*$. Hence, her period-1 continuation payoff equals

$$U_{TC}(w) = \begin{cases} \delta v_{2,TC}^{\text{search}}(w), & w \leq w_{2,TC}^*, \\ \delta v_{2,TC}^{\text{no-search}}(w), & w > w_{2,TC}^*. \end{cases}$$

Since $v_{2,TC}^{\text{search}}(w)$ and $v_{2,TC}^{\text{no-search}}(w)$ are strictly increasing and continuous in w, and they are equal at $w_{2,TC}^*$, it follows that U_{TC} is strictly increasing and continuous on $[w_l, w_h]$. Therefore, it is optimal to accept any offer with $w' > w_1$ and reject all other offers.

Naive worker. From the perspective of period 1, the naive worker evaluates future payoffs as β times the time-consistent continuation payoff. Hence, her perceived continuation payoff is also strictly increasing in w, and she follows the same acceptance rule as the time-consistent worker.

Sophisticated worker. The sophisticated worker anticipates that her period-2 self searches if and only if $w \leq w_{2,S}^*$. From the perspective of period 1, payoffs in periods 2 and 3 are discounted exponentially relative to each other. Hence, her continuation payoff is

$$U_S(w) = \begin{cases} \beta\delta v_{2,TC}^{\text{search}}(w), & w \leq w_{2,S}^*, \\ \beta\delta v_{2,TC}^{\text{no-search}}(w), & w > w_{2,S}^*. \end{cases}$$

Since $w_{2,S}^* < w_{2,TC}^*$, it holds that

$$\beta\delta v_{2,TC}^{\text{search}}(w_{2,S}^*) > \beta\delta v_{2,TC}^{\text{no-search}}(w_{2,S}^*),$$

so U_S exhibits a downward jump at $w_{2,S}^*$. Define $\tilde{w} < w_{2,S}^*$ as follows. If

$$\beta\delta v_{2,TC}^{\text{no-search}}(w_{2,S}^*) \geq \beta\delta v_{2,TC}^{\text{search}}(w_l),$$

let $\tilde{w}$ be the unique solution to

$$\beta\delta v_{2,TC}^{\text{search}}(\tilde{w}) = \beta\delta v_{2,TC}^{\text{no-search}}(w_{2,S}^*).$$

Such $\tilde{w}$ exists and is unique because $v_{2,TC}^{\text{search}}$ is strictly increasing and continuous. Otherwise set $\tilde{w} = w_l$. Now fix $w_1 \in [\tilde{w},\, w_{2,S}^*]$. Then

$$\beta\delta v_{2,TC}^{\text{search}}(w_1) \geq \beta\delta v_{2,TC}^{\text{no-search}}(w_{2,S}^*),$$

with strict inequality for $w_1 \in (\tilde{w},\, w_{2,S}^*]$. Moreover,

$$\beta\delta v_{2,TC}^{\text{search}}(w_1) < \beta\delta v_{2,TC}^{\text{search}}(w_{2,TC}^*) = \beta\delta v_{2,TC}^{\text{no-search}}(w_{2,TC}^*).$$

Since $v_{2,TC}^{\text{no-search}}$ is strictly increasing and continuous, there exists a unique $\psi(w_1) \geq 0$, with $\psi(w_1) > 0$ for all $w_1 \in (\tilde{w},\, w_{2,S}^*]$, such that

$$\beta\delta v_{2,TC}^{\text{search}}(w_1) = \beta\delta v_{2,TC}^{\text{no-search}}(w_{2,S}^* + \psi(w_1)).$$

Then offers in $(w_{2,S}^*,\, w_{2,S}^* + \psi(w_1))$ yield a strictly lower continuation payoff than staying at w_1 and are thus rejected, while offers in $(w_1, w_{2,S}^*]$ and offers above $w_{2,S}^* + \psi(w_1)$ yield a strictly higher continuation payoff and are thus accepted. This proves (ii)(b).

Statement (ii)(c) immediately follows by applying the same argument in the reverse direction to wages $w_1 \in (w^*_{2,S},\, w^*_{2,S} + \psi(w^*_{2,S}))$.

It remains to prove (ii)(a). If $w_1 < \tilde{w}$, then

$$\beta\delta v^{\text{search}}_{2,TC}(w_1) < \beta\delta v^{\text{no-search}}_{2,TC}(w^*_{2,S}).$$

Since $v^{\text{no-search}}_{2,TC}$ is strictly increasing on $(w^*_{2,S}, w_h]$, every offer $w' > w^*_{2,S}$ yields a strictly higher continuation payoff than staying at w_1. Together with strict monotonicity of $v^{\text{search}}_{2,TC}$ on $[w_l, w^*_{2,S}]$, it follows that the worker accepts any offer with wage $w' > w_1$ and rejects all other offers. Finally, suppose $w_1 \geq w^*_{2,S} + \psi(w^*_{2,S})$. Then

$$U_S(w_1) \geq U_S(w^*_{2,S}) = U_S(w^*_{2,S} + \psi(w^*_{2,S})).$$

Since $U_S(w) = \beta\delta v^{\text{search}}_{2,TC}(w)$ is strictly increasing on $[w_l, w^*_{2,S}]$, every offer $w' \leq w^*_{2,S}$ is rejected. Since moreover $U_S(w) = \beta\delta v^{\text{no-search}}_{2,TC}(w)$ is strictly increasing for $w > w^*_{2,S}$, it follows that the worker accepts an offer if and only if $w' > w_1$, proving (ii)(a). □

C.1.3. Proof of Proposition 3.2

Lemma C.3 *The perception-perfect period-1 search strategies are characterized as follows:*

(i) *The time-consistent type searches if and only if*

$$w_1 \leq w_h - \sqrt{\frac{2(w_h - w_l)c}{\delta(1+\delta)}} \coloneqq w^*_{1,TC},$$

where $w^*_{1,TC} > w^*_{2,TC}$.

(ii) *If* $\beta < \frac{1}{1+\delta}$, *the present-biased types* $\phi \in \{S, N\}$ *search if and only if*

$$w_1 \leq w^*_{1,\phi},$$

where $w^*_{1,S} > w^*_{1,N}$ *and* $w^*_{1,N}, w^*_{1,S} \in (w^*_{2,\phi},\, w^*_{2,TC})$.

(iii) *If* $\beta \geq \frac{1}{1+\delta}$, *the present-biased types* $\phi \in \{S, N\}$ *search if and only if*

$$w_1 \leq w_h - \sqrt{\frac{2(w_h - w_l)c}{\beta\delta(1+\delta)}} \coloneqq w^*_{1,\phi},$$

where $w^*_{1,\phi} \in [w^*_{2,TC},\, w^*_{1,TC})$.

Proof. For $\phi \in \{S, N, TC\}$, let $v^{\text{search}}_{1,\phi}(w)$ and $v^{\text{no-search}}_{1,\phi}(w)$ denote the period-1 payoffs from searching and not searching at wage w, given the period-1 acceptance behavior from

Lemma C.2 and the anticipated period-2 search behavior as described in Lemma C.1. Define the gain from search in period 1 as

$$H_{1,\phi}(w) := v_{1,\phi}^{\text{search}}(w) - v_{1,\phi}^{\text{no-search}}(w).$$

Time-consistent worker. We first show that $H_{1,TC}(w) > 0$ for all $w \leq w_{2,TC}^*$.

Fix $w \leq w_{2,TC}^*$. Then,

$$v_{2,TC}^{\text{search}}(w) = w - c + \delta\Big((1 - F(w))\,\mathbb{E}[w'' \mid w'' > w] + F(w)w\Big) \geq w(1+\delta) = v_{2,TC}^{\text{no-search}}(w).$$

Subtract δc from the right-hand side and add $\delta^2\Big((1 - F(w))\,\mathbb{E}[w'' \mid w'' > w] + F(w)w\Big)$ to both sides. Then,

$$\begin{aligned} &w - c + \delta(1+\delta)\Big((1 - F(w))\,\mathbb{E}[w'' \mid w'' > w] + F(w)w\Big) \\ &> w + \delta\Big(w - c + \delta(1 - F(w))\,\mathbb{E}[w'' \mid w'' > w] + \delta F(w)w\Big). \end{aligned}$$

The right-hand side is exactly $v_{1,TC}^{\text{no-search}}(w)$. The left-hand side is the period-1 payoff from searching under the assumption that the worker never searches again in period 2. Since this differs from the actual period-2 behavior from Lemma C.1, which under time consistency is also optimal from the period-1 perspective, the actual payoff from searching must be weakly larger than the left-hand side. Hence

$$v_{1,TC}^{\text{search}}(w) > v_{1,TC}^{\text{no-search}}(w).$$

Thus, the time-consistent worker searches for all wages $w \leq w_{2,TC}^*$ in period 1.

Now consider $w > w_{2,TC}^*$. By Lemma C.2, conditional on searching, the time-consistent worker accepts any higher offer and rejects all other offers. Moreover, by Lemma C.1, any wage strictly above $w_{2,TC}^*$ induces no search in period 2. Hence,

$$H_{1,TC}(w) = -c + \delta(1+\delta)\Big((1 - F(w))\,\mathbb{E}[w' \mid w' > w] + F(w)w - w\Big).$$

This function is continuous and strictly decreasing in w. Solving $H_{1,TC}(w_{1,TC}^*) = 0$ yields

$$w_{1,TC}^* = w_h - \sqrt{\frac{2(w_h - w_l)c}{\delta(1+\delta)}} > w_{2,TC}^*$$

such that the time-consistent worker searches if and only if $w_1 \leq w_{1,TC}^*$. This proves (i).

Auxiliary results for present-biased types. Recall $w_{2,N}^* = w_{2,S}^*$. For the remainder of this proof we use subscript ϕ to denote present-biased types. We show three auxiliary results.

Result 1: $H_{1,\phi}(w) > 0$ for every $w \leq w_{2,\phi}^*$. For $w \leq w_{2,\phi}^*$, both present-biased types search in period 2. Hence,

$$w - c + \beta\delta\Big((1 - F(w))\,\mathbb{E}[w' \mid w' > w] + F(w)w\Big) \geq w + \beta\delta w.$$

Subtracting $\beta\delta c$ from the right-hand side and adding $\beta\delta^2\Big((1-F(w))\,\mathbb{E}[w' \mid w' > w] + F(w)w\Big)$ to both sides yields

$$\begin{aligned} &w - c + \beta\delta(1+\delta)\Big((1-F(w))\,\mathbb{E}[w' \mid w' > w] + F(w)w\Big) \\ &> w + \beta\delta\Big(w - c + \delta(1-F(w))\,\mathbb{E}[w' \mid w' > w] + \delta F(w)w\Big) = v_{1,\phi}^{\text{no-search}}(w). \end{aligned}$$

As above, the left-hand side is a lower bound on $v_{1,\phi}^{\text{search}}(w)$, so $H_{1,\phi}(w) > 0$.

Result 2: $H_{1,N}$ is continuous on $[w_l, w_h]$. This follows because $H_{1,N}(w)$ is obtained from the continuous time-consistent continuation payoff U_{TC} (see proof of Lemma C.2) through integration and continuity-preserving algebraic operations.

Result 3: for both present-biased types, $H_{1,\phi}(w)$ is strictly decreasing and continuous on $(w_{2,\phi}^*,\, w_{2,TC}^*]$. Define the perceived continuation payoff from entering period 2 at wage x by

$$U_\phi(x) := \begin{cases} \beta\delta\, v_{2,TC}^{\text{search}}(x), & \text{if } \phi = N \text{ and } x \le w_{2,TC}^*, \\ \beta\delta\, v_{2,TC}^{\text{no-search}}(x), & \text{if } \phi = N \text{ and } x > w_{2,TC}^*, \\ \beta\delta\, v_{2,TC}^{\text{search}}(x), & \text{if } \phi = S \text{ and } x \le w_{2,S}^*, \\ \beta\delta\, v_{2,TC}^{\text{no-search}}(x), & \text{if } \phi = S \text{ and } x > w_{2,S}^*. \end{cases}$$

Then, for every $w \in (w_{2,\phi}^*,\, w_{2,TC}^*]$, we can write

$$H_{1,\phi}(w) = -c + \mathbb{E}\Big[\big(U_\phi(x) - U_\phi(w)\big)_+\Big],$$

where $(z)_+ := \max\{z, 0\}$. Continuity follows by dominated convergence: on $(w_{2,\phi}^*, w_{2,TC}^*]$, $U_\phi(w)$ is continuous in w. Hence, for any sequence $w_n \to w$ in this interval,

$$\big(U_\phi(x) - U_\phi(w_n)\big)_+ \to \big(U_\phi(x) - U_\phi(w)\big)_+$$

for every offer x. Since U_ϕ is bounded, the integrand in the expectation term is bounded by $2\|U_\phi\|_\infty$, which is integrable. Thus $H_{1,\phi}$ is continuous on $(w_{2,\phi}^*, w_{2,TC}^*]$.

Now take $w_a < w_b$ in $(w_{2,\phi}^*,\, w_{2,TC}^*]$. Since U_ϕ is strictly increasing on $(w_a, w_h]$, it follows that $U_\phi(w_a) < U_\phi(w_b)$ and that, for every x,

$$\big(U_\phi(x) - U_\phi(w_b)\big)_+ \le \big(U_\phi(x) - U_\phi(w_a)\big)_+.$$

The inequality is strict for all $x > w_a$, which has positive probability. Taking expectations yields

$$H_{1,\phi}(w_b) < H_{1,\phi}(w_a).$$

Hence $H_{1,\phi}$ is continuous and strictly decreasing on $(w^*_{2,\phi},\, w^*_{2,TC}]$ for both present-biased types.

Case $\beta \geq \frac{1}{1+\delta}$. Recall from Result 1 that $H_{1,\phi}(w) > 0$ for every $w \leq w^*_{2,\phi}$. For $w \in (w^*_{2,\phi},\, w^*_{2,TC}]$, define a lower bound on the payoff from searching by

$$\underline{v}^{\text{search}}_{1,\phi}(w) := w - c + \beta\delta(1+\delta)\Big((1-F(w))\,\mathbb{E}[w' \mid w' > w] + F(w)w\Big),$$

obtained by assuming that, after searching in period 1, the worker accepts all higher offers, rejects all lower offers, and does not search again in period 2. This is weakly below the payoff from searching for both present-biased types. Also define an upper bound on the payoff from not searching by

$$\overline{v}^{\text{no-search}}_{1,\phi}(w) := w + \beta\delta\Big(w - c + \delta(1-F(w))\,\mathbb{E}[w' \mid w' > w] + \delta F(w)w\Big).$$

This upper bound is obtained by assuming that the worker searches at wage w in period 2. For the naive worker it equals $v^{\text{no-search}}_{1,N}(w)$, since she expects to search in period 2. For the sophisticated worker it is an upper bound on $v^{\text{no-search}}_{1,S}(w)$, since she anticipates that she will not search in period 2 at such wages, which is suboptimal from the perspective of period 1. Hence search is optimal whenever $\underline{v}^{\text{search}}_{1,\phi}(w) \geq \overline{v}^{\text{no-search}}_{1,\phi}(w)$, that is, if

$$\beta\delta\Big((1-F(w))\,\mathbb{E}[w' \mid w' > w] + F(w)w - w + c\Big) \geq c. \tag{C.2}$$

For every $w \leq w^*_{2,TC}$, the time-consistent period-2 search condition gives

$$\delta\Big((1-F(w))\,\mathbb{E}[w' \mid w' > w] + F(w)w - w\Big) \geq c.$$

Together with $\beta \geq 1/(1+\delta)$, this implies (C.2). Therefore, both present-biased types search at all wages $w \leq w^*_{2,TC}$.

For $w > w^*_{2,TC}$, both present-biased types anticipate no period-2 search, accept all higher offers, and reject all lower offers. Thus

$$H_{1,\phi}(w) = -c + \beta\delta(1+\delta)\Big((1-F(w))\,\mathbb{E}[w' \mid w' > w] + F(w)w - w\Big). \tag{C.3}$$

At $w = w^*_{2,TC}$, the same expression applies because $v^{\text{search}}_{2,TC}(w^*_{2,TC}) = v^{\text{no-search}}_{2,TC}(w^*_{2,TC})$. This expression is continuous and strictly decreasing in w, and solving $H_{1,\phi}(w^*_{1,\phi}) = 0$ yields

$$w^*_{1,\phi} = w_h - \sqrt{\frac{2(w_h - w_l)c}{\beta\delta(1+\delta)}}.$$

Moreover,

$$w^*_{1,\phi} \geq w^*_{2,TC} \quad \Longleftrightarrow \quad \beta \geq \frac{1}{1+\delta}. \tag{C.4}$$

Hence the present-biased worker searches if and only if $w_1 \leq w^*_{1,\phi}$, with $w^*_{1,\phi} \in [w^*_{2,TC}, w^*_{1,TC})$. This proves (iii).

Case $\beta < \frac{1}{1+\delta}$. For the present-biased types, Result 1 and the expression for $H_{1,\phi}(w)$ on $[w^*_{2,TC}, w_h]$ in (C.3), together with (C.4), imply

$$H_{1,\phi}(w) > 0 \quad \forall\, w \leq w^*_{2,\phi}, \qquad H_{1,\phi}(w) < 0 \quad \forall\, w \geq w^*_{2,TC}.$$

For the naive worker, since $H_{1,N}$ is continuous on $[w_l, w_h]$ by Result 2 and strictly decreasing on $(w^*_{2,\phi}, w^*_{2,TC}]$ by Result 3, there exists a unique

$$w^*_{1,N} \in (w^*_{2,\phi},\, w^*_{2,TC})$$

such that the naive worker searches if and only if $w_1 \leq w^*_{1,N}$. It remains to show that $w^*_{1,S} > w^*_{1,N}$. Fix $w \in (w^*_{2,\phi},\, w^*_{2,TC})$. For the naive worker,

$$\begin{aligned} H_{1,N}(w) = -c + \beta\delta \Bigg(& -(1-F(w))\, v^{\text{search}}_{2,TC}(w) \\ & + \big(F(w^*_{2,TC}) - F(w)\big)\, \mathbb{E}\big[v^{\text{search}}_{2,TC}(w') \mid w' \in (w, w^*_{2,TC}]\big] \\ & + \big(1 - F(w^*_{2,TC})\big)\, \mathbb{E}\big[v^{\text{no-search}}_{2,TC}(w') \mid w' > w^*_{2,TC}\big] \Bigg). \end{aligned}$$

For the sophisticated worker, consider the lower bound on the gain from search obtained by imposing the potentially suboptimal rule that after searching in period 1 she accepts all higher offers, rejects all lower offers and never searches in period 2:

$$\begin{aligned} \underline{H}_{1,S}(w) = -c + \beta\delta \Bigg(& -(1-F(w))\, v^{\text{no-search}}_{2,TC}(w) \\ & + \big(F(w^*_{2,TC}) - F(w)\big)\, \mathbb{E}\big[v^{\text{no-search}}_{2,TC}(w') \mid w' \in (w, w^*_{2,TC}]\big] \\ & + \big(1 - F(w^*_{2,TC})\big)\, \mathbb{E}\big[v^{\text{no-search}}_{2,TC}(w') \mid w' > w^*_{2,TC}\big] \Bigg) \leq H_{1,S}(w). \end{aligned}$$

Subtracting $H_{1,N}(w)$ yields

$$\begin{aligned} \underline{H}_{1,S}(w) - H_{1,N}(w) = \beta\delta \Bigg(& (1-F(w)) H_{2,TC}(w) \\ & - \big(F(w^*_{2,TC}) - F(w)\big)\, \mathbb{E}\big[H_{2,TC}(w') \mid w' \in (w, w^*_{2,TC}]\big] \Bigg). \end{aligned}$$

Since $H_{2,TC}$ is strictly decreasing (see $H_{2,\phi}$ in the proof of Lemma C.1 with $\beta = 1$),

$$H_{2,TC}(w) > \mathbb{E}\big[H_{2,TC}(w') \mid w' \in (w, w^*_{2,TC}]\big]\,.$$

Since also $1 - F(w) > F(w^*_{2,TC}) - F(w)$ and $H_{2,TC}(w) > 0$ on $(w^*_{2,\phi},\, w^*_{2,TC})$,

$$\underline{H}_{1,S}(w) > H_{1,N}(w) \qquad \forall\, w \in (w^*_{2,\phi},\, w^*_{2,TC}).$$

Since $\underline{H}_{1,S}$ is continuous on $(w^*_{2,\phi},\, w^*_{2,TC})$ and $\underline{H}_{1,S}(w^*_{1,N}) > H_{1,N}(w^*_{1,N}) = 0$, there exists $\varepsilon_1 > 0$ such that

$$\underline{H}_{1,S}(w) > 0 \qquad \forall\, w \in [w^*_{1,N},\, w^*_{1,N} + \varepsilon_1].$$

It follows that

$$H_{1,S}(w) > 0 \qquad \forall\, w \in [w^*_{1,N},\, w^*_{1,N} + \varepsilon_1].$$

Finally, for all $w \le w^*_{2,TC}$ sufficiently close to $w^*_{2,TC}$, the sophisticated worker accepts all higher offers, rejects all lower offers, and anticipates no period-2 search. Hence, on a left neighborhood of $w^*_{2,TC}$,

$$H_{1,S}(w) = -c + \beta\delta(1+\delta)\Big((1 - F(w))\,\mathbb{E}[w' \mid w' > w] + F(w)w - w\Big).$$

This expression is continuous in w, and for $\beta < \frac{1}{1+\delta}$ it is strictly negative at $w = w^*_{2,TC}$. Therefore, $H_{1,S}(w) < 0$ for all w sufficiently close to $w^*_{2,TC}$ from below. Since moreover $H_{1,S}$ is continuous and strictly decreasing on $(w^*_{2,\phi},\, w^*_{2,TC}]$ by Result 3, $H_{1,S}(w) > 0$ on $[w^*_{1,N}, w^*_{1,N} + \varepsilon_1]$, and $H_{1,S}(w) > 0$ for all $w \le w^*_{2,\phi}$ by Result 1, there exists a unique threshold $w^*_{1,S} \in (w^*_{1,N},\, w^*_{2,TC})$ such that the sophisticated worker searches if and only if $w_1 \le w^*_{1,S}$. This proves (ii) and completes the proof. □

C.2. Preliminary Lemmas and Results for the Infinite-Horizon Model

Roadmap of the Mathematical Appendix for the Infinite-Horizon Model

This section gives a short overview of how the infinite-horizon model's results are derived. For time-consistent and sophisticated workers, a perception-perfect strategy has a fixed point structure: current search and acceptance decisions must be optimal given the continuation value generated by the same strategy in all future periods. For sophisticated present-biased workers, this fixed point problem introduces an intertemporal wedge between the current self's evaluation of a strategy and the continuation value generated by future selves: the current self evaluates future gains from search with the additional present-bias discount factor β, while the continuation value does not include this additional discounting. This wedge drives the nonexistence result. For naive workers, the problem is simpler. They choose optimally today under the belief that future selves behave like time-consistent workers. Hence their problem does not require solving for a self-consistent strategy.

Step 1: Solving the time-consistent worker's problem. Lemma C.4 first establishes existence and uniqueness of the perception-perfect strategy of the time-consistent worker. The proof applies a standard contraction mapping argument to the Bellman operator that compares the payoff from not searching with the payoff from searching, where the latter incorporates optimal offer acceptance. The fixed point of this operator gives the unique value function. The proof then defines the strategy that chooses the optimal current behavior at each wage, given this value function, and verifies that this strategy is indeed the unique perception-perfect strategy.

The proof then denotes the unique value function by V_{TC} and shows that it is continuous and strictly increasing in the current wage. Strict monotonicity implies that, under the unique perception-perfect strategy, any higher wage offer is accepted. Continuity and strict monotonicity also imply that the current gain from search is continuous and strictly decreasing in the wage. Hence, given the future behavior induced by the existing perception-perfect strategy, current optimality takes a threshold form: search is optimal below a unique cutoff, no search is optimal above it, and the worker is indifferent at the cutoff. Since the strategy is perception-perfect, current and future behavior must be consistent. Therefore, the perception-perfect strategy itself must have a threshold form. Imposing indifference at the threshold yields the characterization of the search threshold.

Step 2: Solving the naive worker's problem. Since naive workers believe that future selves behave like time-consistent workers, the proof of Lemma C.5 uses the time-

consistent continuation value V_{TC} from Step 1 as their perceived continuation value. Since V_{TC} is strictly increasing, naive workers accept any higher offer. Their gain from search is also continuous and strictly decreasing, which yields a unique threshold search strategy.

Step 3: Establishing that no perception-perfect strategy exists for the sophisticated worker. Showing this result, stated in Lemma 3.2, consists of two steps. First, Lemma C.6 shows that if a perception-perfect strategy existed, it would have to take a threshold form. The proof of Lemma 3.2 then shows that no such threshold strategy can be perception-perfect. The reason is that the current self evaluates future gains from search with the additional present-bias discount factor β, while the continuation value generated by future selves does not apply this additional discounting. Hence, whenever search is weakly optimal for the current self at a given wage, the induced continuation value at that wage strictly exceeds the value of keeping the wage forever. Thus, when a threshold strategy switches from search to no search, the continuation value drops discontinuously. This discontinuity creates a profitable deviation to search just above the threshold.

Step 4: Characterizing the refined local equilibrium object. Lemma 3.3 constructs the maximal locally perception-perfect pair $(w_S^*; \bar{w})$, as defined in Definition 3.3. The proof starts from a candidate search threshold k and defines the permanent-wage equivalent of the continuation value at the cutoff: $z_k := (1-\delta)U_S^k(k)$. Then, the continuation value at k gives one equation linking k and z_k, while imposing indifference on the gain from search at k gives a second equation. Solving these two equations yields $(w_S^*; \bar{w})$.

The proof concludes by verifying consistency and maximality of both regions. The instability result in Corollary 3.1 follows immediately from maximality. Finally, Lemma C.7 compares thresholds and shows that the naive search threshold lies below the sophisticated no-search boundary, $w_N^* < \bar{w}$.

The non-monotone acceptance behavior in Lemma C.8 also follows immediately. Because $U_S(w_S^*) = \bar{w}/(1-\delta)$ and $\bar{w} > w_S^*$, the sophisticated worker near the top of the stable search region may reject higher offers that move her into the conjectured no-search region but deliver a lower continuation value than continued search.

Step 5: Deriving reservation-wage results. An unemployed worker compares the value of accepting an offer with the value of remaining unemployed, given the perceived on-the-job search behavior of her type as derived in the previous steps. If $b+c$ lies in the relevant search region, then accepting wage $b+c$ gives exactly the same value as remaining unemployed. Hence the lowest accepted wage is $b+c$. If $b+c$ lies above the relevant search cutoff, the reservation wage must compensate for the value of giving up search. This logic underlies Proposition 3.3. Corollary 3.2 then compares the resulting reservation wages across types.

Lemma C.4 *There exists a unique perception-perfect strategy of the time-consistent worker. Its associated payoff function is continuous and strictly increasing in w. Conditional on searching, the worker accepts a wage offer w' if and only if $w' > w$. Moreover, the worker searches if and only if*

$$w \leq w^*_{TC} := w_h - \sqrt{\frac{2(w_h - w_l)(1-\delta)c}{\delta}}.$$

PROOF. We proceed as follows: first, we show existence of a perception-perfect strategy. Then, we show uniqueness of this strategy and characterize the associated value function. Finally, we characterize the perception-perfect strategy.

Existence. Define the mapping T on the space $B([w_l, w_h])$ of bounded measurable functions $V : [w_l, w_h] \to \mathbb{R}$ by

$$(TV)(w) := \max\left\{w + \delta V(w),\ w - c + \delta \mathbb{E}\left[\max\{V(w), V(x)\}\right]\right\}, \qquad w \in [w_l, w_h], \tag{C.5}$$

where x denotes a wage offer. We first show that T maps $B([w_l, w_h])$ into itself. Let $V \in B([w_l, w_h])$ and set $M := \|V\|_\infty < \infty$. Then, for every $w \in [w_l, w_h]$,

$$|w + \delta V(w)| \leq w_h + \delta M \qquad \text{and} \qquad |w - c + \delta \mathbb{E}\left[\max\{V(w), V(x)\}\right]| \leq w_h + c + \delta M.$$

Hence, $|(TV)(w)| \leq w_h + c + \delta M$ for every $w \in [w_l, w_h]$, so TV is bounded. Moreover, TV is measurable. In particular, $(w, x) \mapsto \max\{V(w), V(x)\}$ is jointly measurable. Thus, integrating over x preserves measurability, so $w \mapsto \mathbb{E}[\max\{V(w), V(x)\}]$ is measurable. Hence, T maps the space of bounded measurable functions on $[w_l, w_h]$ into itself.

Next, we verify Blackwell's sufficient conditions to show that T is a contraction mapping. First, T is monotone. Let $V, W \in B([w_l, w_h])$ satisfy $V(w) \leq W(w)$ for all $w \in [w_l, w_h]$. Then, for every w,

$$\begin{aligned} w + \delta V(w) &\leq w + \delta W(w), \\ w - c + \delta \mathbb{E}\left[\max\{V(w), V(x)\}\right] &\leq w - c + \delta \mathbb{E}\left[\max\{W(w), W(x)\}\right]. \end{aligned}$$

Hence $(TV)(w) \leq (TW)(w)$ for every $w \in [w_l, w_h]$. Thus T is monotone. Second, T satisfies discounting for discount factor $\delta \in (0, 1)$. Let $a \geq 0$ be constant. Then,

$$\begin{aligned} (T(V + a))(w) &= \max\left\{w + \delta(V(w) + a),\ w - c + \delta \mathbb{E}\left[\max\{V(w) + a, V(x) + a\}\right]\right\} \\ &= (TV)(w) + \delta a. \end{aligned}$$

Hence T satisfies Blackwell's sufficient conditions. Therefore, T is a contraction on $B([w_l, w_h])$. Since $B([w_l, w_h])$ equipped with the sup norm is a complete metric space, the Banach Fixed Point Theorem implies that T has a unique fixed point V^* in $B([w_l, w_h])$ such that

$$V^*(w) = \max\left\{w + \delta V^*(w),\ w - c + \delta \mathbb{E}\left[\max\{V^*(w), V^*(x)\}\right]\right\}. \tag{C.6}$$

Define a strategy σ^* as follows. At each current wage w, search if and only if

$$w - c + \delta \mathbb{E}\left[\max\{V^*(w), V^*(x)\}\right] \geq w + \delta V^*(w).$$

Conditional on searching and observing an offer x, prescribe acceptance if and only if $V^*(x) > V^*(w)$.

We verify that σ^* constitutes a perception-perfect strategy. First, given continuation value V^*, the search decision prescribed by σ^* attains the maximum between searching and not searching, and the acceptance decision attains the maximum between accepting and rejecting an offer. Hence, the choices prescribed by σ^* are optimal in the current period. Moreover, these choices generate the current payoff V^* when the continuation value is V^*. Second, because V^* is a fixed point of T, following σ^* in all future periods generates the same continuation payoff function V^*. Thus the continuation values anticipated by the worker coincide with the payoffs induced by the strategy itself, so beliefs about future behavior are consistent with the optimal behavior prescribed by the strategy in every period. Therefore σ^* constitutes a perception-perfect strategy of the time-consistent worker.

Uniqueness of the perception-perfect strategy. Let σ be any perception-perfect strategy with associated payoff function V_{TC}. By definition of perception-perfectness, any current action prescribed by σ must be optimal given the continuation value it induces. Hence, V_{TC} must satisfy for every $w \in [w_l, w_h]$,

$$V_{TC}(w) = \max\left\{w + \delta V_{TC}(w),\ w - c + \delta \mathbb{E}\left[\max\{V_{TC}(w), V_{TC}(x)\}\right]\right\} \qquad \forall w \in [w_l, w_h].$$

Because $w \in [w_l, w_h]$, the per-period payoff is bounded, and since $\delta \in (0, 1)$, the discounted payoff generated by any feasible strategy is bounded. Hence, $V_{TC} \in B([w_l, w_h])$. Therefore V_{TC} is a bounded fixed point of the contraction mapping T. By uniqueness of the bounded fixed point, it follows that $V_{TC} = V^*$. Given this value function, the optimal action is unique except at indifference points, which are resolved by the maintained tie-breaking convention. Hence the induced perception-perfect strategy is unique.

We henceforth denote the value function associated with the unique perception-perfect strategy of the time-consistent worker as V_{TC}.

V_{TC} is continuous and strictly increasing in w. We first show that the fixed point V_{TC} is continuous on $[w_l, w_h]$. Let $C([w_l, w_h])$ denote the space of continuous bounded functions $V : [w_l, w_h] \to \mathbb{R}$. Take any $V \in C([w_l, w_h])$. Then TV is bounded and continuous. In particular, since $\max\{V(w), V(x)\}$ is continuous and bounded by an integrable function, the expectation over $\max\{V(w), V(x)\}$ is continuous by the dominated convergence theorem. All remaining components are continuous by standard arguments.

Thus T maps $C([w_l, w_h])$ into itself. Since by the previous argument T is a contraction and by completeness of $C([w_l, w_h])$ under the sup norm, the Banach Fixed Point Theorem implies that T has a unique fixed point in $C([w_l, w_h])$. As $C([w_l, w_h]) \subset B([w_l, w_h])$, by uniqueness of the fixed point in $B([w_l, w_h])$, this fixed point must coincide with V_{TC}. Hence V_{TC} is continuous on $[w_l, w_h]$.

We now show that V_{TC} is strictly increasing. Let V be any weakly increasing function on $[w_l, w_h]$. Then TV is strictly increasing, since w and $w - c$ are strictly increasing in w, while $\delta V(w)$ and $\max\{V(w), V(x)\}$ are weakly increasing in w.

Now define $V_0(w) \coloneqq 0$ for all $w \in [w_l, w_h]$, and recursively $V_{n+1} \coloneqq TV_n$. Since V_0 is weakly increasing, it follows that $V_1, V_2, \ldots$ are strictly increasing. Hence, each V_n is strictly increasing. Because V_n converges to the unique fixed point V_{TC},

$$V_{TC}(w') = \lim_{n\to\infty} V_n(w') \geq \lim_{n\to\infty} V_n(w) = V_{TC}(w) \qquad \text{for every } w' > w$$

since strict inequalities for each $n \geq 1$ imply weak inequalities in the limit. Thus, V_{TC} is weakly increasing in w. Since $V_{TC} = TV_{TC}$, the same argument used above implies that TV_{TC} and thus V_{TC} is strictly increasing.

Characterization of the perception-perfect strategy. Regarding offer acceptance behavior, the worker accepts an offer if and only if $w' > w$, since $V_{TC}(w)$ is strictly increasing in w.

We next show that the search behavior is characterized by a unique threshold $w^*_{TC} \in (w_l, w_h)$ such that the worker searches if and only if $w \leq w^*_{TC}$. We can write

$$V_{TC}(w) = \max\{V_{TC}^{\text{no-search}}(w), V_{TC}^{\text{search}}(w)\},$$

where

$$\begin{aligned} V_{TC}^{\text{no-search}}(w) &:= w + \delta V_{TC}(w), \\ V_{TC}^{\text{search}}(w) &:= w - c + \delta \left(F(w) V_{TC}(w) + (1 - F(w)) \mathbb{E}\left[V_{TC}(x) \mid x > w\right]\right). \end{aligned}$$

The expression for $V_{TC}^{\text{search}}(w)$ follows from the worker accepting an offer if and only if it exceeds the current wage. It holds that

$$V_{TC}^{\text{search}}(w) \geq V_{TC}^{\text{no-search}}(w) \iff -c + \delta(1 - F(w)) \left(\mathbb{E}\left[V_{TC}(x) \mid x > w\right] - V_{TC}(w)\right) \geq 0.$$

Define

$$H(w) := -c + \delta(1 - F(w)) \left(\mathbb{E}\left[V_{TC}(x) \mid x > w\right] - V_{TC}(w)\right). \tag{C.7}$$

We now show that $H(w)$ is strictly decreasing in w. For any $w \in [w_l, w_h]$, we can write

$$H(w) = -c + \delta \mathbb{E}\left[\left(V_{TC}(x) - V_{TC}(w)\right) \mathbf{1}\{x > w\}\right]. \tag{C.8}$$

Now take $w' > w$. Since V_{TC} is strictly increasing, for every realization of x,

$$(V_{TC}(x) - V_{TC}(w'))\,\mathbf{1}\{x > w'\} \leq (V_{TC}(x) - V_{TC}(w))\,\mathbf{1}\{x > w\},$$

with strict inequality for all $x > w$. Since $\mathbb{P}(x > w) > 0$ for every $w < w_h$, taking expectations yields

$$\mathbb{E}\left[(V_{TC}(x) - V_{TC}(w'))\,\mathbf{1}\{x > w'\}\right] < \mathbb{E}\left[(V_{TC}(x) - V_{TC}(w))\,\mathbf{1}\{x > w\}\right].$$

Therefore $H(w') < H(w)$, so $H(w)$ is strictly decreasing in w.

Moreover, $H(w)$ is continuous by the dominated convergence theorem, using continuity of V_{TC}, almost everywhere continuity of the indicator function, and boundedness of V_{TC}.

By the maintained parameter restriction, search is optimal at w_l, so $H(w_l) > 0$. At w_h, no strictly higher offer can arrive and therefore $H(w_h) = -c < 0$. By H being continuous and strictly decreasing, there exists a unique threshold $w^*_{TC} \in (w_l, w_h)$ such that

$$H(w^*_{TC}) = 0, \quad H(w) > 0 \iff w < w^*_{TC}, \quad H(w) < 0 \iff w > w^*_{TC}.$$

Hence, the worker searches if and only if $w \leq w^*_{TC}$.

It only remains to determine this threshold. At $w = w^*_{TC}$ the worker is indifferent between searching and not searching. Hence, $V_{TC}(w^*_{TC})$ is equal under searching and not searching. Specifically,

$$V_{TC}(w^*_{TC}) = w^*_{TC} + \delta V_{TC}(w^*_{TC}) \iff V_{TC}(w^*_{TC}) = \frac{w^*_{TC}}{1-\delta}.$$

For any $w > w^*_{TC}$, the worker does not search at wage w and thus receives w forever,

$$V_{TC}(w) = \frac{w}{1-\delta} \qquad \forall w > w^*_{TC}.$$

Substituting this into the indifference condition between not searching and searching at w^*_{TC} yields

$$w^*_{TC} + \delta\frac{w^*_{TC}}{1-\delta} = w^*_{TC} - c + \delta\left(F(w^*_{TC})\frac{w^*_{TC}}{1-\delta} + (1 - F(w^*_{TC}))\mathbb{E}\left[\frac{x}{1-\delta} \mid x > w^*_{TC}\right]\right).$$

Solving for w^*_{TC} yields

$$w^*_{TC} = w_h - \sqrt{\frac{2(w_h - w_l)(1-\delta)c}{\delta}}.$$

□

Lemma C.5 *There exists a unique perception-perfect strategy of the naive worker. Conditional on searching, the worker accepts a wage offer w' if and only if $w' > w$. Moreover, the worker searches if and only if $w \leq w^*_N$ for a unique threshold $w^*_N \in (w_l, w_h)$ satisfying $w^*_N < w^*_{TC}$.*

PROOF. By naivety, the worker evaluates future continuation payoffs as if future selves followed the perception-perfect strategy of the time-consistent worker. Hence, conditional on accepting an offer w', the perceived continuation payoff is $\beta\delta V_{TC}(w')$, whereas conditional on rejecting the offer and remaining at the current wage w, the perceived continuation payoff is $\beta\delta V_{TC}(w)$, where V_{TC} is the payoff function induced by the perception-perfect strategy of the time-consistent worker. Since $\beta\delta > 0$ and V_{TC} is strictly increasing in w by Lemma C.4, it follows that the worker accepts a wage offer if and only if $w' > w$.

It remains to characterize search behavior. At current wage w, the perceived payoff from not searching and from searching are given by

$$\begin{aligned} V_N^{\text{no-search}}(w) &:= w + \beta\delta V_{TC}(w), \\ V_N^{\text{search}}(w) &:= w - c + \beta\delta \left(F(w)V_{TC}(w) + (1 - F(w))\mathbb{E}\left[V_{TC}(x) \mid x > w\right]\right). \end{aligned}$$

Therefore, it holds that

$$V_N^{\text{search}}(w) \geq V_N^{\text{no-search}}(w) \iff -c + \beta\delta(1 - F(w))\left(\mathbb{E}\left[V_{TC}(x) \mid x > w\right] - V_{TC}(w)\right) \geq 0.$$

Define

$$H_N(w) := -c + \beta\delta(1 - F(w))\left(\mathbb{E}\left[V_{TC}(x) \mid x > w\right] - V_{TC}(w)\right).$$

Using the function H from the proof of Lemma C.4,

$$H_N(w) = \beta H(w) - (1 - \beta)c. \tag{C.9}$$

Since H is continuous and strictly decreasing, the same holds for H_N.

Under the maintained parameter restriction, search is optimal at w_l. Hence, it holds that $H_N(w_l) > 0$. Moreover, $H_N(w_h) = -c < 0$. Since H_N is continuous and strictly decreasing, there exists a unique threshold $w_N^* \in (w_l, w_h)$ such that $H_N(w_N^*) = 0$, and the naive worker searches if and only if $w \leq w_N^*$.

It remains to show that $w_N^* < w_{TC}^*$. Using (C.9) and that $H(w_{TC}^*) = 0$ by Lemma C.4,

$$H_N(w_{TC}^*) = \beta H(w_{TC}^*) - (1 - \beta)c = -(1 - \beta)c < 0.$$

Since H_N is strictly decreasing and satisfies $H_N(w_N^*) = 0$, it follows that $w_N^* < w_{TC}^*$. □

Lemma C.6 *Assume that the sophisticated worker follows a perception-perfect strategy. Then there exists $w_S^* \in (w_l, w_h)$ such that the worker searches for all $w < w_S^*$, does not search for all $w > w_S^*$, and at $w = w_S^*$ may either search or not search. Moreover, the continuation payoff is strictly increasing on the search region. Conditional on searching at wage w, any offer $w' \leq w$ is rejected, while any offer $w' \in (w, w_S^*)$ is accepted. If search is prescribed at w_S^*, then an offer $w' = w_S^*$ is accepted as well.*

PROOF. Assume there is a perception-perfect strategy of the sophisticated worker. We first show that in any such strategy, the worker must follow a threshold search rule.

Let $U_S : [w_l, w_h] \to \mathbb{R}$ denote the continuation payoff from the next period onward induced by the strategy. At current wage w, if the worker does not search, the current payoff is

$$V_S^{\text{no-search}}(w) := w + \beta\delta U_S(w).$$

If instead the worker searches, then after observing an offer x she compares the continuation payoff from rejecting the offer, namely $U_S(w)$, with the continuation payoff from accepting it, namely $U_S(x)$. Hence the current payoff from searching is

$$V_S^{\text{search}}(w) := w - c + \beta\delta\mathbb{E}\left[\max\{U_S(w), U_S(x)\}\right].$$

Therefore, search is optimal at wage w if and only if

$$H_S(w) := V_S^{\text{search}}(w) - V_S^{\text{no-search}}(w) = -c + \beta\delta\mathbb{E}\left[(U_S(x) - U_S(w))_+\right] \geq 0, \tag{C.10}$$

where $(z)_+ := \max\{z, 0\}$. We first establish two auxiliary claims.

Claim 1. If the worker does not search at wage w, then

$$U_S(w) = \frac{w}{1-\delta}.$$

This follows immediately from stationarity of the worker's behavior. If the strategy prescribes no search at wage w, then every future self at wage w also does not search, so

$$U_S(w) = w + \delta U_S(w) \iff U_S(w) = \frac{w}{1-\delta}.$$

Claim 2. If the worker searches at wage w, then

$$U_S(w) > \frac{w}{1-\delta}.$$

If the worker searches at wage w, stationarity implies

$$U_S(w) = w - c + \delta\mathbb{E}\left[\max\{U_S(w), U_S(x)\}\right] = w + \delta U_S(w) - c + \delta A(w), \tag{C.11}$$

where

$$A(w) := \mathbb{E}\left[(U_S(x) - U_S(w))_+\right].$$

Search optimality (C.10) for the current sophisticated self implies $-c+\beta\delta A(w) \geq 0$. Since $\beta \in (0,1)$, this yields $-c + \delta A(w) > 0$. Hence, it follows by (C.11) that

$$U_S(w) > \frac{w}{1-\delta}.$$

We now show that if search is prescribed at a given wage, then search must be prescribed at any lower wage. Suppose, toward a contradiction, that there exist $w < w'$ such that the worker does not search at w but does search at w'. By Claims 1 and 2,

$$U_S(w) = \frac{w}{1-\delta} \quad \text{and} \quad U_S(w') > \frac{w'}{1-\delta},$$

so $U_S(w) < U_S(w')$. Since search is optimal at w', we have $H_S(w') \geq 0$, which implies

$$\mathbb{E}\big[(U_S(x) - U_S(w'))_+\big] > 0.$$

Because $U_S(w) < U_S(w')$, it follows that

$$(U_S(x) - U_S(w))_+ \geq (U_S(x) - U_S(w'))_+$$

for all x, with strict inequality on a set of positive probability. Hence

$$\mathbb{E}\big[(U_S(x) - U_S(w))_+\big] > \mathbb{E}\big[(U_S(x) - U_S(w'))_+\big],$$

and therefore $H_S(w) > H_S(w') \geq 0$. Thus search is optimal at w, a contradiction. Hence search at w' implies search at all $w < w'$. Let

$$\mathcal{S} := \{w \in [w_l, w_h] : \text{search is prescribed at } w\}, \qquad w_S^* := \sup \mathcal{S}.$$

Then search is prescribed for all $w < w_S^*$ and no search is prescribed for all $w > w_S^*$, while at $w = w_S^*$ either action may be prescribed. This establishes the threshold structure.[1]

It remains to characterize acceptance behavior within the search region. For notational convenience, write the search region as $[w_l, w_S^*]$; if the strategy prescribes no search at w_S^*, the same argument applies with $[w_l, w_S^*)$ instead. For wages in the search region, stationarity and optimality of the acceptance decision in the assumed perception-perfect strategy imply

$$U_S(w) = w - c + \delta \mathbb{E}\left[\max\{U_S(w), \widetilde{U}_S(x)\}\right],$$

where

$$\widetilde{U}_S(x) := \begin{cases} U_S(x), & \text{if } x \leq w_S^*, \\ \dfrac{x}{1-\delta}, & \text{if } x > w_S^*. \end{cases}$$

Define the operator T on $B([w_l, w_S^*])$ by

$$(TV)(w) := w - c + \delta \mathbb{E}\left[\max\{V(w), \widetilde{V}(x)\}\right], \qquad w \in [w_l, w_S^*],$$

[1]Note that because continuity of H_S has not been established, the argument does not imply that necessarily $H_S(w_S^*) = 0$. Thus the threshold need not be an indifference point.

where

$$\widetilde{V}(x) := \begin{cases} V(x), & \text{if } x \le w_S^*, \\ \dfrac{x}{1-\delta}, & \text{if } x > w_S^*. \end{cases}$$

By a contraction mapping argument analogous to the proof of Lemma C.4, T is a contraction on the set of bounded measurable functions $B([w_l, w_S^*])$. Hence, U_S is the unique fixed point of T.[2]

We next show that U_S is strictly increasing on the search region. If V is weakly increasing on $[w_l, w_S^*]$ and $w' > w$, then

$$\begin{aligned}(TV)(w') - (TV)(w) &= w' - w + \delta\Big(\mathbb{E}[\max\{V(w'), \widetilde{V}(x)\}] - \mathbb{E}[\max\{V(w), \widetilde{V}(x)\}]\Big) \\ &\ge w' - w > 0.\end{aligned}$$

Thus TV is strictly increasing whenever V is weakly increasing. Starting from $V_0 \coloneqq 0$ and defining $V_{n+1} := TV_n$, each V_n is strictly increasing. Since V_n converges to the unique fixed point U_S, U_S is weakly increasing. Because $U_S = TU_S$, the same argument implies that U_S is strictly increasing on the search region.

Now fix a wage w in the search region and suppose the worker draws an offer w'. If $w' < w$, then $U_S(w') < U_S(w)$ so w' is rejected. If instead w' lies in the search region and satisfies $w' > w$, then $U_S(w') > U_S(w)$, so w' is accepted. So, any offer $w' \in (w, w_S^*)$ is accepted, and if search is prescribed at w_S^*, then an offer $w' = w_S^*$ is accepted as well. □

Lemma C.7 *Let w_N^* denote the perception-perfect search threshold of the naive worker and let $(w_S^*; \bar{w})$ be the unique maximal locally perception-perfect pair of the sophisticated worker. Then $w_N^* < \bar{w}$.*

Proof. By the proof of Lemma C.5, define the payoff gain from searching for the naive worker

$$H_N(w) := -c + \beta\delta\,\mathbb{E}\left[(V_{TC}(x) - V_{TC}(w))\,\mathbf{1}\{x > w\}\right],$$

where V_{TC} denotes the payoff function of the time-consistent worker. The naive worker searches at wage w if and only if $H_N(w) \ge 0$. We show that $H_N(\bar{w}) < 0$. By Lemma 3.3

[2]Hence, any other bounded (measurable) candidate function $G \neq U_S$ cannot be a self-consistent continuation payoff under the threshold rule and optimal acceptance. If the optimal acceptance rule were defined through G, the threshold rule and that acceptance rule would generate some actual payoff function W. If the acceptance rule is not optimal given W, the strategy is not perception-perfect; if it is optimal given W, then W must solve the fixed-point equation and therefore $W = U_S$. Thus, there is exactly one acceptance rule that is optimal given the payoff it induces, namely the rule defined by U_S. The same argument applies to any fixed candidate threshold k: the threshold-specific operator has a unique fixed point, which uniquely pins down the continuation payoff and optimal acceptance rule for that threshold. I use this in the proof of Lemma 3.3.

and Lemma C.4, $\bar{w} < w^*_{TC}$. Fix any $x \in (\bar{w}, w^*_{TC}]$. Since $\bar{w}$ and x lie in the search region of the time-consistent worker,

$$V_{TC}(x) - V_{TC}(\bar{w}) = (x - \bar{w}) + \delta\Big(\mathbb{E}[\max\{V_{TC}(x), V_{TC}(x')\}] - \mathbb{E}[\max\{V_{TC}(\bar{w}), V_{TC}(x')\}]\Big).$$

Since V_{TC} is strictly increasing and $x > \bar{w}$,

$$\max\{V_{TC}(x), V_{TC}(x')\} - \max\{V_{TC}(\bar{w}), V_{TC}(x')\} \leq V_{TC}(x) - V_{TC}(\bar{w}),$$

with strict inequality whenever $x' > \bar{w}$, an event of positive probability. Taking expectations yields

$$V_{TC}(x) - V_{TC}(\bar{w}) < (x - \bar{w}) + \delta\big(V_{TC}(x) - V_{TC}(\bar{w})\big),$$

and hence

$$V_{TC}(x) - V_{TC}(\bar{w}) < \frac{x - \bar{w}}{1 - \delta}.$$

If instead $x > w^*_{TC}$, then

$$V_{TC}(x) = \frac{x}{1 - \delta} \qquad \text{and} \qquad V_{TC}(\bar{w}) \geq \frac{\bar{w}}{1 - \delta},$$

so

$$V_{TC}(x) - V_{TC}(\bar{w}) \leq \frac{x - \bar{w}}{1 - \delta}.$$

Thus

$$V_{TC}(x) - V_{TC}(\bar{w}) \leq \frac{x - \bar{w}}{1 - \delta} \qquad \forall x > \bar{w},$$

and the inequality is strict for all $x \in (\bar{w}, w^*_{TC}]$, which is an event of positive probability. Therefore,

$$H_N(\bar{w}) < -c + \frac{\beta\delta}{1 - \delta}\mathbb{E}\left[(x - \bar{w})\mathbf{1}\{x > \bar{w}\}\right].$$

By construction of $\bar{w}$ (see proof of Lemma 3.3), the right side equals 0. Hence $H_N(\bar{w}) < 0$.

Since H_N is strictly decreasing (see proof of Lemma C.5), it follows that $H_N(w) < 0$ for all $w \geq \bar{w}$. Therefore the naive worker does not search at $\bar{w}$ or at any wage above $\bar{w}$. Since the naive worker searches exactly at wages $w \leq w^*_N$, this implies $w^*_N < \bar{w}$. □

Lemma C.8 *Under the unique perception-perfect strategies of the time-consistent and naive workers, conditional on searching, both workers accept any higher wage offer and reject any weakly lower wage offer. Under the unique maximal locally perception-perfect pair $(w^*_S; \bar{w})$, conditional on searching within the stable search region, the sophisticated worker rejects any weakly lower wage offer and, for wages sufficiently close to w^*_S from below, also rejects a nonempty interval of higher wage offers, while all sufficiently high wage offers, in particular all offers $w' > \bar{w}$, are accepted.*

PROOF. For the time-consistent and naive worker, the claim follows immediately from Lemma C.4 and Lemma C.5, respectively. It remains to consider the sophisticated worker.

Let U_S denote the continuation payoff induced by the conjecture that future selves search if and only if $w \leq w_S^*$, where $(w_S^*; \bar{w})$ is the unique maximal locally perception-perfect pair. By the proof of Lemma 3.3, U_S is continuous and strictly increasing on $[w_l, w_S^*]$. Hence, conditional on searching at wage $w \leq w_S^*$, every offer $x \leq w$ is rejected and every offer $x \in (w, w_S^*]$ is accepted.

It thus remains to study offers $x > w_S^*$. By the proof of Lemma 3.3,

$$U_S(w_S^*) = \frac{\bar{w}}{1-\delta} > \frac{w_S^*}{1-\delta}.$$

Since U_S is continuous and strictly increasing on $[w_l, w_S^*]$, there exists $\varepsilon > 0$ such that

$$U_S(w) > \frac{w_S^*}{1-\delta} \qquad \forall w \in (w_S^* - \varepsilon, w_S^*).$$

Fix any such w and define

$$\hat{w}(w) := (1-\delta)U_S(w).$$

Since U_S is strictly increasing on $[w_l, w_S^*]$,

$$\frac{w_S^*}{1-\delta} < U_S(w) < U_S(w_S^*) = \frac{\bar{w}}{1-\delta},$$

and therefore $\hat{w}(w) \in (w_S^*, \bar{w})$.

Now let $x > w_S^*$. Under the conjecture that future selves search if and only if $w \leq w_S^*$, accepting such an offer leads future selves not to search, so the continuation payoff of x is

$$U_S(x) = \frac{x}{1-\delta}.$$

Hence:

$$\begin{aligned} x \in (w_S^*, \hat{w}(w)) \quad &\Longrightarrow \quad U_S(x) = \frac{x}{1-\delta} < \frac{\hat{w}(w)}{1-\delta} = U_S(w), \\ x \in (\hat{w}(w), w_h] \quad &\Longrightarrow \quad U_S(x) = \frac{x}{1-\delta} > \frac{\hat{w}(w)}{1-\delta} = U_S(w). \end{aligned}$$

Because $\hat{w}(w) > w_S^*$, the interval $(w_S^*, \hat{w}(w))$ is nonempty. Thus, conditional on searching at wage $w \in (w_S^* - \varepsilon, w_S^*)$, the sophisticated worker rejects every offer $x \in (w_S^*, \hat{w}(w)]$ and accepts every offer $x \in (\hat{w}(w), w_h]$.

Finally, since $U_S(w_S^*) = \frac{\bar{w}}{1-\delta}$ and U_S is strictly increasing on $[w_l, w_S^*]$, it follows that for any $w \leq w_S^*$ and any offer $x > \bar{w}$,

$$U_S(x) = \frac{x}{1-\delta} > \frac{\bar{w}}{1-\delta} = U_S(w_S^*) \geq U_S(w),$$

so all offers $x > \bar{w}$ are always accepted. □

C.3. Proofs of the Infinite-Horizon Model

C.3.1. Proof of Lemma 3.2

Suppose, toward a contradiction, that a pure perception-perfect strategy exists. By Lemma C.6, it must be a threshold strategy with threshold $w_S \in (w_l, w_h)$. Let U_S denote the continuation payoff from the next period onward induced by this threshold strategy. Denote by $V_S^{\text{search}}(w), V_S^{\text{no-search}}(w)$ the current payoffs of searching and not searching, respectively. Let

$$D(w) := V_S^{\text{no-search}}(w) - V_S^{\text{search}}(w).$$

Assume first that search is prescribed at w_S. Then, *Claim 2* in the proof of Lemma C.6 implies $U_S(w_S) > \frac{w_S}{1-\delta}$. Hence there exists $\varepsilon_0 > 0$ such that for every $\varepsilon \in (0, \varepsilon_0)$,

$$\frac{w_S + \varepsilon}{1-\delta} < U_S(w_S).$$

At the threshold wage w_S, we have

$$D(w_S) = c + \beta\delta \left(U_S(w_S) - \mathbb{E}[\max\{U_S(x), U_S(w_S)\}] \right).$$

Now fix $\varepsilon \in (0, \varepsilon_0)$ such that $w := w_S + \varepsilon < w_h$. The threshold rule prescribes no search at w. So,

$$D(w) = c + \beta\delta \left(\frac{w}{1-\delta} - \mathbb{E}\left[\max\left\{ U_S(x), \frac{w}{1-\delta} \right\} \right] \right).$$

Subtracting gives

$$D(w) - D(w_S) = \beta\delta \left(\mathbb{E}\left[\max\{U_S(x), U_S(w_S)\} - \max\left\{ U_S(x), \frac{w}{1-\delta} \right\} \right] - \left(U_S(w_S) - \frac{w}{1-\delta} \right) \right).$$

For every realization of x,

$$\max\{U_S(x), U_S(w_S)\} - \max\left\{ U_S(x), \frac{w}{1-\delta} \right\} \leq U_S(w_S) - \frac{w}{1-\delta},$$

with strict inequality whenever

$$U_S(x) > \frac{w}{1-\delta}.$$

Since $w < w_h$, the event $\{x > w\}$ has positive probability, and for such x,

$$U_S(x) = \frac{x}{1-\delta} > \frac{w}{1-\delta}.$$

Hence the above inequality is strict on a set of positive probability, so $D(w) < D(w_S)$. Perception-perfectness at w_S requires $D(w_S) \leq 0$. Therefore $D(w) < 0$, so deviating to

searching in the current period is strictly optimal at $w = w_S + \varepsilon$. This contradicts the threshold rule, which prescribes no search at all wages above w_S.

It remains to consider the case in which no search is prescribed at w_S. Then

$$U_S(w_S) = \frac{w_S}{1-\delta}.$$

For any wage $w < w_S$, the proof of Claim 2 in Lemma C.6 gives

$$U_S(w) \geq \frac{w}{1-\delta} + \frac{c(1-\beta)}{\beta(1-\delta)}.$$

Indeed, search at w implies

$$U_S(w) = w + \delta U_S(w) - c + \delta A(w), \qquad -c + \beta\delta A(w) \geq 0,$$

where

$$A(w) := \mathbb{E}\left[(U_S(x) - U_S(w))_+\right].$$

Thus $-c + \delta A(w) \geq c(1-\beta)/\beta$, which yields the stated bound. Hence, for $w < w_S$ sufficiently close to w_S, $U_S(w) > U_S(w_S)$. Applying the same comparison as above with w and w_S in place of w_S and $w_S + \varepsilon$ yields

$$D(w_S) < D(w).$$

But search at w requires $D(w) \leq 0$. Therefore $D(w_S) < 0$, so deviating to searching in the current period is strictly optimal at w_S, a contradiction. Hence, no pure perception-perfect strategy exists. □

C.3.2. Proof of Lemma 3.3

Fix a conjectured candidate threshold $k \in (w_l, w_h)$ and suppose future selves search if and only if $w \leq k$. Applying the contraction mapping argument from the proof of Lemma C.6 to this threshold k induces a unique continuation payoff function satisfying

$$U_S^k(w) = w - c + \delta\mathbb{E}\big[\max\{U_S^k(w), \widetilde{U}_S^k(x)\}\big] \qquad \forall w \in [w_l, k],$$

where

$$\widetilde{U}_S^k(x) := \begin{cases} U_S^k(x), & \text{if } x \leq k, \\ \dfrac{x}{1-\delta}, & \text{if } x > k. \end{cases}$$

The proof of Lemma C.6 implies that U_S^k is strictly increasing on $[w_l, k]$. Moreover, continuity follows from an argument analogous to the proof of Lemma C.4: defining the

contraction mapping operator on the space of continuous functions and applying the dominated convergence theorem to the expectation term shows that the operator maps continuous functions into continuous functions. Hence its unique fixed point U_S^k is continuous on $[w_l, k]$. Define

$$z_k := (1-\delta)U_S^k(k).$$

Equivalently, $U_S^k(k) = z_k/(1-\delta)$. If search is consistent at k, then

$$U_S^k(k) > \frac{k}{1-\delta}.$$

Indeed, let

$$A_k := \mathbb{E}\left[\left(\widetilde{U}_S^k(x) - U_S^k(k)\right)_+\right].$$

Since k belongs to the conjectured search region,

$$U_S^k(k) = k + \delta U_S^k(k) - c + \delta A_k.$$

Search consistency at k implies $-c + \beta\delta A_k \geq 0$. Since $\beta < 1$, this implies $-c + \delta A_k > 0$. Hence

$$U_S^k(k) > k + \delta U_S^k(k),$$

and therefore $U_S^k(k) > k/(1-\delta)$ for any candidate threshold k at which search is consistent. Thus also $z_k = (1-\delta)U_S^k(k) > k$. Moreover, $z_k \leq w_h$, since the continuation payoff cannot exceed receiving w_h forever.

For any such candidate threshold k, we can thus evaluate $U_S^k(k)$ by replacing $\widetilde{U}_S^k(x)$ by $x/(1-\delta)$:

$$U_S^k(k) = k - c + \delta\mathbb{E}\left[\max\left\{U_S^k(k), \frac{x}{1-\delta}\right\}\right].$$

Since

$$\max\left\{U_S^k(k), \frac{x}{1-\delta}\right\} = \frac{z_k}{1-\delta} + \frac{1}{1-\delta}(x - z_k)\mathbf{1}\{x > z_k\},$$

it follows that

$$z_k = k - c + \frac{\delta}{1-\delta}\mathbb{E}\big[(x - z_k)\mathbf{1}\{x > z_k\}\big]. \tag{C.12}$$

Moreover, the current gain from search at the threshold is

$$H_S^k(k) = -c + \frac{\beta\delta}{1-\delta}\mathbb{E}\big[(x - z_k)\mathbf{1}\{x > z_k\}\big]. \tag{C.13}$$

Since the wage distribution is uniform,

$$\mathbb{E}\big[(x - z_k)\mathbf{1}\{x > z_k\}\big] = \frac{(w_h - z_k)^2}{2(w_h - w_l)}.$$

Thus (C.12) and (C.13) become

$$z_k = k - c + \frac{\delta}{1-\delta}\frac{(w_h - z_k)^2}{2(w_h - w_l)}, \tag{C.14}$$

$$H_S^k(k) = -c + \frac{\beta\delta}{1-\delta}\frac{(w_h - z_k)^2}{2(w_h - w_l)}. \tag{C.15}$$

To construct the candidate maximal search-consistent threshold, impose indifference at the cutoff, $H_S^k(k) = 0$. Using (C.15), this indifference condition yields

$$z^* = w_h - \sqrt{\frac{2(w_h - w_l)(1-\delta)c}{\beta\delta}}.$$

Substituting $z_k = z^*$ into (C.14) gives

$$w_S^* = z^* + c - \frac{\delta}{1-\delta}\frac{(w_h - z^*)^2}{2(w_h - w_l)} = z^* + c - \frac{c}{\beta}.$$

Define

$$\bar{w} := z^* \qquad \text{and hence} \qquad \bar{w} = w_S^* + \frac{c(1-\beta)}{\beta} > w_S^*.$$

Note that $\bar{w} = z^* < w_h$. By the maintained parameter restriction, $w_l < w_S^*$. We now verify that $(w_S^*; \bar{w})$ is the unique maximal locally perception-perfect pair.

Search consistency and maximality of the search region. Fix the conjectured threshold w_S^* and write $U_S(w) := U_S^{w_S^*}(w)$. For $w \le w_S^*$, the gain from search is

$$H_S(w) = -c + \beta\delta\mathbb{E}\big[(\tilde{U}(x) - U_S(w))_+\big],$$

where

$$\tilde{U}(x) := \begin{cases} U_S(x), & \text{if } x \le w_S^*, \\ \dfrac{x}{1-\delta}, & \text{if } x > w_S^*. \end{cases}$$

Since U_S is strictly increasing on $[w_l, w_S^*]$, it follows (analogously to the proof of Lemma C.6) that $H_S(w)$ is strictly decreasing on $[w_l, w_S^*]$. At $w = w_S^*$,

$$H_S(w_S^*) = -c + \beta\delta\mathbb{E}\big[(\tilde{U}(x) - U_S(w_S^*))_+\big].$$

Since $U_S(x) \le U_S(w_S^*)$ for all $x \le w_S^*$ and $U_S(w_S^*) = \frac{\bar{w}}{1-\delta}$, acceptance requires $x > \bar{w}$. Hence only offers $x > \bar{w}$ are accepted, and for such offers $\tilde{U}(x) = \frac{x}{1-\delta}$. Hence,

$$H_S(w_S^*) = -c + \frac{\beta\delta}{1-\delta}\mathbb{E}\big[(x - \bar{w})\mathbf{1}\{x > \bar{w}\}\big] = 0,$$

where we used that $U_S(w_S^*) = \frac{\bar{w}}{1-\delta}$ and that, by construction of $\bar{w}$, the right-hand side equals 0. Since $H_S(w_S^*) = 0$, it follows that $H_S(w) > 0$ for every $w < w_S^*$. Thus search is optimal for all $w \leq w_S^*$.

To show maximality, consider any conjectured threshold k such that search is consistent for all $w \leq k$. It thus must hold that $H_S^k(k) \geq 0$. By (C.15), this implies $z_k \leq \bar{w}$. Using (C.14),

$$k = z_k + c - \frac{\delta}{1-\delta}\frac{(w_h - z_k)^2}{2(w_h - w_l)}.$$

The right-hand side is strictly increasing in z_k. Hence $z_k \leq \bar{w}$ implies $k \leq w_S^*$. Therefore, w_S^* is the largest threshold such that search is consistent for all wages below the threshold.

No-search consistency and maximality of the no-search region. First consider any wage $w \geq \bar{w}$ under the conjectured threshold w_S^*. Since U_S is strictly increasing on $[w_l, w_S^*]$,

$$U_S(y) \leq U_S(w_S^*) = \frac{\bar{w}}{1-\delta} \leq \frac{w}{1-\delta} \qquad \forall y \leq w_S^*.$$

Thus, when deviating to search at wage w, accepting an offer in the search region is not optimal. Offers $x \in (w_S^*, w]$ yield $x/(1-\delta) \leq w/(1-\delta)$ and are also rejected. Hence only offers $x > w$ are accepted, and the gain from deviating to search is

$$\widehat{H}_S(w) := -c + \frac{\beta\delta}{1-\delta}\mathbb{E}\big[(x-w)\mathbf{1}\{x > w\}\big]. \tag{C.16}$$

The function $\widehat{H}_S$ is strictly decreasing, and by construction of $\bar{w} = z^*$,

$$\widehat{H}_S(\bar{w}) = 0.$$

Thus $\widehat{H}_S(w) < 0$ for all $w > \bar{w}$ and $\widehat{H}_S(\bar{w}) = 0$, so no search is indeed optimal for every $w \geq \bar{w}$.[3]

It remains to show that no smaller wage can define a stable no-search region. Consider any $w < \bar{w}$ and any conjectured threshold that assigns w to the no-search region. If the worker deviates to search at w, one feasible but potentially suboptimal acceptance rule is to accept only offers $x > w$. This feasible rule alone yields the gain

$$\widehat{H}_S(w) = -c + \frac{\beta\delta}{1-\delta}\mathbb{E}\big[(x-w)\mathbf{1}\{x > w\}\big].$$

Since $\widehat{H}_S$ is strictly decreasing and $\widehat{H}_S(\bar{w}) = 0$, we have $\widehat{H}_S(w) > 0$. Hence a deviation to search is strictly profitable even under this possibly suboptimal acceptance rule. Therefore no-search can never be optimal at any wage $w < \bar{w}$ under any conjecture assigning that wage to the no-search region. Thus $\bar{w}$ is the smallest wage such that no search is consistent for all $w \geq \bar{w}$.

So, $(w_S^*; \bar{w})$ is locally perception-perfect and maximal. Uniqueness follows immediately: maximality of the consistent search and no-search regions pin down $\underline{w} = w_S^*$ and $\bar{w}$. □

[3]Strictly speaking, the tie-breaking convention selects search in the event of indifference. Still, no-search is weakly optimal at $\bar{w}$ and is therefore in the stable no-search region.

C.3.3. Proof of Corollary 3.1

Fix any $w \in (w_S^*, \bar{w})$. First, suppose that w is conjectured to lie in the no-search region. By the proof of Lemma 3.3, any $w \in (w_S^*, \bar{w})$ assigned to the no-search region admits a strictly profitable deviation to search, so search is optimal at wage w.

Second, suppose that w is conjectured to lie at the upper boundary of the search region, so that the conjectured threshold is $k = w$. Let $H_S^k(k)$ denote the current gain from searching rather than not searching at the cutoff under conjectured threshold k. The proof of Lemma 3.3 shows that

$$H_S^k(k) \geq 0 \quad \Longrightarrow \quad k \leq w_S^*.$$

Taking the contrapositive, any conjectured threshold $k > w_S^*$ satisfies $H_S^k(k) < 0$. Since here $k = w > w_S^*$, no-search is strictly optimal at wage w. □

C.3.4. Proof of Proposition 3.3

Fix any type $\phi \in \{TC, N, S\}$ and define the perceived on-the-job search thresholds

$$w_\phi^* = \begin{cases} w_{TC}^*, & \text{if } \phi \in \{TC, N\}, \\ w_S^*, & \text{if } \phi = S. \end{cases}$$

For this proof, let $V_\phi(w)$ denote the perceived continuation value from next period onward when employed at wage w, and let U_ϕ denote the corresponding continuation value when unemployed.

An unemployed worker receives b in the current period and observes an offer w'. Hence

$$U_\phi = b + \delta \mathbb{E}\big[\max\{V_\phi(w'), U_\phi\}\big]. \tag{C.17}$$

The worker accepts an offer w if and only if

$$V_\phi(w) \geq U_\phi.$$

We first consider the case $b + c \leq w_\phi^*$. Define

$$\Gamma_\phi(z) := b + \delta \mathbb{E}\big[\max\{V_\phi(w'), z\}\big].$$

Then U_ϕ is a fixed point of Γ_ϕ by (C.17). Moreover, for any $z, \tilde{z} \in \mathbb{R}$,

$$|\Gamma_\phi(z) - \Gamma_\phi(\tilde{z})| \leq \delta \mathbb{E}\left[|\max\{V_\phi(w'), z\} - \max\{V_\phi(w'), \tilde{z}\}|\right] \leq \delta |z - \tilde{z}|,$$

so Γ_ϕ is a contraction and therefore has a unique fixed point.

Now consider the wage $b+c$. Since $b+c \leq w^*_\phi$, this wage lies in the conjectured search region of type ϕ. Hence

$$V_\phi(b+c) = (b+c) - c + \delta \mathbb{E}\big[\max\{V_\phi(w'), V_\phi(b+c)\}\big] = \Gamma_\phi(V_\phi(b+c)).$$

Thus $V_\phi(b+c)$ is also a fixed point of Γ_ϕ. By uniqueness of the fixed point, $U_\phi = V_\phi(b+c)$. Therefore an unemployed worker accepts wage w if and only if

$$V_\phi(w) \geq V_\phi(b+c).$$

For $\phi \in \{TC, N\}$, the function V_ϕ is strictly increasing by Lemma C.4. For $\phi = S$, the function V_ϕ is strictly increasing on the search region (see proof of Lemma 3.3). Since $b+c \leq w^*_\phi$, it follows that

$$V_\phi(w) \geq V_\phi(b+c) \iff w \geq b+c \qquad \text{for all } w \leq w^*_\phi.$$

Hence the reservation wage is $r^*_\phi = b+c$. This proves part (i)(a) and (ii)(a).

We now turn to the case $b+c > w^*_\phi$. Equation (C.17) still holds. We first show that no wage $w \leq w^*_\phi$ is accepted. For any $w \leq w^*_\phi$, the wage lies in the conjectured search region, so

$$V_\phi(w) = w - c + \delta \mathbb{E}\big[\max\{V_\phi(w'), V_\phi(w)\}\big].$$

Hence

$$\Gamma_\phi(V_\phi(w)) = b + \delta \mathbb{E}\big[\max\{V_\phi(w'), V_\phi(w)\}\big] = V_\phi(w) + b + c - w > V_\phi(w). \tag{C.18}$$

Suppose, toward a contradiction, that there exists $w \leq w^*_\phi$ such that $V_\phi(w) \geq U_\phi$. Then

$$\Gamma_\phi(V_\phi(w)) - \Gamma_\phi(U_\phi) \leq \delta\big(V_\phi(w) - U_\phi\big),$$

so

$$\Gamma_\phi(V_\phi(w)) \leq V_\phi(w),$$

where we used that $\Gamma_\phi(U_\phi) = U_\phi$. This contradicts $\Gamma_\phi(V_\phi(w)) > V_\phi(w)$, i.e., (C.18). Hence

$$V_\phi(w) < U_\phi \qquad \forall w \leq w^*_\phi,$$

so no such wage is accepted. Therefore, if a reservation wage exists in this case, it must satisfy $r^*_\phi > w^*_\phi$. Accepting such a reservation wage places the worker in the conjectured no-search region, where

$$V_\phi(r^*_\phi) = \frac{r^*_\phi}{1-\delta}.$$

Moreover, the reservation wage must make the worker indifferent. If instead $V_\phi(r^*_\phi) > U_\phi$, continuity of V_ϕ on the conjectured no-search region and $r^*_\phi > w^*_\phi$ would imply that some wage $\tilde{w} \in (w^*_\phi, r^*_\phi)$ is also accepted, contradicting that r^*_ϕ is the minimum accepted wage. Hence

$$U_\phi = V_\phi(r^*_\phi) = \frac{r^*_\phi}{1-\delta}.$$

Since $r^*_\phi > w^*_\phi$, every accepted offer $w' \geq r^*_\phi$ lies in the conjectured no-search region, so $V_\phi(w') = w'/(1-\delta)$ for all $w' \geq r^*_\phi$. Thus (C.17) becomes

$$\begin{aligned}\frac{r^*_\phi}{1-\delta} &= b + \delta \mathbb{E}\left[\max\left\{\frac{w'}{1-\delta}, \frac{r^*_\phi}{1-\delta}\right\}\right] \\ &= b + \delta\left(\mathbb{P}(w' < r^*_\phi)\frac{r^*_\phi}{1-\delta} + \mathbb{P}(w' \geq r^*_\phi)\mathbb{E}\left[\frac{w'}{1-\delta} \mid w' \geq r^*_\phi\right]\right),\end{aligned}$$

which rearranges to

$$r^*_\phi = b + \frac{\delta}{1-\delta}\mathbb{P}(w' \geq r^*_\phi)\big(\mathbb{E}[w' \mid w' \geq r^*_\phi] - r^*_\phi\big).$$

It remains to show that such a reservation wage exists uniquely and lies above the relevant cutoff, w^*_ϕ. Define

$$\begin{aligned}F(r) &:= b + \frac{\delta}{1-\delta}\mathbb{P}(w' \geq r)\big(\mathbb{E}[w' \mid w' \geq r] - r\big) - r \\ &= b - r + \frac{\delta}{1-\delta}\mathbb{E}\big[(w'-r)\mathbf{1}\{w' \geq r\}\big].\end{aligned}$$

This function is continuous and strictly decreasing in r on $[w_l, w_h]$.

If $\phi \in \{TC, N\}$, then $b + c > w^*_{TC}$ together with the threshold condition (see proof of Lemma C.4) imply

$$F(w^*_{TC}) = b + c - w^*_{TC} > 0.$$

Since F is continuous, strictly decreasing, $F(w^*_{TC}) > 0$, and $F(w_h) = b - w_h < 0$, there exists a unique $r^*_\phi \in (w^*_{TC}, w_h)$ such that $F(r^*_\phi) = 0$. In particular, $r^*_\phi > w^*_\phi$.

If $\phi = S$, then by construction of $\bar{w}$ (see proof of Lemma 3.3),

$$-c + \frac{\beta\delta}{1-\delta}\mathbb{E}\big[(w'-\bar{w})\mathbf{1}\{w' \geq \bar{w}\}\big] = 0,$$

so

$$\frac{\delta}{1-\delta}\mathbb{E}\big[(w'-\bar{w})\mathbf{1}\{w' \geq \bar{w}\}\big] = \frac{c}{\beta}.$$

Hence

$$F(\bar{w}) = b - \bar{w} + \frac{\delta}{1-\delta}\mathbb{E}\big[(w'-\bar{w})\mathbf{1}\{w' \geq \bar{w}\}\big] = b - \bar{w} + \frac{c}{\beta}.$$

Using $\bar{w} = w_S^* + \frac{c(1-\beta)}{\beta}$, we obtain

$$F(\bar{w}) = b + c - w_S^* > 0.$$

Since F is continuous and strictly decreasing, $F(\bar{w}) > 0$, and $F(w_h) = b - w_h < 0$, there exists a unique solution $r_S^* \in (\bar{w}, w_h)$ such that $F(r_S^*) = 0$. In particular, the reservation wage lies strictly above the unstable region. □

C.3.5. Proof of Corollary 3.2

First, suppose $b + c \leq w_S^*$. By Lemma C.4 and Lemma 3.3, then also $b + c < w_{TC}^*$, so by Proposition 3.3,

$$r_S^* = r_{TC}^* = r_N^* = b + c.$$

Next, suppose $b + c > w_{TC}^*$. Then $b + c > w_S^*$, so for all types $\phi \in \{TC, N, S\}$,

$$r_\phi^* = b + \frac{\delta}{1-\delta}\mathbb{P}(w' \geq r_\phi^*)\big(\mathbb{E}[w' \mid w' \geq r_\phi^*] - r_\phi^*\big).$$

By uniqueness of the solution established in Proposition 3.3, it follows that

$$r_S^* = r_{TC}^* = r_N^*.$$

Next, suppose $b + c = w_{TC}^*$. By Proposition 3.3,

$$r_{TC}^* = r_N^* = w_{TC}^* = b + c.$$

Since $w_{TC}^* > w_S^*$, the sophisticated worker is in case (ii)(b) of Proposition 3.3. By the proof of Proposition 3.3 and since $b + c = w_{TC}^*$,

$$F(w_{TC}^*) = b - w_{TC}^* + c = 0.$$

Because r_S^* is the unique zero of F, $r_S^* = w_{TC}^* = b + c$.

Finally, suppose $b + c \in (w_S^*, w_{TC}^*)$. Then by Proposition 3.3,

$$r_{TC}^* = r_N^* = b + c.$$

Consider again the function F defined in the proof of Proposition 3.3,

$$F(r) = b - r + \frac{\delta}{1-\delta}\mathbb{P}(w' \geq r)\big(\mathbb{E}[w' \mid w' \geq r] - r\big).$$

Since $b + c < w_{TC}^*$ and the second term in F is strictly decreasing and equals c at w_{TC}^*,

$$F(b + c) > b - (b + c) + c = 0.$$

Since F is strictly decreasing by the proof of Proposition 3.3 and r_S^* is its unique zero,

$$r_S^* > b + c = r_{TC}^* = r_N^*.$$

□

C.4. Model Appendix

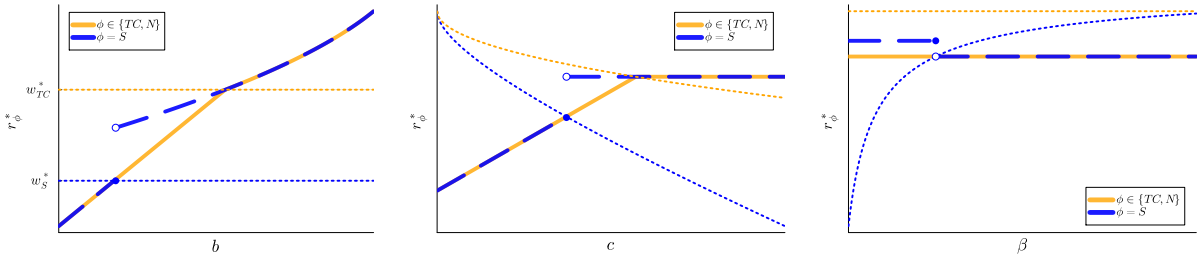
C.4.1. Comparative Statics of Reservation Wages

Differences in expected on-the-job search translate into differences in reservation wages. Figure C.1 illustrates how these differences evolve with unemployment benefits b, search costs c, and the present-bias parameter β.

Panel (a) shows how reservation wages vary with unemployment benefits b. If $b + c < w_S^*$, all types have reservation wage $r = b + c$. If $b + c \in (w_S^*, w_{TC}^*)$, the reservation wage of the sophisticated worker increases discontinuously and exceeds that of the other types. This discontinuity arises because, at wages just above $b + c$, accepting implies giving up future search, leading to a discontinuous drop in continuation payoffs. The sophisticated worker anticipates this and therefore requires higher wages to compensate for the foregone gains from future search. By contrast, time-consistent and naive workers continue to set $r = b + c$, as they expect to continue searching on the job. Once $b + c \geq w_{TC}^*$, all types anticipate no further search and reservation wages coincide again.

Figure C.1: Comparative Statics of Reservation Wages

(a) Unemployment Benefits b (b) On-the-job Search Cost c (c) Present Bias β



Note: Panels (a)–(c) plot reservation wages as a function of unemployment benefits b, on-the-job search costs c, and the present-bias parameter β, respectively. The orange solid line shows the reservation wage of the naive and time-consistent type, and the blue long-dashed line shows the reservation wage of the sophisticated type. The finely dashed lines indicate the type-specific expected on-the-job search thresholds.

Panel (b) shows how reservation wages vary with the on-the-job search cost c. An increase in c raises $b+c$ while lowering the on-the-job search thresholds. The same mechanism as in panel (a) applies: once $b+c > w_S^*$, the sophisticated reservation wage increases discontinuously to compensate for the foregone gains from future search. Thereafter, it remains flat, since accepting implies no further on-the-job search and c thus no longer affects payoffs. By contrast, time-consistent and naive workers continue to set $r = b + c$ until $b + c \geq w_{TC}^*$, after which all types coincide again.

Lastly, Panel (c) shows how reservation wages vary with the present-bias parameter β

for $b + c < w^*_{TC}$.[4] For time-consistent and naive workers, the reservation wage $r = b + c$ is independent of β. For the sophisticated worker, an increase in β raises the on-the-job search threshold w^*_S. By the same logic as above, as long as $w^*_S < b + c$, the reservation wage exceeds $b + c$. Once $w^*_S \geq b + c$, the reservation wage drops discontinuously to $b + c$.

C.5. Empirical Appendix

C.5.1. Details on Type Classification

This appendix describes the exact classification procedure from Cobb-Clark et al. (2024). Let w_{17} denote reported actual body weight in 2017, w_{18} reported actual body weight in 2018, w^E_{18} expected body weight in 2018 as reported in 2017, and w^I_{18} ideal body weight in 2018 as reported in 2017. All weights are measured in kilograms. The classification is restricted to respondents with complete information on these four variables. Following Cobb-Clark et al. (2024), the sample is further restricted to individuals who want to maintain or lose weight, $w^I_{18} \leq w_{17}$; whose absolute weight change between 2017 and 2018 does not exceed 10 kg, $|w_{18} - w_{17}| \leq 10$; whose realized weight does not fall below their ideal weight, $w_{18} \geq w^I_{18}$; whose expected weight does not fall below their ideal weight, $w^E_{18} \geq w^I_{18}$; and whose expected weight does not exceed their realized weight, $w^E_{18} \leq w_{18}$.

These restrictions reproduce the estimation sample of Cobb-Clark et al. (2024). They exclude individuals who want to gain weight, large weight changes that may reflect health shocks, and cases in which respondents expect or realize a future weight below their stated ideal weight. The latter cases are inconsistent with the framework in which weight loss is costly, since utility maximization would imply stopping once the ideal weight has been reached. The restrictions also exclude overly pessimistic beliefs, $w^E_{18} > w_{18}$, which would correspond to $\hat{\beta} < \beta$ and therefore fall outside the standard framework of present bias.

The classification allows for a two percent margin of error around realized weight to account for rounding and uncertainty in reported body weight. Define the lower bound $\underline{w}_{18} = 0.98 w_{18}$. Since the preceding restrictions already impose $w^E_{18} \leq w_{18}$, equality between expected and realized weight is coded as $w^E_{18} \geq \underline{w}_{18}$. Individuals are then classified as follows, reproducing the classification of 1,155 individuals in Cobb-Clark et al. (2024):

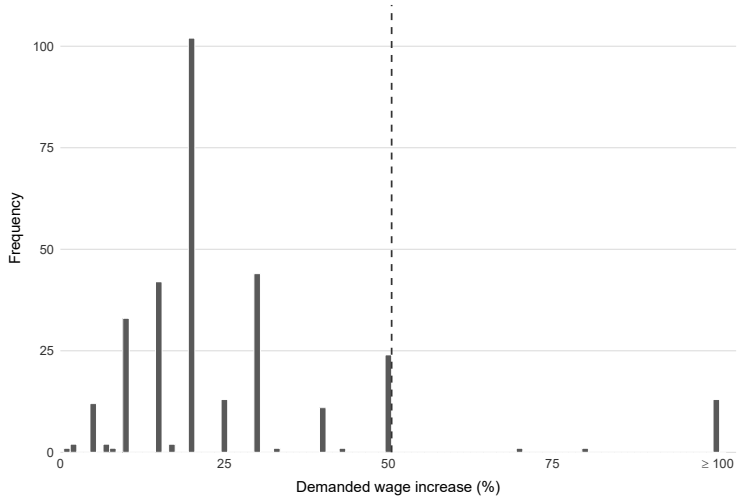
$$
\begin{aligned}
w^I_{18} = w^E_{18} \text{ and } w^E_{18} \geq \underline{w}_{18} &\implies \text{Time-consistent,} \\
w^I_{18} = w^E_{18} \text{ and } w^E_{18} < \underline{w}_{18} &\implies \text{Naive,} \\
w^I_{18} < w^E_{18} \text{ and } w^E_{18} \geq \underline{w}_{18} &\implies \text{Fully sophisticated,} \\
w^I_{18} < w^E_{18} \text{ and } w^E_{18} < \underline{w}_{18} &\implies \text{Partially sophisticated.}
\end{aligned}
$$

[4] For $b + c \geq w^*_{TC}$, reservation wages of all types coincide and are independent of β.

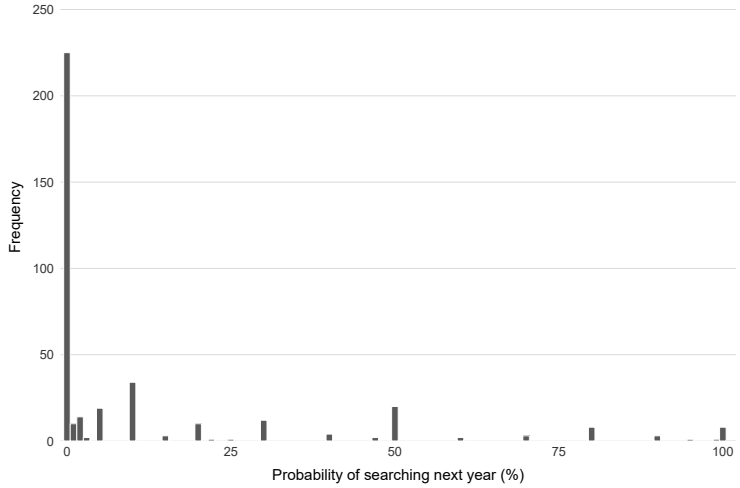
C.5.2. Additional Figures and Tables

Figure C.2: Distributions of Outcome Variables

(a): Demanded Wage Increases



(b): Probability of Searching



Note: Panel (a) plots the distribution of the wage increase demanded to be willing to change jobs, measured in percent. The vertical dashed line marks the exclusion threshold: responses strictly above 50% are excluded from the analysis. Panel (b) plots the distribution of workers' reported probability of engaging in on-the-job search during the next twelve months, also measured in percent. The plotted distributions consist of all observations with an assigned type and non-missing value for the respective outcome.

Table C.1: Separating Fully and Partially Sophisticated Workers

	Belief to Search (%)				Demanded Wage Increase (%)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	16.029***	8.407	17.320***	12.512	20.143***	22.964***	21.753***	35.157***
	(5.585)	(11.200)	(6.407)	(14.803)	(2.343)	(5.313)	(2.609)	(7.007)
S	−4.188	−4.021	-8.752	−8.519	6.667**	6.625**	7.399**	7.258**
	(5.628)	(5.495)	(6.055)	(5.899)	(3.103)	(3.118)	(3.324)	(3.315)
PS	−4.625	−4.000	-7.885	−7.233	3.931*	3.799*	2.776	2.253
	(4.949)	(4.893)	(5.433)	(5.428)	(2.254)	(2.244)	(2.626)	(2.583)
N	−1.921	−0.265	-2.303	−0.778	2.135	1.935	0.305	−0.701
	(5.049)	(4.969)	(5.607)	(5.658)	(2.161)	(2.124)	(2.445)	(2.440)
Demographic controls		✓		✓		✓		✓
Individual-trait controls			✓	✓			✓	✓
p-value: S = 0	0.458	0.465	0.150	0.150	0.033	0.035	0.027	0.030
p-value: PS = 0	0.351	0.414	0.148	0.184	0.083	0.092	0.292	0.384
p-value: N = 0	0.704	0.957	0.682	0.891	0.324	0.363	0.901	0.774
p-value: S = N	0.587	0.365	0.169	0.103	0.103	0.097	0.018	0.009
p-value: PS = N	0.412	0.271	0.155	0.101	0.310	0.300	0.276	0.185
Observations	383	382	325	325	305	304	255	255
R^2	0.006	0.025	0.049	0.067	0.037	0.039	0.077	0.093

Note: Columns (1)–(4) report results from weighted least squares regressions of workers' reported probability of engaging in on-the-job search in the next year, measured in percent, on worker type dummies. Columns (5)–(8) report results from weighted least squares regressions of the wage increase required to be willing to change jobs, also measured in percent, on worker type dummies. PS denotes partially sophisticated workers. The omitted category is time-consistent workers. The baseline specifications control for the gross wage in the previous month. Demographic and individual-trait controls are included as indicated. Standard errors are clustered at the individual level and reported in parentheses. The reported p-values are based on two-sided t tests. The p-values for $S = N$ and $PS = N$ are obtained from identical regressions in which naive workers are the omitted category. $^{*}: p < 0.1$; $^{**}: p < 0.05$; $^{***}: p < 0.01$.

Table C.2: Excluding Self-Reported Patience as a Control Variable

	Belief to Search (%)		Demanded Wage Increase (%)	
	(1)	(2)	(3)	(4)
Constant	17.035***	11.375	21.430***	33.528***
	(6.445)	(14.448)	(2.639)	(6.970)
S	−8.090	−7.555	4.538*	4.213*
	(5.281)	(5.202)	(2.535)	(2.521)
N	−2.334	−0.727	0.442	−0.456
	(5.673)	(5.693)	(2.451)	(2.431)
Demographic controls		✓		✓
Individual-trait controls	✓	✓	✓	✓
P-value: S = 0	0.127	0.148	0.075	0.097
P-value: N = 0	0.681	0.899	0.857	0.851
P-value: S = N	0.114	0.064	0.052	0.027
Observations	325	325	255	255
R^2	0.040	0.059	0.060	0.074

Note: Columns (1)–(2) report results from weighted least squares regressions of workers' reported probability of engaging in on-the-job search in the next year, measured in percent, on worker type dummies. Columns (3)–(4) report results from weighted least squares regressions of the wage increase required to be willing to change jobs, also measured in percent, on worker type dummies. The omitted category is time-consistent workers. All specifications include individual-trait controls, excluding self-reported patience. Demographic controls are included as indicated. Standard errors are clustered at the individual level and reported in parentheses. The reported p-values are based on two-sided t tests. The p-values for $S = N$ are obtained from identical regressions in which naive workers are the omitted category. $^{*}: p < 0.1$; $^{**}: p < 0.05$; $^{***}: p < 0.01$.

Table C.3: Extensive Margin of Search Beliefs

	Positive Search Belief			
	(1)	(2)	(3)	(4)
Constant	35.651***	34.360*	36.937***	66.666**
	(8.831)	(20.802)	(9.585)	(27.197)
S	−1.611	−1.056	−4.660	−5.345
	(8.104)	(7.841)	(8.943)	(8.602)
N	7.244	9.161	3.910	3.987
	(8.537)	(8.492)	(9.194)	(9.317)
Demographic controls		✓		✓
Individual-trait controls			✓	✓
p-value: S = 0	0.843	0.893	0.603	0.535
p-value: N = 0	0.397	0.282	0.671	0.669
p-value: S = N	0.169	0.112	0.231	0.195
Observations	383	382	325	325
R^2	0.009	0.070	0.056	0.093

Note: The table reports results from weighted least squares regressions of an indicator for positive on-the-job search beliefs on worker type dummies. The dependent variable equals 100 if the respondent reports a strictly positive probability of engaging in on-the-job search during the next twelve months and 0 otherwise. The omitted category is time-consistent workers. Demographic and individual-trait controls are included as indicated. Standard errors are clustered at the individual level and reported in parentheses. The p-values for $S = N$ are obtained from identical regressions in which naive workers are the omitted category. $^{*}: p < 0.1$; $^{**}: p < 0.05$; $^{***}: p < 0.01$.

Bibliography

Abberger, K., Funk, A.-K., Lamla, M., Lein, S., and Siegrist, S. S. (2025). Inflation-expectation pass-through into prices and wages: Evidence from an RCT survey. *Working Paper*.

Adda, J., and Dustmann, C. (2023). Sources of wage growth. *Journal of Political Economy*, *131*(2), 456–503.

Albrecht, J., Gautier, P. A., and Vroman, S. (2006). Equilibrium directed search with multiple applications. *The Review of Economic Studies*, *73*(4), 869–891.

Albrecht, J., Tan, S., Gautier, P., and Vroman, S. (2004). Matching with multiple applications revisited. *Economics Letters*, *84*(3), 311–314.

Arad, A., and Rubinstein, A. (2012). The 11-20 money request game: A level-k reasoning study. *American Economic Review*, *102*(7), 3561–73.

Armantier, O., Sbordone, A., Topa, G., van der Klaauw, W., and Williams, J. C. (2022). A new approach to assess inflation expectations anchoring using strategic surveys. *Journal of Monetary Economics*, *129*, 82–101.

Babecký, J., Du Caju, P., Kosma, T., Lawless, M., Messina, J., and Rõõm, T. (2010). Downward nominal and real wage rigidity: Survey evidence from European firms. *The Scandinavian Journal of Economics*, *112*(4), 884–910.

Bachmann, R., Born, B., Elstner, S., and Grimme, C. (2019). Time-varying business volatility and the price setting of firms. *Journal of Monetary Economics*, *101*, 82–99.

Bachmann, R., Carstensen, K., Lautenbacher, S., and Schneider, M. (2021). Uncertainty and change: Survey evidence of firms subjective beliefs. *NBER Working Paper No. 29430*.

Bachmann, R., Elstner, S., and Sims, E. R. (2013). Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics*, *5*(2), 217–49.

Bandiera, O., Parekh, N., Petrongolo, B., and Rao, M. (2022). Men are from mars, and women too: A bayesian meta-analysis of overconfidence experiments. *Economica*, *89*, 38–70.

Banfi, S., and Villena-Roldán, B. (2019). Do high-wage jobs attract more applicants? Directed search evidence from the online labor market. *Journal of Labor Economics*, *37*(3), 715–746.

Barattieri, A., Basu, S., and Gottschalk, P. (2014). Some evidence on the importance of sticky wages. *American Economic Journal: Macroeconomics*, *6*(1), 70–101.

Barber, B. M., and Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics*, *116*(1), 261–292.

Baumann, U., Ferrando, A., Georgarakos, D., Gorodnichenko, Y., and Reinelt, T. (2024). Safe to update inflation expectations? New survey evidence on euro area firms. *NBER Working Paper No. 32504*.

Belot, M., Kircher, P., and Muller, P. (2022). How wage announcements affect job search—a field experiment. *American Economic Journal: Macroeconomics*, *14*(4), 1–67.

Belot, M., Kircher, P., and Muller, P. (2025). Eliciting time preferences when income and consumption vary: Theory, validation, and application to job search. *American Economic Journal: Microeconomics*, *17*(1), 130–70.

Bernanke, B. S. (2007). *Inflation expectations and inflation forecasting* (Speech No. 306). Board of Governors of the Federal Reserve System (U.S.)

Bock, O., Baetge, I., and Nicklisch, A. (2014). Hroot: Hamburg registration and organization online tool. *European Economic Review*, *71*(C), 117–120.

Bronchetti, E. T., Kessler, J. B., Magenheim, E. B., Taubinsky, D., and Zwick, E. (2023). Is attention produced optimally? Theory and evidence from experiments with bandwidth enhancements. *Econometrica*, *91*(2), 669–707.

Bryan, G., Karlan, D., and Nelson, S. (2010). Commitment devices. *Annual Review of Economics*, *2*, 671–698.

Buchanan, J., and Houser, D. (2022). If wages fell during a recession. *Journal of Economic Behavior & Organization*, *200*, 1141–1159.

Buchheim, L., Link, S., and Möhrle, S. (2024). Inflation and wage expectations of firms and employees. *Working Paper*.

Buchheim, L., Dovern, J., Krolage, C., and Link, S. (2022). Sentiment and firm behavior during the COVID-19 pandemic. *Journal of Economic Behavior & Organization*, *195*, 186–198.

Burdett, K., and Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, *39*(2), 257–273.

Caliendo, M., Cobb-Clark, D. A., and Uhlendorff, A. (2015). Locus of control and job search strategies. *The Review of Economics and Statistics*, *97*(1), 88–103.

Candia, B., Coibion, O., and Gorodnichenko, Y. (2024). The inflation expectations of U.S. firms: Evidence from a new survey. *Journal of Monetary Economics*, *145*, 103569.

Cason, T. N., and Noussair, C. (2007). A market with frictions in the matching process: An experimental study. *International Economic Review*, *48*(2), 665–691.

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

Cobb-Clark, D., Dahmann, S., Kamhöfer, D., and Schildberg-Hörisch, H. (2024). Sophistication about self-control. *Journal of Public Economics*, *238*.

Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018). How do firms form their expectations? New survey evidence. *American Economic Review*, *108*(9), 2671–2713.

Cortés, P., Pan, J., Pilossoph, L., Reuben, E., and Zafar, B. (2023). Gender differences in job search and the earnings gap: Evidence from the field and lab. *The Quarterly Journal of Economics*, *138*(4), 2069–2126.

Crawford, V. P., Costa-Gomes, M. A., and Iriberri, N. (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, *51*(1), 5–62.

Davis, S. J., and Krolikowski, P. M. (2025). Sticky wages on the layoff margin. *American Economic Review*, *115*(2), 491–524.

DellaVigna, S., Lindner, A., Reizer, B., and Schmieder, J. F. (2017). Reference-dependent job search: Evidence from Hungary. *The Quarterly Journal of Economics*, *132*(4), 1969–2018.

DellaVigna, S., and Paserman, D. (2005). Job search and impatience. *Journal of Labor Economics*, *23*(3), 527–588.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., and Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, *9*(3), 522–550.

Druant, M., Fabiani, S., Kezdi, G., Lamo, A., Martins, F., and Sabbatini, R. (2012). Firms' price and wage adjustment in europe: Survey evidence on nominal stickiness. *Labour Economics*, *19*(5), 772–782.

Elsby, M. (2008). Evaluating the economic significance of downward nominal wage rigidity. *Journal of Monetary Economics*, *56*(2), 154–169.

Faberman, J., and Menzio, G. (2018). Evidence on the relationship between recruiting and the starting wage. *Labour Economics*, *50*, 67–79.

Faberman, J., Mueller, A. I., and Sahin, A. (2025). Job search and the gender wage gap. *Working Paper*.

Faberman, J., Mueller, A. I., Şahin, A., and Topa, G. (2022). Job search behavior among the employed and non-employed. *Econometrica*, *90*(4), 1743–1779.

Fahn, M., and Seibel, R. (2022). Present bias in the labor market – when it pays to be naive. *Games and Economic Behavior*, *135*, 144–167.

Fehr, E., and Goette, L. (2005). Robustness and real consequences of nominal wage rigidity. *Journal of Monetary Economics*, *52*(4), 779–804.

Fluchtmann, J., Glenny, A. M., Harmon, N. A., and Maibom, J. (2024). The gender application gap: Do men and women apply for the same jobs? *American Economic Journal: Economic Policy*, *16*(2), 182–219.

Galenianos, M., and Kircher, P. (2009). Directed search with multiple job applications. *Journal of Economic Theory*, *144*(2), 445–471.

Gautier, E., Savignac, F., and Coibion, O. (2025). Firms' inflation and wage expectations during the inflation surge. *NBER Working Paper*, *33799*.

Gillen, B., Snowberg, E., and Yariv, L. (2019). Experimenting with measurement error: Techniques with applications to the caltech cohort study. *Journal of Political Economy*, *127*(4), 1826–1863.

Gödl, M., and Gödl-Hanisch, I. (2024). Wage setting in times of high and low inflation. *CESifo Working Paper*, 1–42.

Goldin, C., and Rouse, C. (2000). Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review*, *90*(4), 715–741.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association*, *1*(1), 114–125.

Gul, F., and Pesendorfer, W. (2001). Temptation and self-control. *Econometrica*, *69*(6), 1403–1435.

Gürkaynak, R. S., Levin, A., and Swanson, E. (2010). Does inflation targeting anchor long-run inflation expectations? Evidence from the U.S., UK, and Sweden. *Journal of the European Economic Association*, *8*(6), 1208–1242.

Haaland, I., Roth, C., and Wohlfart, J. (2023). Designing information provision experiments. *Journal of Economic Literature*, *61*(1), 3–40.

Hajdini, I., Knotek, E., Leer, J., Pedemonte, M., Rich, R., and Schoenle, R. (2023). Low passthrough from inflation expectations to income growth expectations: Why people dislike inflation. *Federal Reserve Bank of Cleveland Working Paper Series No. 22-21R*.

Harris, C., and Laibson, D. (2001). Dynamic choices of hyperbolic consumers. *Econometrica*, *69*(4), 935–957.

He, H., Neumark, D., and Weng, Q. (2022). 'I still haven't found what I'm looking for': Evidence of directed search from a field experiment. *The Economic Journal*, *133*(649), 258–280.

Helland, L., Moen, E. R., and Preugschat, E. (2017). Information and coordination frictions in experimental posted offer markets. *Journal of Economic Theory*, *167*, 53–74.

Hospido, L., Laeven, L., and Lamo, A. (2022). The gender promotion gap: Evidence from central banking. *The Review of Economics and Statistics*, *104*(5), 981–996.

Jäger, S., Roth, C., Roussille, N., and Schoefer, B. (2024). Worker beliefs about outside options. *The Quarterly Journal of Economics*, *139*(3), 1505–1556.

Jain, M., Olena, K., and Xu, Z. (2024). How do people view price and wage inflation? *Journal of Monetary Economics*, *145*, 103552.

Kiessling, L., Pinger, P., Seegers, P., and Bergerhoff, J. (2024). Gender differences in wage expectations and negotiation. *Labour Economics*, *87*, 102505.

Kircher, P. (2009). Efficiency of simultaneous search. *Journal of Political Economy*, *117*(5), 861–913.

Kloosterman, A. (2016). Directed search with heterogeneous firms: An experimental study. *Experimental Economics*, *19*(1), 51–66.

Koenig, F., Manning, A., and Petrongolo, B. (2024). Reservation wages and the wage flexibility puzzle. *The Review of Economics and Statistics*, 1–32.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, *112*(2), 443–477.

Laibson, D. (2015). Why don't present-biased agents make commitments? *American Economic Review*, *105*(5), 267–72.

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2021). Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, *136*(1), 381–426.

Le Bihan, H., Montornès, J., and Heckel, T. (2012). Sticky wages: Evidence from quarterly microeconomic data. *American Economic Journal: Macroeconomics*, *4*(3), 1–32.

Link, S., Peichl, A., Pfäuti, O., Roth, C., and Wohlfart, J. (2026). Attention to the macroeconomy. *CESifo Working Paper No. 10858*.

Link, S., Peichl, A., Roth, C., and Wohlfart, J. (2023). Information frictions among firms and households. *Journal of Monetary Economics*, *135*, 99–115.

Lochner, B., and Merkl, C. (2026). Gender-specific application behaviour, matching, and the residual gender earnings gap. *The Economic Journal*, *136*, 97–124.

Maestas, N., Mullen, K. J., Powell, D., von Wachter, T., and Wenger, J. B. (2023). The value of working conditions in the United States and the implications for the structure of wages. *American Economic Review*, *113*(7), 2007–47.

Maskin, E., and Tirole, J. (2001). Markov perfect equilibrium: I. Observable actions. *Journal of Economic Theory*, *100*(2), 191–219.

McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, *84*(1), 113–126.

McGee, A., and McGee, P. (2016). Search, effort, and locus of control. *Journal of Economic Behavior & Organization*, *126*, 89–101.

McGee, A., and McGee, P. (2025). Gender differences in reservation wages in search experiments. *Labour Economics*, *94*, 102698.

Meyer, C. (2018). In search of a better life: Self-control in the Ethiopian labor market. *Working Paper*.

Mishkin, F. S. (2007). Inflation dynamics. *International Finance*, *10*(3), 317–334.

Moore, D. A., and Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, *115*(2), 502–517.

Mueller, A. I., Spinnewijn, J., and Topa, G. (2021). Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias. *American Economic Review*, *111*(1), 324–63.

Niederle, M., and Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *The Quarterly Journal of Economics*, *122*(3), 1067–1101.

O'Donoghue, T., and Rabin, M. (1999). Doing it now or later. *The American Economic Review*, *89*(1), 103–124.

O'Donoghue, T., and Rabin, M. (2001). Choice and procrastination. *The Quarterly Journal of Economics*, *116*(1), 121–160.

Pannenberg, M. (2010). Risk attitudes and reservation wages of unemployed workers: Evidence from panel data. *Economics Letters*, *106*(3), 223–226.

Paserman, D. (2008). Job search and hyperbolic discounting: Structural estimation and policy evaluation. *The Economic Journal*, *118*(531), 1418–1452.

Pilossoph, L., and Ryngaert, J. M. (2025). Job search, wages, and inflation. *Working Paper*.

Raven, J. (2000). The Raven's progressive matrices: Change and stability over culture and time. *Cognitive Psychology*, *41*(1), 1–48.

Richter, D., and Schupp, J. (2015). The SOEP innovation sample (SOEP IS), *Journal of Contextual Economics – Schmollers Jahrbuch*, *135*(3), 389–399.

Sauer, S., Schasching, M., and Wohlrabe, K. (2023). Handbook of ifo surveys. *ifo Beiträge zur Wirtschaftsforschung No. 100*.

Savignac, F., Gautier, E., Gorodnichenko, Y., and Coibion, O. (2024). Firms inflation expectations: New evidence from France. *Journal of the European Economic Association*, 1–34.

Schaefer, D., and Singleton, C. (2023). The extent of downward nominal wage rigidity: New evidence from payroll data. *Review of Economic Dynamics*, *51*, 60–76.

Schmitt-Grohé, S., and Uribe, M. (2016). Downward nominal wage rigidity, currency pegs, and involuntary unemployment. *Journal of Political Economy*, *124*(5), 1466–1514.

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association*, *13*(1), 130–167.

Stüber, H., and Beissinger, T. (2012). Does downward nominal wage rigidity dampen wage increases? *European Economic Review*, *56*(4), 870–887.

Topel, R. H., and Ward, M. P. (1992). Job mobility and the careers of young men. *The Quarterly Journal of Economics*, *107*(2), 439–479.

Weber, M., D'Acunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, *36*(3), 157–184.

Weibull, J. W., Mattsson, L.-G., and Voorneveld, M. (2007). Better may be worse: Some monotonicity results and paradoxes in discrete choice under uncertainty. *Theory and Decision*, *63*, 121–151.

Werning, I. (2022). Expectations and the rate of inflation. *N.B.E.R. Working Paper No.30260*.

Wright, R., Kircher, P., Julien, B., and Guerrieri, V. (2021). Directed search and competitive search equilibrium: A guided tour. *Journal of Economic Literature*, *59*(1), 90–148.

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, *110*(2), 337–61.

Erklärungen

Erklärung zur Verwendung von KI als Hilfsmittel

Hiermit erkläre ich, dass ich bei der Erstellung dieser Dissertation KI-basierte Hilfsmittel verwendet habe, insbesondere verschiedene Versionen von ChatGPT bis einschließlich GPT-5.5 sowie refine.ink. Nach der Nutzung dieser Hilfsmittel habe ich den Inhalt sorgfältig überprüft und überarbeitet. Ich übernehme die vollständige Verantwortung für den gesamten Inhalt dieser Dissertation.

Ort, Datum Gregor Bodo Jost

Eidesstattliche Versicherung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Arbeit selbstständig und ohne unzulässige fremde Hilfe verfasst habe. Ich habe alle in Anspruch genommenen Quellen und Hilfsmittel vermerkt. Die den herangezogenen Werken wörtlich oder sinngemäß entnommenen Stellen sind als solche gekennzeichnet. Diese Dissertation ist weder in der gegenwärtigen noch in einer anderen Fassung oder in Teilen an der Technischen Universität Dortmund oder an einer anderen Hochschule in Zusammenhang mit einer staatlichen oder akademischen Prüfung bereits vorgelegt worden.

Ort, Datum Gregor Bodo Jost