Year : 2023

Three Essays in Labor and Public Economics

Zuchuat Jeremy

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

Three Essays in Labor and Public Economics

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de

Docteur en Économie

par

Jeremy ZUCHUAT

Directeur de thèse
Prof. Rafael Lalive

Jury

Prof. Paul André, président
Prof. Marius Brülhart, expert interne
Prof. Josef Zweimüller, expert externe
Prof. Thomas Le Barbanchon, expert externe

Lausanne
2023

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

Three Essays in Labor and Public Economics

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de

Docteur en Économie

par

Jeremy ZUCHUAT

Directeur de thèse
Prof. Rafael Lalive

Jury

Prof. Paul André, président
Prof. Marius Brülhart, expert interne
Prof. Josef Zweimüller, expert externe
Prof. Thomas Le Barbanchon, expert externe

Lausanne
2023

UNIL | Université de Lausanne
HEC Lausanne
Le Décanat
Bâtiment Internef
CH-1015 Lausanne

IMPRIMATUR

Sans se prononcer sur les opinions de l'auteur, la Faculté des Hautes Etudes Commerciales de l'Université de Lausanne autorise l'impression de la thèse de Monsieur Jeremy ZUCHUAT, titulaire d'un bachelor et d'un master en Économie politique de l'Université de Lausanne, en vue de l'obtention du grade de docteur en Economie.

La thèse est intitulée :

THREE ESSAYS IN LABOR AND PUBLIC ECONOMICS

Lausanne, le 5 juin 2023

La Doyenne

Marianne SCHMID MAST

Members of the thesis committee

Prof. Rafael Lalive

Full Professor, University of Lausanne

Thesis Supervisor

Prof. Marius Brülhart

Full Professor, University of Lausanne

Internal member of the thesis committee

Prof. Josef Zweimüller

Full Professor, University of Zürich

External member of the thesis committee

Prof. Thomas Le Barbanchon

Associate Professor, University of Bocconi

External member of the thesis committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Jérémy ZUCHUAT

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: R. Lalive Date: April 19, 2023

Prof. Rafael LALIVE

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Jérémy ZUCHUAT

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________________ Date: 28 March 2023

Prof. Marius BRÜLHART
Internal member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Jérémy ZUCHUAT

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________________ Date: March 28, 2023

Prof. Josef ZWEIMÜLLER
External member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Jérémy ZUCHUAT

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________________ Date: March 28th, 2023

Prof. Thomas LE BARBANCHON
External member of the doctoral committee

Three Essays in Labor and Public Economics

Jeremy Zuchuat
University of Lausanne

2023

Acknowledgments

My gratitude goes first and foremost to my supervisor, Rafael Lalive. Since we first met in 2016 in the context of his empirical research class, Rafael has taught me how to think critically, how to ask relevant questions, how to conduct rigorous analyses, how to structure and communicate ideas and how to remain connected to policy-relevant matters. In brief, he has instructed me how to become a proficient scientist and researcher in Economics. In addition, Rafael has offered me unique opportunities to enhance my professional experience, whether through teaching or policy consulting. I enormously benefited from his knowledge in that matters, and will surely continue to do so. Most importantly, Rafael has consistently demonstrated himself to be a highly supportive mentor, who cares about the well-being and fulfillment of all members of his research group, to which I had the chance to belong during my time as a PhD student in Lausanne.

Special thanks also go to the members of my thesis committee. I am very grateful to Prof. Marius Brülahrt, internal expert, for his helpful guidance all along the writing of the third chapter of this thesis. His precious advice and challenging questions proved instrumental in structuring and improving this project. I also want to warmly thank Prof. Josef Zweimüller, my co-author on the first chapter: his established expertise and insightful feedback undoubtedly had a significant impact on the amelioration of my work. Finally, I am thankful to Thomas Le Barbanchon, external expert, for all his relevant comments and recommendations that helped me polish the content of my thesis.

I would also like to thank all members of the academic community, who participated to improving the content of this thesis either way. Especially, I would like to express my gratitude to my two co-authors, Prof. Aderonke Osikominu and Lorenzo Pesaresi, for their enriching contribution to our joint project all along its development and realization, and to Prof. Mathias Thoenig and Prof. Aurélien Eyquem, for our pleasant professional collaboration and regular exchanges over the last five years.

As I started my post-graduate studies five years ago, I thought the PhD would be a long and solitary journey. I have never been so wrong! I found myself surrounded by very smart, friendly and stimulating people, without whom my PhD experience would have been so much more monotonous. *Kudos* to Nicola, my co-author, for all those nice years spent in-between Internef, IDHEAP and outdoors: not only does he know the way on our cycling tours, but with his quick and curious mind, he has taught me a great deal on how to be a competent researcher (and on “fiscal perequation” in Bern and Fribourg). I deeply thank Elio and Brendan, my two inseparable Luzerner and Jurassien floor-mates: our shared journey through the MScE did not only lead us to the “fifth”, but also bought us some memorable times whether at local festivities or around Europe. Maxime, Anis and Mathias have taught me that Econ PhD students in Gerzensee do not only have brilliant mind, but that they also excel at table football and Mario Kart. Noëmi and Andrea, have contributed tremendously to making the 556 the coolest office of Internef, where skiing and the outdoors sometimes prevail over work. I am also very grateful to all my friends and colleagues from the “fifth” and UNIL, with whom I have shared countless discussions and good memories, whether around a coffee or a beer: Adrien, Fléance, Pauline, Laia, Guillaume, Jeremias, Tobias, Fabrizio, John, Sofia, Eloi, Paula, Pascal, Alice, Bernhard, Andrea, Christoph, Giacomo, Katja, Resuf, Marius, Noah, and all others I would have forgotten. Many thanks also to Gerzensee crew 2018-2019, that made my time at the doctoral school memorable: Florian, Hyacinthe Georges, Evert, Juliette, Flavia, Tabea, Anja, Pia, Sarah.

I am of course very thankful to my family. My parents Yvette and Philippe, who have always been encouraging and welcoming whenever I felt the urge of recharging my batteries in my home mountains. Thanks to my siblings Elise and Laurent, for our shared journeys at the university in "Lôz" and for their support all over those years. I am also very grateful to all my friends from Valais, Lausanne and elsewhere, for their unwavering support over the last five years. Finally, I would like to thank Estelle, who has accompanied me all along this demanding but fulfilling journey. Not only did she exhibit extraordinary patience and support as I was talking about labor market and taxation, coding at late hours or expressing doubts about my research progress. She also taught me that intelligence is emotional on top of being cognitive and participated deeply to broadening my view on the environment surrounding us, beyond economic-thinking. My gratitude for these years as a PhD student is deeply confounded by the fact that I met her on the first day of this journey.

Contents

Introduction

In modern societies, life is structured around work, with individuals devoting a significant amount of time and energy to their professional life. Considering that work has a fundamental impact on our daily routines, social interactions, and overall lives, understanding how professional choices are made and outcomes achieved in labor markets is essential. The aim of this thesis is to analyze different facets of individuals' working life, with a special emphasis on the impact of public policies on the decisions and outcomes of workers, so as to provide new insights on the design of these polices. Specifically, the three chapters contribute to the bodies of research on the functioning and dynamics of labor markets, the effects of job search assistance programs on job search activity and success, as well as the role of fiscal policy in workers' residential and commuting decisions.

In the first chapter, joint with Rafael Lalive, Aderonke Osikominu, Lorenzo Pesaresi and Josef Zweimüller, we study the dynamics of job search and the causes of duration dependence in job finding, *i.e.* the reasons why unemployed job finding chances decrease with elapsed unemployment duration. Examining this subject is essential to our understanding of labor markets, in particular to determine whether the decline in job finding chances during unemployment is caused by genuine duration dependence in job seekers and/or firms' behaviors, or whether it is mechanically due to some form of dynamic selection.

We analyze job search dynamics empirically, using a granular and comprehensive source of data. Our empirical base material originates from job search diaries filled in by Swiss unemployed, which contain detailed records on all applications they send out, including firms' responses to applications (callbacks to interviews and job offers). On the job seeker side, we find evidence of a slight empirical decline in the number of applications sent out per month, a measure of job seekers' application/search effort (-1 application over 15 months). Accounting for dynamic selection through fixed effects, we find that the net decline in application effort is much more marked (-3 application over 15 months). These results reveal the existence of positive dynamic selection in terms of application effort. Studying firms' responses to applications, we report first-sight evidence of (1) a sharp decline over time in the probability that an application results in a callback (from 5% in month 1 to 2.5% in month 15) and (2) an increase in the probability that a job interview converts into a job offer (from 20% in month 1 to 30% in month 15). Accounting for job seekers' heterogeneity through a control function approach, we find that the genuine decline in callback chances is actually less marked than what raw data suggest (from 5% in month 1 to 3.5% in month 15). This indicates that dynamic selection is negative at the callback phase. In contrast, the duration profile of the job offer conversion probability remains unchanged, even when accounting for individual heterogeneity. All in all, our analysis of firms' responses to job applications suggest that elapsed unemployment duration affects negatively unemployed job seekers' chances at the callback stage of the recruiting process, but not beyond.

We next develop a job search model with statistical discrimination by firms towards long-term unemployed to rationalize our empirical findings. Both sides of the labor market are assumed to be heterogeneous, hence creating room for imperfect information on workers' ability and suitability of workers-firms matches. Just like in our empirical analysis, the model assumes that firms recruit following a two-step procedure: they first decide whether to call back applicants based on imperfect information, before discovering applicants' true ability during costly interviews. Only at that moment do they decide whether to make job offers. In the callback phase, recruiters use elapsed unemployment duration as a signal of workers' productivity. Since high-ability-type job seekers exit the pool of unemployed more rapidly and given that firms are heterogeneous, some recruiters stop calling back job seekers who have been unemployed for too long, as they infer them to be unsuitable matches. This generates negative duration dependence in the callback probability. In contrast, as the pool of re-contacted job seekers get more homogeneous at longer unemployment duration, and since callbacks get more targeted, the probability of a job offer conditional on a job interview increases over time. Finally, as job seekers face declining job finding chances over time, they endogenously respond by lowering their search effort during unemployment.

The second chapter, single-authored, studies a specific job search assistance policy: vacancy referrals by caseworkers. This policy, which consists in the transmission of information on vacancies that are considered as suitable by caseworkers to job seekers, has been proven to be effective at fighting unemployment. However, little is known about the mechanisms through which the policy operates as well as its effects on other dimensions of job finding, such as jobs quality. In this second chapter, we use the same comprehensive and granular source of data as in the first one, to shed light on the effects of caseworker referrals on job search activity and success in Switzerland.

We first analyze the effects of the policy at the application-level, both on applications success and vacancies characteristics. In order to deal with endogeneity issues, we follow a high dimensional fixed effects approach with individual-by-duration fixed effects. This approach boils down to comparing treated and non-treated applications, sent by the same individual, in the same month of unemployment. Therefore, it enables us to easily handle diverse sources of heterogeneity, and solves the issue of assigning the policy to workers based on their expected outcomes in the labor market. Our empirical results reveal that referred job applications face significantly higher success rates compared to non-referred applications, both in terms of callbacks (+ 4 pp. or + 103%) and job offers (+ 1.3 pp. or + 142%). This effect seems to be partially driven by referrals targeting lower-paying positions: referred vacancies are more likely to be part-time and characterized by higher (lower) physical (cognitive) skills requirements.

In a second empirical section, we aggregate our data to the individual-monthly level, to study the effect of referrals on application effort and labor matching. In order to deal with endogeneity issues, we follow an instrumental variable approach based on caseworker stringency. This approach relies on (1) caseworkers' leverage over the use of the referral policy and (2) the quasi-random allocation of job seekers to caseworkers, within Public Employment Service offices. In practice, we build a time-varying leave-one-out stringency instrument, which provides us with an exogenous source of job seekers' exposure to the referral policy. According to our empirical findings, vacancy referrals imply a significant increase in the total application effort exerted by job seekers: one additional referral per month leads to an increase of 0.72 job applications sent. This increase is concurrent to a slight and insignificant decrease in the number of personal applications sent out by job seekers, a result that points towards the imperfect substituability of referred and non-referred vacancies. Applying the same methodology to our labor matching indices (number of job interviews and job offers per month), we find that referrals lead to a large increase in matching opportunities between job-seekers and firms. In light of our previous empirical results, these positive effects are possibly realized at the cost of accessing lower-paying positions.

The third chapter, joint with my fellow PhD student Nicola Mauri, examines another aspect of individuals' working life: residential and commuting decisions, and how those are affected by local fiscal policies. Specifically, we study how decentralized progressive income taxation impacts the location decisions of workers settling down in urban areas, accounting for commuting to work. Studying how individuals of different income levels trade off taxes, housing prices and commuting and whether taxation directly affects their willingness to commute is crucial to understand the causes of spatial income sorting.

In that respect, we develop a structural model of individual residential decisions in a monocentric city with decentralized progressive income taxation. Individuals choose a municipality of residence, which is featured with local level of (residence-based) taxation, housing prices, and commuting distance to the productive center. We explicitly include commuting costs in individuals' optimization program, hence acknowledging that within urban areas location decisions essentially translate into commuting choices. Our model predicts high-income earners to be less sensitive to housing prices compared to taxes and commuting. They are hence predicted to locate in jurisdictions featured with higher housing prices. Similarly, high-income individuals are predicted to be more sensitive to taxes relative to housing prices and commuting, if the progressivity of the tax schedule proves to be sufficiently high. Whether high incomes value centrality more than lower incomes remains ambiguous though. This ambiguity arises from the effect that the progressivity of the income tax schedule has on high incomes' opportunity cost of time, making them more willing to commute if they get compensated with lower taxes.

Based on a random utility framework, we estimate an empirical equivalent to our structural model, to measure the effects of municipality characteristics on individuals' indirect utility along the income distribution. The model is estimated on micro-level data on moving decisions within Switzerland. Those include detailed records of individuals' income, a measure of the location-specific tax burden they face as well as other socio-demographic information. Our estimation results confirm our model's predictions: high-income agents are found to be more sensitive to taxes and less sensitive to housing prices than low incomes. Regarding commuting, we find that high incomes tend to place less value on living in central municipalities. Additionally interacting distance with taxes and income, we show that lower taxation levels directly increase individuals' willingness to commute, consistently with the intuition that local taxes affect the opportunity cost of commuting. Finally, we conduct an illustrative counterfactual exercise where we impose a homogeneous local income tax rate within the urban areas of study. We show that shutting off fiscal decentralization leads to an increase in the demand for central municipalities from high-income movers. Shutting down fiscal federalism would thus entail a gentrification of city centers and a decline in the income levels of peripheral jurisdictions.

Chapter 1

What Drives the Decline in Job Offers?

Jeremy Zuchuat
University of Lausanne

Rafael Lalive
University of Lausanne

Aderonke Osikominu
University of Hohenheim

Lorenzo Pesaresi
University of Zürich

Josef Zweimüller
University of Zürich

Abstract

Job finding is a complex process involving sequential decisions taken by job seekers and firms. The decline in job finding chances over the course of unemployment is well established, but we still know little about the dynamics of these intermediary decisions. We use longitudinal monthly data from job search diaries for several thousand job seekers, which record all applications they send, along with employer callbacks and job offers. Accounting for heterogeneity, we document a net decline in the number of applications and chances of a callback over time, whereas the chances of receiving a job offer after being interviewed are slightly increasing. Individual heterogeneity is found to play a role mostly before the interview stage. Our empirical findings are rationalized in a job search model with statistical discrimination. Firms use elapsed unemployment duration as a signal of applicants' productivity when making callbacks decisions and decide whether to make job offers based on new information revealed during the job interviews. As duration negatively impacts job seekers' chances of finding a job, those react endogenously by lowering their search effort over time.

Keywords: Job search, job finding, duration dependence, dynamic selection, search effort, job application, callback, job interview, job offer.

JEL: J24, J64

1. Introduction

The labor market is characterized by a high degree of informational frictions. In this search and matching context, job seekers and firms do not match directly, but rather encounter each other through a complex job search process, during which information is revealed on both sides of the market. Labor matching has for long been documented as exhibiting negative duration dependence: the longer a job seeker remains unemployed, the lower the probability that she transits from unemployment to employment. Conceptually, part of the decline in job finding chances is due to the direct effect elapsed unemployment duration has on agents' behaviors, both on the supply and demand sides of the labor market. Alternatively, this pattern can be partially explained by workers' heterogeneity and dynamic selection, as unemployed faced with systematically lower employment prospects tend to be over-represented at later stages of unemployment.

Studies that have sought to disentangle the sources of duration dependence in job search have typically focused on the ultimate job finding outcome. However, job finding is far from being a simple one-step process: it involves several distinct stages from the sending of applications by the job seeker, to the decision by the firm whether to make a job offer to the applicant, through the intermediary screening stages. Given the current state of the literature, we lack empirical evidence and theoretical understanding about the effect duration has on these sequential decisions, the role of heterogeneity in their dynamics and how they eventually contribute to the decline in job finding chances.

In this paper, we use granular administrative data on 600,000 job applications made by Swiss unemployed to study the different stages of the job search process. For each of these stages, we document how job seekers or firms' behaviors evolve with respect to elapsed unemployment duration. We also assess the contribution of workers' heterogeneity to the observed duration patterns, and measure the net effect of duration on each of these sequential decisions.

On the firm's side, we find evidence of a large and marked decline in the chances that an application results in a job interview, a "callback", for applications sent out later in the spell, compared to those sent earlier on. To address the compositional change in the pool of job seekers sending applications in late and early duration periods, we control for applications' baseline chances of leading to a job interview, by conditioning on the likely information set that is available to firms when making callbacks decisions. Once controlling for these *ex-ante* chances of leading to an interview, applications sent to firms still face lower callback chances as unemployment duration elapses, but this decline is only about half as strong compared to the observed callback probability. The observed callback probability declines more strongly than the controlled one due to negative dynamic selection, *i.e.* the pool of applications observed at late duration periods have substantially lower *ex-ante* callback chances. We further show that this change in composition arises because job seekers with the highest callback chances tend to exit the observed sample.

We apply the same methodology to the choice of the firm whether to convert a job interview into a job offer, focusing on those applications that first lead to a job interview. In contrast to the callback decision, we find that the job offer conversion decision is positively correlated with duration in the raw data: the probability with which job interviews are converted into job offers is slightly increasing over time. This pattern is rationalizable in a context with dynamic selection : as time passes by, the pool of workers who remain unemployed and who get interviewed becomes more homogeneous, hence leading to a higher probability of securing a job offer conditional on an interview. In line with this argument, we find that observed workers' heterogeneity plays a limited role at this stage of the recruitment process, as controlling for observed unemployed characteristics does not significantly affect the estimated duration profile of the job offer conversion probability.

On the job seeker's side, we find that search effort, as measured by the number of applications sent out per month, decreases slightly and steadily over the course of unemployment. Just like for firms' decisions, this descriptive pattern encompasses both the net effect of duration and the effect of dynamic selection due to workers' heterogeneity. As application effort is observed repeatedly for each job seeker in our data, we can use fixed effects models to control for observed and unobserved individual characteristics. Once heterogeneity is accounted for, we find a much stronger net effect of duration on application effort. Heterogeneity hence entails positive dynamic selection at the application phase: high-application-effort job seekers tend to remain unemployed longer, hence contributing to flattening the duration profile of application effort in the raw data.

Taken altogether, our empirical findings are consistent with a statistical learning view of the labor market, with imperfect information on workers' productivity. We extend Jarosch and Pilossoph (2019) framework, who develop this view for firms, by adding statistical learning on the labor supply side. In this framework, firms use elapsed unemployment duration in addition to observed heterogeneity to infer the productivity of applicants, when deciding whether to call them back for a job interview. Unemployment duration however does not contain additional information and does not reduce applicants' chances at the job offer conversion stage, as their productivity level is revealed during the interview. In turn, the overall decline in firms' responses entails a decrease in application effort by job seekers, who anticipate firms' discriminating behavior towards long unemployment duration in a forward-looking manner. In this framework, net duration patterns on both sides of the market are rooted in statistical discrimination from firms towards longer-term unemployed.

Some of the patterns we document are also coherent with competing theories of job search, such as application targeting, stock-flow matching or social capital exhaustion. However, we find little supportive evidence for these alternative mechanisms. Not only do we observe few changes in application targeting over time, whether in terms of occupational targeting or skills requirements, but we also find little evidence of adjustments in applications quality, notably in terms of the channels used when contacting firms. In particular, applications in person, which are shown to be significantly more successful, represent a relatively stable share of all applications at each unemployment duration.

Our paper addresses various strands of the vast literature on job search and labor markets dynamics. We first contribute to the literature on the role of workers' heterogeneity in explaining duration dependence in job finding. Since the seminal study of Van den Berg (1990), which generalized non-stationarity in job search models, a long standing debate has emerged on whether the decline in job finding chances is due to dynamic selection and workers' heterogeneity, or to genuine duration dependence in agents' behaviors. Focusing on the exit rate out of unemployment, Van den Berg and Van Ours (1996) find that heterogeneity explains most of the decline in job finding, for most socio-demographic groups in the US. Alvarez et al. (2016) find similar results based on a dynamic model of transitions in and out of employment, with arbitrary heterogeneity across workers. The authors estimate their model using social security data for Austrian workers who experience two or more unemployment spells, and find that dynamic selection is a critical source of duration dependence. The same type of approach is taken by Ahn and Hamilton (2020), who find similar results for the US. More recently, Mueller et al. (2021) exploits job seekers' elicited beliefs about job finding to shed light on the sources of negative duration dependence in the job finding rate. They find evidence of substantial heterogeneity across job seekers, resulting into substantial dynamic selection that explains most of the observed decline in job finding chances. Consistent results are found by Mueller and Spinnewijn (2023), who exploit rich administrative data from Sweden to study the predictability and determinants of long-term unemployment. The authors find large amount of heterogeneity in job seekers' employability, and argue that at least half of the decline in job finding over the unemployment spell is driven by dynamic selection.

Our study further contributes to two additional strands of the job search literature, which examine changes in agents' behavior as potential drivers of job finding dynamics. Firstly, we address the literature on firms' decisions during the recruitment process, and how those evolve with respect to elapsed unemployment duration. In their seminal study, Kroft et al. (2013) study duration dependence in the probability that a job application ends up into a callback for a job interview. Based on an experimental audit setup, the authors find that the chances of getting a callback from firms decrease sharply in the early months of unemployment, this decline varying according to local labor market conditions. Their results are consistent with screening models in which employers use unemployment duration as a signal of workers' unobserved productivity (Vishwanath, 1989; Lockwood, 1991). A similar experimental approach has been followed by Oberholzer-Gee (2008), Eriksson and Rooth (2014) and Nüß (2018) for Switzerland, Sweden and Germany respectively. All studies come up with consistent results: the duration of the contemporaneous unemployment spell reduces the probability of being called back by firms. Again, the authors interpret these results as elapsed unemployment duration conveying a stigmatic signal about job seekers' productivity. In contrast, Farber et al. (2016) find no negative effect of unemployment duration on the callback probability for the sub-market of experienced college-educated females applying for administrative support jobs. The authors argue that these results might be due to the specific group of workers they analyse. More recently, Jarosch and Pilossoph (2019) have taken a more structural view at the negative duration dependence in the callback probability. Using a frictional job search model where employers endogenously discriminate against longer-term unemployed, the authors are able to replicate the decline in job interviews over the course of unemployment. Their structural approach allows to further assess the consequences of firms' callback behavior for job finding. They find that the decline in callbacks has a limited impact for unemployment exit, as job interviews lost in the first place would have had little chance to be converted into job offers, and eventually hirings. The authors hence interpret negative duration dependence in the callback probability as being largely driven by dynamic selection and statistical discrimination.

Secondly, we add up to the increasing literature on the dynamics of job search effort provision by job seekers. Recent papers have shown that the decline in search effort might be a major driver of decreasing job finding chances. Exploiting online job application data, Faberman and Kudlyak (2019) study the dynamics of search effort along job search spells in the US. Using the weekly number of applications sent out per month as proxy, they show that search effort is decreasing with the duration of the job search spell. Their main result is robust and even accentuated when controlling for individual heterogeneity through individual fixed effects. Fluchtmann et al. (2021) find consistent results for the universe of Danish job seekers. The authors use information from an online job search monitoring platform, which is legally constraining for unemployed. Their evidence show that the monthly average number of applications recorded on the platform is relatively constant over the duration of unemployment, in the raw data. However, once job seekers' heterogeneity is accounted for through individual fixed effects, application effort is found to decline markedly within unemployment spells. Using similar administrative data, Marinescu and Skandalis (2021) analyse the dynamics of application effort in France, with a specific focus on unemployment benefits exhaustion. Their descriptive analysis shows a net empirical decline in the number of applications sent out per month, both for job seekers that are eligible and non-eligible to unemployment benefits. However, after accounting for compositional changes through spells fixed effects, the authors find limited evidence of decreasing search effort for non-eligible unemployed, whereas application effort is found to be increasing for eligible unemployed until the end of the eligibility period. These mixed results using applications as job search effort proxy echoe those relying on survey data: Krueger et al. (2011) report evidence that the time devoted to job search decreases sharply over the course of unemployment, while DellaVigna et al. (2022) show that search effort is flat in early months of unemployment, increasing before and decreasing after benefit exhaustion.

The contributions of our paper to the existing job search literature are manifold. First of all, we document new duration dependence patterns for the entire sequence of decisions of the job search process, from applications to job offers, through callbacks, all in a unified empirical framework. Especially, we provide unseen empirical evidence on firms' behavior beyond callback decisions, *i.e.* what are the chances of job seekers to obtain a job offer after having been interviewed.

Since our empirical analysis is based on real-world data, we are also able to discuss how heterogeneity and dynamic selection on the one side, and net duration dependence on the other side, contribute to the observed duration patterns. Our results corroborate the previous finding that individual heterogeneity is a major driver of duration dependence in job search. They further highlight the mulit-dimensional role of heterogeneity, which affects differently the sequentaial phases of job search, as well as their dynamics. Specifically, dynamic selection is found to be positive for application effort, leading to an attenuation of its negative duration profile in the raw data. In contrast, dynamic selection turns out to be negative for callback decisions, while it does not play much role at the time when interviews are converted into job offers. In addition, when accounting for heterogenity, we find evidence of a net effect of duration at each phase of the job search process, whether negative for job seekers' application and firms' callback decisions, or positive for job offer conversion.

Our comprehensive empirical evidence allow to better identify which type of model is susceptible to generate the observed duration dependence relationships. The picture that emerges from our empirical exercise is surprisingly consistent with a statistical learning view of the labor market. Our augmented version of Jarosch and Pilossoph (2019), with an additional job application phase, is able to explain most of the patterns we observe empirically. According to this framework, unemployment duration is used by firms at the callback stage to infer applicants' quality, but plays no negative role when firms decide whether to convert interviews into job offers. Faced with declining callback and job offer chances, job seekers endogenously reduce their job application effort, which further dampens job finding chances.

Our findings have important implications for our understanding of labor markets functioning. They suggest that the dynamics in job finding chances can be attributed both to workers' heterogeneity and to the net effect of duration. In particular, the latter is found to be negative for firms' callback decisions and for job seekers' application effort. According to our structural model, these net duration effects are generated endogenously and result from statistical discrimination by firms. In a context with informational frictions with respect to job seekers' productivity, elapsed unemployment duration might have detrimental effects on the two sides of the labor market, if the signal conveyed by duration is an important input of firms' recruitment decisions. From a policy perspective, our study emphasizes the importance of providing better information on job seekers skills and experience already in the early phases of the job search process, so as to reduce the detrimental effects of longer unemployment duration on job seekers' and firms' behaviors, and to improve labor matching efficiency.

The rest of this paper is organized as follows. Section 2 describes the institutional context of our study. In section 3, we present the data we use in our empirical exercise and show how those can be exploited to measure job search outcomes. We also provide descriptive evidence on job offers as proxy for job finding, and discuss how those can be decomposed using the granular information contained in our data. Sections 4 and 5 represent the core of our empirical analysis. In section 4, we study the dynamics of application effort, as measured by the number of applications sent out per month. After presenting descriptive evidence on its decline, we study its dynamics net of compositional effects. We proceed the same way in section 5 for firms' responses, *i.e.* callback and job offer conversion decisions. In section 6, we present a job search model with statistical discrimination that rationalizes the patterns we find empirically. We also discuss alternative mechanisms that might explain our findings. Section 7 concludes.

2. Institutional context

Swiss workers are entitled to unemployment benefits if they contribute at least twelve months within two years prior to the beginning of their unemployment spells.[1] The typical potential benefit duration amounts to 12 or 18 months and is a function of the contribution period, age and family situation of unemployed. The replacement rate ranges from 70% to 80%, depending on the level of the insured salary and the presence of children in the household.

Job seekers who intend to claim unemployment benefits have to register at a regional Public Employment Service (PES) office. Offices are organized at the cantonal level and exert some discretion over the implementation of unemployment policies. Once registered at a regional PES center, unemployed are assigned to a caseworker, either based on caseworkers' caseload, their occupation or at random (Behncke et al., 2010).

One of the main tasks conducted by caseworkers consists in monitoring the unemployed's job search activity. According to the legal framework, unemployment benefit recipients "*must be*

[1]This section overlaps considerably with a similar section in Zuchuat (2023b).

able to demonstrate [their] effort [to find a job]".[2] To make this assessment by caseworkers possible, the unemployed have to document their search activity in job search diaries using pre-defined forms.[3] These forms contain detailed information on all applications made by job seekers, in each month of their unemployment spells. They include information on applications dates, application channels (written, personal or by phone), the work-time percentage of targeted positions (full-time or part-time), whether applications result from caseworker referrals, as well as short job vacancy descriptions. Most importantly, these documents report information on applications' outcomes (callback, job offer, negative or still open), which can be used to quantify applications' success.

Job search diaries are filled in and submitted to PES offices on a monthly basis, together with copies of job applications. These documents serve as a basis for monitoring by caseworkers, who make sure that minimum search requirements are met. Those are defined both in quantitative terms, *i.e.* a minimum number of job applications to be made per month, and in qualitative terms, as caseworkers review copies of job applications to assess their truthfulness (Arni and Schiprowski, 2019).[4] In case of non-compliance with search requirements, job seekers are notified and potentially sanctioned by a benefit cut.[5] Caseworkers and job seekers update information on the success of job applications for up to two months after the job application has been sent out. This provides detailed and accurate information on success of job applications, *i.e.* whether they lead to a callback for a job interview, or a job offer.

3. Data sources

3.1 Data and empirical measurements

Our empirical investigation of job search dynamics relies on various Swiss administrative data sources. Our main source of information stems from search diaries filled in by unemployed at PES. For the purpose of this study, paper-format documents were transcripted into numeric format a two different occasions, based on a stock-flow sampling design.[6] The main large-scale data collection took place between April 2012 and March 2013, in five different Swiss

[2]*Loi fédérale du 25 juin 1982 sur l'assurance-chômage obligatoire et l'indemnité en cas d'insolvabilité (LACI)*; RO 1982 2184. Retrieved 19th October 2022 from https://www.admin.ch/opc/fr/classified-compilation/19820159/index.html.

[3]A copy of the standardized form in French can be found in the Appendix, in Figure A1.

[4]This systematic control by caseworkers aims to prevent over-reporting by job seekers who would over-report to meet search requirements. Under-reporting is for its part unlikely: not only is the cost of reporting an additional application relatively low, but job seekers have few financial incentive to under-report their application effort.

[5]The average size of a sanction amounts to 5.5 days of unemployment benefits, around CHF 900.- on average (Arni and Schiprowski, 2019).

[6]This means that we observe all job seekers that were unemployed at the start of the study, and also all those job seekers who entered unemployment during the observation period.

cantons (Zürich, Bern, Vaud, Zug and St-Gallen), and provides us with our *Main sample* of analysis. This sample covers several hundred thousand job applications and contains most information reported in the job search diaries, with the exception of the information on the posted occupation or on the firm. We supplement this *Main sample* with an *Auxiliary sample*, which originates from a smaller-scale data collection. This one took place from July 2007 to March 2008 in the canton of Zurich only. This additional sample contains all information recorded in the job search diaries, including the occupation targeted by each application and the posting firm. Due to its limited size, the *Auxiliary sample* is principally used in the context of analyses requiring information on occupations. Taken together, these two data sources provide information on job applications and their success at a highly granular level.

We complement search diaries data with information on job seekers' characteristics, which we retrieve from PES registers. Those contain demographic (*e.g.* age, education level, residence status, etc.), job-search related (*e.g.* desired occupation), as well as unemployment institutions-related (*e.g.* caseworker and PES identifiers) information. In addition, we collect information on job seekers employment status and labor market history from social security registers, for the *Main sample* exclusively. This enables us to track job seekers' labor income and unemployment benefits flows before, during and after their unemployment spells. All complementary information are available on an individual-monthly basis and are merged with job search diaries data using individual social security identifiers and calendar months.

We restrict our analysis samples to job applications during months with benefits receipt. This choice is motivated by data reliability: only unemployment benefits beneficiaries have the legal obligation to fill in search diaries, and for those the truthfulness of recorded information is diligently checked by caseworkers. Additionally, we focus on individuals for whom socio-demographic and employment history information are non-missing, these pieces of information playing an important role in our identification strategy. In the end, our *Main sample* of analysis contains 600'323 applications sent by 14'798 individuals, while the *Auxiliary sample* is made of 24'770 applications sent by 655 unemployed.[7]

3.2 Job search outcomes

Job search diaries provide a unique, granular and comprehensive source of information to study the sequential phases of the job search process.[8]

We observe the universe of applications a_{ijt} sent out by job seeker i to job vacancy j in month t of her unemployment spell, where $a_{ijt} = 1$ if the job seeker sends the application, while

[7] Socio-demographic and labor market history-related information on job seekers belonging to each of the two samples are reported in Table A1.

[8] The initial idea of using this type of information to analyse labor market outcomes originate from Falk et al. (2005).

$a_{ijt} = 0$ if she does not apply to the vacancy. The total number of applications made by job seeker i in unemployment month t is then $A_{it} = \sum_{j=1} a_{ijt}$. This provides us with a direct quantitative measure of search or application effort, recently used in the empirical job search literature (Faberman and Kudlyak, 2019; Marinescu and Skandalis, 2021; Fluchtmann et al., 2021). Furthermore, we know firms' responses to each application a_{ijt}, by means of two binary and sequential success indicators. First, we observe whether a job application is followed by a callback for job interview c_{ijt}, a common measure of applications success in audit studies (Oberholzer-Gee, 2008; Kroft et al., 2013; Eriksson and Rooth, 2014; Farber et al., 2016; Nüß, 2018), where $c_{ijt} = 1$ if the job seeker receives a callback and $c_{ijt} = 0$ otherwise. Second, we have access to a never-seen piece of information that goes beyond callbacks: we know whether an application that first led to a job interview eventually ends up in a job offer o_{ijt}, with $o_{ijt} = 1$ if the job seeker receives a job offer for vacancy j.[9] As for applications, we can construct the numbers of callbacks, $C_{it} = \sum_{j=1} a_{ijt} \cdot c_{ijt}$, and the number of job offers $O_{it} = \sum_{j=1} a_{ijt} \cdot c_{ijt} \cdot o_{ijt}$, obtained by individual i in unemployment month t. These monthly aggregates of search activity and firms' responses measure how applications translate into interviews and job offers, and serve as labor matching proxies.

We report descriptive statistics on our empirical job search measures for the *Main sample* in Table 3.1.[10] Panel A reports statistics for application-level outcomes, while panel B focuses on outcomes measured at the individual-monthly level.

As shown in Panel A, the average probability of getting a callback after sending out an application amounts to 4%. This number is slightly lower compared to evidence from audit studies, possibly because of different callback definitions.[11] The probability of obtaining a job offer out of an application amounts to 0.9%. This corresponds to an average job offer conversion probability, *i.e.* the probability of obtaining a job offer conditional on a callback, of 22.5%. Panel B shows that the average monthly number of job applications sent out by unemployed equals 10.55, while the monthly average numbers of callbacks and job offers obtained per month amount to 0.399 and 0.075 respectively. Translated in terms of extensive margins, the probabilities of obtaining at least one job offer or at least one job interview in a given month

[9]In our conceptual framework, the two success indicators are sequential, *i.e.* a job offer can only occur conditional on a callback. This sequence is sometimes not verified empirically, *i.e.* a job offer is recorded without a preceding callback. This might for instance be the case if the interview and job offer took place within the same month of unemployment (inbetween two meetings at the PES office), and the caseworker only requires the job seeker to fill in the job offer tickbox. In those cases, we impute a job interview.

[10]Corresponding figures for the *Auxiliary sample* can be found in the appendix, in Table A2.

[11]In audit studies, callbacks are sometimes refer to as any reply from firms, from "asking for additional information" to "inviting for a job interview". In our case, a callback is registered only when firms invite the applicants for a job interview.

Table 3.1: Descriptive statistics, *Main sample*

	Mean	SDV	Min	Median	Max	N
A. By application						
$\mathbb{P}(c_{ijt} = 1)$, callback prob. [in %]	4.013	19.626	0.000	0.000	100.000	600323
$\mathbb{P}(o_{ijt} = 1)$, job offer prob. [in %]	0.905	9.468	0.000	0.000	100.000	600323
$\mathbb{P}(o_{ijt} = 1 \mid c_{ijt} = 1)$, job offer conversion prob. [in %]	22.515	41.769	0.000	0.000	100.000	22422
B. By monthly-individual						
A_{it}, nbr. applications	10.553	4.698	1.000	10.000	30.000	58755
C_{it}, nbr. callbacks	0.399	0.961	0.000	0.000	9.000	58755
O_{it}, nbr. job offers	0.075	0.334	0.000	0.000	9.000	58755
$\mathbb{P}(C_{it} > 0)$, prob. a.l. one interview [in %]	22.551	41.792	0.000	0.000	100.000	58755
$\mathbb{P}(O_{it} > 0)$, prob. a.l. one job offer [in %]	6.108	23.947	0.000	0.000	100.000	58755
C. Sample structure						
Time-period			04.2012 - 03.2013			
Region			BE, SG, VD, ZG, ZH			
Nbr. applications			600323			
Nbr. monthly-individual			58755			
Nbr. individuals			14798			

Note: This table reports descriptive statistics about our *Main sample* of study. Panels A and B report descriptives on application-level and individual-monthly-level job search outcomes respectively. Panel C provides information about the sample structure.

of unemployment are equal to 22.6% and 6.1%.[12]

3.3 Decomposing job offers using job search diaries information

Job search diaries are designed to provide a reliable and in-depth description of the job search process leading to job offers. Based on search diaries, we identify all individuals who receive at least one job offer out of the applications they send in month t, or $\mathbb{1}(O_{it} > 0)$. We define the (empirical or theoretical) expectation of this outcome for individuals who are still applying for jobs in month t of unemployment as $\mathbb{E}_t\big[\mathbb{1}(O_{it} > 0)\big] = \mathbb{P}_t(O_{it} > 0)$. This expression defines the probability that an individual still unemployed in unemployment month t receives at least one job offer in that same month.

Existing empirical analyses focus on duration dependence in the job finding rate (or unemployment-to-employment transition rate). Conceptually, this one is closely related to the probability of obtaining at least one job offer, computable from job search diaries data. This can be seen in Figure 3.1, where the probability $\mathbb{E}_t\big[\mathbb{1}(O_{it} > 0)\big]$ is plotted together with the average

[12] Summary statistics obtained on the *Auxiliary sample* are qualitatively similar, even though the callback and job offer probabilities (and consequently the numbers of callbacks and job offers) are relatively higher. These discrepancies might be due to differences in data recording across institutions, local labor market conditions or macroeconomic conditions at the time of data collection.

Figure 3.1: Monthly probability of at least one job offer and job finding rate

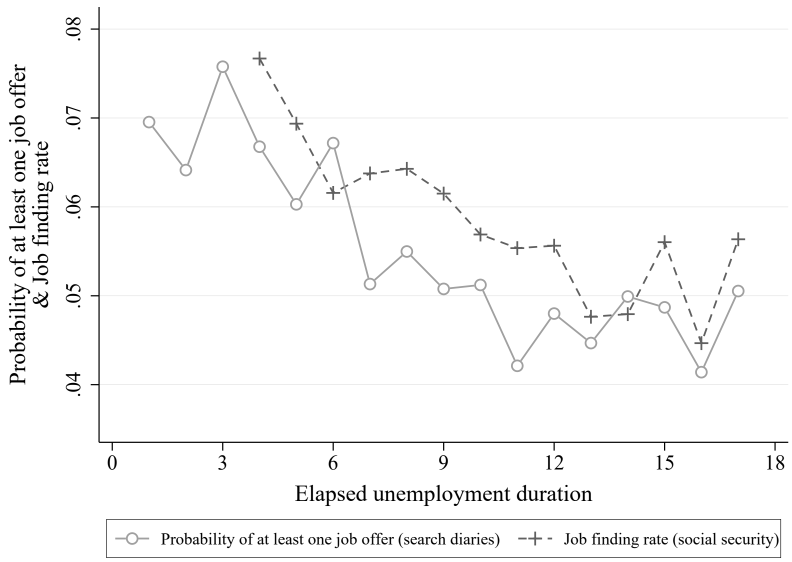

Note: This figure depicts the empirical duration dependence in the monthly probability of obtaining at least one job offer (computed on search diaries data) and the monthly job finding rate (computed on social security data).

monthly job finding rate.[13] The job finding rate exhibits a typical negative duration dependence behavior: in early months of unemployment, the instantaneous chances of transiting from unemployment to employment are relatively high (around 8%), before decreasing to levels close to 4-5% after 12 months. The probability of getting at least one job offer follows a similar dynamics. However, the two curves do not overlap perfectly, the job finding rate typically corresponding to a right-shift of the job offer curve. This pattern is due to the timing of data recording: job search diaries report applications dates and not dates at which job offers are made. Given that recruiting processes can expand over a certain period and that job starting dates might be delayed, such shift between the two curves is expected. As an additional check, we proceed to an event-study that tracks the evolution of monthly labor income flows after the recording of a job offer. Corresponding results are reported in Figure A2 in the Appendix and show that job offers are strongly predictive of increases in labor income.[14]

[13] The job finding rate is computed using social security data, that are only available for the *Main sample*. Conceptually, the job finding rate is defined using a binary indicator which takes the value 1 if job seeker i leaves unemployment for a job after t months of unemployment, and zero otherwise. In our context, a job seeker is considered as having found a job if her monthly labor earnings exceed 2,000 CHF. This is equivalent to 50% of the unofficial minimum wage in Switzerland (4,000 CHF monthly). In most cases, the job finding rate and the probability of getting at least one job offer coincide, with a certain time-lag.

[14] Figure A2 shows that labor income increases after the recording of the last job search diary, even in the absence of a job offer. This pattern is expected, given that we observe social security data up to 2015, while job search diaries were only collected until March 2013. Put differently, the graph also shows the evolution of labor income flows for individuals who have been unemployed after March 2013, a period during which we did not collect

We show above that job search diaries provide reliable information on job offers, which are predictive of job finding. The innovative aspect of our data is that they do not restrict us to the sole analysis of the final outcome of job search. Their granularity allows to analyze duration dependence at various stages of this process, and facilitates the assessment of the contribution of heterogeneity to this phenomenon.

At a basic level, the probability of obtaining at least one job offer (and the job finding rate eventually) are driven by the number of job offers obtained per month, O_{it}.[15] The latter depends on the number of job applications sent in a month multiplied by the average success probability of job applications sent in that month. In our context, a job application is successful if it first leads to a job interview, and then to a job offer. Using our notation above, the expected number of job offers for individuals sending applications in month t is[16]

$$\mathbb{E}_t\left[O_{it}\right] = \mathbb{E}_t\left[A_{it} \cdot \mathbb{P}(o_{ijt} = 1)\right] = \mathbb{E}_t\left[A_{it} \cdot \mathbb{P}(c_{ijt} = 1) \cdot \mathbb{P}(o_{ijt} = 1 | c_{ijt} = 1)\right]. \tag{3.1}$$

This simple conceptual framework highlights three stages of the process which leads to a job offer: job seekers' application effort (measured by the total number of applications sent to firms in a month, A_{it}), the probability that an application ends up in an interview (or callback probability, $\mathbb{P}(c_{ijt} = 1)$), and the chances of having a job offer if interviewed (or job offer conversion probability, $\mathbb{P}(o_{ijt} = 1 | c_{ijt} = 1)$). All three steps of this process are observed empirically thanks to the job search diaries data, as described schematically in Figure A4, in the Appendix.

We further decompose the empirical decline in job offers between month t and the initial month 0 as follows

$$\mathbb{E}_t[O_{it}] - \mathbb{E}_0[O_{i0}] = \underbrace{\mathbb{E}_t[O_{it} - O_{i0}]}_{\text{Net duration effect}} + \underbrace{\mathbb{E}_t\left[O_{i0}\right] - \mathbb{E}_0\left[O_{i0}\right]}_{\text{Compositional change}}. \tag{3.2}$$

This standard equation states that the observed decrease in job offers is made of two components. First, the decline in job offers for the set of job seekers who are unemployed in month t, $\mathbb{E}_t\left[O_{it} - O_{i0}\right]$, which we refer to as the net duration dependence. Second, the change in the composition of the pool of individuals looking for a job between t and 0, $\mathbb{E}_t\left[O_{i0}\right] - \mathbb{E}_0\left[O_{i0}\right]$. This compositional change creates the challenge for empirical identification of net duration dependence.

any diary, and who eventually found a job.

[15] Figure A3 in the Appendix shows that the empirical average number of job offers per month $\mathbb{E}_t\left[O_{it}\right]$ is closely related to the empirical monthly probability of obtaining at least on job offer $\mathbb{E}_t\left[\mathbb{1}(O_{it} > 0)\right]$, which in turn directly affects the job finding rate.

[16] The sequential aspect of the process implies that $\mathbb{P}(o_{ijt} = 1 | c_{ijt} = 0) = 0$.

Information on job offers O_{it} alone is not sufficient to fully understand how duration dependence and compositional changes participate to the reduction in job finding chances. Although, search diaries data allow dissecting the reasons for reduced labor market matching success in a uniquely powerful manner. To see why, let us rewrite equation (3.2) as

$$\begin{aligned}\mathbb{E}_t\big[O_{it}\big] - \mathbb{E}_0\big[O_{i0}\big] = \; & \mathbb{E}_t\big[(A_{it} - A_{i0}) \cdot \mathbb{P}(o_{ijt} = 1)\big] \\ & + \mathbb{E}_t\big[A_{i0} \cdot \big(\mathbb{P}(o_{ijt} = 1) - \mathbb{P}(o_{ij0} = 1)\big)\big] \\ & + \mathbb{E}_t\big[A_{i0} \cdot \mathbb{P}(o_{ij0} = 1)\big] - \mathbb{E}_0\big[A_{i0} \cdot \mathbb{P}(o_{ij0} = 1)\big], \end{aligned} \tag{3.3}$$

where $\mathbb{P}(o_{ijt} = 1) = \mathbb{P}(c_{ijt} = 1) \cdot \mathbb{P}(o_{ijt} = 1 | c_{ijt} = 1)$. The decline in job offers is now expressed as a sum of products between applications sent by job seekers, A_{it}, and the probability that applications result in a job offer, $\mathbb{P}(o_{ijt} = 1)$. This latter component can again be decomposed into the product of two probabilities characterizing firms' sequential decisions, *i.e.* callback and job offer conversion decisions. Learning about duration dependence hence requires understanding how job seekers' application effort and firms' responses to applications evolve along unemployment.

Going forward, in Section 4, we will discuss an empirical approach to model job seekers' decision, who control the number of job applications they send out per month. Specifically, we study the dynamics of application effort, A_{it}, with respect to elapsed unemployment duration. Since we observe applications for each job seeker repeatedly, we address individual (un-)observed heterogeneity with fixed effects models. Thanks to this powerful approach, we are able to uncover the genuine decline in applications for a fixed set of individuals, *i.e.* those still sending applications in month t, *i.e.* $\mathbb{E}_t\big[A_{it} - A_{i0}\big]$. We can also readily discuss compositional changes between the pool of job seekers looking for a job in month t compared to the initial period, *i.e.* $\mathbb{E}_t\big[A_{i0}\big] - \mathbb{E}_0\big[A_{i0}\big]$. This first analysis provides us with a unique opportunity to discuss the role of unobserved individual heterogeneity in shaping the first stage of the job offer arrival process.

In Section 5, we take firms' perspective and study their responses to applications. Addressing heterogeneity and duration dependence in the context of firms' decisions is more challenging. Positive values of c_{jit} and o_{jit} are not repeatedly observed along unemployment spells, but are rather concentrated at their end. This data specificity prevents us from adopting the same fixed effects approach as for applications.

To study how the probability of calling back an applicant, $\mathbb{P}(c_{jit} = 1)$, changes with elapsed unemployment duration, we rely on the nature of this first decision by the firm. This one is arguably based on job seeker's observed characteristics, but the recruiter may use elapsed unemployment duration to infer the likely unobserved characteristics of the applicant, as models of statistical discrimination would imply (Jarosch and Pilossoph, 2019). To deal with individual heterogeneity, we condition our estimates on the *ex-ante* probability of each application to end up in a positive callback. This application-specific *ex-ante* chance is computed using the information set observed by firms when making callback decisions in the very early stage of the unemployment spell and aims to control for endogeneity of sample selection, in the spirit of a control function approach (Matzkin, 2003). As an alternative approach, we also directly condition on a large set of characteristics that are observed both to firms and us. We then examine the net duration dependence in the callback probability by contrasting the chances of callbacks for applications sent in month t compared to the chances these same applications would have had, when sent in month 0, *i.e.* $\mathbb{E}_t\big[\mathbb{P}(c_{jit}=1)-\mathbb{P}(c_{ji0}=1)\big]$. Conversely, we study how the composition of job applications changes by contrasting the *ex-ante* chances of a callback between applications sent in month t and month 0, *i.e.* $\mathbb{E}_t\big[\mathbb{P}(c_{ji0}=1)\big]-\mathbb{E}_0\big[\mathbb{P}(c_{ji0}=1)\big]$. Finally, we analyze the dynamics of the job offer conversion probability, $\mathbb{P}(o_{ait}=1|c_{ait}=1)$. We address heterogeneity through the same control function approach as in the callback phase. Importantly, the decision whether to make a job offer to the interviewee differs conceptually from the callback decision, as the interview provides new information on the suitability of the applicant for the position. Any information already known at the callback stage is therefore likely to play a less important role at the job offer conversion stage.[17]

4. Job seekers: job applications

We first study how application effort changes over the course of unemployment. After providing descriptive evidence on job applications dynamics, we present our identification strategy to measure net duration dependence in the number of applications, accounting for individual heterogeneity, before discussing our main results.

4.1 Descriptive analysis

Figure 4.1 describes the change in the average monthly number of job applications A_{it} with respect to elapsed unemployment duration, in red. On average, job seekers send around 11 applications in their first month of unemployment, a number that decreases down to 9.75 after fifteen months spent unemployed. These *prima facie* evidence suggest that the number of applications sent by job seekers is slightly decreasing with respect to elapsed unemployment duration. However, since applications are observed on a rapidly changing pool of job seekers, this pattern encompasses both dynamic selection and the net effect of duration. To visualize the role played by individual heterogeneity in the raw-data pattern, we plot the same relationship for different subsamples, defined according to the (individual) maximal

[17] Ideally, we would also condition on information obtained during the job interviews, but those remain unobserved to us.

Figure 4.1: Empirical duration dependence in application effort

Note: This figure describes the empirical duration dependence in application effort, measured by the monthly number of job applications A_{it}. 95%-confidence intervals are reported. The graph also depicts the average duration profiles in A_{it} for spells subsamples, defined based on the maximal unemployment duration observed for each spell (1-3, 4-6, 7-9, 10-12, 12-15, > 15 months).

unemployment duration observed in the sample.[18] Corresponding duration profiles are also reported in Figure 4.1, as gray lines.

Two key facts emerge from this graphical subsample analysis, in line with previous findings in the literature (Faberman and Kudlyak, 2019; Fluchtmann et al., 2021). First, there seems to exist differences in levels between individuals who remain unemployed for a short and long period of time: job seekers who are observed at later stages of unemployment tend to write more applications, at any duration. Second, when accounting for differences in levels, the net duration profiles computed on the various subsamples appear to be steeper than their counterpart based on the full sample, and essentially parallel.[19] Taken altogether, these descriptive patterns suggest that the monthly number of applications sent out by unemployed declines more strongly within individuals than across unemployed, indicating that the net effect of duration on application effort is potentially stronger that what raw data suggest.

[18]The subsamples are defined according to the following within-spells maximal unemployment duration intervals: 1-3, 4-6, 7-9, 10-12, 12-15 and > 15.

[19]Faberman and Kudlyak (2019) find a similar pattern for applications in the context of an online job search platform.

4.2 Empirical approach

We develop an empirical strategy to identify the net effect of duration on application effort and to assess how individual heterogeneity contributes to its empirical decline. Exploiting the longitudinal aspect of our data, we follow a within-estimation approach with fixed effects at the individual level. Our baseline specification writes as

$$A_{it} = \alpha_i + f^A(t;\phi^A) + X_{it}\beta + \delta_{mk} + \varepsilon_{it}, \tag{4.1}$$

where i stands for individuals and t for elapsed unemployment duration. The function $f^A(t;\phi^A)$ corresponds to the parametric estimate of the net effect of duration on the monthly number of applications. This effect is estimated net of individual observed characteristics X_{it} and local labor market conditions (in occupational sector m and calendar quarter k).[20] Moreover, it is obtained conditional on individual fixed effects α_i. ε_{it} represents an idiosyncratic error term.

The main strength of our specification lies in the within-individual identification of the net duration effect. The set of individual fixed effects α_i controls for any form of (time-constant) observed and unobserved individual heterogeneity. This rules out spurious duration dependence generated by dynamic selection based on hard-to-quantify individual characteristics, such as job seekers productivity, labor market history, professional network and intrinsic motivation. This approach has already been applied for application effort by recent studies (Faberman and Kudlyak, 2019; Marinescu and Skandalis, 2021; Fluchtmann et al., 2021) and delivers reliable estimates of net duration dependence when the dependent variable is not directly related to exits from the observed sample (Zuchuat, 2023a).

Our approach for applications also resembles Mueller and Spinnewijn (2023), who discuss the role of heterogeneity in shaping long-term unemployment on the basis of observed pre-unemployment characteristics of individuals. We can go, however, beyond their setting by focusing on applications, which are repeated within individuals. This provides us with a means to discuss the role of observed and unobserved heterogeneity at the application phase, while they assume unobserved heterogeneity to be orthogonal to observed heterogeneity.

4.3 Results

4.3.1 Main results

We report step-by-step estimates of equation (4.1) using OLS in Table 4.1, where the net effect of duration is specified linearly, *i.e.* $f^A(t;\phi^A) = \phi^A t$. Standard errors are clustered at the individual level and reported in parentheses. Coefficients in relative terms are reported in

[20] As a matter of simplification, indices m and k are omitted for other elements than δ_{mk} in the regression equation.

Table 4.1: Duration dependence in application effort, linear specification

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable : application effort A_{it}						
Elapsed unemployment duration	-0.078***	-0.053***	-0.035***	-0.040***	-0.214***	-0.217***
	(0.008)	(0.008)	(0.007)	(0.007)	(0.010)	(0.021)
	[-0.718%]	[-0.487%]	[-0.326%]	[-0.367%]	[-1.976%]	[-2.003%]
Individual controls	No	Yes	Yes	Yes	No	Yes
Policy controls	No	No	Yes	Yes	No	Yes
LLMC	No	No	No	Yes	No	Yes
Individual FE	No	No	No	No	Yes	Yes
Mean outcome 1^{st} month	10.846	10.846	10.846	10.846	10.846	10.846
adj.-R^2	0.005	0.038	0.179	0.192	0.486	0.498
N. observations	58755	58755	58755	58755	58755	58755

Note: This table reports empirical estimates of equation (4.1) using OLS, where the parametric duration function $f^A(t;\phi^A)$ is specified linearly. Each column sequentially adds a set of controls or FE. Errors are clustered at the individual level and reported in parentheses. Coefficients in relative terms (standardized with respect to the average in the first month of unemployment) are indicated in squared brackets. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

squared brackets. Column (1) reports estimates from a bivariate model, where application effort is regressed on elapsed unemployment duration only. The estimated duration coefficient is negative and strongly significant, just like descriptive evidence suggest. Adding sequentially individual, policy and local labor market conditions controls in columns (2) to (4) tends to attenuate the duration parameter value, which remains negative and strongly significant though. In column (5), we regress the monthly number of applications on elapsed unemploy-

Figure 4.2: Duration dependence in application effort, saturated specification

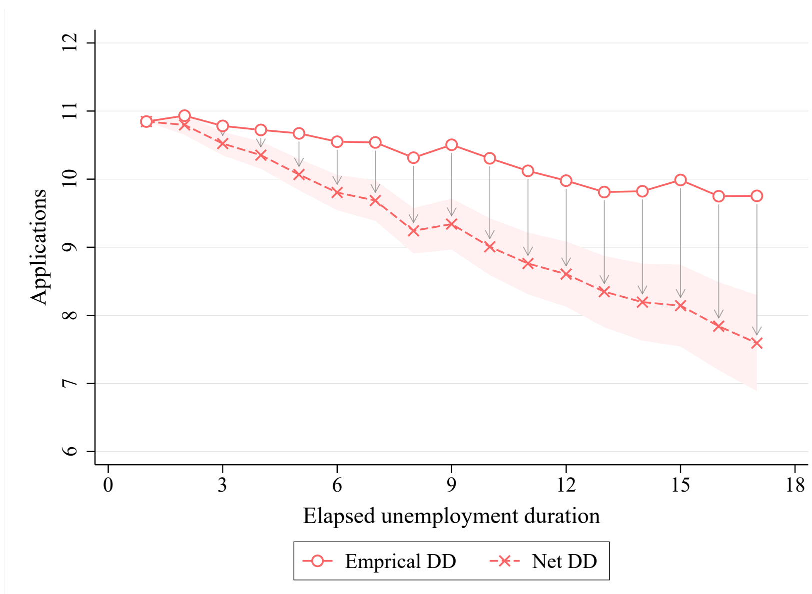

Note: This figure reports empirical estimates of equation (4.1), where the parametric duration function $f^A(t;\phi^A)$ is specified in a saturated manner. Estimated are obtained based on the full specification, including controls and individual fixed effects. 95%-confidence intervals based on standard errors clustered at the individual level are reported. The average empirical duration dependence pattern in A_{it} is also reported, as a solid line.

ment duration, including individual fixed effects only. In contrast to controlling for observed characteristics, this leads to a dramatic increase in the estimated duration parameter, in absolute terms. Further adding individual controls in column (6) virtually does not affect the estimated value of the interest coefficient: according to our full specification, one additional month of unemployment leads to a decrease of -0.22 applications per month on average, or about one application less after 5 months.

We assess non-linearity in the net duration dependence in application effort by estimating our model with a saturated specification of $f^A(t;\phi^A)$, *i.e.* one dummy per month of elapsed unemployment. Corresponding results are reported graphically in Figure 4.2. They show that the decrease in the number of application due to elapsed unemployment duration is essentially linear.

The above results emphasize the crucial role played by individual unobserved heterogeneity and dynamic selection in the decline of application effort with respect to elapsed unemployment. At the application phase, dynamic selection is positive: job seekers who remain unemployed for a longer period of time send systematically more applications, at any duration. Heterogeneity hence tends to attenuate the net effect of duration on application effort in the raw data; once it is accounted for, applications' duration profile becomes much steeper. To further explore this point, we plot the distribution of estimated $\hat{\alpha}_i$ in Figure 4.3. This graph shows that the distribution of $\hat{\alpha}_i$ shifts rightwards as we consider individuals who leave unem-

Figure 4.3: Unobserved heterogeneity in application effort, empirical distribution

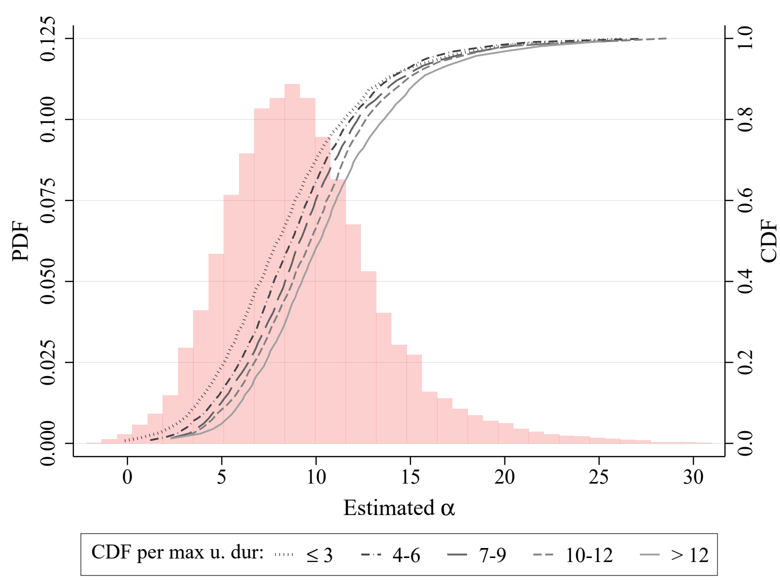

Note: This figure plots the distribution of the estimated α_i in equation (4.1). The global empirical density is represented by the histogram in red. The cumulative distributions for different subsamples, defined based on the maximal unemployment duration observed per spell, are reported as gray lines.

ployment after 12 months or longer, compared to the fixed effect of individuals who leave in the first three months, hence confirming positive dynamic selection with respect to application effort.[21]

To better understand who hide behind those longer-term high-application-effort individuals, we regress estimated $\hat{\alpha}_i$ on observed individual characteristics. Partial correlation coefficients are reported in Table B1 in the Appendix. They reveal that job seekers who provide more application effort are more likely to be women, younger, not to be native and not to hold a tertiary education degree. They also tend to have higher past labor earnings and are less likely to have experienced unemployment previously. The table also reveals the existence of substantial variation in job application effort according to occupational sectors and local labor market institutions (canton, PES or caseworkers, depending on the specification).

4.3.2 Robustness

We proceed to several robustness checks, to assess the validity of our baseline findings of a negative net effect of elapsed unemployment duration on application effort.

First, we consider an alternative model specification. Given the count data nature of the dependent variable, we estimate a Poisson-pseudo maximum likelihood model with fixed effects. Corresponding results are reported in Table B2 in the Appendix, and are very close to our baseline OLS estimates, both qualitatively and quantitatively. Specifically, accounting for unobserved heterogeneity through fixed effects consistently leads to a marked steepening in the estimated effect of duration (from a semi-elasticity of -0.9% to -2.1%).

Second, we consider alternative measures of application effort. In Switzerland, job search effort is monitored by caseworkers based on job search diaries and search requirements. The latter are defined in terms of minimal number of job applications to be sent per month. As a result, it is common to observe some bunching around the standard minimal search requirements values, $\underline{A} = 8, 10$.[22] Also, some applications might not directly result from job seekers' private search activity, but rather from intervention by caseworkers. For instance, caseworkers may refer job seekers to apply to jobs (Zuchuat, 2023b). For that reasons, it might be argued that the total number of applications sent per month, A_{it}, does not validly measure application effort. As a robustness check, we re-estimate our model using alternative search effort measures as dependent variables: excess application effort $\bar{A}_{it} = \max(0, A_{it} - \underline{A})$, above the search requirements thresholds $\underline{A} = 8, 10$, and the monthly number of applications which do not result from caseworker referrals. Corresponding estimates are reported in Table B3

[21] Positive dynamic selection at the application phase can also be seen in Figure B1, which plots the average estimated $\hat{\alpha}_i$ in Figure 4.3 for all individuals observed at duration t.

[22] This point can be observed in Figure B3, where we plot the empirical distribution of application effort. Modes are typically observed at $A_{it} = 8$ and 10.

and are very much in line with our baseline findings.
Third, we discuss the existence of a potential within-estimation duration bias in our baseline estimates. As shown in Zuchuat (2023a), using fixed effects models to estimate duration dependence relationships with data subject to attrition might entail a strong bias in the estimated duration parameters. This is notably the case if the dependent variable is closely related to the attrition mechanism, as this generates a mechanical and undesirable correlation between the within-time regressor and the within-error term. Such situation shall not occur in our context, given that applications are observed repeatedly within unemployment spells, and do not directly translate into exits from the observation sample. As an additional check, we re-estimate our baseline specification on a subsample that does not include the last observation of each non-right-censored spell, *i.e.* on individual-monthly observations that are not contemporaneous to an unemployment exit. Corresponding estimation results are reported in Table B4 and turn out to be highly similar to our baseline estimates.
Fourth, we examine whether duration patterns in application effort vary across subgroups of unemployed. If job search behavior differs across different types of unemployed, our pooled estimates of the net duration effect would possibly be biased. We re-estimate equation (4.1) on different sub-samples, based on various sample-split variables. Corresponding results are reported in Figure B4 in the Appendix. They reveal that the net effect of duration on application effort is consistently found to be negative across most sub-groups.
All in all, this first empirical section emphasizes the existence of a sizeable net effect of duration on job application effort. Such pattern might arise from stock-flow job postings (Salop, 1973), discouragement (Falk et al., 2006), or changes in application strategy (Galenianos and Kircher (2009); Wright et al. (2021); Lehmann (2023)). Our reduced-from analysis does not enable us to determine directly which exact mechanism is at play in our context. Nevertheless, studying how firms react to applications over time might help us to pine down which explanation is susceptible to explain the observed dynamics in job application effort.

5. Firms' responses: callbacks and job offers

We now study how firms' responses to job applications change over the course of unemployment, exploiting our data at the application level. We first present descriptive evidence about the dynamics of the callback and job offer conversion probabilities, before turning to our empirical approach to disentangle the net effect of duration from dynamic selection.

5.1 Descriptive analysis

Our conceptual framework formalizes firm's recruitment process as two sequential decisions. First, the firm chooses whether to call back an applicant for a job interview; second, conditional on a callback, it decides whether to make a job offer to the interviewee. The empirical relationships between those two decisions and elapsed unemployment duration are depicted in Figure 5.1.

Figure 5.1: Empirical duration dependence in firms' responses

(A) Callback probability

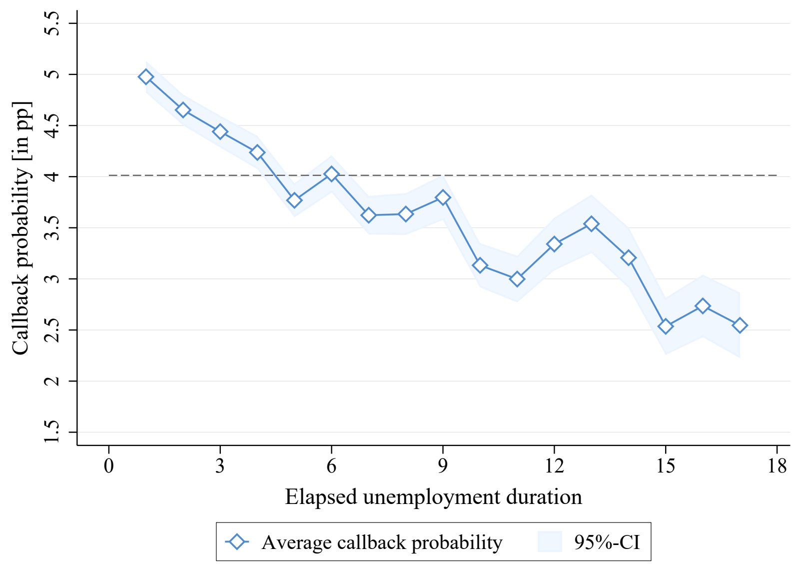

(B) Job offer conversion probability

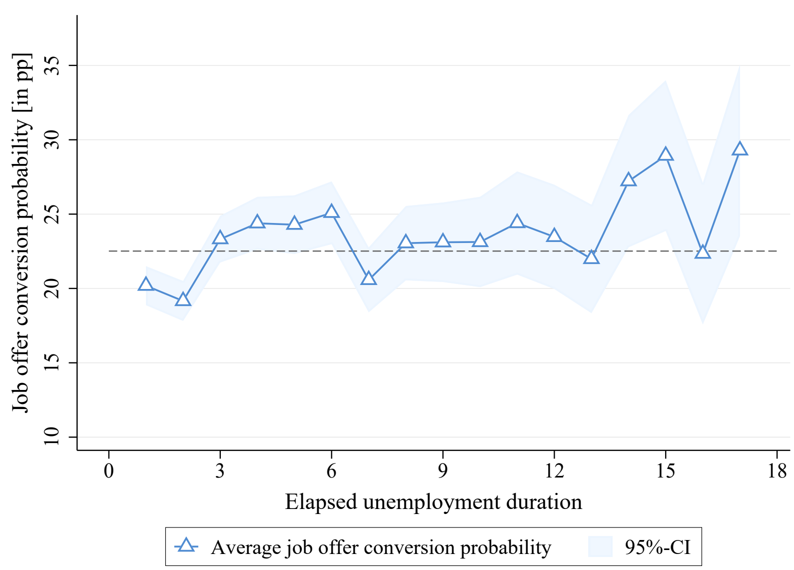

Note: These graphs depict the empirical duration patterns in the callback probability (panel A) and in the job offer conversion probability (panel B). Panel A is based on all applications, while panel B is based exclusively on applications that previously led to a job interview. Application-level observations are weighted according to the inverse of the number of applications sent by individual i in month t, so as to put equal weight on all individual-monthly units.

Panel A represents the average (application-level) callback probability for each month of elapsed unemployment, computed on all applications observed in the sample. The graph emphasizes a substantial decrease in the chances of getting a positive reply from firms at the first stage of the recruitment process, from 5% in the first month to 2.5% after fifteen months. This pattern echoes results from audit studies, which finds evidence of negative duration dependence in the callback probability (Oberholzer-Gee, 2008; Kroft et al., 2013; Eriksson and Rooth, 2014; Nüß, 2018). Unlike experimental audit studies, we observe this pattern in an empirical setup; just like for application effort, it might capture some form of dynamic selection. This empirical aspect of our data represents a valuable addition to audit studies, as it enables us to examine the role played by individual heterogeneity in the decreasing chances of going through the first stage of firms' recruitment process.[23]

In panel B, we report descriptive evidence of duration dependence in the next step of the hiring process. The graph plots the average monthly (application-level) probability with which callbacks are converted into job offers, for applications that previously led to job interviews. In contrast to callback decision's dynamics, the empirical duration profile of the job offer conversion probability is non-negative, and even slightly increasing. In early months of unemployment, 20% of callbacks are converted into job offers. From month three onwards, the job offer conversion probability stabilizes around 25%, before reaching 30% after fifteen months. In spite of this slight positive duration dependence in the job offer conversion probability, we still observe a strong negative relationship between the application-level probability of getting a job offer (out of sending an application) and elapsed unemployment duration, as shown in Figure C2 in the Appendix.

Taken altogether, these *prima facie* evidence suggest that elapsed unemployment duration enters firms' decision process negatively mostly at the callback stage of the screening process. In contrast, duration does not seem to reduce unemployed chances of obtaining a job offer, once interviews have taken place. These evidence are in line with previous findings from experimental studies, but are novel in important dimensions. Given that our data are not limited to the firm's first response by design, we are able to discuss the dynamics of firms' behavior beyond the mere job interviews. Moreover, contrary to audit studies, which are based on standardized fake applications, our results are obtained in a real-life setup. Consequently, they encompass both the net effect of duration and dynamic selection, which we seek

[23] Another difference between the pattern we observe in our study and those from previous related work lies in the shape of the callback probability duration profile. In their seminal study, Kroft et al. (2013) find evidence of a sharp decline in callback chances in early months of unemployment, typically up to month six. In our context, we observe a gradual, almost linear decrease in the callback probability with respect to elapsed unemployment, up to month eighteen. This discrepancy could be explained by differences in institutional settings, notably in terms of potential duration of unemployment benefits (six months in the US, eighteen months in Switzerland). This interpretation is corroborated by evidence of a flattening in the callback probability for month at or beyond benefits exhaustion (see Figure C1 in the Appendix).

to disentangle in the following.

5.2 Empirical approach

Distinguishing the role played by individual heterogeneity and dynamic selection from the net effect of duration for firms' responses requires an alternative identification strategy. Applying the same fixed effects approach as for application effort would indeed be misleading. Since the dependent variables, either callbacks or job interviews, represent direct proxies for unemployment exits and sample attrition, the fixed effects estimator of duration parameters is subject to a within-estimation duration bias (Zuchuat, 2023a). Intuitively, as positive values of c_{ijt} and o_{ijt} tend to be concentrated at the end of the spells (see Figure C3 in the Appendix), positive realizations of the within-error term are more likely to occur for positive values of the within-time regressor. Consequently, using the fixed effects approach for this type of outcomes entails a sizeable positive mechanical correlation between the error term and time regressor, which translates into a large bias in the estimated net duration parameters.

Our alternative identification strategy to measure the net duration profiles of firm's decisions is based on characteristics of individuals and job applications which are known by the firm prior to its callback decision. Inspired by Mueller and Spinnewijn (2023), we exploit our detailed application-level data to condition on the same set of information a firm has on an applicant, when it receives her application and decides whether to call her back for an interview. Specifically, we use our rich data to construct an index capturing the *ex-ante* propensity that an application sent out early during the unemployment spell receives a positive response from the firm. We include this index in the specification of the callback and job offer conversion probabilities to control for dynamic selection when measuring net duration dependence in firms' responses, in the spirit of a control function approach (Matzkin, 2003).

Our index is based on all those variables that capture the information that are typically contained in the job seeker's CV and in the application itself. As CV characteristics, we consider age, education, residential status, sex, and targeted occupational sector of the applicant, all provided through the unemployment office data, as well as additional information on labor market history, retrieved from the social security registers. Further, we consider information on the caseworker or PES office to which the job seeker is affiliated. As application characteristics, we consider the application channel (*i.e.* in person, by phone, written), an indicator for whether the application results from a caseworker referral, the within-month rank of the application and a measure of search intensity (the estimated individual application fixed effect $\hat{\alpha}_i$).

To construct the application-specific *ex-ante* propensity of success, we start by using the first month in which individual i's job search behavior is documented in the data. We denote this month τ_i. For the callback stage, this individual-specific reference month corresponds to the first month when individual i starts recording applications, *i.e.* $\tau_i = \tau_i^A$. For the job offer conversion stage, it is equal to the first month when job seeker i records an interview, as job offers are conditional on having been called back, *i.e.* $\tau_i = \tau_i^C$. For each of the two stages, we estimate a binary outcome model for the application-level probability of obtaining a callback or a job offer (conditional on a callback) in the corresponding reference month. We model the latent propensities $\tilde{c}_{ij\tau_i}$ and $\tilde{o}_{ij\tau_i}$ in month τ_i as

$$\tilde{c}_{ij\tau_i} = \vartheta_0 + X^1_{i\tau_i}\vartheta_1 + X^2_{ij\tau_i}\vartheta_2 + \delta^c_{mk} - \nu_{ij\tau_i} \tag{5.1a}$$

$$\tilde{o}_{ij\tau_i} = \varphi_0 + X^1_{i\tau_i}\varphi_1 + X^2_{ij\tau_i}\varphi_2 + \delta^o_{mk} - \eta_{ij\tau_i} \tag{5.1b}$$

where $\tau_i = \tau_i^A$ in equation (5.1a) and $\tau_i = \tau_i^C$ in equation (5.1b). The row vector $X^1_{i\tau_i}$ contains the individual-level characteristics, the row vector $X^2_{ij\tau_i}$ the application-level characteristics, δ^c_{mk} and δ^o_{mk} are fixed effects capturing the conditions in local labor market m in calendar quarter k, $\nu_{ij\tau_i}$ and $\eta_{ij\tau_i}$ are idiosyncratic error terms. The conditional *ex-ante* probabilities of obtaining a callback and a job offer (conditional on a callback) in month τ_i are given by

$$\gamma^c(X_{ji\tau_i}) = \mathbb{P}(c_{ij\tau_i} = 1 | X_{ij\tau_i}, a_{ij\tau_i} = 1) = \mathbb{P}(\vartheta_0 + X^1_{i\tau_i}\vartheta_1 + X^2_{ij\tau_i}\vartheta_2 + \delta^c_{mk} > \nu_{ij\tau_i}) \tag{5.2a}$$

$$\gamma^o(X_{ij\tau_i}) = \mathbb{P}(o_{ij\tau_i} = 1 | X_{ij\tau_i}, c_{ij\tau_i} = 1) = \mathbb{P}(\varphi_0 + X^1_{i\tau_i}\varphi_1 + X^2_{ij\tau_i}\varphi_2 + \delta^o_{mk} > \eta_{ij\tau_i}) \tag{5.2b}$$

where $X_{ij\tau_i} = (X^1_{i\tau_i}, X^2_{ij\tau_i}, \delta^y_{mk})$, with $y = c, o$. Again, $\tau_i = \tau_i^A$ in the first equation and $\tau_i = \tau_i^C$ in the second.

We estimate equations (5.2a) and (5.2b) using logit models, respectively on applications in the reference month $t = \tau_i^A$, and on callbacks in the reference month $t = \tau_i^C$. We retrieve parameters estimates $\hat{\vartheta}' = (\widehat{\vartheta}_0 \ \widehat{\vartheta}_1 \ \widehat{\vartheta}_2)$ and $\widehat{\varphi}' = (\widehat{\varphi}_0 \ \widehat{\varphi}_1 \ \widehat{\varphi}_2)$ and predict the conditional *ex-ante* probabilities for all subsequent months $t \geq \tau_i$, that is to say the probability of a callback for all applications in $t \geq \tau_i^A$, and the probability of a job offer conversion for all callbacks in $t \geq \tau_i^C$. The resulting predicted probabilities, $\widehat{\gamma}^c_{ijt} = \widehat{\gamma}^c(X_{ijt})$ and $\widehat{\gamma}^o_{ijt} = \widehat{\gamma}^o(X_{ijt})$, capture the propensity with which an application sent out in month t receives a positive response from the firm, if the firm's behavior was kept as it was early in the unemployment spell, in month τ_i.

Finally, we estimate the net duration dependence in firms' responses using the logarithm of the *ex-ante* probabilities to control for dynamic sorting based on observables. Specifically, we estimate the following two binary outcome models using respectively all applications in months $t \geq \tau_i^A$ and all callbacks in months $t \geq \tau_i^C$:

$$\mathbb{P}(c_{ijt} = 1 | \, \widehat{\gamma}^c_{ijt}, a_{ij\tau_i} = 1) = \mathbb{P}\Big(\alpha^c + f^c(t; \phi^c) + \beta^c \ln(\widehat{\gamma}^c_{ijt}) > \varepsilon^c_{ijt}\Big) \tag{5.3a}$$

$$\mathbb{P}(o_{ijt} = 1 | \, \widehat{\gamma}^o_{ijt}, c_{ij\tau_i} = 1) = \mathbb{P}\Big(\alpha^o + f^o(t; \phi^o) + \beta^o \ln(\widehat{\gamma}^o_{ijt}) > \varepsilon^o_{ijt}\Big) \tag{5.3b}$$

where $\beta^c \ln(\widehat{\gamma}^c_{ijt})$ and $\beta^o \ln(\widehat{\gamma}^o_{ijt})$ control for dynamic selection, whereas $f^c(t; \phi^c)$ and $f^o(t; \phi^o)$ measure the net duration dependence in the callback and job offer conversion probabilities.[24]

5.3 Results

5.3.1 Main results

We report logit estimates for the *ex-ante* callback and job offer conversion probabilities defined in equations (5.2a) and (5.2b) in Table C1, in the Appendix. Observed characteristics are found to predict the *ex-ante* callback probability significantly, with older job seekers and those writing many applications receiving lower callbacks, while those with high education, and a high wage receiving callbacks with higher probability. Applications in person, and those referred by caseworkers tend to receive more callbacks. Job offers (conditional on callbacks) are not significantly predicted by age or residence permit, but job seekers with higher education and high previous wage are found to stand a lower chance of receiving a job offer. Overall *ex-ante* callbacks are predicted with a higher pseudo-R^2 (of around 11 percent) than *ex-ante* job offer conversions (pseudo-R^2 is 6 percent).

We now present evidence on the relationships between the *ex-ante* chances of positive responses by firms and elapsed unemployment duration in Figure 5.2. This provides insight on the role of heterogeneity and how the pool of applications evolves dynamically.

Figure 5.2A shows evidence for callbacks. The left graph depicts the empirical average callback probabilities and average *ex-ante* callback probabilities for all applications, in months $t \geq \tau_i^A$. As previously emphasized, the observed chances of being interviewed decrease strongly over the course unemployment, from around 5 percent to 2.5 percent (solid line). The *ex-ante* prediction of a positive callback (dashed line) also exhibits a substantial decline with respect to elapsed unemployment duration, from around 5 percent to less than 4 percent. This suggests that a sizeable part of the reduction in the callback chances is related to the quality of job applications. However, duration itself still seems to directly affect firms' callback decisions, as suggested by the steeper profile in the empirical callback probability compared to its *ex-ante* counterpart.

[24]Alternatively, we directly control for observables X_{ijt} in equations (5.3a) and (5.3b) instead of controlling for the logarithms of the *ex-ante* chances $\widehat{\gamma}^c_{ijt}$ and $\widehat{\gamma}^o_{ijt}$. These alternative results are reported along our *ex-ante* chances specification in the next subsection.

The right graph of Figure 5.2A further characterizes the dynamic selection process in the callback phase. It provides additional information on the distribution of the *ex-ante* chances of a callback for the pool of job applications sent out at each duration of the unemployment spell. The quality of the pool of job applications deteriorates substantially, with the best applications (95$^{\text{th}}$ percentile) having an almost 15 percent callback chance in the first month, and around 10.5 percent after fifteen months, whereas the lowest quality applications have very low callback chances throughout the spell (approximately 1 percent). This suggests that high-quality applications tend to disappear from the applications pool, as job seekers who write those get invited to job interviews, receive job offers and exit unemployment. Controlling for heterogeneity in the quality of applications hence seems to be crucial, if we seek to have a precise idea on whether there truly exists net duration dependence in callback decisions.

Figure 5.2B depicts a completely different picture for the job offer conversion stage. As previously emphasized, the observed duration profile in the job offer conversion probability is slightly increasing over the whole duration of the unemployment spell. It also appears to be essentially flat for intermediary durations: from month three to thirteen, the average probability that an interview converts into a job offer amounts to approximately 24 percent, regardless of when the application that led to the interview takes place. This is also the case for the *ex-ante* chances of the interview to be transformed into a job offer, which are globally constant over time. Importantly, the *ex-ante* probability of a job interview to convert into a job offer is statistically related to job seekers' and applications' characteristics. However, the distinguishing difference between callbacks and job offer conversions is that firms do not reduce the chances of offering a job to an interviewee whose application arrives late in the spell. If anything, the probability of obtaining a job offer after being interviewed is somewhat higher for the long-term unemployed than for the short-term unemployed.

The limited role of observable characteristics at this stage of the recruitment process is also visible on the right graph of Figure 5.2B, which presents additional evidence on the distribution of the *ex-ante* job offer conversion chances, for each month of unemployment. The graph shows that the *ex-ante* quality of interviewees, as predicted from their observed characteristics, remains constant over the spell of unemployment. The median chances of securing the job conditional on an interview are around 21 percent, with both substantially higher and lower chances of getting a job offer. Unlike callbacks, the *ex-ante* quality of interviewees measured by observable characteristics remains constant, regardless of the time when the interview occurs. This second decision by the firm is thus likely to be based on information that are unobserved to us, or that are idiosyncratic to each worker-firm match. In the end, firms seem to make limited use of elapsed unemployment duration to infer applicant's quality when deciding whether to make her a job offer.

Figure 5.2: Job seekers' *ex-ante* chances and elapsed unemployment duration

(A) Callback probability

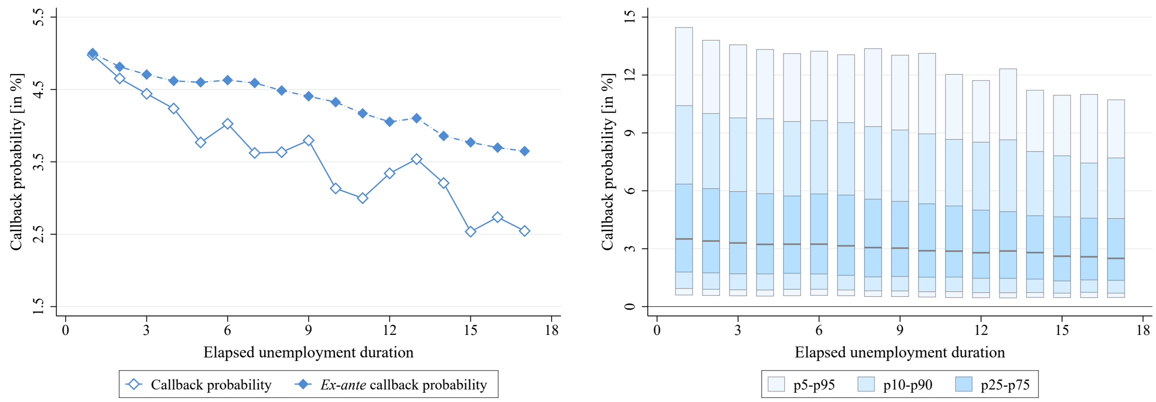

(B) Job offer conversion probability

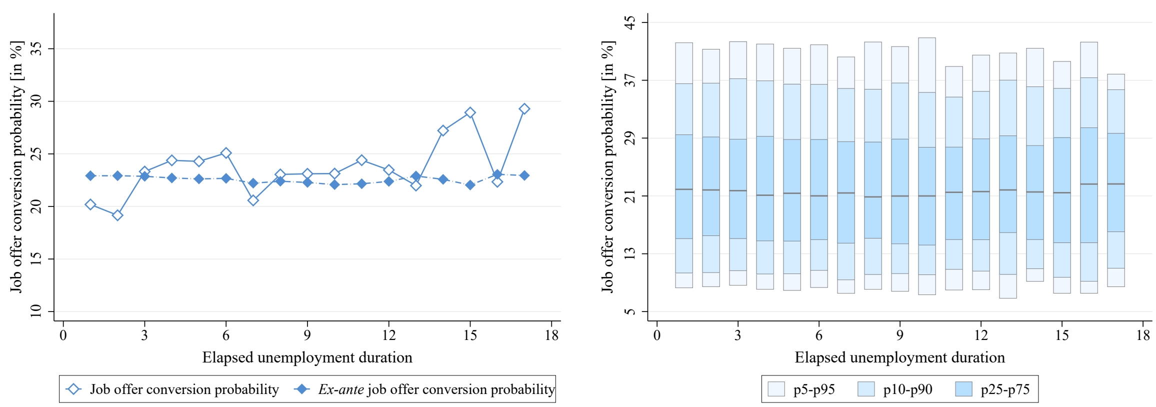

Note: This figure reports evidence on the relationships between *ex-ante* applications success chances $\widehat{\gamma}^c_{ait}$, $\widehat{\gamma}^o_{ait}$ and elapsed unemployment duration. Panel A depicts descriptive evidence for the callback probability, while panel B focuses on the job offer conversion probability. For each panel, the left graph represents the average empirical duration profiles of the probability of a positive response by firms (solid line) and of its *ex-ante* counterpart (dashed line). The right graph depicts summary statistics on the distribution of the *ex-ante* chances of a callback and job offer conversion, for each month of elapsed unemployment.

We next provide estimates of the net effects of duration on the probability of a callback or a job offer conversion, formalized in equations (5.3a) and (5.3b). Table 5.1 presents estimates for a linear specification of the duration effects. Columns (1)-(3) report results for callbacks, while columns (4)-(6) focus on job offer conversions.

Column (1) shows that the probability of a callback decreases by approximately 0.15 percentage points per additional month spent unemployed, in the raw data. Directly, controlling for individual and applications characteristics, policy controls and local labor market conditions dampens the decline in the callback probability to less than 0.1 percentage points per additional month in unemployment, as shown in column (2). Alternatively, controlling for the logarithm of the *ex-ante* callback chances of job applications in column (3) delivers a simi-

Table 5.1: Duration dependence in firms' responses, linear specification

	Callback probability			Job offer conversion probability		
	(1)	(2)	(3)	(4)	(5)	(6)
Elapsed unemp. dur.	-0.155***	-0.097***	-0.096***	0.350***	0.429***	0.380***
	(0.015)	(0.015)	(0.015)	(0.099)	(0.097)	(0.094)
	[-3.117%]	[-1.945%]	[-1.921%]	[1.736%]	[2.123%]	[1.883%]
ln(Ex-ante chance)			3.364***			18.805***
			(0.094)			(0.863)
Individual controls	No	Yes	No	No	Yes	No
Policy controls	No	Yes	No	No	Yes	No
LLMC	No	Yes	No	No	Yes	No
Control for ex-ante pr.	No	No	Yes	No	No	Yes
Pseudo R^2	0.003	0.094	0.075	0.001	0.050	0.044
N. observations	600323	600323	600323	22422	22422	22422

Note: This table reports estimates of net duration dependence in the callback and job offer conversion probabilities. Columns (1)-(3) report estimates for callbacks and correspond to equation (5.3a). Columns (4)-(6) focus on job offer conversions and correspond to equation (5.3b). Application-level observations are weighted by the inverse of the monthly number of applications sent out by individual i in month t, so as to put equal weight on all monthly-individual observations. Coefficients correspond to average marginal effects and are reported in percentage points. Errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

lar role for prolonged unemployment duration, a reduction of 0.1 percentage point for each additional month spent unemployed.

As the callback probability declines in a somewhat non-linear fashion with elapsed unemployment duration, we probe results from a model that leaves duration dependence of the callback probability fully unrestricted, in Figure 5.3A. Recall that the observed callback probability decreases from 5 percent to 2.5 percent. The callback probability adjusted for dynamic selection, which we obtain from a model like column (3) of Table 5.1, but where the $f^c(t; \phi^c)$ function is specified in a saturated manner, is substantially higher: it decreases from 5 percent in the first month to 3.5 percent after fifteen months.

After accounting for individual heterogeneity, the decline in callback chances for long-term unemployed appears real, and might be responsible for part of the decrease in job finding chances over time. Quantitatively, the net decline in the callback probability amounts to about two thirds of the observed decline. Our estimates are surprisingly similar to Kroft et al. (2013), who find a 3.7 percentage points decline for 36 months of unemployment, equivalent to a monthly decrease of 0.1 percentage points in callback chances.

Turning to the results for the conversion of interviews into job offers, we find significantly positive duration dependence in the raw data, consistently with our descriptive evidence. Column (4) of Table 5.1 shows that the job offer conversion probability increases by 0.35 percentage points for each additional month spent unemployed. Directly controlling for observed heterogeneity slightly increases the duration dependence parameter in column (5), whereas controlling for *ex-ante* job offer conversion chances does not affect it in column (6). Again, we probe for non-linear duration dependence patterns in a version of the model used in column

Figure 5.3: Net duration dependence in firms' responses, saturated specification

(A) Callback probability

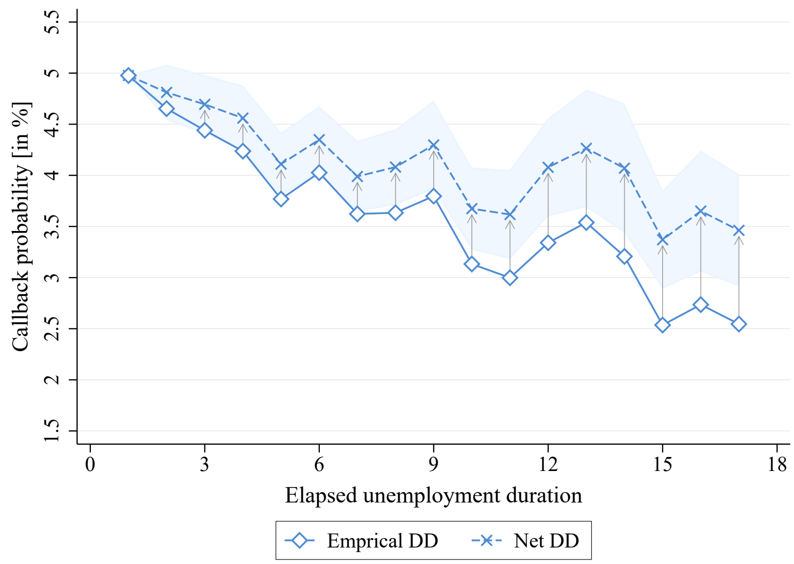

(B) Job offer conversion probability

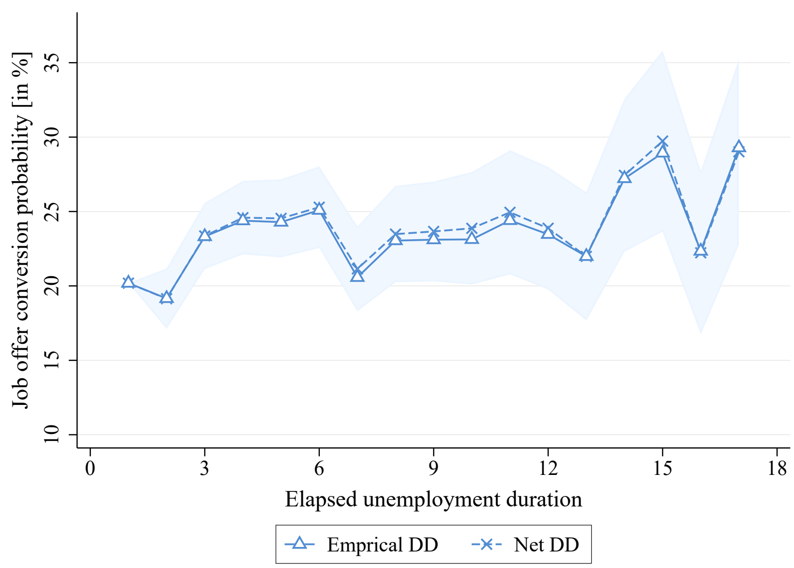

Note: This figure depicts estimates of equations (5.3a) and (5.3b), where the duration functions $f^c(t;\phi^c)$ and $f^o(t;\phi^o)$ are fully saturated. Panel A relates to the callback probability, while panel B focuses on the job offer conversion probability. The solid lines depict empirical average duration profiles computed on the raw data, while the dashed lines represent the corrected average duration profiles that control for the *ex-ante* chances of job seekers to obtain a callback or a job offer conversion. 90% confidence intervals for the corrected duration profiles based on clustered standard errors at the individual level are reported.

(6) of Table 5.1, where duration dependence is left unrestricted. Results in Figure 5.3B show graphically that observed characteristics do not affect estimates of the job offer conversion duration profile at all, which remains positive and aligned on the pattern obtained on the raw

data.

Our empirical analyses suggest that the net effect of duration on job offer conversion is positive, meaning that this second decision by the firm does not exacerbate the decline in job finding chances. Intuitively, positive duration dependence at the job offer conversion stage might arise from the reduction in applicants' heterogeneity in the callback phase: as interviewees tend to become more homogeneous over time, their chance of obtaining a job offer after being interviewed might be increasing with elapsed unemployment duration. Alternatively, this positive duration profile might arise from workers' higher availability at longer unemployment duration, or to some form of learning on the job seeker's or employer's side. Job seekers' learning, for instance, may occur because unemployed participate in several interviews, and learn about the recruiting process from previous failed interviews. Figure C4 in the Appendix tends to invalidate this mechanism: the duration profile of job offer conversion for job seekers who attempt their first interview does not differ much from its counterpart for follow-up interviews, with different employers on different jobs. In both cases, the conversion probabilities are higher for interviews undertaken at a later stage of the unemployment spell. Even though we find evidence that contradict part of these alternative mechanisms, we cannot fully rule them out based on our observations.

When contrasting results for the callback and job offer conversion stages, it appears that the two differ tremendously. At the callback stage, firms pick job applications from the available pool of applicants, and this pool depletes leaving those applicants that are less likely and able to receive callbacks. This process hence affects the composition of the pool of job seekers strongly. The situation differs for job offer conversion: firms make offers to interviewees, but the chances of a job offer, which are still related to individual characteristics, do not decline with elapsed duration of unemployment, but rather increase.

These results highlight two interesting insights: the duration of unemployment negatively affects the callback stage, as emphasized by audit studies (Oberholzer-Gee, 2008; Kroft et al., 2013; Eriksson and Rooth, 2014; Nüß, 2018), but does not reduce job seekers' chances in the job offer conversion phase (Jarosch and Pilossoph, 2019). Job seekers lose out on some interviews through duration dependence in callbacks, but long-term unemployed do not seem to be further discriminated when firms make their job offer decisions. All told, our results are coherent with a statistical learning view of the labor market, a view we explore and outline more fully in the next section of the paper.

5.3.2 Robustness

Our identification of the net effect of duration on firms' responses is based on a conditional independence assumption: we suppose we observe all relevant information to the recruiting firm, at the moment when it evaluates applications sent by job seekers. This assumption is reasonably met for callbacks, as we observe most information that is relevant for this decision in our data. Our approach is thus similar in essence to Mueller and Spinnewijn (2023), both being based on rich administrative data sources to proxy job seekers' chances of obtaining positive responses from firms.

In the previous section, we consistently show that our set of conditioning variables mostly play a role in the first phase of firms' screening process, which is typically based on CV information. To further assess the relevance and predictability of our conditioning variables, we re-estimate equations (5.3a) and (5.3b) using additional controls from our administrative data, which are supposedly unobserved by firms when they first screen applications. Those additional variables consist in information collected by the caseworker at the occasion of her first meeting with the job seeker at the PES office (job seeker's employability, job seeker's degree of mobility) and additional information that is not disclosed to the firm by the job seeker when applying (experience of sick days during the unemployment spell). Given that these variables are not directly observed by firms when they screen applications, we expect

Table 5.2: Duration dependence in firms' responses, control for non-CV information

	(1)	(2)	(3)	(4)
A. Callback probability				
Elapsed unemp. duration	-0.097***	-0.094***	-0.096***	-0.092***
	(0.015)	(0.015)	(0.015)	(0.014)
ln(Ex-ante chance)			3.365***	3.372***
			(0.094)	(0.092)
Individual controls	Yes	Yes	No	No
Policy controls	Yes	Yes	No	No
LLMC	Yes	Yes	No	No
Control for ex-ante pr.	No	No	Yes	Yes
Info not on CV	No	Yes	No	Yes
Pseudo R^2	0.094	0.099	0.075	0.079
N. observations	600323	600323	600323	600323
B. Job offer conversion probability				
Elapsed unemp. duration	0.430***	0.407***	0.381***	0.362***
	(0.097)	(0.097)	(0.094)	(0.094)
ln(Ex-ante chance)			18.829***	19.637***
			(0.868)	(0.810)
Individual controls	Yes	Yes	No	No
Policy controls	Yes	Yes	No	No
LLMC	Yes	Yes	No	No
Control for ex-ante pr.	No	No	Yes	Yes
Info not on CV	No	Yes	No	Yes
Pseudo R^2	0.050	0.060	0.044	0.054
N. observations	22422	22422	22422	22422

Note: This table reports estimates of net duration dependence in the callback and job offer conversion probabilities, for our baseline set of conditioning variables (columns 1-3) and the extended set of conditioning variable, including non-CV characteristics (columns 2-4). Panel A reports estimates for callbacks, whereas panel B reports estimates for job offer conversion. Application-level observations are weighted by the inverse of the monthly number of applications sent out by individual i in month t, so as to put equal weight on all monthly-individual observations. Coefficients correspond to average marginal effects and are reported in percentage points. Errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

their role to be minor when measuring net duration dependence in the callback phase. In contrast, their role is possibly greater when estimating the net duration profile in the job offer conversion probability, as more information on the job seeker is revealed to the firm through the interview.

Duration dependence estimates for this extended model together with our baseline results are reported in Table 5.2. As expected, adding these controls leads to a more marked change in the pseudo-R^2 for the job offer conversion stage (+20%) compared to the callback stage (+5%). This can also be seen through the parameter associated with the *ex-ante* application success probability, which increases more markedly for job offer conversion compared to callbacks (in columns 3 and 4). Consequently, the change in the estimated net duration effect is virtually zero for callbacks, while it is slightly larger for job offer conversion (in relative terms). These results support the idea that our baseline set of conditioning variables capture most individual heterogeneity that is relevant at the callback stage, and that we truly capture the net duration profile of callback chances. They also point towards the fact that additional information is revealed during the job interview, and that those might affect firm's decisions whether to make a job offer to the interviewee. Even though such information remains unobserved in our context, it should have a limited impact on the net duration profile of job offer conversion we estimate.

6. Rationalizing the evidence: a job search model with statistical discrimination

6.1 What do we know?

So far, we have examined the dynamics of the different phases of the job search process. From a purely descriptive view, we find evidence of a slight decline in the monthly number of applications sent out by job seekers with respect to elapsed unemployment duration. The decrease in application effort is concurrent to a large reduction in the probability of each application to lead to a callback for a job interview. In contrast, the probability with which interviews are converted into job offers is slightly increasing over time.

The contribution of individual heterogeneity to these empirical patterns is contrasted. On the job seeker's side, we find evidence of positive dynamic selection with respect to application effort. As longer-term unemployed apply more at any duration, heterogeneity tends to attenuate the net effect of duration in the raw data. The opposite is true for callback decisions by firms: individuals who remain unemployed longer tend to face lower chances of being called back, at any duration. Consequently, dynamic selection is negative at this stage, and tends to exacerbate the duration profile of the callback probability in the raw data.[25] Finally, workers' characteristics play a minor role in the last phase of the job search process, when

[25]Consequent to dynamic selection being positive with respect to job application effort, and negative with respect

interviews are converted into job offers: controlling for observed heterogeneity virtually does not affect the duration profile of the job offer conversion probability, which is still estimated to be positive.

All things considered, our results corroborate the findings that individual heterogeneity is an important driver of duration dependence in job finding. However, our analysis suggests that the role of heterogeneity is not uni-dimensional, as it affects the duration profiles of job seekers and firms' behaviors differently. Moreover, even after accounting for heterogeneity, we find that duration itself still affects directly and negatively application effort provision by job seekers, and firms' callback decisions. In contrast, duration is found to have a limited positive effect on the final decision of the job search process, *i.e.* firms' choice of converting job interviews into job offers. In the next section, we present a job search model with statistical discrimination, that rationalizes the evidence we have emphasized so far.

6.2 Duration based discrimination in job search

6.2.1 The model

We develop a job search model with statistical discrimination to rationalize our empirical findings in an equilibrium framework. The model builds on Jarosch and Pilossoph (2019), which constructs a frictional labor market characterized by two-sided heterogeneity, positive assortative matching, and a multi-stage hiring process (callback and interview/job offer conversion stage). This setting lends itself naturally to generating negative duration dependence at the callback stage as an endogenous response to negative dynamic selection of lower-ability workers at longer unemployment duration. Intuitively, unemployment duration conveys a signal to firms about the average ability of the applicant: the longer the unemployment spell, the higher the likelihood that the applicant is of low ability and, therefore, the smaller the pool of firms that is willing to interview her.

To study duration dependence in workers' and firms' decisions jointly, we augment Jarosch and Pilossoph (2019) framework with endogenous search effort by workers.[26] In practice, we add a preliminary stage at the beginning of the hiring process in which workers decide how much search effort to exert. Crucially, optimal search effort is increasing in the job offer probability (per unit of search effort). As long as the latter declines with unemployment dura-

to the callback probability, we find that high application-effort individuals face lower callback chances. This can be seen in Figure C5A, where we plot the relationship between the α_i estimated from section 4 and the individual-specific average of the empirical (or *ex-ante*) callback probability, from section 5. The same relationship is plotted for the job offer conversion probability, in Figure C5B.

[26]The possibility of augmenting the model with endogenous search effort is already hinted at in footnote 47, in Jarosch and Pilossoph (2019). However, the authors did not pursue such research address since it would allegedly have led to mitigating the impact of firms' discrimination on the job finding rate, the authors' goal being to quantify an upper bound to such an impact.

tion, *e.g.* because of firms discriminating against longer unemployment duration, job seekers find it optimal to reduce their search effort over the unemployment spell. This allows us to rationalize net duration dependence both on the worker's and the firm's side as equilibrium responses to negative dynamic selection of workers at longer unemployment duration. In what follows, we lay down the main elements of the model and then highlight their role to rationalize our empirical findings.

Environment

We consider a discrete-time economy populated by a unit mass of workers, who differ by their permanent ability $x \sim \mathcal{L}(x),\ x \in \mathcal{X} = (\underline{x}, \overline{x})$, and a continuum of firms, which differ by their productivity $y \sim F(y),\ y \in \mathcal{Y} = \big(\underline{y}, \overline{y}\big)$. Workers have preferences exhibiting constant relative risk aversion $\gamma \geq 1$, firms are risk-neutral and all agents discount the future at common rate $\beta \in (0, 1)$.

Workers and firms interact in a frictional labor market under a random search protocol. Search frictions are represented by an exogenous separation probability δ_H and the endogenously determined job finding probability $O(x, \tau)$, where $\tau \in \mathbb{N}$ stands for elapsed unemployment duration.[27] The exogenous separation probability δ_H comprises both quits to unemployment with probability δ_L and job-to-job transitions with complementary probability $\delta_H - \delta_L$.

Job finding comes as the result of a three-step hiring process. First, workers decide how many job applications to send out, *i.e.* how much search effort s to exert, subject to an increasing and convex search cost function $\sigma(s),\ \sigma'(s) > 0,\ \sigma''(s) > 0$ (Pissarides, 2000). Second, workers' job applications come together with firms' vacancies with exogenous probability λ. Upon meeting, the only relevant information released to firms from workers' applications is the length of their unemployment spell. Based on this unique piece of information, firms decide whether to call the applicant back for a job interview at cost κ. Finally, conditional on interviewing the applicant, the firm gets to know her true ability type and decides whether to offer her a job.

Match Chapters/Chapter1/Output is governed by a production technology $p(x, y)$ characterized by positive assortative matching, *i.e.* the most productive firms are the most selective in terms of workers' ability:

$$p(x, y) = \begin{cases} x + y & \text{if } x \geq y \\ 0 & \text{else} \end{cases} \tag{6.1}$$

[27]The job finding probability in our structural model is the theoretical counterpart of the expected (number of) job offers in the empirical part of the paper, for reasons that will become clear later on.

A worker is hence qualified for a job if her ability x exceeds firms' productivity y.[28] Absent any discrimination, the job offer probability for a x-type worker reads

$$o^{ND}(x) = \lambda F(x) \tag{6.2}$$

Workers enjoy a flow value of leisure b while unemployed. Following Hall (2005), wages are rigid and fixed at $\omega \in \big(b, p(x,y)\big)$ for the entire duration of the match. Unlike in Jarosch and Pilossoph (2019), employed workers are therefore strictly better-off than unemployed, thus providing a motive for exerting search effort.

Workers

Workers are assumed to be risk-averse and make intertemporal consumption-savings choices. They find it optimal to perfectly smooth consumption q during their lifetime, so that their consumption equals their permanent income, which is function of their ability types.[29] Regarding labor supply decisions, workers are either matched to a firm or unemployed. Unemployed workers choose how many applications to send out over their unemployment spell, *i.e.* how much search effort s to exert at each unemployment duration τ, so as to maximize the value of unemployment. Risk aversion implies that workers value their current income according to their marginal utility of consumption, thus giving rise to a wealth effect in optimal search effort decisions. As a result, the marginal benefit of exerting search effort differs across workers' ability types x.

The values of unemployment and employment can be expressed recursively as

$$U(\tau;x) = \max_{\hat{s}\geq 0} \; \frac{b}{q(x)^{\gamma}} - \sigma_x(\hat{s}) + \beta\Big[U(\tau+1;x) + \hat{s}\cdot o(x,\tau)\,\big(W(x) - U(\tau+1;x)\big)\Big] \tag{6.3}$$

$$W(x) = \frac{\omega}{q(x)^{\gamma}} + \beta\Big[W(x) + \delta_L\big(U(0;x) - W(x)\big)\Big]. \tag{6.4}$$

In words, the value of unemployment at duration τ is made up by the flow value of leisure net of search effort costs and a continuation value, which equals the discounted value of unemployment at duration $\tau+1$ plus the expected capital gain upon finding a job.[30] The value of employment equals the flow value of the wage rate and a continuation value, which accounts for possible separations into the zero-duration unemployment state. Since wages

[28] Throughout we assume that any qualified worker is profitable for the firm.

[29] Formally, permanent income is defined as I, and consumption $q = I(x)$. See Section D for more details on worker's consumption-saving choices.

[30] We model the search effort cost as a utility cost. Alternatively, it can be modelled as a monetary cost. The two formulations are equivalent up to the fact that the marginal utility cost of search effort is unitary in the former formulation, equal to the marginal utility of consumption in the latter.

are rigid, productivity heterogeneity across firms does not translate into wage dispersion, so the standard option value embedded in the unemployment state disappears. Consequently, workers' reservation wage boils down to b and every job offer is accepted in equilibrium.[31]
Optimal search effort balances the marginal cost of exerting higher search effort to the expected marginal benefit of meeting a firm. The latter is made up by the discounted capital gain upon employment multiplied by the marginal increase in the matching probability from harder search effort:

$$s(\tau;x): \quad \underbrace{\sigma'(s)}_{\text{marginal cost}} = \underbrace{\beta\, o(x,\tau)\, \Big[W(x) - U(\tau+1;x)\Big]}_{\text{marginal benefit}} \tag{6.5}$$

The optimal search effort at unemployment duration τ depends on the value of unemployment at duration $\tau+1$, thus making the application decision non-stationary. Specifically, the model generates negative duration dependence in search effort if and only if the marginal benefit of search effort decreases with elapsed unemployment duration τ. This is the case if the job offer probability exhibits negative duration dependence in equilibrium.[32] Moreover, the capital gain upon employment features a wealth effect due to heterogeneity in marginal utility of consumption across workers, as the flow capital gain upon employment is discounted differently depending on the worker's permanent income. As long as workers' permanent income is increasing in ability x, more productive workers enjoy lower marginal benefit from finding a job because they can live off their larger savings while unemployed.[33] As a result of such wealth effect, workers with higher ability exert less search effort over their entire spell, all else equal.[34] If differences in wealth and marginal utility of consumption dominate those in job offer probabilities in the cross section, workers with low ability would exert higher search effort than high-ability ones, thus providing a potential explanation for the positive dynamic selection of workers with higher search effort over the unemployment spell detected in the data.[35]

[31] As job offers and hirings coincide in the model, $O(x,\tau)$ characterizes both the job finding probability and the probability that a worker obtains a job offer, accounting for her effort provision.

[32] Formally, this condition writes $\frac{\partial o(x,\tau)}{\partial \tau} < 0$. A sufficient and necessary condition for negative duration dependence in search effort to arise is that the direct effect of the reduction in the job offer probability, *i.e.* $\frac{\partial o(x,\tau)}{\partial \tau} < 0$, dominates the indirect effect from the increase in the capital gain upon employment due to the depletion of the value of unemployment as the unemployment spell lengthens, *i.e.* $U(\tau+1;x) < U(\tau;x)$.

[33] Given that workers with higher ability find jobs more quickly, their permanent income is higher than that of lower ability types, even in the absence of any wage differential.

[34] Faberman and Kudlyak (2019) alludes in footnote 27 to a dominant wealth effect in search effort as a possible explanation for the evidence that workers with lower job prospects or located in less tight labor markets search more intensely.

[35] Alternatively, differences in search effort levels across job seekers might be due to type-specific search costs

Firms

Firms can either be matched with one worker or not. In the latter case, the value of the firm, J_v, is assumed to be zero. As a result, the value of a filled job, $J(x, y)$, is simply given by the present discounted value of flow profits.

$$J_v = 0 \tag{6.6}$$

$$J(x, y) = \frac{p(x, y) - \omega}{1 - \beta(1 - \delta_H)} \tag{6.7}$$

Upon receiving a worker's application, the firm decides whether to call her back for a job interview, based on her elapsed unemployment duration τ only, at cost κ. After the interview takes place, the firm discovers the worker's true ability x and decides whether to offer her a job.

In the first phase of the recruitment process, the callback set is defined as

$$\mathcal{C}(\tau) = \left\{ \tau \in \mathbb{N} \;\middle|\; \int \max\left\{ J(x, y), 0 \right\} \mu(x|\tau)\, dx \geq \kappa \right\} \tag{6.8}$$

where $\mu(x|\tau)$ is the conditional density of workers' ability at unemployment duration τ, the key equilibrium object governing statistical discrimination. In words, a firm calls back an unemployed worker with elapsed unemployment duration τ if the expected value of matching to this worker exceeds the interview cost κ. Denoting $y^*(\tau)$ the productivity of the firm that is just indifferent between calling back an unemployed worker with duration τ or not, the callback probability writes[36]

$$c(\tau) \equiv \lambda \mathbb{P}\big(y \leq y^*(\tau)\big) = \lambda F\big(y^*(\tau)\big) \tag{6.9}$$

As a result of negative dynamic selection, the callback set weakly shrinks with unemployment duration in equilibrium, thereby resulting in negative duration dependence in the callback probability.

In the second phase of the recruitment process, conditional on calling back the applicant, the firm gets to know worker's true ability during the interview. As a result, the firm makes a job offer to any worker who is qualified for its production technology, according to equation (6.1), regardless of unemployment duration. Since any job offer is accepted by the worker, the job

functions, $\sigma(s; x)$.

[36]Formally, the productivity of the marginal firm is such that $\int \max\{J\left(x, y^*(\tau)\right), 0\}\mu(x|\tau)dx = \kappa$.

offer (or hiring) set writes

$$\mathcal{O}(x) = \Big\{x \in \mathcal{X} \mid p(x,y) \geq \omega\Big\} = \Big\{x \in \mathcal{X} \mid p(x,y) > 0\Big\} \tag{6.10}$$

In words, a firm makes a job offer to any worker that grants it positive flow profits, after discovering her type during the interview. The job offer conversion probability, or hiring probability conditional on being interviewed, is then defined as

$$o(x)|c(\tau) \equiv \mathbb{P}\big(y \leq x | y \leq y^*(\tau)\big) = \frac{F\big(\tilde{y}(x,\tau)\big)}{F\big(y^*(\tau)\big)} \tag{6.11}$$

where $\tilde{y}(x,\tau) = \min\big\{x, y^*(\tau)\big\}$. Since $\frac{\partial y^*(\tau)}{\partial \tau} < 0$ due to the decline in callback chances, duration dependence in the job offer conversion probability is non-negative. In particular, there exists positive duration dependence in the job offer conversion probability if $y^*(0) > x$ for at least some x. Intuitively, positive duration dependence at the job offer conversion stage arises because the pool of job applicants become increasingly more homogeneous as unemployment duration lengthens, with low-ability workers accounting for the lion's share of job seekers. As a result, the signal embedded in unemployment duration becomes increasingly more informative about worker's ability, thus making firms callbacks more targeted. Finally, in the presence of statistical discrimination, the job offer probability is defined by:

$$o(x,\tau) \equiv c(\tau) \cdot o(x)|c(\tau) = \lambda\, F\big(\tilde{y}(x,\tau)\big) \tag{6.12}$$

Contrasting equations (6.2) and (6.12), we notice that statistical discrimination affects a worker's job offer probability if and only if $y^*(\tau) < x$, that is, if the worker is denied an interview for a job she would have been qualified for.

Stationary equilibrium

Closing the model requires to specify the equilibrium conditions for the measure of unemployed, as well as the unemployment composition across ability types and unemployment duration. To do so, we solve the model in stationary equilibrium.

Type-specific unemployment rate equals the sum of the measure of unemployed across all durations

$$u(x) = \sum_{\tau=0}^{\infty} u(x,\tau) \tag{6.13}$$

$$\text{where } u(x,\tau) = \begin{cases} \delta_L\big(1 - \sum_{t=0}^{\infty} u(x,t)\big) & \text{if } \tau = 0 \\ u(x,\tau-1)\cdot\big[1 - s(\tau-1;x)\; o(x,\tau-1)\big] & \text{if } \tau > 0 \end{cases}$$

Since the wage rate is fixed, differences in consumption or permanent income across workers are totally accounted for by differential labor market attachment.[37] As a result, high-ability workers search less intensely for a given job offer probability than low-ability ones, due to a wealth effect.

The key equilibrium object of the model is the conditional density of worker types x at duration τ, which is defined as

$$\mu(x|\tau) = \frac{s(\tau;x)\ u(x,\tau)\ \ell(x)}{\int s(\tau;\tilde{x})\ u(\tilde{x},\tau)\ d\mathcal{L}(\tilde{x})} \tag{6.14}$$

where $\ell(x) = \mathcal{L}'(x)$ is the probability density function of x. A stationary equilibrium of this economy is a triple $\{s(x;\tau), o(x,\tau), u(x,\tau)\}$, where search effort satisfies equation (6.5), the job offer probability satisfies equation (6.12), and the unemployment rate satisfies equation (6.13).

6.2.2 What drives duration dependence?

We can draw a clear parallel between our structural model and equation (3.2), describing the empirical decomposition of duration dependence in job offers. Noting that the job finding probability for a worker of type x at duration τ writes $O(x,\tau) = s(\tau;x)\ o(x,\tau)$, duration dependence in $O(x,\tau)$ can be decomposed into a net duration effect and a compositional change:

$$\begin{aligned} \mathbb{E}_\tau\big[O(x,\tau)\big] - \mathbb{E}_0\big[O(x,0)\big] = &\ \underbrace{\mathbb{E}_\tau\big[s(\tau;x)\ o(x,\tau) - s(0;x)\ o(x,0)\big]}_{\text{Net duration effect}} \\ &+ \underbrace{\mathbb{E}_\tau\big[s(0;x)\ o(x,0)\big] - \mathbb{E}_0\big[s(0;x)\ o(x,0)\big]}_{\text{Compositional change}} \end{aligned} \tag{6.15}$$

where $\mathbb{E}_\tau[\cdot]$ denotes the expectation with respect to the conditional density in equation (6.14). Again, the job offer probability can be decomposed into two stages as $o(x,\tau) \equiv c(\tau)\cdot o(x)|c(\tau)$. Following the steps in equation (3.3), we further write

$$\begin{aligned} \mathbb{E}_\tau\left[O(x,\tau)\right] - \mathbb{E}_0[O(x,0)] = &\ \mathbb{E}_\tau\big[\big(s(\tau;x) - s(0;x)\big)\ o(x,\tau)\big] \\ &+ \mathbb{E}_\tau\big[s(0;x)\ \big(o(x,\tau) - o(x,0)\big)\big] \\ &+ \mathbb{E}_\tau\big[s(0;x)\ o(x,0)\big] - \mathbb{E}_0\big[s(0;x)\ o(x,0)\big]. \end{aligned} \tag{6.16}$$

Just like in our empirical exercise, the decline in job offers and in the job finding rate are explained in part by a compositional change in the pool of unemployed and by the net effect

[37]Formally, this writes $q(x) = b\ u(x) + \omega\ \big(1 - u(x)\big)$.

of duration. The negative net duration effect can further be decomposed into a reduction of the job offer probability, driven by the decline in the callback probability, and a reduction in the number of applications sent out by unemployed.

We are now in the position to discuss the mechanism behind the duration dependence patterns observed in the data through the lens of our structural model. Upon meeting a worker with unemployment duration τ, firms form an expectation about her ability based on $\mu(x|\tau)$. Since workers with high ability x match more easily according to the production technology (equation (6.1)), the density $\mu(x|\tau)$ is featured by negative dynamic selection, with low-ability workers being over-represented at longer unemployment duration with respect to the unconditional density. Such negative dynamic selection entails a net negative duration dependence in the callback probability (equation (6.9)), as firms use elapsed unemployment duration as a screening device when choosing whether to call back an applicant for an interview.[38] This generates in turn negative duration dependence in the job offer probability (equation (6.12)). In equilibrium, workers optimally respond to the negative duration dependence in the job offer probability by scaling down their search effort over the unemployment spell, which therefore exhibits net negative duration dependence as well.

6.3 Competing explanations

Our structural model presents one possible explanation for the net duration patterns we observe in the data. Nevertheless, some of the dynamics we observe empirically might be due to other mechanisms, which we discuss in the following.

On the job seeker's side, a first alternative explanation for the downwards-sloping net duration profile in application effort lies in stock-flow sampling (Salop, 1973; Ebrahimy and Shimer, 2010). The basic idea behind this theory is that suitable jobs to which a job seeker might apply originate both from the initial stock of vacancies and the inflow of new vacancies in each period. In this setup, the number of applications is decreasing over the unemployment spell because workers initially apply to the stock of existing vacancies, before applying to the inflow of new vacancies in the subsequent periods. This mechanism entails a non-gradual decline in application effort with respect to elapsed unemployment duration. In our context, we find that application effort decreases gradually and linearly over time, which tends to

[38]Note that the signal conveyed by elapsed unemployment duration might be considered as more or less important by firms, depending on the sub-markets in which they are recruiting. Consequently, the negative effect of elapsed unemployment duration possibly affects job seekers' chances of receiving a callback heterogeneously. For instance, local labor market tightness has been shown to affect the extent to which firms consider this signal (Kroft et al., 2013). Even though this analysis cannot be replicated in our context for data limitation reasons, we find evidence of differentiated effects of unemployment duration on callback chances according to job seekers' characteristics, *e.g.* with respect to education level (see Figure D1). However, the effect of duration remains globally negative for all sub-categories of job seekers.

contradict the stock-flow sampling hypothesis. [39] Moreover, this hypothesis can directly be tested by estimating the duration dependence profile in application effort, controlling for the stock and flow of vacancies in the relevant labor market. If the stock-flow mechanism prevails, the net duration profile estimated by this augmented model should be flat. This approach is not applicable in our context given that we have no access to job vacancies data, but it has been followed by Faberman and Kudlyak (2019), who do not find supportive evidence for the stock-flow hypothesis.

Another competing hypothesis for the decline in application effort relates to the depletion of job seeker's personal network, which has been shown to play an important role in job finding (Beaman and Magruder, 2012; Burks et al., 2015; Hensvik and Skans, 2016). This mechanism can be seen as a form of personal stock-flow sampling, where the decline in total application effort is entailed by the exhaustion of job seekers' personal contacts. In contrast, applications sent out through other channels ought to remain constant, throughout the unemployment spell. We assess this alternative explanation by estimating equation (4.1) for three application effort measures, corresponding to the three application channels (personal, phone, written). Corresponding results are reported in Table 6.1. They show that the number of applications sent out in person, per phone and in writing all decrease with respect to elapsed unemployment, hence providing little support to the personal contacts exhaustion

Table 6.1: Duration dependence in application effort per channel

	Written		**Phone**		**Personal**	
	(1)	(2)	(3)	(4)	(5)	(6)
Elapsed unemployment duration	-0.037***	-0.132***	-0.003	-0.046***	-0.038***	-0.075***
	(0.009)	(0.020)	(0.006)	(0.012)	(0.005)	(0.012)
	[-0.535%]	[-1.899%]	[-0.146%]	[-2.278%]	[-2.034%]	[-4.034%]
Constant	7.224***		2.025***		1.771***	
	(0.071)		(0.039)		(0.039)	
Individual controls	No	Yes	No	Yes	No	Yes
Policy controls	No	Yes	No	Yes	No	Yes
LLMC	No	Yes	No	Yes	No	Yes
Individual FE	No	Yes	No	Yes	No	Yes
adj.-R^2	0.001	0.631	0.000	0.614	0.003	0.615
N. observations	58755	58755	58755	58755	58755	58755

Note: This table reports empirical estimates of equation (4.1) using OLS, where the parametric duration function $f^A(t;\phi^A)$ is specified linearly. The dependent variables are the number of applications sent out through the written (columns 1-2), phone (columns 3-4) and personal channel (columns 5-6). For each dependent variable, we report estimation results from a simple binary regression (on duration only) and from the full specification described in equation (4.1). Errors are clustered at the individual level and reported in parentheses. Coefficients in relative terms (with respect to the average in the first month of unemployment) are indicated in squared brackets. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

[39] In the most stylized framework, job seekers apply to the stock of vacancies only in the first month of unemployment. This results in a large discrete jump from the first period to the subsequent ones. In a more refined version of the model, this discontinuity can be smoothed out by assuming convex application costs, which would make it optimal for the job seekers not to exhaust the stock of existing vacancies in the first period. However, if stock-flow sampling were to prevail, we would still expect a larger decline in application effort in the early periods of unemployment, compared to latter ones.

mechanism .

On the firm's side, several alternative stories might explain the empirical net decline in the callback probability. A first candidate explanation relates to changes in application quality over time: part of the downwards-sloping duration profile in firms' callback probability could be due to the gradual downgrading of job applications characteristics. In our context, we observe an important qualitative aspect of applications: the channel used when contacting firms. As seen in section 5, this characteristic is strongly predictive of applications' success at the callback stage, and captures an important dimension of applications' quality (Beaman and Magruder, 2012; Burks et al., 2015; Hensvik and Skans, 2016).[40] Even though we control for this characteristic in our regressions, we still find evidence of a marked net decline in callback chances, which is unrelated to changes in applications quality. As shown in Figure 6.1, this is because the relative share of each channel in the pool of applications is relatively constant over time, even when controlling for individual heterogeneity. Changes in applications' quality, and more precisely in application channels, are hence unlikely to represent the main driver of our results.[41]

Another potential explanation for the decline in callback chances lies in applications targeting (Galenianos and Kircher, 2009; Wright et al., 2021; Lehmann, 2023). Initially, job seekers

Figure 6.1: Changes in the shares of application channels

(A) No control for individual heterogeneity — (B) Control for individual heterogeneity

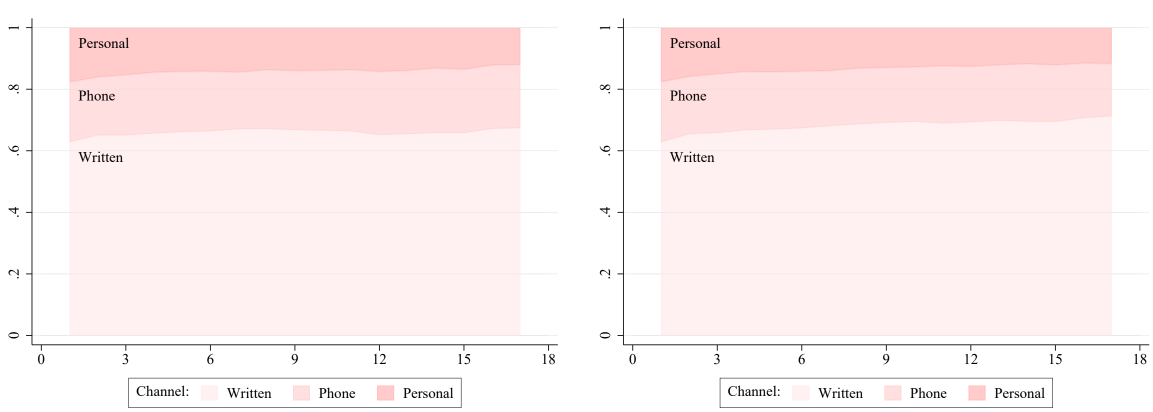

Note: This figure represents the share of applications sent out through the written, phone and personal channels, per month of elapsed unemployment. Panel A corresponds to the patterns in the raw data, without accounting for changes in the pool of applicants. Panel B corresponds to the results of a fixed effects regression, that accounts for the evolution of the pool of applicants.

[40] See Table C1 in the Appendix for the role of channels in callback and job offer conversion chances.

[41] If any, we would expect (unobserved) changes in applications' quality to be actually increasing overt time, as job seekers learn how to write better applications over time. Such omitted factor would entail an upward bias in the net duration profile we estimate, meaning that the true net duration dependence in the callback probability would actually be more negative.

might target a specific occupation, before starting to search broader and apply to a wider set of job ads, as spell duration increases. This may reduce callback chances, as job seekers are potentially less suited to the positions they newly apply to. If this mechanism is at play, we should observe adjustments in job search targets over time. We assess this point by studying how occupational targeting changes over time in the *Auxiliary sample*, for which information on occupations is available. Specifically, we construct two measures that characterize the types of occupations job seekers target: a binary variable indicating whether the targeted occupation is the same as the occupation desired by the job seeker, and a measure of net cognitive requirements of targeted occupations.[42] As depicted in Figure 6.2A, we find very little change in occupation search breadth, as measured by the same-occupation indicator. This result is robust to control for heterogeneity through individual fixed effects. Regarding skills requirements, Figure 6.2B shows that the average value of our net cognitive requirements measure decreases substantially along unemployment, from 0.18 to 0.14 over seventeen months. However, this decline is strongly attenuated when job seekers' fixed effects are added. This suggests that the above-mentioned change is largely due to a compositional change, and not to a change in application targeting within spells. Altogether, these evidence

Figure 6.2: Changes in application targeting

(A) Same occupation (B) Skills requirements

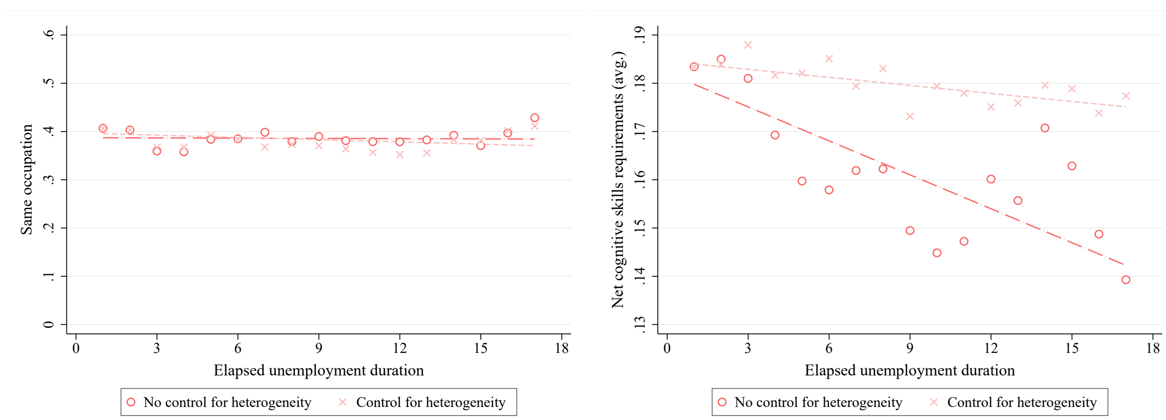

Note: This figure describes the evolution of applications characteristics with respect to elapsed unemployment duration. The two panels are based on the *Auxiliary sample*. Panel A shows results for the share of targeted positions that are the same as occupations desired by the job seekers. Panel B reports evidence for the net-cognitive skills requirements of targeted occupations. Both panels show evidence based on the raw data (circle) and evidence controlling for individual heterogeneity, through individual fixed effects (x-cross).

[42]In our data, occupations are categorized according to the *Swiss Standard Classification of Occupations 2000* (SSCO 2000). This job nomenclature follows a hierarchical structure, and presents 5 different levels of occupational groups. The binary indicator for occupational similarity between the desired occupation (at the spell level) and targeted occupation (at the application level) can be constructed for the different levels of the SSCO 2000. As for the net cognitive requirements measure, we use *O*Net* skill and ability requirements for each occupation. *O*Net* provides 52 abilities and skills, grouped into cognitive and physical. Our net cognitive measure is based on the difference between weighted importance of cognitive skills requirements and physical requirements.

point towards a limited role of application targeting in the decline of callbacks and job finding chances.

Finally, the net decline in callback chances might be due not to statistical discrimination, but rather to taste-based discrimination. In that case, firms would discriminate against long-term unemployed for preference motives rather than because of statistical inference. Such mechanism is typically formalized in models with multiple applications per job opening, which gives rise to coordination frictions. This requires firms to select a tie-breaking rule when calling back *ex-ante* homogeneous workers. In this framework, negative duration dependence in the callback probability would thus arise in the absence of screening motives, as long as firms call back applicants with lower unemployment duration first (Blanchard and Diamond, 1994). In our context, taste-based discrimination would imply a flat duration profile in the average *ex-ante* probabilities of applications to end up in callbacks, since duration is not statistically related to job seekers' productivity. This hypothetical pattern is however at odds with our empirical observation of dynamic selection in the callback phase and negative duration dependence in the *ex-ante* callback probability (see Figure 5.2A), hence contradicting the taste-based discrimination hypothesis.

7. Conclusion

The decline in job finding chances due to prolonged unemployment has for long been documented in the job search literature. The reasons behind this decline remain relatively misunderstood though. Recent studies have highlighted individual heterogeneity and dynamic selection as key drivers of this negative duration dependence pattern. However, those usually focus on the ultimate outcome of job search, and tend to overlook that job search is a multi-step process, encompassing sequential decisions taken by job seekers and firms. Moreover, these studies are somewhat at odds with related work, that shows that elapsed unemployment duration itself has a marked net effect on job seekers' and firms' behaviors.

In this paper, we use a unique empirical data source to shed light on how heterogeneity and unemployment duration affect the dynamics of the different phases of the job search process. We collect longitudinal granular information on Swiss unemployed job search activity, stemming from job search diaries filled in at PES. These documents contain information on all applications sent out by a job seeker in each month of her unemployment spell. Specifically, we know for each application whether the contacted firm calls back the applicant, and eventually makes her a job offer.

We examine the dynamics of job seekers' behavior using the monthly number of job applications as a search effort proxy. We document a slight gradual decrease in application effort over the course of unemployment in the raw data. Accounting for individual heterogeneity through individual fixed effects, we find evidence of a much sharper decline in our measure of search effort. This steepening of the duration profile in job seekers' behavior is due to positive dynamic selection: job seekers who experience longer unemployment spells send systematically more applications at any duration.

We study firms' responses to job applications exploiting our rich dataset at its most granular level, *i.e.* at the application level. We show descriptively that the (application-level) callback probability is strongly decreasing over time, whereas the rate at which interviews are converted into job offers, *i.e.* the job offer conversion probability, exhibits slight positive duration dependence. To disentangle the contribution of heterogeneity and duration to these patterns, we use an alternative identification strategy. For each application, we compute its *ex-ante* chance of getting a positive response from firms, based on firms' responses in the very early periods of unemployment, and conditioning on all information that is relevant to firms when they make callback decisions. Controlling for heterogeneity through these *ex-ante* chances, we find that dynamic selection is negative at the callback stage, hence accentuating the duration profile of the callback probability in the raw data. In contrast, our conditioning approach has little effect at the job offer conversion stage, whose duration profile remains slightly positive. These findings are consistent with our intuition that observed job seekers' characteristics are mostly used in the first phase of firms' screening process.

Building on our empirical evidence of duration dependence in job seekers' and firms' behaviors, we develop a job search model with informational frictions and statistical discrimination à la Jarosch and Pilossoph (2019). We expand their baseline model by adding an application phase, in which heterogeneous job seekers choose the optimal level of search effort to exert when contacting heterogeneous firms. Upon receiving an application, a firm decides whether to call back the applicant for a costly job interview. As job seeker's type is unknown to the firm in the callback phase, its callback decision is based on the observed time the job seeker has spent unemployed, which conveys a signal about her productivity. During the interview, the firm discovers the type of the candidate, and decides whether to hire her or not. Statistical discrimination by firms towards long-term unemployed generates negative net duration dependence in the callback probability, as well as non-negative duration dependence in the job offer conversion probability, due to the pool of interviewees getting more homogeneous. The overall reduction in job finding chances entails in turn an endogenous net decline in search effort on the job seeker's side, due to the reduction in the marginal benefit from applying.

Our study enriches our understanding of the dynamics of the job search process, by providing granular and comprehensive evidence on how job seekers' and firms' behaviors evolve along unemployment. Our results corroborate the previous finding that individual heterogeneity is a major driver of these dynamics. However, they also highlight that the role of heterogeneity is not uni-dimensional: dynamic selection is either positive or negative, depending on which side of the labor market we consider. In particular, job seekers with different employment prospects might choose to exert different levels of search effort endogenously, hence generating complex dynamics in the pool of unemployed that are observed all along the duration of unemployment.

Furthermore, our study shows that dynamic selection does not account for the whole story: elapsed unemployment duration itself directly affects agents' behaviors in the labor market. Specifically, the net effects of duration on job seekers' and firms' behaviors tend to further dampen the chances that a match between labor supply and labor demand materializes over time. In light of our structural model, these net duration effects are rooted in firms' statistical discrimination against long-term unemployed, who use the time a job seeker has spent unemployed as a key information to infer her true productivity. If this signal is of primary importance in firms' screening process, prolonged unemployment might have detrimental implications for job seekers' employment prospects, not only in terms of successfully getting through firms' recruiting process, but also with respect to their provision of search effort.

Our work points towards clear policy recommendations for reducing informational frictions in the labor market, already in the early phases of the job search process. The provision of detailed information on workers' productivity, skills or labor market history to firms, already at the moment when they make callback decisions, might reduce the weight put on the negative signal conveyed by prolonged unemployment. Such attenuation of informational frictions would not only makes firms' recruiting decisions more accurate, but would also prevent negative endogenous responses by job seekers, faced with declining job finding chances.

Bibliography

Ahn, H. J. and Hamilton, J. D. (2020). Heterogeneity and unemployment dynamics. *Journal of Business & Economic Statistics*, 38(3):554–569.

Alvarez, F. E., Borovičková, K., and Shimer, R. (2016). Decomposing duration dependence in a stopping time model. Technical report, National Bureau of Economic Research.

Arni, P. and Schiprowski, A. (2019). Job search requirements, effort provision and labor market outcomes. *Journal of Public Economics*, 169:65–88.

Beaman, L. and Magruder, J. (2012). Who gets the job referral? Evidence from a social networks experiment. *American Economic Review*, 102(7):3574–93.

Behncke, S., Frölich, M., and Lechner, M. (2010). A caseworker like me - Does the similarity between the unemployed and their caseworkers increase job placements? *The Economic Journal*, 120(549):1430–1459.

Blanchard, O. J. and Diamond, P. (1994). Ranking, unemployment duration, and wages. *The Review of Economic Studies*, 61(3):417–434.

Burks, S. V., Cowgill, B., Hoffman, M., and Housman, M. (2015). The value of hiring through employee referrals. *The Quarterly Journal of Economics*, 130(2):805–839.

DellaVigna, S., Heining, J., Schmieder, J. F., and Trenkle, S. (2022). Evidence on job search models from a survey of unemployed workers in germany. *The Quarterly Journal of Economics*, 137(2):1181–1232.

Ebrahimy, E. and Shimer, R. (2010). Stock–flow matching. *Journal of Economic Theory*, 145(4):1325–1353.

Eriksson, S. and Rooth, D.-O. (2014). Do employers use unemployment as a sorting criterion when hiring? Evidence from a field experiment. *American Economic Review*, 104(3):1014–39.

Faberman, R. J. and Kudlyak, M. (2019). The intensity of job search and search duration. *American Economic Journal: Macroeconomics*, 11(3):327–57.

Falk, A., Huffman, D., and Sunde, U. (2006). Do I have what it takes? Equilibrium search with type uncertainty and non-participation. *Working Paper*.

Falk, A., Lalive, R., and Zweimüller, J. (2005). The success of job applications: A new approach to program evaluation. *Labour Economics*, 12(6):739–748.

Farber, H. S., Silverman, D., and Von Wachter, T. (2016). Determinants of callbacks to job applications: An audit study. *American Economic Review*, 106(5):314–18.

Fluchtmann, J., Glenny, A. M., Harmon, N., and Maibom, J. (2021). The dynamics of job search in unemployment: Beyond search effort and reservation wages. *Working Paper*.

Galenianos, M. and Kircher, P. (2009). Directed search with multiple job applications. *Journal of economic theory*, 144(2):445–471.

Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. *American Economic Review*, 95(1):50–65.

Hensvik, L. and Skans, O. N. (2016). Social networks, employee selection, and labor market outcomes. *Journal of Labor Economics*, 34(4):825–867.

Jarosch, G. and Pilossoph, L. (2019). Statistical discrimination and duration dependence in the job finding rate. *The Review of Economic Studies*, 86(4):1631–1665.

Kroft, K., Lange, F., and Notowidigdo, M. J. (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *The Quarterly Journal of Economics*, 128(3):1123–1167.

Krueger, A. B., Mueller, A., Davis, S. J., and Şahin, A. (2011). Job search, emotional well-being, and job finding in a period of mass unemployment: Evidence from high frequency longitudinal data. *Brookings Papers on Economic Activity*, pages 1–81.

Lehmann, T. (2023). How directed is search in the labor market? Evidence from an online job board. *Working paper*.

Lockwood, B. (1991). Information externalities in the labour market and the duration of unemployment. *The Review of Economic Studies*, 58(4):733–753.

Marinescu, I. and Skandalis, D. (2021). Unemployment insurance and job search behavior. *The Quarterly Journal of Economics*, 136(2):887–931.

Matzkin, R. L. (2003). Nonparametric estimation of nonadditive random functions. *Econometrica*, 71(5):1339–1375.

Mueller, A. I. and Spinnewijn, J. (2023). Predicting long-term unemployment risk. *Working paper*.

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.

Nüß, P. (2018). Duration dependence as an unemployment stigma: Evidence from a field experiment in germany. Technical report, Economics Working Paper.

Oberholzer-Gee, F. (2008). Nonemployment stigma as rational herding: A field experiment. *Journal of Economic Behavior & Organization*, 65(1):30–40.

Pissarides, C. A. (2000). *Equilibrium Unemployment Theory, 2nd Edition*, volume 1 of *MIT Press Books*. The MIT Press.

Salop, S. C. (1973). Systematic job search and unemployment. *The Review of Economic Studies*, 40(2):191–201.

Van den Berg, G. J. (1990). Nonstationarity in job search theory. *The Review of Economic Studies*, 57(2):255–277.

Van den Berg, G. J. and Van Ours, J. C. (1996). Unemployment dynamics and duration dependence. *Journal of Labor Economics*, 14(1):100–125.

Vishwanath, T. (1989). Job search, stigma effect, and escape rate from unemployment. *Journal of Labor Economics*, 7(4):487–502.

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.

Zuchuat, J. (2023a). Estimating duration dependence in job search: the within estimation duration bias. *Working paper*.

Zuchuat, J. (2023b). How vacancy referrals affect job search: Evidence from job applications data. *Working paper*.

Appendix

A. Data and empirical measurements

Figure A1: Pre-defined job search forms

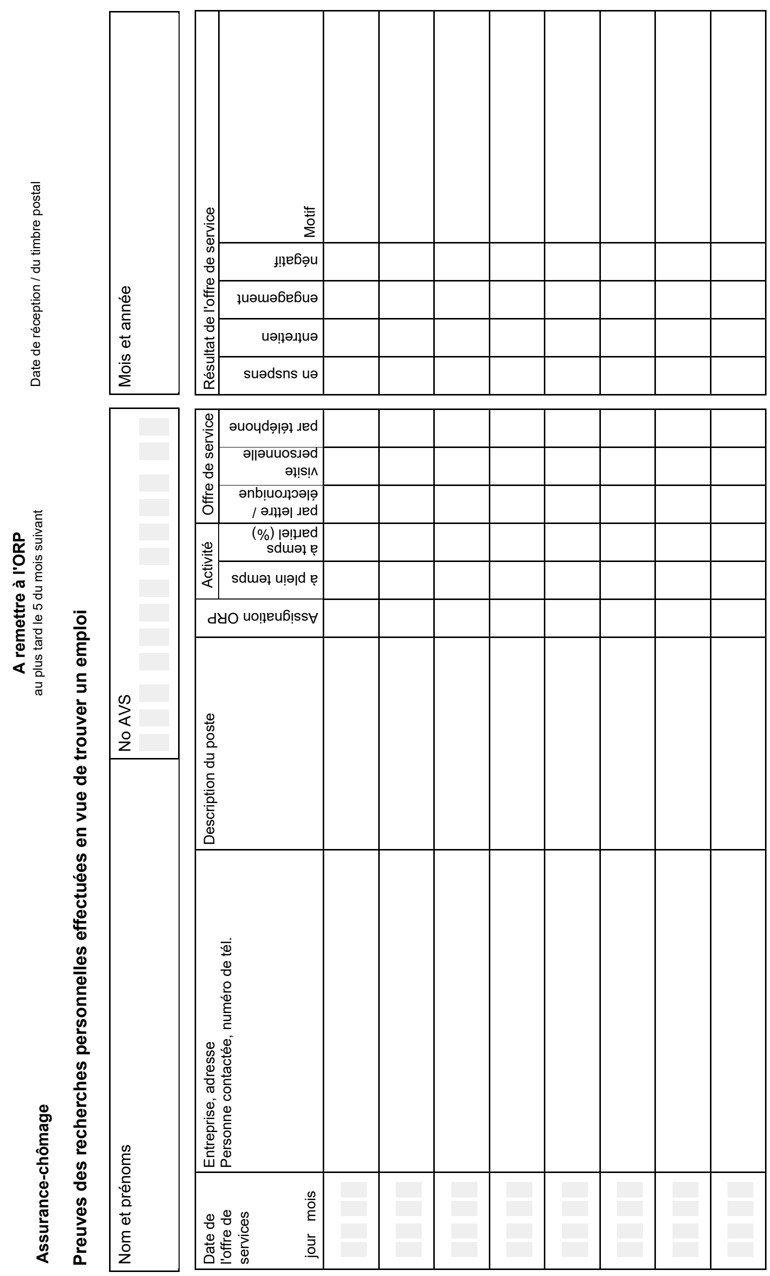

Assurance-chômage

A remettre à l'ORP
au plus tard le 5 du mois suivant

Date de réception / du timbre postal

Preuves des recherches personnelles effectuées en vue de trouver un emploi

Nom et prénoms	No AVS	Mois et année

Date de l'offre de services		Entreprise, adresse Personne contactée, numéro de tél.	Description du poste	Assignation ORP	Activité		Offre de service			Résultat de l'offre de service				
jour	mois				à plein temps	à temps partiel (%)	par lettre / électronique	visite personnelle	par téléphone	en suspens	entretien	engagement	négatif	Motif

Note: This figure presents the pre-defined job search diary form, in which unemployed record their job search activity.

Table A1: Job seekers' observed characteristics

	Main sample			***Auxiliary sample***		
	Mean	SDV	N	Mean	SDV	N
Individual characteristics						
Age	39.372	11.898	14798	39.307	10.651	655
1 = Female	0.458	0.498	14798	0.487	0.500	655
1 = Swiss	0.545	0.498	14798	0.539	0.499	655
1 = Primary education	0.269	0.444	14798	0.351	0.478	655
1 = Secondary education	0.588	0.492	14798	0.377	0.485	655
1 = Tertiary education	0.143	0.350	14798	0.189	0.392	655
1 = Manager	0.054	0.225	14798	0.092	0.289	655
1 = Specialist	0.598	0.490	14798	0.475	0.500	655
1 = Auxiliary	0.331	0.471	14798	0.423	0.494	655

Note: This table reports descriptive statistics on job seekers' socio-demographic characteristics, for the *Main sample* and *Auxiliary sample* of study.

Table A2: Descriptive statistics, *Auxiliary sample*

	Mean	SDV	Min	Median	Max	N
A. By application						
$\mathbb{P}(c_{ijt} = 1)$, callback prob. [in %]	7.396	26.171	0.000	0.000	100.000	24770
$\mathbb{P}(o_{ijt} = 1)$, job offer prob. [in %]	1.514	12.213	0.000	0.000	100.000	24770
$\mathbb{P}(o_{ijt} = 1 \mid c_{ijt} = 1)$, job offer conversion prob. [in %]	20.567	40.432	0.000	0.000	100.000	1559
B. By monthly-individual						
A_{it}, nbr. applications	8.900	4.597	1.000	9.000	36.000	2783
C_{it}, nbr. callbacks	0.560	1.259	0.000	0.000	21.000	2783
O_{it}, nbr. job offers	0.089	0.330	0.000	0.000	3.000	2783
$\mathbb{P}(C_{it} > 0)$, prob. a.l. one interview [in %]	28.926	45.350	0.000	0.000	100.000	2783
$\mathbb{P}(O_{it} > 0)$, prob. a.l. one job offer [in %]	7.833	26.874	0.000	0.000	100.000	2783
C. Sample structure						
Time-period			04.2012 - 03.2013			
Region			ZH			
Nbr. applications			24770			
Nbr. monthly-individual			2699			
Nbr. individuals			655			

Note: This table reports descriptive statistics about our *Auxiliary sample* of study. Panels A and B report descriptives on application-level and monthly-individual-level job search outcomes respectively. Panel C provides information about the sample structure.

Figure A2: Job offers and income trajectories

(A) Observed average income trajectories

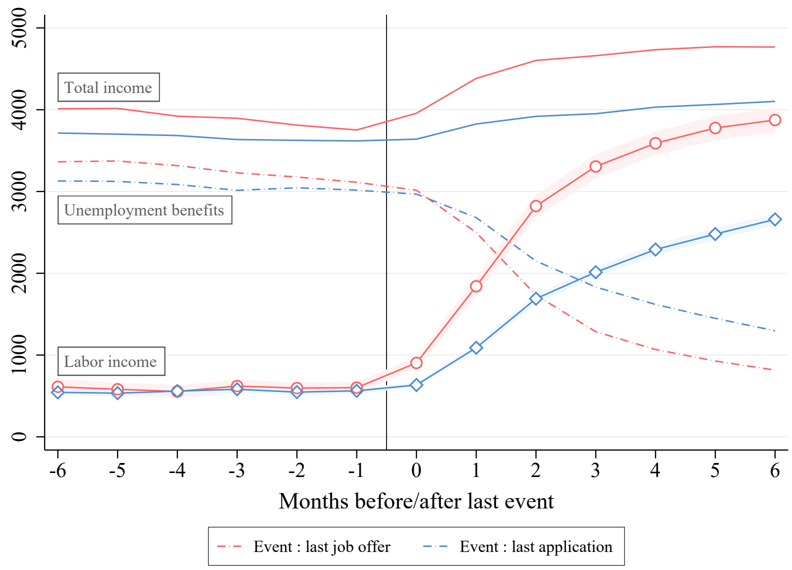

(B) Δ in labor income trajectories (accounting for heterogeneity)

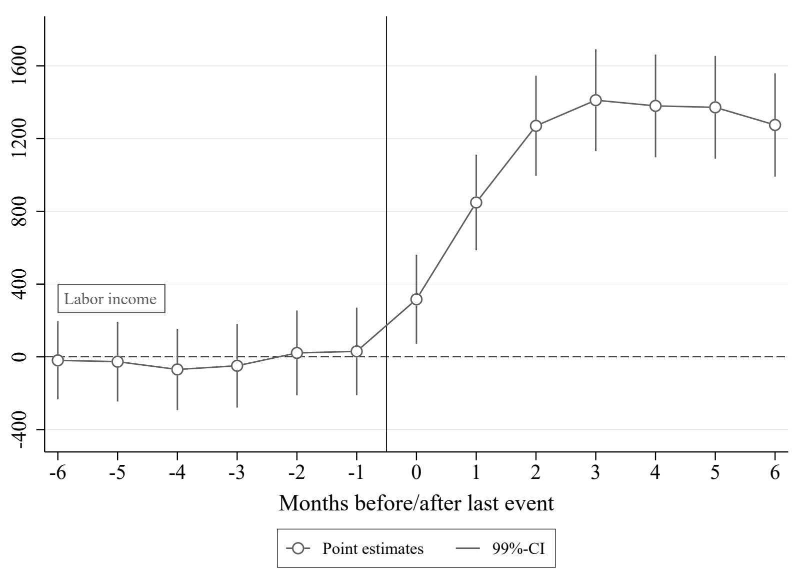

Note: This figure presents an event-study analysis, crossing information from the search diaries and the social security data. It highlights the informational content of the diaries data. Panel A shows the average evolution of total income, labor income and unemployment benefits in months before and after individual-specific events. For each individual, the event is either the last month when a job offer is recorded (in red, if at least one job offer is recorded in the observed data) or the last month when search diaries are collected (in blue, in the absence of job offer recorded). Panel B presents the results of a two-way fixed effects specification, to measure the differences in the labor income trajectories of the two above mentioned groups.

Figure A3: Monthly probability of at least one job offer and number of job offers

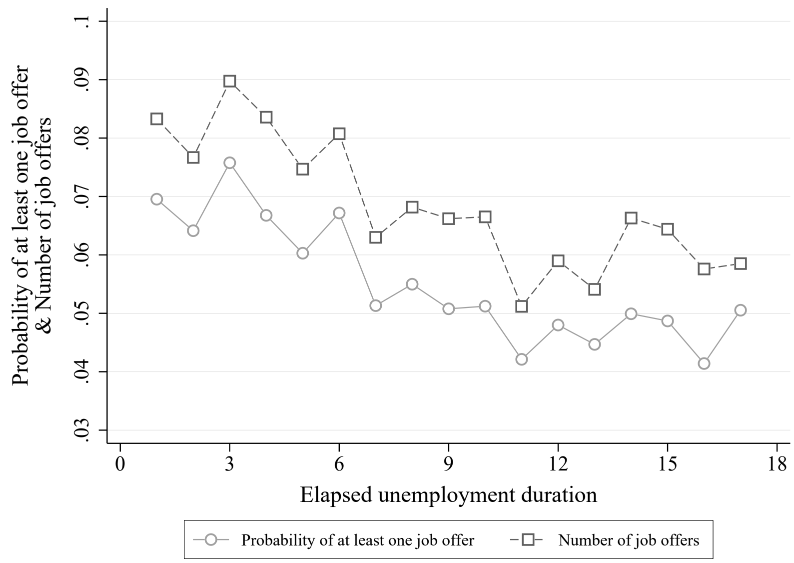

Note: This figure plots the average monthly probability of obtaining at least one job offer ($\mathbb{1}(O_{it} > 0)$, solid line) together with the average monthly number of job offers (O_{it}, dashed line).

Figure A4: Conceptual job search framework and empirical measurements of job search

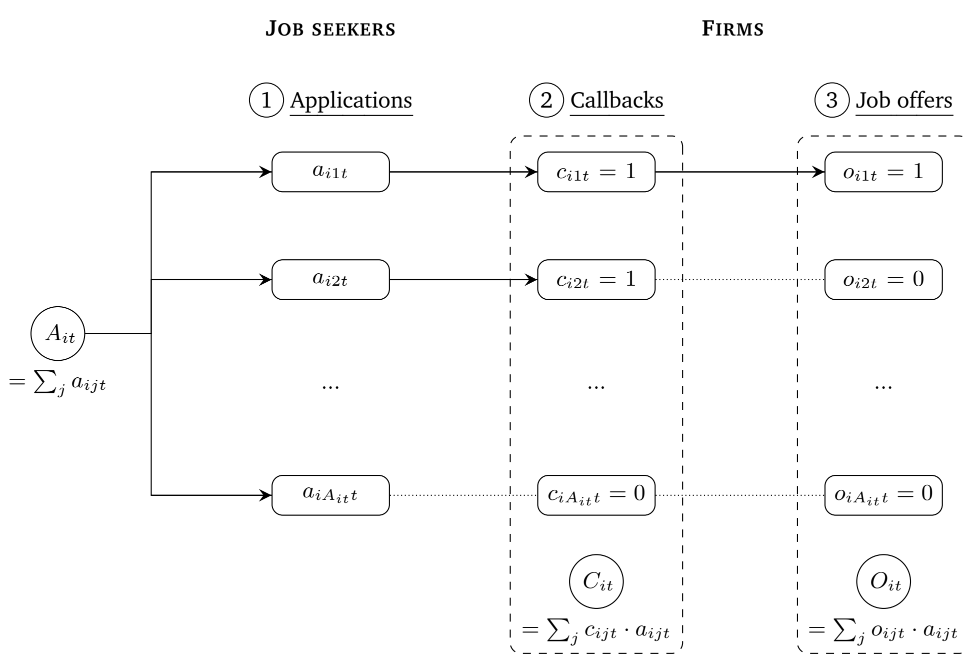

Note: This figure presents our conceptual job search framework, together with data measurements from job search diaries.

Table A3: Job search diaries, information coding

			Callback	Still open	Job offer	Negative	Freq.	%	Cum. %
Application	Callback	Job offer	☑	☑	☑	☑	93	0.015	0.015
			☑	☑	☑	☐	557	0.093	0.108
			☑	☐	☑	☑	55	0.009	0.117
			☑	☐	☑	☐	958	0.160	0.277
			☐	☑	☑	☑	78	0.013	0.290
			☐	☐	☑	☑	540	0.090	0.380
			☐	☑	☑	☐	500	0.083	0.463
			☐	☐	☑	☐	1506	0.251	0.714
		No job offer	☑	☑	☐	☑	3772	0.628	1.342
			☑	☐	☐	☑	5727	0.954	2.296
			☑	☑	☐	☐	4670	0.778	3.074
			☑	☐	☐	☐	3966	0.661	3.735
	No callback		☐	☑	☐	☑	125234	20.861	24.596
			☐	☐	☐	☑	277523	46.229	70.825
			☐	☑	☐	☐	141224	23.525	94.350
			☐	☐	☐	☐	33920	5.650	100.000
Total							600323	100.000	

Note: This table reports the distribution of the "types" of applications that are recorded in the job search diaries data, based on information contained in the 4 tick-boxes *Interview*, *Still open*, *Job offer* and *Negative*. Given the sequential decisions made by the firm (callbacks first, and then job offers), we impute a job interview for each application that records a job offer.

B. Job seekers: job applications

Figure B1: Dynamic selection with respect to application effort

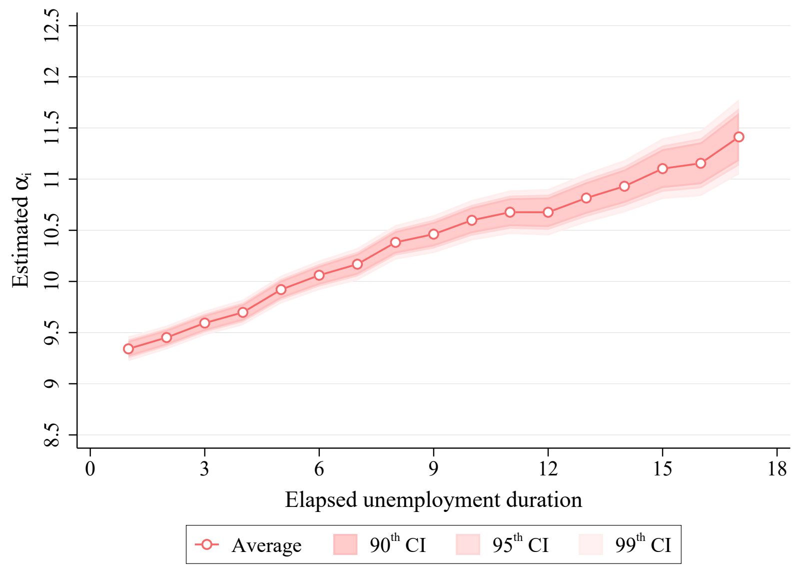

Note: This figure reports evidence on positive dynamic selection with respect to application effort. It plots the average estimated α_i from equation (4.1), per month of elapsed unemployment. For each month, only individuals that are observed in the raw data are considered when computing the average α_i.

Figure B2: Empirical duration dependence in excess application effort

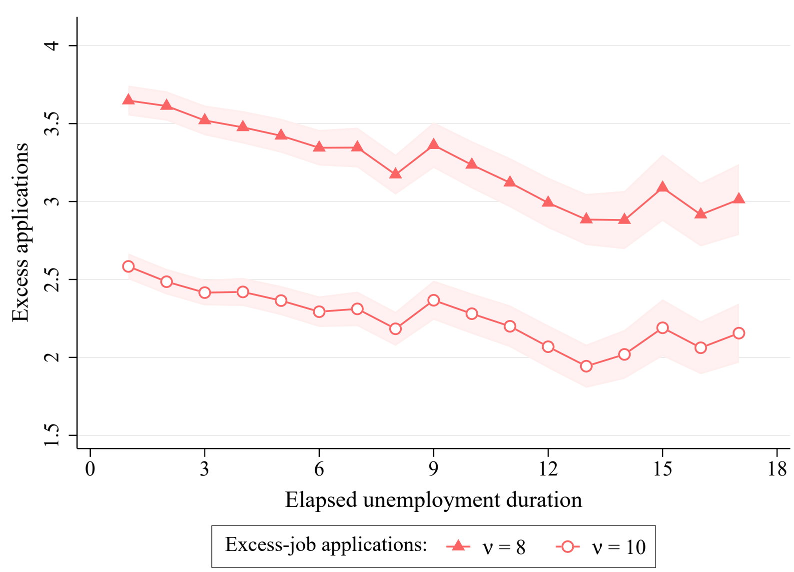

Note: This figure reports empirical evidence of duration dependence in our excess application effort measures.

Figure B3: Empirical distribution of application effort

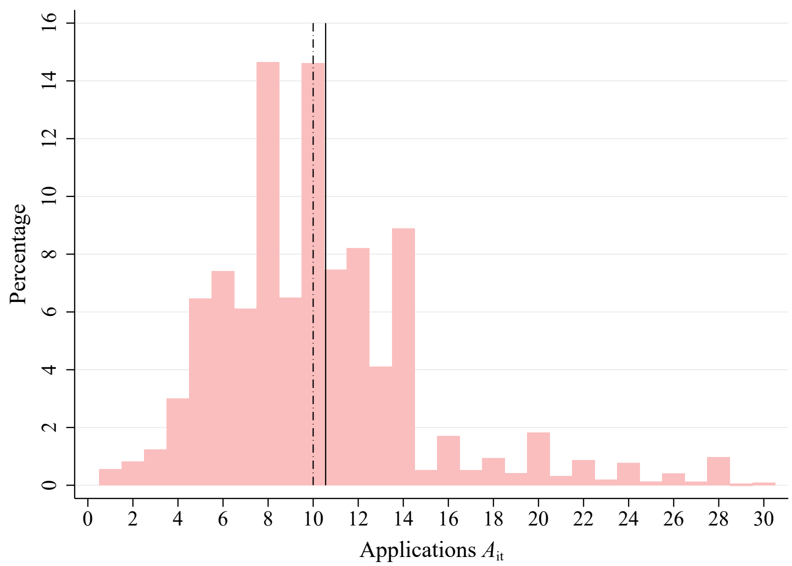

Note: This figure shows the empirical distribution of application effort, as measured by the number of applications sent out per month.

Figure B4: Duration dependence in application effort, subgroup analysis

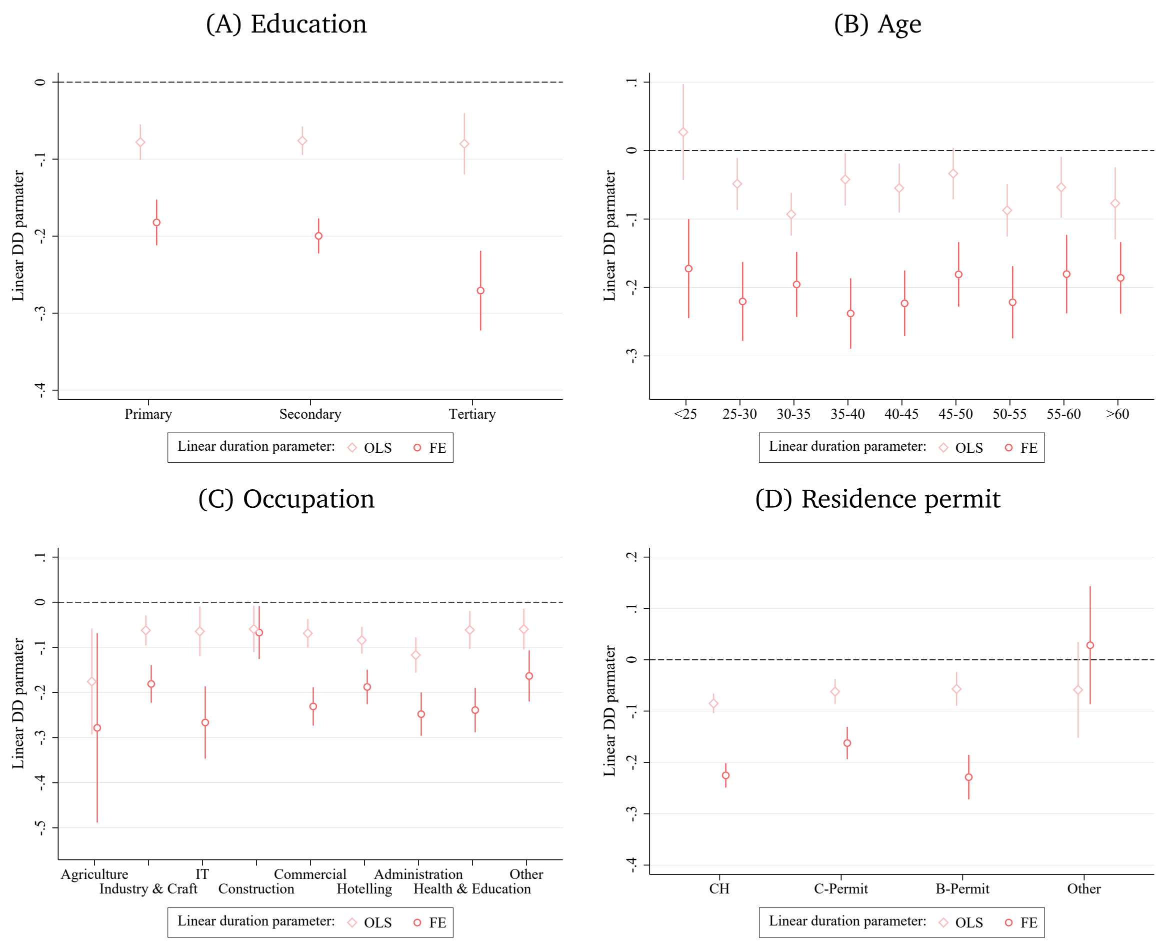

Note: This figure reports evidence of the empirical and net duration dependence for different subgroups of unemployed job seekers. We report estimates of equation (4.1) where the duration function $f^A(t;\phi^A)$ is specified linearly. We present both OLS and FE estimates (only including FE), for different sub-samples based on different sample-split variables.

Table B1: Partial correlations between estimated α_i and observed characteristics

	(1)		(2)		(3)	
Dependent variable: estimated α_i						
Age						
$25-30$	0.021	(0.148)	0.020	(0.147)	-0.002	(0.146)
$30-35$	-0.106	(0.148)	-0.150	(0.148)	-0.147	(0.145)
$35-40$	-0.231	(0.153)	-0.271*	(0.152)	-0.253*	(0.150)
$40-45$	-0.285*	(0.153)	-0.329**	(0.151)	-0.332**	(0.149)
$45-50$	-0.302*	(0.154)	-0.298*	(0.153)	-0.259*	(0.153)
$50-55$	-0.382**	(0.157)	-0.385**	(0.155)	-0.388**	(0.155)
$55-60$	-0.688***	(0.168)	-0.685***	(0.167)	-0.661***	(0.166)
> 60	-2.496***	(0.194)	-2.518***	(0.192)	-2.478***	(0.191)
Residential status						
C-permit	0.479***	(0.093)	0.442***	(0.092)	0.446***	(0.092)
B-permit	0.505***	(0.110)	0.477***	(0.109)	0.467***	(0.108)
Other permit	0.242	(0.235)	0.188	(0.232)	0.320	(0.230)
Education						
Apprentice.	0.134	(0.093)	0.106	(0.092)	0.099	(0.091)
High school	0.453**	(0.185)	0.357*	(0.186)	0.345*	(0.183)
Professional mat.	-0.050	(0.173)	-0.054	(0.171)	-0.007	(0.172)
UAS	-0.125	(0.220)	-0.237	(0.220)	-0.198	(0.224)
University	-0.491***	(0.158)	-0.665***	(0.162)	-0.664***	(0.160)
Female	0.175**	(0.082)	0.177**	(0.081)	0.109	(0.082)
Labor market history						
$\ln$(previous wage)	0.379***	(0.069)	0.378***	(0.069)	0.312***	(0.068)
Unemployment history	-0.495*	(0.268)	-0.486*	(0.266)	-0.502*	(0.267)
Occupation						
Industry & Craft	-2.235***	(0.284)	-2.190***	(0.283)	-2.271***	(0.282)
IT	-2.730***	(0.311)	-2.691***	(0.310)	-2.815***	(0.310)
Construction	-1.580***	(0.296)	-1.565***	(0.295)	-1.526***	(0.294)
Commercial	-1.181***	(0.282)	-1.162***	(0.281)	-1.310***	(0.282)
Hotelling	-1.224***	(0.282)	-1.275***	(0.281)	-1.361***	(0.284)
Administrative	-1.231***	(0.290)	-1.215***	(0.289)	-1.403***	(0.289)
Health & Educ.	-2.599***	(0.293)	-2.608***	(0.293)	-2.573***	(0.294)
Other	-2.451***	(0.300)	-2.435***	(0.298)	-2.578***	(0.298)
Canton						
SG	-2.460***	(0.093)				
VD	2.349***	(0.132)				
ZG	1.305***	(0.137)				
ZH	0.623***	(0.094)				
Constant	7.851***	(0.614)	7.891***	(0.633)	6.183***	(0.761)
Instituitions	Canton		PES		CW	
F-stat. instituitons	709.821		135.287		11.049	
p-value instituitons	0.000		0.000		0.000	
Observations	14798		14798		14798	
Adj-R^2	0.173		0.187		0.227	

Note: This table reports estimates of a multivariate linear regression, where we regress the estimated α_i from equation (4.1) on observed individual characteristics. Three models are reported, differing with respect to the policy controls included as regressors (cantons, PES offices or caseworkers fixed effects).

Table B2: Duration dependence in application effort, Poisson pseudo maximum likelihood

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable: application effort A_{it}						
Elapsed unemployment duration	-0.009***	-0.006***	-0.004***	-0.004***	-0.020***	-0.021***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
	[-0.097]	[-0.069]	[-0.048]	[-0.050]	[-0.226]	[-0.230]
Individual controls	No	Yes	Yes	Yes	No	Yes
Policy controls	No	No	Yes	Yes	No	Yes
LLMC	No	No	No	Yes	No	Yes
Individual FE	No	No	No	No	Yes	Yes
Mean outcome 1st month	11.107	11.107	11.107	11.107	11.107	11.107
Pseudo-R^2	0.002	0.014	0.066	0.071	0.201	0.206
N. observations	55559	55559	55559	55559	55559	55559

Note: This table reports empirical estimates of equation (4.1) using a Poisson pseudo maximum likelihood estimator, where the parametric duration function $f^A(t;\phi^A)$ is specified linearly. Models are estimated on a restricted sample, that discards individuals who do not exhibit within-variation in application effort. Each column sequentially adds a set of controls or FE. Errors are clustered at the individual level and are reported in parentheses. Absolute coefficients (measuring the monthly decrease in application effort) are indicated in squared brackets and are directly comparable to our OLS baseline estimates. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table B3: Duration dependence in application effort, alternative application effort measures

	Excess applications				Private applications	
	$\underline{A} = 8$		$\underline{A} = 10$			
	(1)	(2)	(3)	(4)	(5)	(6)
A. OLS						
Elapsed unemployment duration	-0.069***	-0.201***	-0.058***	-0.179***	-0.099***	-0.202***
	(0.008)	(0.022)	(0.007)	(0.022)	(0.009)	(0.022)
Individual controls	No	Yes	No	Yes	No	Yes
Policy controls	No	Yes	No	Yes	No	Yes
LLMC	No	Yes	No	Yes	No	Yes
Individual FE	No	Yes	No	Yes	No	Yes
Mean outcome 1st month	3.707	3.707	2.754	2.754	10.452	10.452
Adjusted-R^2	0.005	0.393	0.004	0.338	0.008	0.468
N. observations	45901	45901	39563	39563	51305	51305
B. Poisson						
Elapsed unemployment duration	-0.019***	-0.057***	-0.022***	-0.070***	-0.010***	-0.020***
	(0.002)	(0.006)	(0.003)	(0.008)	(0.001)	(0.002)
	[-0.071]	[-0.213]	[-0.060]	[-0.193]	[-0.101]	[-0.205]
Individual controls	No	Yes	No	Yes	No	Yes
Policy controls	No	Yes	No	Yes	No	Yes
LLMC	No	Yes	No	Yes	No	Yes
Individual FE	No	Yes	No	Yes	No	Yes
Mean outcome 1st month	3.707	3.707	2.754	2.754	10.452	10.452
Pseudo-R^2	0.004	0.328	0.004	0.334	0.003	0.200
N. observations	45901	45901	39563	39563	51305	51305

Note: This table reports empirical estimates of equation (4.1) for our alternative job search effort measures (excess application effort and private applications), where the parametric duration function $f^A(t;\phi^A)$ is specified linearly. Models are estimated using OLS (panel A) or Poisson pseudo maximum likelihood (panel B). For each independent variable, we consider either a bivariate model or the full specification. Errors are clustered at the individual level and are reported in parentheses. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table B4: Duration dependence in application effort, no exit months

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable : application effort A_{it}						
Elapsed unemployment duration	-0.082***	-0.056***	-0.037***	-0.041***	-0.190***	-0.215***
	(0.008)	(0.008)	(0.007)	(0.007)	(0.010)	(0.021)
	[-0.750%]	[-0.518%]	[-0.343%]	[-0.378%]	[-1.747%]	[-1.975%]
Individual controls	No	Yes	Yes	Yes	No	Yes
Policy controls	No	No	Yes	Yes	No	Yes
LLMC	No	No	No	Yes	No	Yes
Individual FE	No	No	No	No	Yes	Yes
Mean outcome 1$^{\text{st}}$ month	10.846	10.846	10.846	10.846	10.846	10.846
adj.-R^2	0.006	0.035	0.179	0.193	0.495	0.502
N. observations	56646	56646	56646	56646	56646	56646

Note: This table reports empirical estimates of equation (4.1), where the parametric duration function $f^A(t; \phi^A)$ is specified linearly. Models are estimated on a restricted sample, that discards individual-monthly observations when an unemployment exit is observed. Each column sequentially adds a set of controls or FE. Errors are clustered at the individual level and reported in parentheses. Coefficients in relative terms (with respect to the average in the first month of unemployment) are indicated in squared brackets. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

C. Firms' responses: callbacks and job offers

Figure C1: Duration dependence in the callback probability, at or beyond UB exhaustion

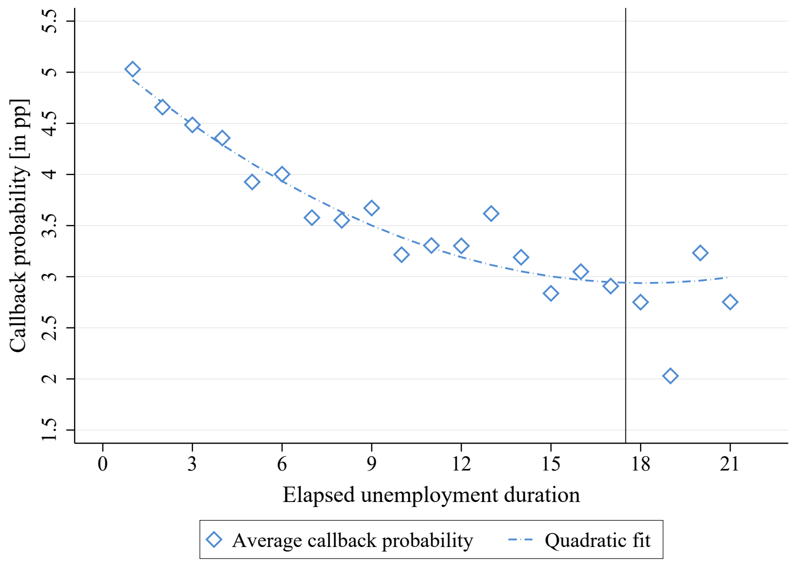

Note: This figure plots the average callback probability per month of elapsed unemployment, including months at or beyond unemployment benefits exhaustion. A quadratic fit regression line is added on the graph to emphasize the duration profile of the callback probability.

Figure C2: Duration dependence in the job offer probability

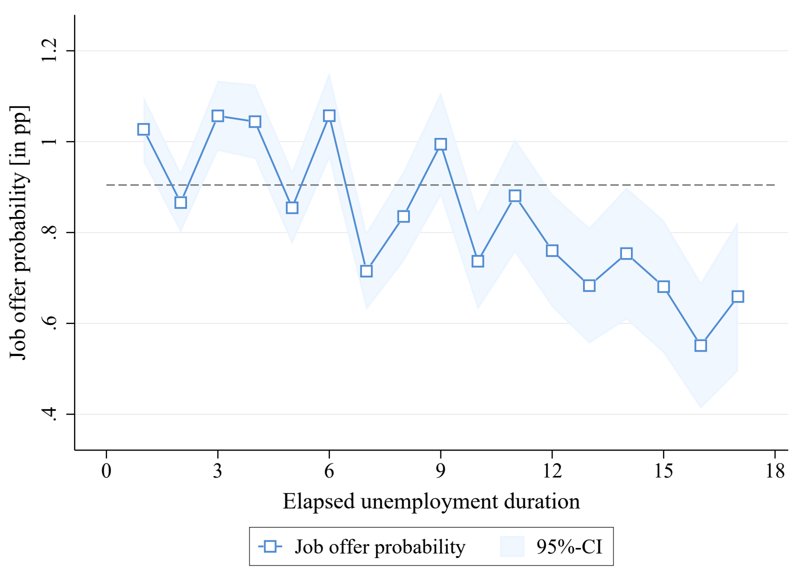

Note: This figure plots the average job offer probability, *i.e.* the probability that an application leads to a job offer, against elapsed unemployment duration. Application-level observations are weighted by the inverse of the monthly number of applications sent out by individual i in month t, so as to put equal weight on all monthly-individual observations.

Figure C3: Concentration of callbacks and job offers at the end of the spell

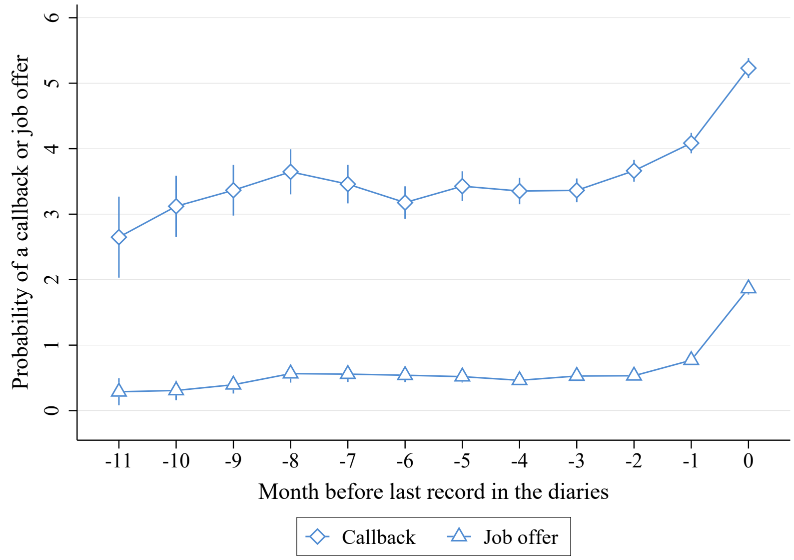

Note: This figure plots the average probability of a callback and a job offer (computed on all applications), for each month prior to the last month of record in the job search diaries. Application-level observations are weighted by the inverse of the monthly number of applications sent out by individual i in month t, so as to put equal weight on all monthly-individual observations. 95% confidence intervals for the average probabilities are reported.

Figure C4: Duration dependence in job offer conversion probability
First interview vs. Second and more interviews

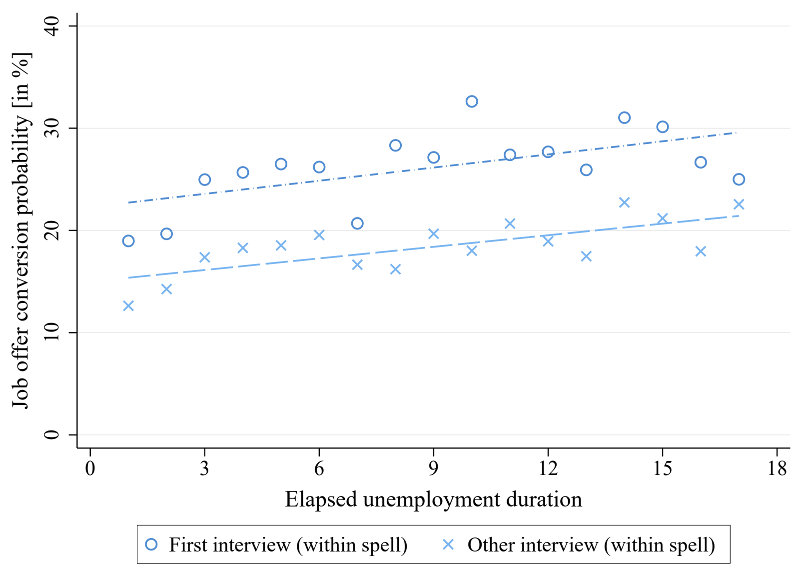

Note: This figure plots the average job offer conversion probability, distinguishing between the first interview recorded per unemployment spell, and all subsequent interviews.

Figure C5: Job seekers' heterogeneity
Application effort and *ex-ante* chances of positive replies by firms

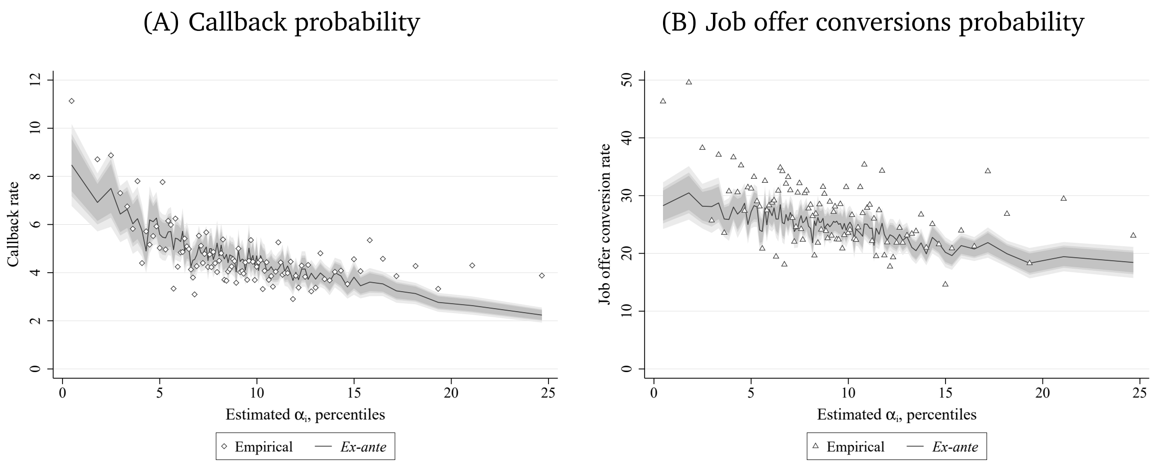

Note: This figure reports evidence on the relationship between (i) the estimated application effort fixed effects α_i and (ii) job seekers' average chances of getting a positive response from firms. The latter are computed for the callback (panel A) and job offer conversion (panel B) probabilities, by averaging application-level information across all applications sent by a job seeker (only those which led to a callback for panel B). Those job seekers' average chances of a callback or job offer conversion are based either on firms' responses that are directly observed in the data (empirical), or on their *ex-ante* counterparts (*ex-ante*).

Table C1: *Ex-ante* probabilities, estimation in the reference month

	Callback probability		**Job offer conversion probability**	
	Marginal effects	SE	Marginal effects	SE
Age				
$25-30$	-0.158	(0.469)	3.739*	(2.095)
$30-35$	-0.812*	(0.447)	2.052	(2.068)
$35-40$	-0.561	(0.474)	0.841	(2.145)
$40-45$	-0.997**	(0.455)	-1.084	(2.113)
$45-50$	-0.889*	(0.458)	1.113	(2.144)
$50-55$	-1.241***	(0.466)	-0.189	(2.268)
$55-60$	-2.241***	(0.513)	3.656	(2.588)
> 60	-3.693***	(0.494)	5.189	(3.605)
Residential status				
C-permit	-1.164***	(0.270)	-1.548	(1.340)
B-permit	-1.260***	(0.292)	1.624	(1.617)
Other permit	-1.293*	(0.765)	-2.316	(3.820)
Education				
Apprentice.	2.313***	(0.249)	-4.904***	(1.618)
High school	1.571***	(0.508)	-6.619**	(2.934)
Professional mat.	3.786***	(0.517)	-5.035**	(2.457)
UAS	3.784***	(0.612)	-5.054*	(2.714)
University	3.326***	(0.486)	-11.269***	(2.201)
Female	0.360	(0.234)	-0.569	(1.171)
Labor market history				
$\ln$(previous wage)	1.225***	(0.207)	-2.459**	(0.982)
Unemployment history	-3.471***	(0.822)	2.117	(4.179)
Application process				
Phone	-0.044	(0.215)	6.177***	(1.513)
Personal	7.580***	(0.429)	6.900***	(1.202)
CW referral	3.670***	(0.383)	1.068	(2.180)
Search effort α_i	-0.269***	(0.041)	-0.205	(0.154)
Policy controls	CW		PES	
LLMC	Yes		Yes	
Observations	153316		12060	
Pseudo-R^2	0.107		0.057	

Note: This table reports the empirical estimates of equations (5.2a) and (5.2b), in the references months τ_i^A and τ_i^C. Coefficients are reported as average marginal effects (in pp). Errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table C2: *Ex-ante* probabilities with non-CV characteristics, estimation in the reference month

	Callback probability		**Job offer conversion probability**	
	Marginal effects	SE	Marginal effects	SE
Age				
$25-30$	0.086	(0.444)	3.914*	(2.057)
$30-35$	-0.573	(0.423)	2.265	(2.033)
$35-40$	-0.229	(0.453)	1.362	(2.111)
$40-45$	-0.604	(0.435)	-0.605	(2.077)
$45-50$	-0.441	(0.441)	1.994	(2.117)
$50-55$	-0.718	(0.455)	0.622	(2.236)
$55-60$	-1.652***	(0.516)	5.181**	(2.605)
>60	-2.935***	(0.517)	7.413**	(3.721)
Residential status				
C-permit	-1.118***	(0.269)	-1.091	(1.344)
B-permit	-1.224***	(0.290)	1.978	(1.616)
Other permit	-1.287*	(0.755)	-2.622	(3.778)
Education				
Apprentice.	2.190***	(0.253)	-5.283***	(1.618)
High school	1.487***	(0.502)	-7.147**	(2.914)
Professional mat.	3.657***	(0.511)	-5.344**	(2.442)
UAS	3.598***	(0.610)	-5.254*	(2.695)
University	3.033***	(0.472)	-11.134***	(2.215)
Female	0.339	(0.234)	-0.707	(1.169)
Labor market history				
$\ln$(previous wage)	1.122***	(0.204)	-2.173**	(0.982)
Unemployment history	-3.576***	(0.820)	2.076	(4.188)
Application process				
Phone	-0.003	(0.215)	6.317***	(1.513)
Personal	7.687***	(0.426)	6.829***	(1.194)
CW referral	3.597***	(0.383)	1.387	(2.164)
Search effort α_i	-0.261***	(0.040)	-0.186	(0.151)
Non-CV characteristics				
Employability grade CW	0.977***	(0.245)	0.512	(1.137)
1 = Experienced sickness	-1.817***	(0.245)	-8.432***	(1.179)
Mobility degree				
Daily commute	-6.945	(5.489)	10.017	(8.214)
Part of the country	-6.367	(5.506)	10.830	(8.771)
Whole country	-5.613	(5.560)	5.499	(9.178)
International	-3.217	(5.720)	-3.506	(9.095)
Policy controls	CW		PES	
LLMC	Yes		Yes	
Observations	153316		12060	
Pseudo-R^2	0.112		0.065	

Note: This table reports the empirical estimates of equations (5.2a) and (5.2b), in the references months τ_i^A and τ_i^C, adding characteristics that are not observed on the CV. Coefficients are reported as average marginal effects (in pp). Errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

D. Rationalizing the evidence: a job search model with statistical discrimination

Model's details

A worker with permanent ability x maximizes the present-discounted value of her utility streams by solving the following utility maximization problem in recursive form:

$$V\left(a, n, u(\tau); x\right) = \max_{q, a', s(\tau)} \frac{q^{1-\gamma}}{1-\gamma} - \sigma\Big(s(\tau)\Big) + \beta V\Big(a', n', u'(0), u'(\tau+1); x\Big) \tag{D.1}$$

$$\text{s.t.} \quad q + a' = Ra + \omega n + bu(\tau) \tag{D.2}$$

$$n' = (1-\delta_L)n + s(\tau)o(x,\tau)u(\tau) \tag{D.3}$$

$$u'(\tau+1) = \big(1 - s(\tau)o(x,\tau)\big)u(\tau) \tag{D.4}$$

$$u'(0) = \delta_L n \tag{D.5}$$

$$1 = n + u(\tau) \tag{D.6}$$

where $n = \mathbb{1}\{\text{employed}\}$ and $u(\tau) = \mathbb{1}\{\text{unemployed at duration } \tau\}$ are indicator functions for the employment and unemployment states, respectively. As a result, at each time t the worker is either employed ($n = 1$) or unemployed at duration τ ($u(\tau) = 1$), as spelled out by equation (D.6). Equation (D.2) represents the flow budget constraint. Equations (D.3) to (D.5) express the probability of being employed, unemployed at longer duration and unemployed with zero duration tomorrow.

Figure D1: Callback probability by education level

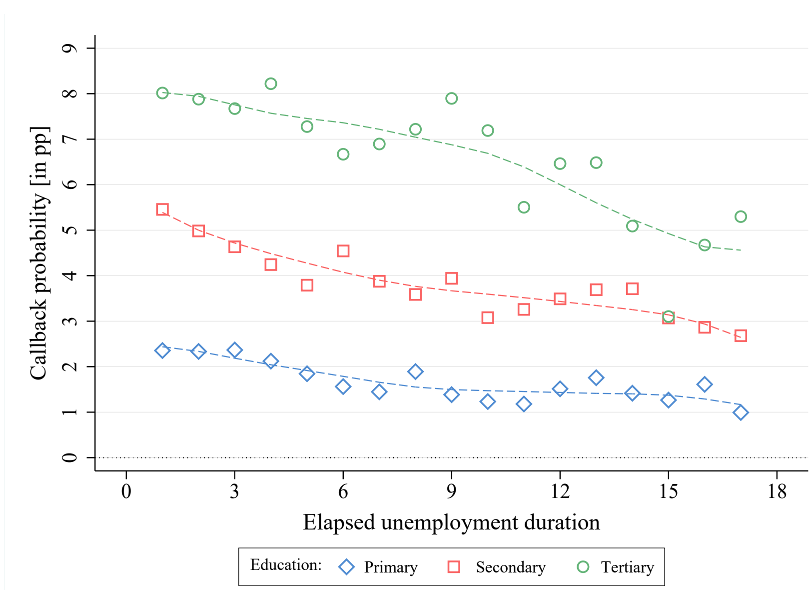

Note: This figure plots the average callback probability of applications sent by job seekers, distinguishing between three levels of education (primary, secondary, tertiary).

CHAPTER 2

How Vacancy Referrals affect Job Search: Evidence from Job Applications Data

Jeremy Zuchuat

University of Lausanne

Abstract

Vacancy referrals by caseworkers represent an effective and widely-used job search assistance tool. However, little is known about the mechanisms through which they affect job finding. This paper uses granular longitudinal job applications data with firms' responses to shed light on those mechanisms. Controlling for individual-by-unemployment duration fixed effects, we find that referred applications have a much higher probability of success (callback or job offer). These effects are driven by the fact that referred vacancies exhibit different and potentially lower-paying characteristics. Aggregating data to the monthly-individual level and using caseworker stringency as an instrument, we show that job seekers apply more in months when the referral policy is used, in spite of a slight reduction in their private application effort. This eventually translates into a large increase in the number of job interviews and offers they obtain per month.

Keywords: vacancy referrals, job search, active labor market policies, program evaluation, applications data, stringency instrument.

JEL: J63, J64, J68

1. Introduction

Job Search Assistance (JSA) represents a key element of the Active Labor Market Policies (ALMP) toolbox developed in western economies to reduce unemployment (Card et al., 2010, 2018). This widespread set of policies seeks to attenuate job search frictions, so as to increase matching efficiency between unemployed workers and recruiting firms. A crucial aspect of the JSA architecture lies in the intermediary assisting job seekers. This can take the form of automated informational tools (*e.g.* an informational brochure in Altmann et al., 2018; an online recommendation algorithm in Belot et al., 2019 or Hensvik et al., 2023), but in most cases JSA services are provided by caseworkers at Public Employment Service (PES) offices.

One of the main JSA instruments at caseworkers disposal is job vacancy referrals. This policy consists in the transmission of information on vacancies from caseworkers to job seekers. Depending on the context, it can be accompanied by the obligation to apply to the referred positions, which confers both an informational and coercive aspect on this policy. Several studies have emphasized the positive effect of caseworker referrals on job finding (Fougère et al., 2009; Bollens and Cockx, 2017; Van Den Berg et al., 2019; Cheung et al., 2019). However, none has yet delivered a comprehensive picture of their effects on the job search process. Not only does most of the literature remain silent about the mechanisms through which referrals affect job search outcomes, but this policy has never been analyzed at the relevant treatment unit, *i.e.* job applications. Furthermore, the dual aspect of referrals, in-between a purely informational and a sanctioning tool, has only been discussed to a limited extent.

This paper addresses these gaps simultaneously: it represents the first empirical analysis based on job applications data covering firms' responses, that provides an in-depth analysis of caseworker referrals effects. Our study does not uniquely examine the ultimate impact of this JSA tool on job finding, but it also sheds light on its underlying mechanisms to better understand how the policy affects the various steps of job search. Our analysis relies on a unique dataset stemming from monthly job search diaries, first used in Zuchuat et al. (2023). Those data provide granular information on unemployed job search process. Not only, do they cover all applications made by job seekers in each month of their unemployment spells, but they also contain information on applications treatment status (referred, non-referred) and outcomes (callback, job offer). Following the intuitions developed in the theoretical framework of Fougère et al. (2009), we build empirical indices from diaries data to measure the effects of referrals on job search activity and success.

We first proceed to an application-level analysis to study whether caseworker referrals affect applications success probability and the characteristics of jobs unemployed apply to. Identification is ensured by a within-estimation procedure with high-dimensional fixed effects, at the individual-by-duration level. This approach is made possible by the three-dimensional nature of our data (individuals, unemployment duration and applications), and accounts for both time-varying individual heterogeneity and allocation to treatment based on individuals' expected outcomes. Our results show that referred applications have a much higher probability of success compared to non-referred applications, whether applications success is measured via the callback (+4.0 pp or + 105%) or job offer probability (+ 1.3 pp or + 143%). Our analysis also reveals the existence of sizable heterogeneous treatment effects, especially with respect to education level. We apply the same methodology to study how referrals affect two vacancies characteristics that are revealing of jobs remuneration: work-time percentage and occupational sector. Our findings show that referred applications have a higher probability of targeting part-time jobs, specifically among job seekers who previously held full-time positions. They also reveal that referrals focus more on positions in job seekers' desired occupational sectors, which are characterized by lower cognitive and higher physical skills requirements.

In a second step, we study how caseworker referrals affect application effort and matching in the labor market. We aggregate applications data to the individual-monthly level and build search effort (number of applications made per month) and labor matching (number of callbacks and job offers obtained per month) indices. Data aggregation prevents us from applying the same high-dimensional fixed effects approach as previously. To circumvent endogeneity issues, we develop an alternative identification strategy which relies on the quasi-random allocation of unemployed to caseworkers within PES offices (Arni and Schiprowski, 2019). Given that caseworkers exert some leverage over the use of referrals, this quasi-random allocation generates exogenous variation in the exposure to the policy across unemployed. Exploiting this institutional feature, we build a leave-one-out caseworker stringency instrument, which we use to causally estimate the effects of the policy on monthly-individual outcomes. Our results show that caseworker referrals significantly increase job seekers' total application effort in treatment months (0.70 more applications per month per additional referral). This increase is accompanied by a slight reduction in the monthly number of non-referred or private applications (0.30 less applications per month). This limited substitution effect suggests that positions targeted by referred and non-referred applications are imperfectly substitutable, consistently with our application-level results. The positive effects of referrals on applications success probability and on total application effort eventually translate into a large increase in the monthly matching rate: one additional referral leads to 0.15 more interviews and 0.03 more job offers per month on average.

Our empirical findings confirm that caseworker referrals represent an effective tool to improve unemployed job finding chances. They additionally provide better understanding of the mechanisms driving this positive effect: the higher matching rate entailed by referrals is due to both a higher success probability of job applications and a higher contacting rate between workers and firms, due to higher application effort. Moreover, our results suggest that improvements in employment prospects are possibly realized at the cost of accessing less well-paid jobs, which points towards the coercive aspect of the policy.

Our paper contributes to three main strands of the literature. It first addresses the growing literature on the effects of caseworker referrals on job search outcomes. This ALMP has first been studied by Fougère et al. (2009), who develop a partial equilibrium job search model with two different search channels (private and vacancy referrals). Their framework suggests that applications sent through the referral channel might have a disincentivizing effect on private search effort, and eventually on job finding. Estimating their model on French unemployment spells data, the authors find that vacancy referrals increase unemployment exit rate, especially in the strata of low-educated and low-skilled unemployed workers. Using Danish data, Van den Berg et al. (2012) find a positive impact of repeated meetings between unemployed and caseworkers on job finding. The authors argue that this effect is explained by an increase in the number of referrals assigned during the meetings. More direct evidence of referrals effects are presented by Bollens and Cockx (2017), who explore the effects of this policy in the Flemish labor market. Based on a timing-of-event approach, the authors find that three different types of referrals have positive effects on the transition rate to employment, both in the short and long-run. Using the same type of approach, Van Den Berg et al. (2019) find similar positive effects in the German context, where non-compliance to referrals can be sanctioned. Moreover, the authors provide evidence that this policy impacts the types of jobs eventually taken up by job seekers, as it leads to less stable employment spells and lower wages. Finally, Cheung et al. (2019) provide further indirect evidence of referrals effects in the context of Sweden. Using an experimental approach, they study the direct and displacement effects of JSA programs. The authors find that a more intensive use of JSA services tends to reduce unemployment among treated individuals, at the cost of substantial displacement effects for the non-treated population. Based on their observation of caseworkers' behavior, they argue that the positive effect on treated job seekers is mostly driven by an increase in the use of vacancy referrals.

Our study differs from previous work on caseworker referrals in several dimensions. First, it is innovative with respect to its methodological approach. Exploiting within-individual variation in treatment assignment and job search outcomes, we are able to account for individual heterogeneity and other potential endogeneity issues in a straightforward manner. Second, our analysis is not limited to the study of the overall effect of referrals. Using granular applications data, we are able to dig deeper into the mechanisms explaining why this policy has been found to successfully affect job finding. Finally, our study goes beyond providing indirect evidence of caseworker referrals effects: applications data enable us to measure treatment and outcomes at the relevant treatment unit, *i.e.* job applications, and to emphasize the policy effects in a direct manner.

Our paper additionally contributes to the general literature on JSA services. Starting from the seminal meta-studies by Card et al. (2010, 2018), several papers have explored more thoroughly how effective these policy instruments are at improving job seekers' employment prospects. Part of this literature has focused on job search assistance provided by public institutions (*e.g.* Wunsch, 2013; Manoli et al., 2018), while another has examined the effects of services offered by private companies (*e.g.* Crépon et al., 2013; Battisti et al., 2019). Those studies usually find positive direct effects of JSA on employment perspectives, with limited differences between publicly and privately provided services (Bennmarker et al., 2013; Krug and Stephan, 2013; Cottier et al., 2018). Accounting for general equilibrium effects leads to more contrasted conclusions though (Cahuc and Le Barbanchon, 2010; Crépon et al., 2013). In addition to classical JSA services provided by public or private counselors, researchers have also explored the potential of alternative automated JSA tools (Altmann et al., 2018; Belot et al., 2019; Briscese et al., 2020; Hensvik et al., 2023). Those have been found to affect markedly job search scope and effort provision, while their effect on job finding has not yet been clearly established.

From a methodological perspective, our paper also contributes to the large literature using judges or caseworkers' stringency as instrumental variables to estimate the treatment effects of court decisions or welfare programs (*e.g.* Kling, 2006; French and Song, 2014; Aizer and Doyle Jr, 2015; Bhuller et al., 2020). In the context of Swiss labor markets, this methodological approach has been followed by Arni and Schiprowski (2019), who study the impacts of search requirement and effort provision on job search outcomes. The authors combine the quasi-random allocation of unemployed to caseworkers within PES offices and variation in search requirements set by caseworkers to build a stringency instrument for the required and actual monthly number of job applications. Their findings show that both higher search requirements and higher provision of search effort lead to a substantial reduction in completed unemployment and non-employment duration.

The rest of the paper is organized as follows. Section 2 introduces the institutional context of our study. In section 3, we develop a theoretical framework to formalize caseworker referrals effects on unemployed job search process. In section 4, we present the data used in our empirical investigation and discuss how those can be related to our conceptual framework. We also provide descriptive evidence of the effects of caseworker referrals on job search activity and success. The two following sections represent the core of our empirical analysis. In section 5, we proceed to an application-level analysis to study how the referral policy affects the success probability of applications and the type of positions job seekers apply to. In section 6, we consider data aggregated to the individual-monthly level to examine whether the policy impacts application effort and labor matching. Section 7 concludes.

2. Institutional context

Swiss workers are entitled to unemployment benefits if they contribute at least twelve months within two years prior to the beginning of their unemployment spells. The typical potential benefit duration amounts to 12 or 18 months and is a function of the contribution period, age and family situation of unemployed. The replacement ratio ranges from 70% to 80%, depending on unemployed characteristics. Job seekers who claim unemployment benefits have to register at a regional PES office. Offices are organized at the cantonal level and exert some power over the implementation of unemployment policies. Once registered at a regional PES center, unemployed are assigned to a caseworker, either based on caseworkers' caseload, on occupation sector or at random (Behncke et al., 2010). Caseworkers carry out multiple tasks, among which monitoring and job search assistance.[1] Their incentives are relatively well aligned on those of job seekers: caseworkers' main objective consists in placing unemployed into jobs quickly, under the legal restriction that jobs must be suitable for them.

One important job search assistance tools at disposal of caseworkers is vacancy referrals. These consist in the transmission of information on suitable job vacancies, from caseworkers to job seekers, either in the context of meetings at PES offices or remotely. Information on vacancies originate from job postings directly announced to PES or from publicly posted adds. These two information sources characterize two types of vacancy referrals (standard and quick referrals), which follow slightly different processes.[2] In the Swiss context, the notification of a vacancy referral is accompanied by the obligation to apply to the referred position. Non-compliance to this obligation is potentially sanctioned by an unemployment benefits cut. This confers a coercive aspect on referrals, on top of their informational content. The exact implementation of the policy is decided at the cantonal or PES level, but caseworkers have some discretion over its use: some choose to refer vacancies frequently, while others make use of the policy more parsimoniously. As shown in Figure 2.1, referrals represent a widely used JSA tool in the Swiss labor market: approximately 30% of job seekers are referred at least one vacancy during their unemployment spells.[3] Referrals are also used

[1] On top of referring vacancies and monitoring unemployed search effort, caseworkers carry out several other tasks, from meeting and sanctioning job seekers, through assigning them to Active Labor Market Programs. As it will be emphasized in our empirical analyses, those additional tasks are only weakly correlated with the use of the referral policy.

[2] The only major difference between the two types of vacancy referrals lies in the involvement of the job posting firms during the referral process. In the former, recruiters are informed by caseworkers on whether suitable candidates were found, and can provide feedback on the candidates. In the latter, such preliminary information is not shared with the firms.

[3] These 30% are likely to be underestimated, as these figures are obtained from our data, which are partially made

Figure 2.1: Caseworker referrals (CwR) in the Swiss labor market

(A) Unemployment spells with CwR

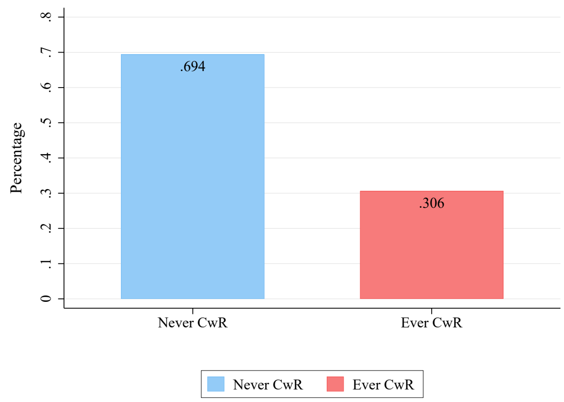

(B) CwR and elapsed unemployment duration

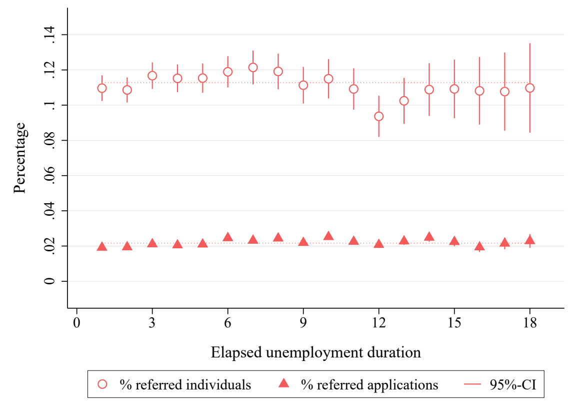

Note: This figure reports information on the use of caseworker referrals in the Swiss labor market. Panel A reports the share of individuals who ever get referred by their caseworkers. Panel B reports for each month of elapsed unemployment the share of individuals who get referred and the share of applications that result from a referral.

consistently over the course of unemployment: 10 to 12% of unemployed get treated in each month of elapsed unemployment, which corresponds to approximately 2% of all applications resulting from caseworker referrals.

In parallel to providing job search assistance services, caseworkers monitor unemployed job search activity. According to the legal framework, unemployment benefits recipients "*must be able to demonstrate [their] effort [to find a job]*".[4] To make this assessment by caseworkers possible, unemployed have to document their search activity in job search diaries.[5] These forms contain detailed information on all applications made by job seekers, in each month of their unemployment spells. They include information on applications dates, application modes (written, personal or by phone), the work-time percentage of targeted positions (full-time or part-time), as well as short vacancies descriptions. Most importantly, they report information on applications outcomes (callback, job offer, negative or still open) and on their treatment status with respect to the referral policy (referred or non-referred). Job search diaries are filled in and submitted to PES offices on a monthly basis, together with copies of job applications. These documents serve as a basis for monitoring by caseworkers, who make sure that minimum search requirements are met, both in quantitative and qualitative terms.[6]

of right-censored, left-censored and left-and-right censored spells.

[4] *Loi fédérale du 25 juin 1982 sur l'assurance-chômage obligatoire et l'indemnité en cas d'insolvabilité (LACI)*; RO 1982 2184. Retrieved 19th October 2022 from https://www.admin.ch/opc/fr/classified-compilation/19820159/index.html.

[5] A copy of the standardized form in French can be found in the Appendix, in Figure B1.

[6] Quantitatively, search requirements are measured in terms of applications per months. Qualitatively, copies of job applications are reviewed in order to assess their veracity.

In case of non-compliance with those requirements, job seekers are notified and potentially sanctioned by a benefit cut.[7] Caseworkers and job seekers update information on the success of job applications for up to two months after the job application has been sent out. This provides detailed and accurate information on success of job applications, *i.e.* whether they lead to a callback for a job interview, or a job offer.

3. Theoretical mechanisms

In order to apprehend the mechanisms through which caseworker referrals affect unemployed job search process, we rely on the theoretical framework developed in Fougère et al. (2009). We assume job search to take place in continuous time t, with individual discount rate ρ. The model assumes the existence of two different search channels: a private and a job search assistance channel. The private search channel, denoted P, is costly: it involves some utility cost $\zeta(e)$, which depends on the private application effort e. Search costs are assumed to be strictly increasing and convex in e, $\zeta'(e) > 0$ and $\zeta''(e) > 0$.[8] The job search assistance channel, denoted R, does not involve any cost. It can be thought as vacancy referrals passed on by caseworkers to job seekers.

The job search sequence takes place in three stages. First, unemployed workers receive information on and apply to vacancies posted by firms through channels $k = P, R$. This job application stage is formalized by means of a Poisson process of intensity α_k. Application rate through the private channel is endogenously determined and equal to search effort $\alpha_P = e$. For the referral channel, this rate is exogenous and equal to α_R. The total application rate is defined as $\alpha = \alpha_P + \alpha_R$. Second, applications are converted into job offers with channel-specific probabilities ψ_k. Differences in job applications success probabilities are considered as exogenous and might for instance be due to signaling (Van Belle et al., 2019) or differentiated characteristics of positions targeted by the two channels (Van Den Berg et al., 2019). This last point is formalized via channel-specific wage distributions $\Gamma_k(\omega)$, from which the wage ω offered by each vacancy is drawn. Third, unemployed accept a job offer if the offered wage is higher than their reservation wage. In that case, newly employed workers earn a constant wage ω and their working relationship terminates with probability σ.[9]

[7] The average size of a sanction amounts to 5.5 days of unemployment benefits, around CHF 900.- on average (Arni and Schiprowski, 2019).

[8] Increasing and convex search costs come from the ever-increasing difficulty in collecting additional information on vacancies via the private channel.

[9] When referring to the model of Fougère et al. (2009), our goal is to emphasize how private application effort and job finding chances change in the presence of an additional JSA (referral) application channel, accounting for differentiated vacancies and applications success probabilities across channels. For simplicity, we assume that the arrival rate of referred vacancies (and hence application rate through the referral channel) is exogenous and that full compliance prevails. The model hence makes abstraction of the effects of sanctions following non-compliance

Job seekers seek to maximize their expected life-time utility. While unemployed, they are eligible to benefits b and face two alternatives: (i) applying only through the referral channel and providing zero private search effort, a costless strategy, or (ii) applying through both channels. In the general setup, job seekers choose the alternative that yields the highest utility. In our context however, we assume that the two-channels strategy is always opted for by unemployed. This choice is motivated by the empirical observation that job seekers always apply through the private channel.[10]

Fougère et al. (2009) show that there exists a unique reservation wage for the private and referral channels when the two-channels strategy is adopted. The solutions for this reservation wage $\underline{\omega}^*$ and for the optimal level of search effort α_P^* are given by

$$\zeta'(\alpha_P^*) = \frac{\psi_P}{\sigma + \rho}\, \Lambda_P(\underline{\omega}^*) \tag{3.1}$$

$$\underline{\omega}^* = b - \zeta(\alpha_P^*) + \alpha_R \frac{\psi_R}{\rho + \sigma}\, \Lambda_R(\underline{\omega}^*) + \alpha_P^*\, \zeta'(\alpha_P^*), \tag{3.2}$$

where $\Lambda_k(w) = \int_w^\infty (\chi - w)\, d\Gamma_k(\chi)$ is the surplus function associated with $\Gamma_k(\omega)$.

We totally differentiate equation (3.2) to determine the effect of an exogenous increase in the use of referrals $d\alpha_R > 0$ on the optimum:

$$\left(1 - \alpha_R \frac{\psi_R}{\rho + \sigma}\, \Lambda'_R(\underline{\omega}^*)\right) d\underline{\omega}^* = \frac{\psi_R}{\rho + \sigma}\, \Lambda_R(\underline{\omega}^*)\, d\alpha_R + \zeta''(\alpha_P^*)\alpha_P^*\, d\alpha_P^*. \tag{3.3}$$

Equation (3.3) reveals that an exogenous increase in the application rate through referrals $d\alpha_R > 0$ is accompanied either by a reduction in the personal application effort $d\alpha_P^* < 0$ (due to search cost convexity, $\zeta''(e) > 0$) or by an increase in the reservation wage $d\omega^* > 0$ (due to the surplus function being monotonically weakly decreasing, $\Lambda'_R(\omega) \leq 0$).[11] This equation also suggests that a left-shift in the wage distribution of referred positions, *i.e.* the surplus $\Lambda_R(\underline{\omega})$ decreases $\forall \omega$, entails less adjustment in α_P^* (or $\underline{\omega}^*$): the less attractive the referred positions in terms of wage, the less pronounced the re-optimization by job seekers. The extent of the adjustments in α_P^* (or $\underline{\omega}^*$) following a change in α_R is thus revealing of the relative attractiveness of positions obtained via the referral channel compared to those accessed via the private one.

To analyze the impact of caseworker referrals on unemployment exit chances, we compute the derivative of the two-channels strategy unemployment exit rate $h(\alpha_P^*, \underline{\omega}^*; \alpha_R)$ with respect to

to the referral policy, *e.g.* on the reservation wage or on job finding.

[10] As shown in Figure B2 in the Appendix, most applications correspond to non-referred (or private) applications, hence the focus on the two-channels strategy.

[11] This result is consistent with the total derivation of equation (3.1), which shows that $d\alpha_P^*$ and $d\underline{\omega}^*$ are of opposite sign, as $\zeta''(e) > 0$ and $\Lambda'_P(\omega) < 0$.

the policy parameter α_R:

$$h(\alpha_P^*, \underline{\omega}^*; \alpha_R) = \alpha_R\ \psi_R\ \Big(1 - \Gamma_R(\underline{\omega}^*)\Big) + \alpha_P^*\ \psi_P\ \Big(1 - \Gamma_P(\underline{\omega}^*)\Big) \tag{3.4}$$

$$\begin{aligned}\frac{dh(\alpha_P^*, \underline{\omega}^*; \alpha_R)}{d\alpha_R} = &\ \psi_R\Big(1 - \Gamma_R(\underline{\omega}^*)\Big) \\ &+ \psi_P\Big(1 - \Gamma_P(\underline{\omega}^*)\Big)\ \frac{d\alpha_P^*}{d\alpha_R} \\ &- \Big(\alpha_R\ \psi_R\ \Gamma_R'(\underline{\omega}^*) + \alpha_P^*\ \psi_P\ \Gamma_P'(\underline{\omega}^*)\Big)\frac{d\underline{\omega}^*}{d\alpha_R}.\end{aligned} \tag{3.5}$$

The referral policy affects unemployment exit in three different manners. First, it has a direct positive effect following the increase in the applications sent through the referral channel. This effect depends positively on referred applications success probability ψ_R. Second, it entails a negative substitution effect on private application effort $\frac{d\alpha_P^*}{d\alpha_R} < 0$, which reduces exit chances through the private channel. Which of these two effects dominates depends on the application success probability differential $\psi_R - \psi_P$ and the size of the substitution effect. Third, it increases job seekers' reservation wage $\frac{d\underline{\omega}^*}{d\alpha_R} > 0$, which reduces the set of acceptable job offers and impacts negatively unemployment exit chance.
In the next sections, we use a reduced-form approach based on granular applications data to study caseworker referrals effects on job search, guided by the framework of Fougère et al. (2009).

4. Data and empirical measurements of job search

4.1 Data

Our empirical investigation of caseworker referrals effects relies on various Swiss administrative data. Our primary source of information consists of job search diaries collected by PES offices (Zuchuat et al., 2023). Paper-format documents were transcripted into numeric format at two different occasions. The main data collection took place in 2012-2013, in five different Swiss cantons (Zürich, Bern, Vaud, Zug and St-Gallen). This first dataset contains most application-level information that were mentioned hereinbefore, with the exception of vacancies textual descriptions. It is supplemented with a smaller job applications sample containing the missing textual piece of information. This auxiliary dataset, gathered in 2007-2008 in Zürich, is mostly used in the context of occupational mobility analyses. Our main analysis sample contains 699'652 applications sent by 15'616 individuals, which corresponds to 66'882 individual-monthly units. Our auxiliary sample covers 22'928 applications, sent by

692 individuals, or equivalently 2'858 individual-monthly units. [12]

We complement information from job search diaries with administrative data on job seekers' characteristics, employment status and labor market history. These information are available on an individual-monthly basis and merged with job search diaries data via individual social security identifiers. Part of these supplementary data originate from PES. They contain information on socio-demographic and unemployment-related characteristics (*e.g.* age, residential status, education level, gender, desired occupation), as well as unique identifiers for PES offices and caseworkers to which unemployed are affiliated. The remaining data concern job seekers' earnings and are retrieved from social security records, exclusively for the main applications sample. Among others, social security information enable us to precisely define unemployment spells.

Table 4.1 reports descriptive statistics on job seekers' observable characteristics for the main analysis sample.[13] The table distinguishes job seekers who ever get referred job vacancies from those who are never referred. On the one side, no systematic differences are observed between the two groups at the 10% level with respect to past labor earnings and work-hour percentage prior to unemployment. On the other side, ever-referred individuals tend to be older on average, are more likely to be women, to hold a secondary school degree and to have the Swiss nationality. In spite of being statistically significant, these differences remain quantitatively small. This corroborates the idea that referrals represent a broadly-used policy within the unemployed Swiss population. The table also shows that ever-referred individuals tend to be observed at later stages of their unemployment spells. This pattern is expected as right-censored spells are over-represented in the never-treated group: for those, caseworkers might not have yet made use of the referral policy at the time of observation. The table finally reveals substantial variability in the use of vacancy referrals according to local labor market institutions: some cantons are over-represented in the ever-treated group and conversely. This point substantiates the crucial role played by local institutions, and by extrapolation caseworkers, in the implementation of the referral policy.

In Table B3 in the Appendix, we proceed to the same comparative analysis at the application-level. We study whether referred applications differ from non-referred ones with respect to observed job seekers' characteristics. Again, some statistically significant differences are to be noted across the two sets of applications, notably in terms of age, labor market attachment and education level. However, those differences remain quantitatively small overall. This confirms that the referral policy is widely used across all categories of unemployed, and suggests

[12] Table B1 in the Appendix contains more details on the structure of our analysis samples.

[13] Similar descriptive statistics for the auxiliary sample are reported in Table B2 in the Appendix.

Table 4.1: Job seekers' characteristics in the main sample, Ever CwR vs. Never CwR

	Never CwR		Ever CwR		Comparison	
A. Numeric variables	Mean	SDV	Mean	SDV	t-stat.	p-value
Age	38.773	(11.782)	39.460	(11.481)	-3.386	0.001
Female	0.460	(0.498)	0.480	(0.500)	-2.345	0.019
Swiss	0.525	(0.498)	0.566	(0.493)	-4.745	0.000
Past labor income [in kCHF]	4.643	(5.108)	4.600	(5.039)	0.489	0.625
Unemployment history	0.087	(0.140)	0.091	(0.141)	-1.720	0.085
Previous work-time %	90.057	(18.231)	90.588	(17.448)	-1.450	0.147
Elasped unemp. duration	4.946	(3.789)	6.279	(3.922)	-20.041	0.000
B. Categorical variables	Freq.	%	Freq.	%	χ^2-stat.	p-value
Education levels						
Primary	2990	27.598	1104	23.087	52.039	0.000
Secondary	6336	58.483	3087	64.555		
Tertiary	1508	13.919	591	12.359		
Civil status						
Married/Separated	5274	48.680	2263	47.323	6.508	0.089
Single	4107	37.908	1807	37.788		
Divorced	1348	12.442	660	13.802		
Widow(er)	105	0.969	52	1.087		
Canton						
Bern	2316	21.377	934	19.532	51.342	0.000
St. Gallen	2633	24.303	1027	21.476		
Vaud	1919	17.713	808	16.897		
Zug	978	9.027	435	9.097		
Zürich	2988	27.580	1578	32.999		
Observations	10834		4782			

Note: This table reports descriptive statistics on job seekers' characteristics for the main sample of analysis. It distinguishes between the group of individuals who never get referred referrals (Never CwR) and those who receive at least one caseworker referral (Ever CwR). Panel A reports descriptive statistics on numeric variables, together with a test for the equality of the means across the two groups. Panel B reports absolute and relative frequencies for the two groups, together with a χ^2-independence test (between the group indicator and the categorical variable of interest).

that referred vacancies are found in the same sub-markets as their non-referred equivalents.[14]

4.2 Empirical measurements of job search activity and success

Job search diaries offer a comprehensive data source to study the effects of caseworker referrals on job search activity and success. They provide granular information on each application a out of the A_{it} sent by individual i at unemployment duration t.[15] In addition to knowing the treatment status of applications $CwR_{ait} = 0, 1$, we observe two sequential measures of

[14] The comparability of job seekers' characteristics across referred and non-referred applications suggests that the pool of referred applications is unlikely to be endogenously determined, as it does not concentrate on a specific type of worker or occupation sector. However, this does not prevent referred vacancies to exhibit different qualitative characteristics within specific sub-markets, *e.g.* lower quality jobs, characteristics which we observe partially and will analyze later on.

[15] Given the monthly record of job search diaries, months represent the "natural" time unit at which application-level data can be aggregated. In all our baseline analyses, t will thus be equal to a month of unemployment.

applications success. First, we know whether applications lead to a callback or a job interview $c_{ait} = 0, 1$, a common measure in the empirical job search literature (Kroft et al., 2013; Eriksson and Rooth, 2014; Farber et al., 2016). Second, conditional on a callback, we observe whether job seekers receive a job offer $o_{ait} = 0, 1$.[16] These two indicators enable us to study job applications success probability at different stages of the recruitment process. Aggregating these indicators to the individual-monthly level, we obtain two proxies for labor matching: the numbers of callbacks $C_{it} = \sum_a c_{ait}$ and job offers $O_{it} = \sum_a o_{ait}$ obtained at duration t. In addition, we observe two application-level characteristics which are revealing of the type and remuneration of jobs targeted by unemployed. First, we know whether targeted vacancies concern part-time positions, $PT_{ait} = 0, 1$. Second, for the auxiliary sample, we have information on targeted vacancies' occupational sectors, $\mathit{Targeted}_{ait}^{Occ}$. This piece of information can be combined with information on occupations desired by job seekers, $\mathit{Desired}_{it}^{Occ}$, to measure occupational search breadth (Belot et al., 2019).[17] The detailed structure of our tri-dimensional data together with job search indices definitions are reported schematically in Figure 4.1.

These empirical indices can be used to assess the assumptions and predictions developed in the theoretical section. First, the model assumes that referred and private applications differ in terms of success probability, ψ_R and ψ_P. This assumption can be tested by means of the two application-level success indicators c_{ait} and o_{ait}, measuring applications success probability.[18] Second, theory assumes that referred applications target different types of jobs. In the model, this is formalized through differentiated wage distributions, $\Gamma_R(\omega)$ and $\Gamma_P(\omega)$. As we do not observe vacancy-specific wages, we assess empirically the effects of referrals on job search targeting by studying two observable vacancies characteristics: work-time percentage and occupational sector. Third, the model predicts that referrals have a detrimental effect on exit chances via a negative substitution effect on private search effort, $\frac{d\alpha_P^*}{d\alpha_R} < 0$. Consequently, the effect on the total application rate $\alpha^* = \alpha_R + \alpha_P^*$ is ambiguous. Our data provide direct empirical counterparts to both α^* and α_P^*: the total number of applications A_{it} and the number of non-referred applications $A_{it}^P = A_{it} - \sum_a CwR_{ait}$ sent out in unemployment month t. Those represent typical measures of search effort (Faberman and Kudlyak, 2019; Arni and

[16] In our analysis, these two indicators are assumed to be sequential, *i.e.* a job offer can only occur if a callback took place previously. In the raw data, this condition is not always met: in some cases, a job offer is recorded without a preceding callback. This discrepancy between our analytical framework and the coding in the raw data it mostly due to caseworkers' practice. To ensure consistency in our data, we impute a callback any time a job offer is reported.

[17] Information on the work-time percentage is available for both analysis samples, while information on occupational sector is only available for the auxiliary sample.

[18] The job offer indicator o_{ait} corresponds to a direct counterpart to the theoretical ψ_k, while c_{ait} measures the success probability of applications in earlier stages of the recruitment process.

Figure 4.1: Job search diaries and empirical job search indices

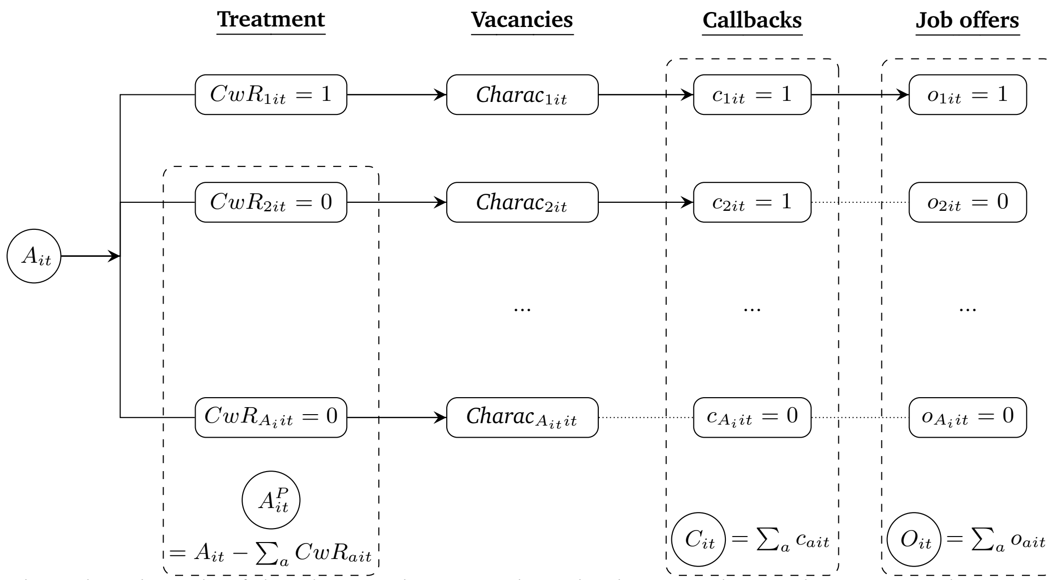

Note: This diagram depicts the recording of job search activity and success in our data. Each application a out the A_{it} applications sent out at duration t can be treated ($CwR_{ait} = 0, 1$). Applications outcomes are measured by means of two sequential success indicators: callbacks ($c_{ait} = 0, 1$) and job offers ($o_{ait} = 0, 1$). Matching between firms and job seekers is proxied through the monthly number of callbacks ($C_{it} \in \mathbb{N}$) and job offers ($O_{it} \in \mathbb{N}$). Additional vacancies characteristics $Charac_{ait}$ are also observed: work-time percentage PT_{ait} (part-time indicator) and the occupational sector of the targeted position $Occupation^T_{ait}$ (coded according to the Swiss Standard Classification of Occupations 2000 (SSCO 2000)).

Schiprowski, 2019; Fluchtmann et al., 2021). Fourth, theory predicts an ambiguous global effect of referrals on unemployment exit. To proxy job finding based on the information contained in the job search diaries data, we consider the two labor matching measures C_{it} and O_{it}. Those two proxies do not correspond one-to-one to unemployment exits, as they do not go beyond the job offer stage, *i.e.* they do not capture workers' acceptance of job offers.[19] Nevertheless, they are strongly predictive of increase in labor income as shown in Figure B3 in the Appendix, and can thus be thought of valid proxies for job finding.

4.3 Caseworker referrals effects on job search: descriptive evidence

We present graphical evidence of caseworker referrals effects on application-level outcomes in Figure 4.2. In all graphs, we distinguish between three mutually exclusive groups of applications: (i) applications sent by never-treated individuals, (ii) non-referred applications sent by ever-treated individuals and (iii) referred applications.

Panels A and B investigate the existence of success probability differentials between referred

[19] Since callbacks and job offers measure labor matching before job seekers' acceptance of the jobs, our reduced-form approach does not encompass the theoretical effect of the referral policy on job seekers' reservation wage, $\frac{d\bar{\omega}^*}{d\alpha_R} > 0$, and by repercussion on job finding. Nor does it capture the likely effect of sanctions following non-compliance to the referral policy. These effects, which we cannot rigorously analyze based on a reduce-form approach due to data limitation, should be negligible in any case since referred applications only represent a small fraction of the total pool of applications sent out by Swiss job seekers (see Figure 2.1B).

and non-referred applications, measured respectively by the callback (c_{ait}) and job offer (o_{ait}) probabilities. Both panels show that in the absence of referrals, never-treated individuals have a higher probability of receiving a positive response from firms, compared to ever-treated individuals. The two graphs also suggest the existence of a positive treatment effect on applications success: for ever-treated individuals, referred applications have a much higher probability of receiving a positive callback or job offer.

Panels C and D of Figure 4.2 show the impact of the policy on job vacancies characteristics. Panel C describes referrals effects on work-time percentage, measured by the part-

Figure 4.2: CwR effects on applications success probability and vacancies characteristics

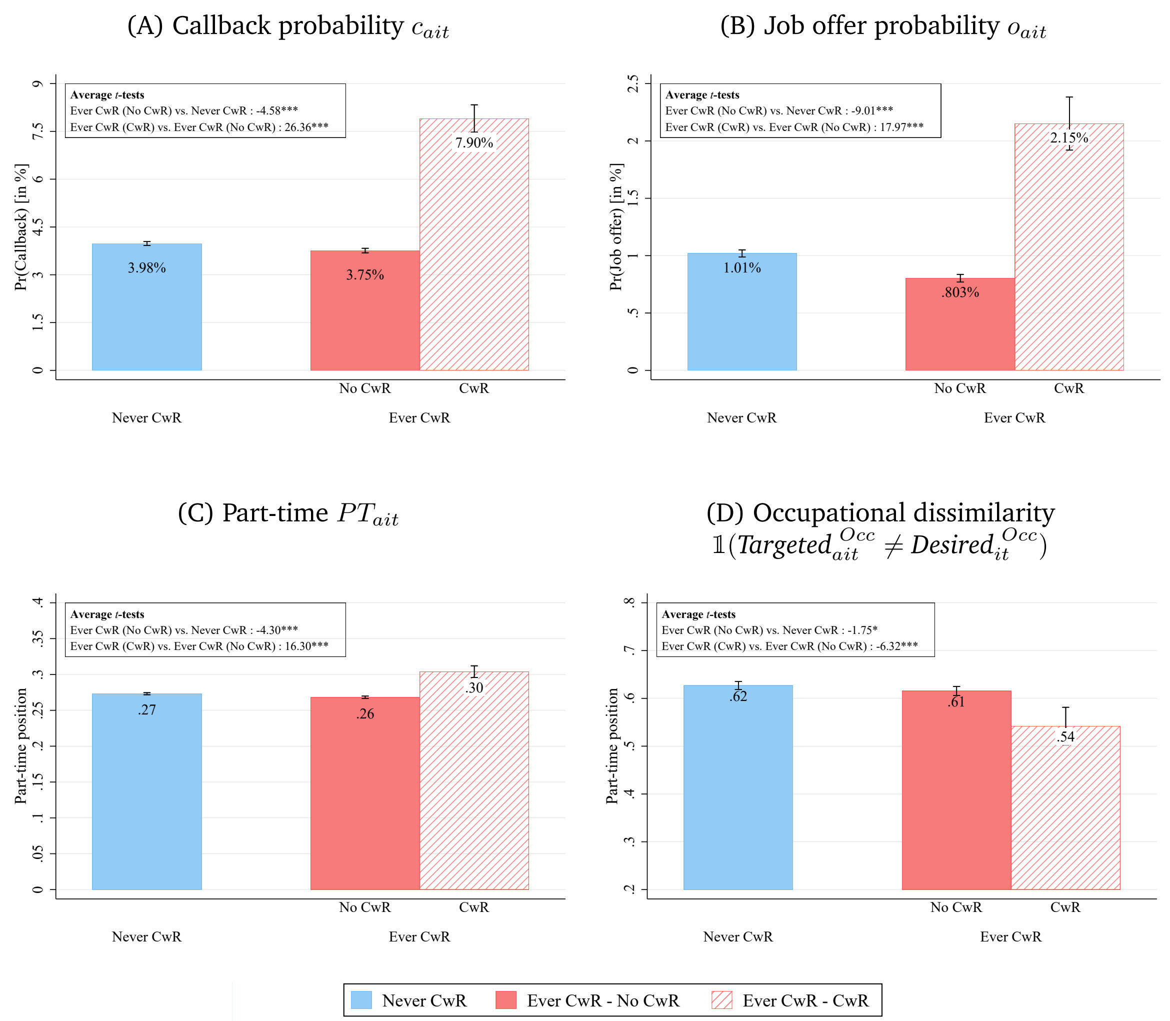

Note: This figure reports descriptive evidence on the effects of caseworker referrals on applications success probability and vacancies characteristics. Panels A and B reports evidence of referrals effects on the callback and job offer probabilities. Panels C and D reports similar evidence for vacancies characteristics, either measured by the part-time indicator (for work-hour percentage) or an occupational dissimilarity measure between targeted and desired occupations. All panels are based on the main sample of analysis, but panel D, which is obtained based on the auxiliary sample. All graphs distinguish between three groups of applications: (i) applications sent by individuals who are never treated (Never CwR), (ii) applications sent by individuals who are treated at least once but which do not result from a referral (Ever CwR - No CwR) and (iii) treated applications resulting from a referral (Ever CwR - CwR). 95%-confidence intervals for the mean values of each group are reported. Pairwise t-tests for the equality of the averages across groups are reported in boxes. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

time indicator PT_{ait}. Panel D shows evidence of their effect on occupational search breadth, measured by the extensive dissimilarity indicator between targeted and desired occupations $\mathbb{1}(\textit{Targeted}_{ait}^{Occ} \neq \textit{Desired}_{it}^{Occ}) = 0, 1$. Both panels show that referred and non-referred applications tend to target different types of vacancies: ever-treated individuals are more likely to apply to part-time jobs and to positions in their desired occupational sectors after receiving a referral from their caseworkers. The graphs also suggest the existence of slight systematic differences between characteristics of positions targeted by ever-treated and never-treated job seekers, in the absence of referrals.

In Figure 4.3, we examine the effects of caseworker referrals on job search effort and labor matching, based on data aggregated to the monthly-individual level. We distinguish between three mutually-exclusive groups of observations: (i) individual-monthly observations of never-treated job seekers, (ii) observations of ever-treated job seekers when no referral is used and (iii) observations when at least one referral is used.

Panel A reports the average number of applications sent per month (in total A_{it} and through the private channel A_{it}^{P}) for the three different groups of observations. Again, there exist systematic differences between ever-treated and never-treated individuals: in the absence of intervention by caseworkers, never-treated job seekers send less applications per month than ever-treated.[20] The graph also reports evidence of a positive treatment effect on total application effort: ever-treated individuals tend to send more applications in months when they receive referrals from their caseworkers. As suggested by theory, this effect is damped by a reduction in the private search effort.

Similar patterns are observed for matching indices C_{it} and O_{it}, in panels B and C of Figure 4.3. In the absence of caseworkers intervention, never-treated individuals face slightly more matching opportunities than ever-treated. Moreover, the graphs show evidence of a strong positive treatment effect of caseworker referrals on labor matching chances: the average numbers of interviews and job offers increase markedly in treatment months. In light of previous graphical evidence, these effects are likely due to the combined increases in applications success probability and in total job application effort.

Descriptive results are globally consistent with the theoretical mechanisms we emphasized previously. In the next two sections, we develop rigorous empirical strategies to causally estimate caseworker referrals effects on job search activity and success. In section 5, we study whether the referral policy affects job applications success probability and the characteristics of vacancies targeted by job seekers. We exploit our data at the application-level and rely on a within-identification strategy, with high-dimensional individual-by-duration fixed effects. In

[20]For those two subgroups, which focus on individual-monthly observations when the referral policy is not used, the personal and total application rates are equal by construction: $A_{it}^{P} = A_{it}$.

Figure 4.3: CwR effects on application effort and matching

(A) # applications A_{it} and A_{it}^P

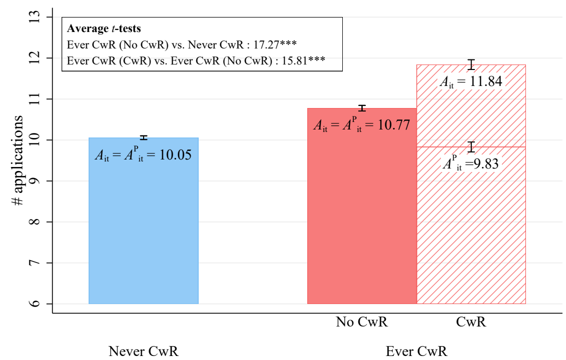

(B) # interviews C_{it}

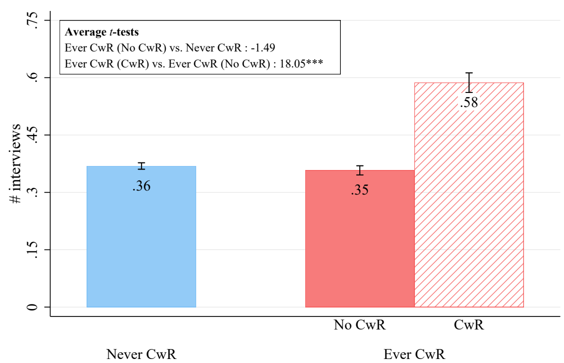

(C) # job offers O_{it}

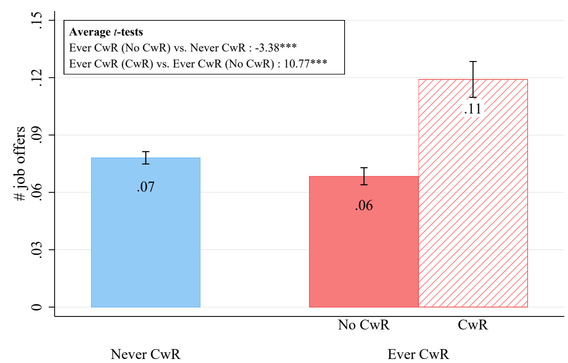

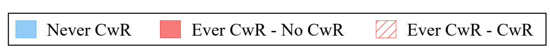

Note: This figure reports descriptive evidence on the effects of caseworker referrals on search effort and matching, based on data aggregated at the individual-monthly level. Panel A reports evidence for the monthly number of applications sent (total A_{it} and personal A_{it}^P). Panels B and C focus on matching proxies, respectively the monthly number of interviews C_{it} and of job offers O_{it}. The three graphs distinguish between three groups of individual-monthly observations: (i) observations concerning individuals who are never referred (Never CwR), (ii) observations related to individuals who are treated at least once but that concern months when the referral policy is not used (Ever CwR -No CwR) and (iii) months when the referral policy is used (Ever CwR - CwR). 95%-confidence intervals for the mean values of each group are reported. Pairwise t-tests for the equality of the averages across groups are reported in boxes. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

section 6, we consider our data aggregated to the monthly-individual level. Following an instrumental variable approach based on caseworker stringency, we evaluate the impacts of referrals on application effort and labor matching.

5. Application-level analysis: applications success and vacancies characteristics

This first empirical section examines the effects of caseworker referrals on application-level outcomes. We first discuss identification issues and present our within-identification strategy. We then apply this methodology to the callback and job offer indicators to assess whether referrals affect applications success probability. We also examine whether the policy modifies the type of jobs unemployed apply to (work-time percentage and occupational sector).

5.1 Identification strategy and econometric specification

The identification of caseworker referrals treatment effects faces two main empirical challenges. First, caseworkers decisions to refer vacancies are likely based on individual confounders, potentially unobserved. Second, even if the use of referrals is exogenously driven by the arrival rate of vacancies at PES, assignment to the policy is not necessarily fully exogenous: it might partially be based on job seekers' potential outcomes at the time of treatment. In an ideal setup, these endogeneity issues would be solved by comparing outcomes of similar applications sent by the same individual at a given point of her unemployment spell, differing only with respect to their treatment status.

Such empirical approach is usually inapplicable due to data limitations (*e.g.* see Fougère et al., 2009; Bollens and Cockx, 2017; Van Belle et al., 2019, for indirect evidence based on spell-level data). In contrast, we have access to granular job search data, where application-level treatment is directly linkable to job applications success outcomes and vacancies characteristics. This enables us to mimic the ideal empirical setup described above very closely.

Our baseline econometric specification to estimate caseworker referrals effects on application-level outcomes y_{ait} writes as

$$y_{ait} = \mu(i,t) + \beta \cdot \mathit{CwR}_{ait} + X^1_{ait}\gamma_1 + X^2_{it}\gamma_2 + \varepsilon_{ait}, \tag{5.1}$$

where a stands for applications, i for individuals and t for elapsed unemployment duration. The variable of interest is the treatment indicator $\mathit{CwR}_{ait} = 0, 1$. ε_{ait} corresponds to an idiosyncratic term. This specification controls for two types of observables. First, it conditions estimates on application-specific characteristics X^1_{ait}, which encompass application mode (written, phone, in person) and applications rank (1$^{\text{st}}$, 2$^{\text{d}}$, ... application sent by individual i at duration t). Second, it controls for a large set of time-varying individual characteristics X^2_{it}, measured on a monthly basis. Those range from socio-demographic characteristics to occupational features, through institutional background.[21] In addition, our specification includes a large set of individual and duration fixed effects, $\mu(i,t)$. These parameters are specified either separately as two different sets of fixed effects, μ_i and μ_t, or jointly as individual-by-duration fixed effects, μ_{it}.[22]

The main strength of our identification approach lies in the within-estimation of treatment

---

[21] All observables used as controls are reported in Table C1 in the Appendix.

[22] Note that in the later case, all time-constant and time-varying individual controls X^2_{it} disappear from the equation. As they are measured on the same time basis as the baseline individual-duration fixed effects, *i.e.* on a monthly basis, their effects are captured by μ_{it}.

effects. In the most conservative specification, where $\mu(i,t) = \mu_{it}$, treatment parameters are estimated within individual-duration units. This addresses the first identification challenge straightforwardly: high-dimensional fixed effects μ_{it} control for a large spectrum of observed and unobserved individual characteristics. Those parameters account for individual characteristics that are fixed in time, such as work experience, unemployment history or labor productivity. Further, they control for any time-varying confounders, as long as those remain constant within time interval t, *e.g.* gradual changes in job search strategy, in the quality of applications or time-varying socio-demographic characteristics.

Our identification strategy also addresses the second empirical challenge: on top of being exogenously determined by the arrival rate of vacancies, the treatment indicator is not systematically correlated with job seekers' potential outcomes, as we observe both treated and non-treated applications within individual-duration units.[23] Put differently, given that we measure treatment at the application-level, the probability of a positive callback or job offer faced by individual i in period t is orthogonal to caseworkers' decisions of referring vacancies. Our within-estimation procedure hence delivers estimates of average treatment effects for the subpopulation of ever-treated individuals, which does not extensively differ from the overall population of study according to our sample description.

5.2 Applications success

5.2.1 Main results

We examine the effect of caseworker referrals on applications success in Table 5.1. This table reports step-by-step estimates of equation (5.1) for callbacks c_{ait} (panel A) and job offers o_{ait} (panel B), based on a linear probability model. The baseline duration dimension t corresponds to elapsed unemployment months. We weight observations by the inverse of the number of applications made by individual i in t, to put equivalent weights on all individual-monthly units. Regression coefficients are measured in percentage points; they are reported in relative terms in squared brackets.[24] Standard errors are clustered by individual, adjusted for the presence of singletons (Correia, 2019), and reported in parentheses.

Column (1) in panel A shows that referred applications are more likely to lead to job interviews compared to non-referred applications, by 4.0 pp (+103%) on average. Adding individual and application-level controls in columns (2) and (3) does not affect the estimated treatment coefficient significantly. In column (4), we add individual μ_i and unemployment months μ_t fixed effects. This leads to a small yet insignificant reduction in the estimated

[23] Figure C1 shows that in months where the policy is used, we observe both referred and non-referred applications within individual-monthly units.

[24] The relative coefficients are obtained by dividing the percentage points coefficients by the outcome average of non-referred applications

Table 5.1: CwR effects on applications success probability

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Callback probability c_{ait}						
Caseworker referral	4.007***	4.103***	3.826***	4.284***	4.042***	4.005***
	(0.339)	(0.325)	(0.326)	(0.318)	(0.330)	(0.329)
	[102.846%]	[105.293%]	[98.187%]	[109.938%]	[103.739%]	[102.774%]
Individual controls	No	Yes	Yes	Yes	No	No
Application controls	No	No	Yes	Yes	Yes	Yes
Individual FE	No	No	No	Yes	No	No
Unemployment duration FE	No	No	No	Yes	No	No
Individual × Unemp. dur. FE	No	No	No	No	Yes	Yes
# identifying observations	699652	699652	699652	699631	699385	699385
# singleton observations	0	0	0	21	267	267
# total observations	699652	699652	699652	699652	699652	699652
# parameters	2	437	454	16021	66633	66634
# individuals	15616	15616	15616	15616	15616	15616
Adj.-R^2	0.001	0.016	0.033	0.111	0.170	0.170
Panel B: Job offer probability o_{ait}						
Caseworker referral	1.214***	1.204***	1.140***	1.297***	1.336***	1.322***
	(0.167)	(0.165)	(0.165)	(0.166)	(0.186)	(0.186)
	[129.581%]	[128.458%]	[121.686%]	[138.411%]	[142.53%]	[141.034%]
Individual controls	No	Yes	Yes	Yes	No	No
Application controls	No	No	Yes	Yes	Yes	Yes
Individual FE	No	No	No	Yes	No	No
Unemployment duration FE	No	No	No	Yes	No	No
Individual × Unemp. dur. FE	No	No	No	No	Yes	Yes
# identifying observations	699652	699652	699652	699631	699385	699385
# singleton observations	0	0	0	21	267	267
# total observations	699652	699652	699652	699652	699652	699652
# parameters	2	437	454	16021	66633	66634
# individuals	15616	15616	15616	15616	15616	15616
Adj.-R^2	0.000	0.005	0.011	0.084	0.118	0.119

Note: This table reports baseline estimation results of equation (5.1), where the dependent variable corresponds to applications success indicators. The duration dimension t corresponds to month of elapsed unemployment. Panel A reports results for the callback indicator c_{ait}, while panel B reports equivalent results for the job offer indicator o_{ait}. Application-level observations are weighted according to the inverse of the number of applications sent within individual-monthly units. Regression coefficients and standard errors (in parentheses) are reported in percentage points. Coefficients in relative terms are reported in squared brackets. They are standardized with respect to the average value of the dependent variable computed on the set of non-referred applications. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

treatment coefficient. Column (5) includes individual-by-duration fixed effects μ_{it}. This demanding specification still exhibits a positive and strongly significant treatment effect, which does not systemically differ from the unconditional estimate. Finally, in column (6), we additionally control for the work-hour rate of the targeted position, an application characteristic that is presumably affected by the referral policy. In spite of getting slightly smaller, the estimated treatment coefficient remains quantitatively large and highly significant. Those results confirm that referred vacancies face a higher success probability, even though this additional vacancy characteristic is controlled for. They should nevertheless be taken with caution, due to work-hour rate being endogenously determined by the policy.

Panel B reports a similar positive treatment effect for the job offer indicator. Depending on the specification, applications resulting from referrals have between 1.1 and 1.3 pp more chance (+121% to + 143%) to lead to job offers. As the two success indicators are sequential, a large share of this effect is due to the higher callback probability. Nonetheless, as parameters in panel B are estimated to be larger in relative terms compared to those in panel A, it seems that the referral policy does not only increase the chances of being interviewed, but also the probability with which interviews are converted into job offers.[25] These first results thus provide strong evidence of a positive success probability differential in favor of referred applications, compared to non-referred ones.

We examine the existence of heterogeneous treatment effects by estimating equation (5.1), including an additional interaction term between treatment and job seekers' observed characteristics. As applications success probability might vary extensively across different subgroups, we focus both on treatment coefficients in absolute and relative terms. Estimation results for the callback and job offer probabilities are reported graphically in Figures C3 and C4 in the Appendix. We find evidence of substantial heterogeneous treatment effects for both success outcomes. In particular, education seems to be a major mediating characteristic: caseworker referrals are found to benefit mostly unemployed who hold a primary level of education. For those, relative treatment effects are estimated to be much larger than for job seekers who hold a secondary or tertiary education degree (*e.g.* + 300%, +100% and + 75% respectively for the probability of a job offer). Estimated treatment effects also vary substantially across occupational sectors and seem to be mostly effective for people who work in the Hotelling and Industrial sectors. Our results also point towards heterogeneous effects with respect to nationality, even thought those are not found to significantly differ from each other.

In Figures C5 and C6, we assess whether treatment effects vary with respect to the local institutions implementing the policy. We estimate equation (5.1) on sub-populations defined either by the cantons or caseworkers to which job seekers are affiliated. All canton-specific treatment coefficients are found to be positive and statistically different from zero. These results also reveal the existence of substantial heterogeneity with respect to local institutions, a point that is confirmed by the caseworker-level analysis. However, the effectiveness of the policy does not seem to depend upon the intensity with respect to which it is used, as suggested by Figure C7 for caseworker-specific treatment effects.

Finally, we assess whether treatment effects vary with respect to elapsed unemployment duration, as caseworker referrals might be more or less successful depending on whether they are used early or late in the unemployment spells. We re-estimate Equation (5.1) interacting the treatment indicator with elapsed unemployment duration. Corresponding results are re-

[25] Graphical evidence of the higher conversion probability of interviews into job offers following caseworker referrals are reported in Figure C2 in the Appendix.

ported in Figure C8 and show that the policy effects on the callback and job offer probability remain relatively constant as elapsed unemployment duration increases.

5.2.2 Robustness

In the following, we address potential concerns regarding the identification of caseworker referrals effects on job applications success.

First, we tackle the question of information reporting. According to the data collection protocol, caseworkers were asked to follow up applications until two months after the submission of the forms at PES offices. This supposedly ensures that all interviews and job offers get recorded in job search diaries. From a data coding perspective, applications are usually characterized by a positive (callback or job interview) or negative response (negative). In some cases though, explicit information on applications outcomes is censored.[26] Given the data collection procedure, we interpret censored applications as implicit negative replies by firms. However, we cannot rule out that those correspond to missing interviews or job offers. If this censoring rate were to differ between referred and non-referred applications, previous treatment effects estimates would potentially be biased.[27] Figure C9 presents graphical evidence of discrepancies in information recording between referred and non-referred applications. The graph reports the censoring rates for the two types of applications, for c_{ait} and o_{ait}. It shows that callbacks are slightly less likely to be censored in the pool of referred applications (-2.1 pp), while no discrepancy is observed for job offers (-0.3 pp). In order to account for heterogeneity in information reporting, we estimate equation (5.1) using the censoring indicators as dependent variables. Corresponding results are reported in Table C2 in the Appendix. In specifications controlling for institutional and individual heterogeneity, callbacks are found to have the same censoring rate in the pool of referred and non-referred applications. No bias due to coding differentials should thus occur for callbacks. In contrast, job offers are more likely to be censored when applications result from referrals. As potentially missing job offers are more common among referred applications, any bias in the estimated treatment effect on job offers would be negative. The issue of information censoring hence remains relatively minor in our context.

Second, we address the assumption of comparability of applications sent by individual i in unemployment month t. Our choice of a monthly time window for t results from the monthly recording of search diaries. This assumption is questionable if search behavior evolves at

[26] The exact coding procedure used to define which application is considered as censored and which is considered as non-censored is reported in Table B4, in the Appendix.

[27] For instance, caseworkers might have a higher incentive to follow up job applications to vacancies that they refer. If this is the case, referred applications will be less likely to be censored, which will mechanically lead to a higher success probability for treated applications.

a higher frequency. As a robustness check, we estimate our empirical model with a higher frequency for the individual-by-duration fixed effects, either bi-weekly or weekly.[28] Corresponding results are reported in Table C3 in the Appendix. Point estimates and statistical significance of treatment coefficients turn out to be stable across all specifications, whether we consider the callback or job offer success indicator. This corroborates our baseline findings of strong positive effects of caseworker referrals on applications success, even when finer gradual changes in the job search process are allowed.

Third, we discuss potential changes in job seekers' behavior following the first assignment to the policy. As caseworker referrals are more likely to be unexpected in early stages of unemployment, they might affect applications outcomes differently at the beginning and at the end of unemployment spells. Our baseline estimates of referrals treatment effects might thus reflect some form of learning effects. To address this point, we estimate our empirical model based on two different sub-samples. First, for each individual we restrict the sample to applications sent before or in the same month as the first referral recorded in the data. Second, we further restrict the sample to observations from non-left-censored spells. This ensures that treatment effects are estimated on individuals who have never been confronted to the policy. Estimated treatment coefficients for these two subsamples are reported graphically in Figure C10 in the Appendix, together with our baseline estimates. All estimated coefficients do not systematically differ from each other, which suggests that learning effects are not driving our baseline results.

Finally, we elaborate on our choice of a linear probability model. This model offers the major advantage of allowing the inclusion of high-dimensional individual-by-duration fixed effects μ_{it}, a parametrisation that is at the core of our identification strategy. It also comes with the simplifying assumption of a linear relationship between the probability of a positive outcome and the regressors. In our context, this linear approximation might be questionable as callback and job offer probabilities are relatively low, *i.e.* identification comes mostly from regions where the cumulative distributions of outcomes are likely to be non-linear. We address these concerns by estimating logit and probit models. Despite their appropriate structure for this binary outcome analysis, those do not allow for the inclusion of high-dimensional fixed effects, due to the incidental parameter problem (Lancaster, 2000). For that reason, we specify models including only non-interacted parameters, μ_i and μ_t.[29] Estimates from non-linear models are reported in Table C4 in the Appendix, together with the baseline linear probability results. They tend to corroborate our baseline findings of positive and significant treatment effects of referrals on applications success, even if they deliver slightly lower estimates in

[28] We do not consider higher frequency duration dimensions, due to data limitation.

[29] Heterogeneity parameters are specified either under the form of fixed effects in the linear and (conditional) logit specification, or as correlated random effects in the probit specification.

comparison to the linear model. The latter nonetheless represents our preferred specification, as it accounts for high-dimensional fixed effects.

5.3 Vacancies characteristics

5.3.1 Work-time percentage

Our application-level analysis has emphasized positive treatment effects of caseworker referrals on applications success. Our theoretical framework implicitly assumes that this is due to referred and non-referred applications targeting different types of vacancies. We assess this conjecture by first studying the effect of the policy on targeted positions working hours. This information is encoded by means of the binary part-time indicator PT_{ait} and is informative about jobs remuneration.[30] To formally evaluate the impact of caseworker referrals on this variable, we estimate equation (5.1) with PT_{ait} as dependent variable.

Estimation results from our most robust specification including individual-by-duration fixed effects μ_{it} are shown in Table 5.2. Column (1) reports results based on the pooled sample,

Table 5.2: CwR effects on targeted work-time percentage

		Subsamples		
Previous work-time percentage	All	Part-time	Full-time	Missing
	(1)	(2)	(3)	(4)
Dependent variable : Part-time PT_{ait}				
Caseworker referral	2.902***	0.805	3.203***	4.012***
	(0.449)	(1.022)	(0.638)	(0.785)
Application controls	Yes	Yes	Yes	Yes
Individual × Unemp. duration FE	Yes	Yes	Yes	Yes
# identifying observations	616414	127290	318458	170666
# singleton observations	862	244	386	232
# total observations	617276	127534	318844	170898
# parameters	62542	13600	31414	17538
# individuals	15216	15216	15216	15216
Adj.-R^2	0.732	0.745	0.604	0.738

Note: This table reports the estimation results of equation (5.1), where the dependent variable corresponds to the part-time indicator PT_{ait}. All estimates are based on a specification with individual-duration fixed effects μ_{it}. Application-level observations are weighted according to the inverse of the number of applications sent within individual-monthly units. Regression coefficients and standard errors (in parentheses) are reported in percentage points. Columns (1) refers to the whole sample of applications for which the occupation rate is unambiguous. Columns (2), (3) and (4) correspond to estimates obtained on subsamples based on the previous occupation rate held by the job seekers, which can either be Part-time, Full-time or Missing. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

[30] According to the Swiss *Federal Statistical Office* a part-time job typically corresponds to a job whose work-time percentage is below 90%, *i.e.* approximately 36-38 hours per week.

containing all applications for which information on working hours is unambiguous.[31, 32] Those are in line with our descriptive evidence: applications following caseworker referrals are more likely to target part-time positions, even when controlling for individual heterogeneity through μ_{it}. Quantitatively, referred applications are 3.0 pp more likely to target part-time positions.

We estimate the same model separately on three sub-samples defined according to unemployed previous work-time rate (Part-time, Full-time or Missing). This information is retrieved from PES registers and represents a proxy for job seekers' desired work-time percentage. Corresponding results are reported in columns (2) to (4) of Table 5.2. In column (2), the treatment coefficient for the subsample of individuals who previously held part-time positions is not found to be statistically significant. In contrast, columns (3) and (4) report positive and strongly significant coefficients for the two other subsamples. These results suggest that caseworker referrals mostly modify the scope of job applications with respect to work-time percentage among unemployed who search for full-time jobs.

As identification is made within individual-duration units, the above results can be interpreted as evidence that caseworker referrals target lower-paid, part-time positions. In particular, individuals who are willing to work full-time and who apply to full-time jobs might end up targetting part-time positions following the intervention of their caseworkers. This interpretation is coherent with descriptive evidence on labor income before and after unemployment: job seekers who get a job offer through a caseworker referral tend to get paid less relative to their previous wage than other job seekers (see Figure C11 in the Appendix).[33] This story is also in line with the findings of Van Den Berg et al. (2019), who show that vacancy referrals lead to less well-paid and less stable jobs.

5.3.2 Occupational sector

The second vacancy characteristic we examine relates to the occupational sectors of targeted positions. As emphasized by the literature, online job search assistance can impact occupa-

[31] From a data-coding perspective, job search diaries indicate whether positions are part-time or full-time by means of two separate tickboxes. In most cases, this piece of information is unambiguous, *i.e.* the part-time and full-time indicators are mutually exclusive. In a minority of cases, both binary indicators are equal to one or to zero. Figure C12 shows that this form of miscoding concerns a small fraction of observed applications, less than 20%. As it is not clear how to impute these miscoded information, ambiguous applications are discarded for this analysis.

[32] As a robustness check, we re-estimate the same model including μ_{it} fixed effects exclusively on individual-duration units for which there is no miscoded information. Corresponding results are reported in Table C5 and are qualitatively the same as our baseline results.

[33] Given that we do not have access to application-specific wage information, we cannot apply the same methodology for wage as for the work-time percentage piece of information. For that reason, we only report descriptive analyses on labor income before and after unemployment (the ratio of the latter over the former). Those exhibit consistent patterns with our interpretation of lower-paying jobs obtained through caseworker referrals.

Table 5.3: CwR effects on occupations

	Occupational dissimilarity		Skills requirements	
Dependent variable	Extensive	Intensive	Physical	Cognitive
	(1)	(2)	(3)	(4)
Caseworker referral	-0.046*	-0.210*	0.683*	-0.590*
	(0.024)	(0.114)	(0.353)	(0.337)
Application controls	Yes	Yes	Yes	Yes
Individual × Unemp. dur. FE	Yes	Yes	Yes	Yes
# identifying observations	22772	22772	22772	22772
# singleton observations	156	156	156	156
# total observations	22928	22928	22928	22928
# parameters	2863	2863	2863	2863
# individuals	22928	22928	22928	22928
Adj.-R^2	0.559	0.551	0.674	0.720

Note: This table reports the estimation results of equation (5.1), where the dependent variable corresponds to occupational dissimilarity or skills requirements measures. All estimates are based on a specification with individual-duration fixed effects μ_{it}. Application-level observations are weighted according to the inverse of the number of applications sent within individual-monthly units. Columns (1) and (2) refer to the two occupational dissimilarity measures, between targeted and desired occupations: $\mathbb{1}(\textit{Targeted}^{Occ}_{ait} \neq \textit{Desired}^{Occ}_{it})$ and $\delta(\textit{Targeted}^{Occ}_{ait}, \textit{Desired}^{Occ}_{it})$. Columns (3) and (4) correspond to the two skills requirements indices: $\textit{Skills}^{Cogn.}_{ait}$ and $\textit{Skills}^{Phys.}_{ait}$. Standard errors are clustered at the individual level and reported in parentheses. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

tional mobility by modifying unemployed search scope (Belot et al., 2019). This might also be the case of caseworker referrals.

Empirically, we quantify occupational mobility by means of two dissimilarity measures between the occupation targeted by each application and the occupation desired by the job seeker. First, the extensive measure we previously introduced $\mathbb{1}(\textit{Targeted}^{Occ}_{ait} \neq \textit{Desired}^{Occ}_{it})$. Second, an intensive measure of occupational distance $\delta(\textit{Targeted}^{Occ}_{ait}, \textit{Desired}^{Occ}_{it})$.[34] Both application-level variables are based on the 5-digits hierarchical SSCO 2000 job nomenclature, the national standard at the time of data collection. To further characterize occupations, we complement our data with information on skills requirements obtained from *O*Net* repository. This additional data source contains information on specific abilities that are required per occupation, distinguishing between cognitive and physical skills. Those two sets of skills requirements are summarized by means of two application-level indices $\textit{Skills}^{Cogn.}_{ait}$ and $\textit{Skills}^{Phys.}_{ait}$.[35]

We estimate equation (5.1) on the auxiliary sample with the four above measures as dependent variables, to quantify caseworker referrals effects on occupational mobility and skills

[34] This intensive measure is defined as the number of different digits between the "targeted" and "desired" SSCO 2000 codes, taking the hierarchical structure of the nomenclature into account. For instance, occupations 85101 and 85102 have an occupational distance of $\delta(85101, 85102) = 1$, while occupations 85101 and 84101 have an occupational distance of $\delta(85101, 84101) = 4$.

[35] Skills requirements indices are based on the levels associated with skills in the *O*Net* repository, which are listed in Figure C13 in the Appendix. For each category of skills (Cognitive and Physical), the index is simply defined as the maximum level across all skills.

requirements. Estimation results are reported in Table 5.3. Columns (1) and (2) report results for the two occupational dissimilarity measures, while columns (3) and (4) relate to skills requirements indices. All reported estimates are based on specifications including individual-by-duration fixed effects. The two first columns confirm our descriptive evidence that applications resulting from caseworker referrals target specific types of occupations: they are more likely to concern positions in the same occupational sectors as the ones desired by job seekers. Translated into skills requirements, referred applications are found to focus on occupations involving more physical and less cognitive skills. This point is confirmed by Figure C13 in the Appendix, which shows that occupations targeted by referrals require more Stamina, Extent flexibility, Static and dynamic strength, while they involve less Oral expression, Speech recognition and Speech clarity.

These results point towards a potential narrowing effect of referrals on occupational mobility: the policy tends to reduce unemployed search breadth. Also, referred applications tend to focus on positions requiring more low-valued (physical) and less high-valued skills (cognitive). These findings echoe our analysis on work-hour percentage: caseworkers intervention might push job seekers to pick "lower-hanging fruits" in the labor market, a story that is consistent with the coercive aspect of the policy. These additional results and their interpretation should nevertheless be considered with caution, due to the reduced-size of the auxiliary sample on which they are based.

6. Individual-monthly analysis: application effort and labor matching

In this second empirical section, we move away from our application-level analysis to investigate whether caseworker referrals affect individual-monthly outcomes. Due to data aggregation, we develop an alternative identification strategy that exploits variation in caseworker stringency as an instrument, to causally estimate the effects of the policy on application effort and labor matching.

6.1 Identification strategy and econometric specification

Measuring caseworker referrals impacts on individual-monthly variables requires an alternative identification strategy. Because of data aggregation, no straightforward counterfactual can be built for individual i at duration t as in section 5, *i.e.* μ_{it} cannot be included in the model. Moreover, a simple comparison between individuals' outcomes in treatment and non-treatment months would be misleading. Not only is the assignment to treatment likely to depend upon potential outcomes at a given duration, but the job search process is subject to substantial dynamics (Faberman and Kudlyak, 2019; Fluchtmann et al., 2021). This might limit the comparability between individual outcomes in t and $t' \neq t$.

To address endogeneity issues in the context of our individual-monthly analysis, we exploit an exogenous variation in the exposure to the referral policy entailed by the institutional framework. This one is due to the combination of (i) conditional random assignment of job seekers to caseworkers within PES centers and (ii) the discretion let to caseworkers over the implementation of the policy. These two features imply that job seekers are exogeneously referred more or less regularly, depending on the caseworkers they are assigned to.

Following Arni and Schiprowski (2019), whose analysis is based on a long-standing literature on judge and caseworker stringency (*e.g.* Kling, 2006; French and Song, 2014; Aizer and Doyle Jr, 2015; Bhuller et al., 2020), we build a time-varying instrument $\overline{CwR}_{-it}$ measuring the stringency of caseworkers with respect to the use of referrals. For each observation unit, our instrument is defined as the leave-one-out average of the number of caseworker referrals assigned by the caseworker supervising individual i in unemployment month t (excluding observation it).[36] We causally estimate the policy effects on job search effort and matching proxies using the following empirical model:

$$CwR_{it} = \psi_{1,i} + \delta_{1,t} + \beta_1 \, \overline{CwR}_{-it} + X'_{it}\gamma_1 + \phi_{1,rt} + \varepsilon_{1,it} \tag{6.1}$$

$$Y_{it} = \psi_{2,i} + \delta_{2,t} + \beta_2 \, CwR_{it} + X'_{it}\gamma_2 + \phi_{2,rt} + \varepsilon_{2,it}, \tag{6.2}$$

where i stands for job seekers, t for months of elapsed unemployment and r for PES offices. The treatment variable $CwR_{it} \in \mathbb{N}$ corresponds to the number of vacancies referred to individual i in month t. It is instrumented using the time-varying stringency instrument $\overline{CwR}_{-it}$, in equation (6.1). The structural effect of the policy on the dependent variable is formalized in equation (6.2). Time-varying outcomes Y_{it} correspond either to application effort measures (A_{it} and A^P_{it}) or to matching proxies (I_{it} and O_{it}). Given our access to longitudinal data and a time-varying instrument, our specification allows the inclusion of both individual and unemployment duration fixed effects (ψ_i and δ_t). Standard socio-demographic characteristics and labor market-related factors are controlled for in X_{it}. Crucially, the two equations contain PES centers-by-quarter fixed effects, captured by the ϕ_{rt} parameters.[37] The inclusion of this

[36]Formally, let $\mathcal{C}(i,t)$ be the set of all individual-monthly observations $i't'$ associated with the caseworker supervising individual i in month t, *i.e.* the set of relevant observations for the computation of the (leave-one-out) stringency index for observation it.
The leave-one-out instrument for individual-monthly observation it is then defined as

$$\overline{CwR}_{-it} = \frac{\left[\sum_{i'}\sum_{t'} \mathbb{I}\big(i't' \in \mathcal{C}(i,t)\big) \cdot CwR_{i't'}\right] - CwR_{it}}{\left[\sum_{i'}\sum_{t'} \mathbb{I}\big(i't' \in \mathcal{C}(i,t)\big)\right] - 1}.$$

[37]Note that we omit the calendar quarter index in the notation, as a matter of simplification. Given that we know for each individual i the starting month of her unemployment spell, the elapsed duration dimension t is enough to retrieve the calendar quarter in which the observation occurs. Also, the r index is omitted for other variables,

specific set of controls ensures that we exploit variation in caseworker stringency within PES centers and calendar quarters, where allocation of unemployed to caseworkers is presumably quasi-random (Behncke et al., 2010). This set of fixed effects moreover rules out potential confounding effects of center-specific policy implementation or particular local labor market conditions.

6.2 Validity of the instrument

In the following, we assess the validity of our instrumental variable approach. First, we discuss the relevance of our instrument by providing evidence on the first stage relationship. Second, we discuss additional conditions underlying our identification strategy. This assessment procedure is closely related to the analyses conducted in Arni and Schiprowski (2019) and Bhuller et al. (2020).

6.2.1 Relevance of the instrument

Our empirical strategy makes the implicit assumption that allocation to treatment is affected by caseworker stringency. Intuitively, unemployed will more likely be referred vacancies if they are supervised by caseworkers who are more prone to use this policy. Figure 6.1 offers a visual representation of the first-stage relationship between our caseworker stringency measure and the number of vacancies referred to job seekers. Both variables are reported net of the effects of PES centers-by-quarters fixed effects. The figure emphasizes a strong positive relationship between the instrument and the treatment variable: within PES offices, the more stringent a caseworker, the higher the number of referrals she assigns to job seekers on average. The graph also shows the existence of substantial residual variation in caseworker stringency within PES, which can be exploited empirically.[38]

Table 6.1 reports empirical estimates of the first-stage linear relationship described in equation (6.1). All estimates are conditional on PES-by-quarter fixed effects. The various specifications differ with respect to the inclusion of controls, or individual and duration fixed effects. Columns (1) and (2) do not include fixed effects. They correspond to the analytical counterparts of Figure 6.1. Both estimates ot the first-stage coefficient are found to be positive and strongly significant: a unit increase in caseworker stringency implies an increase of 0.6 in the number of vacancies referred to job seekers on average. Columns (3) and (4) report estimation results when individual and duration fixed effects are accounted for. In these specifications, first-stage parameters are estimated to be negative and highly significant. This seemingly counterintuitive result is explained by the low within-individual variation in the

as a matter of simplification.

[38]Conditional on PES $\times$ quarter fixed effects, caseworker stringency standard deviation amounts to 0.10, for an average value of 0.25.

Figure 6.1: Conditional first stage relationship

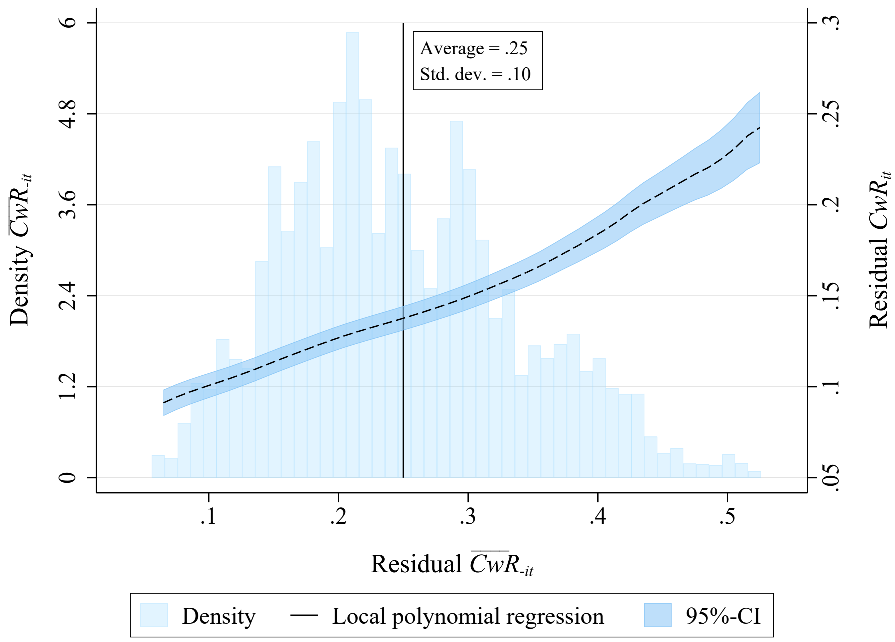

Note: This figure is inspired by Arni and Schiprowski (2019) and Bhuller et al. (2020). It represents graphically the first stage relationship between the monthly number of referrals and the caseworker stringency instrument, conditional on PES $\times$ quarter fixed effects. Residual values of CwR_{it} and $\overline{CwR}_{-it}$ are obtained by regressing those variables on PES-by-quarter fixed effects. The estimated relationship corresponds to a local polynomial regression, with 95%-confidence intervals. The graph additionally reports the distribution of the stringency index (the 0.5 bottom percentile and 2 top percentiles of the distribution are omitted).

instrumental variable: most variation in the stringency index is observed across individuals (see Figure D1 in the Appendix). This pattern is presumably due to the low frequency of job seekers experiencing a change in caseworker in our sample (approximately 10%), and by the fact that caseworkers are consistently stringent in their use of referrals across all job seekers they supervise. Consequently, the main effect of caseworker stringency is in great part captured by individual fixed effects $\psi_{1,i}$ in the first stage. Net of this main effect, higher stringency values are associated with lower number of caseworker referrals.

The above results suggest that our caseworker stringency measure is made of two different components, both affecting the actual use of referrals: part of it is determined by the intrinsic stringency of the caseworker, whereas the remaining part results from short-term fluctuations in the exogenous arrival rate of vacancies at PES offices. The exact source of identification used in the first-stage relationship differs depending on the specification we consider. In the specification without individual fixed effects, identification is made trough the overall stringency of caseworkers, and relies on the plausibly exogenous allocation of job seekers to caseworkers that are more or less prone to using the referral policy. In contrast, in the specification with individual fixed effects, the effect is identified mostly through the short-term fluctuations in supply of vacancies to be referred, as the intrinsic part of our caseworker

Table 6.1: Conditional first stage relationship, regressions

	(1)	(2)	(3)	(4)
Dependent variable: # of caseworker referrals CwR_{it}				
Caseworker stringency	0.614***	0.603***	-1.618***	-1.614***
	(0.072)	(0.072)	(0.180)	(0.180)
Individual controls	No	Yes	No	Yes
Unemployment duration FE	No	No	Yes	Yes
Individual FE	No	No	Yes	Yes
PES $\times$ Quarter FE	Yes	Yes	Yes	Yes
# identifying observations	66882	66882	64026	64026
# singleton observation	0	0	2856	2856
# total observations	66882	66882	66882	66882
F-stat (Kleibergen-Paap)	68.790	67.743	115.185	114.539
adj.-R^2	0.012	0.015	0.310	0.310

Note: This table reports estimation results for the first-stage relationship defined in equation (6.1). Columns (1) and (2) do not account for the longitudinal dimension of our data, while columns (3) and (4) include both individual and duration fixed effects. Errors are clustered at the individual level. Stars indicate the following significance levels: * 0.10, ** 0.05 and *** 0.01.

stringency measure is captured to a great extent by individual fixed effects. [39]As it will be shown later one, the two approaches yield very similar results when we estimate the effects of referrals on job search outcomes.

6.2.2 Additional assumptions

We proceed to several additional identification checks, to make sure that caseworker stringency represents a valid instrument for assignment to caseworker referrals.

First, we assess the conditional independence assumption (CIA) of the instrument with respect to job seekers' characteristics. This assumption is presumably met given the conditional random assignment of job seekers to caseworkers within PES (Behncke et al., 2010). We report descriptive evidence for the CIA in Table D1 in the Appendix, in columns (1) and (2). Those report partial correlations between the instrument and job seekers' socio-demographic characteristics. The partial correlations are obtained from a multivariate linear regression, where individual-monthly observations are aggregated to the individual level to put equal weights on job seekers. The two columns differ with respect to the inclusion of PES centers fixed effects. Column (1) shows that there exist substantial correlations between caseworker stringency and specific socio-demographic characteristics (age and education notably). Once we condition on PES centers fixed effects in column (2), all previously significant coefficients

[39]This interpretation is consistent with the negative sign of the first-stage relationship between the use of referrals and our leave-one-out instrument, when individual fixed effects are accounted for. Precisely, months in which the referral policy is used are likely to be concurrent to a positive idiosyncratic shock on the supply of vacancies to be referred. By construction, the leave-one-out approach introduces a negative correlation between vacancies supply shocks and the stringency measure we compute. This negative relationship only shows up in the specification with individual fixed effects, as the intrinsic component of caseworker stringency is captured by the fixed effects.

become insignificant. Also, the strength of the correlations tends to get attenuated, as shown in Figure D2 in the Appendix. These results provide supporting evidence that the CIA is met in our setting.

Second, we discuss the usual exclusion restriction (ER) in instrumental variables regressions. In our context as in Arni and Schiprowski (2019), this assumption implies that caseworker stringency is not correlated with other decisions typically made by caseworkers. To gauge the validity of this assumption, we proceed in the same way as for the CIA: we compute partial correlations between our instrument and other decisions by caseworkers (assignment to individual or collective training programs, sanctions, minimal search requirements), using multivariate linear regressions.[40] Estimation results are reported in Table D1 in the Appendix, in columns (3) and (4).[41] Without conditioning on PES centers, caseworker stringency is found to significantly correlate with assignment to individual training programs, while coefficients associated with other caseworkers' decisions are insignificant. Once PES fixed effects are accounted for, most caseworkers' decisions remain uncorrelated with the stringency measure, with the expectation of assignment to collective programs (significant at the 10% level). Even though these evidence do not fully support the ER assumption, this one is presumably met in our context. Not only is the size of the partial correlation between caseworker stringency and assignment to collective training programs small, but those programs are relatively scarce in comparison to caseworker referrals.[42] Consequently, our measure of caseworker stringency should mostly affect job search outcomes through the use of caseworker referrals.

Third, we evaluate the monotonicity assumption behind our stringency instrumental variable approach. Bhuller et al. (2020) argue that monotonicity can be violated in contexts where caseworker or judge stringency is used as an instrument, because public servants might behave differently towards different types of individuals. In our setting, similarly to Arni and Schiprowski (2019), the monotonicity assumption is verified as long as caseworkers assign vacancies consistently across different subpopulations of unemployed. As pointed out in Bhuller et al. (2020), there exists no formal test to assess this assumption. Supporting evidence can nevertheless be provided by estimating the first stage regression on different subsamples and

[40] Search requirements, defined in terms of number of applications to be sent per month, are set by caseworkers (Arni and Schiprowski, 2019). They are not observed in our data. Instead, we consider the minimum monthly number of applications made by individual i as a proxy for this decision.

[41] The two last columns (5) and (6) of Table D1 report partial correlations between on the one side our instrument, and on the other side job seekers' characteristics and other decisions taken by caseworkers . These regressions provide evidence for the CIA and ER simultaneously.

[42] Columns (7) and (8) of Table D1 measure partial correlations between on the one side our instrument and on the other side job seekers' characteristics, other decisions taken by caseworkers and the individual-level average of the number of caseworker referrals. These results show that our instrument is mostly correlated with the assignment to caseworker referrals. Regarding the prevalence of ALMPs in our sample data, approximately 30% of sampled individuals ever get referred, while only 15% ever get assigned to a collective training program.

showing that all estimated parameters are non-negative. We follow this approach and consider various subsamples, defined according to different sample-split variables. Corresponding first-stage estimates, without individual and duration fixed effects, are reported in Figure D3 in the Appendix. Estimated parameters all appear to be positive and relatively stable across subsamples. Moreover, none of the subpopulation coefficients differs significantly from the baseline first-stage estimate obtained on the pooled sample. In Figure D4, we consider finer subsamples based on the combination of one, two or three observable characteristics. Again, all first-stage parameters are estimated to be positive. Therefore, the monotonicity assumption seems to be verified in our context.

6.3 Results

6.3.1 Main results

We report baseline estimates of caseworker referrals effects on the four individual-monthly outcomes A_{it}, A_{it}^{P}, C_{it} and O_{it} in Table 6.2. Columns (1) and (2) report estimates based on OLS without and with fixed effects, while columns (3) and (4) report equivalent 2SLS results where referrals are instrumented with caseworker stringency.

Panel A reports the results for the total number of applications sent per month A_{it}. All columns emphasize a strong and positive effect of caseworker referrals on total application effort. Depending on the specification, this one ranges from 0.46 to 0.74 additional application per month, per additional referral. The 2SLS estimator delivers slightly larger point estimates than OLS, even though coefficients do not significantly differ from each other. These findings are consistent with estimates for the private application effort A_{it}^{P}, reported in Panel B: the positive impact of referrals on total application effort coincides with a decrease in private search effort. Quantitatively, estimated treatment coefficients are relatively low, ranging from -0.54 to -0.26. They significantly differ from -1, a value that would indicate full substitution of job applications.

In light of our theoretical framework, these results point towards imperfect substitutability between referred and non-referred applications. The fact that job seekers proceed to limited adjustment of their optimal private application effort following caseworkers' intervention suggests that they do not access the same types of positions via the referral and private channels. The above results are hence consistent with our application-level analysis, which emphasizes that referred applications target different and potentially less attractive positions than those obtained through the private channel.

Panels C and D report estimation results for our two labor matching proxies C_{it} and O_{it}. Treatment coefficients for both outcomes are estimated to be positive, across all specifications. For callbacks, all parameters are found to be statistically significant at the 1% level. For

Table 6.2: CwR effects on individual-monthly outcomes

	OLS	OLS FE	2SLS	2SLS FE
	(1)	(2)	(3)	(4)
A. Dependent variable: # applications A_{it}				
# Caseworker referrals	0.515***	0.461***	0.740*	0.723**
	(0.034)	(0.029)	(0.437)	(0.282)
	[5.007%]	[4.485%]	[7.194%]	[7.034%]
B. Dependent variable: # non-referred applications A_{it}^P				
# Caseworker referrals	-0.485***	-0.539***	-0.260	-0.277
	(0.034)	(0.029)	(0.437)	(0.282)
	[-4.715%]	[-5.237%]	[-2.528%]	[-2.688%]
C. Dependent variable: # interviews I_{it}				
# Caseworker referrals	0.061***	0.073***	0.210***	0.151***
	(0.006)	(0.007)	(0.073)	(0.054)
	[16.608%]	[20.072%]	[57.381%]	[41.266%]
D. Dependent variable: # job offers O_{it}				
# Caseworker referrals	0.013***	0.016***	0.033	0.027
	(0.002)	(0.003)	(0.025)	(0.020)
	[17.944%]	[21.607%]	[43.362%]	[35.898%]
Individual controls	Yes	Yes	Yes	Yes
Unemployment duration FE	No	Yes	No	Yes
Individual FE	No	Yes	No	Yes
PES $\times$ Quarter FE	Yes	Yes	Yes	Yes
# individuals	15616	15616	15616	15616
# identifying observations	66882	64026	66882	64026
# singleton observation	0	2856	0	2856
# total observations	66882	66882	66882	66882

Note: This table reports the empirical estimates of equation (6.2). Panel A considers the monthly total application effort A_{it} as dependent variable, panel B the monthly private application effort A_{it}^P, panel C the monthly number of job interviews I_{it} and panel D the monthly number of job offers O_{it}. Columns (1) and (2) report OLS estimates, with or without individual fixed effects. Columns (3) and (4) report similar 2SLS estimates, based on the caseworker stringency instrument. Coefficients in relative terms are reported in brackets. They are standardized with respect to the average of the outcome variable in non-treatment months. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

job offers, only coefficients estimated by OLS appear to be significant, while 2SLS delivers relatively imprecise estimates.[43] Quantitatively, our most robust specification shows that one additional referral leads to 0.151 more interviews and 0.027 more job offers per month on average.

The results for individual-monthly labor matching proxies are consistent with our application-level analysis. They also corroborate previous findings in the literature: vacancy referrals represent an effective policy tool to increase unemployed job finding chances (Fougère et al., 2009; Bollens and Cockx, 2017; Van Den Berg et al., 2019; Cheung et al., 2019). Our analysis goes one step further though, as it sheds light on the mechanisms through which this positive

[43] This lack of significance for the monthly number of job offers might arise from the relatively scarce number of job offers we observe in our data and their skewed distribution. Re-estimating the model using a Poisson specification, we find strong and positive treatment effects, also for job offers (see Table D2 in the Appendix).

effect operates: it arises both from the higher success probability of referred applications and from an increase in total application effort from job seekers.

6.3.2 Robustness

We report specification checks for our individual-monthly analysis in the Appendix.
In Table D2, we use the Poisson pseudo-maximum likelihood estimator (PPMLE) instead of the baseline 2SLS estimator. This alternative specification is motivated by the count data nature of the dependent variables, and by the properties of the PPMLE estimator, which turns out to be reliable even in situations where the outcome distribution is strongly asymmetric (Silva and Tenreyro, 2011). Table D2 reports the results of three different Poisson specifications: a standard Poisson model, a model including fixed effects and a model where we account for the endogeneity of the treatment variable, using the two-steps GMM estimator. Estimated coefficients can be interpreted as semi-elasticities and are to be compared with relative coefficients in the baseline table. All specifications deliver results that are very close to our linear estimates. Moreover, previously insignificant treatment coefficients for job offers now turn out to be strongly significant. Small discrepancies across the two sets of results arise because of the count data specification and the change in sample composition needed for the estimation of the Poisson fixed effects model.[44]
In Table D3, we present log-lin specifications for the application effort measures. This alternative modeling choice is motivated by the slight positive skewness of the dependent variables. These results tend to confirm the limited substitution effect entailed by the referral policy: job seekers reduce their private search effort to a limited extent, leading to an increase in their total application effort in months when they are referred vacancies by caseworkers.

7. Conclusion

Vacancy referrals by caseworkers represent one of the most common JSA tools used at PES. Recent studies have shown that this instrument is effective at fighting unemployment, as it leads to significant improvements in job finding chance. However, existing literature has not yet studied the mechanisms through which caseworker referrals affect job search. Moreover, little is known about their consequences on the types of jobs unemployed apply and eventually access to. Understanding these mechanisms is crucial for policy-makers to have a thorough picture of this policy's implications and potential limitations.
This paper provides the first comprehensive analysis of referrals effects on multiple aspects of job search. Based on the intuitions developed in a partial equilibrium job search model à la

[44]Only individuals who exhibit variability in the dependent variable A_{it}, A^P_{it}, C_{it} or O_{it} are kept in this specification.

Fougère et al. (2009), where referred and non-referred applications coexist and application effort is endogenously determined, we use granular longitudinal data on job applications sent by Swiss unemployed to quantify the impacts of the policy on job search activity and success. Using data at the application-level, we study how referrals affect applications success indicators and targeted vacancies characteristics. We address endogeneity concerns via a within-identification approach with high-dimensional individual-by-duration fixed effects. We find that job applications resulting from caseworker referrals have a much higher probability of leading to callbacks and job offers. Referred applications also appear to focus on part-time positions and on job seekers' desired occupations, which require less cognitive and more physical skills. This suggests that caseworker referrals do not only affect job search success, but also the type of positions unemployed apply to.

Aggregating data to the individual-monthly level, we examine whether referrals impact job search effort and, eventually, matching between firms and unemployed workers. Because of data aggregation, we develop an alternative identification strategy based on an instrumental variable capturing caseworker stringency. Exogenous variation in the exposure to the policy is induced by quasi-random allocation of unemployed to caseworkers and by the discretion let to caseworkers over the use of the policy. Our results show that caseworker referrals impact positively the total number of applications job seekers send out per month. This effect coincides with a slight decrease in the number of non-referred applications. This limited substitution effect is consistent with the two application channels targeting different types of jobs. Consequent to the positive effects on applications success probability and total application effort, we find evidence of a strong positive impact of referrals on both the monthly number of callbacks and job offers, two labor matching proxies.

Our results are overall consistent with existing related studies: they confirm that caseworker referrals are an effective tool to improve unemployed job finding chances. However, our analysis nuances previous findings on referrals' positive impact: this policy might have direct adverse effects for job seekers, on top of its general equilibrium consequences (see Cahuc and Le Barbanchon, 2010 or Crépon et al. (2013) for job search assistance programs in general, and Cheung et al., 2019 for the specific referral policy). In particular, the positive impact on job finding chances might be realized at the cost of accessing lower-quality jobs. Such detrimental consequences of referrals have already been pointed out in Van Den Berg et al. (2019) in a more indirect manner: jobs taken up shortly after receiving referrals offer significantly lower wages and are less stable than jobs found in the absence of caseworkers' mediating intervention. Our study substantiates this idea by providing direct and granular evidence of referrals effects on the type of positions job seekers apply to. In light of our analysis, part of the positive effect of referrals on job finding might be due to job seekers sending more applications to part-time and less cognitively-demanding positions, which are likely to offer lower remuneration. Findings of limited adjustment in the optimal private application effort in months when the policy is used also hint towards this point.

From a labor market policy design perspective, our study is informative of the dual nature of caseworker referrals: their sanctioning aspect might prevail over their purely informational component, as they force unemployed to apply to less attractive positions. Consequently, this policy might have longer-term detrimental effects on individuals' labor market outcomes. Those ought to be considered by policy makers who sketch out the referral policy and by caseworkers who implement it on the field, and put into balance against the shorter-term benefits from increased employment.

Bibliography

Aizer, A. and Doyle Jr, J. J. (2015). Juvenile incarceration, human capital, and future crime: Evidence from randomly assigned judges. *The Quarterly Journal of Economics*, 130(2):759–803.

Altmann, S., Falk, A., Jäger, S., and Zimmermann, F. (2018). Learning about job search: A field experiment with job seekers in germany. *Journal of Public Economics*, 164:33–49.

Arni, P. and Schiprowski, A. (2019). Job search requirements, effort provision and labor market outcomes. *Journal of Public Economics*, 169:65–88.

Battisti, M., Giesing, Y., and Laurentsyeva, N. (2019). Can job search assistance improve the labour market integration of refugees? Evidence from a field experiment. *Labour Economics*, 61:101745.

Behncke, S., Frölich, M., and Lechner, M. (2010). A caseworker like me-does the similarity between the unemployed and their caseworkers increase job placements? *The Economic Journal*, 120(549):1430–1459.

Belot, M., Kircher, P., and Muller, P. (2019). Providing advice to jobseekers at low cost: An experimental study on online advice. *The Review of Economic Studies*, 86(4):1411–1447.

Bennmarker, H., Gronqvist, E., and Ockert, B. (2013). Effects of contracting out employment services: Evidence from a randomized experiment. *Journal of Public Economics*, 98:68–84.

Bhuller, M., Dahl, G. B., Løken, K. V., and Mogstad, M. (2020). Incarceration, recidivism, and employment. *Journal of Political Economy*, 128(4):1269–1324.

Bollens, J. and Cockx, B. (2017). Effectiveness of a job vacancy referral scheme. *IZA Journal of Labor Policy*, 6(1):1–24.

Briscese, G., Zanella, G., and Quinn, V. (2020). Improving job search skills: A field experiment on online employment assistance. *IZA Discussion Paper No. 13170*.

Cahuc, P. and Le Barbanchon, T. (2010). Labor market policy evaluation in equilibrium: Some lessons of the job search and matching model. *Labour Economics*, 17(1):196–205.

Card, D., Kluve, J., and Weber, A. (2010). Active labour market policy evaluations: A meta-analysis. *The Economic Journal*, 120(548):F452–F477.

Card, D., Kluve, J., and Weber, A. (2018). What works? A meta analysis of recent active labor market program evaluations. *Journal of the European Economic Association*, 16(3):894–931.

Cheung, M., Egebark, J., Forslund, A., Laun, L., Rodin, M., and Vikström, J. (2019). Does job search assistance reduce unemployment? Experimental evidence on displacement effects and mechanisms. *FN Working Paper No. 1307*.

Correia, S. (2019). Reghdfe: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects.

Cottier, L., Fluckiger, Y., Kempeneers, P., and Lalive, R. (2018). Does job search assistance really raise employment? *IZA Discussion Paper No. 11766*.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. (2013). Do labor market policies have displacement effects? Evidence from a clustered randomized experiment. *The Quarterly Journal of Economics*, 128(2):531–580.

Eriksson, S. and Rooth, D.-O. (2014). Do employers use unemployment as a sorting criterion when hiring? Evidence from a field experiment. *American Economic Review*, 104(3):1014–39.

Faberman, R. J. and Kudlyak, M. (2019). The intensity of job search and search duration. *American Economic Journal: Macroeconomics*, 11(3):327–57.

Farber, H. S., Silverman, D., and Von Wachter, T. (2016). Determinants of callbacks to job applications: An audit study. *American Economic Review*, 106(5):314–18.

Fluchtmann, J., Glenny, A. M., Harmon, N., and Maibom, J. (2021). The dynamics of job search in unemployment: Beyond search effort and reservation wages. *Working Paper*.

Fougère, D., Pradel, J., and Roger, M. (2009). Does the public employment service affect search effort and outcomes? *European Economic Review*, 53(7):846–869.

French, E. and Song, J. (2014). The effect of disability insurance receipt on labor supply. *American Economic Journal: Economic Policy*, 6(2):291–337.

Hensvik, L., Le Barbanchon, T., and Rathelot, R. (2023). Do algorithmic job recommendations improve search and matching? Evidence from a large-scale randomised field experiment in sweden. *Working paper*.

Kling, J. R. (2006). Incarceration length, employment, and earnings. *American Economic Review*, 96(3):863–876.

Kroft, K., Lange, F., and Notowidigdo, M. J. (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *The Quarterly Journal of Economics*, 128(3):1123–1167.

Krug, G. and Stephan, G. (2013). Is the contracting-out of intensive placement services more effective than provision by the PES? Evidence from a randomized field experiment. *IZA*

Discussion Paper No. 7403.

Lancaster, T. (2000). The incidental parameter problem since 1948. *Journal of Econometrics*, 95(2):391–413.

Manoli, D. S., Michaelides, M., and Patel, A. (2018). Long-term effects of job-search assistance: Experimental evidence using administrative tax data. *NBER Working Paper 24422*.

Silva, J. S. and Tenreyro, S. (2011). Further simulation evidence on the performance of the poisson pseudo-maximum likelihood estimator. *Economics Letters*, 112(2):220–222.

Van Belle, E., Caers, R., De Couck, M., Di Stasio, V., and Baert, S. (2019). The signal of applying for a job under a vacancy referral scheme. *Industrial Relations: A Journal of Economy and Society*, 58(2):251–274.

Van Den Berg, G. J., Hofmann, B., and Uhlendorff, A. (2019). Evaluating vacancy referrals and the roles of sanctions and sickness absence. *The Economic Journal*, 129(624):3292–3322.

Van den Berg, G. J., Kjærsgaard, L., and Rosholm, M. (2012). To meet or not to meet (your caseworker) - That is the question. *IZA Discussion Paper No. 6476*.

Wunsch, C. (2013). Optimal Use of Labor Market Policies: The Role of Job Search Assistance. *The Review of Economics and Statistics*, 95(3):1030–1045.

Zuchuat, J., Lalive, R., Osikonimu, A., Pesaresi, L., and Zweimüller, J. (2023). What drives the decline in job offers? *Working paper*.

Appendix

A. Theoretical mechanisms

The value function for the two-channels strategy writes

$$U = \frac{1}{1+\rho\Delta t}\Big[\big(b-\zeta(e)\big)\Delta t + (1-\alpha\,\Delta t)\,U \\ + \alpha_R\,\psi_R\,\Delta t\,\mathbb{E}_{\Gamma_R}\left[\max\big(U, V_R(w)\big)\right] \\ + \alpha_P\,\psi_P\,\Delta t\,\mathbb{E}_{\Gamma_P}\left[\max\big(U, V_P(w)\big)\right]\Big], \tag{A.1}$$

where $V_k(w)$ is the discounted value of a job obtained through channel k.
The system of equations characterizing the optimal search strategy reads

$$\zeta'(\alpha_P^*) = \frac{\psi_P}{\sigma+\rho}\,\Lambda_P(\underline{\omega}^*) \tag{A.2}$$

$$\underline{\omega}^* = b - \zeta(\alpha_P^*) + \alpha_R\frac{\psi_R}{\rho+\sigma_R}\,\Lambda_R(\underline{\omega}^*) + \alpha_P^*\,\zeta'(\alpha_P^*) \tag{A.3}$$

where α_P^* and $\underline{\omega}^*$ are both endogenously determined.
We are interested in studying how the private search effort α_P^* and the reservation wage $\underline{\omega}^*$ change when the application rate through referrals α_R increases exogenously.
Applying the implicit function theorem to the first equation, we obtain:

$$\overbrace{\zeta''(\alpha_P^*)}^{>0}\;d\alpha_P^* = \overbrace{\frac{\psi_P}{\rho+\sigma_P}}^{>0}\;\overbrace{\Lambda_P'(\underline{\omega})}^{<0}\;d\underline{\omega}^*$$

$\zeta''(e) > 0$, as $\zeta(e)$ is convex in e. Also, the surplus function $\Lambda_P(\omega) = \int_\omega^\infty(\kappa-\omega)\;d\Gamma_P(\kappa)$ is decreasing in ω, *i.e.* $\Lambda_P'(\omega) < 0$. This first equation thus shows that $d\alpha_P^*$ and $d\underline{\omega}^*$ are of opposite sign.
Applying the implicit function theorem to the second equation, we get:

$$\begin{aligned} d\underline{\omega}^* &= -\zeta'(\alpha_P^*)\;d\alpha_P^* + \alpha_R\frac{\psi_R}{\sigma+\rho}\Lambda_R'(\underline{\omega}^*)\;d\underline{\omega}^* + \frac{\psi_R}{\sigma+\rho}\Lambda_R(\underline{\omega}^*)\;d\alpha_R + \zeta'(\alpha_P^*)\;d\alpha_P^* + \alpha_P^*\,\zeta''(\alpha_P^*)\;d\alpha_P^* \\ &= \alpha_R\,\frac{\psi_R}{\sigma+\rho}\Lambda_R'(\underline{\omega}^*)\;d\underline{\omega}^* + \frac{\psi_R}{\sigma+\rho}\Lambda_R(\underline{\omega}^*)\;d\alpha_R + \alpha_P^*\,\zeta''(\alpha_P^*)\;d\alpha_P^* \end{aligned}$$

Rearranging the above equation, we obtain:

$$\underbrace{\left(1-\alpha_R\frac{\psi_R}{\rho+\sigma}\Lambda_R'(\underline{\omega}^*)\right)}_{>0}\;d\underline{\omega}^* = \underbrace{\frac{\psi_R}{\rho+\sigma}\;\Lambda_R(\underline{\omega}^*)}_{>0}\;d\alpha_R + \underbrace{\zeta''(\alpha_P^*)\alpha_P^*}_{>0}\;d\alpha_P^*. \tag{A.4}$$

It can be seen from the above equation that an increase in the arrival rate of referrals $d\alpha_R > 0$ is offset either by (i) an increase in the reservation wage, $d\underline{\omega}^* > 0$, or by (ii) a reduction in the personal search effort, $d\alpha_P^* < 0$. Note that the substitution effect between private search effort α_P^* and α_R arises through the convexity of the search cost function $\zeta''(e) > 0$.

B. Data and descriptive statistics

Figure B1: ob search diaries

Assurance-chômage

Preuves des recherches personnelles effectuées en vue de trouver un emploi

A remettre à l'ORP
au plus tard le 5 du mois suivant

Date de réception / du timbre postal

Nom et prénoms	No AVS	Mois et année

Date de l'offre de services (jour mois)	Entreprise, adresse Personne contactée, numéro de tél.	Description du poste	Assignation ORP	Activité: à plein temps	Activité: à temps partiel (%)	Offre de service: par lettre / électronique	Offre de service: visite personnelle	Offre de service: par téléphone	Résultat de l'offre de service: en suspens	entretien	engagement	négatif	Motif

Note: This figure presents the pre-defined job search diary form, in which unemployed record their job search activity.

Figure B2: All applications and non-referred (private) applications
Cumulative distribution functions

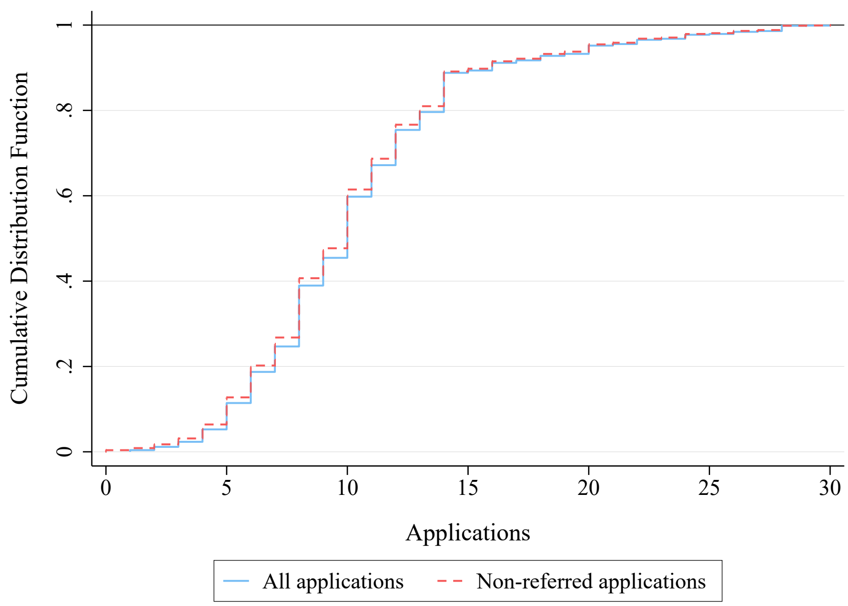

Note: This figure reports the Cumulative Distribution Functions (CDFs) of the monthly number of all and of non-referred (private) job applications sent by job seekers.

Figure B3: Job offers and labor income trajectories

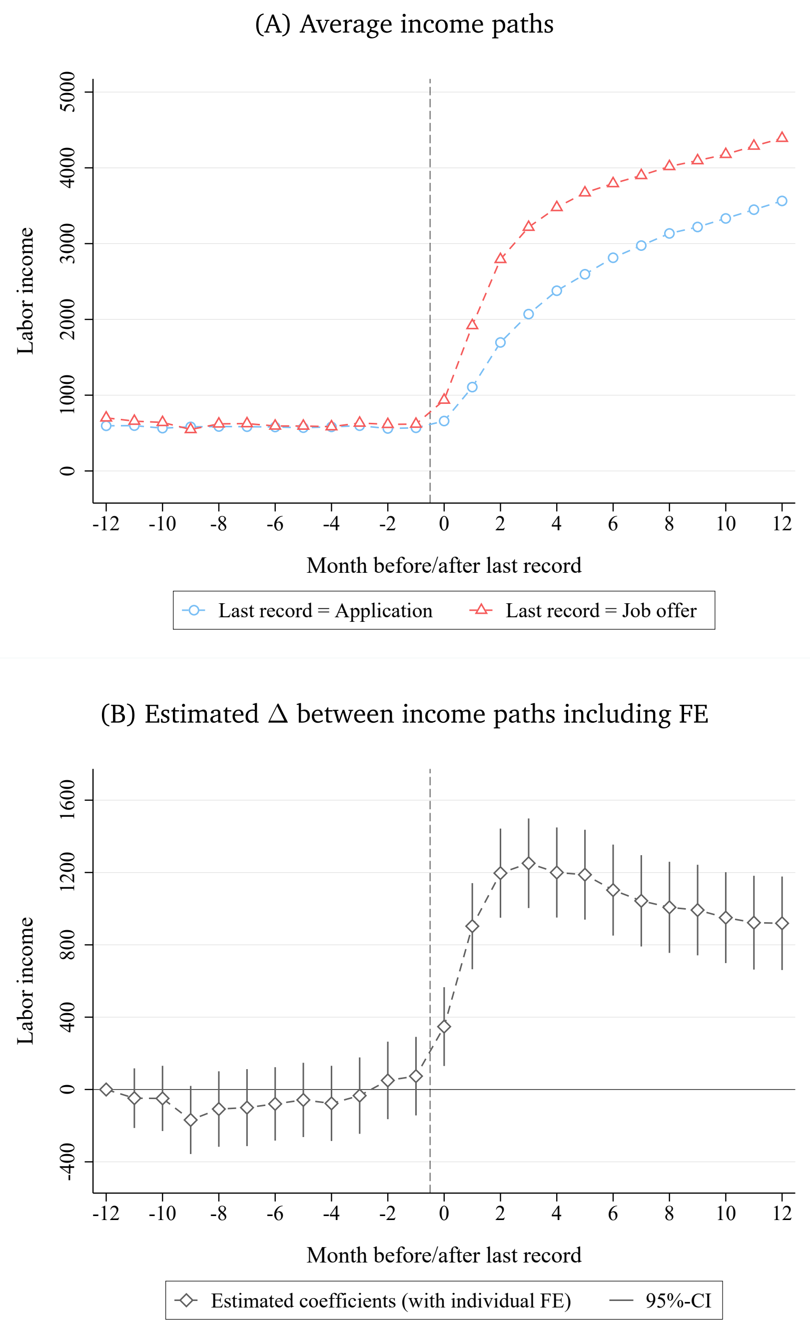

Note: This figure presents evidence on the validity of the job offer indicator contained in the job search diaries as a proxy for job finding. Panel A represents the average monthly income trajectories for two subgroups, before and after reference month $t = 0$. For the first group, made of individuals for whom we do not observe a job offer, the reference month corresponds to the last month of record of the job search diaries (*Last record = Application*). For the other group, the reference month corresponds to the month of the last job offer recorded in the forms (*Last record = Job offer*). Panel B presents the estimation result of parameters Δ_t of the following event-study equation: $\textit{Labor income}_{it} = \alpha_i + \delta_t + \Delta_t \cdot \mathbb{I}(i \in \textit{Last record} = \textit{Job offer}) + \varepsilon_{it}$. The Δ_t capture the difference in the average income trajectories between the two groups of individuals, accounting for individual fixed effects α_i. Errors are clustered at the individual level and 95%-CI are reported.

Table B1: Databases structure

	# applications		# ind.-monthly		# individuals	
A. Main sample						
No CwR	684556	[97.842]	59366	[88.762]	10869	[69.602]
CwR	15096	[2.158]	7516	[11.238]	4747	[30.398]
Total	699652	[100.000]	66882	[100.000]	15616	[100.000]
B. Auxiliary sample						
No CwR	22305	[97.283]	2494	[87.264]	473	[68.353]
CwR	623	[2.717]	364	[12.736]	219	[31.647]
Total	22928	[100.000]	2858	[100.000]	692	[100.000]

Note: This table reports details on the structure of our main and auxiliary samples. For each of the two samples, it reports the number of applications we observe (distinguishing between referred and non-referred applications), the number of individual-monthly units (distinguishing between units when at least one referral is assigned and units when the referral policy is not used) and the number of unemployed (distinguishing between ever and never-referred individuals). Relative frequencies are reported in brackets.

Table B2: Job seekers' characteristics in the auxiliary sample, Ever CwR vs. Never CwR

	Never CwR		Ever CwR		Comparison	
A. Numeric variables	Mean	SDV	Mean	SDV	t-stat.	p-value
Age	39.513	(10.595)	39.550	(10.970)	-0.043	0.965
Female	0.501	(0.501)	0.460	(0.499)	1.022	0.307
Swiss	0.583	(0.494)	0.477	(0.501)	2.668	0.008
Elasped unemp. duration	5.196	(5.495)	6.852	(5.394)	-3.793	0.000
B. Categorical variables	Freq.	%	Freq.	%	χ^2-stat.	p-value
Education levels						
Primary	134	33.005	105	47.297	15.874	0.000
Secondary	175	43.103	87	39.189		
Tertiary	97	23.892	30	13.514		
Observations	406		222			

Note: This table reports descriptive statistics on job seekers' characteristics for the auxiliary sample of analysis. It distinguishes between the group of individuals who never get referred referrals (Never CwR) and those who receive at least one caseworker referral (Ever CwR). Panel A reports descriptive statistics on numeric variables, together with a test for the equality of the means across the two groups. Panel B reports absolute and relative frequencies for the two groups, together with a χ^2-independence test (between the group indicator and the categorical variable of interest).

Table B3: Applications and job seekers' characteristics in the main sample, CwR vs. no CwR

	No CwR		CwR		Comparison	
A. Numeric variables	Mean	SDV	Mean	SDV	t-stat.	p-value
Age	39.803	(11.348)	40.221	(11.260)	-4.475	0.000
Female	0.473	(0.499)	0.477	(0.499)	-1.018	0.309
Swiss	0.520	(0.500)	0.526	(0.499)	-1.588	0.112
Past labor income [in kCHF]	4.954	(5.465)	4.509	(4.489)	9.928	0.000
Unemployment history	0.088	(0.141)	0.095	(0.146)	-6.106	0.000
Previous work-time %	90.609	(17.738)	90.828	(17.094)	-1.276	0.202
Elapsed unemp. duration	6.410	(4.496)	6.585	(4.397)	-4.716	0.000
B. Categorical variables	Freq.	%	Freq.	%	χ^2-stat.	p-value
Education levels						
Primary	187618	(27.407)	4201	(27.829)	52.048	0.000
Secondary	398507	(58.214)	9037	(59.864)		
Tertiary	98431	(14.379)	1858	(12.308)		
Civil status						
Married/Separated	349421	(51.043)	7635	(50.576)	21.057	0.000
Single	237260	(34.659)	5112	(33.863)		
Divorced	90497	(13.220)	2186	(14.481)		
Widow(er)	7378	(1.078)	163	(1.080)		
Canton						
BE	106092	(15.498)	2881	(19.085)	674.489	0.000
SG	108117	(15.794)	3113	(20.621)		
VD	168774	(24.655)	2635	(17.455)		
ZG	73401	(10.722)	1422	(9.420)		
ZH	228172	(33.331)	5045	(33.419)		
Desired occupation						
Agriculture	7914	(1.156)	122	(0.808)	168.538	0.000
Industry & Craft	100685	(14.708)	2158	(14.295)		
IT	40135	(5.863)	870	(5.763)		
Construction	45953	(6.713)	1124	(7.446)		
Commercial	127150	(18.574)	2884	(19.104)		
Hotelling	121599	(17.763)	2970	(19.674)		
Adminsistrative	117216	(17.123)	2515	(16.660)		
Health & Educ.	72563	(10.600)	1209	(8.009)		
Other	51341	(7.500)	1244	(8.241)		
Observations	684556		15096			

Note: This table reports descriptive statistics on job seekers' characteristics at the application-level for the main sample of analysis. It reports evidence on how the pool of referred applications (CwR) compare to non-referred ones (no CwR). Panel A reports descriptive statistics on numeric variables, together with a test for the equality of the means across the two groups. Panel B reports absolute and relative frequencies for the two groups, together with a χ^2-independence test (between the group indicator and the categorical variable of interest).

Table B4: Job search diaries, coding of the information and application types

			Callback	Still open	Job offer	Negative	Freq.	%	Cum. %
Non-censored application	Non-censored interview	Job offer	☑	☑	☑	☑	88	0.013	0.013
			☑	☑	☑	☐	681	0.097	0.110
			☑	☐	☑	☑	62	0.009	0.119
			☑	☐	☑	☐	1209	0.173	0.292
			☐	☑	☑	☑	90	0.013	0.304
			☐	☐	☑	☑	619	0.088	0.393
			☐	☑	☑	☐	631	0.090	0.483
			☐	☐	☑	☐	1968	0.281	0.764
		No job offer	☑	☑	☐	☑	4154	0.594	1.358
			☑	☐	☐	☑	6456	0.923	2.281
	Censored interview		☑	☑	☐	☐	5232	0.748	3.029
			☑	☐	☐	☐	4910	0.702	3.730
	No interview		☐	☑	☐	☑	140497	20.081	23.811
			☐	☐	☐	☑	315995	45.165	68.976
Censored application			☐	☑	☐	☐	176118	25.172	94.148
			☐	☐	☐	☐	40942	5.852	100.000
Total							699652	100.000	

Note: This table reports the distribution of the $2^4 = 16$ types of applications that are recorded in the database, based on the information contained in the 4 tick boxes *Interview*, *Still open*, *Job offer* and *Negative*. Given the sequential callback and job offer indicators, we impute a job interview for all applications that have a job offer recorded.

C. Application-level analysis: applications success and vacancies characteristics

Figure C1: Share of referred applications within individual-monthly units

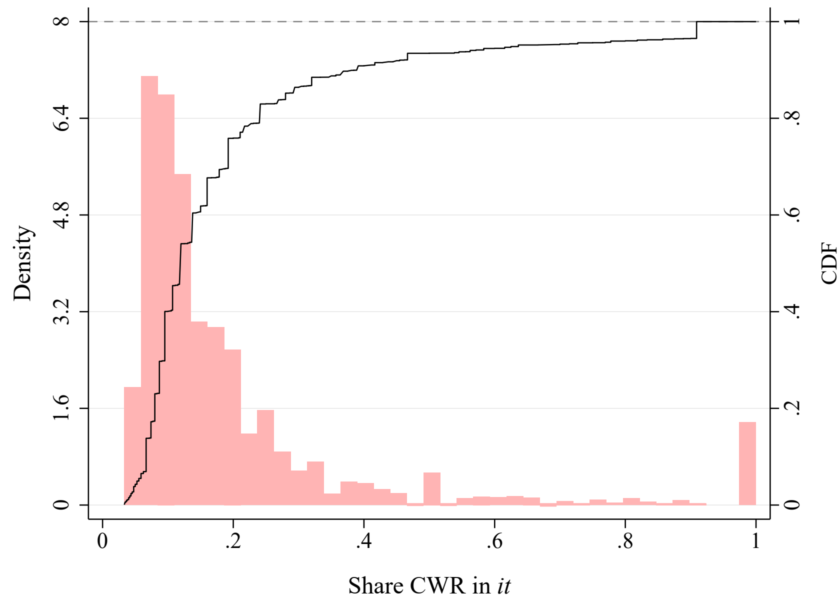

Note: This figure displays the distribution of the share of referred applications in the total number of applications sent by individual i in unemployment month t. The figure focuses on individual-monthly units when at least one referral is assigned.

Figure C2: CwR effects on job offer conversion probability

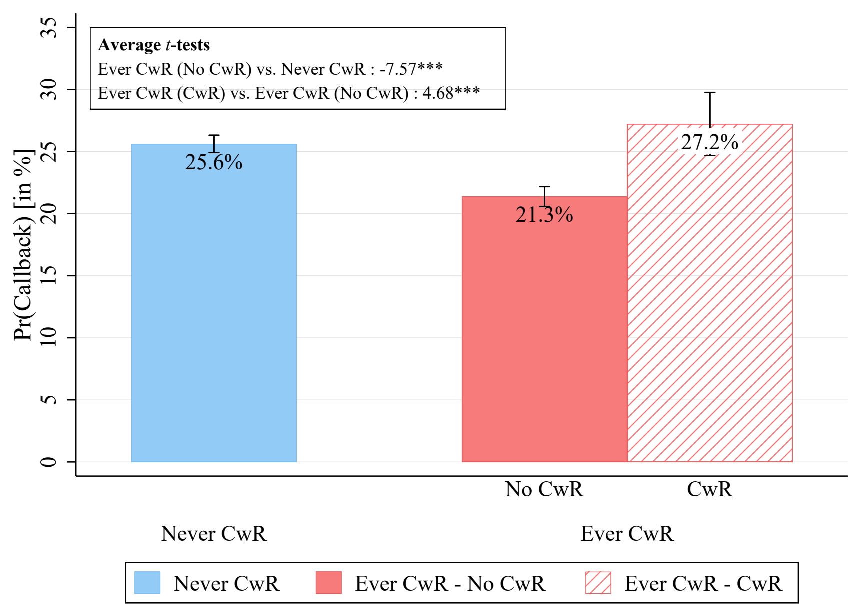

Note: This figure reports descriptive evidence of the effects of referrals on the conversion probability of interviews into job offers. Each graph reports the average share of interviews that end up in a job offer. The graph distinguishes between three groups of interviews: (i) interviews obtained by individuals who are never treated (Never CwR), (ii) interviews obtained by individuals who are treated at least once but who do not directly result from a referral (Ever CwR & No CwR) and (iii) interviews resulting from a referral (Ever CwR & CwR). 95%-confidence intervals for the average probability in each group are reported. Pairwise t-tests for the equality of the average probabilities across groups are reported in boxes. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Figure C3: CwR effects on the callback probability, heterogeneous effects

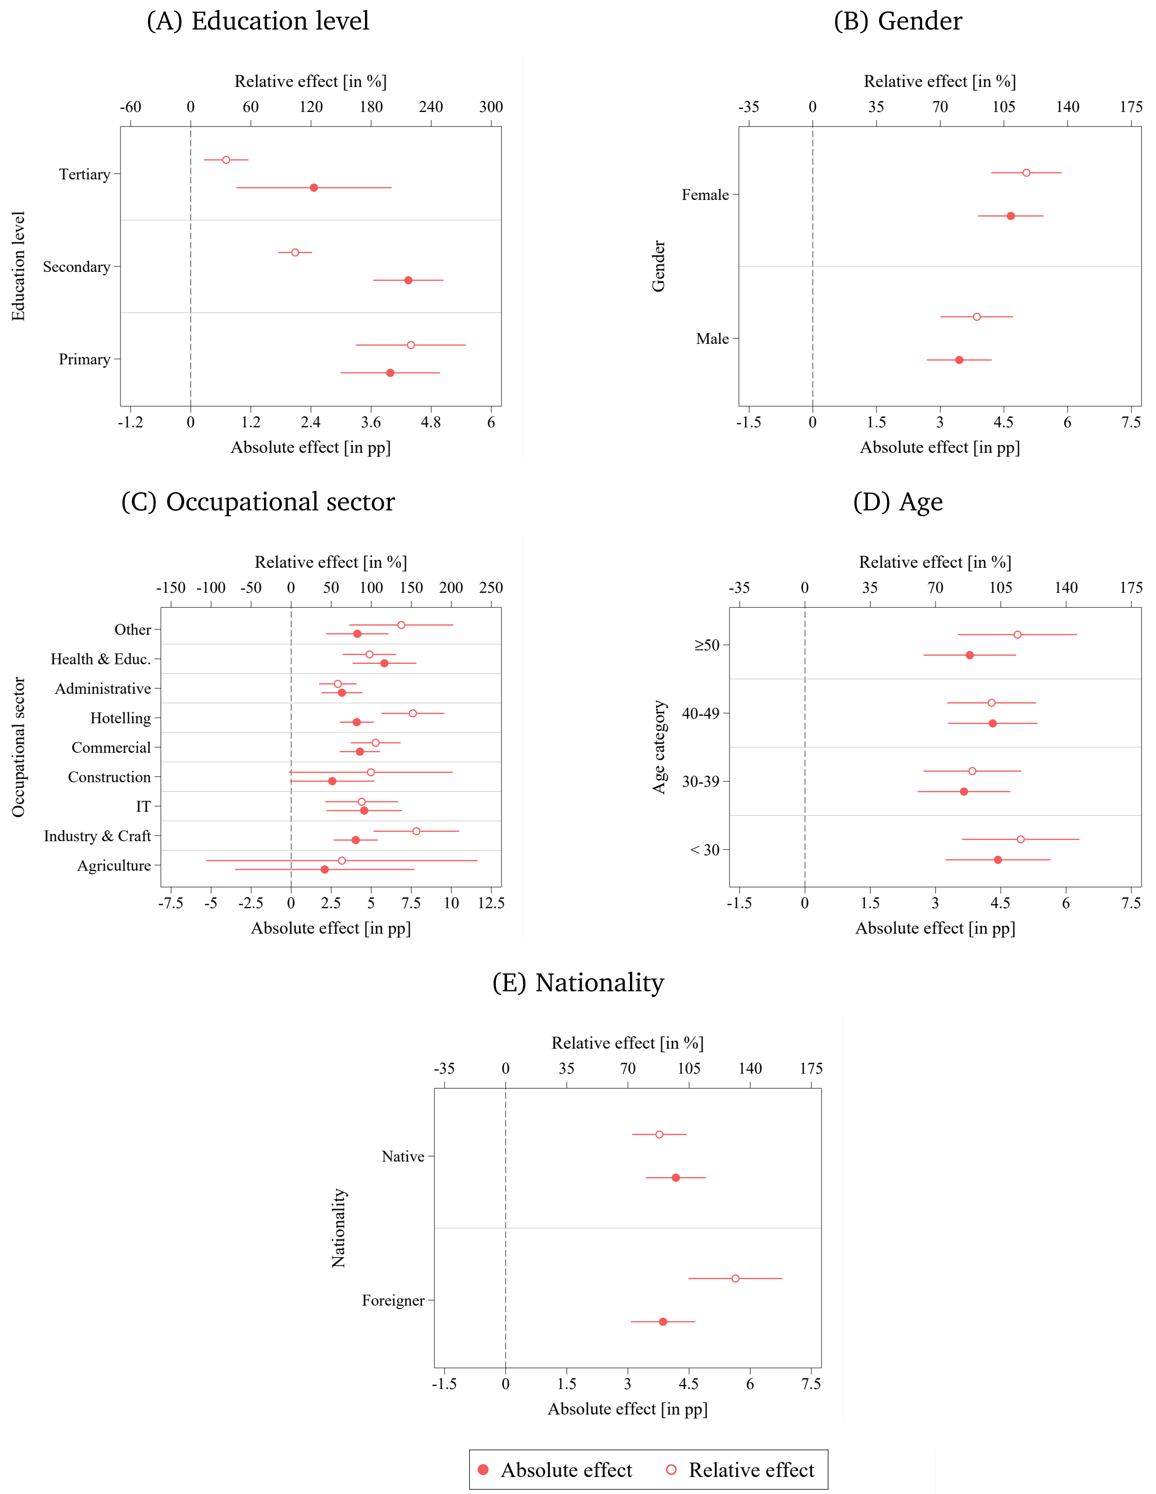

Note: This figure presents evidence of heterogeneous treatment effects of caseworker referrals on the callback probability. It reports graphically the estimation results from equation (5.1), including individual-duration fixed effects, where the treatment variable is interacted with the categorical characteristic indicated in each panel. Effects are reported both in absolute terms (in pp) and in relative terms (in %, relative to the baseline probability of a positive callback in each group in the absence of referrals). 90% confidence intervals based on clustered standard errors at the individual level are reported.

Figure C4: CwR effects on the job offer probability, heterogeneous effects

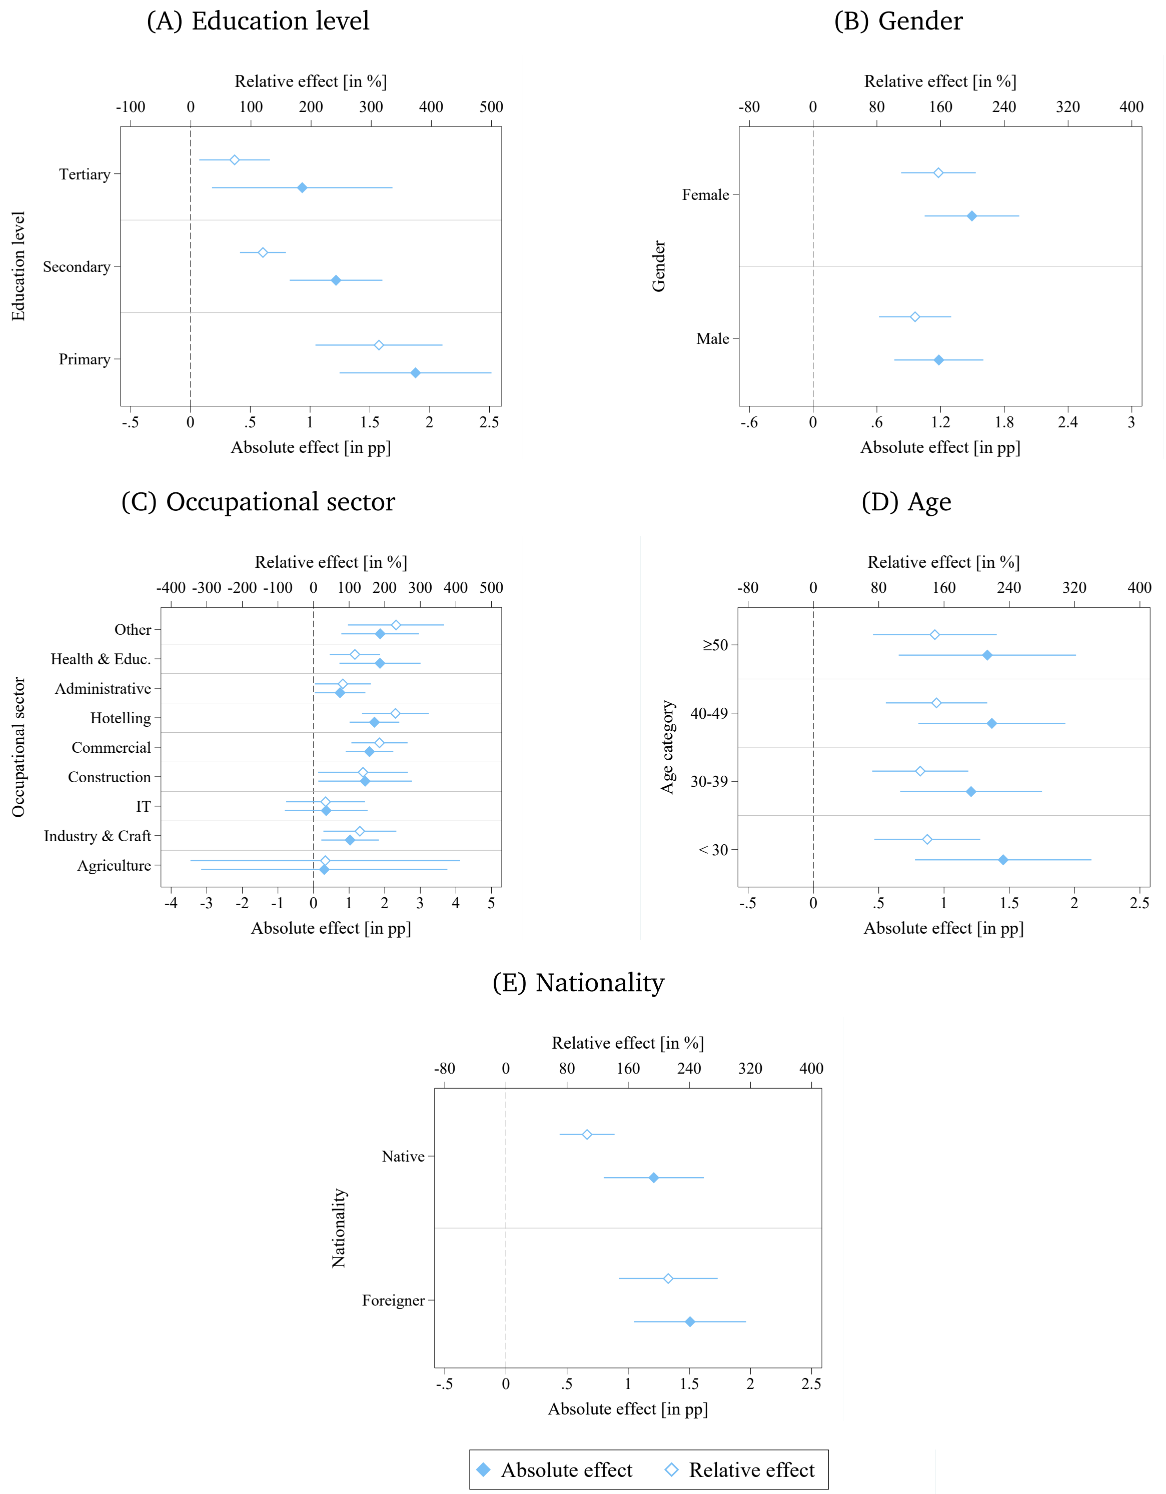

Note: This figure presents evidence of heterogeneous treatment effects of caseworker referrals on the job offer probability. It reports graphically the estimation results from equation (5.1), including individual-duration fixed effects, where the treatment variable is interacted with the categorical characteristic indicated in each panel. Effects are reported both in absolute terms (in pp) and in relative terms (in %, relative to the baseline probability of a job offer in each group in the absence of referrals). 90% confidence intervals based on clustered standard errors at the individual level are reported.

Figure C5: CwR effects per canton

(A) Callback probability

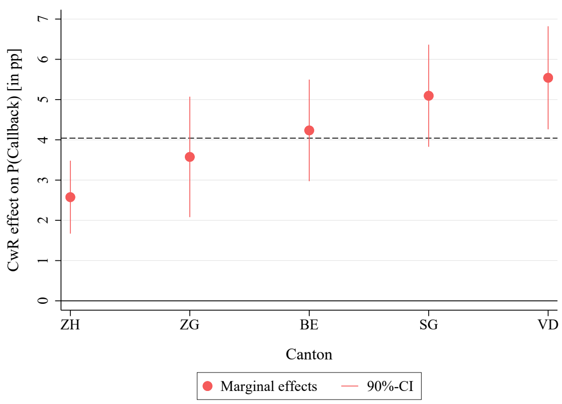

(B) Job offer probability

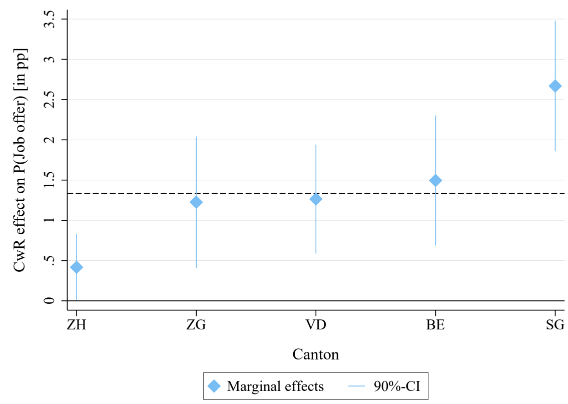

Note: This figure presents evidence of heterogeneous treatment effects of caseworker referrals with respect to local public institutions, *i.e.* the canton where unemployed job seekers are registered. Panel A reports evidence for the callback probability c_{ait}, while panel B reports evidence for the job offer probability o_{ait}. Each coefficient is obtained based on the subsample of individuals registered in the canton. 90%-confidence intervals based on clustered standard errors at the individual level are reported.

Figure C6: CwR effects per caseworker

(A) Callback probability

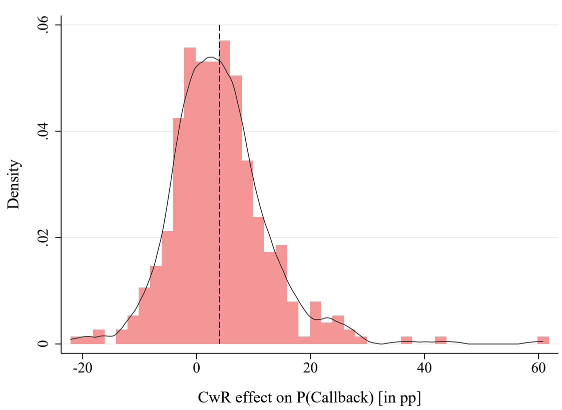

(B) Job offer probability

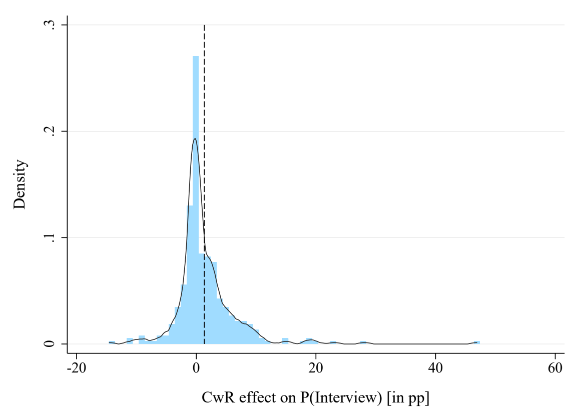

Note: This figure depicts the distribution of caseworker-specific treatment effects of referrals on application success probability. Panel A reports evidence for the callback probability c_{ait}, while panel B reports evidence for the job offer probability o_{ait}. Each coefficient is obtained based on the subsample of individuals supervised by a specific caseworker. The average of the treatment effects is reported as a dashed line.

Figure C7: CwR effects per caseworker and intensity in the use of the policy

(A) Callback probability

(B) Job offer probability

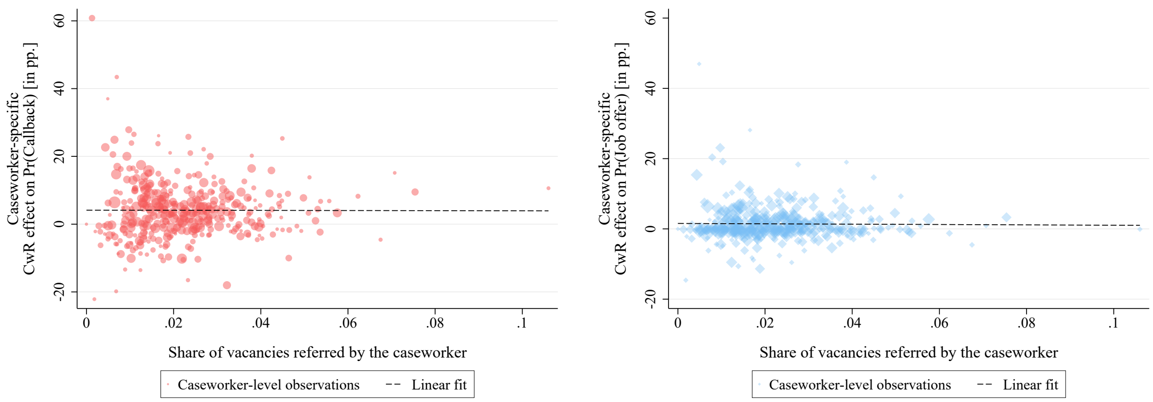

Note: This figure reports the relationship between the caseworker-specific policy treatment effect and the intensity in the use of the policy by caseworkers (measured by the share of referred applications, among all applications associated to each caseworker). Dots are weighted according to the number of applications associated to each caseworker. Panel A reports the relationship for the callback indicator c_{ait}, while panel B reports results for the job offer indicator o_{ait}.

Figure C8: Duration dependence in referrals' effects on applications success probability

(A) Callback probability

(B) Job offer probability

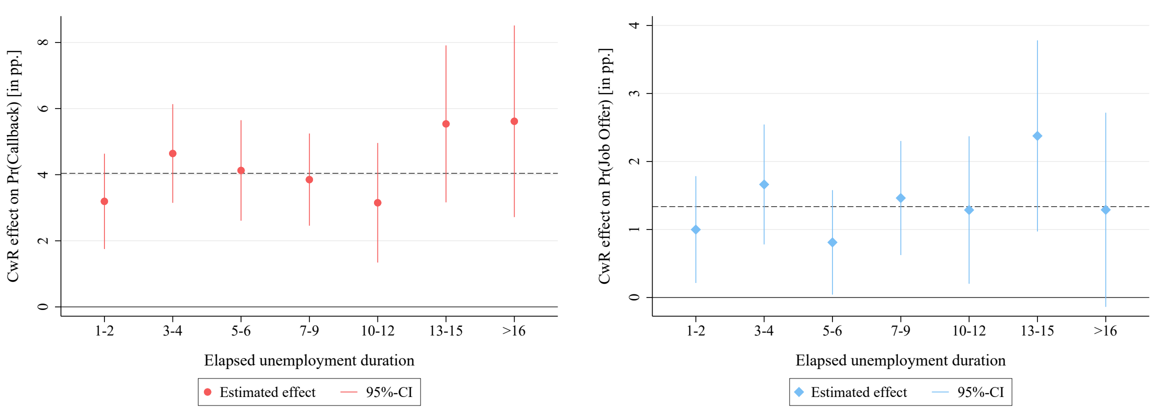

Note: This figure reports the estimation result of equation (5.1) where the treatment indicator is interacted with (categorized) elapsed unemployment duration. Panel A reports results for the callback indicator c_{ait}, while panel B reports results for the job offer indicator o_{ait}. 95%-CI intervals based on clustered standard errors at the individual level are reported.

Figure C9: Information censoring, share of censored c_{ait} and o_{ait}

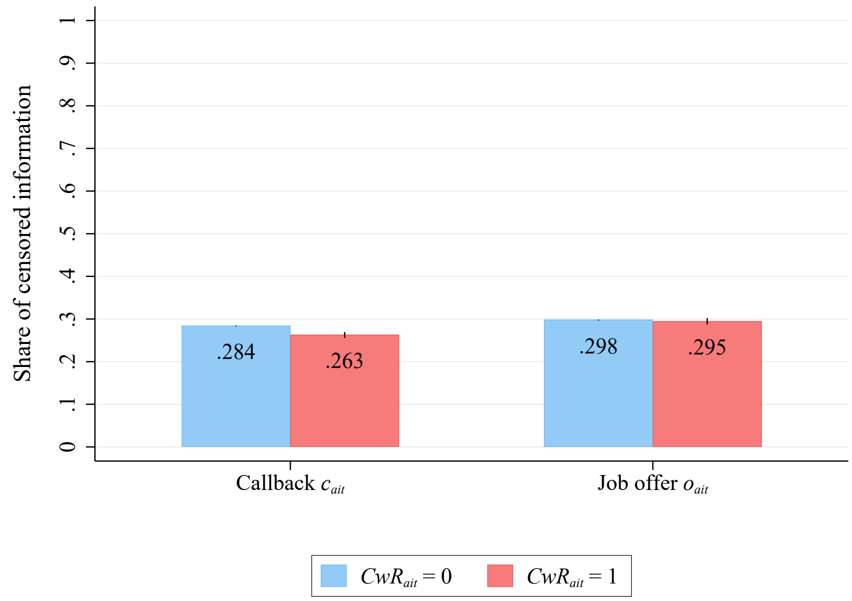

Note: This figure reports the censoring rate of the callback and job offer indicators. Averages are weighted according to the inverse number of applications sent in each individual-monthly unit.

Figure C10: CwR effects restricted to the first use of the policy

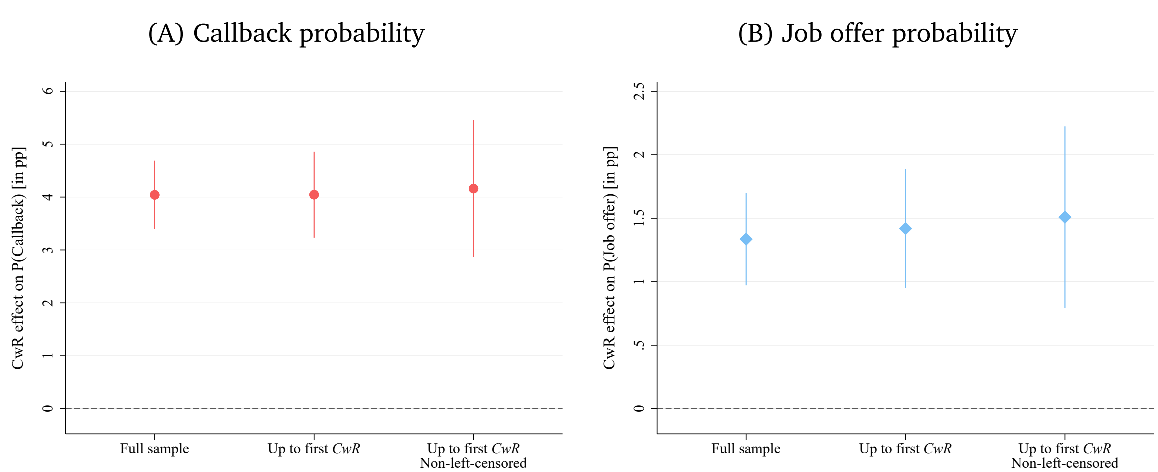

Note: This figure reports the estimation result of equation (5.1) on restricted samples. For each individual, we only consider observations up to the first use of the referral policy. Baseline estimates (Full sample) are reported together with estimates based on two restricted samples: Up to first CwR and Up to first CwR & Non-left censored spells. Panel A reports results for the callback indicator c_{ait}, while panel B reports results for the job offer indicator o_{ait}. 95%-CI intervals based on clustered standard errors at the individual level are reported.

Figure C11: CwR and labor income ratio after/before (25$^{\text{th}}$, 50$^{\text{th}}$ and 75$^{\text{th}}$ percentiles)

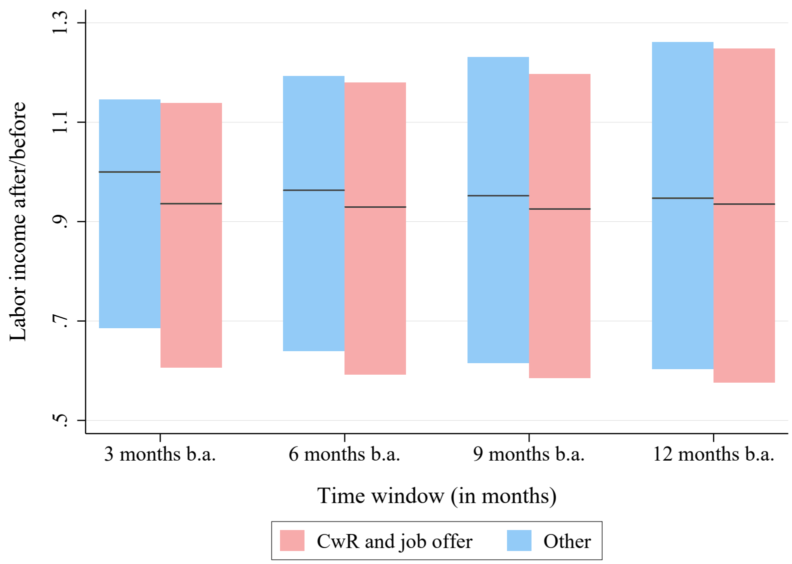

Note: This figure reports descriptive evidence on the labor income ratio after/before unemployment, at the job seeker level. It distinguishes between job seekers who get a job offer through a caseworker referral and other job seekers. The graph reports summary statistics on the distribution of this ratio (25$^{\text{th}}$, 50$^{\text{th}}$ and 75$^{\text{th}}$ percentiles). The labor income ratio after/before is computed based on social security data, for different time windows (3,6,9,12 months before and after unemployment).

Figure C12: Work-time percentage of job vacancies

(A) CwR and work-time percentage, all applications

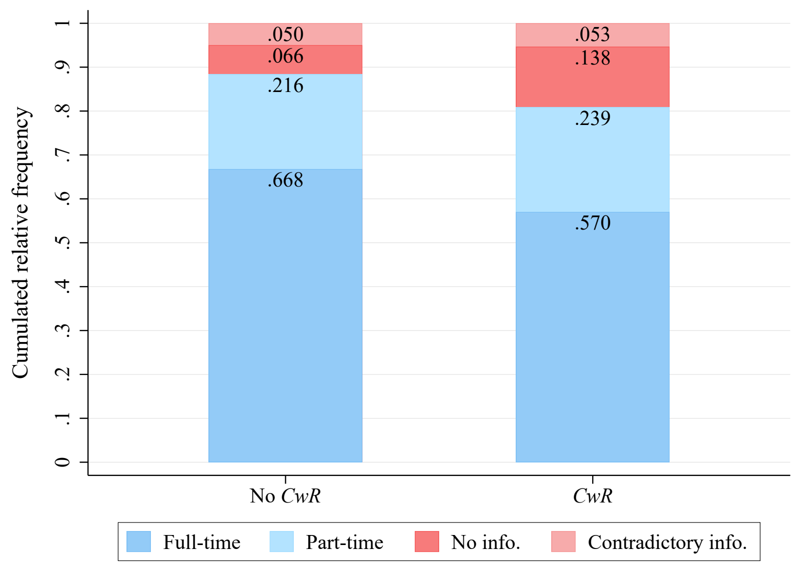

(B) CwR and work-time percentage, unambiguous applications

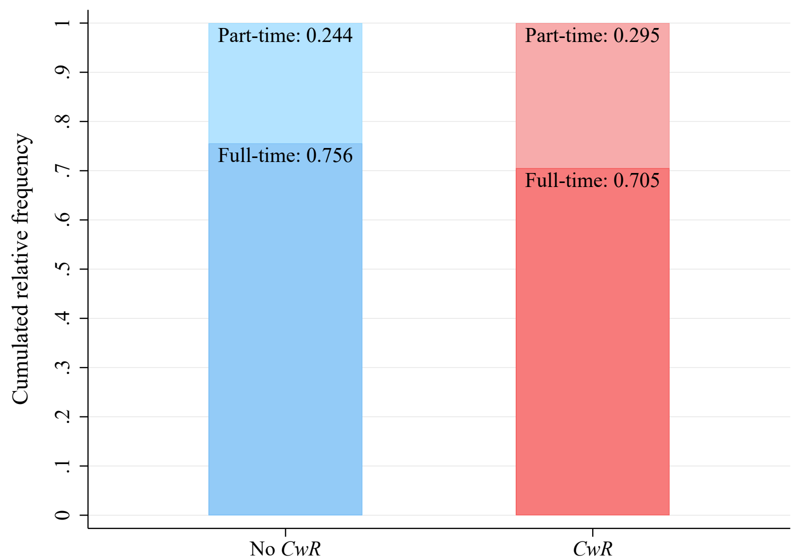

Note: This figure reports descriptive evidence of the share of applications that target vacancies with different work-time percentage. Panel A reports the information for all applications, where the information on work-time percentage can be ambiguous. Panel B reports the information for unambiguous applications, *i.e.* applications that can unambiguously be characterized as part-time or full-time. The two graphs distinguish between referred and non-referred applications.

Figure C13: CwR effects on skills requirements

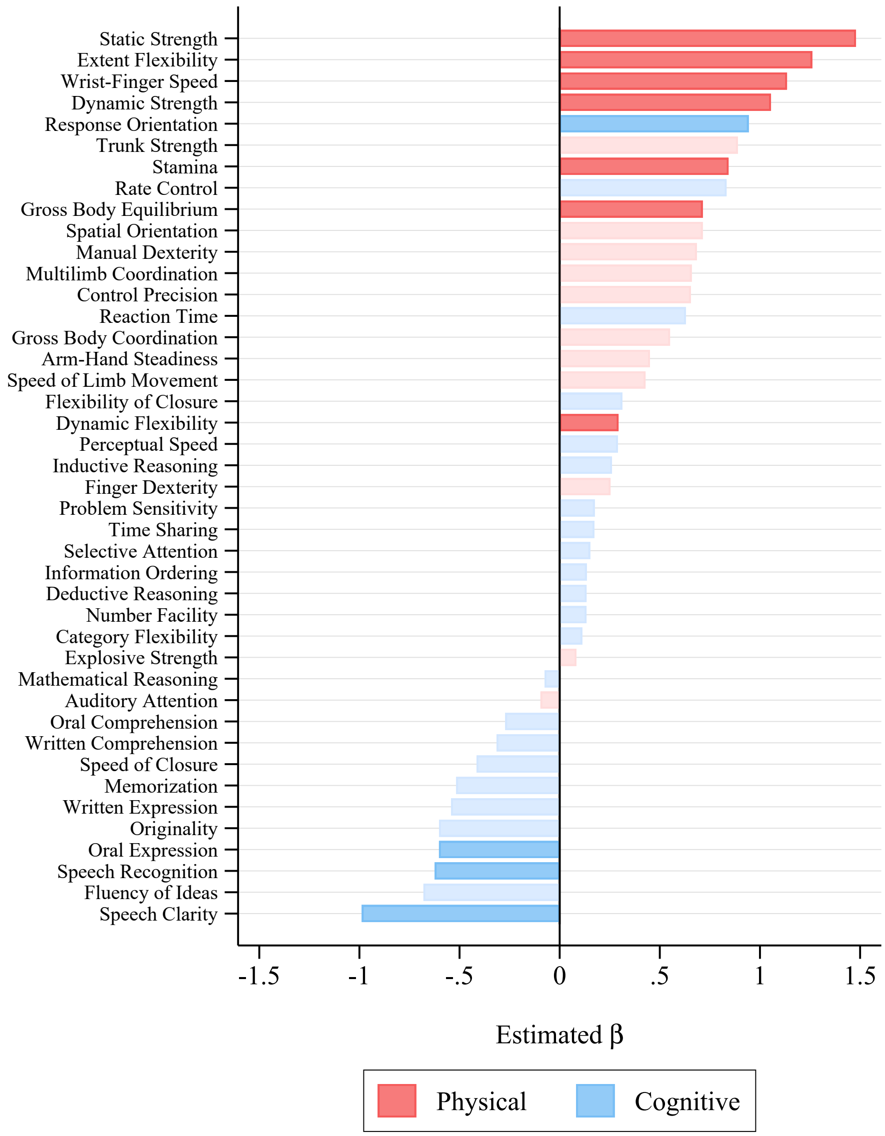

Note: This figure reports the estimated effects of caseworker referrals on each skill (requirement) in the *O*Net* repository. Reported coefficients are obtained by estimating equation (5.1), where the dependent variable corresponds to one of the skills on the ordinate. The graph distinguishes between cognitive skills (in blue) and physical skills (in red). Darker colors indicate significance of the effect at the 10% significance level.

Table C1: List of controls

Variable	Description	Type
A. Application characteristics X^1_{ait}		
Mode$_{ait}$	Mode of application (written, personal, phone, multiple)	Categorical
Application rank$_{ait}$	Within-month rank of application	Categorical
B. Individual characteristics X^2_{it}		
Gender$_{it}$	Gender	Binary
Age$_{it}$	Age category	Categorical
Education$_{it}$	Education level	Categorical
Function$_{it}$	Professional function	Categorical
Residence$_{it}$	Residence permit	Categorical
Civil$_{it}$	Civil status	Categorical
LLMC$_{it}$	Local Labor Market Conditions (Occupational)	Categorical
Caseworker$_{it}$	Caseworker in charge	Categorical
$\ln$(Previous income$_i$)	Log of the (average) previous monthly labor income	Numerical

Note: This table reports the list of variables used as controls in our analyses.

Table C2: Caseworker referrals and information censoring

	(1)	(2)	(3)	(4)
Panel A: Censored callback c_{ait}				
Caseworker referral	-2.108***	-0.074	0.120	-0.459
	(0.789)	(0.549)	(0.487)	(0.430)
Individual controls	No	Yes	Yes	No
Application controls	No	Yes	Yes	Yes
Individual FE	No	No	Yes	No
Unemployment duration FE	No	No	Yes	No
Individual $\times$ Unemp. duration FE	No	No	No	Yes
# identifying observations	699652	699652	699631	699385
# singleton observations	0	0	21	267
# total observations	699652	699652	699652	699652
# parameters	2	419	15986	66633
# individuals	15616	15616	15616	15616
Adj.-R^2	0.000	0.273	0.425	0.543
Panel B: Censored job offer o_{ait}				
Caseworker referral	-0.274	1.798***	2.115***	1.273***
	(0.832)	(0.577)	(0.491)	(0.437)
Individual controls	No	Yes	Yes	No
Application controls	No	Yes	Yes	Yes
Individual FE	No	No	Yes	No
Unemployment duration FE	No	No	Yes	No
Individual $\times$ Unemp. duration FE	No	No	No	Yes
# identifying observations	699652	699652	699631	699385
# singleton observations	0	0	21	267
# total observations	699652	699652	699652	699652
# parameters	2	419	15986	66633
# individuals	15616	15616	15616	15616
Adj.-R^2	-0.000	0.279	0.435	0.555

Note: This table reports the estimation of equation (5.1), where the dependent variable corresponds to the censoring indicators for callbacks and job offers. Panel A refers to callbacks, while panel B report results for job offers. Application-level observations are weighted according to the inverse of the number of applications sent within individual-monthly units. Regression coefficients and standard errors (in parentheses) are reported in percentage points. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table C3: CwR effects on applications success probability, alternative duration dimensions

Duration dimension	Monthly		Bi-weekly		Weekly	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Callback probability c_{ait}						
Caseworker referral	4.298***	4.042***	4.459***	3.794***	4.062***	3.668***
	(0.320)	(0.330)	(0.351)	(0.347)	(0.326)	(0.370)
Individual controls	Yes	No	Yes	No	Yes	No
Application controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	No	Yes	No	Yes	No
Unemployment duration FE	Yes	No	Yes	No	Yes	No
Individual $\times$ Unemp. dur. FE	No	Yes	No	Yes	No	Yes
# identifying observations	699631	699385	699631	694284	699631	658777
# singleton observations	21	267	21	5368	21	40875
# total observations	699652	699652	699652	699652	699652	699652
# parameters	15658	66633	15682	121824	15730	186084
# individuals	15616	15616	15616	15616	15616	15616
Adj.-R^2	0.110	0.170	0.114	0.227	0.111	0.263
Panel B: Job offer probability o_{ait}						
Caseworker referral	1.303***	1.336***	1.383***	1.078***	1.179***	1.015***
	(0.166)	(0.186)	(0.196)	(0.188)	(0.169)	(0.202)
Individual controls	Yes	No	Yes	No	Yes	No
Application controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	No	Yes	No	Yes	No
Unemployment duration FE	Yes	No	Yes	No	Yes	No
Individual $\times$ Unemp. dur. FE	No	Yes	No	Yes	No	Yes
# identifying observations	699631	699385	699631	694284	699631	658777
# singleton observations	21	267	21	5368	21	40875
# total observations	699652	699652	699652	699652	699652	699652
# parameters	15658	66633	15682	121824	15730	186084
# individuals	15616	15616	15616	15616	15616	15616
Adj.-R^2	0.083	0.118	0.081	0.177	0.070	0.207

Note: This table reports the estimation of equation (5.1), for different duration dimensions: Monthly, Bi-weekly and Weekly. Panel A reports results for the callback probability, while panel B report results for the job offer probability. Application-level observations are weighted according to the inverse of the number of applications sent within individual-duration units. Regression coefficients and standard errors (in parentheses) are reported in percentage points. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table C4: CwR effects on applications success probability, non-linear specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Callback probability c_{ait}						
Caseworker referral	6.034***	6.208***	4.482***	5.565***	4.677***	4.686***
	(0.468)	(0.490)	(0.260)	(0.323)	(0.283)	(0.281)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Application controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	No	Yes	No	Yes	No	Yes
Unemployment duration FE	No	Yes	No	Yes	No	Yes
Individual heterogeneity	∅	FE	∅	FE	∅	CRE
Observations	434837	434837	434837	434837	434837	434837
Panel B: Job offer probability o_{ait}						
Caseworker referral	4.361***	4.719***	2.877***	3.411***	3.082***	3.108***
	(0.537)	(0.531)	(0.253)	(0.296)	(0.274)	(0.272)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Application controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	No	Yes	No	Yes	No	Yes
Unemployment duration FE	No	Yes	No	Yes	No	Yes
Individual heterogeneity	∅	FE	∅	FE	∅	CRE
Observations	188530	188530	188530	188530	188530	188530

Note: This table reports estimation results from equation (5.1), for the callback (panel A) and job offer (panel B) probabilities. Columns (1) and (2) correspond to linear probability models, columns (3) and (4) to logit models while columns (5) and (6) report probit results. Specifications within a class of models differ with respect to the inclusion of individual heterogeneity components (either under the form of fixed effects FE or correlated random effects CRE). The estimation sample is common across all specifications and is restricted to individuals who exhibit variability in the dependent variable. This restriction is required for the identification of the conditional logit model. Average marginal effects in percentage points are reported together with standard errors (in parentheses). For the conditional logit specification, we use the following transformation: $ME = \widehat{\beta} \cdot \widehat{y}_{ait}$, where $y_{ait} = c_{ait}, o_{ait}$. Application-level observations are weighted according to the inverse of the number of applications sent within individual-duration units. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table C5: CwR effects on targeted work-time percentage
Restriction to individual-duration units without miscoded information

Previous work-time percentage	All	Part-time	Full-time	Missing
	(1)	(2)	(3)	(4)
Dependent variable : Part-time PT_{ait}				
Caseworker referral	3.449***	2.154**	4.103***	3.223***
	(0.472)	(1.061)	(0.681)	(0.815)
Application controls	Yes	Yes	Yes	Yes
Individual × Unemp. duration FE	Yes	Yes	Yes	Yes
# identifying observations	471306	93419	247243	130644
# singleton observations	221	55	121	45
# total observations	471527	93474	247364	130689
# parameters	46135	9512	23562	13071
# individuals	13266	13266	13266	13266
Adj.-R^2	0.785	0.808	0.644	0.791

Note: This table reports the estimation results of equation (5.1), where the dependent variable corresponds to the part-time indicator PT_{ait}. Only individual-duration units it without miscoding in the job vacancies work-time percentage are kept for estimation. All estimates are based on a specification with individual-duration fixed effects μ_{it}. Application-level observations are weighted according to the inverse of the number of applications sent within individual-monthly units. Regression coefficients and standard errors (in parentheses) are reported in percentage points. Columns (1) refers to the whole sample of applications for which the occupation rate is unambiguous. Columns (2),(3) and (4) correspond to estimates obtained on subsamples based on the previous occupation rate held by the job seekers, which can either be Part-time, Full-time or Missing. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table C6: Baseline outcomes results for the auxiliary sample

	Callback prob.		Job offer prob.		Part-time	
	(1)	(2)	(3)	(4)	(5)	(6)
Caseworker referral	5.555***	5.243***	1.725*	1.388*	4.301	5.085**
	(1.704)	(1.759)	(0.899)	(0.711)	(2.713)	(2.271)
Application controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	No	Yes	No	Yes	No
Unemployment duration FE	Yes	No	Yes	No	Yes	No
Individual $\times$ Unemp. duration FE	No	Yes	No	Yes	No	Yes
# identifying observations	22913	22772	22913	22772	20580	20409
# singleton observations	15	156	15	156	14	185
# total observations	22928	22928	22928	22928	20594	20594
# parameters	709	2833	709	2833	695	2696
# individuals	692	692	692	692	692	692
Adj.-R^2	0.199	0.310	0.107	0.159	0.598	0.666

Note: This table reports the estimation results of equation (5.1) for the auxiliary sample. The application-level dependent variable is either the callback indicator c_{ait}, the job offer indicator o_{ait} or the part-time indicator PT_{ait}, in columns (1)-(2), (3)-(4) and (5)-(6) respectively. All specifications include either individual and duration fixed effects (μ_i and μ_t) or individual-duration fixed effects (μ_{it}). Application-level observations are weighted according to the inverse of the number of applications sent within individual-monthly units. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

D. Individual-monthly analysis: search effort and labor matching

Figure D1: Distribution of caseworker stringency
Unconditional and conditional on individual fixed effects

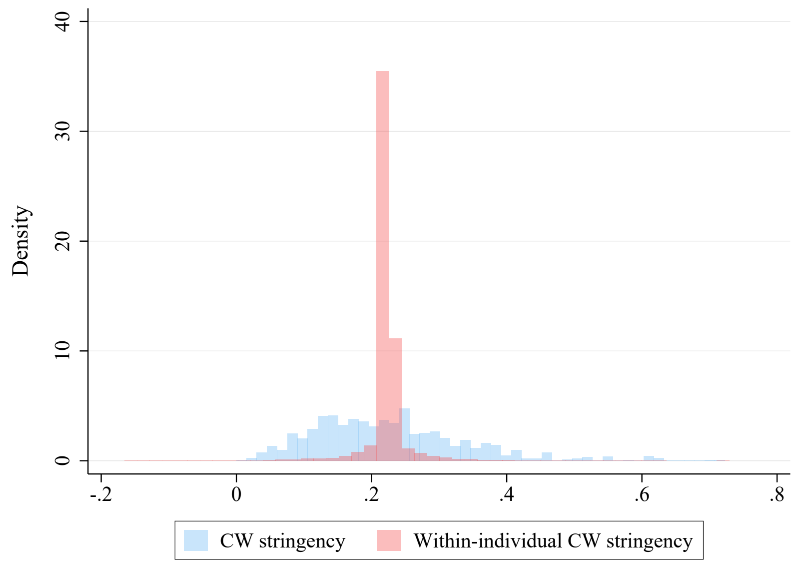

Note: This figure reports the empirical distribution of the time-varying caseworker stringency measure, both unconditional and conditional on individual fixed effects. In the second case, residual caseworker stringency is obtained by means of a linear regression on individual fixed effects.

Figure D2: CIA and ER, distribution of the partial correlation t-stats

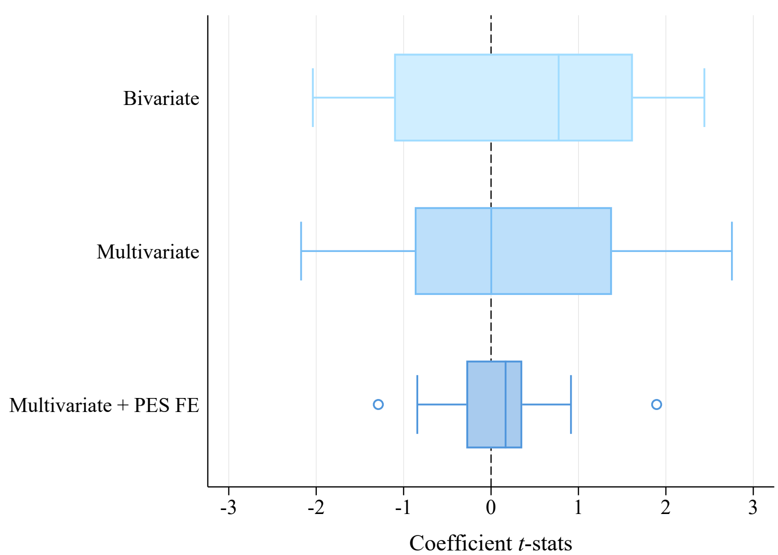

Note: This figure plots the distribution of the t-stats associated with the coefficients from the CIA and ER assessment, *i.e.* the regression of our instrument on job seekers characteristics and caseworkers' decisions. The plotted t-stats are associated with the coefficients of the independent variables reported in Table D1. The first boxplot (Bivariate) reports the distribution of the t-stats obtained through bivariate models, where the caseworker stringency instrument is regressed on each job seekers characteristic/caseworkers' decision separately. The second boxplot (Multivariate) reports the distribution of t-stats associated with partial correlation coefficients obtained from the regression of our instrument on all job seekers characteristics and caseworkers' decisions altogether. The third boxplot (Multivariate + PES FE) reports the distribution of the t-stats obtained on additionally conditioning on PES fixed effects.

Figure D3: Monotonicity assumption, first-stage estimates for different subsamples

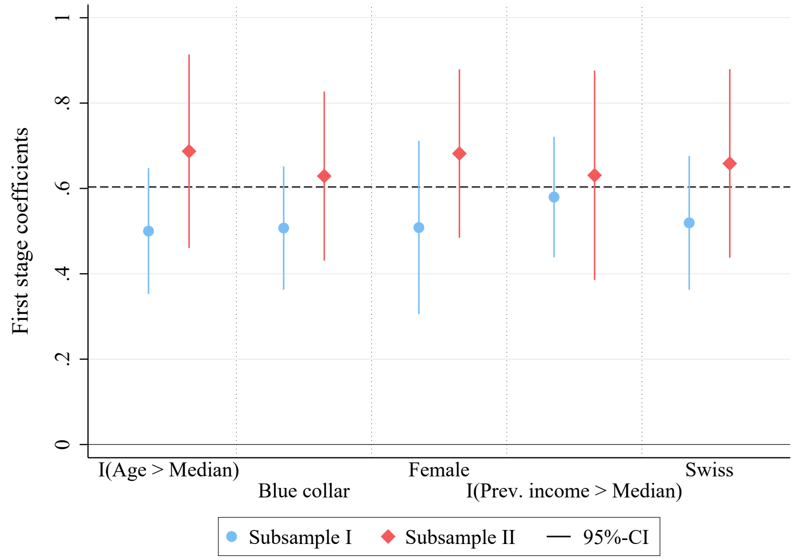

Note: This figure reports the estimated coefficients associated with caseworker stringency in the first stage equation (without individual fixed effects), across various subsamples. The sample-split variables used to define the subsamples are reported on the abscissa. The pooled-sample point estimate is reported as a dashed line on the graph. 95% confidence intervals based on clustered standard errors at the individual level are reported.

Figure D4: Monotonicity assumption
First-stage estimates for combination of job seekers' characteristics

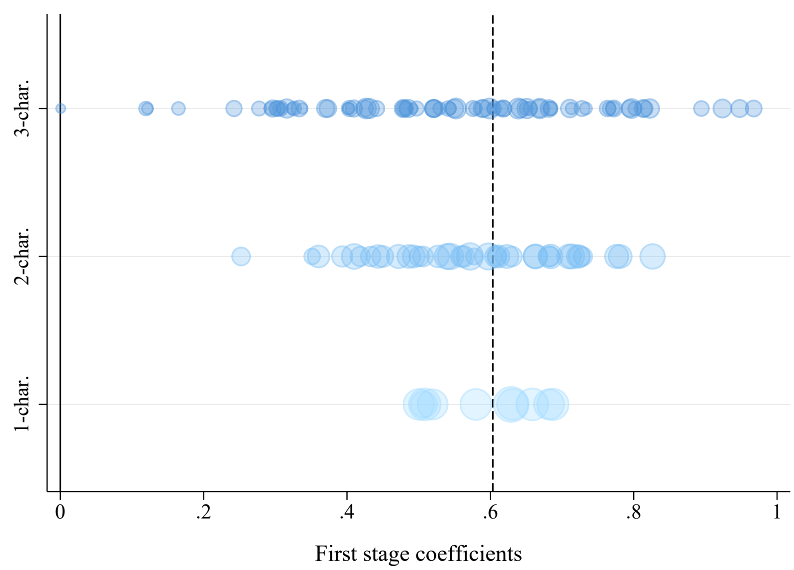

Note: This figure reports the distributions of the estimated coefficients associated with caseworker stringency in the first stage equation (without individual fixed effects), across various subsamples. Each subsample is based on the combination of either one, two or three (binary) job seekers characteristics (age, blue collar, female, previous income, Swiss). Markers size represent the subsamples size. The pooled-sample point estimate is reported as a dashed line on the graph.

Table D1: Conditional Independence Assumption (CIA) and Exclusion Restriction (ER)
Correlations between caseworker stringency instrument and job seekers' characteristics/caseworkers' decisions

	CIA		ER		CIA + ER		CIA + ER + Referral policy	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable: Caseworker stringency								
Job seekers' characteristics								
Age	0.0001*	0.0000			0.0002**	0.0000	0.0002**	0.0000
	(0.0001)	(0.0001)			(0.0001)	(0.0001)	(0.0001)	(0.0001)
Female	0.0023	0.0006			0.0024	0.0006	0.0024	0.0007
	(0.0018)	(0.0017)			(0.0018)	(0.0017)	(0.0018)	(0.0017)
Married	0.0026	0.0014			0.0026	0.0014	0.0027	0.0014
	(0.0019)	(0.0018)			(0.0019)	(0.0018)	(0.0019)	(0.0018)
Swiss	0.0014	-0.0023			0.0016	-0.0024	0.0017	-0.0022
	(0.0020)	(0.0018)			(0.0020)	(0.0018)	(0.0020)	(0.0018)
Primary education	-0.0050**	0.0007			-0.0050**	0.0006	-0.0048**	0.0007
	(0.0023)	(0.0021)			(0.0023)	(0.0021)	(0.0023)	(0.0021)
Blue collar worker	-0.0010	0.0022			-0.0014	0.0021	-0.0036	0.0003
	(0.0085)	(0.0084)			(0.0085)	(0.0084)	(0.0084)	(0.0083)
White collar worker	-0.0071	-0.0009			-0.0075	-0.0010	-0.0094	-0.0026
	(0.0086)	(0.0084)			(0.0086)	(0.0084)	(0.0085)	(0.0083)
ln(Previous income)	0.0008	0.0005			0.0007	0.0005	0.0007	0.0005
	(0.0005)	(0.0005)			(0.0005)	(0.0005)	(0.0005)	(0.0005)
Unemployment history	-0.0085	-0.0057			-0.0088	-0.0052	-0.0111*	-0.0070
	(0.0066)	(0.0061)			(0.0066)	(0.0061)	(0.0066)	(0.0061)
Caseworkers' decisions								
Individual program			-0.0107*	0.0006	-0.0111*	0.0009	-0.0121**	-0.0000
			(0.0060)	(0.0055)	(0.0060)	(0.0055)	(0.0059)	(0.0055)
Collective program			0.0001	0.0045**	0.0000	0.0043*	-0.0005	0.0038*
			(0.0025)	(0.0023)	(0.0025)	(0.0023)	(0.0024)	(0.0023)
Sanction			-0.0050	-0.0015	-0.0033	-0.0018	-0.0028	-0.0013
			(0.0067)	(0.0062)	(0.0068)	(0.0062)	(0.0067)	(0.0061)
Minimal search requirement			0.0043**	-0.0009	0.0050***	-0.0011	0.0070***	0.0008
			(0.0018)	(0.0017)	(0.0018)	(0.0017)	(0.0018)	(0.0017)
Referral policy								
Caseworker referrals							0.0187***	0.0143***
							(0.0015)	(0.0014)
PES FE	No	Yes	No	Yes	No	Yes	No	Yes
adj.-R^2	0.001	0.157	0.001	0.157	0.002	0.158	0.016	0.165
N. Individuals	15616	15616	15616	15616	15616	15616	15616	15616

Note: This table presents evidence on the CIA and ER. It reports the estimation results of the regression of our caseworker stringency instrument on job seekers characteristics and caseworkers' decisions. Estimates are based on observations aggregated to the individual level, so as to put equivalent weights on all job seekers. Robust standard errors are reported. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table D2: CwR effects on individual-monthly outcomes, count data models

	Poisson (1)	Poisson FE (2)	IV Poisson (3)
A. Dependent variable: # applications A_{it}			
# Caseworker referrals	0.045***	0.036***	0.061*
	(0.003)	(0.002)	(0.032)
# identifying observations	66882	64026	66882
# singleton observation	0	2856	0
# total observations	66882	66882	66882
B. Dependent variable: # non-referred applications A_{it}^P			
# Caseworker referrals	-0.052***	-0.056***	-0.027
	(0.005)	(0.004)	(0.048)
# identifying observations	66882	63972	66882
# singleton observation	0	2910	0
# total observations	66882	66882	66882
C. Dependent variable: # interviews I_{it}			
# Caseworker referrals	0.102***	0.155***	0.304***
	(0.009)	(0.012)	(0.065)
# identifying observations	66882	39157	66882
# singleton observation	0	27718	0
# total observations	66882	66882	66882
D. Dependent variable: # job offers O_{it}			
# Caseworker referrals	0.111***	0.183***	0.258**
	(0.014)	(0.026)	(0.112)
# identifying observations	66882	17151	66882
# singleton observation	0	49695	0
# total observations	66882	66882	66882
Individual controls	Yes	Yes	Yes
Unemployment duration FE	No	Yes	No
Individual FE	No	Yes	No
PES $\times$ Quarter FE	Yes	Yes	Yes

Note: This table reports the empirical estimates of equation (6.2), based on count data (Poisson) models. Panel A considers the monthly total application effort A_{it} as dependent variable, panel B the monthly private application effort A_{it}^P, panel C the monthly number of job interviews I_{it} and panel D the monthly number of job offers O_{it}. Column (1) reports estimates from a standard Poisson model, column (2) fixed effects Poisson estimates, while column (3) reports results from a Poisson model that accounts for the endogeneity of the treatment variable, using the two-steps GMM estimator. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table D3: CwR effects on application effort, log-transformation

	OLS	OLS FE	2SLS	2SLS FE
	(1)	(2)	(3)	(4)
A. Dependent variable: log-# applications $\ln(A_{it})$				
# Caseworker referrals	0.047***	0.039***	0.108**	0.056**
	(0.003)	(0.002)	(0.044)	(0.026)
Individual controls	Yes	Yes	Yes	Yes
Unemployment duration FE	No	Yes	No	Yes
Individual FE	No	Yes	No	Yes
PES $\times$ Quarter FE	Yes	Yes	Yes	Yes
# individuals	15616	15616	15616	15616
# identifying observations	66882	64026	66882	64026
# singleton observation	0	2856	0	2856
# total observations	66882	66882	66882	66882
B. Dependent variable: log-# non-ref. applications $\ln(A_{it}^P)$				
# Caseworker referrals	-0.047***	-0.065***	0.057	-0.054*
	(0.004)	(0.004)	(0.063)	(0.032)
Individual controls	Yes	Yes	Yes	Yes
Unemployment duration FE	No	Yes	No	Yes
Individual FE	No	Yes	No	Yes
PES $\times$ Quarter FE	Yes	Yes	Yes	Yes
# individuals	15616	15616	15616	15616
# identifying observations	66620	63758	66620	63758
# singleton observation	0	2862	0	2862
# total observations	66620	66620	66620	66620

Note: This table reports the empirical estimates of equation (6.2), for the application effort measures. The natural logarithms of A_{it} and A_{it}^P are used as dependent variables. Panel A reports results for the monthly total application effort A_{it}, while panel B relates to the monthly private application effort A_{it}^P. Columns (1) and (2) report OLS estimates, with or without individual fixed effects. Columns (3) and (4) report similar 2SLS estimates, using caseworker stringency as an instrument. Standard errors are clustered at the individual level. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Chapter 3

Progressive Taxation, Commuting Costs and Residential Decisions in Fiscally Decentralized Cities

Nicola Mauri
University of Lausanne

Jeremy Zuchuat
University of Lausanne

Abstract

Decentralized taxation and commuting affect numerous dimensions of economic activity, but how they shape jointly residential decisions and spatial income sorting is still not well understood. We develop a structural model of individual location choices in a fiscally decentralized monocentric city with explicit commuting costs. Spatial income sorting arises from decentralized progressive income taxes and non-homothetic housing preferences. Progressive taxation impacts high-income individuals' utility more markedly, hence reducing their commuting disutility. This induces high incomes to be more willing to commute to low-tax jurisdictions. We estimate our model using micro-level data on moving decisions in Switzerland. Based on a random utility framework, we find that richer individuals are more sensitive to taxes and less sensitive to commuting. Further, our results show that lower local taxes directly increase individuals' willingness to commute. In a counterfactual exercise, we show that uniform taxation within urban areas would lead to more high-income individuals locating in central jurisdictions, hence enduring shorter commutes.

Keywords: income sorting; commuting; fiscal decentralization; tax differentials; tax progressivity; non-homothetic preferences; regional science.

JEL: H71, H73, R23, R41

1. Introduction

Most cities are composed of multiple municipalities which enjoy some degree of autonomy over their taxing power. This has important consequences for the economic activity in urban areas (Brülhart et al., 2015). In particular, local taxes play an important role in individuals' residential decisions within fiscally decentralized economies. Progressive income taxation coupled with the decentralization of taxing powers have been shown to be important drivers of income sorting: households with similar income levels tend to live in jurisdictions which offer similar local tax rates (Schmidheiny, 2006; Schaltegger et al., 2011; Basten et al., 2017; Brülhart et al., 2021).

However, existing studies have generally omitted to explicitly account for the spatial structure of the economies they analyze and the induced commuting patterns. In fact, the historical development of transport infrastructure and concentration of job opportunities in urban centers have led to a dissociation between work and residence places (Heblich et al., 2020; Fretz et al., 2022). This implies that residential decisions inherently translate into commuting behaviors, entailing monetary or psychological costs (Brownstone and Small, 2005; Roberts et al., 2011). Ignoring the role played by commuting in those decisions possibly means omitting a key location factor to understand the spatial distribution of income in fiscally decentralized areas.

We address this gap in the literature by investigating how local taxes affect residential decisions and the willingness to commute of households with different income levels. We first set up a structural location choice model with non-homothetic preferences for housing and decentralized progressive income taxation. In our setup, individuals choose a residential location within a monocentric city, where everyone works in the central business district (CBD) and earns an exogenous wage. Jurisdictions are characterized by their commuting distance to the CBD, local tax rates and level of housing prices. Given that these three features are undesirable, we show that individuals trade off these characteristics and that these trade-offs differ according to individuals' income level. On the one hand, taxes make up a larger share of total expenditures of high-income earners in a progressive tax system. On the other hand, non-homothetic preferences for housing make higher incomes spend less on housing, as a proportion of their total expenditure. These two elements generate income sorting across jurisdictions, as people of different income levels value jurisdictional characteristics differently. Our model shows that high-income earners are less sensitive to housing prices compared to both taxes and commuting. They are hence predicted to locate in jurisdictions with higher housing prices. Whether high-income individuals are more sensitive to taxes compared to commuting (or conversely) remains ambiguous. We show that this ambiguity comes from the progressivity of income taxation and high-income individuals' larger opportunity cost of commuting. If the progressivity of the tax schedule is sufficiently marked, high incomes end up being more sensitive to taxes relative to commuting (and to housing prices in any case), and choose to locate in low-tax jurisdictions, at the cost of longer commutes. Intuitively, the progressivity of the income tax schedule reduces the opportunity cost of commuting of higher incomes to a greater extent, making them more willing to commute if they get compensated with lower taxes.

We exploit administrative fragmentation and fiscal decentralization in Switzerland to assess empirically how local taxes, commuting and housing prices impact individual location decisions along the income distribution. Using individual-level data, we study residential decisions of in-moving households working in the CBDs of four of the largest Swiss urban areas. This analysis is conducted conditional on the existing spatial equilibrium, which is characterized by low-tax suburbs. From the point of view of atomistic in-moving households, aggregate housing prices, local tax rates and distance to the CBD are taken exogenously. Based on a random utility framework, we estimate an empirical equivalent to our structural location choice model, where we interact municipal characteristics (taxes, housing prices and commuting measures) with individual's household income. This econometric set-up allows us to determine how agents of different income levels value the jurisdictional characteristics of interest and reveals how those might influence the spatial income sorting of Swiss residents.

Our estimation results show that high-income agents are indeed more sensitive to taxes compared to low incomes. This confirms local taxation as an important driver of income sorting in fiscally decentralized economies. We also observe that high-income earners tend to sort themselves into municipalities where housing prices are high. Finally, centrality is valued less strongly by higher incomes: our estimates show that well-off individuals are willing to endure longer commutes to work. To test formally whether this higher willingness to commute results from tax differentials, we additionally interact commuting distance with taxes and individual's income. Those additional results show that lower taxation levels reduce directly individuals' disutility from commuting, consistently with our theoretical predictions.

We assess the role of fiscal decentralization in the urban spatial distribution of income by conducting an illustrative counterfactual exercise. Conditioning on the existing spatial equilibrium, we impose a homogeneous progressive tax schedule within the urban areas of study and compute individual location probabilities in the absence of municipal tax differentials. Our results show that shutting off fiscal decentralization leads to an increase in demand for central locations from high-income individuals, who cannot self select into low-tax suburbs any more.

The main innovations of our study are threefold. First, we complement existing studies on residential location decisions in fiscally decentralized economies by explicitly accounting for commuting as a feature of the alternatives in individuals' choice set. Most related literature has put aside or not modeled formally the commuting cost entailed when location decisions are made (Schmidheiny, 2006; Basten et al., 2017; Brülhart et al., 2021). Acknowledging that a large majority of people live and work in different places, we add a monocentric structure to the canonical model of location decisions. We also specify a commuting disutility similar to Ahlfeldt et al. (2015) or Heblich et al. (2020), so as to derive the trade-offs made by individuals between taxes, commuting and housing prices in fiscally decentralized environments. This addition allows us to predict the residence places of agents with different incomes and to investigate the spatial consequences of fiscal federalism on income sorting in monocentric urban areas.

Second, we are the first to analyse jointly local income taxation and commuting in a set-up with a large number of small jurisdictions. While Agrawal and Hoyt (2018) look at metropolitan areas which sprawl over two states in the U.S., no previous research has studied the relationship between commuting and local taxes in places which present a high degree of spatial administrative fragmentation. At the aggregate level, we document a so-called "tax-gradient", *i.e.* the relationship between distance to the city center and local taxation. This one is featured with lower taxes in the suburbs compared to central municipalities. Just like the widely studied housing price gradient in urban economics (for a recent overview, see Duranton and Puga, 2015), it may have important implications for residential decisions of individuals, who are able to trade off commuting with taxes.[1] Exploiting the high fragmentation of the Swiss tax system within a small geographical space, we are hence able to study the role of taxes and commuting in individual decisions in a finer manner than studies focusing on the sub-federal level for instance.

Third, our study is the first to explicitly relate tax progressivity and decentralization to the opportunity cost of commuting. Accounting explicitly for the necessity of commuting to work, we show theoretically that income tax progressivity combined with decentralization can affect the willingness to commute of high incomes, in a positive manner. Our counterfactual exercise hence echoes recent findings by Su (2022), who shows that gentrification of city centers can be imputed to the rising opportunity cost of time. Given the existing decentralized equilibrium with lower taxes in the suburbs, the removal of fiscal decentralization increases high incomes' opportunity cost of commuting, as their longer commutes are not compensated through lower taxes any more. This modification of the fiscal system might eventually lead to more gentrified central municipalities than under a fiscally decentralized economy. This finding helps us to make sense of a pattern that is observed in our data and that was suggested in previous studies (Schmidheiny, 2006; Fretz et al., 2022): high-income individuals commute longer distances and time than lower incomes in Switzerland.

[1]Note that in this study, our formulation of the "tax-gradient" diverges from that of Agrawal (2015), which studies state-borders discontinuities and their effect on local sales taxes. The gradient is here making allusion to the distance to the CBD and not to the state-border.

Our study is part of two main strands of literature. First, we address the literature on the role of taxation in individuals' location decisions, and more specifically studies that investigate the role of local income taxation and tax progressivity in spatial income sorting. The seminal study of Schmidheiny (2006) was the first to show theoretically and empirically that the progressivity of local income taxes is a crucial factor to explain the observed income segregation patterns in fiscally decentralized areas. This result is confirmed by Schaltegger et al. (2011), who look at income sorting patterns in the agglomeration of Zürich using a municipal-level IV approach to instrument local taxes. Additional causal evidence of the impact of taxes on residential choices is provided by Schmidheiny and Slotwinski (2018), who make use of an institutional rule inducing a large variation in taxes within the choice-set of foreign employees in Switzerland. Another recent study by Basten et al. (2017), which exploits precise micro-data and institutional border discontinuities, provides causal estimates of the capitalization effect of local income taxes into housing prices and new evidence of tax-based income sorting in Switzerland. Adopting a more structural approach to study the welfare implications of local taxes, Brülhart et al. (2021) provide reduced-form IV estimates for the elasticity of the count of tax-payers residing in a municipality with respect to local taxes. The authors show a negative causal relationship between changes in local tax-rates and changes in tax-payer counts, the magnitude of the effect increasing with income classes.

Second, we add up to the small literature that addresses the topics of local taxes and commuting in a joint framework. Theoretically, De Bartolome and Ross (2003) extend the canonical monocentric model of Alonso (1964) and Muth (1969) by considering a fiscally decentralized environment. The authors use their stylized model to prove the existence of perfect and imperfect income sorting equilibria, where local public good works as a compensating factor for commuting distance. More recently, Agrawal and Hoyt (2018) investigate theoretically and empirically how local taxation impacts commuting time. In their theoretical model, the authors consider the context of metropolitan statistical areas (MSA) divided by state borders to examine the impact of tax discontinuities on employment and residential location. They differentiate between residence-based taxation (reciprocity) and both residence and employment-based tax schemes (no reciprocity). In their empirical application, the authors make use of border discontinuities within a large number of MSA in the U.S. to estimate the impact of tax differentials on commuting. They report evidence that, under the reciprocity regime, an increase in the tax differential leads to a change in residential mobility, resulting in more mobility from the low-tax side of the MSA.

The work of Agrawal and Hoyt (2018) is the closest to our study, but both differ in several important dimensions. From a theoretical perspective, our paper focuses on the equivalent of a reciprocal tax scheme, not in a two-sided MSA, but in a highly fragmented fiscal environment with a larger choice set. As opposed to the general equilibrium model developed in Agrawal and Hoyt (2018), our structural model of individual location decisions takes the spatial equilibrium as given. In contrast to their stylized model, ours explicitly accounts for income sorting mechanisms that have been highlighted in the literature: non-homothetic housing consumption and decentralized progressive taxation. On the empirical side, the high variability of local taxes in Switzerland represents a particularly interesting case to study the impact of income tax differentials on optimal commuting behavior. Larger individual choice sets enable us to investigate how local taxes impact location and commuting decisions more finely, hence providing better understanding of how those differ with respect to income level. All in all, our results are consistent with the findings of Agrawal and Hoyt (2018), who observe an increase in the commuting flows from the low-tax states, particularly marked for high-income individuals.

The rest of this paper is organized as follows. In section 2, we introduce the urban and institutional context of our research before turning to the structural location choice model in section 3. In section 4 we present our data and descriptive statistics. Section 5 describes the empirical approach we follow to assess the predictions of our theoretical model and presents our empirical findings. In section 6, based on our empirical estimates, we conduct a counterfactual exercise to illustrate the role played by fiscal decentralization in income sorting and commuting decisions. Section 7 concludes.

2. Institutional and urban setting

Switzerland is a multi-level fiscal federation divided in 26 states called cantons, which in turn are divided in around 2,500 municipalities.[2] Each administrative level collects its own taxes and possesses a certain degree of tax autonomy. The 26 cantons have their own tax laws and decide autonomously on their tax schedules, rates and progressivity. As there exists an almost perfect overlap of tax bases at the cantonal and municipal levels, municipalities choose one or multiple tax multipliers that shift the cantonal tax schedules, hence generating variability in municipal-level taxes. At the individual level, Swiss residents are liable to paying municipal, cantonal and federal taxes, both on their income and wealth.[3] For most taxpayers, income taxes represent the largest share of their annual tax bill.[4] Personal taxation is progressive, and based on households' residence place and characteristics. Tax rates hence vary substantially

[2]The exact number of municipalities depends on the year of observation, due to the ongoing municipal merger process across the country.

[3]For foreigners, only people holding a C-permit (permanent residency) are subject to the so-called "ordinary taxation". Others are subject to a withholding tax, which does not vary within cantons. For that reason, foreigners with other permits than C are not considered in this study.

[4]As an example, the median tax payer in Switzerland (single) with an annual income of CHF 80,0000.- and a wealth below CHF 100,000.- would pay approximately CHF 10,000.- of income tax and less than CHF 100.- of wealth tax annually, in 2015.

Figure 2.1: Urban areas of study

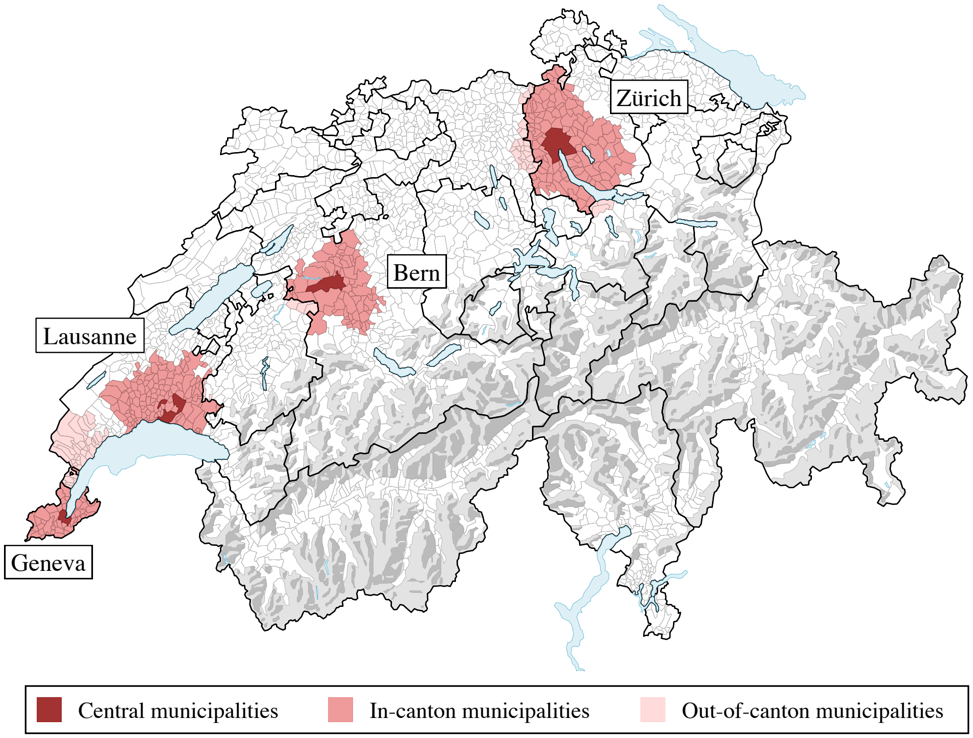

Note: This figure displays the four urban areas of study. For each agglomeration, it distinguishes between the central municipality, the municipalities located in the same canton as the central municipality and municipalities which are located outside this canton.

across and within cantonal borders, as well as along the household income distribution.[5] Half of the Swiss population resides in the five largest agglomerations of the country: Zürich, Geneva, Basel, Bern and Lausanne. Those represent more than 50% of total employment in Switzerland (OFS, 2019). With the exception of the city of Basel, which we do not consider in this study, Swiss cities have developed circularly around a single central municipality and mostly within a single canton.[6] Agglomerations of study are depicted in Figure 2.1. They exhibit a typical monocentric structure: within a radius of 30 km around the city center, approximately half of the residing population works in the central municipalities (see Figure A1 in the Appendix). Within agglomerations, approximately 50% of the people working in the

[5]Figures A3 to A5 in the Appendix exhibit the consolidated tax burden faced respectively by a single tax payer, a household with no children and a household with 2 children, within each of the agglomerations of study. The tax burden measure refers to the consolidated federal, cantonal and municipal tax bill in CHF divided by gross income.

[6]The high degree of local fiscal autonomy and the heterogeneity in cantonal-level tax schedules makes intercantonal comparison rather tedious. For that reason, we will focus on intra-cantonal tax variation in this study, and consequently omit the agglomeration of Basel, due to its low intra-cantonal administrative fragmentation (only three municipalities).

urban center reside in the central jurisdiction, while the remaining half lives in the rest of the agglomeration (see Figure A2 in the Appendix).[7]

To give a sense of the spatial distribution of taxes within our four sample urban areas, we report standardized tax multipliers and municipal housing price indices as a function of distance to the central municipalities, in Figure 2.2. These figures highlight the variability of taxes and housing prices within a small geographical area: for a given distance to the productive centers, one can face substantial differences in taxes and housing prices. In urban centers, tax rates and housing prices levels are typically high. As we move away from central municipalities, we observe a steady decline in the level of housing prices. This matches the classic bid-rent theory from Alonso (1964) or Muth (1969) and prior empirical evidence of negative housing price gradients (Duranton and Puga, 2015). In addition, moving slightly outside of city-centers entails a sharp drop in local tax rates, which increase again slightly as we move further outside the close ring of municipalities around the centers. Nevertheless, taxes in remote suburbs remain lower than in central municipalities on average.[8] The empirical spatial distributions of taxes and housing prices point towards the existence of a spatial equilibrium

Figure 2.2: Taxes, housing prices and distance to agglomerations centers

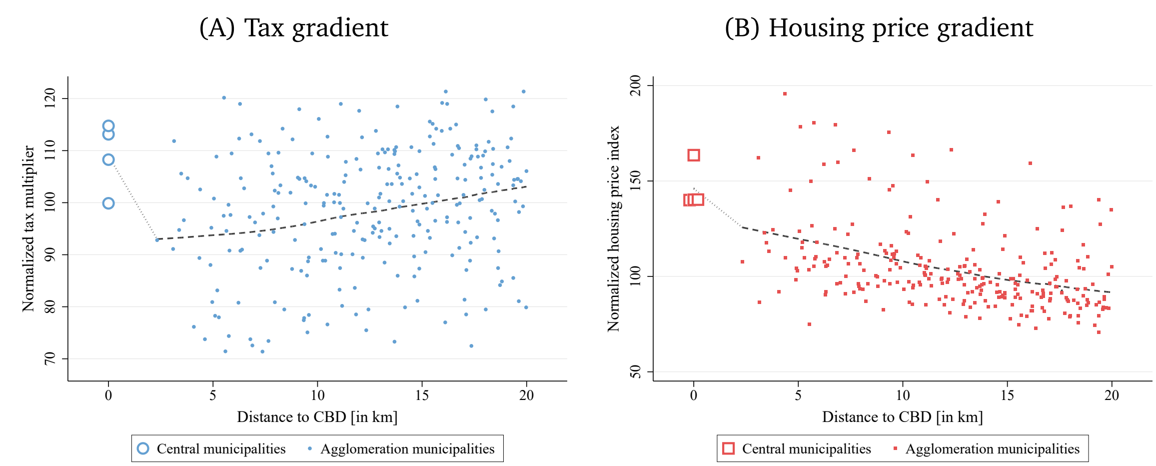

Note: This figure shows how municipal tax multipliers and housing price indices change with distance to the central municipality for the agglomerations of Zürich, Bern, Lausanne and Geneva. To make variables comparable across agglomerations, we standardize them by dividing the local multiplier or housing price index by its within-urban area average.

[7]We borrow the definition of central municipalities from the Swiss Statistical Office. A central municipality is coined as such based on its density of population, employment and hotel stays. The rest of the agglomerations made of peripheral or "suburban" municipalities are defined based on mobility patterns of commuters (OFS, 2019).

[8]All those points are represented visually in Figure C1, where we illustrate how taxes, housing prices and distance to the CBD relate to each other, in a single graph. The graph shows that central municipalities are characterized by higher taxes and housing prices. Suburban jurisdictions exhibiting low taxes and higher housing prices tend to be closer to the CBD. On the contrary, high-tax and low-housing prices jurisdictions tend to be more distant.

where taxes (and housing prices) are lower in non-central municipalities. Such equilibrium might arise in a fiscally decentralized monocentric environment if taxes and housing prices act as compensating differentials for longer commutes, as suggested in De Bartolome and Ross (2003).

At the individual level, the shape of this "tax gradient" can generate an additional arbitrage between commuting and taxes when choosing where to locate in the agglomeration, on top of the usual trade-off between commuting and housing prices. As an example, a single individual earning a yearly gross income of CHF 100,000.- in the municipality of Zürich will save approximately CHF 2,100.- in income taxes by deciding to live in the neighboring municipality of Kilchberg, located 7.5 km away from Zürich center. This is equivalent to a 19% decrease in her yearly tax bill, at the cost of a 12-minute commute with the regional train or a 15-minute car ride. In the next section, we develop a structural model of individual location decisions to formally characterize these trade-offs and to study how those change with income level.

3. A model of residential location decisions in a fiscally decentralized and monocentric urban area

3.1 Model setup

We develop a model of individual residential location decisions within a monocentric urban area, where all individuals work in the central business district (CBD). We consider an urban area that is composed of J municipalities, including the central one where the CBD is located. The J municipalities are heterogeneous in various dimensions: they differ in terms of distance to the central municipality δ_j, local tax rate τ_j, housing prices Q_j and additional exogenous amenities B_j (topographical amenities, historical monuments, etc.). The urban area is populated by atomistic individuals who choose their residence municipality j. Individual i derives utility from a numeraire C_i, housing floor space H_i and a publicly provided good G.[9] Agent i's private consumption is financed through her wage W_i, which she earns at the CBD, and which is assumed to be set exogenously.

When making residential location decisions, atomistic agents take municipal characterisics as given. Individual i's decision to reside in municipality j impacts her utility in four different manners. First, she is subject to a residence-based income tax $\tau_j r(W_i)$. The function $r(W_i)$ corresponds to the supra-municipal (cantonal-level) progressive tax schedule, with $r'(W_i) > 0$ and $r''(W_i) < 0$, whereas τ_j represents the municipality-specific tax multiplier that shifts the

[9]Following Schmidheiny (2006), we assume that the level of government spending does not vary across municipalities. G can therefore be interpreted as the targeted minimal level of public spending. This choice is also motivated by our empirical study context, where government spending is homogenized across municipalities through fiscal equalization.

cantonal schedule.[10] Second, commuting costs depend on the distance between worker's place of residence and the CBD, δ_j. Third, the local level of housing prices Q_j affects individual's budget constraint through housing consumption. Finally, local exogenous amenities B_j may directly affect the utility of the individual.

We assume individual i's utility function to be Stone-Geary, which implies non-homothetic preferences for housing (Schmidheiny, 2006; Brülhart et al., 2021). As for commuting, we follow the recent literature and specify exponential commuting utility costs $\theta_j = e^{\kappa\delta_j}$ (Ahlfeldt et al., 2015; Heblich et al., 2020).[11] Formally, the consumer problem writes

$$\max_{C_i, H_i} \quad U_{ij} = \frac{z_{ij} B_j}{\theta_j} \Big(C_i\Big)^{1-\beta} \Big(H_i - \upsilon_H\Big)^{\beta} \Big(G\Big)^{\gamma} \tag{3.1}$$
$$\text{s.t. } C_i + Q_j H_i = \big[1 - \tau_j r(W_i)\big] W_i,$$

where z_{ij} is the idiosyncratic utility taste of individual i for municipality j, while β and γ are preference parameters, and υ_H a Stone-Geary parameter. Based on the consumer's program, we derive Marshallian demand functions for the numeraire C_i^* and housing H_i^*, which we plug back into the objective function to obtain the indirect utility function

$$\begin{aligned} V_{i|j} = \Omega + \ln(z_{ij}) + \ln(B_j) - \kappa\delta_j + \ln\left(\big[1 - \tau_j r(W_i)\big] W_i - \upsilon_H Q_j\right) \\ - \beta \ln(Q_j) + \gamma \ln(G), \end{aligned} \tag{3.2}$$

where $\Omega = (1-\beta)\ln(1-\beta) + \beta\ln(\beta)$ is a constant term.[12]

3.2 Non-homothetic preferences and income sorting

We next study the trade-offs made by individuals between the municipal characteristics of interest, τ_j, δ_j and Q_j. As a matter of simplification, we set $B_j = B$ for all municipalities, where B can be interpreted as agglomeration-level amenities to which all individuals have access. Such hypothesis seems reasonable in our context, given that municipalities are all located in the same agglomeration, which expands over a limited geographical area. Municipalities hence presumably offer relatively homogeneous amenities B_j within agglomerations,

[10] This formulation of the progressive federalist tax schedule is inspired by Schmidheiny (2006).

[11] Note that the choice of considering commuting utility costs rather than monetary ones is also motivated by empirical observations. It can be seen in Figure B1 that longer-distance commuters to agglomeration centers in Switzerland work just as much and commute almost as often as shorter-commute individuals. It therefore appears reasonable to assume that commuting is essentially a disutility cost in our context, and that it does not reduce income. As a matter of completeness, we derive the model with commuting costs specified as monetary costs in Section B and show that our results are qualitatively unchanged.

[12] Formally, Marshallian demand functions are given by $H^* = \frac{\beta\left[(1-\tau_j r(W))W - Q_j \upsilon_H\right]}{Q_j} + \upsilon_H$ and $\mathrm{C}^* = (1-\beta)\big[1 - \tau_j r(W)\big]W - (1-\beta)Q_j\upsilon_H$.

the latter playing a minor role in individual location decisions.[13]
We compute the marginal rates of substitution between municipal characteristics τ_j, δ_j and Q_j by totally differentiating equation (3.2), holding utility $V_{i|j} = V$ constant:

$$MRS_{Q,\tau} = \frac{dQ_j}{d\tau_j}\Big|_{dV=0} = -\frac{\partial V/\partial \tau_j}{\partial V/\partial Q_j} = -\frac{r(W_i)W_i}{H_i^*} < 0 \tag{3.3a}$$

$$MRS_{Q,\delta} = \frac{dQ_j}{d\delta_j}\Big|_{dV=0} = -\frac{\partial V/\partial \delta_j}{\partial V/\partial Q_j} = -\frac{\kappa\big(\big[1-\tau_j r(W_i)\big]W - Q_j\upsilon_H\big)}{H_i^*} < 0 \tag{3.3b}$$

$$MRS_{\tau,\delta} = \frac{d\tau_j}{d\delta_j}\Big|_{dV=0} = -\frac{\partial V/\partial \delta_j}{\partial V/\partial \tau_j} = -\frac{\kappa\big(\big[1-\tau_j r(W_i)\big]W_i - Q_j\upsilon_H\big)}{r(W_i)W_i} < 0 \tag{3.3c}$$

where $H_i^* = \frac{\beta\left[(1-\tau_j r(W_i))W_i - Q_j\upsilon_H\right]}{Q_j} + \upsilon_H$ is the optimal housing quantity. Those three equations give rise to Proposition 1.

Proposition 1: Individual trade-offs between municipal characteristics

At the individual level, the three MRS between the local level of taxation τ_j, commuting distance δ_j and housing prices Q_j are all negative.

Proof. Proposition 1 directly follows from equations equations (3.3a) to (3.3c). □

This first result is straightforward: given that the three jurisdictional characteristics τ_j, δ_j and Q_j are all negatively valued by individuals, *i.e.* $\frac{\partial V}{\partial \tau_j} < 0$, $\frac{\partial V}{\partial \delta_j} < 0$ and $\frac{\partial V}{\partial Q_j} < 0$, an increase in one of those must necessarily be compensated by a decrease in another.
Given the mix of jurisdictional characteristics, households choose the residential location k that yields the highest utility level:

$$V(\tau_k, Q_k, \delta_k; W_i) \geq V(\tau_l, Q_l, \delta_l; W_i) \quad \forall k \neq l. \tag{3.4}$$

Empirically, all municipalities are populated. We assume the same in our theoretical model. Consequent to individuals valuing negatively the three jurisdictional characteristics, it must be that there exists some compensation between those, to rule out empty municipalities. This

[13] Our assumption of homogeneous amenities should not only hold for reasons related to the small geographical scale of the urban areas of study. It is also reasonable given the centralized organization of agglomerations and planification of their development at the supra-municipal level, which prevent the emergence of degenerated places. For instance, all municipalities must offer public primary schooling at the local level. This factor, which has been shown to be an important source of location decisions in the U.S. context, should play a minor role in the Swiss context. Exogenous amenities, but also possibly endogenous ones (which are taken as given by individuals when choosing their residence place), are thus presumably homogeneous within agglomerations. On the empirical side, we further address this simplifying assumption in our econometric analyses by controlling for different observable municipal amenities.

point is formalized in the next lemma.[14]

Lemma 1: Compensating municipal characteristics

If every location j is populated, then at least one municipal characteristic must compensate for the two others. A high-tax jurisdiction must exhibit either low housing prices and/or low commuting distance. Jurisdictions with high housing prices must have either low taxes and/or low commuting distances. Jurisdictions located far from the city-center must exhibit low taxes and/or low housing prices.

Proof. The proof is by contradiction. If a jurisdiction exhibits lower level of taxes, housing prices and commuting distance than other jurisdictions, then all households will move to this jurisdiction. This contradicts the assumption that every location is inhabited. □

Next, we study how preferences towards municipal characteristics vary with income level W_i. In that respect, we take the derivative of the MRS with respect to W_i:

$$\frac{\partial MRS_{Q,\tau}}{\partial W_i} = -\frac{[(1-\beta)Q_j\upsilon_H[r(W_i) + r'(W_i)W_i] + \beta r'(W_i)W_i^2]}{Q(H_i^*)^2} < 0 \tag{3.5a}$$

$$\frac{\partial MRS_{Q,\delta}}{\partial W_i} = -\frac{\kappa\big(1 - \tau_j\big[r(W_i) + r'(W_i)W_i\big]\big)\upsilon_H}{(H_i^*)^2} < 0. \tag{3.5b}$$

$$\frac{\partial MRS_{\tau,\delta}}{\partial W_i} = -\frac{\kappa\big[Q_j\upsilon_H[r(W_i) + r'(W_i)W_i] - W_i^2 r'(W_i)\big]}{W_i^2[r'(W_i)]^2} \lesseqgtr 0 \tag{3.5c}$$

The three equations above show that low and high-income individuals have different sensitivity towards municipal characteristics, *i.e.* the pairwise marginal rates of substitution between municipal characteristics are more or less steep depending on the income level W_i. Equations (3.5a) and (3.5b), whose directions are unambiguous, show that rich individuals are less sensitive to housing prices relative to both taxes and commuting. Intuitively, high-income earners are willing to pay fewer taxes (to avoid higher progressive tax rates) and to commute less (due to their higher opportunity cost of commuting), at the cost of higher housing prices. Equation (3.5c), for its part, shows that it is not clear how the trade-off between taxes and commuting distance varies with income.

Taken together, equations (3.5a) and (3.5b) show that high-income individuals are always less sensitive to housing prices. Without any further restriction, we cannot clearly state how preferences for taxes and commuting vary with income, due to the ambiguity of equation (3.5c).

[14]This assumption is presumably met in our empirical context, as shown in Figure C1.

This ambiguity is removed if the following condition holds:

$$0 < \frac{Q_j \upsilon_H}{W_i} < \frac{1}{1+\gamma(W_i)^{-1}} < 1, \tag{3.6}$$

where $\gamma(W_i) \equiv \frac{r'(W_i)W_i}{r(W_i)} > 0$ captures the degree of progressivity of the tax schedule. The above condition states that if the share of a marginal increase in the income devoted to taxes paid in municipality j is larger than the share of income devoted to minimal housing consumption ν_H valued at the local housing price Q_j, then the sign of equation (3.5c) is strictly positive. Consequently, high-income individuals will be more averse to taxes than to commuting. This condition is more likely to hold for tax schemes that are strongly progressive, *i.e.* for high $r'(W_i)$.

The predictions of our model on how the valuation of jurisdictional characteristics changes with income is summarized in Proposition 2.

Proposition 2: Individual sensitivity to municipal characteristics

High-income individuals are less sensitive to housing prices. If additionally equation (3.6) holds, high-income individuals are unambiguously more sensitive to taxes. In that case, how preferences for commuting change with income level is ambiguous.

Proof. The proof directly follows from equations (3.5a) to (3.5c) and equation (3.6). □

What are the consequence of our model for individual location decisions and spatial income sorting? As high-income individuals are less sensitive to housing prices, they are less reluctant to locate in municipalities with high level of housing prices. Similarly, if the condition described in equation (3.6) is met, rich individuals are more sensitive to taxes and prefer to locate in municipalities where taxes are low.[15] Finally, equation (3.6) does not remove the ambiguity towards commuting. We thus remain agnostic about how richer individuals value centrality.

[15] If the condition is not met, the prediction with respect to taxes is ambiguous. However, the higher sensitivity of higher income towards taxes has been shown empirically in different geographical and time contexts in Switzerland, notably in Schmidheiny (2006), Basten et al. (2017) or Brülhart et al. (2021). Based on these past findings, we expect richer individuals to be more sensitive towards taxes in spite of our ambiguous theoretical predictions. For that reason, assuming equation (3.6) seems reasonable in our context, so as to impose more structure on our predictions. However, it is worth mentioning that equation (3.6) is a *sufficient* but not *necessary* condition to state that high-income individuals are globally more sensitive to taxes. For different parameter values, this may occur even despite condition equation (3.6) not being met.

3.3 Fiscal decentralization and tax progressivity

The reason behind the ambiguity in the non-homotheticity in preferences for commuting is related to two counteracting forces in our model: non-homothetic preferences for housing and decentralized progressive taxation. In order to emphasize the role played by progressivity, we switch off tax progressivity by setting $r(W_i) = \bar{r}$, *i.e.* we consider a system of proportional taxation. In this new setting, the changes in relative preferences with respect to income are modified as follows:[16]

$$\frac{\partial MRS_{Q,\tau}}{\partial W_i} = -\frac{(1-\beta)Q_j \upsilon_H \bar{r}}{Q(H_i^*)^2} < 0 \tag{3.7a}$$

$$\frac{\partial MRS_{Q,\delta}}{\partial W_i} = -\frac{\kappa\big(1-\tau_j\bar{r}\big)\upsilon_H}{(H_i^*)^2} < 0. \tag{3.7b}$$

$$\frac{\partial MRS_{\tau,\delta}}{\partial W_i} = -\frac{\kappa Q_j \upsilon_H}{\bar{r}W_i^2} < 0 \tag{3.7c}$$

The sign of equation (3.7c), describing the non-homotheticity in the relative valuation of commuting and taxes, is now strictly negative. This affects the spatial sorting predictions of the model: higher-income individuals now have a higher willingness to pay for shorter commutes than for lower taxes. The intuition behind this result is the following: in the absence of tax progressivity, the gain from living in a low-tax jurisdiction is no longer over-proportionally larger for high-income earners. At the same time, the cost of living farther is still higher for high incomes and increases exponentially with distance. In this setup, richer individuals locate closer to the city center and have no particular preferences for low-tax municipalities.

What if instead of shutting-off progressive taxation, we remove income taxes as a municipal characteristic? This corresponds to a fiscally centralized case, where taxes are the same wherever individuals choose to locate. In this context, the residential decision of the individual boils down to a trade-off between commuting distance δ_j and housing prices Q_j. The corresponding MRS between the two jurisdictional characteristics is equivalent to equation (3.3b), only that the local multiplier loses its subscript j. Shutting off the decentralized structure of taxation then leads to the prediction that richer agents are more sensitive to commuting than to housing prices. Given our assumption that all municipalities are populated, we expect a fiscally centralized city structure where income is monotonically decreasing with distance to the CBD.

[16] Note that the signs of the marginal rate of substitutions between jurisdictional characteristics remain unchanged, as all three characteristics τ_j, δ_j and Q_j are still undesirable to the consumer.

Taken together, these results show that progressive taxation combined with fiscal decentralization are at the core of the ambiguous trade-off between commuting and taxes. By “hurting” high-income individuals disproportionately more, tax progressivity makes them potentially more sensitive to taxes, and reduces their opportunity cost of commuting. Given high incomes’ ability to self-select into low-tax jurisdictions, those end up willing to commute longer distances so as to reduce their tax bills. Fiscal decentralization, together with tax progressivity, thus affects the optimal commuting distances chosen by individuals, and reduces higher incomes’ disutility from commuting compared to the fiscally centralized case. In the remaining of the paper, we confront our theoretical model to the data, to assess our predictions on the trade-offs made by individuals when they make residential location decisions in fiscally decentralized cities.

4. Data and descriptive evidence

4.1 Data

Our empirical analysis relies on four main sources of data. First, we collect information on Swiss residents’ mobility through the *Structural Survey*, for the years 2010 to 2016. This annual cross-sectional survey is conducted by the *Swiss Federal Statistical Office* and covers approximately 200,000 different individuals every year (approximately 2.5% of the residing population). It contains detailed socio-demographic information on individuals who are interviewed, as well as some information on the members and composition of their households. The survey also contains information on residence municipality up to five years prior to the year of observation, which enables us to identify individuals who move and choose a new municipality of residence within our agglomerations of study. The *Structural Survey* additionally contains various measures of commuting such as information on individuals’ workplace, the exact commuting distance (accounting for the mean of transportation) and self-reported commuting time.

We complement this first dataset with information on individuals’ gross income. Those are obtained from the *Swiss Federal Compensation Office* and are matched by means of social security unique identifiers. This second dataset contains information on gross income sources and amounts that are liable to social security contributions, for all household members related to the individuals we observe in the *Structural Survey*.

Third, we compute individual-level measures of taxation for each municipality and year. These data are provided by the *Federal Tax Administration* under the form of a consolidated tax burden (encompassing municipal, cantonal and federal taxes), which is available for different generic tax payers. Specifically, this individual-specific measure is computed for all Swiss municipalities (of residence), for different family types (single, married without children, married with children) and for different income classes.[17] As the computation by the *Federal Tax Administration* takes into account any type of fiscal deductions, the tax burden measure represents a fairly good approximation of the annual tax bill paid by the households in their municipality of residence.

Finally, we complement household-level information with municipality-level variables. We collect information on municipal public finances from the cantonal statistical offices, information on housing vacancies, housing supply and land use from the *Swiss Federal Statistical Office*, data on housing prices from *Fahrländer Partner AG*, and additional geocoded data from the *Federal Office of Topography Swisstopo*.[18]

4.2 Sample of analysis

In order to empirically assess the income sorting predictions of our theoretical model, we impose several sample restrictions. First, we restrict our sample to individuals who recently moved to one of the urban areas of study and who actively made a location decision. For each in-moving individual, we observe precisely her characteristics and the municipal features she faces in the year when she chooses her residence place. In addition, we exclude people who settle outside the canton where the agglomeration center is located, to ensure comparability in the tax schedules and public services they face.[19] Consequent to this within-agglomeration restriction, differentials in tax burdens between two municipalities arise exclusively from differentials in municipal taxation, for a given income class.

Second, in line with our theoretical framework, we study exclusively individuals who work in the central municipalities of sample areas. This monocentric assumption appears to be reasonable in the Swiss urban context: as shown in Figure A1 in the Appendix, the majority of people living in the four urban areas of study actually work in the central municipalities. This restriction also fixes the productivity at the workplace to the same level for all individuals, in line with our assumption of exogenous individual wages.

Our baseline sample of analysis, which we refer to as the *All movers* sample, is composed of 11'031 individual locations decisions (or "moves") in the four agglomerations of study,

---

[17] Given that the information reported in the tax burden dataset is only available for generic households, we compute individual-specific tax burden by means of linear interpolation, *e.g.* a married couple with one child (in-between 0 and 2), and a gross income of CHF 110'000.- (in-between CHF 100'000.- and CHF 120'000.-).

[18] The housing prices variable stems from a hedonic regression at the micro-level, which controls for various characteristics of housing lots. Hedonic housing prices are then aggregated to the municipality level. The resulting index measures the relative expensiveness of a municipality in terms of housing prices. As, it is only available up to year 2013, we impute the housing price index based on the average municipal-level growth rate 5 years prior to 2013, for the remaining missing years 2014-2016.

[19] This "intra-cantonal" restriction does not crucially affect the composition of our sample. The urban areas of Zürich, Bern, Lausanne and Geneva being located principally within a single canton, most individuals commuting to agglomerations' CBDs work and live in the same canton, as shown in Figure A2.

between 2010 and 2016. Additional details on the geographical and time structure of our analysis sample can be found in Table C1, in the Appendix.
We assess the representativity of our sample of analysis in Table 4.1. This one reports descriptive statistics on individuals in the *All movers* sample in columns (1)-(2), *i.e.* individuals who work in the city center and who have moved into the urban areas during the observation period. This sub-population might differ from the overall population living in the agglomerations of study. As matter of comparison, we report descriptive information on two broader samples: columns (3)-(4) cover all workers who commute to the central municipalities (including non-movers), whereas columns (5)-(6) refer to all workers living in the areas of study

Table 4.1: Characteristics of sampled individuals and representativity of the analysis sample

	Analysis sample		Comparison samples			
	All movers		**All commuters CBD**		**All ind. in agglo.**	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A : Numerical var.	Mean	SDV	Mean	SDV	Mean	SDV
Location and commuting						
Live in CBD	0.36	(0.48)	0.48	(0.50)	0.39	(0.49)
Commuting distance [in km]	7.42	(5.30)	7.57	(6.43)	12.53	(17.10)
Commuting time [in min.]	28.23	(16.39)	28.70	(16.86)	32.04	(24.50)
Taxation						
TB residence place [in pp.]	10.41	(4.44)	10.45	(4.76)	10.04	(4.99)
TB workplace [in pp.]	10.78	(4.59)	10.72	(4.87)	10.11	(5.08)
Δ TB [in pp.]	-0.37	(0.67)	-0.27	(0.58)	-0.08	(1.01)
Employment						
Income contributor [in kCHF]	87.73	(68.30)	94.45	(77.01)	88.96	(101.82)
Household income [in kCHF]	111.41	(91.54)	129.98	(101.09)	123.58	(129.82)
Hours worked per week	38.64	(9.43)	36.85	(10.71)	34.58	(14.40)
Sociodemographic						
Age	34.59	(9.69)	42.00	(11.37)	41.55	(12.14)
Female	0.54	(0.50)	0.53	(0.50)	0.51	(0.50)
# dependent children	0.39	(0.75)	0.67	(0.95)	0.63	(0.94)
Panel B : Categorical var.	Frequency	%.	Frequency	%.	Frequency	%.
Education						
Primary	613	5.56	7838	9.78	18725	9.53
Secondary	4209	38.16	32821	40.96	83704	42.60
Tertiary	6209	56.29	39461	49.25	94048	47.87
Marital status						
Single	6343	57.50	31167	38.90	77423	39.41
Married	3741	33.91	40397	50.42	99189	50.48
Divorced and other	947	8.58	8556	10.68	19865	10.11
Residence status						
Swiss resident	9172	83.15	66569	83.11	173352	88.23
C-permit	1859	16.85	13525	16.89	23125	11.77

Note: This table reports descriptive statistics on the characteristics of the individuals in our data. Statistics on three different samples are reported. Columns (1)-(2) correspond to the sample of analysis, made of individuals who work in the central municipalities and who have moved recently. Columns (3)-(4) correspond to all individuals who commute to the central municipalities. Columns (5)-(6) encompass all workers in the intra-cantonal agglomeration. Panel A reports summary statistics on numeric variables, while panel B reports information on categorical characteristics.

(including non-movers and those who do not work in the CBDs).

In terms of socio-demographic characteristics, people in the analysis sample are 34.59 years old on average. They are slightly younger than those in the two comparison samples (42.00 and 41.55 years old). This pattern is not surprising given that individuals who make moving decisions tend to be younger. Similarly, they are more likely to be single and have less children. They also tend to have a higher level of education and are somewhat less likely to hold the Swiss nationality, compared to the average resident of the agglomerations.

Individuals in the analysis sample commute an average distance of 7.42 km to work, equivalent to an average commuting time of 28.23 min. They work 38.64 hours per week and earn a yearly individual gross income of CHF 87,730.- on average. This represents the major source of their household income (CHF 111,410.- on average), which is the basis for income taxation in Switzerland. Those figures are somehow similar for the two comparison samples, with slightly more marked differences for the sample covering all workers in the agglomerations. The average tax burden paid at the residence place (10.41%) is smaller on average than at the workplace (10.78%) in the analysis sample: this corresponds to an average tax burden differential (between the residence and workplace) of -0.37 percentage points. Similar patterns, yet less pronounced, are observed for the two comparison samples.

4.3 Descriptive evidence

Figure 4.1 provides descriptive evidence on the trade-offs made by individuals in the analysis sample between taxes and housing prices on the one side, and commuting distance on the other side. Solid lines in the two panels show the unconditional average of the (municipal-level) standardized tax multiplier or housing price index, for each decile of individual-level commuting distance. Dashed lines display the same relationship, net of various municipal-level amenities and individual characteristics.[20]

Panel A shows that longer commutes are offset by lower taxes, whether we consider the conditional or unconditional relationship. This pattern turns out to be slightly convex: individuals living further from the central municipalities are compensated through lower taxes, but this compensation starts decreasing for very long commutes. This picture is consistent with Figure 2.2 in section 2, which describes individuals' choice set and shows that outer-ring municipalities of the urban areas tend to have higher tax rates than suburban municipalities

[20]This conditional relationship is obtained by regressing the municipal level of taxes or price on distance (deciles) and several municipal amenities: government expenditure, transportation network, distance to the next motorway, share of green space, share of 3-4 rooms flats, housing vacancy rate, a binary indicator for the presence of a secondary school, a binary indicator for the presence of lake border and the share of non-Swiss citizens in the municipality. The local level of taxation is also added as a regressor in the case where housing prices is the dependent variable, and conversely. Individual characteristics are also controlled for (size of the household, age, gender, civil status, number of children, income).

Figure 4.1: Taxes and housing prices per decile of commuting distance

(A) Taxes

(B) Housing prices

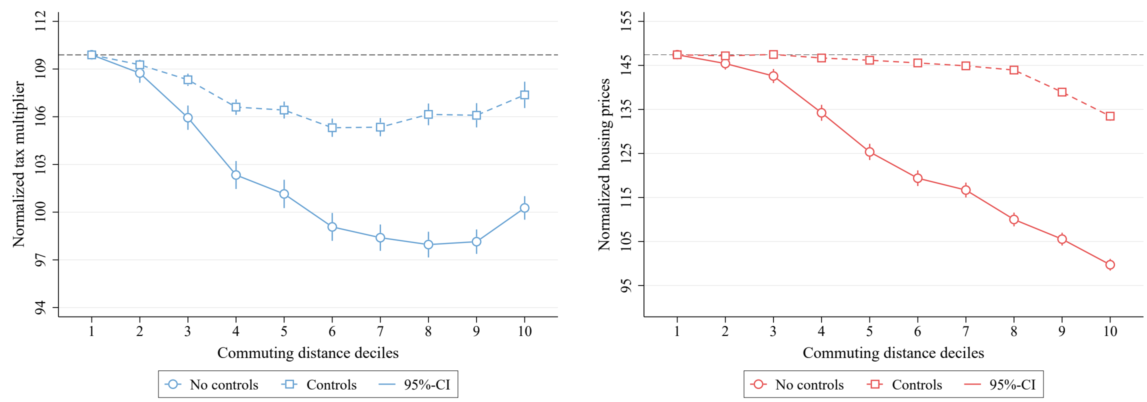

Note: This figure reports descriptive evidence on the trade-offs made by in-moving individuals in our analysis sample between taxes, housing prices and commuting. The figure reports the average (standardized) municipal tax multiplier and municipal-level housing prices per decile of individual-level commuting distance. Deciles of commuting distance are computed within agglomerations. Panel A displays the relationship for the tax multiplier, while panel B shows the same relationship for the housing price index. Both municipal-level amenities are standardized with respect to the agglomeration average. Solid lines correspond to the unconditional tax or housing price level per income decile. Dashed lines report the same relationship, conditional on several municipal and individuals controls, obtained through a multivariate linear regression. 99%-confidence intervals for the average are reported.

located next to city centers. Further, as the tax expensiveness measure only weakly increases for the last deciles of commuting distance in Figure 4.1, this graph suggests that most individuals who commute to CBDs reside in municipalities nearby the city centers, where taxes are lowest. In a similar fashion, panel B shows that longer commutes to work are compensated through lower housing prices, that are ever-decreasing with respect to distance-to-CBD. This pattern also holds when we account for municipal and individual characteristics, as illustrated by the dashed line.

We next provide preliminary evidence on income sorting patterns in Figure 4.2. For each income decile, we report the average level of the tax multiplier, commuting distance and housing prices of the chosen residence municipalities. This graph first reveals the existence of substantial income sorting based on local taxation: the average standardized tax multiplier is monotonically decreasing with income. This pattern suggests that higher incomes tend to self-select into low-tax jurisdictions.[21] Second, high-income individuals tend to locate further away from city centers: the average distance commuted by individuals belonging to the 1st income decile amounts to 6.0 km, while those belonging to the 10th income decile commute an average distance of 8.4 km. A similar pattern is observed for commuting time, as illustrated in Figure C2 in the Appendix. Third, housing prices faced by individuals decrease with the

[21] In quantitative terms, this sorting phenomenon can lead to substantial individual tax savings, as shown in Table C2 in the Appendix. Commuters to the CBDs save between 2.2% to 5.9% of their tax bills on average compared to what they would pay if they were to reside in urban centers, depending on whether they belong to bottom or top income deciles.

Figure 4.2: Taxes, housing prices and commuting distance per income decile

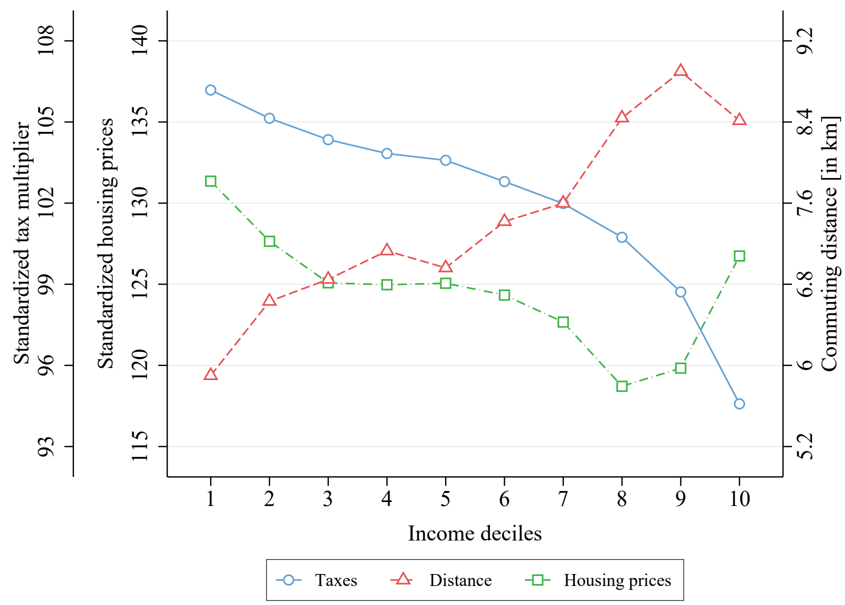

Note: This figure shows graphical evidence on income sorting patterns, based on taxes, housing prices and commuting. The figure plots the average standardized municipal tax multiplier, average standardized municipal-level housing prices and average individual-level commuting distance per income decile.

level of income, up to the 8th decile. This reflects the longer commutes endured by higher incomes and the negative housing price gradient that prevails in our sample urban areas. However, this pattern reverses for individuals belonging to the two last income deciles, who tend to live in jurisdictions featured with substantially higher average housing prices.

These *prima facie* evidence suggest that taxes, housing prices and commuting are valued differently by agents with different income levels. Specifically, high-income earners seem to live outside the city centers, where taxes are low, and commute to CBDs to work. This pattern echoes the income mixing spatial equilibrium described in De Bartolome and Ross (2003), where high incomes get compensated for longer commutes through their preferred level of local public services. In our context, we remain agnostic about the generation of the equilibrium that prevails in the fiscally decentralized agglomerations of study. We rather focus on the empirical evaluation of the trade-offs made by individuals when making residential decisions, given the current spatial distribution of municipal characteristics.

5. Empirical analysis

5.1 Econometric approach

In order to empirically assess our theoretical model's predictions, summarized in Proposition 2, we rely on the indirect utility formulation derived in equation (3.2). We specify an empirical location choice model in a random utility framework to study how individuals with different incomes value local taxes, housing prices and commuting when making residential decisions. Our baseline empirical model reads

$$V_{ijt} = \alpha_j + \psi_\tau \cdot \widehat{\tau}_{ijt}(W_{it}) + \psi_Q \cdot \ln(Q_{jt}) \times \widehat{\ln(W_{it})} + \psi_\delta \cdot \delta_j \times \widehat{\ln(W_{it})} \\ + \boldsymbol{X_{it}}\gamma_j + \eta_{ijt}, \tag{5.1}$$

where V_{ijt} represents the random utility obtained by individual i in moving-year t, if she chooses to settle in municipality j.[22] The parameter α_j is a municipality- (or alternative-) specific constant capturing the utility that is derived from all municipal characteristics which are fixed in time, averaged across all individuals i. This coefficient does not only encompass the average utility obtained from municipal amenities B_j, but it also captures the "main effect" of taxes, commuting and housing prices on the indirect utility. Consequently, the average (marginal) utility derived from each specific municipal characteristic is not directly identified in the model.[23],[24]

To capture differentials in the valuation of the municipal characteristics of interest with respect to income, we introduce interaction terms between on the one side municipal-level taxes, housing prices and commuting, and on the other side individual's (household) income. For housing prices, we specify a multiplicative interaction term between the log-housing price index of municipality j and the log-income of individual i demeaned with respect to the urban-area specific yearly income average, all in year t: $\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$.[25] This interaction term captures the mediating effect of income on the marginal disutility from an increase in housing prices. We proceed in a similar fashion for commuting, with $\delta_j \times \widehat{\ln(W_{it})}$ measuring how commuting is valued across the income distribution. Specifically, the municipal-level commuting measure δ_j is defined either as the *Centroid distance* (from center to center), *Av-*

[22]Given that we pool our sample across agglomerations, we specify disjoint choice sets for each year and agglomeration. As a matter of computational ease, we restrict the choice set to municipalities that are located less than 20 km from the central municipality (as measured by the centroid distance), and for which we observe at least 10 moves. The estimated coefficients correspond to a weighted average of the effects across agglomeration areas.

[23]This limitation comes from the fact that the model is identified on differences in utility. More specifically, only the average difference in utility between two jurisdiction is identified, provided that the alternative specific constant of one alternative is standardized (to zero).

[24]Given that we observe moving decisions over several years, the "main effect" of (time-varying) taxes and housing prices are theoretically identified in our context. However, this identification would rely on the time-variability of these variables, which is very limited over the sample period, as described in Figures D2 and D3. This makes such identification infeasible in practice.

[25]Formally, $\widehat{\ln(W_{it})}$ is defined as $\ln(W_{it}) - \ln(\overline{W}_t)$, where W_{it} is the household income of individual i in year t and $\overline{W}_t$ is the urban-area specific yearly income average. The agglomeration index is omitted as a matter of simplicity.

erage commuting distance or *Average commuting time* between municipality j and the central jurisdiction of the agglomeration. The two last measures are computed based on the information reported in the *Structural Survey*, and account for non-zero commuting of individuals residing in the central municipalities.[26] Regarding taxes, we consider an alternative interaction term, namely the tax burden to which individual i would be subject if she were to move to municipality j, demeaned with respect to the average tax burden she faces in the agglomeration (both in moving-year t). This variable $\widehat{\tau}_{ijt}(W_{it})$ depends on the municipal tax multiplier τ_j, individual's income W_{it} and tax-payer's type. It is equivalent to an interaction term between municipal taxation level and individual income, and does not need to be further interacted with income.[27]

In addition to the above-mentioned interaction terms, we control for individual-specific characteristics in our specification through the vector $\boldsymbol{X_{it}}$ (sex, age, marital status, household size), whose coefficients are municipal-specific. The idiosyncratic taste error term η_{ijt} is assumed to follow a generalized extreme value (GEV) distribution of type I (McFadden 1974, 1978).[28] Under this assumption, the conditional probability of individual i choosing residential location j can be written as $P_{ijt}(\xi_j) = \frac{\exp V_{ijt}}{\sum_k \exp(V_{ikt})}$, where $\xi_j = \{\alpha_j, \psi_\tau, \psi_Q, \psi_\delta, \gamma_j\}$. This specification makes the implicit assumption of *Independence of Irrelevant Alternatives* (IIA), which we further discuss in the context of our robustness checks.

Our empirical strategy relies on the identifying assumption that municipal characteristics are taken as given by atomistic individuals when they make residential decisions. This assumption is presumably met in our context, as we study urban areas that are densely populated and where each in-moving individual has a marginal effect on the overall residing population. Each atomistic agent is therefore unlikely to have a direct effect on municipal housing prices or local tax rates where she decides to settle down.[29] Moreover, our sample period is relatively short (2010-2016), which limits the impacts of aggregate in-moving population dynamics on local public finances and housing markets. This can be seen in Figures D2 and D3 in the

[26]The two latter variables account for the fact that individuals living in the central municipality also commute to work, while the former assume zero commuting for those individuals. The relationship between the three municipal-level commuting measures are reported in Figure D4.

[27]The tax burden measure $\widehat{\tau}_{ijt}(W_{it})$ accounts for tax-payer type and the exact cantonal tax schedule. We thus favor this measure over the multiplicative interaction term between (log-) municipal tax multiplier and income $\ln(\tau_{jt}) \times \widehat{\ln(W_{it})}$, the latter representing a more constrained specification of the interaction term between local taxes and income.

[28]In the ensuing estimations, we follow the literature on residential location decisions and do not cluster the error term by municipality (see *e.g.* Schmidheiny, 2006; Basten et al., 2017; Mulalic and Rouwendal, 2020). Given that we observe every individual making a location decision once and that the "treatments" of interest (interactions between municipal characteristics and individual's income) are assigned at the individual level, clustering is not warranted (Abadie et al., 2023).

[29]Taken altogether, in-moving individuals represent a relatively small proportion of the overall residing population in sample urban areas (approximately 5% to 9%, as shown in Figure D1).

Appendix, which show that local tax rates and housing prices growth rates have remained relatively stable during the period covered by our study.

Relating our theoretical predictions on individuals' preferences towards municipal characteristics to empirics, we expect ψ_Q to be positive, as the mediating effect of income on the marginal disutility from housing prices is predicted to positive. Conversely, the parameter ψ_τ is expected to be negative, *i.e.* the mediating effect of income on the marginal disutility from taxes is negative. Finally, the sign of ψ_δ remains *a priori* undetermined, due to the counteracting forces emphasized in the theoretical section.

5.2 Income sorting results

5.2.1 Main results

We report estimation results of equation (5.1) in Table 5.1. Odd columns report the baseline estimates including only the municipal characteristics of interest, for our three different measures of commuting. Results in the even columns add interaction terms between other municipal amenities on the one side (municipal-level share of 3-4 bedroom apartments, share of foreigners, share of green space and presence of a lake or river), and individual's income on the other side. These amenities could be additional drivers of income sorting and confound the estimates of the municipal characteristics of interest. As regressors are not measured in the same unit within models and since commuting is not measured in the same unit across models, the utility metrics and the quantitative interpretation of coefficients differ across our estimates. Consequently, the focus is set on the signs of the coefficients rather than on their magnitude.

The estimated $\widehat{\psi}_\tau$ parameter associated with demeaned tax burden turns out to be negative and strongly significant across all specifications. This result confirms that high-income individuals are effectively more sensitive to taxes due to tax progressivity, in line with previous findings in the literature (Schmidheiny, 2006; Schmidheiny and Slotwinski, 2018; Brülhart et al., 2021). Our estimation results also show that richer individuals are less deterred by housing prices: the coefficient $\widehat{\psi}_Q$ for the interaction term between housing prices and individual income is found to be positive and strongly significant in all columns. This reveals that high incomes' utility is less affected by housing prices compared to low-income individuals, a consequence of non-homothetic preferences for housing in light of our theoretical model. Finally, the *a priori* ambiguous change in the valuation of commuting along the income distribution turns out to be positive: across all specifications, the coefficient $\widehat{\psi}_\delta$ is estimated to be positive and significant, whatever the measure of commuting δ_j we consider. Those findings reveal that high income-individuals are less sensitive to commuting, and corroborate the observation first made by Schmidheiny (2006) for the agglomeration of Basel. According to

Table 5.1: Location choices and income sorting

Measure of δ_j	Centroid distance		Avg. commut. distance		Avg. commut. time	
	(1)	(2)	(3)	(4)	(5)	(6)
Location Choice Model						
Municipal characteristics						
$\widehat{\tau}_{ijt}(W_{it})$	-0.619***	-0.554***	-0.673***	-0.593***	-0.636***	-0.559***
	(0.055)	(0.055)	(0.053)	(0.054)	(0.053)	(0.055)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	0.830***	0.849***	0.628***	0.661***	0.745***	0.742***
	(0.126)	(0.160)	(0.110)	(0.150)	(0.111)	(0.142)
$\delta_j \times \widehat{\ln(W_{it})}$	0.037***	0.023***	0.038***	0.018***	0.044***	0.029***
	(0.005)	(0.006)	(0.006)	(0.007)	(0.005)	(0.006)
Other municipal amenities B_j						
$\ln(\text{Housing}_{jt}) \times \widehat{\ln(W_{it})}$		-0.502***		-0.485**		-0.546***
		(0.192)		(0.192)		(0.192)
$\ln(\text{Foreigners}_{jt}) \times \widehat{\ln(W_{it})}$		-0.428***		-0.457***		-0.420***
		(0.077)		(0.081)		(0.076)
$\ln(\text{Green Space}_{jt}) \times \widehat{\ln(W_{it})}$		0.077**		0.061*		0.059*
		(0.036)		(0.035)		(0.036)
$\mathbb{I}(\text{Lake}_{jt} = 1) \times \widehat{\ln(W_{it})}$		-0.093*		-0.101*		-0.059
		(0.056)		(0.056)		(0.055)
Average ASC						
Central mun. $\hat{\alpha}_j$	0.000	0.000	0.000	0.000	0.000	0.000
Average suburban $\hat{\alpha}_j$	-6.130	-6.077	-6.168	-6.104	-6.130	-6.077
Indiv.-specific controls	Yes	Yes	Yes	Yes	Yes	Yes
Nbr. municipalities	166	166	166	166	166	166
Nbr. individuals	11031	11031	11031	11031	11031	11031
Log-likelihood	-30645.480	-30610.090	-30650.330	-30615.097	-30640.527	-30607.002

Note: This table reports the pooled estimates of equation (5.1) with disjoint choice sets across agglomerations. The error term η_{ijt} is assumed to follow a Type I GEV distribution. In each specification, the alternative-specific constants of the central municipalities are standardized to 0. The three sets of estimates in columns (1)-(2), (3)-(4) and (5)-(6) are based on different measures of commuting δ_j between municipality j and the center of the agglomeration (*Centroid distance*, *Average commuting distance* and *Average commuting time*). Columns (1)-(3)-(5) report the estimation results for the baseline specification, while columns (2)-(4)-(6) introduce additional interactions between income and municipal amenities. Robust standard errors are reported. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

our theoretical model, this result is explained by the fact that the benefit from lower taxes, exacerbated by tax progressivity, more than outweighs the commuting disutility for high incomes. Consequently, high-income individuals tend to self-select into low-taxes municipalities located in the suburban area, at the cost of longer commutes.

Does this pattern however hold for any commuting distance or time? In Table D1 in the Appendix, we additionally interact income with the squared measure of commuting in order to assess non-linearities in the relationship between commuting and income sorting. Across our different specifications a similar pattern emerges: the parameter associated with the interaction of income with the commuting measure is found to be positive and significant, whereas the interaction with squared commuting distance or time consistently exhibits a statistically significant negative coefficient. This indicates that high-income earners are more willing to commute from the close suburb of city-centers than lower incomes, but this relation reverses for longer commutes.[30]

5.2.2 Robustness

We assess the robustness of our income sorting results in several ways.

First, we address the question of the exogeneity of the municipal characteristics for individuals in our analysis sample, which contains all movers into the agglomerations of study. Specifically, those people who already lived in the agglomerations before moving might have social bonds within the sample urban areas, and might base their residence choices on other location factors that are unobserved to us. As a robustness check, we re-estimate our baseline empirical location choice model on two restricted samples, that condition on movers' previous residence place: the *Outside urban area* and *Outside canton* samples. These are respectively made of individuals coming from outside the urban area and of those moving from another canton than the one where the agglomeration is located. The residential decisions of individuals in these restricted samples are presumably less prone to be affected by unobserved components that are correlated with taxes, housing prices and commuting, given that they may have limited knowledge about the agglomeration where they move.[31] Coefficients estimated on these subsamples are reported in Table D2 and turn out to be qualitatively the same across all estimation samples, hence confirming our baseline income sorting results.

Second, we address the IIA assumption, which states that the unobserved component η_{ijt} of the random utility is independent across locations j. Intuitively, this assumption implies that if municipality j' were to be removed from the choice set of mover i, individual's choice probabilities for the remaining alternatives would shift in a proportional manner. In our context, one could rightfully presume that densely populated municipalities, or even the central municipalities themselves, are less comparable to more remote and sparsely populated suburban jurisdictions. Correlation between preferences for municipalities of the same type would then contradict the IIA assumption. To assess the validity of this assumption and the robustness of our baseline results, we proceed in two ways. First, we implement nested logit models, which relax the IIA assumption by allowing correlations of the error term within nests of alternatives. Second, we run the baseline conditional logits excluding the central municipalities.

In Tables D3 and D4, we report nested logit estimates, with two nests specified according to municipal-level population density measures (high and low-density). Those are based on

[30] This pattern was already noted by Schmidheiny (2006) for the case of Basel. In our study, we propose a structural interpretation of this observation through, among others, the role of decentralized progressive taxation and exponentially increasing commuting costs. Intuitively, for high-income individuals, the disutility from too long commutes might outweigh potential gains in local taxes from locating farther, even with strong income tax progressivity.

[31] This approach is similar in essence to the study of Schmidheiny and Slotwinski (2018), who rely on moving decisions made by foreigners living in Switzerland and who obtain a permanent residence permit.

two official municipality classifications from the *Swiss Statistical Office* and *EUROSTAT*.[32] The nested logit model identifies additional parameters $\lambda_{\text{high-density}}$ and $\lambda_{\text{low-density}}$, each capturing the degree of dissimilarity between the alternatives within the corresponding nest. Those parameters indicate whether the nested model should be favored over the conditional logit specification.[33] We estimate this model on each agglomeration separately, and additionally report agglomeration-specific conditional logit results as benchmarks and further checks of our pooled estimates obtained on all agglomerations simultaneously. Agglomeration-specific conditional logit results are found to be mostly consistent with our pooled estimates. Also, taking the nest structure into account does not lead to a noticeable change in the estimated coefficients and standard-errors. However, for most specifications the dissimilarity parameters $\lambda_{\text{high-density}}$ and $\lambda_{\text{low-density}}$ are either bigger than one or not significantly different from one. As no nest structure generally meets the assumption of the random utility framework, we favor the conditional logit specification over the nested model.

In Table D5, we report conditional logit results based on the pooled sample, where we remove the central municipalities. This exercise can be considered as an additional check for the IIA assumption. Our baseline estimation results remain qualitatively unchanged: $\widehat{\psi}_{\tau}$ is estimated to be negative, while estimates of $\widehat{\psi}_{Q}$ and $\widehat{\psi}_{\delta}$ are both positive. Parameters remain statistically significant across all specifications.

Third, we elaborate on another competing explanation for the lower valuation of centrality by high-income earner: the mode of transport, which might depend on the income class (Fretz et al., 2022). If higher-income workers are more likely to commute by car, longer commutes will not necessarily entail a stronger disutility for them. Inversely, if low-income workers use public transport more often, they might endure stronger disutility from commuting. Figure D5 displays the self-reported main mode of transport for commuting in our sample, per income decile. In spite of a slight increase in car use and a slight decrease in public transit with income level, the proportions of the different modes of transportation remain relatively stable across income deciles, with public transports accounting for the lion's share in each income class. As an additional check, we re-estimate our baseline regression controlling for individuals' main mode of transportation. This case-specific variable is endogenously determined as it is concurrent to the residential location choice, hence its omission in our baseline specification.

[32] Firstly we use the 9-categories classification from the *Swiss Statistical Office* where we assign to the high-density nest municipalities categorized as "urban and part of a large agglomeration" (OFS, 2017). The other municipalities are assigned to the low-density nest. Secondly, we use the *DEGURBA* classification from *EUROSTAT* which separates municipalities in 3 categories: densely populated areas, suburbs and rural areas (EUROSTAT, 2018). We assign municipalities in the first category to the high-density nest and the others to the low-density nest.

[33] Formally, if at least one dissimilarity parameters is significantly smaller than 1, then the nested specification is preferred. Otherwise, it is equivalent to the conditional logit. Finally, if the parameter is bigger than unity, the nested logit is not consistent with random utility theory (Hensher et al., 2005).

Nevertheless, estimates accounting for this additional control remain the same in essence as those obtained in our baseline specification, as shown in Table D6. Transportation mode alone thus cannot explain entirely why high-income earners locate further away from city centers.

5.3 Local taxes and the willingness to commute

Our baseline results show that high-income earners value centrality less than low incomes. According to our theoretical framework, this pattern is explained by tax progressivity and low levels of taxes in the suburbs: progressive decentralized taxation directly impacts individuals' willingness to commute, and this more markedly for high incomes. To formally test this hypothesis, we estimate the following equation:

$$\begin{aligned} V_{ijt} = \mu_j + \beta_\tau \cdot \widehat{\tau}_{ijt}(W_{it}) + \beta_Q \cdot \ln(Q_{jt}) \times \widehat{\ln(W_{it})} + \beta_\delta \cdot \delta_j \times \widehat{\ln(W_{it})} \\ + \beta_{\delta\tau} \cdot \delta_j \times \widehat{\tau}_{ijt}(W_{it}) + \boldsymbol{X_{it}}\zeta_j + \nu_{ijt}. \end{aligned} \tag{5.2}$$

This specification adds an interaction term between the municipality-level commuting measure δ_j and the demeaned tax burden $\widehat{\tau}_{ijt}(W_{it})$ to the baseline equation (5.1). This additional regressor is equivalent to a triple interaction between the municipal-level commuting measure, local tax multiplier and individual income. The coefficient $\beta_{\delta\tau}$ associated with this new regressor measures how local taxes impact the disutility from commuting along the income distribution. Our theoretical framework predicts $\beta_{\delta\tau}$ to be negative as individuals are more willing to commute to a peripheral municipality if it offers lower taxes. Formally, taking the derivative of the above equation with respect to δ_j yields $\frac{\partial V_{ijt}}{\partial \delta_j} = \beta_\delta \cdot \widehat{\ln(W_{it})} + \beta_{\delta\tau} \cdot \widehat{\tau}_{ijt}(W_{it})$. A negative $\beta_{\delta\tau}$ hence implies that the marginal disutility from commuting is less negative as $\widehat{\tau}_{ijt}(W_{it})$ decreases, *i.e.* when local taxes are lower (see Figure D6 for a visual representation). Also, as $\widehat{\tau}_{ijt}(W_{it})$ increases with income due to tax progressivity, the attenuating effect of lower taxes on commuting disutility is stronger for high-income individuals.

Empirical estimates of equation (5.2) are reported in Table 5.2. The coefficients for the tax burden and housing prices interactions remain qualitatively unchanged compared to our baseline estimates. The estimated parameter of the interaction term between the commuting measure and income remains for its part positive and significant, across all specifications. This implies that even if municipal taxes were harmonized within agglomerations, higher incomes would still be more willing to commute than lower incomes. Additionally, the coefficient on the interaction $\delta_j \times \widehat{\tau}_{ijt}(W_{it})$ is estimated to be negative and strongly statistically significant. Those results are robust to the inclusion of additional municipal amenities interacted with income, as shown in even-numbered columns.

Decentralized and progressive income taxation therefore seems to directly affect individuals' willingness to commute. This additional result is coherent with the intuitions developed in

Table 5.2: Location choices and income sorting
Commuting disutility and decentralized taxation

Measure of δ_j	Centroid distance		Avg. commut. distance		Avg. commut. time	
	(1)	(2)	(3)	(4)	(5)	(6)
Location Choice Model						
Municipal characteristics						
$\widehat{\tau}_{ijt}(W_{it})$	-0.655***	-0.593***	-0.686***	-0.608***	-0.551***	-0.480***
	(0.055)	(0.056)	(0.053)	(0.054)	(0.056)	(0.057)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	0.810***	0.849***	0.620***	0.663***	0.746***	0.789***
	(0.127)	(0.161)	(0.110)	(0.151)	(0.112)	(0.143)
$\delta_j \times \widehat{\ln(W_{it})}$	0.051***	0.038***	0.054***	0.034***	0.067***	0.051***
	(0.006)	(0.007)	(0.008)	(0.008)	(0.007)	(0.008)
$\delta_j \times \widehat{\tau}_{ijt}(W_{it})$	-0.003***	-0.003***	-0.003***	-0.003***	-0.004***	-0.004***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Average ASC						
Central mun. $\hat{\alpha}_j$	0.000	0.000	0.000	0.000	0.000	0.000
Average suburban $\hat{\alpha}_j$	-6.098	-6.035	-6.140	-6.069	-6.079	-6.022
Indiv.-specific controls	Yes	Yes	Yes	Yes	Yes	Yes
Amenities B_j	No	Yes	No	Yes	No	Yes
Nbr. municipalities	166	166	166	166	166	166
Nbr. individuals	11031	11031	11031	11031	11031	11031
Log-likelihood	-30636.681	-30599.908	-30644.476	-30608.237	-30627.948	-30595.655

Note: This table reports the pooled estimates of equation (5.2), with disjoint choice sets across agglomerations. The error term ν_{ijt} is assumed to follow a Type I GEV distribution. In each specification, the alternative-specific constants of the central municipalities are standardized to 0. The three sets of estimates in columns (1)-(2), (3)-(4) and (5)-(6) are based on different measures of commuting δ_j between municipality j and the center of the agglomeration (*Centroid distance*, *Average commuting distance* and *Average commuting time*). Columns (1)-(3)-(5) report the estimation results for the baseline specification, while columns (2)-(4)-(6) introduce additional interactions between income and municipal amenities. Robust standard errors are reported. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

the theoretical framework: high incomes, whose opportunity cost of commuting is affected more strongly by taxation, are more likely to commute from peripheral low-tax municipalities to save on their tax bill.

6. The spatial distribution of in-movers' income under harmonized tax rates

In order to better understand and quantify the role played by decentralized taxation in residential and commuting decisions, we conduct an illustrative counterfactual analysis based on the empirical preference parameters estimated in section 5, conditional on the prevailing spatial equilibrium. This counterfactual exercise corresponds to a fiscal harmonization scenario: we shut down fiscal decentralization at the municipal level by setting the individual-specific tax burden to its within-agglomeration average: $\tau_{jit} = \bar{\tau}_{it}$. We compute individual location probabilities under the counterfactual scenario, and investigate how this affects the income distributions of in-moving individuals in the central and suburban municipalities. We show that conditional on the current spatial distribution of other municipal characteristics, shutting off fiscal decentralization leads to more high-income movers choosing to locate in the central municipalities, hence shifting the income distribution in agglomeration centers to the right, and conversely for suburban municipalities.

6.1 Methodology

Given our utility specification and our empirical estimates from equation (5.2), spatially invariant taxes imply that $\beta_\tau \cdot \tau_{ijt} = \beta_\tau \cdot \overline{\tau}_{it}$ and $\beta_{\delta\tau} \cdot \delta_j \times \tau_{ijt} = \beta_{\delta\tau} \cdot \delta_j \times \overline{\tau}_{it}$. Taxation does not affect directly location choices of individuals any more, nor does it affect their willingness to commute. Formally, we use our empirical estimates of the parameters in Table 5.2, where commuting is measured through the *Average commuting distance*, to compute baseline and counterfactual individual location probabilities for each municipality j. Those are computed using the logistic link function, $P_{ijt}(\widehat{\xi}) = \frac{\exp(V_{ijt})}{\sum_k \exp(V_{ikt})}$, where P_{ijt} is the probability that individual i chooses jurisdiction j in year t.

In the benchmark scenario with fiscal decentralization, we simply predict probabilities using the deterministic part of utility V_{ijt}. This component is a function of the estimated parameter vector $\hat{\xi} = \{\hat{\mu}_j, \hat{\beta}_\tau, \hat{\beta}_\delta, \hat{\beta}_Q, \hat{\beta}_{\delta\tau}, \hat{\zeta}_j; Q_t\}$ evaluated at τ_{ijt}. We proceed the same way for the counterfactual scenario without decentralized taxation, except that we evaluate V_{ijt} at the individual-specific average tax burden $\overline{\tau}_{it}$ within the agglomeration. The resulting probabilities are denoted respectively $\widehat{P}_{ijt}$ for the benchmark case, and $\overline{P}_{ijt}$ for the counterfactual scenario.

Before turning to the results, we discuss the assumptions and limitations behind our counterfactual exercise. First, this approach can be considered as a partial equilibrium one: it emphasizes demand-shifts following fiscal harmonization, ignoring responses from the supply side. In particular, it does not account for the fact that housing prices may adjust at the new equilibrium. Specifically, lower taxes are usually associated with higher housing prices through the bid-rent mechanism.[34] To account for this phenomenon, we additionally compute another set of location probabilities $\widetilde{P}_{ijt}$, based on a housing-price vector that accounts for this capitalization effect.[35] Second, the removal of fiscal federalism might affect the empirical preference parameters that characterize individuals' location decisions. In particular, the harmonization of tax rates might impact the average marginal disutility from local taxes, which is captured by the alternative-specific constant. Preferences for housing might also be modified, even though, under the fiscally centralized scenario, high incomes remain less sen-

[34]Such capitalization phenomenon is observed in our data, as it can be seen in Figure E1. As shown in Basten et al. (2017), the aggregate relationship between taxes and housing prices however tends to overstate the actual causal capitalization effect.

[35]This new vector of prices is based on the elasticity of housing prices with respect to local income taxes estimated by Basten et al. (2017). This study explores the capitalization of income taxes into housing prices using a border regression discontinuity design and based on Swiss micro-geographic data. The authors estimate a relatively low elasticity of housing prices with respect to local income taxes of -0.277 in their most conservative specification. As this study presents a setting very close to ours, it should enable us to approximate reasonably the housing price levels that would prevail in the absence of fiscal federalism.

sitive to housing prices due to non-homothetic housing preferences. As we do not identify the effect of the removal of decentralized taxation on location choice parameters, we assume those to remain constant under the counterfactual scenario.

6.2 Results

We first assess the fit of our empirical model with respect to spatial income sorting. We use the benchmark predicted location probabilities $\widehat{P}_{ijt}$ to compute the average income of in-moving individuals in each municipality j, under the benchmark case: $\widehat{W}_j = \frac{\sum_t \sum_i \widehat{P}_{ijt} W_{ijt}}{\sum_t \sum_i \widehat{P}_{ijt}}$. Figure E2 in the Appendix shows that the municipal income averages observed in the data are relatively well replicated by our empirical model: the unweighted and population-weighted correlation coefficients between observed and predicted municipal income averages amount to 0.859 and 0.918 respectively.

Next, we use predicted probabilities to compute the cumulative income distribution functions (CDFs) in two subsets of municipalities: the four central municipalities and all the remaining suburban jurisdictions, which we label as the "suburbs".[36] The CDFs for the benchmark case ($\widehat{P}_{ijt}$) and the counterfactual scenario without housing prices adjustment ($\overline{P}_{ijt}$) are reported in Figure 6.1. The additional scenario where housing prices are adjusted for tax capitalization ($\widetilde{P}_{ijt}$) is presented in Figure E3, yielding very similar results.

Figure 6.1 illustrates the shifts in income distributions resulting from changes in demand for locations after the harmonization of taxes across space. In panel A, we observe that under the counterfactual scenario without fiscal decentralization, central municipalities tend to attract richer movers, as shown by the right shift of the red curve. Suburban jurisdictions, on the contrary, tend to attract poorer new residents, as visible through the left shift of the blue curve. Panel B reports the difference in kCHF between the values of the income deciles under the baseline and under the counterfactual scenario. It confirms that city-centers become more attractive for in-moving high incomes upon tax harmonization. The 90$^{\text{th}}$ percentile in city centers goes up by approximately 13,500 CHF, while the 10$^{\text{th}}$ decile increases by 4,000 CHF. In contrast, income deciles in the suburbs decrease relatively homogeneously.

Our counterfactual exercise illustrates that the removal of decentralized taxation would lead to a gentrification of central jurisdictions: under the fiscal harmonization scenario, agglomeration centers attract more high-income new residents, who end up commuting shorter distances compared to the fiscally decentralized case.

[36] Formally, we compute the CDFs as follows. Let $\mathcal{C}$ and $\mathcal{S}$ denote respectively the set of city centers and the set of suburban municipalities. The CDF of centers and of suburbs can be written as:

$$CDF^{\mathcal{K}}(\omega) = \frac{\sum_j \sum_t \sum_i \mathbb{I}\left(\widehat{\ln(W_{it})} \leq \omega\right) \cdot P_{ijt}}{\sum_j \sum_t \sum_i P_{ijt}} \quad \text{with} \quad j \in \mathcal{K} \quad \text{and} \quad \mathcal{K} = \mathcal{C}, \mathcal{S}. \tag{6.1}$$

Figure 6.1: Counterfactual income distribution, central municipalities vs. suburbs

(A) Income CDFs

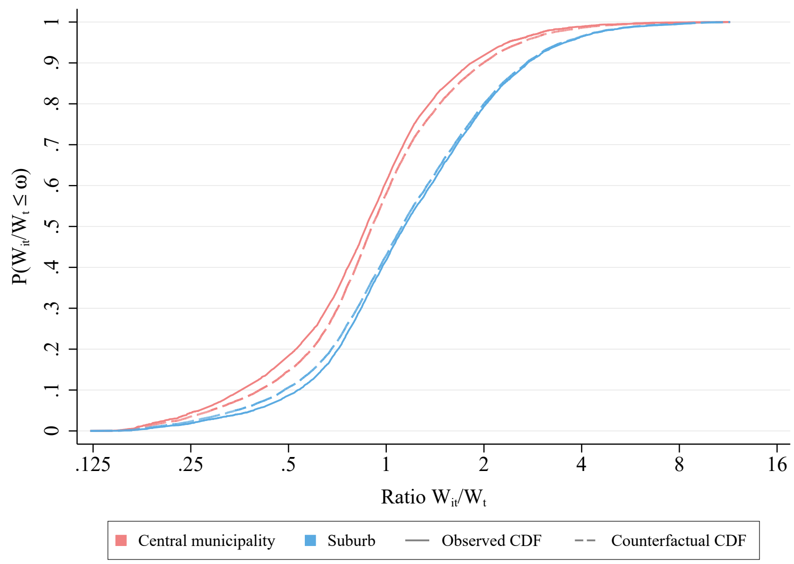

(B) Changes in income deciles

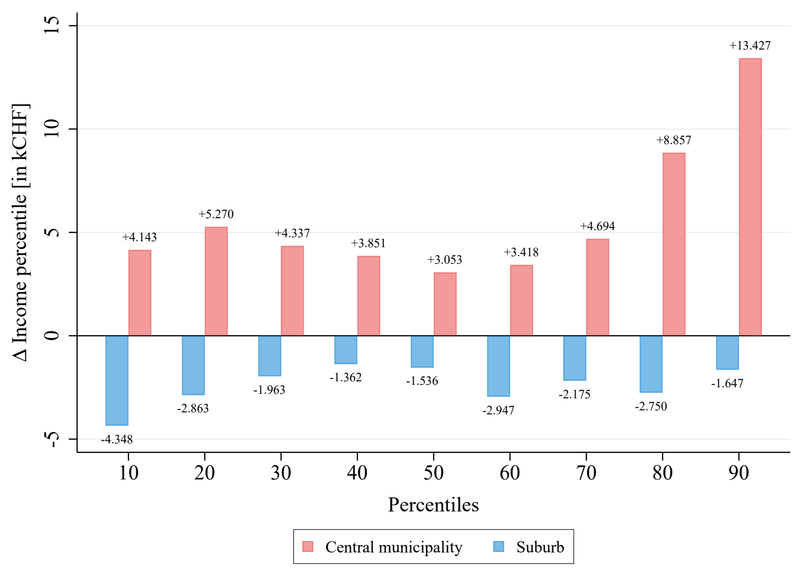

Note: This figure reports graphical results of our counterfactual exercise. Panel A plots the CDFs of income in the central municipalities (in red) and in the suburbs (in blue), for both the observed data (solid line) and the counterfactual scenario without accounting for housing prices capitalization (dashed line). So as to obtain a smooth CDF for the four agglomerations together, we consider the exponentiated demeaned log-income $W_{it}/\overline{W}_t$ (a monotonic transformation of the term $\widehat{\ln(W_{it})}$), reported on a log-scale. Panel B reports corresponding changes in income deciles (reported in kCHF .-), for the central municipalities and the suburbs.

6.3 Discussion

Our counterfactual results echo several recent developments in the urban and public finance literature. First, they complement the findings of Gaigné et al. (2022), who study commuting and income sorting in a fiscally centralized economy. The authors show empirically that commuting is a strong driver of income segregation in the Netherlands. Consistently with the usual bid-rent argument, they find that high-income individuals tend to locate closer to city-centers. In contrast, our findings suggest that the existence of a progressive decentralized tax system might reduce high incomes' incentive to locate in central places. As illustrated in our counterfactual exercise, high-income earners are more likely to choose a central jurisdiction if they cannot self-select into suburban low-tax municipalities.

Second, our results are in line with recent literature on the link between transportation and income sorting. Fretz et al. (2022) show that a municipality whose transportation infrastructure is improved through access to highways attracts more high-income earners. Such broadening of the tax base may lead to lower tax rates, which in turn will attract richer new residents. Following this line of reasoning, Straumann (2012) argues that the expansion of the commuting zone of Zürich led to local tax decreases, notably in the neighboring canton of Schwyz, but also in suburban municipalities located in Zürich. Our empirical results corroborate this idea: in fiscally decentralized urban areas with dense transportation networks, high incomes face incentives to live in well-connected suburbs where taxes are low. Our counterfactual exercise shows that when the fiscal incentive is removed, central municipalities become more attractive for high incomes, even though the suburbs remain richer on average.

Finally, our results relate to the work of Su (2022) who shows that gentrification of American city-centers can mainly be attributed to the rising value of time and the influence of endogenous amenities. Our results substantiate this idea: progressive income taxation, which affects the net-of-tax revenue of higher incomes more markedly, coupled to fiscal decentralization, reduce the opportunity cost of commuting. Our counterfactual analysis illustrates that tax harmonization in fiscally decentralized urban areas might remove a potential counteracting force against the gentrification of city centers.

7. Conclusion

Progressive income taxation has been shown to be a major source of income sorting across space within fiscally decentralized urban areas. However, most studies analyzing individual residential decisions in these contexts have overlooked cities' structure and the fact that workers commute to productive centers to work. In this paper, we explore how progressive taxation coupled with decentralized taxing powers affect the location choices and willingness to commute of individuals with different income levels.

We develop a structural model of individual residential decisions in a fiscally decentralized monocentric urban environment. The model is featured with decentralized progressive taxation at the residence place, non-homothetic preferences for housing and explicit commuting costs. Individuals choosing their residence place are predicted to trade off taxes, housing prices and commuting differently along the income distribution. Specifically, high incomes are predicted to be less sensitive to housing prices and, under certain conditions, more sensitive to local taxes. The prediction on commuting remains ambiguous. This ambiguity arises from progressive income taxation and fiscal decentralization, which reduce high incomes' larger disutility from commuting more markedly, possibly up to the point where they become less sensitive to commuting.

We confront the predictions of our theoretical model to the data, using information on individual residential choices in four of the main agglomerations of Switzerland. Based on a random utility framework, we analyze decisions of in-moving individuals working in agglomerations' central municipalities. Our empirical results show that individuals trade off municipal characteristics differently depending on their income level: high-income earners tend to be more sensitive to taxes, less sensitive to housing prices and less sensitive to commuting distance. Further, we find that lower local taxes directly reduce the marginal disutility from commuting, consistently with the mechanisms emphasized in our theoretical framework. This picture is also consistent with the observation that high incomes are more likely to reside in suburban municipalities, where taxes are lower, hence enduring longer commutes to work.

To illustrate the role of decentralized progressive taxation in individuals' spatial income sorting, we perform a counterfactual exercise based on our empirical location choice model estimates. We impose a homogeneous, centralized tax scheme, and show that in the absence of fiscal federalism, high incomes would be more likely to locate in agglomerations' central municipalities, where taxes were higher. Consequently, high-income (low-income) earners would commute shorter (longer) distances to work.

Our research enriches the literature on the distorting nature of local taxation and its effect on residential and commuting decisions. Even tough our structural and empirical modeling approaches are derived in a partial equilibrium framework, they provide further understanding on how local taxes affect individual decisions within highly fiscally fragmented urban areas. From the perspective of suburban municipalities, taxes can be seen as an instrument to compensate tax payers for the cost of commuting to nearby productive centers. If lower taxes are effectively set in the suburbs, in-moving high-incomes might have a stronger incentive to locate farther from city centers. Progressive taxation coupled with fiscal decentralization might therefore play as a counteracting force against the gentrification of central locations. We delegate the task of fully modeling urban and fiscally decentralized economies, accounting for their spatial structure, commuting flows and fiscal competition aspects to further research.

Bibliography

Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. M. (2023). When should you adjust standard errors for clustering? *The Quarterly Journal of Economics*, 138(1):1–35.

Agrawal, D. R. (2015). The tax gradient: spatial aspects of fiscal competition. *American Economic Journal: Economic Policy*, 7(2):1–29.

Agrawal, D. R. and Hoyt, W. H. (2018). Commuting and taxes: theory, empirics and welfare implications. *The Economic Journal*, 128(616):2969–3007.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The economics of density: evidence from the Berlin Wall. *Econometrica*, 83(6):2127–2189.

Alonso, W. (1964). *Location and land use. Toward a general theory of land rent.* Harvard University Press, Cambridge, MA.

Basten, C., von Ehrlich, M., and Lassmann, A. (2017). Income taxes, sorting and the costs of housing: Evidence from municipal boundaries in switzerland. *The Economic Journal*, 127(601):653–687.

Brownstone, D. and Small, K. A. (2005). Valuing time and reliability: assessing the evidence from road pricing demonstrations. *Transportation Research Part A: Policy and Practice*, 39(4):279–293.

Brülhart, M., Danton, J., Parchet, R., and Schläpfer, J. (2021). Who bears the burden of local taxes? *CEPR Discussion Paper No. 15826*.

Brülhart, M., Bucovetsky, S., and Schmidheiny, K. (2015). Taxes in cities. In Duranton, G., Henderson, J. V., and Strange, W. C., editors, *Handbook of Regional and Urban Economics*, volume 5, pages 1123–1196. Elsevier - North Holland, Amsterdam.

De Bartolome, C. A. and Ross, S. L. (2003). Equilibria with local governments and commuting: income sorting vs. income mixing. *Journal of Urban Economics*, 54(1):1–20.

Duranton, G. and Puga, D. (2015). Urban land use. In Duranton, G., Henderson, J. V., and Strange, W. C., editors, *Handbook of Regional and Urban Economics*, volume 5, pages 467–560. Elsevier - North Holland, Amsterdam.

EUROSTAT (2018). *Methodological manual on territorial typologies*. Office of the European Union, Luxembourg.

Fretz, S., Parchet, R., and Robert-Nicoud, F. (2022). Highways, Market Access and Spatial Sorting. *The Economic Journal*, 132(643):1011–1036.

Gaigné, C., Koster, H. R., Moizeau, F., and Thisse, J.-F. (2022). Who lives where in the city? Amenities, commuting and income sorting. *Journal of Urban Economics*, 128:103394.

Heblich, S., Redding, S. J., and Sturm, D. M. (2020). The making of the modern metropolis: evidence from London. *The Quarterly Journal of Economics*, 135(4):2059–2133.

Hensher, D. A., Rose, J. M., Rose, J. M., and Greene, W. H. (2005). *Applied choice analysis: a primer*. Cambridge university press, Cambridge.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*. Academic Press, New-York.

McFadden, D. (1978). Modeling the choice of residential location. In Karlqvist, A., editor, *Spatial Interaction Theory and Residential Location*. North-Holland, Amsterdam.

Mulalic, I. and Rouwendal, J. (2020). Does improving public transport decrease car ownership? evidence from a residential sorting model for the copenhagen metropolitan area. *Regional Science and Urban Economics*, 83:103543.

Muth, R. F. (1969). *Cities and housing: the spatial pattern of urban residential land use*. University of Chicago Press, Chicago.

OFS (2017). *Typologie des communes et typologie urbain-rural 2012*. Département Fédéral de l'Intérieur, Office fédéral de la Statistique, Neuchâtel.

OFS (2019). *Bassins d'emploi 2018*. Département Fédéral de l'Intérieur, Office fédéral de la Statistique, Neuchâtel.

Roberts, J., Hodgson, R., and Dolan, P. (2011). "It's driving her mad": gender differences in the effects of commuting on psychological health. *Journal of Health Economics*, 30(5):1064–1076.

Schaltegger, C. A., Somogyi, F., and Sturm, J.-E. (2011). Tax competition and income sorting: evidence from the Zurich metropolitan area. *European Journal of Political Economy*, 27(3):455–470.

Schmidheiny, K. (2006). Income segregation and local progressive taxation: empirical evidence from Switzerland. *Journal of Public Economics*, 90(3):429–458.

Schmidheiny, K. and Slotwinski, M. (2018). Tax-induced mobility: evidence from a foreigners' tax scheme in Switzerland. *Journal of Public Economics*, 167:293–324.

Straumann, T. (2012). Die Wirtschaft im 20. Jahrhundert. In des Kantons Schwyz, H. V., editor, *Die Geschichte des Kantons Schwyz, Band 5: Wirtschaft und Gesellschaft (1712 bis 2010)*. Chronos Verlag, Zürich.

Su, Y. (2022). The rising value of time and the origin of urban gentrification. *American Economic Journal: Economic Policy*, 14(1):402–39.

Appendix

A. Institutional and urban setting

Figure A1: Monocentricity assumption

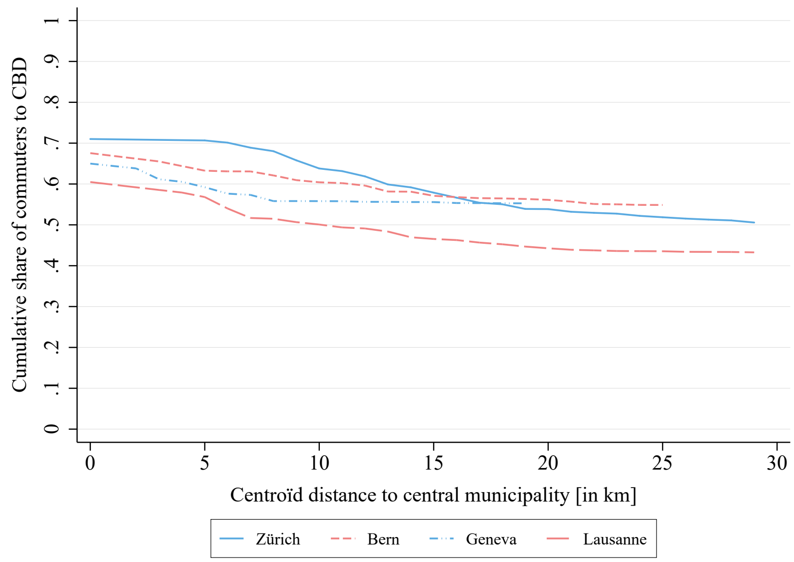

Note: This figure shows the cumulative share of commuters to the CBD within a centroid distance radius for the agglomerations of Zürich, Bern, Geneva and Lausanne.

Figure A2: Share of commuters to CBD per type of municipalities

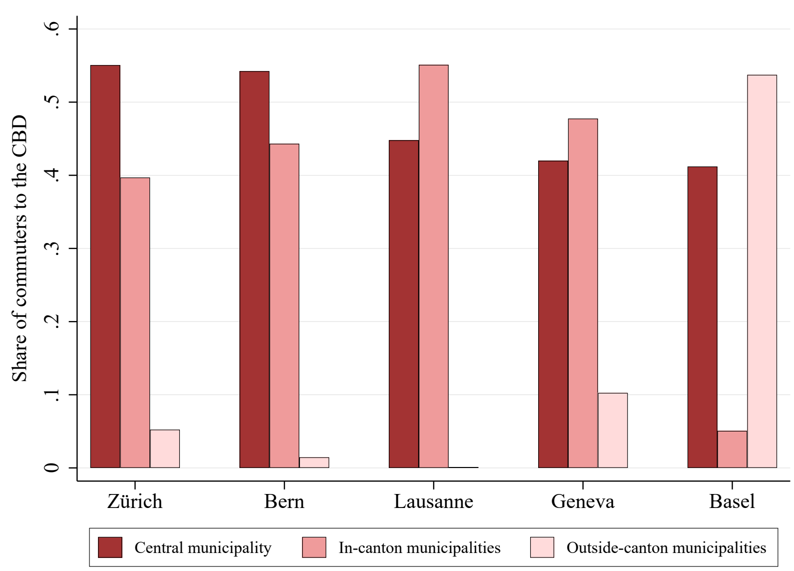

Note: This figure displays the share of commuters living in (i) the central municipality, (ii) the same canton as the central municipality and (iii) a out-of-canton municipality, for each of the five largest agglomerations of Switzerland.

Figure A3: Intra-cantonal tax scheme, single individual, no children

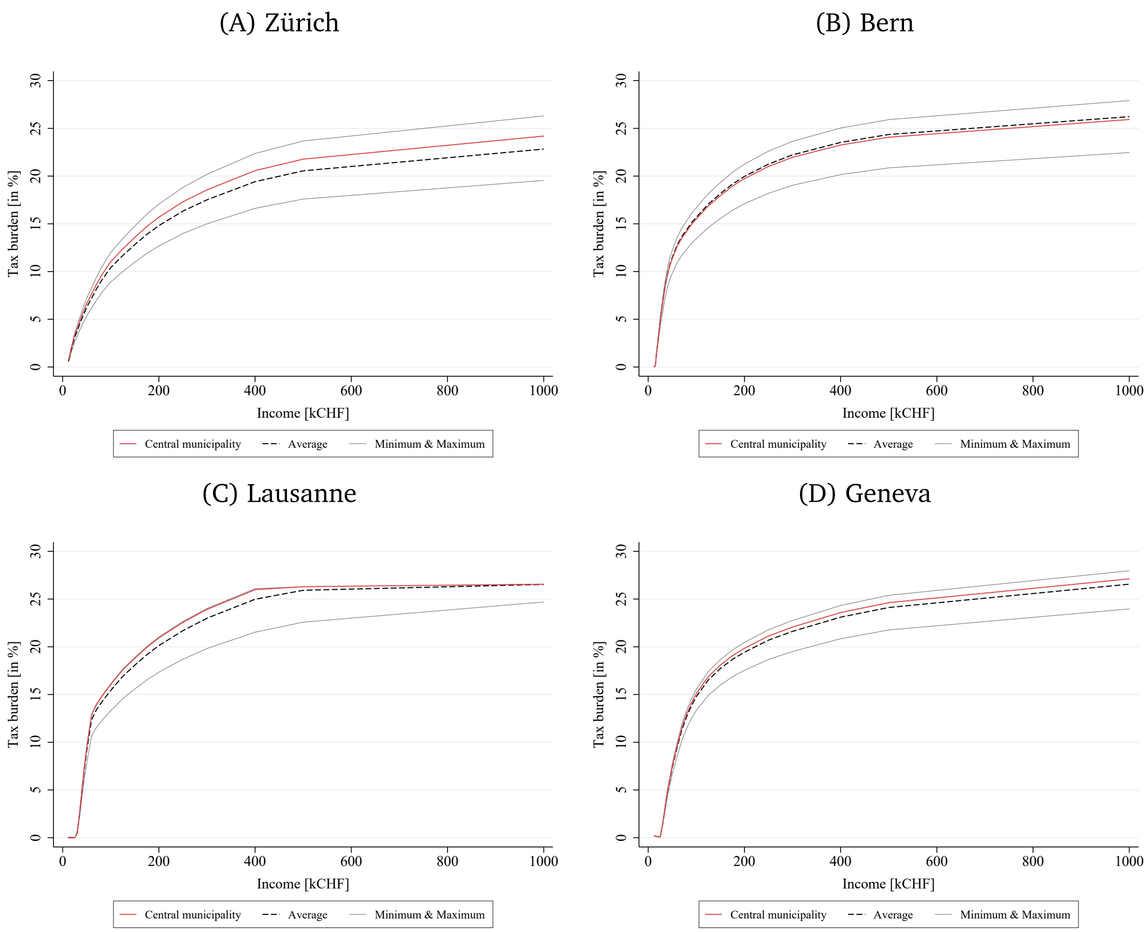

Note: This figure represents the progressive tax scheme faced by an household composed of a single individual, for each of the four intra-cantonal agglomerations of study. The minimum, maximum and average tax burden for each level of gross income as well as the tax burden in the central municipality are reported.

Figure A4: Intra-cantonal tax scheme, married couple, no children

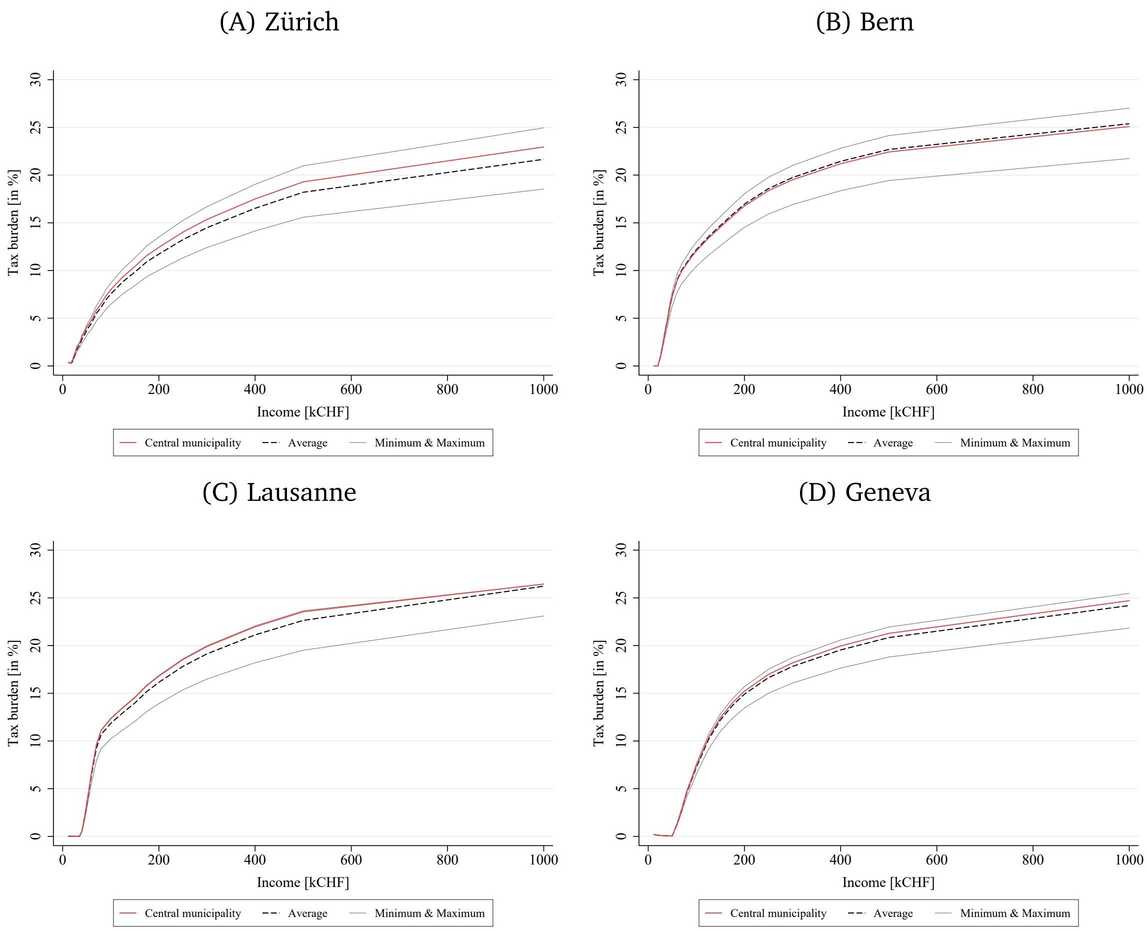

Note: This figure represents the progressive tax scheme faced by an household composed of a a married couple without children, for each of the four intra-cantonal agglomerations of study. The minimum, maximum and average tax burden for each level of gross income as well as the tax burden in the central municipality are reported.

Figure A5: Intra-cantonal tax scheme, married couple, two children

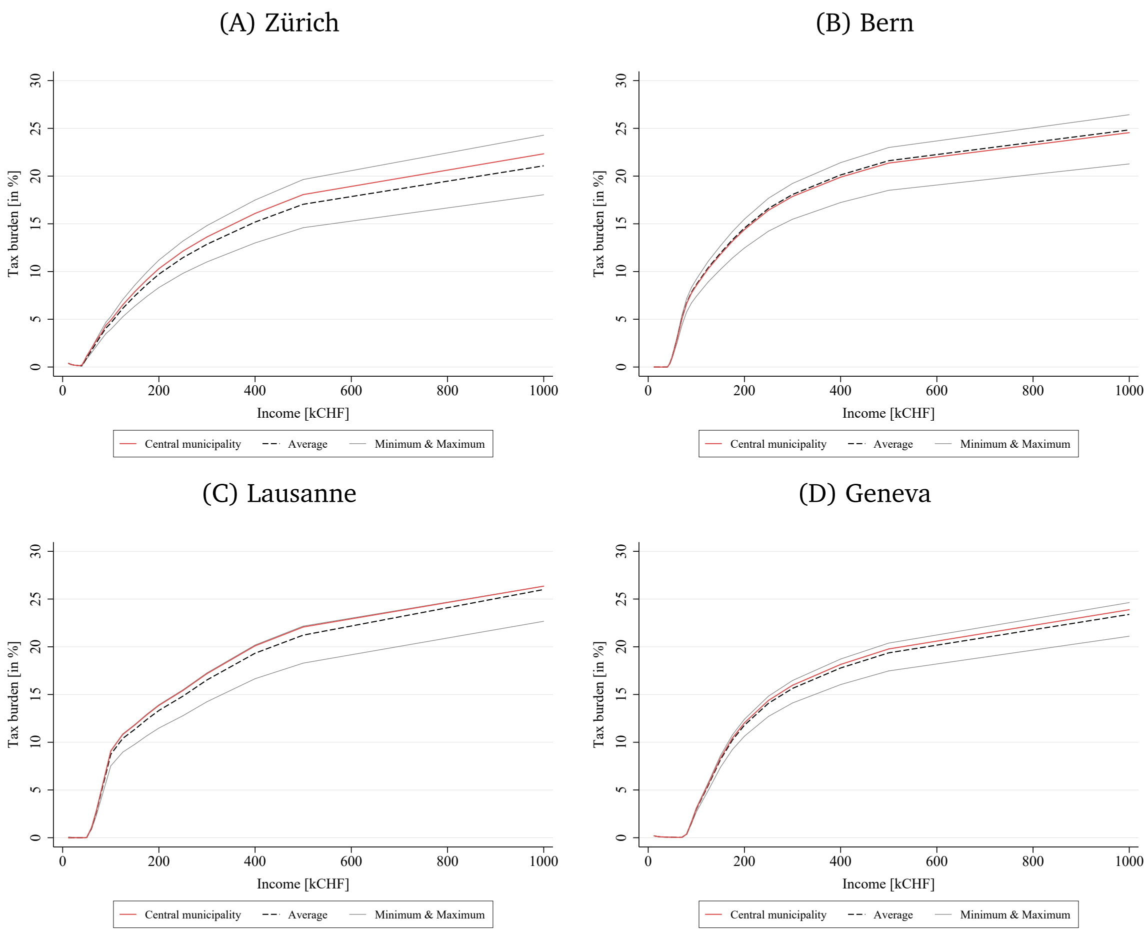

Note: This figure represents the progressive tax scheme faced by an household composed of a a married couple with two children, for each of the four intra-cantonal agglomerations of study. The minimum, maximum and average tax burden for each level of gross income as well as the tax burden in the central municipality are reported.

B. A model of residential location decisions in a fiscally decentralized and monocentric urban area

Figure B1: Commuting as a disutility cost
Number of weekly hours worked & weekly number of commutes

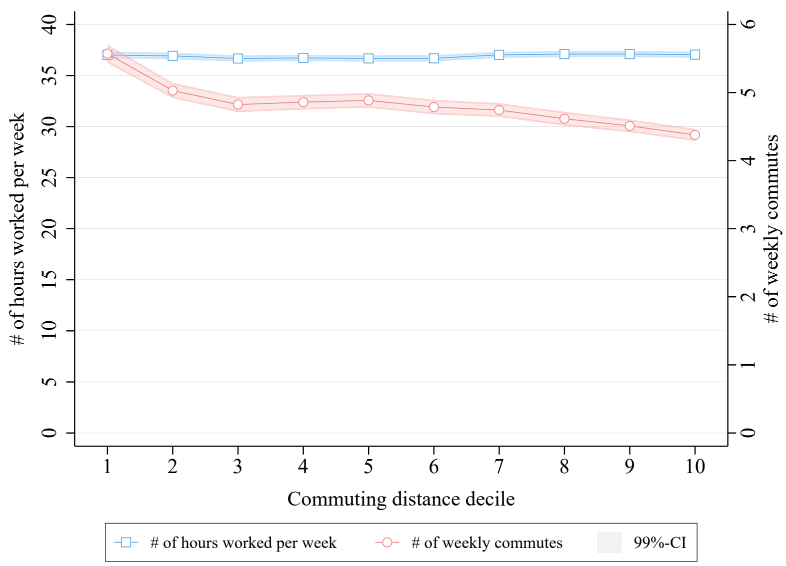

Note: This figure plots the average number of hours worked by the commuters to the central municipalities, per decile of commuting distance. The average number of weekly commutes is also reported. 99%-confidence intervals are reported.

Location choice model with monetary commuting costs

Our baseline model is a model of location decision with decentralized progressive income taxation, non-homothetic preferences towards housing and a commuting disutility. We show that our predictions qualitatively hold when specifying the commuting as monetary cost entering the consumers problem through the budget constraint. In the spirit of Gaigné et al. (2022), we allow the monetary commuting cost $\kappa\delta_j W$ to increase with income. The optimization program hence becomes

$$\max_{C_i,H_i} \quad U_{ij} = z_{ij} B_j \Big(C\Big)^{1-\beta} \Big(H - \upsilon_H\Big)^{\beta} \Big(G\Big)^{\gamma} \tag{B.1}$$
$$\text{s.t. } C_i + Q_j H_i = \big[1 - \tau_j r(W_i)\big] W_i - \kappa\delta_j W_i.$$

From this we derive the marshallian demand functions:

$$
\begin{aligned}
H_i^* &= \frac{\beta\big[(1-\tau_j r(W_i)-\kappa\delta_j)W_i - Q_j\upsilon_H\big]}{Q_j} + \upsilon_H \\
C_i^* &= (1-\beta)\big[1-\tau_j r(W_i)-\kappa\delta_j\big]W_i - (1-\beta)Q_j\upsilon_H,
\end{aligned}
$$

which give the indirect utility function

$$
\begin{aligned}
V_{i|j} = \Omega + \ln(z_{ij}) + \ln(B_j) + \ln\big[(1-\tau_j r(W_i)-\kappa\delta_j)W_i - \upsilon_H Q_j\big] \\
- \beta\ln(Q_j) + \gamma\ln(G).
\end{aligned} \tag{B.2}
$$

The marginal rates of substitution between community characteristics are given by

$$
MRS_{Q,\tau} = \frac{dQ_j}{d\tau_j}\bigg|_{dV=0} = -\frac{W_i}{H_i^*} < 0 \tag{B.3a}
$$

$$
MRS_{Q,\delta} = \frac{dQ_j}{d\delta_j}\bigg|_{dV=0} = -\frac{\kappa W_i}{H_i^*} < 0 \tag{B.3b}
$$

$$
MRS_{\tau,\delta} = \frac{d\tau_j}{d\delta_j}\bigg|_{dV=0} = -\frac{\kappa}{r(W_i)} < 0. \tag{B.3c}
$$

Once again, the three different jurisdictional characteristics are undesirable and thus exhibit negative marginal rates of substitution. Taking derivatives of the latter with respect to income we obtain

$$
\frac{\partial MRS_{Q,\tau}}{\partial W_i} = -\frac{\Big[(1-\beta)Q_j\upsilon_H + \beta\tau_j r'(W_i)W_i^2\Big]}{Q_j(H_i^*)^2} < 0 \tag{B.4a}
$$

$$
\frac{\partial MRS_{Q,\delta}}{\partial W_i} = -\frac{\kappa\Big[(1-\beta)Q_j\upsilon_H + \beta\tau_j r'(W_i)W_i^2\Big]}{Q_j(H_i^*)^2} < 0 \tag{B.4b}
$$

$$
\frac{\partial MRS_{\tau,\delta}}{\partial W_i} = \frac{\kappa r'(W_i)}{(r(W_i))^2} > 0. \tag{B.4c}
$$

Given equations (B.4a) to (B.4c), agents with higher incomes are more sensitive to taxes, less sensitive to housing prices and ambiguous towards commuting distance. High-income individuals are hence more likely to locate in low-tax, high housing-price jurisdictions. The prediction concerning commuting remains unclear, just like in our baseline model.

Indifference curves under proportional and progressive taxation

Figure B2 illustrates how proportional and progressive taxation impact indifference curves in the $(\delta_j;\tau_j)$ plan for two income levels $W^2 > W^1$. In the proportional case, described in panel (a), richer agents are more sensitive to commuting distance than taxes, as described in

Figure B2: Indifference curves, proportional vs. progressive tax schemes

(A) Proportional taxation

(B) Progressive taxation

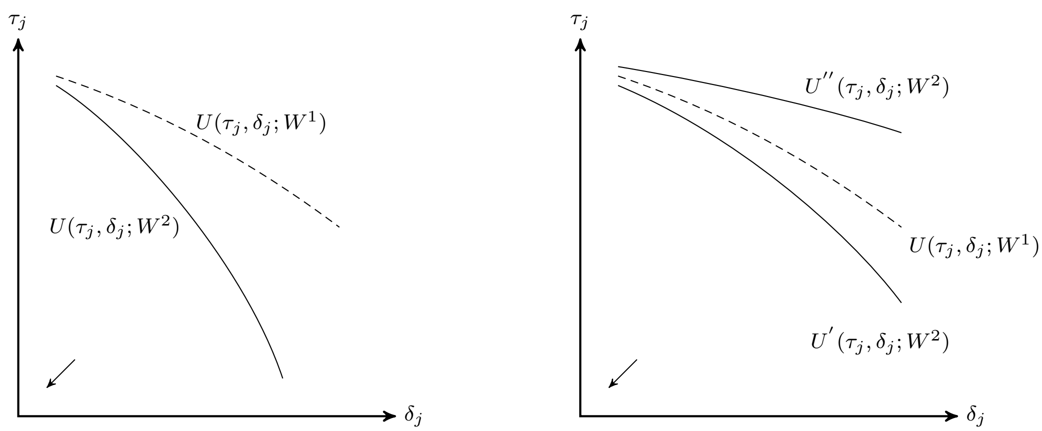

Note: This figure depicts the indifference curves in the $(\delta_j; \tau_j)$ plan for individuals with different levels of income $W_2 > W_1$. Panel (a) depicts the case of a proportional tax system, while panel (b) represents the case of a progressive tax system. Dashed lines represent the indifference curves of the low-income type W_1 while plain lines represent indifference curves of the high-income type W_2. In panel (b), two indifference curves with level of income W_2 are represented: $U'(\cdot)$ represents the case where Equation (3.6) does not hold and $U''(\cdot)$ represents the case where Equation (3.6) holds.

equation (3.7c). Consequently, in the plan $(\delta_j; \tau_j)$, high-income individuals of type W^2 have steeper indifference curves. In the case of progressive taxation, two alternative cases may occur. In one case, represented by $U'(\cdot)$, rich individuals are more sensitive to commuting than to taxes, just like in panel (a). In the other case, depicted by $U''(\cdot)$, rich individuals are less sensitive to commuting, making their indifference curve flatter than the indifference curve of poorer individuals. The key element that actually makes this case possible is tax progressivity; as shown in equation (3.5b), the term $-W_i^2 r'(W_i)$, which is directly related to the progressivity of taxation, might revert the sign of the derivative, eventually making rich individuals more sensitive to taxes relative to commuting. This is typically the case when the (sufficient) condition described in equation (3.6) holds.

C. Data and descriptive evidence

Table C1: Data structure of the analysis sample

	2010	2011	2012	2013	2014	2015	2016	Total
Zürich	1043	863	889	778	641	490	293	4997
Bern	502	465	480	439	168	140	124	2318
Lausanne	381	402	379	297	247	184	114	2004
Geneva	362	335	328	273	208	132	74	1712
Total	2288	2065	2076	1787	1264	946	605	11031

Note: This table reports the frequency of observations per agglomeration and moving year in our sample of analysis.

Table C2: Tax burden differentials and multipliers per income decile

Income decile	Abs. ΔTax Burden [in pp]		Rel. ΔTax Burden [in %]		Normalized tax multiplier	
1st decile	-0.05	(0.00)	-2.30	(0.15)	105.79	(0.30)
2^{d} decile	-0.13	(0.01)	-2.16	(0.12)	105.03	(0.30)
3^{d} decile	-0.20	(0.01)	-2.29	(0.12)	104.37	(0.30)
4th decile	-0.25	(0.01)	-2.54	(0.13)	103.81	(0.31)
5th decile	-0.28	(0.01)	-2.67	(0.13)	103.54	(0.32)
6th decile	-0.33	(0.02)	-3.00	(0.14)	102.54	(0.33)
7th decile	-0.37	(0.02)	-3.23	(0.14)	101.97	(0.34)
8th decile	-0.43	(0.02)	-3.61	(0.15)	101.16	(0.34)
9th decile	-0.56	(0.02)	-4.25	(0.16)	99.48	(0.35)
10th decile	-0.98	(0.03)	-5.85	(0.19)	95.36	(0.39)
Total	-0.36	(0.01)	-3.19	(0.05)	102.31	(0.11)

Note: This table reports the average tax burden differential (ΔTax Burden) and relative tax burden differential (ΔTax Burden/ Tax Burden Workplace) between the residence and workplace, per income decile. The average (normalized) tax multiplier of the residence place per income decile is also reported. Standard errors are reported in parentheses.

Figure C1: Municipal characteristics: taxes, housing prices and commuting distance

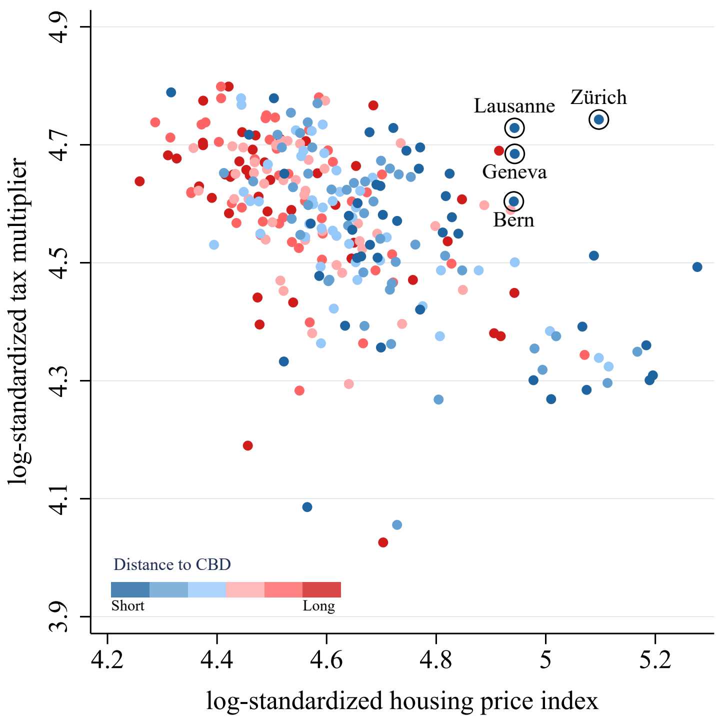

Note: This figure describes the relationship between the three municipal characteristics of interest, measured by the normalized tax multiplier, normalized housing prices and the distance to agglomerations' central municipalities.

Figure C2: Average commuting time per income decile

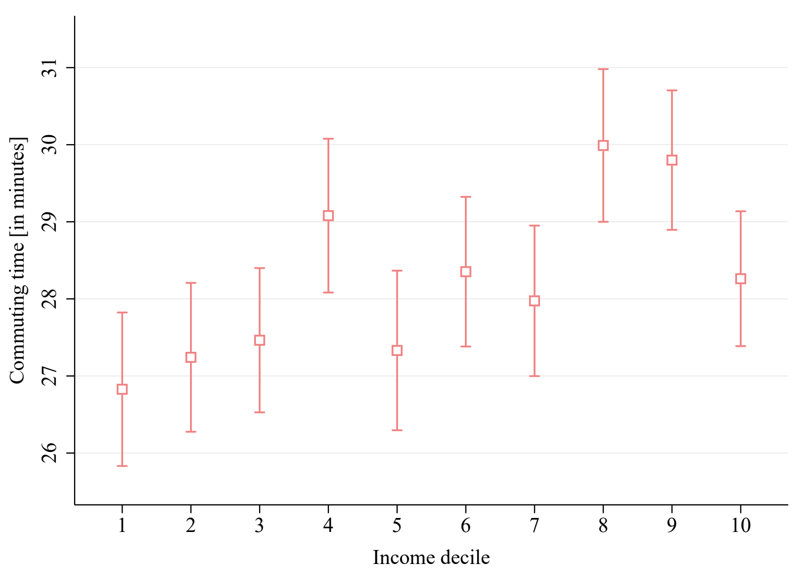

Note: This figure shows the average (self-reported) commuting time per decile of income. 95 % confidence intervals for the average are reported.

D. Empirical analysis

Figure D1: Share of in-moving households over the total population

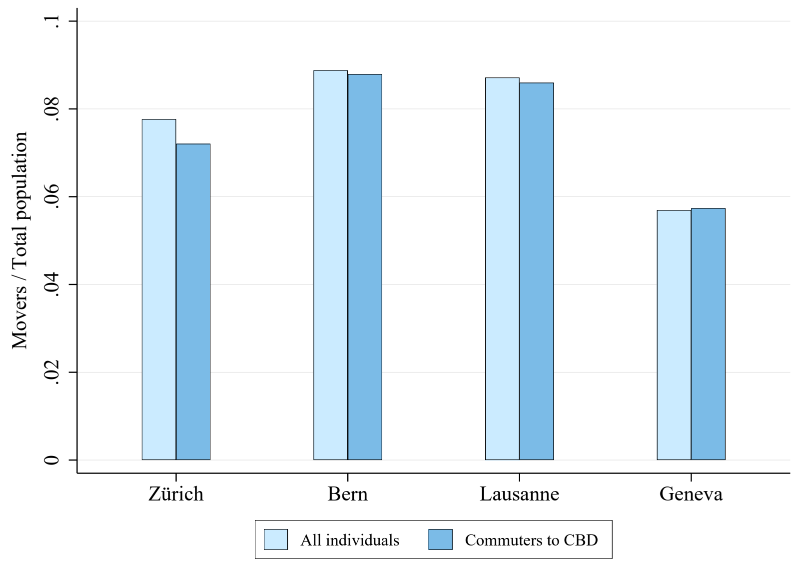

Note: This figure reports the ratio between (i) individuals moving into and (ii) the overall population living in each agglomeration. The populations considered in the computation of the ratio is either the entire population of the agglomerations (*All individuals*) or the sub-population of commuters to the central municipalities (*Commuters to CBD*).

Figure D2: Relative change in municipal multipliers

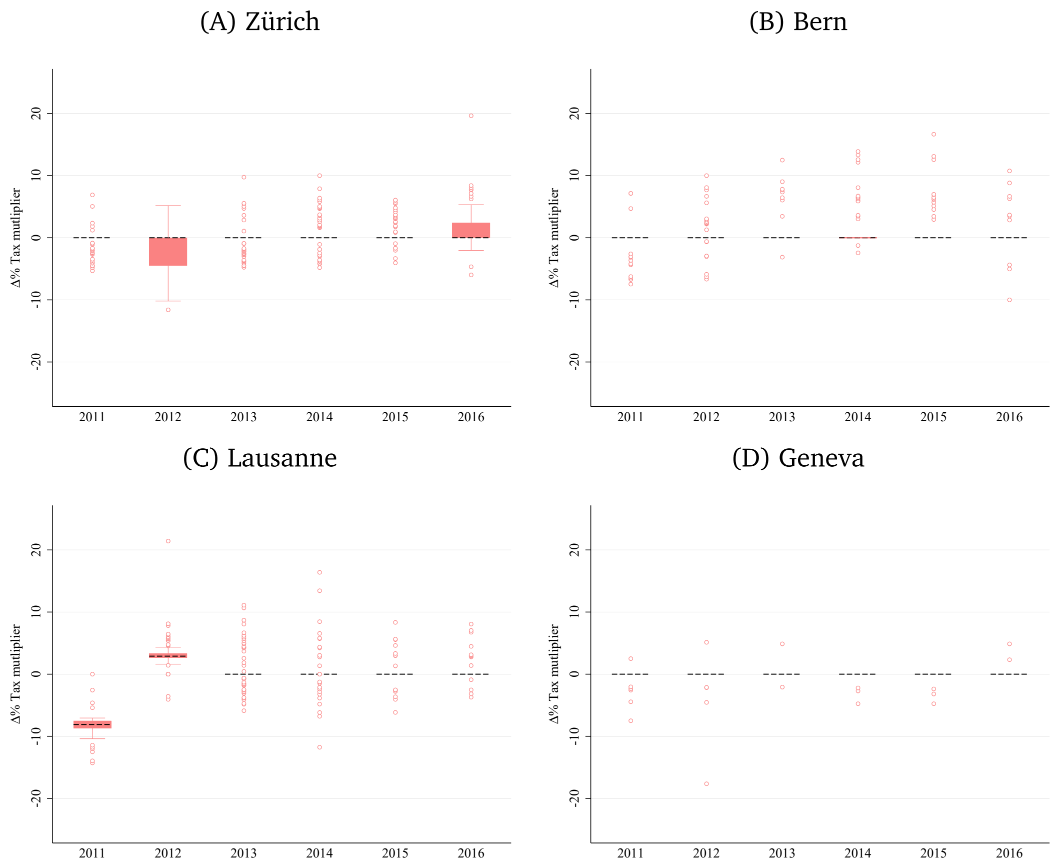

Note: This figure plots the distribution of the relative change in municipal tax multipliers, for each year of observation.

Figure D3: Housing prices growth rates

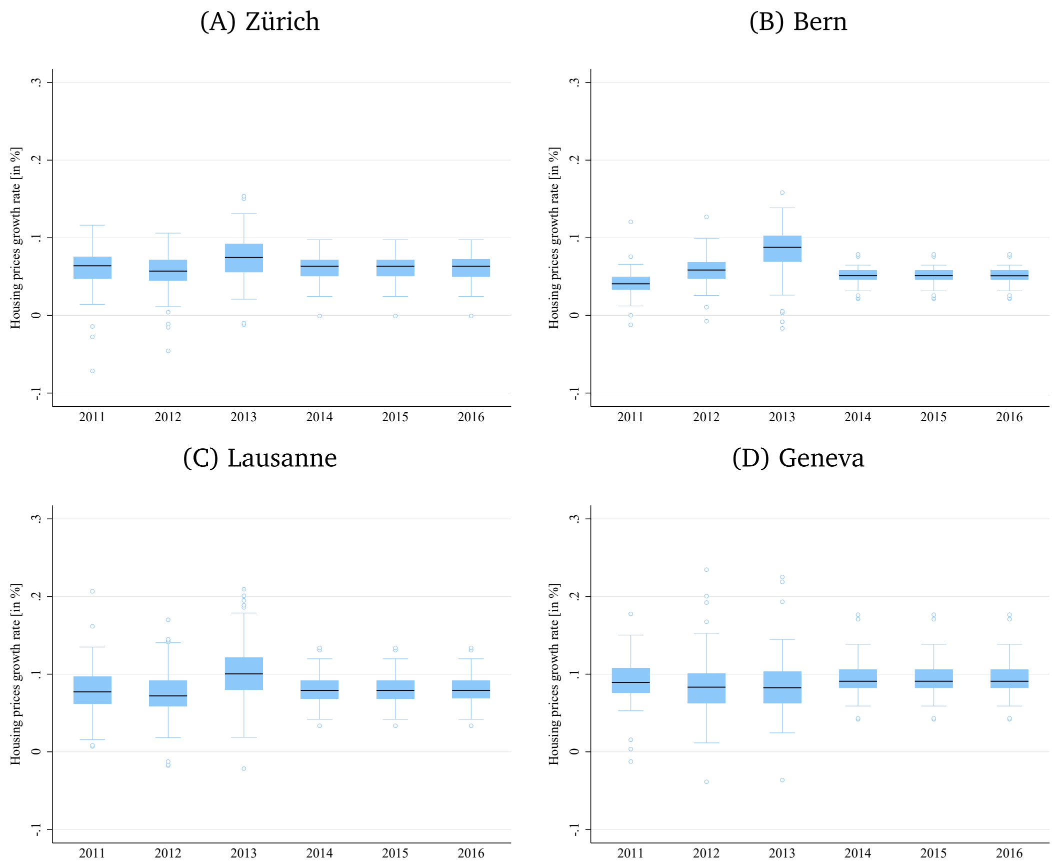

Note: This figure plots the distribution of the municipal-level housing prices growth rate, for each year of observation.

Figure D4: Relationship between municipal-level commuting measures

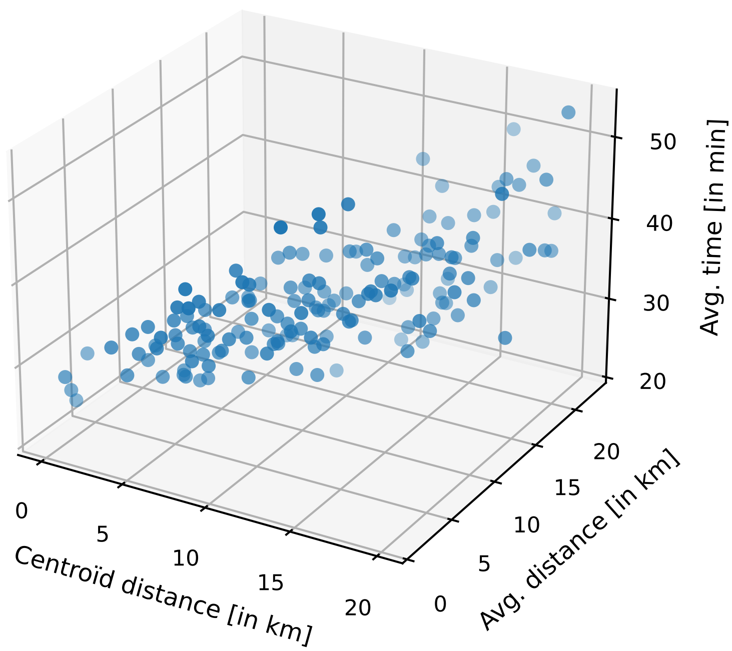

Note: This figure describes graphically the relationship between the three municipal-level commuting measures δ_j: *Centroid distance*, *Average commuting distance* and *Average commuting time*.

Figure D5: Main mode of transportation by income decile

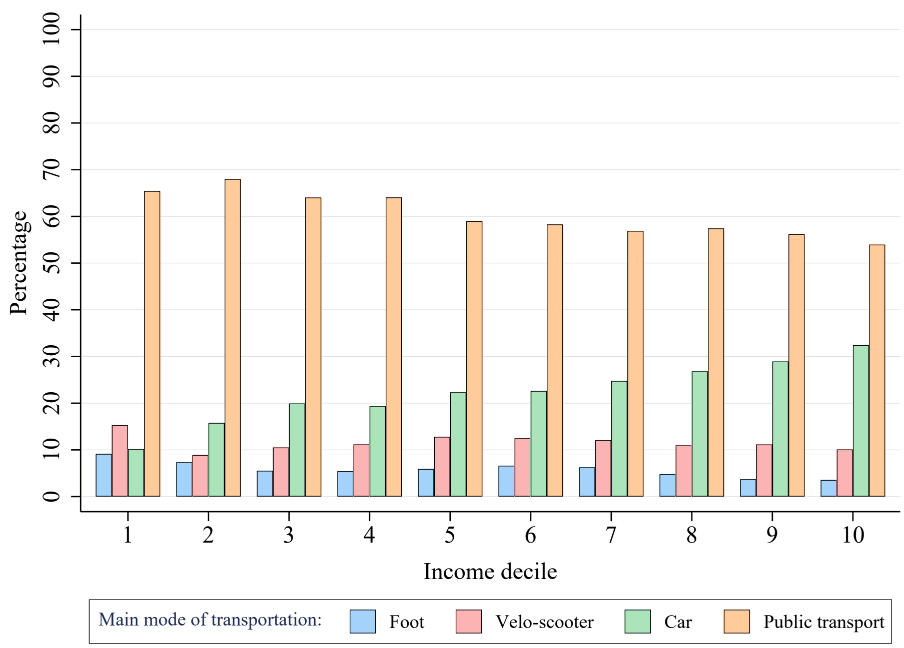

Note: This figure depicts the share of commuters traveling by foot, velo-scooter, car or public transport for each decile of the income distribution.

Figure D6: Theoretical effect of taxes on commuting disutiility

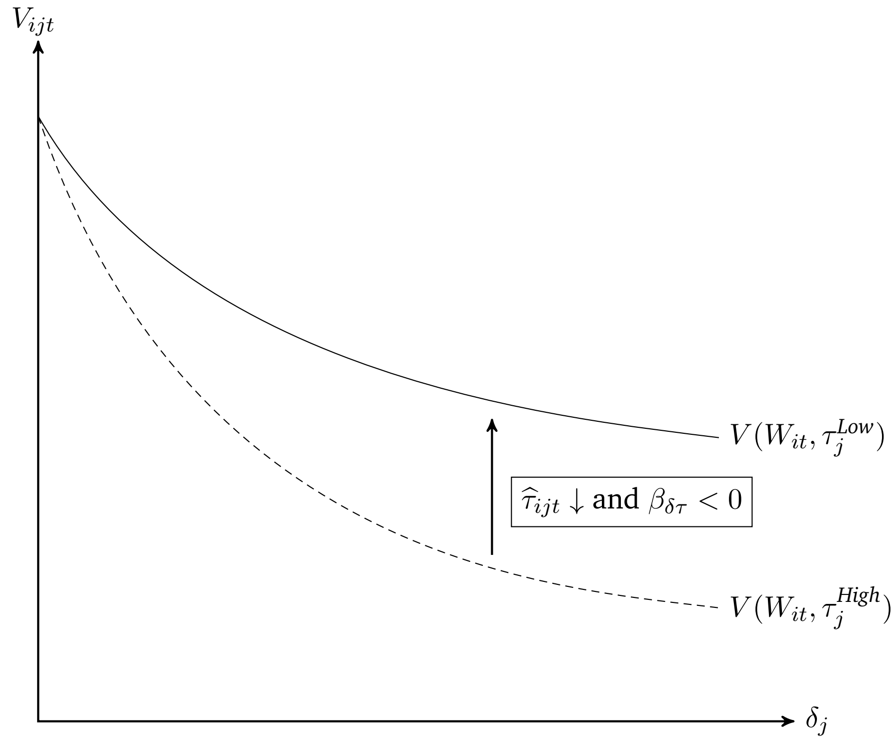

Note: This figure shows the impact of taxes on the willingness to commute, as measured in our empirical framework. The horizontal axis represents commuting distance or time, while the vertical axis indirect utility.

Table D1: Location choices and income sorting, squared commuting measures

Measure of δ_j	Centroid distance		Avg. commut. distance		Avg. commut. time	
	(1)	(2)	(3)	(4)	(5)	(6)
Location Choice Model						
Municipal characteristics						
$\widehat{\tau}_{ijt}(W_{it})$	-0.469***	-0.391***	-0.558***	-0.501***	-0.615***	-0.540***
	(0.058)	(0.059)	(0.056)	(0.057)	(0.055)	(0.056)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	1.222***	1.004***	0.943***	0.832***	0.806***	0.746***
	(0.144)	(0.167)	(0.125)	(0.158)	(0.118)	(0.143)
$\delta_j \times \widehat{\ln(W_{it})}$	0.121***	0.125***	0.159***	0.136***	0.107***	0.097**
	(0.014)	(0.015)	(0.021)	(0.024)	(0.035)	(0.039)
$\delta_j^2 \times \widehat{\ln(W_{it})}$	-0.005***	-0.005***	-0.005***	-0.005***	-0.001*	-0.001*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
Other municipal amenities B_j						
$\ln(\text{Housing}_{jt}) \times \widehat{\ln(W_{it})}$		-0.723***		-0.767***		-0.643***
		(0.197)		(0.201)		(0.201)
$\ln(\text{Foreigners}_{jt}) \times \widehat{\ln(W_{it})}$		-0.381***		-0.319***		-0.389***
		(0.079)		(0.086)		(0.078)
$\ln(\text{Green Space}_{jt}) \times \widehat{\ln(W_{it})}$		0.112***		0.085**		0.067*
		(0.037)		(0.036)		(0.036)
$\mathbb{I}(\text{Lake}_{jt} = 1) \times \widehat{\ln(W_{it})}$		0.078		-0.063		-0.031
		(0.060)		(0.057)		(0.057)
Average ASC						
Central mun. $\hat{\alpha}_j$	0.000	0.000	0.000	0.000	0.000	0.000
Average suburban $\hat{\alpha}_j$	-6.007	-5.976	-6.071	-6.045	-6.109	-6.065
Indiv.-specific controls	Yes	Yes	Yes	Yes	Yes	Yes
Nbr. municipalities	166	166	166	166	166	166
Nbr. individuals	11031	11031	11031	11031	11031	11031
Log-likelihood	-30621.484	-30582.658	-30632.764	-30602.325	-30638.994	-30605.655

Note: This table reports the pooled estimates of equation (5.1) with disjoint choice sets across agglomerations, where we include an additional interaction term between individuals' income and squared municipal commuting measure: $\delta_j^2 \times \widehat{\ln(W_{it})}$. The error term η_{ijt} is assumed to follow a Type I GEV distribution. In each specification, the alternative-specific constants of the central municipalities are standardized to 0. The three sets of estimates in columns (1)-(2), (3)-(4) and (5)-(6) are based on different measures of commuting δ_j between municipality j and the center of the agglomeration (*Centroid distance*, *Average commuting distance* and *Average commuting time*). Columns (1)-(3)-(5) report the estimation results for the baseline specification with squared commuting, while columns (2)-(4)-(6) introduce additional interactions between income and municipal amenities. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table D2: Location choice and income sorting, restricted samples

Measure of δ_j	Centroid distance			Avg. commuting distance			Avg. commuting time		
	All	Outside urb.	Outside cant.	All	Outside urb.	Outside cant.	All	Outside urb.	Outside cant.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\ln(\tilde{\tau}_j) \times \widehat{\ln(W_i)}$	-0.620***	-0.354***	-0.263*	-0.673***	-0.404***	-0.333**	-0.637***	-0.396***	-0.297**
	(0.054)	(0.118)	(0.144)	(0.053)	(0.113)	(0.141)	(0.053)	(0.114)	(0.144)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	0.829***	1.032***	1.049***	0.627***	0.785***	0.655**	0.744***	0.746***	0.785**
	(0.130)	(0.330)	(0.406)	(0.110)	(0.268)	(0.323)	(0.111)	(0.264)	(0.334)
$\delta_j \times \widehat{\ln(W_{it})}$	0.037***	0.031***	0.030**	0.038***	0.030**	0.020	0.044***	0.026**	0.027
	(0.005)	(0.012)	(0.015)	(0.006)	(0.013)	(0.018)	(0.005)	(0.012)	(0.017)
ASC									
Central mun. $\hat{\alpha}_j$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Average suburban $\hat{\alpha}_j^S$	-6.128	-6.128	-5.827	-6.167	-6.160	-5.873	-6.129	-6.151	-5.844
Indiv.-specific controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Nbr. municipalities	166	93	68	166	93	68	166	93	68
Nbr. individuals	11031	3165	2204	11031	3165	2204	11031	3165	2204
Log-likelihood	-30645.978	-5303.071	-3093.370	-30650.810	-5304.122	-3094.712	-30641.053	-5304.607	-3094.243

Note: This table reports the pooled estimates of equation (5.1) with disjoint choice sets across agglomerations. The error term η_{ijt} is assumed to follow a Type I GEV distribution. In each specification, the alternative-specific constants of the central municipalities are standardized to 0. Different measures of commuting δ_j between municipality j and the center of the agglomeration are considered: (*Centroid distance*, *Average commuting distance* and *Average commuting time*). Columns (1)-(4)-(7) report the estimation results for the sample of the *All movers*, columns (2)-(5)-(8) results for the sample of the *Outside urban area* movers, and columns (3)-(6)-(9) results for the sample of the *Outside canton* movers. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table D3: Location choice and income sorting per agglomeration
Nests based on typology OFS 9

Measure of δ_j	Centroid distance		Avg. commut. distance		Avg. commut. time	
	C. Logit	N. Logit	C. Logit	N. Logit	C. Logit	N. Logit
A. Zürich						
$\widehat{\tau_{ijt}}$	-0.641***	-1.248***	-0.735***	-1.321***	-0.630***	-1.185***
	(0.076)	(0.264)	(0.072)	(0.284)	(0.074)	(0.242)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	1.110***	2.339***	0.751***	1.424***	0.969***	2.125***
	(0.180)	(0.671)	(0.148)	(0.419)	(0.151)	(0.614)
$\delta_j \times \widehat{\ln(W_{it})}$	0.048***	0.106***	0.044***	0.087***	0.060***	0.125***
	(0.007)	(0.031)	(0.008)	(0.026)	(0.008)	(0.033)
$\lambda^{\text{High density}}$		2.138***		1.930***		2.035***
		(0.469)		(0.424)		(0.423)
$\lambda^{\text{Low density}}$		1.599***		1.363***		2.183***
		(0.510)		(0.433)		(0.745)
Nbr. municipalities	70	70	70	70	70	70
Nbr. Individuals	4997	4997	4997	4997	4997	4997
Log-likelihood	-14970.536	-14965.827	-14976.091	-14972.092	-14962.948	-14959.032
B. Bern						
$\widehat{\tau_{ijt}}$	-0.469***	-2.796**	-0.515***	-3.631**	-0.518***	-3.032**
	(0.164)	(1.376)	(0.162)	(1.825)	(0.162)	(1.305)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	1.870***	8.857**	1.527***	9.198**	1.390***	7.386**
	(0.542)	(3.668)	(0.516)	(4.478)	(0.467)	(3.050)
$\delta_j \times \widehat{\ln(W_{it})}$	0.055***	0.242**	0.059***	0.337**	0.050***	0.309***
	(0.015)	(0.096)	(0.019)	(0.156)	(0.015)	(0.120)
$\lambda^{\text{High density}}$		5.004***		6.052**		5.435***
		(1.861)		(2.877)		(1.902)
$\lambda^{\text{Low density}}$		6.300		10.158		10.559
		(4.470)		(8.714)		(7.218)
Nbr. municipalities	31	31	31	31	31	31
Nbr. Individuals	2318	2318	2318	2318	2318	2318
Log-likelihood	-5676.920	-5668.904	-5678.857	-5670.309	-5678.393	-5668.624
C. Lausanne						
$\widehat{\tau_{ijt}}$	-0.403***	-1.387***	-0.420***	-1.437***	-0.472***	-1.641***
	(0.111)	(0.481)	(0.106)	(0.479)	(0.108)	(0.525)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	0.638*	2.406	0.668**	2.547	0.365	1.306
	(0.335)	(1.552)	(0.338)	(1.619)	(0.317)	(1.476)
$\delta_j \times \widehat{\ln(W_{it})}$	0.020*	0.083	0.026*	0.095	0.008	0.075
	(0.012)	(0.052)	(0.014)	(0.058)	(0.017)	(0.090)
$\lambda^{\text{High density}}$		4.328***		4.153***		4.660**
		(1.670)		(1.585)		(2.150)
$\lambda^{\text{Low density}}$		3.782		3.772		3.869
		(2.398)		(2.318)		(3.108)
Nbr. municipalities	38	38	38	38	38	38
Nbr. Individuals	2004	2004	2004	2004	2004	2004
Log-likelihood	-5579.750	-5575.019	-5579.584	-5575.089	-5581.099	-5576.656
D. Geneva						
$\widehat{\tau_{ijt}}$	-0.962***	-0.945**	-0.983***	-0.751**	-1.036***	-1.807***
	(0.142)	(0.464)	(0.140)	(0.365)	(0.138)	(0.649)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	0.253	0.245	0.141	0.117	0.037	0.355
	(0.283)	(0.293)	(0.252)	(0.195)	(0.275)	(0.611)
$\delta_j \times \widehat{\ln(W_{it})}$	0.048***	0.047**	0.062***	0.055**	0.022*	0.051*
	(0.018)	(0.021)	(0.022)	(0.023)	(0.013)	(0.030)
$\lambda^{\text{High density}}$		0.983**		0.760**		1.879**
		(0.484)		(0.370)		(0.738)
$\lambda^{\text{Low density}}$		0.925		0.650		24.518
		(0.668)		(0.409)		(43.367)
Nbr. municipalities	27	27	27	27	27	27
Nbr. Individuals	1712	1712	1712	1712	1712	1712
Log-likelihood	-4400.996	-4400.990	-4400.668	-4400.442	-4403.300	-4400.905

Note: This table reports the estimates of equation (5.1) per agglomeration. Columns (1)-(3)-(5) report the conditional logit estimates for different measures of municipal-level commuting δ_j, while columns (2)-(4)-(6) report nested logit estimates, where the nests are defined based on the density of the municipalities. Interaction coefficients between income and municipal characteristics, as well as the estimated dissimilarity parameters for the nested specification are reported. All estimated equations include individual-specific variables. Significance stars indicate the following significance levels: * 0.1, **0.05 and *** 0.01.

Table D4: Location choice and income sorting per agglomeration
Nests based on DEGURBA

Measure of δ_j	Centroid distance		Avg. commut. distance		Avg. commut. time	
	C. Logit	N. Logit	C. Logit	N. Logit	C. Logit	N. Logit
A. Zürich						
$\widehat{\tau_{ijt}}$	-0.641***	-0.557***	-0.735***	-0.603***	-0.630***	-0.509***
	(0.076)	(0.139)	(0.072)	(0.148)	(0.074)	(0.131)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it}})$	1.110***	1.026***	0.751***	0.578***	0.969***	0.755***
	(0.180)	(0.338)	(0.148)	(0.202)	(0.151)	(0.239)
$\delta_j \times \widehat{\ln(W_{it}})$	0.048***	0.045***	0.044***	0.035***	0.060***	0.049***
	(0.007)	(0.014)	(0.008)	(0.011)	(0.008)	(0.013)
$\lambda^{\text{High density}}$		0.846***		0.800***		0.783***
		(0.210)		(0.201)		(0.197)
$\lambda^{\text{Low density}}$		0.964***		0.770***		0.798***
		(0.303)		(0.234)		(0.233)
Nbr. municipalities	70	70	70	70	70	70
Nbr. Individuals	4997	4997	4997	4997	4997	4997
Log-likelihood	-14970.536	-14969.995	-14976.091	-14975.597	-14962.948	-14962.407
B. Bern						
$\widehat{\tau_{ijt}}$	-0.469***	-6.988***	-0.515***	-6.959***	-0.518***	-6.794**
	(0.164)	(2.408)	(0.162)	(2.344)	(0.162)	(3.295)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it}})$	1.870***	1.984	1.527***	1.883	1.390***	1.095
	(0.542)	(3.504)	(0.516)	(3.443)	(0.467)	(4.212)
$\delta_j \times \widehat{\ln(W_{it}})$	0.055***	0.083	0.059***	0.107	0.050***	0.082
	(0.015)	(0.089)	(0.019)	(0.116)	(0.015)	(0.146)
$\lambda^{\text{High density}}$		11.499*		11.130*		11.705
		(6.041)		(5.701)		(8.813)
$\lambda^{\text{Low density}}$		33.575		32.843		34.871
		(22.813)		(21.373)		(31.775)
Nbr. municipalities	31	31	31	31	31	31
Nbr. Individuals	2318	2318	2318	2318	2318	2318
Log-likelihood	-5676.920	-5665.295	-5678.857	-5665.364	-5678.393	-5665.460
C. Lausanne						
$\widehat{\tau_{ijt}}$	-0.403***	-0.682**	-0.420***	-0.690**	-0.472***	-0.794***
	(0.111)	(0.274)	(0.106)	(0.271)	(0.108)	(0.251)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it}})$	0.638*	0.954	0.668**	1.032	0.365	0.489
	(0.335)	(0.791)	(0.338)	(0.807)	(0.317)	(0.606)
$\delta_j \times \widehat{\ln(W_{it}})$	0.020*	0.024	0.026*	0.030	0.008	0.011
	(0.012)	(0.022)	(0.014)	(0.026)	(0.017)	(0.029)
$\lambda^{\text{High density}}$		1.882**		1.841**		2.163**
		(0.883)		(0.877)		(0.886)
$\lambda^{\text{Low density}}$		1.584**		1.590**		1.421**
		(0.780)		(0.776)		(0.675)
Nbr. municipalities	38	38	38	38	38	38
Nbr. Individuals	2004	2004	2004	2004	2004	2004
Log-likelihood	-5579.750	-5579.054	-5579.584	-5578.942	-5581.099	-5579.714
D. Geneva						
$\widehat{\tau_{ijt}}$	-0.962***	-1.365***	-0.983***	-1.358***	-1.036***	-1.587***
	(0.142)	(0.300)	(0.140)	(0.315)	(0.138)	(0.301)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it}})$	0.253	0.351	0.141	0.158	0.037	0.195
	(0.283)	(0.453)	(0.252)	(0.386)	(0.275)	(0.507)
$\delta_j \times \widehat{\ln(W_{it}})$	0.048***	0.052**	0.062***	0.058**	0.022*	0.030
	(0.018)	(0.023)	(0.022)	(0.028)	(0.013)	(0.021)
$\lambda^{\text{High density}}$		1.779***		1.661***		1.997***
		(0.553)		(0.563)		(0.600)
$\lambda^{\text{Low density}}$		1.472***		1.412***		1.680***
		(0.422)		(0.422)		(0.453)
Nbr. municipalities	27	27	27	27	27	27
Nbr. Individuals	1712	1712	1712	1712	1712	1712
Log-likelihood	-4400.996	-4398.959	-4400.668	-4399.322	-4403.300	-4400.157

Note: This table reports the estimates of equation (5.1) per agglomeration. Columns (1)-(3)-(5) report the conditional logit estimates for different measures of municipal-level commuting δ_j, while columns (2)-(4)-(6) report nested logit estimates, where the nests are defined based on the density of the municipalities. Interaction coefficients between income and municipal characteristics, as well as the estimated dissimilarity parameters for the nested specification are reported. All estimated equations include individual-specific variables. Significance stars indicate the following significance levels: * 0.1, **0.05 and *** 0.01.

Table D5: Location choice and income sorting, no central municipalities

Measure of δ_j	Centroïd distance		Avg. commut. distance		Avg. commut. time	
	(1)	(2)	(3)	(4)	(5)	(6)
Municipal characteristics						
$\widehat{\tau_{ijt}}$	-0.620***	-0.693***	-0.673***	-0.682***	-0.637***	-0.703***
	(0.055)	(0.076)	(0.053)	(0.077)	(0.053)	(0.076)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	0.829***	1.252***	0.627***	1.333***	0.744***	1.221***
	(0.126)	(0.171)	(0.110)	(0.172)	(0.111)	(0.164)
$\delta_j \times \widehat{\ln(W_{it})}$	0.037***	0.014**	0.038***	0.024***	0.044***	0.019***
	(0.005)	(0.006)	(0.006)	(0.007)	(0.005)	(0.007)
Average ASC						
Central mun. $\hat{\alpha}_j$	0.000	∅	0.000	∅	0.000	∅
Average suburban $\hat{\alpha}_j$	-6.128	∅	-6.167	∅	-6.129	∅
Indiv.-specific controls	Yes	Yes	Yes	Yes	Yes	Yes
Nbr. municipalities	166	162	166	162	166	162
Nbr. individuals	11031	7047	11031	7047	11031	7047
Log-likelihood	-30645.978	-24058.132	-30650.810	-24053.908	-30641.053	-24056.380

Note: This table reports the pooled estimates of equation (5.1) with disjoint choice sets across agglomerations. The error term η_{ijt} is assumed to follow a Type I GEV distribution. In each specification, the alternative-specific constants of the central municipalities are standardized to 0. Different measures of commuting δ_j between municipality j and the center of the agglomeration are considered: (*Centroid distance*, *Average commuting distance* and *Average commuting time*). Columns (1)-(3)-(5) correspond to the baseline results, while columns (2)-(4)-(6) correspond to the model where central municipalities are removed from the choice sets. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

Table D6: Location choices and income sorting, control for the main mode of transport

Measure of δ_j	Centroid distance		Avg. commut. distance		Avg. commut. time	
	(1)	(2)	(3)	(4)	(5)	(6)
Location Choice Model						
Municipal characteristics						
$\widehat{\tau}_{ijt}(W_{it})$	-0.599***	-0.539***	-0.647***	-0.575***	-0.607***	-0.539***
	(0.055)	(0.056)	(0.054)	(0.055)	(0.054)	(0.056)
$\ln(Q_{jt}) \times \widehat{\ln(W_{it})}$	0.877***	0.862***	0.697***	0.683***	0.830***	0.787***
	(0.128)	(0.161)	(0.112)	(0.152)	(0.113)	(0.144)
$\delta_j \times \widehat{\ln(W_{it})}$	0.033***	0.021***	0.034***	0.015**	0.041***	0.028***
	(0.005)	(0.006)	(0.006)	(0.007)	(0.006)	(0.006)
Other municipal amenities B_j						
$\ln(\text{Housing}_{jt}) \times \widehat{\ln(W_{it})}$		-0.572***		-0.554***		-0.620***
		(0.195)		(0.194)		(0.194)
$\ln(\text{Foreigners}_{jt}) \times \widehat{\ln(W_{it})}$		-0.377***		-0.406***		-0.358***
		(0.079)		(0.083)		(0.078)
$\ln(\text{Green Space}_{jt}) \times \widehat{\ln(W_{it})}$		0.085**		0.070*		0.070*
		(0.036)		(0.036)		(0.036)
$\mathbb{I}(\text{Lake}_{jt} = 1) \times \widehat{\ln(W_{it})}$		-0.078		-0.083		-0.045
		(0.056)		(0.056)		(0.055)
Average ASC						
Central mun. $\hat{\alpha}_j$	0.000	0.000	0.000	0.000	0.000	0.000
Average suburban $\hat{\alpha}_j$	-6.521	-6.476	-6.559	-6.503	-6.514	-6.471
Indiv.-specific controls	Yes	Yes	Yes	Yes	Yes	Yes
Nbr. municipalities	166	166	166	166	166	166
Nbr. individuals	11031	11031	11031	11031	11031	11031
Log-likelihood	-29972.265	-29940.600	-29976.081	-29944.864	-29965.943	-29936.601

Note: This table reports the pooled estimates of equation (5.1) with disjoint choice sets across agglomerations, controlling additionally for individuals' main mode of transportation (private vehicle vs. public transports and others). The error term η_{ijt} is assumed to follow a Type I GEV distribution. In each specification, the alternative-specific constants of the central municipalities are standardized to 0. The three sets of estimates in columns (1)-(2), (3)-(4) and (5)-(6) are based on different measures of commuting δ_j between municipality j and the center of the agglomeration (*Centroid distance*, *Average commuting distance* and *Average commuting time*). Columns (1)-(3)-(5) report the estimation results for the baseline specification, while columns (2)-(4)-(6) introduce additional interactions between income and municipal amenities. Stars indicate the following significance levels: * 0.1, ** 0.05 and *** 0.01.

E. The spatial distribution of in-movers' income under harmonized tax rates

Figure E1: Capitalization of local taxes into housing prices

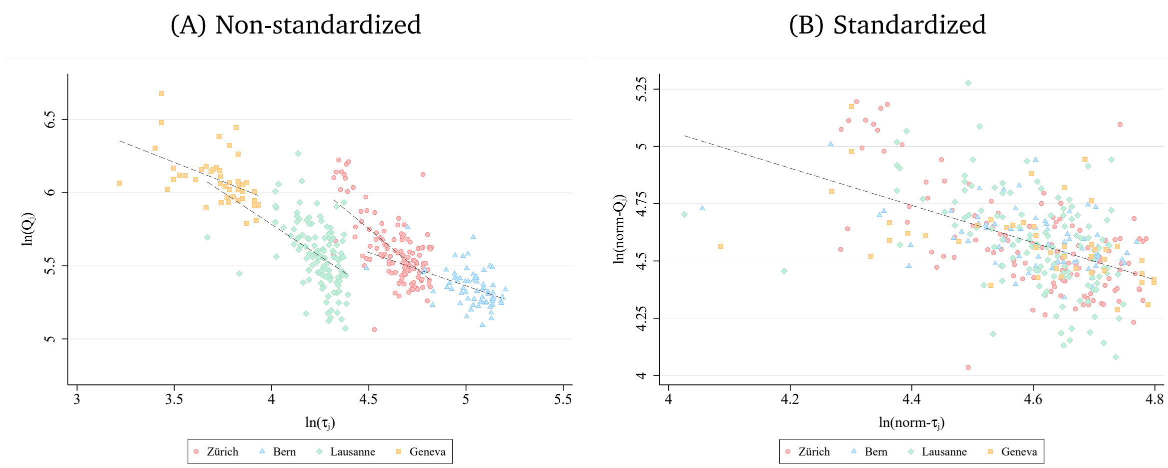

Note: This figure plots the relationship between the local level of taxation (as measured by the log of the local tax multiplier) and the local level of housing prices (as measured by the log of the housing price index). Panel A reports the relationships for the absolute values of the variables, while panel B reports the same relationship for the standardized variables (with respect to the agglomeration average).

Figure E2: Empirical model fit, average income per municipality

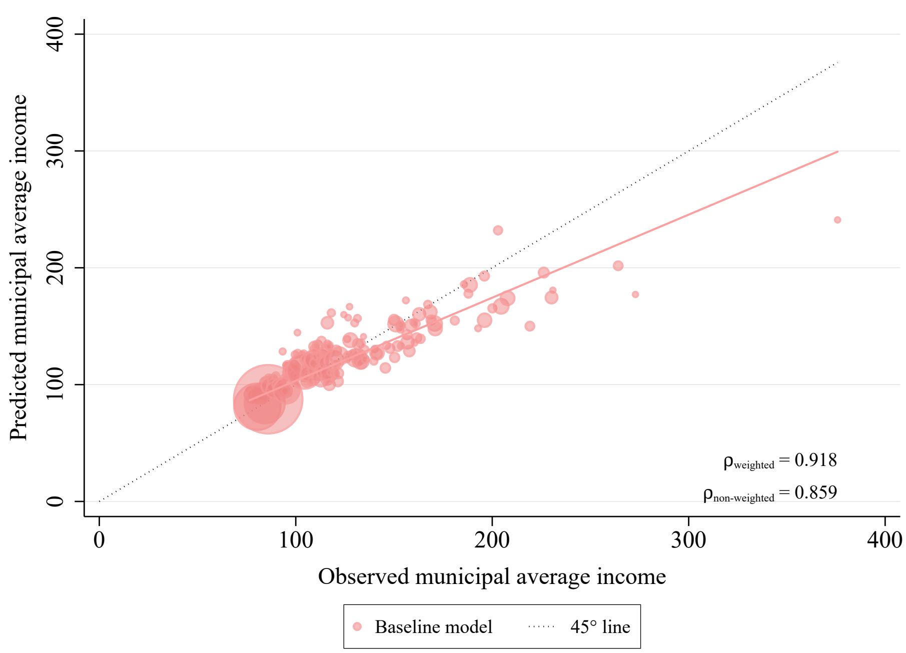

Note: This figure shows the observed vs. predicted average income per municipality based on our empirical estimates from section 5.3. Correlation coefficients are reported on the graph, either weighted according to municipal population or non-weighted.

Figure E3: Counterfactual income distribution accounting for housing prices capitalization
Central municipalities vs. suburbs

(A) Income CDFs

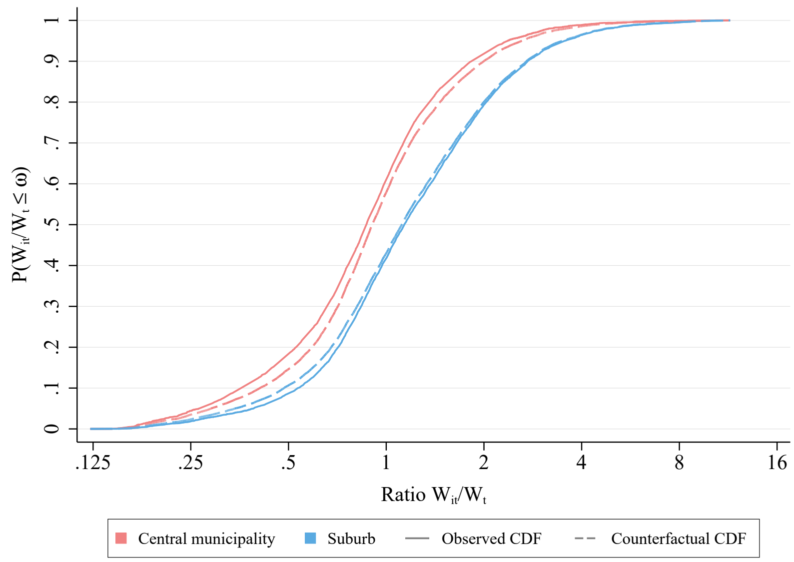

(B) Changes in income deciles

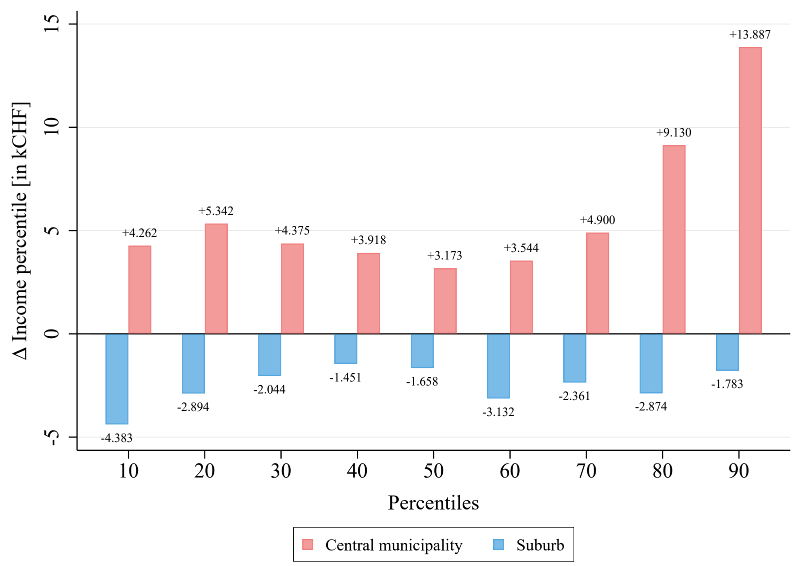

Note: This figure reports graphical results of our counterfactual exercise, where housing prices capitalization is accounted for using the capitalization rate estimates by Basten et al. (2017). Panel A plots the CDFs of income in the central municipalities (in red) and in the suburbs (in blue), for both the observed data (solid line) and the counterfactual scenario that accounts for housing prices capitalization (dashed line). So as to obtain a smooth CDF for the four agglomerations together, we consider the exponentiated demeaned log-income $W_{it}/\overline{W}_t$ (a monotonic transformation of the term $\ln\widehat{(W_{it})}$), reported on a log-scale. Panel B reports corresponding changes in income deciles (reported in kCHF .-), for the central municipalities and the suburbs.