



HAL
open science

Essays in economic geography and urban economics

Maddalena Conte

► **To cite this version:**

Maddalena Conte. Essays in economic geography and urban economics. Economics and Finance. Institut Polytechnique de Paris, 2024. English. NNT : 2024IPPAX090 . tel-04927146

HAL Id: tel-04927146

<https://theses.hal.science/tel-04927146v1>

Submitted on 3 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2024IPPAX090

Thèse de doctorat



Essays in economic geography and urban economics

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École Polytechnique

École doctorale n°626
Ecole Doctorale de l'Institut Polytechnique de Paris (ED IP Paris)
Spécialité de doctorat : Sciences économiques

Thèse présentée et soutenue à Paris, le 22/11/2024, par

MADDALENA CONTE

Composition du Jury :

Gregory Verdugo Professeur, Université Evry Val-d'Essonne	Président
Pierre-Philippe Combes Professeur, Sciences Po (Paris)	Rapporteur
François Fontaine Professeur, Paris School of Economics	Rapporteur
Daphné Skandalis Professeure adjointe, Université de Copenhague	Examinatrice
Benoit Schmutz Professeur agrégé, École polytechnique	Directeur de thèse
Isabelle Mejean Professeure, Sciences Po (Paris)	Co-directrice de thèse

Introduction

Acknowledgments - Remerciements	4
Summary	6
Résumé	6
General introduction	8
Introduction générale	18
1 The volatility advantages of large labor markets	30
1.1 Introduction	31
1.2 Motivating facts on firms' spatial patterns	36
1.2.1 Data	36
1.2.2 Motivating stylized facts	41
1.3 Volatility and firm location: theory	45
1.3.1 Framework	45
1.3.2 Solution	48
1.3.3 Volatility and the sorting of firms	52
1.3.4 Discussion	53
1.4 Volatility and firm location: empirical evidence	54
1.4.1 Empirical strategy	54
1.4.2 Results	56
1.4.3 Robustness: firm birth and demand volatility	59
1.5 Conclusion	61

Appendices	62
A.1 Data appendix	62
A.1.1 Sample selection	62
A.1.2 Total factor productivity	62
A.1.3 Demand volatility	64
A.2 Empirical appendix	65
A.3 Model appendix	71
A.3.1 Definitions	71
A.3.2 Proofs	73
A.3.3 Additional figures	75
A.3.4 Endogenous matching rate	77
A.3.5 Model extension	78
2 Moving to opportunity? Networks, information and skill-biased migration	85
2.1 Introduction	86
2.2 Data and measurement	92
2.2.1 Data	92
2.2.2 Descriptive statistics and sample restrictions	92
2.2.3 Measuring information transmission	96
2.3 Stylized facts	100
2.3.1 Empirical strategy	100
2.3.2 Impact of information on job transitions	102
2.4 A quantitative model of migration with networks	108
2.4.1 Model setup	109
2.4.2 Model estimation	112
2.4.3 Quantification exercises	115
2.4.4 The role of information frictions	120
2.5 Conclusion	122
Appendices	124
B.1 Data appendix	124

B.1.1	Data description	124
B.1.2	Cleaning worker histories	125
B.1.3	Additional sample descriptives	127
B.1.4	Additional descriptives of coworker networks	130
B.1.5	Future networks	135
B.2	Empirical appendix	136
B.2.1	Non-linear models	136
B.2.2	Robustness tests	137
B.3	Model appendix	146
B.3.1	Data construction	146
B.3.2	Empirical implementation of the model	147
B.3.3	Additional quantification exercises	149
3	Affordable housing and the labor market: evidence from the city of Copenhagen	151
3.1	Introduction	152
3.2	Empirical setup and data	157
3.2.1	The quasi-natural experiment	157
3.2.2	Data and descriptive statistics	160
3.3	Effect on labor market participation	163
3.3.1	Empirical strategy	163
3.3.2	Main results	165
3.4	Intergenerational effect	175
3.4.1	Empirical strategy and summary statistics	177
3.4.2	Results for children	179
3.5	Conclusion	182
	Appendices	184
C.1	Data appendix	184
C.2	Empirical appendix	185

Acknowledgments - Remerciements

This dissertation has been inspired by both personal and professional experiences. Over the past 10 years, I have moved between 10 cities and 6 different countries. The curiosity to explore new places and the impulse to say yes to the next adventure fueled my fascination to better understand cities and the diverse opportunities that they can offer. Moving around so much has also made me reflect on the complexity of location decisions and the (sometimes high) costs of moving.

This thesis would not have been possible without the invaluable support of numerous people. First, I would like to thank my two supervisors: Benoit Schmutz and Isabelle Mejean. Their guidance and feedback throughout my PhD has been invaluable. Their presence and encouragement has been all the more precious in light of my numerous external visits. Their support and trust, and the freedom that I have enjoyed during my PhD, have made this experience truly memorable.

I would like to thank Ismir Mulalic, who hosted me in Copenhagen multiple times, making me feel at home in the department and giving me the opportunity to develop a research project in Denmark. Working alongside Ismir and Jos van Ommeren has enabled me to learn new research tools and has motivated me to focus my career on policy evaluation.

I am also grateful for the support and advice of many researchers at CREST who generously dedicated their time and support. I am particularly thankful to Arne Uhlendorff and Roland Rathelot, who have been members of my comité de suivi, and who have provided a reference point throughout each milestone of my PhD. A special mention goes to Clément Malgouyres, whose excitement and passion motivated me to apply to the Institut des Politiques Publiques, which I will join from October 2024.

I am honoured to have met and worked with Carlo Altomonte, whose guidance and mentorship have been central both before starting the PhD and in this new professional chapter. I am also grateful to my defense committee for their precious feedback, in particular Pierre-Philippe Combes, François Fontaine, Gregory Verdugo and Daphé Skandalis.

I would like to express my gratitude to the current and former PhD students both in my home institution and abroad. Our research and non-research discussions have been extremely

helpful during this entire experience. I would like to thank in particular Federica Meluzzi and Marion Brouard, with whom I shared the Californian sun; Thomas Monnier and Léa Bou Sleiman, for the stimulating conversation on urban economics and on life; Jérôme Trinh, Pauline Leveneur, Léa Dubreuil, Riccardo Norbiato. Special thanks also to the Berkeley PhD family, who made the experience truly memorable: Chiara Motta, Andrea Cerrato, Pietro Campa and Eva Davoine.

I also want to thank my friends in Paris and Copenhagen, who made these past three years truly special, in particular Ceci, Angi, and Greta who became my Parisian family, and Henriette, Nele, Leonie and Sara who always motivated me for new Danish adventures. And of course my Genovese rocks and suns, Mimo, Silvia, Ari, Bebe, Sara, Giuli, Laura, and the friends scattered around the globe, in particular Nadine. Thanks for always being there.

I am immensely grateful to my family: my parents, Paolo and Giovanna, who have supported me in every step of the journey, and my siblings, Anna e Pietro, who have filled the physical distance between us with the strongest affection. Lastly, I want to thank Ole, who has always been by my side, when near and when far. Thank you for sharing with me your lightness and your joy for adventures, and may this be the start of many more!

Summary

This thesis studies spatial location decisions of firms and workers, and how these interact with local labor market characteristics. The first chapter focuses on firms and explores a novel mechanism that incentivizes firms to locate in denser cities: the role of volatile demand and its interaction with firm productivity. This channel arises since faster hiring conditions in thicker labor markets attract productive firms that can more swiftly downsize or expand in denser cities. The second chapter explores the location decisions of workers and how regional migration is affected by mobility costs, in particular information frictions. This helps shed light on the mechanisms driving skill-biased migration, namely the empirical regularity that high-skilled workers are substantially more mobile than low-skilled workers. The third chapter studies the interaction of affordable housing policies with incentives for labor market participation. A quasi-natural experimental setting enables to analyze a large public housing privatization event in the city of Copenhagen directed towards low-income households, and to compare the impact of subsidized home purchase versus subsidized rental on long-run labor market outcomes.

Résumé

Cette thèse étudie les décisions de localisation des entreprises et des travailleurs, ainsi que la manière dont celles-ci interagissent avec les caractéristiques locales des marchés du travail. Le premier chapitre se concentre sur les entreprises et s'intéresse à un nouveau mécanisme qui les incite à s'implanter dans les villes les plus denses : le rôle de la volatilité de la demande et de son interaction avec la productivité des entreprises. Ce processus tient à des conditions d'embauche plus fluides qui attirent les entreprises productives pouvant plus rapidement réduire ou augmenter leurs effectifs dans les villes les plus densément peuplées. Le deuxième chapitre s'intéresse aux choix de localisation des travailleurs et à la manière dont les coûts de mobilité, en particulier les frictions informationnelles, affectent les migrations régionales. Cela permet de mettre en lumière les mécanismes à l'origine du biais migratoire lié aux compétences, à savoir la régularité empirique selon laquelle les travailleurs hautement qualifiés

sont significativement plus mobiles que les travailleurs peu qualifiés. Le troisième chapitre porte sur l'interaction des politiques de logement abordable avec les incitations à participer au marché du travail.

General introduction

High economic inequalities between local labor markets are a striking feature of many developed countries (Kline and Moretti, 2013; Kuhn et al., 2021; Bilal, 2023). In 2021 in Versailles, an affluent French city close to Paris, 5.8% of workers were unemployed and the mean after-tax hourly wage was €23.2. By contrast, in Marseille the unemployment rate was almost double and, among employed workers, the mean hourly wage was €15.9.¹ Similar staggering geographical differences arise in most other developed countries and are highly persistent over time (OECD, 2005).

Firms' and workers' location decisions play an important role in explaining the dynamics of geographic inequalities. On the one hand, firm location decisions shape the spatial heterogeneity in productivity, with a profound influence on local employment dynamics and regional economic trajectories (Combes et al., 2012; Gaubert, 2018; Kuhn et al., 2021; Bilal, 2023). On the other hand, worker mobility can, in principle, play a self-equilibrating role in reducing local disparities. However, the propensity to migrate is much higher among the high-skilled (Diamond, 2016; Amior, 2024), while the low-skilled are more dependent on local opportunities.

At the same time, moving to cities that offer better job opportunities is increasingly challenging due to rising housing prices (OECD, 2021). This affects particularly the most disadvantaged social groups, making it difficult to afford quality housing, especially in areas that are close to jobs. Growing housing costs in urban centers have thus put housing support at the forefront of current policy debates (Saiz, 2023), but the impact of affordable housing policies is still not well understood.

This dissertation studies spatial economic inequalities and how they are shaped by firms'

1. Source: 2021 INSEE Census data.

(chapter 1) and workers' (chapter 2) location decisions using French administrative data. In chapter 3, I examine the role of affordable housing policies and their impact on labor market participation, exploiting a quasi-natural experimental setting in the city of Copenhagen. The remainder of this section introduces the key topics of this dissertation. I then provide an overview of each chapter and highlight its contribution to the existing literature.

Agglomeration economies and firm location decisions

Agglomeration economies are the idea that the productivity of firms (and workers) increases with the size of the local economy (Combes and Gobillon, 2015). This is driven by various mechanisms, which the literature to date classifies in three main effects: sharing, matching, and learning (Duranton and Puga, 2004). The first chapter of this thesis focuses on the matching channel, whereby high human densities facilitate the quality and speed of the hiring process.² Locating in denser cities is particularly beneficial for high-productivity firms since their opportunity costs of operating with limited capacity are higher (Combes et al., 2012). The complementarity between employer productivity and local hiring conditions gives rise to a labor market pooling externality (Bilal, 2023), since productive firms are willing to pay more for slack labor markets where they can recruit more rapidly. Besides productivity, a second, less explored, firm characteristic that may be complementary to labor market pooling is firm volatility (Krugman, 1992). When labor demand fluctuates, it is easier for firms to hire when neighboring firms are downsizing. Thus, volatile firms with imperfectly correlated labor demand may also have an incentive to locate in denser cities with thicker labor markets.

Worker regional mobility and skill-biased migration

Spatial disparities in local opportunities and wages especially affect low-skilled workers (Hoynes, 2000; Gregg et al., 2004; Diamond and Moretti, 2021). In principle, regional worker mobility should reduce these geographic inequalities (Blanchard and Katz, 1992). However, empirical evidence points to significantly lower mobility among low-skilled workers (Amior, 2024), who also struggle to relocate in response to local employment shocks (Wozniak, 2010;

2. By contrast, sharing effects describe the gains in larger cities from a greater variety of inputs, from sectoral specialization, from the shared use of indivisible infrastructure, and from the risk-pooling. The learning effects illustrate the increased generation, diffusion, and accumulation of knowledge that is facilitated when a larger number of workers and firms can meet and exchange ideas and experiences.

Amior and Manning, 2018). One key factor driving these different mobility patterns is that, compared to the low-skilled, high-skilled workers see higher returns to job match quality (Amior, 2024), while their job opportunities are highly concentrated in specific high-growth cities. This makes regional migration more salient for workers at the top of the skill distribution and is more likely to justify moving costs for this group of workers. A second aspect that can explain the reduced regional mobility of low-skilled workers is the high barriers to migration (Kennan and Walker, 2011), which are especially elevated among low-skilled workers (Diamond, 2016; Balgova, 2018; Caldwell and Danieli, 2024). Schmutz et al. (2021) in particular highlight spatial search frictions that limit the ability to hear about or apply to remote job opportunities among workers at the lower end of the skill distribution. More specifically, low-skilled workers may experience higher information frictions, namely higher costs of gathering and processing information on the job opportunities that different cities can offer and on other location attributes that can shape migration decisions.

Local public policies and housing support

Public policies can play an important role in reducing spatial inequalities. Among the various policy tools, housing policy stands out as a key instrument to confront increasing geographic income segregation (Saiz, 2023). A key question is whether government efforts should be directed towards supporting home-ownership versus rental. Tax incentives and subsidies for home-ownership are widespread across most developed countries, for example, as tax-deductible mortgage interest payments, or through purchase subsidies for first-time buyers (Schwartz et al., 2006). By contrast, the share of social housing has declined in most OECD countries since 2010 (OECD, 2021). In addition, some local governments have privatized a portion of public housing, for example, in Israel (Hausman et al., 2022), Sweden (Sodini et al., 2023), and in Denmark. While both subsidized home purchase and subsidized rental continue to represent a significant item in public budgets, the effectiveness of these measures is still debated. Of particular interest is the impact that these different policies have on labor supply, and therefore economic well-being, for example by enabling low-skilled workers to live closer to jobs. Empirical evidence is however scarce and ambiguous, also due to the fact that it is challenging to identify settings with plausible exogenous variation.

This dissertation

This dissertation builds on three economic essays on the drivers of firm location decisions (chapter 1), the role of information frictions on worker regional migration (chapter 2) and the interaction of affordable housing policies and labor market supply (chapter 3). The objective of this thesis is twofold. A first aim is to better understand the spatial location decisions of firms (chapter 1) and workers (chapter 2), and how these interact with local labor market characteristics. Secondly, this dissertation attempts to shed more light on the role of affordable housing policies, and how these can support low-income individuals in accessing local job opportunities (chapter 3). In what follows, I summarize the three chapters of this dissertation focusing on the methods used and the contributions to the literature.

Chapter 1 – The volatility advantages of large labor markets

Many studies (Combes et al., 2012; Gaubert, 2018) provide theory and evidence that more productive firms sort into larger cities, highlighting the role of matching economies, whereby larger and denser cities facilitate the quality and speed of the hiring process (Duranton and Puga, 2004). This mechanism benefits particularly high-productivity firms due to their higher opportunity costs of operating under limited capacity.

In this chapter (with Isabelle Mejean, Tomasz Michalski and Benoit Schmutz) we enrich this story by examining another dimension of firm heterogeneity, namely the volatility of firm activity. Firms with volatile and imperfectly correlated labor demand benefit from labor market pooling externalities, which, for example, make it easier for firms to hire when neighboring firms are downsizing.

We study this mechanism both through a stylized model and in the data. Both approaches highlight three key empirical regularities: (i) firms sort positively on productivity and volatility; (ii) these two dimensions reinforce each other; and (iii) the resulting gradient of firm productivity with city density decreases with firm volatility. These findings thus highlight that firms with a more volatile activity benefit from locating in denser locations.

Higher operating costs associated with density create an incentive for volatile firms to adopt a more flexible workforce management strategy. In turn, by frequently releasing workers, these firms generate a positive externality on other firms, which benefit from easier hiring conditions. These results also help understand the non-negative correlation between city size and unemployment rates, and the observation that many low-productivity firms are able to operate in large cities.

Methodology — We study agglomeration patterns when firms are heterogeneous along both the productivity and volatility dimensions. Through a stylized model inspired by Mortensen and Pissarides (1994), we investigate how the complementarity between volatility and local hiring conditions interacts with sorting patterns along the productivity dimension when firms can adjust their size positively or negatively in response to idiosyncratic shocks. High job-filling rates in dense cities reduce the cost of these fluctuations, particularly for high-volatility firms. Using French administrative data, we provide evidence of a systematic correlation between the density of cities and the average volatility of firms, conditional on productivity. We then estimate a model of location choices to quantify the relative importance of productivity and volatility in shaping location decisions. The results are in line with the model’s predictions and corroborate the impact of volatility on firms’ location choices.

Contributions — This chapter contributes to several strands of the spatial economics literature. Among others, Combes et al. (2012) and Gaubert (2018) have shown the central role of the productivity-density nexus. We present a new mechanism for agglomeration economies based on the combination of firm productivity and volatility, which may also help explain why some relatively unproductive firms can survive in denser areas (Combes et al., 2012). From a theoretical perspective, Krugman (1992) already pointed out the potential benefits of labor market pooling for firms with volatile and imperfectly correlated labor demand. The main existing attempt to provide reduced-form evidence for this channel is the study by Overman and Puga (2010), which we integrate by testing our model on firm-level data and exploring the interaction between productivity and volatility on sorting patterns.

This chapter also contributes to the literature on the relationship between city size and

unemployment (Kuhn et al., 2021; Bilal, 2023), as we highlight that firms in large labor markets have higher structural volatility and consequently higher employment volatility, leading to more aggregate labor turnover and also more unemployment. More generally, this paper complements the literature on the spatial dimension of matching in cities, which has so far largely focused on the worker side (Gan and Zhang, 2006; Bleakley and Lin, 2012; Schmutz and Sidibé, 2019; Dauth et al., 2022; Papageorgiou, 2022) with some incomplete evidence on firms (Glaeser et al., 1992; Henderson et al., 1995; Combes, 2000; Duranton, 2007; Findeisen and Südekum, 2008). We also draw inspiration from the macroeconomic literature, which has long discussed heterogeneity across firms in productivity and volatility (Comin and Philippon, 2006; Comin and Mulani, 2006; Davis et al., 2007; di Giovanni et al., 2014). Finally, we enrich the literature on labor market churning (Burgess et al., 2000; Nekoei and Weber, 2020; Weingarden, 2020) with a focus on the spatial dimension.

Chapter 2 – Moving to opportunity? Networks, information and skill-biased migration

A key empirical regularity when studying labor mobility across regions and cities is that high-skilled workers are significantly more likely to move to another local labor market than low-skilled workers (Amior, 2024). One explanation put forward by the literature are high barriers to mobility (Kennan and Walker, 2011), especially among low-skilled workers (Diamond, 2016; Balgova, 2018; Caldwell and Danieli, 2024). To shed light on skill-biased migration, this chapter examines one such type of barrier to mobility, namely information frictions, which capture the costs of gathering and processing information on a range of location attributes.

To measure information frictions, I exploit the spatial distribution of workers' professional networks, which provide key information on the local job opportunities that each worker is likely to hear about (Topa, 2011; Caldwell and Harmon, 2019). Since professional networks are much less developed in cities where workers do not live, their spatial distribution can be used to quantify information frictions across cities. I show that increases in labor demand

at former coworkers' current workplaces in different cities raise the likelihood of migration decisions, especially to cities and plants where the worker knows at least one former coworker. However, this effect is significantly weaker for low-skilled workers, who also tend to be less mobile and have more spatially concentrated networks. This channel contributes substantially to differences in migration propensities between low- and high-skilled workers and to the resulting wage inequality between these two skill groups.

Methodology — Using French matched employer-employee databases, I measure a worker's information about job opportunities by weighting plant-specific changes in labor demand by each worker's coworker network in that plant (Caldwell and Harmon, 2019). In a first set of stylized facts, I estimate the impact of this variable on the migration propensity, using a restrictive set of fixed effects to address the threat of omitted variables. I then develop a quantitative two-step nested logit model of job mobility that enables to disentangle the impact of information frictions from that of other push and pull factors of migration. The model further allows to separate the effect of the spatial distribution of professional networks from that of local opportunities. This reveals that the lower response of low-skilled workers to information on local job opportunities is not only caused by a more spatially homogeneous distribution of job opportunities, which makes migration less salient among this group of workers. Instead, it also reflects their weaker professional network, which co-locates less well with local job opportunities across space, thus giving rise to information frictions that limit the potential wage gains of low-skilled workers.

Contributions — This chapter builds on several existing literatures, starting from studies of mobility between local labor markets (Schmutz and Sidibé, 2019; Monras, 2020; Zerecero, 2021; Bilal, 2023). A strand of this literature has highlighted the high barriers to mobility that limit internal migration (Kennan and Walker, 2011), in particular among low-skilled workers (Topel, 1986; Bound and Holzer, 2000; Wozniak, 2010; Moretti, 2011; Kennan, 2015; Diamond, 2016; Balgova, 2018; Schmutz et al., 2021; Caldwell and Danieli, 2024). However, with few exceptions (Gharad et al., 2014; Wilson, 2021; Porcher, 2022), most studies of internal migration do not consider an explicit information channel, despite the growing evidence that

information shocks affect job application behavior (Skandalis, 2019). In addition, a strand of the literature has highlighted the fact that information transmission on outside options often takes place through social and professional networks (Topa, 2001; Caldwell and Harmon, 2019), but has so far not considered their role as a friction affecting mobility between cities. Conversely, the international migration literature has pointed to social networks as a key source of information on jobs and conditions in the destination country (see Munshi, 2020, for a detailed review). My paper aims to bridge the literatures on regional migration, information frictions and the role of social networks, by shedding light on how inequality in access to information affects migration decisions, and how this differs across the skill distribution of workers.

Chapter 3 – Affordable housing and the labor market: evidence from the city of Copenhagen

Housing support policies have become a key instrument for local governments to confront increasing geographic income segregation and growing housing costs (Saiz, 2023). For example, according to the OECD Affordable Housing Database, public housing accounted for 7% of the total housing stock in OECD countries in 2020, ranging from over 20% in Denmark, to 10% in the US, and 3% in Germany. Nevertheless, the impact of these policies, in particular on long-run labor supply, is still not well understood.

This chapter (with Ismir Mulalic and Jos van Ommeren) exploits a quasi-natural experiment – the 1996 privatization of subsidized public housing owned by the city of Copenhagen – to explore the impact of housing support on labor supply. This intervention offered residents receiving rental subsidies in municipality-owned public housing the opportunity to buy their home at below-market prices. While municipality-owned public housing was privatized, public housing associations continued to provide subsidized public housing. This setup thus enables to compare two key forms of housing support: a large housing purchase subsidy for residents that were offered the option to buy their home (the treatment), versus subsidized rental in public housing (the control). We find that subsidized home-buying had a negative effect on

labor market participation, compared to subsidized rental, as home-buyers experienced large wealth increases, which reduced work incentives. We then examine the impact of this policy on children of home-buyers and of renters in public housing. While effects on their adult labor market participation appear to be limited, we show that public housing privatisation can in the long run increase housing inequalities, due to the high inter-generational persistence of home-ownership.

Methodology — In a difference-in-difference framework with two-way fixed effects, we compare individuals living in public housing owned by the municipality at the time of the policy announcement (the treatment group), with individuals living in public housing owned by public housing associations, who were not offered the possibility to purchase their homes (the control group). This resembles an ideal experiment, in which a large random sample of public housing tenants were offered to buy their home at subsidized prices, and almost all tenants accepted this offer. Crucially, residents could not anticipate the sales policy when they were assigned a specific type of public housing. Our empirical strategy thus enables to overcome a key challenge in studying housing support policies, namely the fact that eligible households are generally a selected non-random sample with distinct socioeconomic characteristics. This is particularly problematic for evaluating long-term labor supply, as the decision to buy a house is typically related to expectations over long-run labor market participation.

Contributions — Despite representing a widely spread policy tool, evidence on the impact of subsidized rental in public housing on labor supply is scarce. One notable exception is van Dijk (2020), who demonstrates that moving into public housing has a positive impact for households that relocate from economically worse to better neighborhoods. Evidence on home purchase subsidies is more developed, with some papers highlighting its positive externalities (Green and White, 1997; Di Pasquale and Glaeser, 1999; Dietz and Haurin, 2003), especially in low-income areas (Shlay, 2006). One additional justification for public involvement in the home-ownership market is that it provides one of the few ways for low-income households to accumulate wealth (Wainer and Zabel, 2020). Considering the extraordinary house price

increases in the last 25 years, these wealth effects are likely to have been substantial, and to have thus reduced labor supply incentives (Henley, 2004; Disney and Gathergood, 2017; Li et al., 2020). In contrast to recent analysis of the privatization of public housing, which has focused on the short run and suggested that home-buyers increase their labor supply (Hausman et al., 2022), we concentrate on the long-run effects of the policy. This enables to account for the large wealth effect implied by home purchase subsidies, that amplified over time, pushed by house price appreciation.

This chapter also contributes to the literature studying the impact of affordable housing policies on children who grew up under differential subsidized housing arrangements. Most existing studies find a positive effect of moving out of public housing on adult labor market outcomes (Chetty and Hendren, 2018; Chyn, 2018; Haltiwanger et al., 2020; Pollakowski et al., 2022). However, they cannot disentangle the location effect from that of public housing as they examine US housing voucher policies that enabled families to move out of public housing located in high-crime low-employment neighbourhoods. Another strand of the literature has studied the general effect of home-ownership on children (Green and White, 1997; Haurin et al., 2002; Boehm and Schlottmann, 1999; Mohanty and Raut, 2009; Bourassa et al., 2015), with mixed conclusions. This chapter enriches this literature by focusing on subsidized home-purchase policies, on which evidence has so far been scarce since it is challenging to find settings with plausible exogenous variation.

Introduction générale

Les fortes inégalités économiques entre les marchés du travail locaux sont une caractéristique marquante de nombreux pays développés (Kline and Moretti, 2013; Kuhn et al., 2021; Bilal, 2023). En 2021 à Versailles, une ville française prospère proche de Paris, 5,8% des travailleurs étaient au chômage et le salaire horaire moyen après impôts était de 23,2 €. En revanche, à Marseille, le taux de chômage était presque le double et, parmi les travailleurs employés, le salaire horaire moyen était de 15,9 €. ³ Des différences géographiques tout aussi significatives apparaissent dans la plupart des autres pays développés et sont très persistantes dans le temps (OECD, 2005).

Les décisions de localisation des entreprises et des travailleurs jouent un rôle important dans l'explication de la dynamique des inégalités géographiques. D'une part, les décisions d'implantation des entreprises façonnent l'hétérogénéité spatiale en termes de productivité, avec une profonde influence sur la dynamique de l'emploi local et des trajectoires économiques régionales (Combes et al., 2012; Gaubert, 2018; Kuhn et al., 2021; Bilal, 2023). D'autre part, la mobilité des travailleurs peut, en principe, jouer un rôle d'auto-équilibre dans la réduction des disparités locales. Cependant, la propension à migrer est beaucoup plus élevée parmi les personnes hautement qualifiées (Diamond, 2016; Amior, 2024), tandis que les personnes peu qualifiées dépendent davantage des opportunités locales.

Dans le même temps, il est de plus en plus difficile de migrer vers des villes qui offrent de meilleures opportunités d'emploi en raison de la hausse des prix de l'immobilier (OECD, 2021). Cela affecte particulièrement les groupes sociaux les plus défavorisés qui peinent à accéder à un logement de qualité, en particulier dans les zones proches des emplois. L'augmentation des coûts du logement dans les centres urbains a ainsi placé l'aide au logement au premier plan

3. Source : données du recensement 2021 de l'INSEE.

des débats politiques actuels (Saiz, 2023), mais l’impact des politiques de logement abordable n’est toujours pas bien compris.

Cette thèse étudie les inégalités économiques spatiales et la manière dont elles sont façonnées par les décisions de localisation des entreprises (chapitre 1) et des travailleurs (chapitre 2) en utilisant des données administratives françaises. Dans le chapitre 3, j’examine le rôle des politiques de logement abordable et leur impact sur la participation au marché du travail, en exploitant un cadre expérimental quasi-naturel dans la ville de Copenhague. Le reste de cette section présente les principaux sujets de cette thèse. Je donne ensuite un aperçu de chaque chapitre et souligne sa contribution à la littérature existante.

Économies d’agglomération et décisions d’implantation des entreprises

Les économies d’agglomération correspondent à l’idée que la productivité des entreprises (et des travailleurs) augmente avec la taille de l’économie locale (Combes and Gobillon, 2015). Cette idée est motivée par divers mécanismes, que la littérature à ce jour classe en trois effets principaux : le partage (“sharing”), l’appariement (“matching”) et l’apprentissage (“learning”)(Duranton and Puga, 2004). Le premier chapitre de cette thèse se concentre sur le canal de l’appariement, par lequel des densités humaines élevées facilitent la qualité et la rapidité du processus d’embauche.⁴ S’implanter dans des villes plus denses est particulièrement bénéfique pour les entreprises à forte productivité puisque leurs coûts d’opportunité à opérer avec une capacité limitée sont plus élevés (Combes et al., 2012). La complémentarité entre la productivité des employeurs et les conditions locales d’embauche donne lieu à une externalité de mutualisation des marchés du travail (“labor market pooling externalities”) entre les entreprises (Bilal, 2023), puisque les entreprises productives sont prêtes à payer plus cher pour opérer dans des marchés du travail plus denses où elles peuvent recruter plus rapidement. Outre la productivité, une deuxième caractéristique des entreprises moins étudiée, qui peut être complémentaire de la mutualisation des marchés du travail est la volatilité de l’activité (Krugman, 1992). Lorsque la demande de main-d’œuvre fluctue, il est plus facile pour les

4. En comparaison, les effets de partage décrivent les gains provenant d’une plus grande variété d’intrants, d’une plus grande spécialisation sectorielle, de l’utilisation partagée d’infrastructures indivisibles et de la mutualisation des risques dans les grandes villes. Les effets d’apprentissage illustrent la génération, la diffusion et l’accumulation accrues de connaissances qui sont facilitées lorsqu’un plus grand nombre de travailleurs et d’entreprises peuvent se rencontrer et échanger idées et expériences.

entreprises d'embaucher lorsque les entreprises voisines réduisent leurs effectifs. Ainsi, des entreprises à l'activité volatile dont la demande de main-d'œuvre est imparfaitement corrélée peuvent également être incitées à s'implanter dans des villes plus denses avec des marchés du travail plus denses.

Mobilité régionale des travailleurs et migrations axées sur les compétences

Les disparités spatiales en termes d'opportunités et de salaires locaux affectent particulièrement les travailleurs peu qualifiés (Hoynes, 2000; Gregg et al., 2004; Diamond and Moretti, 2021). En principe, la mobilité régionale des travailleurs devrait réduire ces inégalités géographiques (Blanchard and Katz, 1992). Cependant, les données empiriques indiquent une mobilité nettement plus faible chez les travailleurs peu qualifiés (Amior, 2024), qui ont également du mal à se relocaliser en réponse aux chocs locaux sur l'emploi (Wozniak, 2010; Amior and Manning, 2018). L'un des facteurs clés à l'origine de ce résultat est que, par rapport aux travailleurs peu qualifiés, les travailleurs hautement qualifiés bénéficient d'un rendement plus élevé en termes d'adéquation de l'emploi (Amior, 2024), alors que leurs opportunités d'emploi sont fortement concentrées dans des villes spécifiques à forte croissance. Cela rend les migrations régionales plus rentables pour ce groupe de travailleurs. Un deuxième aspect qui peut expliquer la mobilité régionale réduite des travailleurs peu qualifiés sont les coûts migratoires (Kennan and Walker, 2011), qui sont particulièrement élevés parmi cet autre groupe de travailleurs (Diamond, 2016; Balgova, 2018; Caldwell and Danieli, 2024). Schmutz et al. (2021) met notamment en évidence les frictions liées à la recherche d'emploi dans l'espace qui limitent la possibilité d'entendre parler ou de postuler à des offres d'emploi lointaines. Plus spécifiquement, les travailleurs peu qualifiés peuvent être confrontés à des frictions informationnelles plus importantes, à savoir des coûts plus élevés de collecte et de traitement d'informations sur les opportunités d'emploi que différentes villes peuvent offrir et sur d'autres attributs de localisation qui peuvent influencer le choix migratoire.

Politiques publiques locales et soutien au logement

Les politiques publiques peuvent jouer un rôle important dans la réduction des inégalités spatiales. Parmi ces différents outils, la politique du logement se distingue comme un

instrument clé pour faire face à la ségrégation géographique croissante des revenus (Saiz, 2023). Une question clé est de savoir si les efforts des gouvernements devraient être orientés vers l’accession à la propriété plutôt que vers la location. Les incitations fiscales et les subventions à l’accession à la propriété sont courantes dans la plupart des pays développés, par exemple sous forme de paiements d’intérêts hypothécaires déductibles des impôts ou de subventions à l’achat pour les primo-accédants (Schwartz et al., 2006). En revanche, la part du logement social a diminué dans la plupart des pays de l’OCDE depuis 2010 (OECD, 2021). En outre, certaines administrations locales ont privatisé une partie du logement public, par exemple en Israël (Hausman et al., 2022), en Suède (Sodini et al., 2023) et au Danemark. Si l’achat et la location subventionnés de logements continuent de représenter un poste important des budgets publics, l’efficacité de ces mesures fait encore l’objet de débats. L’impact que ces différentes politiques ont sur l’offre de travail, et donc sur le bien-être économique, par exemple en permettant aux travailleurs peu qualifiés de vivre plus près de leur emploi, est particulièrement intéressant. Les preuves empiriques sont toutefois rares et ambiguës, notamment en raison du fait qu’il est difficile d’identifier des contextes présentant une variation exogène plausible.

Cette thèse

Cette thèse s’appuie sur trois essais économiques portant sur les facteurs qui influent sur les décisions d’implantation des entreprises (chapitre 1), le rôle des frictions d’information sur la migration régionale des travailleurs (chapitre 2) et l’interaction des politiques de logement abordable et de l’offre du marché du travail (chapitre 3). L’objectif de cette thèse est double. Le premier objectif est de mieux comprendre les décisions d’implantation spatiale des entreprises (chapitre 1) et des travailleurs (chapitre 2), et comment celles-ci interagissent avec les caractéristiques des marchés du travail locaux. Deuxièmement, cette thèse tente de mettre davantage en lumière le rôle des politiques de logement abordable et la manière dont celles-ci peuvent aider les personnes à faible revenu à accéder aux opportunités d’emploi locales (chapitre 3). Dans ce qui suit, je résume les trois chapitres de cette thèse en me concentrant sur les méthodes utilisées et les contributions à la littérature.

Chapitre 1 – Les avantages de volatilité des grands marchés du travail

De nombreuses études (Combes et al., 2012; Gaubert, 2018) fournissent des théories et des preuves selon lesquelles les entreprises les plus productives se regroupent dans les grandes villes, soulignant le rôle des économies d'appariement (“matching economies”), par lesquelles les villes plus grandes et plus denses facilitent la qualité et la rapidité du processus d'embauche (Duranton and Puga, 2004). Ce mécanisme profite particulièrement aux entreprises à forte productivité en raison de leurs coûts d'opportunité plus élevés liés à l'exploitation sous capacité limitée.

Dans ce chapitre (avec Isabelle Mejean, Tomasz Michalski et Benoit Schmutz), nous enrichissons cette vision en examinant une autre dimension de l'hétérogénéité des entreprises, à savoir la volatilité de leur activité. Les entreprises dont la demande de main-d'œuvre est volatile et imparfaitement corrélée bénéficient des externalités de mutualisation des marchés du travail (“labor market pooling externalities”) qui, par exemple, facilitent l'embauche des entreprises lorsque les entreprises voisines réduisent leurs effectifs.

Nous étudions ce mécanisme à la fois à travers un modèle stylisé et dans les données. Les deux approches mettent en évidence trois régularités empiriques clés : (i) les entreprises se regroupent en fonction de leur productivité et de la volatilité de leur activité ; (ii) ces deux dimensions se renforcent mutuellement ; et (iii) le gradient de productivité des entreprises par rapport à la densité des villes diminue avec la volatilité de l'activité des entreprises. Ces résultats soulignent ainsi que les entreprises ayant une activité plus volatile bénéficient de leur localisation dans des endroits plus denses. Les coûts d'exploitation plus élevés associés à la densité incitent les entreprises à l'activité volatile à adopter une stratégie de gestion de la main-d'œuvre plus flexible. En retour, en licenciant fréquemment leurs travailleurs, ces entreprises génèrent une externalité positive sur d'autres entreprises, qui bénéficient de conditions d'embauche plus faciles. Ces résultats aident également à comprendre la corrélation positive entre la taille des villes et le taux de chômage, ainsi que l'observation selon laquelle de nombreuses entreprises à faible productivité parviennent à opérer dans les grandes villes.

Méthodologie — Nous nous intéressons aux phénomènes d'économie d'agglomération avec des entreprises hétérogènes selon des dimensions de productivité et de volatilité de l'activité. Grâce à un modèle stylisé inspiré de Mortensen and Pissarides (1994), nous étudions comment la complémentarité entre la volatilité de l'activité et les conditions d'embauche locales interagissent avec des mécanismes de tri selon la productivité lorsque les entreprises ajustent leur taille en réponse à des chocs idiosyncratiques. Les taux de remplissage élevés des emplois dans les villes denses réduisent le coût de ces fluctuations, en particulier pour les entreprises à l'activité volatile. En utilisant des données administratives françaises, nous montrons une corrélation systématique entre la densité des villes et la volatilité moyenne de la demande des entreprises, conditionnellement à leur productivité. Nous estimons ensuite un modèle de choix de localisation pour quantifier l'importance relative de la productivité et de la volatilité de l'activité dans la détermination des décisions de localisation des entreprises. Les résultats sont conformes aux prédictions du modèle et corroborent l'impact de la volatilité sur les choix de localisation des entreprises.

Contributions — Ce chapitre contribue à plusieurs volets de la littérature en économie spatiale. Entre autres, Combes et al. (2012) et Gaubert (2018) ont montré le rôle central du lien productivité-densité. Nous présentons un nouveau mécanisme pour les économies d'agglomération, basé sur la combinaison de la productivité et de la volatilité de la demande des entreprises, qui peut également aider à expliquer pourquoi certaines entreprises relativement improductives peuvent survivre dans des zones plus denses (Combes et al., 2012). D'un point de vue théorique, Krugman (1992) a déjà souligné les avantages potentiels de la mise en commun des marchés du travail pour les entreprises ayant une demande de travail volatile et imparfaitement corrélée. La principale tentative existante pour fournir des preuves de ce mécanisme sous forme réduite est l'étude de Overman and Puga (2010), que nous intégrons en testant notre modèle sur des données au niveau des entreprises et en explorant l'interaction entre productivité et volatilité dans les modèles de tri.

Ce chapitre contribue également à la littérature sur la relation entre la taille des villes et le chômage (Kuhn et al., 2021; Bilal, 2023), car nous soulignons que les entreprises opérant dans les marchés du travail de grande taille ont une volatilité structurelle plus élevée et

par conséquent une volatilité de l’emploi plus haute, ce qui conduit à une rotation globale de la main-d’œuvre plus importante et à un chômage plus élevé. Plus généralement, cet article complète la littérature portant sur la dimension spatiale de l’appariement des agents économiques dans les villes, qui s’est jusqu’à présent largement concentrée sur les travailleurs (Gan and Zhang, 2006; Bleakley and Lin, 2012; Schmutz and Sidibé, 2019; Dauth et al., 2022; Papageorgiou, 2022), avec quelques données incomplètes sur les entreprises (Glaeser et al., 1992; Henderson et al., 1995; Combes, 2000; Duranton, 2007; Findeisen and Südekum, 2008). Nous nous inspirons également de la littérature macroéconomique qui aborde depuis longtemps l’hétérogénéité de la productivité et de la volatilité de la demande des entreprises (Comin and Philippon, 2006; Comin and Mulani, 2006; Davis et al., 2007; di Giovanni et al., 2014). Enfin, nous enrichissons la littérature sur le renouvellement du marché du travail (Burgess et al., 2000; Nekoei and Weber, 2020; Weingarden, 2020) en mettant l’accent sur la dimension spatiale.

Chapitre 2 – Vers plus d’opportunités ? Réseaux, information et migrations axées sur les compétences

Une régularité empirique clé dans l’étude de la mobilité de la main-d’œuvre entre régions et/ou villes est que les travailleurs hautement qualifiés sont nettement plus susceptibles de se déplacer vers un autre marché du travail local que les travailleurs peu qualifiés (Amior, 2024). L’une des explications avancées par la littérature est l’existence de barrières élevées à la mobilité, en particulier chez les travailleurs peu qualifiés (Diamond, 2016; Balgova, 2018; Caldwell and Danieli, 2024). En se concentrant sur les migrations axées sur les compétences, ce chapitre examine un de ces types d’obstacles à la mobilité, à savoir les frictions informationnelles, qui capturent les coûts de collecte et de traitement des informations portant sur une série d’attributs liés à la localisation des emplois.

Pour mesurer les frictions informationnelles, j’exploite la distribution spatiale des réseaux professionnels personnels (“coworker networks”) des travailleurs, qui fournissent des informations clés sur les opportunités d’emploi locales dont chaque travailleur est susceptible

d’entendre parler (Topa, 2011; Caldwell and Harmon, 2019). Étant donné que les réseaux professionnels sont beaucoup moins développés dans les villes où les travailleurs ne vivent pas, leur distribution spatiale peut être utilisée pour quantifier les frictions informationnelles entre les villes. Je montre que l’augmentation de la demande d’emploi sur les lieux de travail actuels d’anciens collègues augmente la probabilité de migrer, en particulier vers les villes et les usines où les travailleurs connaissent au moins un ancien collègue. Cependant, cet effet est significativement plus faible pour les travailleurs peu qualifiés, qui ont également tendance à être moins mobiles et à avoir des réseaux plus concentrés spatialement. Cela contribue grandement aux différences de propension à migrer entre les travailleurs peu et hautement qualifiés, et à l’inégalité salariale qui en résulte.

Méthodologie — En utilisant des bases de données appariées entre employeurs et employés en France, je mesure les informations dont disposent les travailleurs sur les opportunités d’emploi locales en pondérant la croissance de la demande de travailleurs dans chaque l’établissement par la densité du réseau de collègues de chaque travailleur dans cet établissement (Caldwell and Harmon, 2019). Dans un premier ensemble de faits stylisés, j’estime l’impact de cette variable sur la propension à migrer, en utilisant un ensemble restrictif d’effets fixes pour pallier un éventuel biais de variables omises. J’élabore ensuite un modèle quantitatif de mobilité professionnelle avec une fonction logistique à double imbrication, qui permet de distinguer l’impact des frictions informationnelles de celui des autres facteurs d’attraction et de répulsion. Le modèle permet en outre de séparer l’effet de la distribution spatiale des réseaux professionnels de celui des opportunités locales. Cela révèle que la plus faible réponse des travailleurs peu qualifiés aux informations sur les offres d’emploi locales n’est pas seulement causée par une distribution spatiale plus homogène de leurs opportunités d’emploi, ce qui réduit les flux migratoires parmi ce groupe de travailleurs. Cela reflète également leur réseau professionnel de moins bonne qualité / moins dense, qui se co-localise moins bien avec les opportunités d’emploi locales dans l’espace, donnant ainsi lieu à des frictions informationnelles qui limitent les gains salariaux potentiels de ce groupe de travailleurs.

Contributions — Ce chapitre s’appuie sur plusieurs publications existantes, à commencer par des études portant sur la mobilité entre les marchés du travail locaux (Schmutz and Sidibé, 2019; Monras, 2020; Zerecero, 2021; Bilal, 2023). Un volet de cette littérature a mis en évidence des barrières importantes à la mobilité qui limitent le taux de migration interne (Kennan and Walker, 2011), en particulier parmi les travailleurs peu qualifiés (Topel, 1986; Bound and Holzer, 2000; Wozniak, 2010; Moretti, 2011; Kennan, 2015; Diamond, 2016; Balgova, 2018; Schmutz et al., 2021; Caldwell and Danieli, 2024). Cependant, à quelques exceptions près (Gharad et al., 2014; Wilson, 2021; Porcher, 2022), la plupart des études portant sur les migrations internes ne prennent pas en compte un mécanisme d’information explicite, malgré le nombre grandissant de preuves que les chocs informationnels affectent les stratégies de candidature des travailleurs (Skandalis, 2019). En outre, un autre pan de la littérature a souligné le fait que la transmission d’informations par rapport aux options externes au poste actuel se fait souvent par le biais des réseaux sociaux et professionnels (Topa, 2001; Caldwell and Harmon, 2019), mais n’a jusqu’à présent pas considéré leur rôle en tant que friction affectant la mobilité entre les villes. À l’inverse, la littérature sur les migrations internationales a souligné que les réseaux sociaux étaient une source clé d’information sur les emplois et les conditions de vie dans le pays de destination (voir Munshi, 2020 pour une revue détaillée). Mon article vise à relier les littératures portant sur les migrations régionales, les frictions informationnelles et le rôle des réseaux sociaux et professionnels, en mettant en lumière la manière dont l’inégal accès à l’information affecte le choix migratoire et comment celui-ci diffère selon la répartition des compétences des travailleurs.

Chapitre 3 – Logement abordable et marché du travail : l’exemple de la ville de Copenhague

Les politiques d’aide au logement sont devenues un instrument essentiel pour les pouvoirs publics locaux afin de faire face à la ségrégation géographique croissante des revenus et à la hausse des coûts du logement (Saiz, 2023). Par exemple, selon la base de données de l’OCDE sur le logement abordable, le logement social représentait 7% du parc immobilier total des

pays membres en 2020, allant de plus de 20% au Danemark à 10% aux États-Unis et 3% en Allemagne. Néanmoins, l'impact de ces politiques, en particulier sur l'offre de travail à long terme, n'est toujours pas bien compris.

Ce chapitre (co-écrit avec Ismir Mulalic et Jos van Ommeren) exploite une expérience quasi-naturelle – la privatisation en 1996 des logements sociaux subventionnés appartenant à la municipalité de Copenhague – pour étudier l'impact de l'aide au logement sur l'offre de travail. Cette intervention offrait aux résidents de ces logements la possibilité de devenir propriétaire à des prix inférieurs à ceux du marché. Alors que les logements sociaux appartenant aux municipalités ont été privatisés, les associations de logement social ont continué à fournir des logements subventionnés. Cette configuration permet ainsi de comparer deux formes clés d'aide au logement : une subvention importante à l'achat (le traitement) par rapport à la location subventionnée de logements sociaux (le contrôle). Nous constatons que l'achat subventionné de logements a eu un effet négatif sur le taux de participation au marché du travail par rapport à la location subventionnée, car les acheteurs de logements ont connu de fortes augmentations d'actifs, ce qui a réduit leurs incitations à travailler. Nous examinons ensuite l'impact de cette politique sur les enfants des acheteurs et des locataires. Bien que les effets sur leur participation au marché du travail à l'âge adulte semblent limités, nous montrons que la privatisation du logement social peut à long terme accroître les inégalités d'accès au logement, en raison de la forte persistance intergénérationnelle de la propriété du logement.

Méthodologie — Grâce à une méthode des doubles différences à doubles effets fixes, nous comparons les individus vivant dans des logements sociaux appartenant à la municipalité au moment de l'annonce de la politique (le groupe de traitement) avec les individus vivant dans ceux appartenant à des associations de logement social et à qui l'on n'a pas offert la possibilité d'acheter leur logement (le groupe de contrôle). Le contexte empirique tend vers celui d'une expérience idéale, dans laquelle un large échantillon aléatoire de locataires de logements sociaux se sont vu proposer d'acheter leur logement à des prix subventionnés, et où presque tous ont accepté cette offre. De manière cruciale, les résidents ne pouvaient pas anticiper la politique de vente lorsqu'ils se sont vu attribuer un type spécifique de logement social.

Notre stratégie empirique permet ainsi d'écartier un biais courant dans l'étude des politiques de soutien au logement, à savoir le fait que les ménages éligibles correspondent généralement à un échantillon non-aléatoire sélectionné sur la base de caractéristiques socio-économiques spécifiques. Cela est particulièrement problématique pour évaluer l'offre de travail à long terme, car la décision d'acheter un logement est généralement liée aux attentes concernant la participation au marché du travail à long terme.

Contributions — Bien qu'il s'agisse d'un outil politique largement répandu, les preuves de l'impact des loyers subventionnés dans les logements sociaux sur l'offre de travail sont rares. Une exception notable est van Dijk (2020), qui démontre que l'emménagement dans un logement social a un impact positif pour les ménages qui déménagent d'un quartier défavorisé / aux conditions économiques dégradées vers un quartier plus dynamique. Les résultats portant sur les subventions à l'achat sont plus fréquentes, certains articles soulignant leurs externalités positives (Green and White, 1997; Di Pasquale and Glaeser, 1999; Dietz and Haurin, 2003), en particulier dans les zones à faible revenu (Shlay, 2006). Une justification supplémentaire de l'implication publique dans l'accession à la propriété est qu'elle offre l'un des rares moyens aux ménages à faible revenu d'accumuler des actifs (Wainer and Zabel, 2020). Compte tenu de l'augmentation extraordinaire des prix de l'immobilier au cours des 25 dernières années, ces effets d'accroissement de richesse ont été substantiels et ont ainsi réduit les incitations à participer au marché du travail (Henley, 2004; Disney and Gathergood, 2017; Li et al., 2020). Contrairement aux analyses récentes de la privatisation du logement social, qui se sont concentrées sur le court terme et ont suggéré que les acheteurs de logements augmentent leur offre de travail (Hausman et al., 2022), nous nous concentrons sur les effets à long terme d'une telle politique. Cela permet de tenir compte de l'important effet d'accroissement de richesse impliqué par les subventions à l'achat, lequel s'est amplifié au fil du temps avec l'appréciation des prix de l'immobilier.

Ce chapitre contribue également à la littérature qui étudie l'impact des politiques de logement abordable sur les perspectives socio-professionnelles des enfants qui ont grandi dans ces logements. La plupart des études existantes constatent que le fait de quitter un logement social a un effet positif sur les conditions d'emploi arrivé à l'âge adulte (Chetty and Hendren,

2018; Chyn, 2018; Haltiwanger et al., 2020; Pollakowski et al., 2022). Cependant, elles ne peuvent pas dissocier l'effet de localisation du logement de celui propre au fait d'habiter dans un logement social, par exemple lorsqu'elles examinent les politiques d'aide au logement aux États-Unis qui ont permis aux ménages de quitter des logements sociaux situés dans des quartiers à forte criminalité et à faible taux d'emploi. Un autre volet de la littérature a étudié l'effet général de l'accession à la propriété sur les enfants (Green and White, 1997; Haurin et al., 2002; Boehm and Schlottmann, 1999; Mohanty and Raut, 2009; Bourassa et al., 2015), avec des conclusions mitigées. Ce chapitre enrichit cette littérature en se concentrant sur les politiques d'achat de logements subventionnés sur lesquelles il y a pour l'instant peu de résultats car il est difficile de trouver des contextes empiriques présentant une variation exogène plausible.

Chapter 1

The volatility advantages of large labor markets

With Isabelle Mejean, Tomasz Michalski and Benoit Schmutz

Firms' labor demand is more volatile in larger cities. We propose and test a novel explanation for this finding. Faster hiring conditions attract productive firms with more volatile activity to denser locations where they can swiftly downsize or expand. We estimate a model of firm location choice using French data and show that (i) firm volatility is almost as predictive of location choice as productivity; (ii) both dimensions reinforce each other. This mechanism reduces the productivity–density gradient among volatile firms. Imperfectly correlated firm-level shocks, combined with higher operating costs induced by density, generate matching economies.¹

1. This chapter benefited from comments by Pierre-Philippe Combes, Franck Malherbet, Andrii Parkhomenko, Andrea Weber, and many conference and seminar participants. Maxime Liegey provided research assistance at the start of the project. This research was supported by grants from the Agence Nationale de la Recherche (ANR-10-EQPX-17, Centre d'accès sécurisé aux données, CASD, ANR-11-IDEX-0003-02/Labex ECODEC No. ANR-11-LABEX-0047) and the European Research Council under the European Union's Horizon 2020 research and innovation program (grant agreement No. 714597) and by the Chaire Professorale Jean Marjoulet.

*When I do nothing
I cost less money
Than when I'm working
Or so they tell me.*

Bernard Lavilliers, *Les Mains d'Or*.²

1.1 — Introduction

The productivity-density nexus is a central tenet of economic geography and urban economics (Combes et al., 2012; Gaubert, 2018). One of the channels through which agglomeration economies operate is the matching channel, whereby high human densities facilitate the quality and speed of the hiring process (Duranton and Puga, 2004). This channel is particularly beneficial for high-productivity firms because their opportunity costs of operating with limited capacity are higher. Therefore, the complementarity between employer productivity and local hiring conditions gives rise to a labor market pooling externality (Bilal, 2023). While heterogeneity in firm productivity is a key component of this mechanism, the literature has largely neglected another dimension of heterogeneity that is also pervasive in the data, namely the *volatility* of firm activity. However, Krugman (1992) already pointed out the potential benefits of labor market pooling for firms with volatile and imperfectly correlated labor demand. When labor demand fluctuates, it is easier for firms to hire when neighboring firms are downsizing.

In this paper, we study agglomeration patterns when firms are heterogeneous along two dimensions, productivity and volatility. We do so in the context of a stylized model and in the data. In the model, firms are willing to adjust their size positively or negatively in response to idiosyncratic shocks. High job-filling rates reduce the cost of these fluctuations, more so for high-volatility firms. We investigate how this complementarity between volatility and local hiring conditions interacts with sorting patterns along the productivity dimension when firms can either hold labor demand constant or choose to adjust to shocks. In the data,

2. Song by a popular French blue-collar singer. The original lyrics are “Quand je fais plus rien, moi / Je coûte moins cher / Que quand j’travaillais, moi / D’après les experts.”

we first provide evidence of a systematic correlation between the density of cities and the average volatility of firms there, conditional on productivity. We then estimate a model of location choices to quantify the relative importance of productivity and volatility in shaping location choice decisions. The results are in line with the model's predictions and confirm the role of volatility in firms' location choices.

We begin by documenting new evidence on firm productivity, firm volatility, and local (working-age) population density. We use French administrative data at the worker and firm level over the period 2009-2019, identifying (large) cities with (dense) commuting zones. Our first key finding is that intra-firm employment volatility is higher in denser cities, even after controlling for various relevant firm characteristics such as sector, size, and age. This correlation is quantitatively significant and is hardly reduced when we also control for firm productivity. The second important empirical result is that we find a flatter productivity density gradient among firms with high employment volatility. The elasticity of firms' average productivity with respect to density is reduced by one-third when moving from the first to the last ventile of the volatility distribution.

We propose a simple search model inspired by Mortensen and Pissarides (1994) to rationalize these facts. Firms in the model differ in their productivity and the volatility of their sales. The economy alternates between good and bad states, and the variance of sales induced by these cycles is heterogeneous across firms. Firms can mitigate the impact of volatility by adopting three different employment strategies. The first one aims at maintaining employment levels even in bad times and is chosen by the most productive firms. In the second, the firm freezes hiring in bad states. Freezing hiring avoids facing operating costs in bad states at the cost of entering good states with vacant positions. In the third one, firms "churn": they adopt a turnover strategy, firing workers when hit by bad shocks and hiring only when their demand is high. The latter strategy is preferred by the most volatile firms (fixing productivity) and by the least productive firms (fixing volatility). Compared to a model that does not incorporate volatility and its impact on firms' employment strategies, our model thus shows a weaker selection on productivity.

The model then allows us to analyze where firms choose to locate. The crucial trade-off for firms is that large cities are expensive to operate in — because of higher labor costs or rents —

but allow firms to find workers more quickly when they are needed — when firms experience a positive demand shock. Intuitively, locating in a large city provides “insurance” against volatility because larger cities offer lower adjustment costs for firms. This mechanism is particularly beneficial for high-productivity firms, which have the most to gain from being able to hire more quickly. It is also stronger when a large component of volatility is idiosyncratic, as firms that downsize free up workers that can be hired by expanding firms. Therefore, the model predicts that (i) firms sort positively on productivity and volatility; (ii) these two dimensions reinforce each other; and (iii) the resulting gradient of firm productivity with city density decreases with firm volatility. Since firms are more likely to churn - and thus loosen the market — when they face higher operating costs, this model provides a microfoundation for matching economies based on the existence of urban costs.

Motivated by our theoretical results, we estimate a model of firm location choice and compare the impact of heterogeneity in productivity and volatility on location choice. The model is first estimated on all firms observed in a cross-section of the data, using their productivity and volatility, measured after at least 5 years of existence, as inputs into location decisions. However, this specification does not allow us to distinguish between entry decisions and survival probabilities. In addition, employment volatility is endogenous to the firm’s location choice. According to our model, higher operating costs induced by density will drive some firms to adopt the churning strategy. Outside the model, we cannot rule out the possibility that employment volatility is also driven by a higher probability of job quits in larger cities, where job-to-job transitions are more frequent.

For those reasons, we also estimate a model restricted to firms that were born during the sample period and we propose a novel strategy to measure the exogenous component of a firm’s employment volatility. The measure is similar to a shift-share and combines information on the firm’s product portfolio with the time series of international demand at the product level. Intuitively, the exogenous component of volatility is driven by firms specializing in products for which demand is more or less volatile. We measure the expected volatility of demand resulting from a firm’s decision to produce a given portfolio of products, which we assume is exogenous to its location choice. Both sets of estimates yield consistent results. We show that more volatile firms are more likely to locate in denser commuting

zones. Moreover, firm (demand) volatility is almost as predictive of firm location choice as firm productivity. Finally, consistent with theory, our estimates show that volatility and productivity are complementary in firm location choice.

Relationship to the literature — Many studies, such as Combes et al. (2012) and Gaubert (2018), provide theory and evidence that more productive firms sort into larger cities. Our work presents a new mechanism for agglomeration economies based on the combination of firm productivity and volatility: matching economies arise endogenously from firms' hiring and firing decisions when they face more expensive operating costs. This mechanism may also help explain why relatively unproductive firms can survive in denser areas (Combes et al., 2012), in addition to the mechanisms already proposed in the literature.³ A complementary mechanism that is also consistent with our argument even in the absence of productivity differences is labor market pooling: if demand volatility is uncorrelated across firms, there is a clear advantage for firms to agglomerate because they can hire more workers in good times. This source of agglomeration economies, already recognized by Marshall, was popularized by Krugman (1992). However, this argument has remained largely theoretical.

To the best of our knowledge, the main existing attempt to provide reduced-form evidence for this channel is the study by Overman and Puga (2010). In their static model, firms do not know their productivity before entering a market: productivity is affected by an idiosyncratic shock with known variance. Firms' profits are convex to this shock because firms hire more when the shock is positive, and expected profits thus increase with the variance of the shock. Yet, since wages rise with local demand, firms with higher variance will be all the more profitable when there are many firms, to counteract the effect of individual positive shocks on the local wage level.⁴ Therefore, the model predicts that groups of firms with more variability in labor demand will be more agglomerated, a prediction borne on sector-level data. We instead test our model on individual data and we explore the interaction between productivity

3. Another mechanism, also based on firm entry, is that higher entry costs in larger cities shield unproductive firms from competition from other firms if entry is decided before productivity is realized (Melitz, 2003; Heise and Porzio, 2023).

4. Contrary to our setting, firms do not face hiring frictions. In this respect, we are closer in spirit to the seminal model of Helsley and Strange (1990), which derives agglomeration economies from the matching process of workers to firms.

and volatility on sorting patterns.

By focusing on hiring frictions, this paper also contributes to the literature on the relationship between city size and unemployment. While current leading models of spatial labor markets (Kuhn et al., 2021; Bilal, 2023) posit that more productive firms select into more productive locations, resulting in a negative correlation between average firm productivity and local unemployment rates, they do not directly relate these observations to city size. In the data, large cities are characterized by a higher share of high-productivity firms, but they do not have lower unemployment rates.⁵ One reason for this could be that the mobility of unemployed workers acts as a balancing force in the spatial equilibrium (Gaigne and Sanch-Maritan, 2019). However, churning strategies of firms provide an alternative explanation: if firms in large labor markets have higher structural volatility and consequently higher employment volatility, there may be more aggregate labor turnover and also more unemployment. This mechanism would mitigate the effect of the agglomeration of more productive firms in larger cities.

More generally, this paper complements the literature on the spatial dimension of matching in cities, which has so far largely focused on the worker side (Gan and Zhang, 2006; Bleakley and Lin, 2012; Schmutz and Sidibé, 2019; Dauth et al., 2022; Papageorgiou, 2022) with some incomplete evidence on firms (Glaeser et al., 1992; Henderson et al., 1995; Combes, 2000; Duranton, 2007; Findeisen and Südekum, 2008). In contrast to recent work on the worker side, we abstract from worker heterogeneity. Therefore, we do not address the impact of city size on the level of match assortativeness and we focus on hiring speed as the sole determinant of agglomeration economies.⁶ Moreover, in contrast to the existing literature on the firm side, we do not consider structural characteristics of the economy, such as sectoral composition. Instead, we focus on the heterogeneity of firms, conditional on the sector in which the firm operates. We incorporate two dimensions of heterogeneity that affect the first and second moments of firms' labor demand. In doing so, we draw inspiration from the macroeconomic

5. See Appendix Figure A.3 for the case of French commuting zones.

6. We also abstract from the decision of workers to quit their jobs, which is not observed in our data. Using survey data on U.S. firms, Weingarden (2020) estimates that one-third of firm churning is actually initiated by the employer through layoffs. This figure is arguably a lower bound, since employers may have a financial incentive to get workers to quit rather than lay them off. Weingarden (2020) shows that this component of churning is acyclical, unlike worker quits, which fits well with our modeling assumption of firm-specific shocks.

literature, which has long discussed heterogeneity across firms in productivity *and* volatility.⁷ We enrich this literature by introducing a novel, firm-specific shifter of employment volatility. Finally, we enrich the literature on labor market churning (Burgess et al., 2000; Nekoei and Weber, 2020; Weingarden, 2020) with a focus on the spatial dimension.⁸

The remainder of the paper is organized as follows: in Section 1.2, we provide descriptive evidence that employment volatility increases with city size and that the productivity gradient with respect to city size decreases with employment volatility; in Section 1.3, we present a simple model of firm decisions where employment volatility and location choice are jointly determined. The model predicts that firms sort across space based on the volatility of their activity, and we formally test this prediction in Section 1.4. Section 1.5 concludes.

1.2 — Motivating facts on firms’ spatial patterns

1.2.1. Data

Sample selection — The empirical analysis exploits matched employer-employee data for France over the period from 2009 to 2019 (DADS Postes). This data allows us to characterize the level and volatility of a firm’s labor demand, at the monthly level.⁹ For each employer-employee relationship, we know the type of contract (permanent or short-term), the number of hours and associated earnings, and the worker’s occupation. On the employer’s side, we know the location of each establishment, as well as the sector of activity and date of creation. Finally, the data can be matched with two additional yearly firm-level datasets, namely balance-sheet data used to estimate productivity (FARE) and a production survey (EAP)

7. Comin and Philippon (2006) and Comin and Mulani (2006) document the rise in firm-level volatility among publicly traded US firms in the second half of the 20th century. Davis et al. (2007) instead show diverging trends between public and private firms. In this literature, firm volatility is explained by a combination of aggregate shocks and firm idiosyncratic fluctuations. di Giovanni et al. (2014) provide evidence that a large component of individual firm volatility is driven by idiosyncratic shocks that reflect a combination of demand and supply-side factors.

8. Note that our results also echo some results in the trade literature such as Cuñat and Melitz (2012) showing that countries with more flexible labor markets specialize in sectors with higher volatility.

9. In the rest of the paper, we use a measure of employment equal to the full-time equivalent, based on the number of days worked in each month. We consider a full-time worker as an employee that works 30 days in each month.

that provides additional information on the firm’s portfolio of products.¹⁰

The analysis focuses on firms in manufacturing, construction, and services (including non-tradable services).¹¹ We use information on the address of the establishment to assign each plant to a commuting zone (CZ). Our sample includes plants in mainland France, which is composed of 280 different CZs. Since the focus is on how firms locate across local labor markets, we aggregate plant-level information at the level of a CZ, i.e., firms with multiple plants in the same CZ are treated as a single plant. In the main analysis, we focus on the January 2015 cross-section of the data, which corresponds to the midpoint of our period, but other reference points yield similar results. We focus on firms located in a single CZ because some key variables for our analysis (productivity and demand volatility) can only be calculated at the firm level due to data availability. We also restrict the sample to firms with non-missing data on the key variables of interest (productivity, employment volatility, 2-digit industry, firm age, and firm size). This leaves us with 365,041 firms. Table A.1 in the appendix details the sample selection process.

Each CZ is characterized by its population density, which is defined as the size of its working-age population divided by its area (in square kilometers). The working-age population is taken from the census at the municipality level, where the breakdown of the population by age and municipality is available at a 5-year frequency. For years in which the population is not available, we use data from the previous non-missing year. The area of CZs is based on INSEE 2020 shapefiles (base des zones d’emploi). Table 1.1 provides statistics on the distribution of CZs and the number of firms in each location for the January 2015 cross-section.

Productivity — Firms differ in size, which is typically explained in the literature by some randomness in firm productivity. In the data, we estimate firms’ total factor productivity $\phi_{f,t}$ using the Levinshon-Petrin estimation technique with the Akerberg et al. (2015) correction. Productivity is estimated as the residual of a production function equation including capital

10. The EAP survey is exhaustive for firms in the manufacturing sector above a size threshold of 20 employees. Merging the employer-employee linked data with the this survey introduces severe censoring. The stylized facts discussed in this section exploit the full sample and we restrict the analysis to firms in the EAP survey when we need an exogenous measure of volatility, in section 1.4.3.

11. We exclude the public sector, agriculture, forestry, and fishing, finance and insurance, energy and waste production and distribution, artistic activities, overseas activities, and household services.

Table 1.1 — Population density and firms by commuting zone

	Density	Number of firms
Mean	150.75	1,304
Std. Dev.	496.48	3,723
25th percentile	39.74	447
50th percentile	68.69	682
75th percentile	122.32	1,193

Notes: Summary statistics based on the January 2015 cross-section. Density is measured by working age population divided by the commuting zone’s area in squared kilometers, for the year 2015.

and three types of labor distinguished by their skill levels (Combes et al., 2012). Details of the estimation are provided in the Appendix A.1.2.

Employment volatility — In our model, firms are also heterogeneous in terms of the volatility of their labor demand, due to a combination of structural factors and their endogenous workforce management decisions. We use the panel dimension of the dataset to characterize the volatility of a firm’s labor demand. Following Davis et al. (2006), we define a firm’s volatility as

$$\sigma_{f,t} = \sqrt{\frac{1}{2\omega + 1} \sum_{\tau=-\omega}^{\omega} (\gamma_{f,t+\tau} - \bar{\gamma}_{f,t})^2}, \quad (1.1)$$

where $\gamma_{f,t}$ is the year-on-year monthly growth rate of labor demand and $\bar{\gamma}_{f,t}$ is the mean growth rate computed over the $(2\omega + 1)$ -month period centered around date t . Our baseline measure uses a 35-month window, centered around January 2015. The variable is constructed using the total number of employees as our measure of labor demand. This measure of employment volatility captures second moments in the time-series of labor demand at firm-level, thus treating symmetrically upward and downward adjustments. As our focus is on the potential sources of labor market pooling, we focus on the idiosyncratic component of volatility, and thus residualize $\gamma_{f,t}$ in the sector \times month \times year dimensions.¹² However, our results are robust to using a simpler measure of employment volatility based on actual growth rates. The reason

12. Monthly growth rates of labor demand are winsorized at the 1st and 99th percentile within each of 6 firm size classes. These 6 firm size classes identify firms below 2 employees, between 2 and 9 employees, between 10 and 49 employees, between 50 and 249 employees, between 250 and 4,999 employees, and plants of 5,000 and above employees.

is that the vast majority of firm-level dispersion in volatility is driven by idiosyncratic shocks (di Giovanni et al., 2014).

In Appendix tables A.2 and A.3, we compare our baseline measure with alternatives capturing slightly different aspects of the firm’s employment volatility. While the baseline measure rests on employment, thus on adjustments at the extensive margin, we show that the correlation with the volatility of hours is high, at 0.85. Pure intensive margin adjustments, through the number of hours per employee, are not the main factor at the root of a firm’s labor demand fluctuations. Likewise, one may be concerned that certain type of contracts, most notably short-term contracts, are particularly well-suited to help the firm smooth out the impact of fluctuations in demand. The correlation of our baseline measure with a measure of volatility recovered solely from the growth of permanent contracts is however high, at 0.75. The volatility of open-ended contracts is still substantial, only 5% lower than the volatility of overall employment at the sample mean. Finally, our baseline volatility measure correlates highly with alternatives using slightly different strategies for identifying the idiosyncratic component of volatility. The most sensitive robustness check is obtained from statistics computed on month-on-month, instead of year-on-year, growth rates. Mechanically, the average volatility recovered from month-on-month growth rates is an order of magnitude smaller. However, its cross-sectional correlation with the baseline is still high at 0.70. In our baseline, we neglect month-on-month fluctuations that may to a large extent come from a sector-specific seasonality.

Descriptive statistics — Table 1.2 contains descriptive statistics on the baseline sample of firms. In January 2015, the sample is composed of 365,041 firms that we observe over at least 35 consecutive months. As expected, firms display significant heterogeneity in size, employment volatility, and productivity. Appendix Figure A.1 shows that, conditional on its size, the median firm with 2-10 employees adjusts its labor demand (up or down) by approximately 0.8 employees per month, on average.

Table 1.3 shows how our measure of employment volatility correlates with important covariates. First, as seen in column (1), older firms are less volatile, which is a standard pattern in firm-level data (Davis et al., 2007). This relationship can reflect a form of internal

Table 1.2— Distribution of employment volatility and productivity

	Employment	$\log \sigma$	$\log \phi$
Mean	10.11	-2.03	3.19
Std. Dev.	20.82	1.15	0.72
25th percentile	2.00	-2.45	2.79
50th percentile	5.00	-1.84	3.20
75th percentile	9.87	-1.31	3.61

Notes: The variables are calculated for the January 2015 cross section of the dataset. Employment is the number of employees. Productivity is based on 2015 balance-sheet data. Volatility is computed using a 35-month window centered around January 2015 and the formula in equation (1.1). Statistics are calculated on a cross-section of 365,041 firms.

diversification of risks when firms age and grow.¹³ Appendix Figure A.2 illustrates this pattern in more detail, showing that most of the age variation takes place within the first four years of a firm’s life cycle, while volatility stabilizes afterwards. In Section 1.4, we will use the firm’s volatility after at least five years of existence as a proxy for the firm’s steady state volatility.

Second, more productive firms are also less volatile (column 2). This empirical correlation is then taken into account and we systematically examine the effect of a firm’s volatility, *conditional on its productivity*. Consistent with intuition, firms with a higher share of temporary contracts are more volatile, but controlling for this share does not significantly affect the other estimates (column 3). In addition, controlling for sector fixed effects (FEs) in column (4) and CZ FEs (column 5) does not significantly increase the explanatory power of the model once the other controls are included. In column (6), we further control for the average firm growth over the 35 consecutive periods used in the calculation of employment volatility, corresponding to $\bar{\gamma}$ in equation (1.1). This increases the explanatory power of the model, but does not affect the direction of the other effects described in the previous columns.

Finally, in column (7) we report results from the same model as in column (6), but now focusing on a subset of manufacturing firms for which we have information on demand volatility, a measure that we will use in Section 1.4. Although the cross-section is much smaller, the qualitative patterns are unchanged.

13. Note that in this table we control for firm size class fixed effects for consistency with the rest of the analysis. Firm size is negatively correlated with employment volatility.

Table 1.3 — Firm employment volatility: correlates

	Dependent variable: log Employment volatility						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log Age	-0.323 (0.003)	-0.327 (0.003)	-0.311 (0.002)	-0.290 (0.003)	-0.290 (0.003)	-0.171 (0.002)	-0.137 (0.009)
log Productivity		-0.061 (0.003)	-0.047 (0.003)	-0.024 (0.003)	-0.035 (0.003)	-0.049 (0.003)	-0.051 (0.012)
% fixed-term contracts			1.291 (0.015)	1.229 (0.015)	1.221 (0.015)	0.983 (0.014)	1.604 (0.087)
Size class FEs	✓	✓	✓	✓	✓	✓	✓
Sector FEs				✓	✓	✓	✓
CZ FEs					✓	✓	✓
Average growth						✓	✓
Adjusted R^2	0.112	0.114	0.132	0.148	0.151	0.239	0.251
Sample	Full	Full	Full	Full	Full	Full	EAP
N. firms	365,041	365,041	365,041	365,041	365,041	365,041	20,419

Notes: the table shows the conditional correlation between our baseline measure of employment volatility and the firm’s age, productivity and dependence on fixed-term contracts. The table contains OLS coefficients and their estimated standard errors in parentheses.

1.2.2. Motivating stylized facts

The productivity density gradient — The literature in economic geography has long discussed agglomeration patterns of firms over space. We first reproduce the evidence focusing on the productivity-density correlation (Combes et al., 2012). More precisely, we run the following regression based on the cross-section of firms observed in January 2015

$$\log \phi_f = X_f \beta + FE_{M(f)} + \varepsilon_f, \quad (1.2)$$

where X_f is a set of controls and $FE_{M(f)}$ denotes a set of FEs for each CZ. In this equation, the FEs captures the average productivity of firms in any CZ, once controlling for the heterogeneity that correlates with the control variables, namely the firm’s 2-digit sector of activity, its size class and age.

Figure 1.1 illustrates the correlation between the conditional average productivity of firms and the population density of the CZ. As expected, the correlation is positive and significant, consistent with the view that dense CZs attract more productive firms, on average. As




Figure 1.1 — The productivity advantage of large cities

Notes: The figure shows the correlation between the mean productivity of firms and the density of the CZ where firms locate. Mean productivity is based on 2015 balance-sheet data. The correlation is conditional on the following firm characteristics: sector, size class, age, average growth during the period over which employment volatility is computed ($\bar{\gamma}$). The slope is 0.042 (the adjusted R^2 is 0.3847) and the slope is significantly different from 0 at the 1%.

mentioned in the introduction, there is a vast literature explaining the correlation using various theoretical frameworks. A strand of the literature notably points to the role of matching economies through pooling externalities: locations with higher meeting rates are most beneficial to high-productivity firms that are able to hire more quickly (Bilal, 2023). To the extent that pooling externalities are part of the story, we shall expect that the benefit is also larger for more volatile firms, conditional on productivity. As shown in Section 1.2.1, firms are indeed strongly heterogeneous in terms of the volatility of their labor demand, which may thus affect spatial location patterns.

The volatility density gradient — We provide preliminary evidence for a role of employment volatility in Figure 1.2. As in Figure 1.1, we first recover an estimate of firms' average employment volatility at the CZ level. We run a regression similar to equation (1.2), using the log of employment volatility as the dependent variable. We then correlate this




Figure 1.2 — The volatility advantage of large cities

Notes: The figure shows the correlation between the mean volatility of firms and the density of the CZ where they locate. Volatility is measured by the standard deviation of the firm’s labor demand year-on-year growth. Mean volatility is based on the January 2015 cross-section of firms and is conditional on the following firm characteristics: sector, size class, firm age, firm average growth ($\bar{\gamma}$), log productivity. The slope is 0.023 (the adjusted R^2 is 0.0808) and the slope is significantly different from 0 at the 1%.

measure for conditional average employment volatility with the density of the CZ. Here as well, the conditional correlation is positive and significant, consistent with the intuition that pooling externalities are particularly valuable for volatile firms, which may then agglomerate in denser CZs. Importantly, the set of controls now includes the firm’s productivity, which implies that a positive correlation exists beyond and above the productivity-density nexus that the literature before us has documented.¹⁴

The productivity density gradient by volatility — Finally, Figure 1.3 provides a third motivating stylized fact that directly tackles the joint correlation between density, employment volatility and productivity. Instead of recovering the correlation between firms’

14. This correlation survives if we use the volatility of monthly growth rates instead of measuring the variance of year-on-year growth rates as we do in the rest of the paper.




Figure 1.3 — The productivity-density gradient, along the distribution of volatility

Notes: The figure shows the conditional correlation between log productivity of firms and the density of the CZ where they locate, along the distribution of firms' log employment volatility. Productivity is conditional on the following firm characteristics: sector, size class, firm age and firm average growth of employment. The estimated equation includes the log density of the CZ where the firm is located, log employment volatility, and the interaction of log density and log employment volatility. Volatility is measured by the standard deviation of the firm's residualized labor demand growth. Data is based on the January 2015 cross-section of firms. Whiskers indicate 95% confidence intervals.

attributes and the density of the firm's CZ in two stages, we now directly introduce density in equation (1.2). The downside is that we can no longer control for unobserved heterogeneity between CZs using FEs. However, we can now interact density with a measure of the firm's employment volatility to estimate how the productivity-density correlation varies depending on the firm's volatility. The coefficient on the interaction is negative and strongly significant, which implies that the tendency of high-productivity firms to agglomerate in dense cities is less pronounced within the set of more volatile firms. Quantitatively, the cross-correlation is non-negligible: the elasticity of firms' average productivity to density drops by a third when moving from the first to the last ventile of the volatility distribution.

Overall, the evidence in this Section confirms that denser cities attract a pool of firms that are systematically different from the rest of the population in terms of their productivity

but also the volatility of their labor demand. In the next Section, we build a model that helps understand these agglomeration patterns.

1.3 — Volatility and firm location: theory

We lay out a simple model of the impact of volatility on firms' location decisions. The model provides a micro-foundation of employment volatility based on firms' hiring and firing decisions and helps understand the trade-offs associated with firms' location choice: in particular, it shows why some firms may prefer locating in a denser city, even if that means operating under higher operating costs. The model's main prediction reads as follows: if firms sort across space based on their structural volatility because hiring is faster in denser cities, employment volatility will increase with density and the productivity-density gradient will be lower for firms with higher volatility. Expressions and proofs are provided in Appendix A.3.

1.3.1. Framework

We consider a simplified version of the canonical search-and-matching model proposed by Mortensen and Pissarides (1994), where single-job, risk-neutral, profit-maximizing firms face sales shocks and hiring frictions. The economy operates at a steady state and time is continuous. We focus on a partial equilibrium, leaving the worker problem aside. In particular, workers are homogeneous, their location is fixed, they do not search when employed, and they do not bargain over wages. We also make the simplifying assumption that firms cannot adjust their labor demand at the intensive margin, by paying overtime or using part-time contracts. As discussed in Section 1.2.1, the extensive margin is a quantitatively important source of volatility at firm-level.

Set-up — Firms are heterogeneous in terms of their mean productivity $\phi > 0$ and their volatility $\varepsilon \in [0, 1]$, both known ex ante and independent from each other. We assume that firms are price takers and cannot adjust their price to sales shocks.¹⁵ If we normalize price to

15. See Section 1.3.4 for the discussion of an extended model with an explicit formulation of entry and demand. While in our base setup presented here demand and productivity shocks may be homeomorphic,

1, this means that sales fluctuate in any period between $\phi(1 + \varepsilon)$ in the high state ($t = h$) and $\phi(1 - \varepsilon)$ in the low state ($t = l$) at an exogenous rate ξ that measures the structural volatility of the economy.

Upon entry, firms choose a location or city defined by a density $M > 0$. City choice determines firms' operating costs $R(M) \geq 0$ and job-filling rate $\mu(M) \geq 0$. $R(M)$ is a local index that combines all costs associated with maintaining an active position.¹⁶ Importantly, firms that are not actively producing do not have to pay these costs. For example, if $R(M)$ represents the price of renting capital or real estate, this assumption means that there are no frictions on the capital market. We further assume that $R(M)$ and $\mu(M)$ are both increasing in M .¹⁷

Strategies — Conditional on their location, firms also choose a strategy s , which in this context corresponds to a specific action to take in the low state. Firms can choose between three strategies $s \in \{B, W, C\}$. According to the “Business as usual” strategy (hereafter, denoted by B), if a firm is hit by a bad shock, it will keep paying its workforce or it will keep trying to hire. However, if operating costs are too high, the firm will seek to mitigate them by limiting the amount of time spent active in the low-production state. According to the “Wait-and-see” strategy (hereafter, denoted by W), if an active firm is hit by a bad shock, it will keep paying its workforce and wait for better times; yet, vacant firms, when hit by a bad shock, will postpone hiring until they have reached a high state again. Finally, according to the “Churning” strategy (hereafter, denoted by C), if a firm is hit by a bad shock, it will become idle. This means that it will wait if it is vacant, and it will fire and wait if it is active.

Recursive formulation — Given their choice of city and strategy (M, s) , firms alternate between being vacant (V), active (A) or idle (I). They decide whether to operate or hire

the extended model targets directly demand shocks that are the focus of our empirical work and eschews productivity shocks.

16. It may encompass wages, but those do not depend on firms' individual characteristics (ϕ, ε) in order to keep the focus on hiring decisions.

17. While $R'(M) > 0$ is easily justified by a congestion argument (commuting costs, inelastic housing supply), the sign of $\mu'(M)$ is more contentious because it depends directly on how many firms there are in each location, and what they do. For ease of exposition, we describe the framework in partial equilibrium, whereby those two local factors are not impacted by firms' decisions, and we defer the discussion on the endogenous determination of $\mu(M)$ to Section 1.3.4.

while in a low state or not and this determines the firm's transition to a low state when posting a vacancy (with value $W_s(\phi, \varepsilon, M)$) or when filled (with value $C_s(\phi, \varepsilon, M)$). For any strategy s , firms' value functions are thus summarized as follows

$$\begin{aligned} rV_s^h(\phi, \varepsilon, M) &= -c + \mu(M)[A_s^h(\phi, \varepsilon, M) - V_s^h(\phi, \varepsilon, M)] \\ &+ \xi[W_s(\phi, \varepsilon, M) - V_s^h(\phi, \varepsilon, M)] \end{aligned} \quad (1.3)$$

$$\begin{aligned} rV_s^l(\phi, \varepsilon, M) &= -c + \mu(M)[A_s^l(\phi, \varepsilon, M) - V_s^l(\phi, \varepsilon, M)] \\ &+ \xi[V_s^h(\phi, \varepsilon, M) - V_s^l(\phi, \varepsilon, M)] \end{aligned} \quad (1.4)$$

$$\begin{aligned} rA_s^h(\phi, \varepsilon, M) &= \phi(1 + \varepsilon) - R(M) + \delta[V_s^h(\phi, \varepsilon, M) - A_s^h(\phi, \varepsilon, M)] \\ &+ \xi[C_s(\phi, \varepsilon, M) - A_s^h(\phi, \varepsilon, M)] \end{aligned} \quad (1.5)$$

$$\begin{aligned} rA_s^l(\phi, \varepsilon, M) &= \phi(1 - \varepsilon) - R(M) + \delta[W_s(\phi, \varepsilon, M) - A_s^l(\phi, \varepsilon, M)] \\ &+ \xi[A_s^h(\phi, \varepsilon, M) - A_s^l(\phi, \varepsilon, M)] \end{aligned} \quad (1.6)$$

$$rI_s(\phi, \varepsilon, M) = \xi[V_s^h(\phi, \varepsilon, M) - I_s(\phi, \varepsilon, M)], \quad (1.7)$$

where r is the interest rate, c is the vacancy cost and δ is the exogenous component of the match destruction rate. Both c and δ are assumed to be fixed over time and constant across firms. Strategies determine the values of either posting a vacancy in the low state ($W_s(\phi, \varepsilon, M)$) or being active in the low state ($C_s(\phi, \varepsilon, M)$), as summarized in Table 1.4.

Table 1.4 — Strategies and values of low state

	$W_s(\phi, \varepsilon, M)$	$C_s(p, \varepsilon, M)$
Business as usual	$V_B^l(\phi, \varepsilon, M)$	$A_B^l(\phi, \varepsilon, M)$
Wait-and-see	$I_W(\phi, \varepsilon, M)$	$A_W^l(\phi, \varepsilon, M)$
Churning	$I_C(\phi, \varepsilon, M)$	$I_C(\phi, \varepsilon, M)$

Entry, location choice and employment volatility — Since firms do not know in which state they will enter nor the state in any other period after entry, their expected profit at entry is given by $\mathbb{E}_s(\phi, \varepsilon, M) = 0.5 \times [V_s^h(\phi, \varepsilon, M) + W_s(\phi, \varepsilon, M)]$. Conditional on location, the *preferred strategy* s^* is thus the one that maximizes expected profit: $s^*(\phi, \varepsilon, M) = \underset{s}{\operatorname{argmax}} [\mathbb{E}_s(\phi, \varepsilon, M)]$.

For ease of exposition, we normalize the outside option to zero. Note that even preferred strategies may not be adopted if they yield a negative expected profit. In that case, the firm does not enter. Finally, under some conditions (detailed below), the model delivers a mapping $M^*(\phi, \varepsilon)$ between firms' characteristics and location

$$M^*(\phi, \varepsilon) = \begin{cases} \operatorname{argmax}_M [\mathbb{E}_{s^*(\phi, \varepsilon, M)}(\phi, \varepsilon, M)] & \text{if } \mathbb{E}_{s^*(\phi, \varepsilon, M^*(\phi, \varepsilon))}(\phi, \varepsilon, M^*(\phi, \varepsilon)) \geq 0 \\ \{\emptyset\} & \text{otherwise.} \end{cases} \quad (1.8)$$

Productivity ϕ and volatility ε , together with strategy s and location M determine volatility of employment $\sigma^l(\phi, \varepsilon, M^*(\phi, \varepsilon), s^*(\phi, \varepsilon))$. The model predicts that under reasonable parametric conditions, churning may indeed be associated with higher employment volatility, as summarized in Proposition 1.¹⁸

Proposition 1. *Churning and employment volatility* —Firms that adopt the churning strategy have a higher employment volatility if the structural volatility of the economy is low enough.

1.3.2. Solution

Firms jointly choose s and M . Yet, for exposition purposes, we solve the model in three steps. First, we detail how firms' characteristics determine their strategy choice, for a given location. Then, we compare strategy choices between different cities. Finally, we solve the general model.

Strategy choice — If we solve the system (1.3-1.7), we can make two observations: first, quite naturally, expected profit increases with productivity, regardless of the strategy; second, higher productivity is more profitable under strategy B than under strategy W , and under strategy W than under strategy C . Therefore, strategy choice is determined by

18. For high values of ξ , the model features the degenerate prediction that churning firms will mostly oscillate between the idle and the vacant states, with low associated volatility.

five productivity cutoffs: three *selection cutoffs* that determine whether a given strategy is *feasible*, and two *switching cutoffs* that determine which strategy is preferred.

Strategy s is feasible for a type- (ϕ, ε) firm if ϕ is greater than the selection cutoff $\phi_s(\varepsilon, M)$. Under strategy B , the selection cutoff $\phi_B(\varepsilon, M)$ does not depend on ε and may therefore be denoted $\phi_B(M)$. As is usual in this type of models, sales must cover both operating costs and the vacancy cost at entry and following any exogenous separation. Under strategy W , the selection cutoff $\phi_W(\varepsilon, M)$ is lower than under strategy B if $\varepsilon > 0$, and it decreases with ε . This strategy can therefore accommodate more volatile firms that have lower productivity in the low state compared to less volatile firms: by waiting, the firm mitigates the consequences of being in the low state. Finally, under strategy C , the selection cutoff $\phi_C(\varepsilon, M)$ is even more sensitive to ε than under strategy W : $\partial\phi_C(\varepsilon, M)/\partial\varepsilon < \partial\phi_W(\varepsilon, M)/\partial\varepsilon$. However, the selection cutoff also entails a fixed cost $c\xi/\mu(M)$, which corresponds to the additional time spent vacant. Therefore, only highly volatile firms may be able to churn. In particular, churning only allows for the entry of less productive firms if their volatility exceeds a given cutoff $\tilde{\varepsilon}(M)$, which depends on both local and common parameters.

We then turn to the conditions that determine when firms adopt a churning strategy over alternative strategies. In what follows, an *adopted strategy* is both preferred and feasible. We denote by $\phi_{BW}(\varepsilon, M)$ and $\phi_{WC}(\varepsilon, M)$, with $\phi_{BW}(\varepsilon, M) > \phi_{WC}(\varepsilon, M)$, the corresponding cutoffs. Both cutoffs, as well as the difference between the two, are convex increasing functions of ε . Regarding the B strategy, we can note that $\forall\varepsilon, \phi_{BW}(\varepsilon, M) > \phi_B(M)$. Therefore, if strategy B is preferred, it is also feasible, and therefore, adopted. Conversely, strategies W or C may be preferred, yet unfeasible, if $\phi_{WC}(\varepsilon, M) < \phi_W(\varepsilon, M)$ or $\phi_{WC}(\varepsilon, M) < \phi_C(\varepsilon, M)$. We use a calibration to illustrate the working of the model.¹⁹ The productivity cutoffs are represented in Panel A in Figure 1.4.

Equipped with these definitions, we can fully characterize the distribution of adopted strategies as a function of ϕ and ε . They are represented in Panel B in Figure 1.4 in the form

19. This calibration is somewhat arbitrary, even though we aim for realism for some aspects. The time unit is a year and we set $r = 3\%$. The match destruction rate δ is set to 10%, and the probability of switching between high and low demand states is set to 20%. The vacancy cost is set to 10% of a maximum productivity level $\bar{\phi}$, which is set to 1. We consider a cost function $R(M) = 0.2M^{0.1}$. This 10% elasticity stems from the addition of the 3% of urban costs calibrated by Combes et al. (2019) and 7% elasticity of raw wages (Ahlfeldt and Pietrostefani, 2019). Finally, we consider a worker finding rate given by $\mu(M) = 0.3M^{0.05}$.

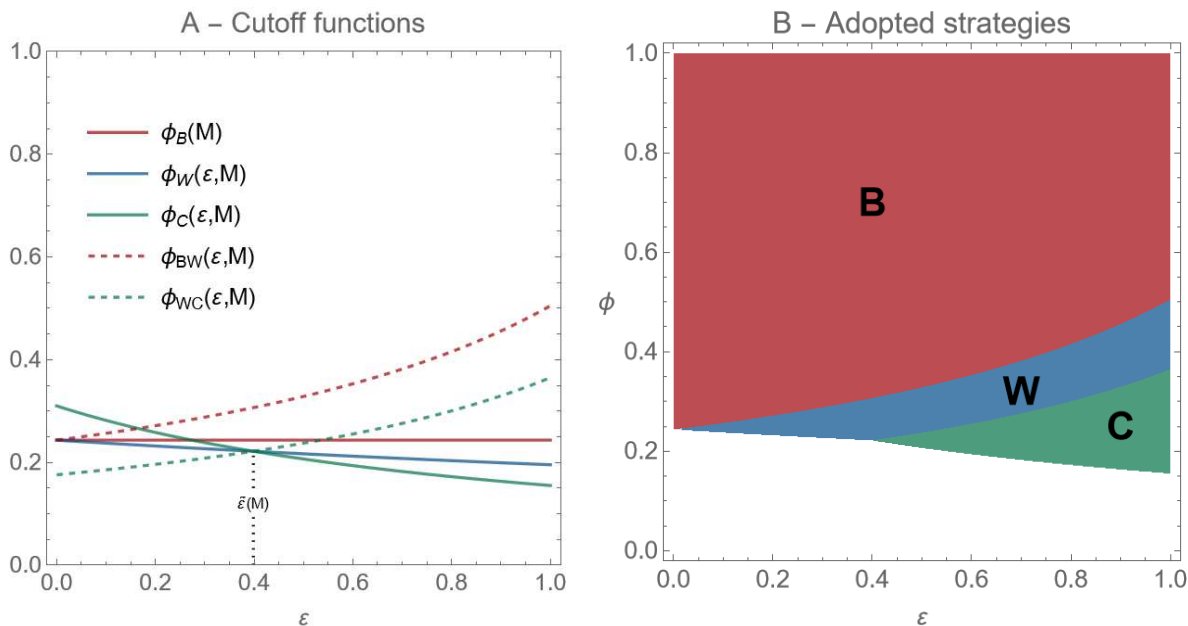


Figure 1.4 — Strategy choice for a given city

Notes: Calibration: $\xi = 0.2$, $r = 0.03$, $\delta = 0.1$, $\mu(M) = 0.3M^{0.05}$, $R(M) = 0.2M^{0.1}$, $c = 0.1$ and $\bar{\phi} = 1$. We set $M = 1$. Panel A: The figure represents the three minimum productivity cutoffs and the two strategy-switching cutoffs as a function of volatility ε . Panel B: The figure represents the set of (ε, ϕ) combinations associated with each adopted strategy. The blank section corresponds to combinations that are not feasible, regardless of the strategy.

of the three regions labeled B, W, and C.²⁰ Panel B highlights our first two key results that hold for a fixed city, as summarized in Proposition 2.

Proposition 2. Strategy choice — In a given city,

2.1 Churning is adopted by more volatile, less productive firms.

2.2 Very volatile firms may churn even if they are quite productive. Conversely, low-productivity firms may be able to operate if they are volatile enough.

The joint strategy/location problem — The next step is understanding how density interacts with firms' productivity, volatility, and strategy choice. To proceed, we make three further assumptions:

20. Note that one strategy may never be adopted, depending on the parameters. In particular, W disappears when $c \rightarrow 0$. Conversely, C disappears for large enough values of c .

Assumption 1. Churning happens in equilibrium.²¹

Assumption 2. For each strategy, selection on productivity does not decrease with density.²²

Assumption 3. For each strategy, there exists an optimal level of density.²³

These assumptions restrict the analysis to cases where the model is both relevant (Assumption 1), realistic (Assumption 2), and analytically well-defined (Assumption 3). Under these assumptions, we can perform comparative statics of strategy choice under different city sizes, which yields the following results, summarized in Proposition 3.

Proposition 3. *Comparative statics* — If cities are heterogeneous in density,

3.1 Denser cities have a higher share of churning firms.

3.2 Low-productivity firms are more volatile in denser cities.

Results 3.1 and 3.2 can also be gauged by comparing adopted strategies in the (ε, ϕ) plane for different levels of density, as we do in Panel A of Appendix Figure A.6. In line with result 2.2, even productive firms may churn in denser cities if they are very volatile. In addition, higher volatility is more conducive to the entry of low-productivity firms, as shown by a steeper lower bound of the colored area.

Finally, we study the firm location choice, and how churning interacts with the spatial sorting of firms based on their productivity. This requires solving a global maximization problem, to identify the density chosen by firms, conditional on their productivity ϕ and volatility ε . While the combinations of (ϕ, ε) associated with strategy choice are only defined

21. This assumption is verified under the condition $\bar{\varepsilon}(M) < 1$, which is equivalent to $c/\mu(M) < R(M)/\xi$. In words, this means that the expected vacancy cost is lower than the operating costs paid by the firm when it is operating in the low state. For simplicity, we will even assume a stronger condition, stating that $\forall M \geq 0, R(M) > c$ and $\mu(M) > \xi$. Note that this assumption means that both $R(M)$ and $\mu(M)$ feature a fixed positive component, or that there is a lower bound for density, as we do in our calibration.

22. The most binding condition is for strategy C , where it is equivalent to: $\forall M \geq 0, R'(M) \geq (r + \delta + \xi)\mu'(M)/\mu(M)^2$. For simplicity, and using Assumption 1, we will even assume a stronger condition on the ratio of the elasticity of each function: $\forall M \geq 0, \epsilon_{R,M}/\epsilon_{\mu,M} > (r + \delta + \xi)/\xi$.

23. This assumption means that $\forall s, \exists M > 0$ s.t. $\partial \mathbb{E}_s(\phi, \varepsilon, M)/\partial M = 0$. As for Assumption 2, this assumption will be met if the ratio of the cost elasticity to the matching elasticity is high enough.

implicitly, the envelope theorem ensures that Proposition 2 is robust to firms' location choice. Panel B in Appendix Figure A.6 illustrates this result. In particular, more volatile firms are more likely to adopt the churning strategy, more productive firms are more likely to adopt the business-as-usual strategy, and low-productivity, high-volatility firms are more likely to be able to operate if they adopt the churning strategy.

1.3.3. Volatility and the sorting of firms

This framework allows us to study how the joint strategy/location optimization problem at the individual firm level translates into aggregate sorting patterns of firms across space. Conditional on selection and strategy choice, the spatial sorting of firms is implicitly defined by the optimal productivity/volatility-density relationship described by *sorting cutoffs* $\phi_s^*(\varepsilon, M) = \underset{\phi}{\operatorname{argmax}} [\mathbb{E}_s(\phi, \varepsilon, M)]$ and $\varepsilon_s^*(\phi, M) = \underset{\varepsilon}{\operatorname{argmax}} [\mathbb{E}_s(\phi, \varepsilon, M)]$.

These sorting cutoffs illustrate how matching economies work in this model. Since high-productivity and high-volatility firms have more to gain from being able to hire more quickly, there is positive sorting with respect to productivity and volatility, for a given strategy. In addition, given the multiplicative structure between ϕ and ε , high-productivity (respectively, high-volatility) firms have all the more to gain from locating in denser cities if they are more volatile (respectively, productive). Formally, we have $\forall \varepsilon \in [0, 1], \partial \phi_B^*(\varepsilon, M)/\partial M \geq \partial \phi_W^*(\varepsilon, M)/\partial M \geq \partial \phi_C^*(\varepsilon, M)/\partial M$. Therefore, even if more productive firms sort into denser cities, the share of churning firms also increases with density and the productivity-density gradient decreases with firm volatility. These patterns are summarized in Proposition 4.

Proposition 4. *Predictions* — If firms choose their location in order to maximize their expected profit upon entry,

4.1 More productive and more volatile firms sort into denser cities.

4.2 Productivity and volatility are complementary in city choice.

4.3 The share of churning firms increases with density and the productivity-density gradient

is flatter for more volatile firms.

Prediction 4.3 echoes the aggregate sorting patterns described in Section 1.2. Appendix Figure A.7 illustrates this prediction for the same calibration of the model, for a density spanning between 1 and 10.²⁴ Panel A, consistent with Figure 1.2, displays the share of churning firms as an increasing function of density. In Panel B, we measure the productivity-density gradient along the distribution of firm volatility. Consistent with Figure 1.3, this gradient decreases with firm volatility. As for predictions 4.1 and 4.2, they will be tested in Section 1.4.

1.3.4. Discussion

In order to maintain analytical tractability, the model rests on several simplifying assumptions. We briefly discuss here the robustness of its conclusion to a more general framework.

Endogenous matching rate — In presenting the model, it was assumed that the worker meeting rate was not affected by firms' location decisions and strategies. However, such assumption is not internally consistent, because the worker meeting rate depends on the local market tightness, which is, itself, an equilibrium outcome. A priori, the impact of churning on market tightness is ambiguous. On the intensive margin, churning firms lay off workers while remaining idle, which loosens the market; conversely, churning may also allow more firms to enter. To recover market tightness, we need to define a fixed point problem that describes steady-state conditions and to specify a process for firm entry, in order to determine the equilibrium firm-to-worker ratio. In Appendix A.3.4, we describe a way of tackling this extended model. Using simulations, we show that, under plausible parametric assumptions on the matching technology, the resulting worker meeting rate is an increasing concave function of density, even if the positive effect of churning on firm entry mitigates the magnitude of agglomeration economies.

24. Note that we need to specify the underlying distributions of ϕ and ε . For simplicity, we assume that they are both uniformly distributed over $[0, 1]$.

Firm size heterogeneity and demand — Up to now, we have not modeled demand nor allowed for firm size heterogeneity linked with productivity. In Appendix A.3.5, we embed this model in a framework à-la Melitz (2003), where monopolistically competitive firms face a CES demand system and draw heterogeneous productivity and demand volatility upon entry. This extended model allows us to better understand the underlying differences in the behavior of firms that face demand shocks and either adopt the Business-as-usual or the Churning strategy. Under the former, firms adjust their prices, while under the latter, prices are independent from individual demand shocks, which are then passed on employment. We show that under plausible parametric restrictions, the main predictions of the base model carry through: churning makes it possible for lower productivity firms to enter, the share of volatile firms increases with density, and the productivity-density gradient is flatter for volatile firms.

1.4 — Volatility and firm location: empirical evidence

In this Section, we turn back to our data and describe two tests of the main predictions of the model. These tests rely on different assumptions and proxies for firm characteristics introduced in Section 1.3.

1.4.1. Empirical strategy

Our empirical framework is based on a location choice model estimated with a conditional logit estimator (CLM). We start with the January 2015 cross-section and select firms that are between 5 and 10 years old in January 2015.²⁵ This leaves us with a sample of 88,168 firms.

While the theoretical model assumed a continuum of densities, we now consider a discrete set of locations $\mathcal{M} = \{M\}$. Conditional on the firm's decision to enter the French market, we model the choice of a location as a function of the firm's and the location's attributes. We

25. The lower age limit excludes firms that are still in the early stages of their life cycle in 2015, when their employment volatility is a poor measure of their long-term (structural) volatility, as discussed in Section 1.2.1 and illustrated in Appendix Figure A.2. The upper age limit is more arbitrary. It is intended to homogenize the sample with respect to the broader economic context in which the location decisions were made.

borrow the notations from Section 1.3 and denote $\mathbb{E}^*(\phi_f, \varepsilon_f, M) = \mathbb{E}_{s^*(\phi_f, \varepsilon_f, M)}(\phi_f, \varepsilon_f, M)$ for brevity. Assuming that the expected inter-temporal profit in each location can be decomposed into a deterministic and a random component e_{fM} , one can write the probability of a firm f choosing a location M as

$$\begin{aligned} \mathbb{P}(f \text{ chooses } M | e_{fM}) &= \mathbb{P}\left(\mathbb{E}^*(\phi_f, \varepsilon_f, M) + e_{fM} > \max_{M' \neq M} \{\mathbb{E}^*(\phi_f, \varepsilon_f, M') + e_{fM'}\}\right) \\ &= \frac{\exp[\mathbb{E}^*(\phi_f, \varepsilon_f, M)]}{\sum_{M' \in \mathcal{M}} \exp[\mathbb{E}^*(\phi_f, \varepsilon_f, M')]} \end{aligned}$$

where the second line uses the assumption that e_{fM} are i.i.d. draws from a type-1 extreme value distribution.

Our model predicts the choice between all commuting zones (CZs) to be a function of the size of operating costs $R(M)$ and the job-filling rate $\mu(M)$ as well as their interaction with firms' productivity ϕ_f and volatility ε_f . Following the theoretical model, the CLM considers the role of CZ density, and its interaction with firms' characteristics, TFP and volatility.²⁶ We obtain productivity and employment volatility from the January 2015 cross-section, focusing on the same variables that we have analysed in Section 1.2.2. We also control for other CZ characteristics that are important for firm location decisions, namely two measures of workforce skill (the share of managers, and the share of college graduates).²⁷ In order to verify that other local characteristics correlated with density do not drive our results, we test the robustness of our estimates to the inclusion of CZ FEs.

Localization economies — The identification assumption behind the CLM is that other firm characteristics that are correlated with volatility and productivity do not interact with density in determining location choice. However, this assumption is unlikely to be true in general. In particular, there is ample evidence that firms benefit from having other firms in the same industry operating in the same area (Combes and Gobillon, 2015). Therefore, we

26. In all rigor, employment volatility should not be denoted by ε , but by σ . We introduce a potential candidate for ε in Section 1.4.3.

27. The share of managers is calculated from INSEE-DADS, where managers are defined by 1-digit occupation (CS1) equal to 2 or 3. The share of college graduates is obtained from Census data. Both of these variables are measured in 2015.

also control for a measure of localization economies measured in 2015. Localization economies measure a firm’s sectoral network based on Mayer et al. (2010). This sectoral network is calculated as the total number of firms in the same industry located in each potential CZ in the year y corresponding to the chosen cross-section.²⁸ More precisely, localization economies are defined as:

$$SectoralNetwork_{i,y}^s = \sum_y \sum_a D_{ai,y}^s \quad (1.9)$$

where $D_{ai,y}^s$ is a dummy variable equal to one for all firms a of sector s located in CZ i and created in year y . The count of firms in each sector and CZ only includes firms with positive employment. Sectors are defined at the 2-digit level of the French sector nomenclature.²⁹

1.4.2. Results

Sorting on volatility and productivity — Estimation results are summarized in Table 1.5. In column (1), we show the coefficient associated with the (log of) density of the CZ, and we confirm the tendency of firms to agglomerate in denser CZs, even after controlling for other CZ characteristics and our measure of localization economies. In columns (2) and (3), we then interact density with the model’s relevant firms characteristics, namely productivity and volatility. Column (2) confirms previous results in the literature, showing that more productive firms are more likely to locate in denser cities. In column (3), we find that volatile firms are also more likely to locate in dense cities. In columns (4), we simultaneously consider the two interaction terms. Finally column (5) further controls for CZ FEs and solely identifies the coefficients on the interaction terms.

Results point to a quantitatively similar impact of productivity and volatility on location patterns, which is stable across specifications. Namely, in our baseline model (column 4) the elasticity of the odds of choosing a specific location to the density of this CZ increases from 0.41 to 0.53 when moving from the first to the ninth decile of the distribution of productivity. The effect is similar (from 0.42 to 0.52), when moving along the distribution of employment

28. In the sample used in Section 1.4.3, we cumulate the sectoral network variables over the years preceding the firm’s creation.

29. Nomenclature d’activités française – NAF rév. 2.

Table 1.5 — Results of the location choice model

	Dependent variable: CZ choice						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
CZ density M	0.464 (0.004)	0.465 (0.004)	0.464 (0.004)	0.464 (0.004)		0.464 (0.004)	
- × Productivity		0.064 (0.002)		0.064 (0.002)	0.061 (0.002)	0.065 (0.002)	0.061 (0.002)
- × Volatility			0.041 (0.002)	0.042 (0.002)	0.040 (0.002)	0.042 (0.002)	0.040 (0.002)
- × Volatility × Productivity						0.005 (0.002)	0.004 (0.002)
CZ characteristics	✓	✓	✓	✓		✓	
Localization economies	✓	✓	✓	✓	✓	✓	✓
CZ FEs					✓		✓
Pseudo R^2	0.1353	0.1364	0.1357	0.1369	0.1569	0.1369	0.1569
Observations	25M	25M	25M	25M	25M	25M	25M

Notes: Coefficient estimates from a conditional logit model with firm FEs. The sample is based firms in the January 2015 cross-section with age 5-10 years (88,168 firms, resulting in $N = 24,687,040$ observations). M is the log of CZ density, volatility is the standardized value of employment volatility, and productivity is the standardized value of productivity. Standard errors in parentheses. All estimates are statistically significant at the 5% level.

volatility. These results are robust to the inclusion of CZ FEs (column 5) and suggest that prediction 4.1 is verified in the data.

The heterogeneity in the determinants of location choices along the distribution of firms is further illustrated in Figure 1.5, which compares the predicted probabilities of locating in a particular CZ, for firms at the 75th percentile relative to the 25th percentile of the distribution of firms' productivity and employment volatility. The patterns recovered from heterogeneous productivity and volatility are very similar. In both cases, the conditional location probabilities are roughly equal at a density of around 150, which corresponds to the level observed in CZs in the top 25th percentile of the population density distribution. Above this level, both high productivity and high volatility firms are more likely to locate in denser cities. However, this figure also suggests that sorting remains higher along the productivity dimension, especially for the highest level of density, which corresponds to the city of Paris.

Combined effects — Finally, we turn to Prediction 4.2, whereby productivity and volatility are complementary in firms' location choices. In other words, more productive (respectively,

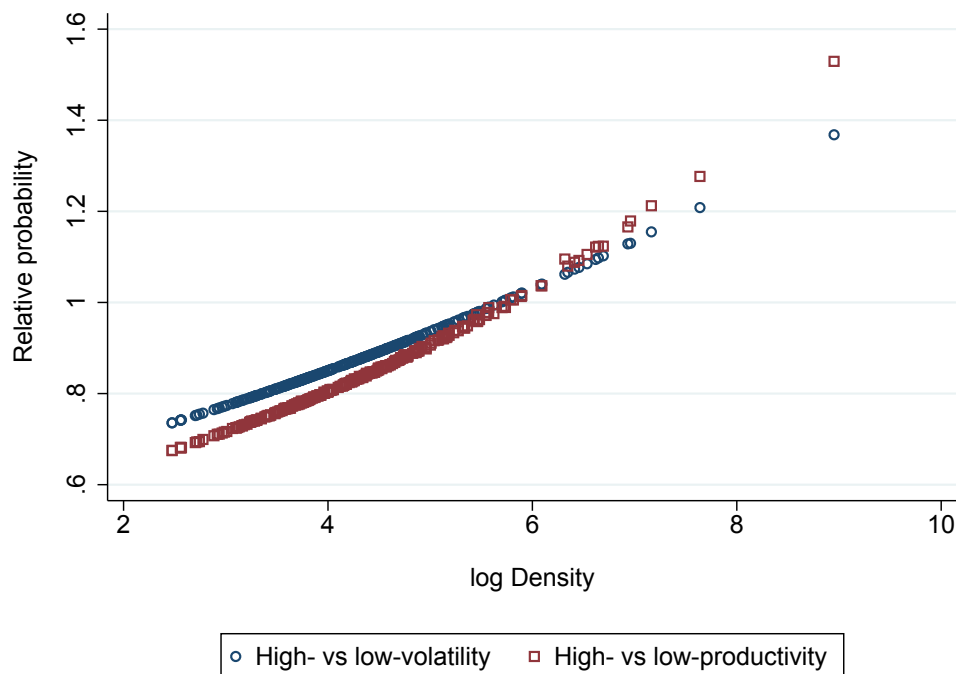


Figure 1.5 — Heterogeneity in location choices, along the distributions of productivity and volatility

Notes: The figure shows the mean probability of locating in each CZ, for high-productivity (respectively, high-volatility) firms in relative terms with low-productivity (respectively low-volatility) firms. The cut-offs are based on firms at the 25th and 75th percentile of each distribution. The probabilities are recovered from the estimation of the model in column (4) of Table 1.5.

volatile) firms are all the more likely to sort into denser locations when they are more volatile (respectively, productive). This prediction stems from the multiplicative structure between ϕ and ε in the model. Therefore, by looking at the impact of the interaction between firm productivity and volatility on firms' location choice, we can gauge the importance of complementarity between those two dimensions.

In Column (6) of Table 1.5, we estimate the impact of the triple interaction between CZ density, firm productivity, and firm demand volatility. The coefficient on the triple interaction is positive and significant and remains so when we further control for CZ FEs (column 7). This suggests that Prediction 4.2 is verified in the data. However, the estimated coefficient is small. In order to get a sense of the magnitude of this complementarity, we compute the elasticity of the odds of locating in a given CZ to CZ density, for different quantiles of productivity and volatility. Results are displayed in Figure 1.6. When going from the

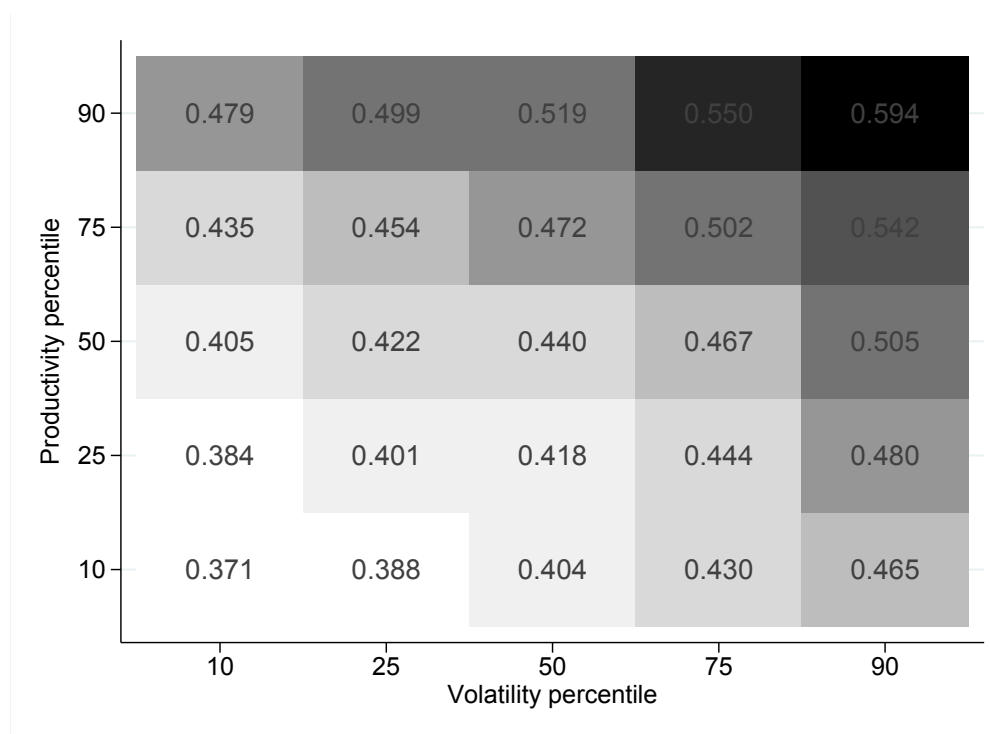


Figure 1.6 — Elasticity of the odds of choosing a CZ to CZ density

Notes: Each cell is computed (in %) as $\exp\{0.01[\hat{\beta}_M + \hat{\beta}_{M\varepsilon}P_\varepsilon + \hat{\beta}_{M\phi}P_\phi + \hat{\beta}_{M\varepsilon\phi}P_\varepsilon P_\phi]\} - 1$. Estimates from model (6) of Table 1.5.

first to the ninth decile of productivity (respectively, volatility) the elasticity increases by 29% (respectively, 25%) for firms at the first decile of volatility (respectively, productivity). Conversely, when going simultaneously from the first to the ninth deciles of both productivity and volatility the elasticity increases by 60%, which is slightly higher than the sum of the two marginal effects.

1.4.3. Robustness: firm birth and demand volatility

One may be concerned that the previous results are biased for two reasons. First, by selecting firms that are old enough to have reached their long-term employment volatility, the analysis does not distinguish between location decisions upon entry and firm survival, which we do not model. In order to verify whether entry matters, one needs to focus on the event of firm creation. Second, and related, using employment volatility as a factor influencing the location decision is contentious, since it is endogenous to location choices. As argued in

Section 1.3, employment volatility varies with the firm’s strategy choice, which depends on the joint impact of the firm’s productivity and *structural* volatility, both directly and indirectly through its impact on location decisions. Conditional on these structural characteristics, a dense location may actually cause an increase in employment volatility. Finally, our measure of unemployment volatility does not allow us to distinguish between job quits and layoffs, which might be an issue if the former are correlated with the interaction of firm characteristics with density — for example, if low-productivity firms in larger cities have a harder time retaining some of their workers. We thus test the robustness of our results to using a different measure of volatility.

Our proposed alternative measure of firm volatility uses exogenous variations in export demand for products in the firm’s portfolio. Therefore, we focus on the volatility of demand that a firm can expect to face, conditional on the nature of its production. Given the structure of the firm’s product portfolio, which we recover from the EAP survey, the volatility of demand in foreign markets can be used as a proxy for the firm’s expected demand volatility. To further rule out endogeneity, we use the firm’s portfolio of products at entry, i.e. no later than two years after its creation, as the firm’s decision to expand its product scope may also be endogenous. For consistency, we also compute productivity over this two-year period. Details are available in Section A.1.3 in the Appendix. As the EAP survey is not available before 2009, we have to restrict the cross-section to firms that were born posterior to 2009. To increase sample size, we include all firms born until 2018, and end up with a sample of 1,682 firm creations.

Results are displayed in Appendix Table A.4, which follows the same structure as Table 1.5. All results are confirmed in this smaller subset of location decisions. Consistent with prediction 4.1, both productivity and demand volatility are significant predictors of firms’ location choices, although the effect of volatility is less precisely estimated. In addition, their predictive power is still quantitatively equivalent.³⁰ Appendix Figure A.4 shows that the resulting sorting patterns are also quantitatively similar. Finally, consistent with prediction 4.2, the coefficient on the triple interaction is positive and significant, confirming that productivity

30. The elasticity of the odds of choosing a specific location to the density of this CZ increases from 0.30 to 0.38 as one moves from the first to the ninth decile of the productivity distribution. The effect is similar (from 0.31 to 0.37), when moving along the distribution of demand volatility.

and volatility are complementary in firms' location choices. Interestingly, however, Appendix Figure A.5 suggests that the effect of volatility and productivity is less symmetric in this selected sample: low-productivity firms hardly sort by demand volatility, while low-volatility firms still sort by productivity. The structure of our model allows for this possibility: When productivity is close to zero, so is expected profit, regardless of location. Accordingly, whether or not such an asymmetry is observed or not in the data depends on the intensity of firm selection by productivity, which our results suggest is lower in this selected sample.

1.5 — Conclusion

This paper shows that firms with a more volatile activity benefit from locating in denser locations. Higher operating costs associated with density create an incentive for volatile firms to adopt a more flexible workforce management strategy. In turn, by frequently releasing workers, those firms generate a positive externality on other firms, which benefit from easier hiring conditions. This finding opens a fruitful avenue for future research on the determinants of the spatial distribution of economic activity that go beyond static characteristics such as productivity. It provides a novel explanation for the non-negative correlation between city size and unemployment rates, and for the observation that many low-productivity firms are able to operate in large cities. However, our partial-equilibrium analysis leaves many effects aside. For example, workers should be compensated for working in more volatile firms. Conversely, a higher share of volatile firms might mitigate urban congestion costs, if firms are able to adjust their operating expenses. We leave these extensions for further research.

A.1 — Data appendix

A.1.1. Sample selection

Table A.1 — Sample selection

	Operating in January 2015	Relevant sectors	Single CZ	Non-missing productivity	Non-missing demand
Number of firms	890,088	770,901 (87%)	609,662 (68%)	365,041 (41%)	20,419 (2%)
Number of plants	1,100,329	934,894 (85%)	672,556 (61%)	408,308 (37%)	22,464 (2%)
Total employment	11,954,963	10,385,953 (87%)	5,474,535 (46%)	3,691,419 (31%)	483,934 (4%)

Notes: Number of firms refers to the concept of firm described in the main text, i.e. we aggregate plants of the same firm within each commuting zone. In column (1), we count firms operating in January 2015 with non-missing employment volatility. Column (2) keeps only the relevant sectors for our analysis: manufacturing, construction and service sectors (including non-tradable services). Column (3) drops firms operating in more than one CZ. In column (4) we keep only firms with non-missing productivity and age (i.e. firms for which we do not have information on the date of creation). Column (5) shows the number of firms for which we have information on demand volatility, a variable used in Section 1.4. Percentages correspond to the share with respect to the numbers in Column (1).

A.1.2. Total factor productivity

In our main results, productivity is calculated using the Levinshon-Petrin estimation technique,³¹ with the Akerberg et al. (2015) correction. We follow Combes et al. (2012) in defining productivity for each firm f and year y as

$$\ln(V_{fy}) = \beta_{0y} + \beta_1 \ln(k_{fy}) + \beta_2 \ln(l_{fy}) + \sum_{s=1}^3 \gamma_s l_{sfy} + \phi_{fy}, \quad (\text{A.1})$$

where V_{fy} is value added, k_{fy} is capital, l_{fy} is employment.³² As in Combes et al. (2012), we distinguish between three skill levels: high, intermediate and low, with l_{sfy} the share of skill

31. We use the Stata `prodest` command that exploits the control function approach.

32. Employment data is taken from INSEE-DADS and refers to mean employment, calculated over the months in each year y .

level s in the firm's overall employment. We estimate the equation separately for each 2-digit sector.³³ To minimize the impact of outliers, we then winsorize productivity at the 1st and 99th percentile within each of the 6 firm size classes described in Section 1.2.1.³⁴

The skill groups are defined following Combes et al. (2012). The low-skill group includes low-skill blue collars (in craft and manufacturing) and low-skill white collars (sales clerk, employees in personal services). The corresponding occupational codes in the French classification are: 55, "employés de commerce"; 67, "ouvriers non qualifiés de type industriel"; 68, "ouvriers non qualifiés de type artisanal". The intermediate-skill group includes high-skill blue collars (in craft, manufacturing, handling, and transport), and intermediate-skill white collars (administrative employees). In the French standard occupational classification, the following two-digit occupations are included: 52, "employés civils et agents de la fonction publique"; 53, "agents de surveillance"; 54, "employés administratifs d'entreprise"; 62, "ouvriers qualifiés de type industriel"; 63, "ouvriers qualifiés de type artisanal"; 64, "chauffeurs"; and 65, "ouvriers qualifiés de la manutention, du magasinage et du transport." Finally, the high-skill group includes managers (in craft, manufacturing or sales), executive and knowledge workers (executives, scientists, engineers), intermediate professions (intermediate professions in administration and sales firms, technicians, foremen). The group covers the following two-digit occupations are included: 21, "artisans (salariés de leur entreprise)"; 22, "commerçants et assimilés (salariés de leur entreprise)"; 23, "chefs d'entreprise de 10 salariés ou plus (salariés de leur entreprise)"; 31, "professions libérales (exercées sous statut de salarié)"; 34, "professeurs, professions scientifiques"; 35, "professions de l'information, des arts et des spectacles"; 37, "cadres administratifs et commerciaux d'entreprises"; 38, "ingénieurs et cadres techniques d'entreprises"; 46, "professions intermédiaires administratives et commerciales des entreprises"; 47, "techniciens"; and 48, "contremaîtres, agents de maîtrise." Finally, we drop the following non-coded occupations: 99, "non codage"; 00, "allocations assedic."

33. We focus on 12 2-digit sectors, which include manufacturing, construction, and services (including non-tradables).

34. Firm size is defined based on the mean employment calculated across months of each year. Productivity percentiles by firm size class are calculated over the 2010-2019 sample. After winsorizing, we further clean the data by only keeping productivity if firm revenues are above the 1st percentile and below the 99th percentile, calculated over the 2010-2019 sample.

A.1.3. Demand volatility

To compute our measure of demand volatility, we use the EAP survey to recover information on the structure of a firm's product portfolio in some base period³⁵

$$w_{fp,0} = \sum_p \frac{Sales_{fp,0}}{\sum_{p' \in P_{f,0}} Sales_{fp',0}},$$

where $Sales_{fp,0}$ is the value of product-level sales and $P_{f,0}$ denotes the set of products in the firm's portfolio in base period 0.³⁶ We then leverage upon trade data to construct a time series of the synthetic demand growth that a firm can expect to face, given the structure of its product portfolio

$$\gamma_{f,t}^D = \sum_{p' \in P_{f,0}} w_{fp',0} \gamma_{p',t}^D,$$

where $\gamma_{p',t}^D$ is the year-on-year growth of the world demand of product p' recovered from Eurostat trade data at monthly-frequency.³⁷ We can finally compute a measure of expected demand volatility

$$\varepsilon_{f,t} = \sqrt{\frac{1}{2\omega + 1} \sum_{\tau=-\omega}^{\omega} (\gamma_{f,t+\tau}^D - \bar{\gamma}_{f,t}^D)^2}. \quad (\text{A.2})$$

In comparison with the baseline measure in equation (1.1), the advantage of $\varepsilon_{f,t}$ is that it is a measure of volatility that is orthogonal to the firm's hiring strategy, or the structural churning rate in a particular location. From this point of view, it is more exogenous to the firm location choice than the volatility of labor demand. However, strict exogeneity requires that the structure of the firm's portfolio is given, at the time of the location decision. To give credibility to the assumption, we use information on the firm's portfolio of products observed

35. The downside is that the use of EAP forces us to focus on a sample of firms in the manufacturing sector, which is not representative of the whole population. See the last Column in Table A.1 for details.

36. Each product p is measured at the 4-digit level of the CPA 2008 product nomenclature, which can be merged to Eurostat data as seen below. Sales are constructed following EAP documentation as the sum of *Ventes de produits industriels* (VS2 + VF1 + VF2), *Ventes de services industriels* (VF3 or VT1), *Installation et pose de produits industriels* (IR1 + IR2 + IR3 = IT1), *Réparation et maintenance* (RR1 + RR2 + RR3 = RT1).

37. The monthly import time-series is smoothed using three-month moving averages. World demand refers to imports from all countries in the world excluding France.

during the first two years of activity. By contrast, world demand used to construct synthetic demand growth is based on the 12 months before and including the firm’s creation month.

A.2 — Empirical appendix

Table A.2 — Firm employment volatility: correlation with alternative measures

	Workers	Hours	Hours per empl.
Baseline	-	0.8492	0.3095
Non-residualized volatility	0.9991	0.8494	0.3111
Residualized by sector×month and CZ×month	0.9999	0.8490	0.3093
Month-on-month growth rates	0.7029	0.6372	0.4392
Growth of permanent contracts	0.7485	0.6513	0.2835

Notes: Correlation coefficients based on various measures of volatility. All measures are computed from the January 2015 cross-section of firms. The baseline measure is described in section 1.2.1. The “non-residualized volatility” measure is computed as in equation (1.1) using the growth rate observed in the data instead of focusing on the idiosyncratic component of growth. The measured which is “residualized by sector×month and CZ×month” the residual of an equation that includes CZ FEs. The “month-on-month growth rates” measure is computed exactly as the baseline except that the raw data are month-on-month (instead of year-on-year) growth rates. Finally, the “growth of permanent contracts” measure is constructed as the baseline on the restricted set of the firm’s permanent contracts. All correlations are statistically significant at the 1% level.

Table A.3 — Firm employment volatility: summary statistics

	Workers		Hours		Hours per empl.	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Baseline	0.210	0.189	0.233	0.219	0.093	0.102
Non-residualized volatility	0.209	0.191	0.231	0.220	0.092	0.103
Residualized by sect×month and CZ×month	0.211	0.189	0.233	0.218	0.093	0.102
Month-on-month growth rates	0.080	0.063	0.085	0.065	0.025	0.019
Growth of permanent contracts	0.199	0.191	0.222	0.215	0.091	0.106

Notes: The statistics are recovered from the cross-section of firms active in January 2015. The baseline measure is described in section 1.2.1. The “non-residualized volatility” measure is computed as in equation (1.1) using the growth rate observed in the data instead of focusing on the idiosyncratic component of growth. The measured which is “residualized by sector×month and CZ×month” the residual of an equation that includes CZ FEs. The “month-on-month growth rates” measure is computed exactly as the baseline except that the raw data are month-on-month (instead of year-on-year) growth rates. Finally, the “growth of permanent contracts” measure is constructed as the baseline on the restricted set of the firm’s permanent contracts.




Figure A.1 — Firm employment volatility: interpretation

Notes: The figure is based on the January 2015 cross-section of single-plant firms. The variable of interest is the expected change in employment computed as the product of the firm's average size and volatility. Results are winsorized at the top 95th percentile.




Figure A.2 — Firm employment volatility: the effect of age

Notes: The figure is based on the panel of firms created in 2010 with non-missing employment volatility (29,168 firms, for a total number of observations of $N = 126,704$). The figure plots age FEs from a regression of log employment volatility on age and firm FEs. The reference age category is two years, because volatility is computed over a 35-month window around the date of observation. Whiskers indicate 95% confidence intervals

Table A.4 — Results of the location choice model: robustness

	Dependent variable: CZ choice						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
CZ Density M	0.339 (0.027)	0.339 (0.027)	0.338 (0.027)	0.338 (0.027)		0.337 (0.027)	
- × Productivity		0.049 (0.019)		0.048 (0.018)	0.050 (0.018)	0.053 (0.015)	0.055 (0.015)
- × Volatility			0.034 (0.017)	0.033 (0.016)	0.030 (0.016)	0.027 (0.017)	0.024 (0.017)
- × Volatility × Productivity						0.030 (0.013)	0.029 (0.013)
CZ characteristics	✓	✓	✓	✓		✓	
Localization economies	✓	✓	✓	✓	✓	✓	✓
CZ FEs					✓		✓
Pseudo R^2	0.0441	0.0446	0.0443	0.0448	0.0451	0.0871	0.0873
Observations	471K	471K	471K	471K	471K	471K	471K

Notes: Coefficient estimates from a conditional logit model with firm FEs. The sample is based on all firm entries from January 2010 to December 2019 (1,682 entries, resulting in $N = 470,960$ observations). M is the log of CZ density, volatility is the standardized value of expected demand volatility, and productivity is the standardized value of productivity. Standard errors in round parentheses.




Figure A.3—Unemployment rate and density

Notes: The figure shows the unemployment rate of the working-age population (aged 15-54) by CZ as a function of its (working-age) population density. In red, we provide the correlation after controlling for the share of university graduates in the population above 15, the share of managers among employed workers, the shares of old and young workers in the working-age population, and 22 region FEs. In the unconditional case, the slope is equal to 0.0064 and the R-squared is equal to 0.0542. In the conditional case, the slope is equal to 0.0060 and the R-squared is equal to 0.0964. Both slopes are statistically significant at the 1% level. Source: 2018 Census




Figure A.4—Heterogeneity in location choices, along the distributions of productivity and volatility: robustness

Notes: The figure shows the mean probability of locating in each CZ, for high-productivity (respectively, high-volatility) firms in relative terms with low-productivity (respectively low-volatility) firms. The cut-offs are based on firms at the 25th and 75th percentile of each distribution. The probabilities are recovered from the estimation of the model in column (4) of Appendix Table A.4.

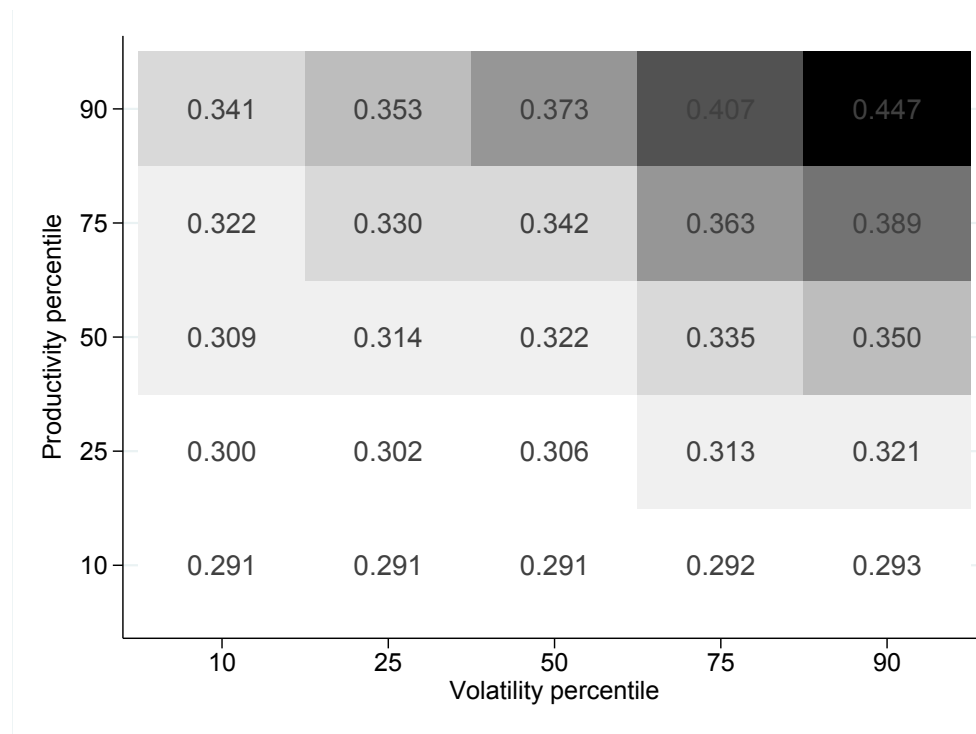


Figure A.5 — Elasticity of the odds of choosing a CZ to CZ density: robustness

Notes: Each cell is computed (in %) as $\exp\{0.01[\hat{\beta}_M + \hat{\beta}_{M\varepsilon}P_\varepsilon + \hat{\beta}_{M\phi}P_\phi + \hat{\beta}_{M\varepsilon\phi}P_\varepsilon P_\phi]\} - 1$. Estimates from model (6) of Appendix Table A.4.

A.3 — Model appendix

A.3.1. Definitions

Expected profits — Expected profits upon entry are given by

$$\mathbb{E}_B(\phi, \varepsilon, M) = \frac{\mu(M)[\phi - R(M)] - c(\delta + r)}{r[r + \delta + \mu(M)]} \quad (\text{A.3})$$

$$\mathbb{E}_W(\phi, \varepsilon, M) = \frac{1}{2} \times \frac{(r + \delta + 2\xi)\{\mu(M)[\phi - R(M)] - c(r + \delta)\} + \mu(M)(r + \delta)\varepsilon\phi}{r[(r + \delta)(r + \delta + 2\xi) + \mu(M)(r + \delta + \xi)]} \quad (\text{A.4})$$

$$\mathbb{E}_C(\phi, \varepsilon, M) = \frac{1}{2} \times \frac{\mu(M)[\phi(1 + \varepsilon) - R(M)] - c(r + \delta + \xi)}{r[r + \delta + \xi + \mu(M)]}. \quad (\text{A.5})$$

These three expressions are increasing in ϕ . In addition, we can show that

$$\begin{aligned} \frac{\partial \mathbb{E}_B(\phi, \varepsilon, M)}{\partial \phi} - \frac{\partial \mathbb{E}_W(\phi, \varepsilon, M)}{\partial \phi} &= \frac{(r + \delta)\mu(M)[(1 - \varepsilon)(r + \delta + \mu(M)) + 2\xi]}{2r(r + \delta + \mu(M))[(r + \delta)(r + \delta + 2\xi) + (r + \delta + \xi)\mu(M)]} > 0 \\ \frac{\partial \mathbb{E}_W(\phi, \varepsilon, M)}{\partial \phi} - \frac{\partial \mathbb{E}_C(\phi, \varepsilon, M)}{\partial \phi} &= \frac{\xi\mu(M)[(1 - \varepsilon)(r + \delta + \mu(M)) + 2\xi]}{2r(r + \delta + \xi + \mu(M))[(r + \delta)(r + \delta + 2\xi) + (r + \delta + \xi)\mu(M)]} > 0. \end{aligned}$$

Selection cutoffs — Solving for $\mathbb{E}_s(\phi, \varepsilon, M) = 0$ we find

$$\phi_B(\varepsilon, M) = \phi_B(M) = R(M) + \frac{c(r + \delta)}{\mu(M)} \quad (\text{A.6})$$

$$\phi_W(\varepsilon, M) = \left(\frac{r + \delta + 2\xi}{(1 + \varepsilon)(r + \delta) + 2\xi} \right) \phi_B(M) \quad (\text{A.7})$$

$$\phi_C(\varepsilon, M) = \frac{1}{1 + \varepsilon} \left(\phi_B(M) + \frac{c\xi}{\mu(M)} \right). \quad (\text{A.8})$$

Note that $\phi_B(M) \geq \phi_W(\varepsilon, M)$ and

$$\phi_C(\varepsilon, M) \leq \phi_W(\varepsilon, M) \iff \varepsilon \geq \tilde{\varepsilon}(M) = \frac{c(r + \delta + 2\xi)}{c(r + \delta) + 2\mu(M)R(M)}.$$

Switching cutoffs — Solving for $\mathbb{E}_B(\phi, \varepsilon, M) = \mathbb{E}_W(\phi, \varepsilon, M)$ and $\mathbb{E}_W(\phi, \varepsilon, M) = \mathbb{E}_C(\phi, \varepsilon, M)$, we find

$$\phi_{BW}(\varepsilon, M) = \left(\frac{r + \delta + \mu(M) + 2\xi}{(1 - \varepsilon)(r + \delta + \mu(M)) + 2\xi} \right) \phi_B(M) \quad (\text{A.9})$$

$$\phi_{WC}(\varepsilon, M) = \frac{R(M)(r + \delta + 2\xi + \mu(M))}{(1 - \varepsilon)(r + \delta + \mu(M)) + 2\xi}. \quad (\text{A.10})$$

Note that $\phi_{BW}(M) \geq \phi_B(M)$. In addition, we can show that

$$\phi_{BW}(\varepsilon, M) - \phi_{WC}(\varepsilon, M) = c \times \frac{(r + \delta)(r + \delta + 2\xi) + (r + \delta + \xi)\mu(M)}{\mu(M)[(1 - \varepsilon)(r + \delta + \mu(M)) + 2\xi]} > 0,$$

so that strategy W is never adopted when $c = 0$. Conversely, since $\partial\phi_C(\varepsilon, M)/\partial c > 0$ and $\partial\phi_{WC}(\varepsilon, M)/\partial c = 0$, there is maximum value of c above which strategy C is never adopted.

Sorting cutoffs — Solving for $\partial\mathbb{E}_s(\phi, \varepsilon, M)/\partial\phi = 0$ and $\partial\mathbb{E}_s(\phi, \varepsilon, M)/\partial\varepsilon = 0$, we find

$$\phi_B^*(\varepsilon, M) = \phi_B^*(M) = R(M) - c + \frac{R'(M)}{\mu'(M)} \left(\frac{\mu(M)(r + \delta + \mu(M))}{(r + \delta)} \right)$$

$$\phi_W^*(\varepsilon, M) = \frac{(r + \delta + 2\xi)R(M) - c(r + \delta + \xi)}{(1 + \varepsilon)(r + \delta) + 2\xi} + \frac{R'(M)}{\mu'(M)} \left(\frac{\mu(M)[(r + \delta)(r + \delta + 2\xi) + (r + \delta + \xi)\mu(M)]}{(r + \delta)[(1 + \varepsilon)(r + \delta) + 2\xi]} \right)$$

$$\phi_C^*(\varepsilon, M) = \frac{1}{1 + \varepsilon} \left[R(M) - c + \frac{R'(M)}{\mu'(M)} \left(\frac{\mu(M)(r + \delta + \xi + \mu(M))}{(r + \delta)} \right) \right]$$

$$\varepsilon_W^*(\phi, M) = \frac{1}{\phi} \left[-(c + \phi - R(M)) - \xi \left(\frac{c + 2(\phi - R(M))}{r + \delta} \right) + \frac{R'(M)}{\mu'(M)} \left(\frac{\mu(M)[(r + \delta)(r + \delta + 2\xi) + (r + \delta + \xi)\mu(M)]}{(r + \delta)^2} \right) \right]$$

$$\varepsilon_C^*(\phi, M) = \frac{1}{\phi} \left[-(c + \phi - R(M)) + \frac{R'(M)}{\mu'(M)} \left(\frac{(r + \delta + \xi + \mu(M))\mu(M)}{r + \delta + \xi} \right) \right].$$

Motion laws — At steady state, the measures of firms in each state are constant, for each strategy. We use bold symbols to distinguish these measures from their corresponding values

$$\text{Strategy B: } \begin{cases} \mu(M)\mathbf{V}_B(M) = \delta\mathbf{A}_B(M) \\ \mathbf{V}_B(M) + \mathbf{A}_B(M) = \mathbf{B}(M) \end{cases} \quad (\text{A.11})$$

$$\text{Strategy W: } \begin{cases} [\xi + \mu(M)]\mathbf{V}_W(M) = \delta\mathbf{A}_W^h(M) + \xi\mathbf{I}_W(M) \\ \xi\mathbf{I}_W(M) = \xi\mathbf{V}_W(M) + \delta\mathbf{A}_W^l(M) \\ (\delta + \xi)\mathbf{A}_W^h(M) = \mu(M)\mathbf{V}_W(M) + \xi\mathbf{A}_W^l(M) \\ (\delta + \xi)\mathbf{A}_W^l(M) = \xi\mathbf{A}_W^h(M) \\ \mathbf{V}_W(M) + \mathbf{I}_W(M) + \mathbf{A}_W^h(M) + \mathbf{A}_W^l(M) = \mathbf{W}(M) \end{cases} \quad (\text{A.12})$$

$$\text{Strategy C: } \begin{cases} [\mu(M) + \xi]\mathbf{V}_C(M) = \delta\mathbf{A}_C(M) + \xi\mathbf{I}_C(M) \\ \xi\mathbf{I}_C(M) = \xi[\mathbf{A}_C(M) + \mathbf{V}_C(M)] \\ (\delta + \xi)\mathbf{A}_C(M) = \mu(M)\mathbf{V}_C(M) \\ \mathbf{V}_C(M) + \mathbf{I}_C(M) + \mathbf{A}_C(M) = \mathbf{C}(M), \end{cases} \quad (\text{A.13})$$

where $\mathbf{B}(M)$, $\mathbf{W}(M)$ and $\mathbf{C}(M)$ are the respective measures of firms that follow strategies B , W and C .

A.3.2. Proofs

Proof of Proposition 1 — Employment is a random variable that follows a Bernoulli distribution. The variance of employment is thus

$$\sigma^l(\phi, \varepsilon, M^*(\phi, \varepsilon), s^*(\phi, \varepsilon)) = \mathbb{P}(\phi, \varepsilon, M^*, s^*) [1 - \mathbb{P}(\phi, \varepsilon, M^*, s^*)],$$

where $\mathbb{P}(\phi, \varepsilon, M^*, s^*)$ is the probability of a position being filled, which follows the steady-state constraints given by equations (A.11)–(A.13) with $\mathbf{B}(M) = \mathbf{W}(M) = \mathbf{C}(M) = 1$. Employment volatility is higher under strategy C than under strategy B if $\xi < (\mu(M) -$

$2\delta)(\delta + \mu(M))/(2\delta)$, which is only possible if $\mu(M) > 2\delta$, and is more likely to be satisfied if M is large, given $\mu'(M) > 0$.

Proof of Proposition 2 — Result 2.1 stems from the facts that $\partial\phi_{WC}(\varepsilon, M)/\partial\varepsilon > 0$ and $\partial\phi_C(\varepsilon, M)/\partial\varepsilon < 0$ and that the selection cutoff for strategy C is the lowest, as long as $\varepsilon \geq \tilde{\varepsilon}(M)$. Result 2.2 stems from the fact that $\partial^2\phi_{WC}(\varepsilon, M)/\partial\varepsilon\partial\varepsilon > 0$.

Proof of Proposition 3 — Result 3.1 is obtained by noticing that, under Assumptions 1 and 2, $\partial\phi_{BW}(\varepsilon, M)/\partial M > 0$. Thus, the area of region B decreases with density. Moreover, under Assumption 2, $\partial\phi_W(\varepsilon, M)/\partial M > 0$. In addition, under no assumption, $\partial[\phi_{BW}(\varepsilon, M) - \phi_{WC}(\varepsilon, M)]/\partial M < 0$. Thus, the area of region W decreases with density. Finally, under no assumption, $\partial\tilde{\varepsilon}(M)/\partial M < 0$. In addition, under Assumption 2, $\partial[\phi_{WC}(\varepsilon, M) - \phi_C(\varepsilon, M)]/\partial M > 0$. Thus, the area of region C increases with density. Result 3.2 can be derived by observing that the productivity-volatility substitution for selection is represented by $\phi_W(\varepsilon, M)$ for $\varepsilon \in [0, \tilde{\varepsilon}(M)]$ and $\phi_C(\varepsilon, M)$ for $\varepsilon \in [\tilde{\varepsilon}(M), 1]$. Then, under Assumption 2, $\partial^2\phi_W(\varepsilon, M)/\partial\varepsilon\partial M < 0$ and $\partial^2\phi_C(\varepsilon, M)/\partial\varepsilon\partial M < 0$. Thus, in denser cities, volatility and productivity are better substitutes for lowering the selection of firms.

Proof of Proposition 4 — Result 4.1 stems from the fact that under Assumption 3, we can show that $\forall s \in \{B, W, C\}, \partial\phi_s^*(\varepsilon, M)/\partial M > 0$ and $\forall s \in \{W, C\}, \partial\varepsilon_s^*(\phi, M)/\partial M > 0$. Therefore, more productive and more volatile firms sort into denser cities. Results 4.2 stems from the fact that $\forall s \in \{W, C\}, \partial^2\phi_s^*(\varepsilon, M)/\partial M\partial\varepsilon < 0$ and $\partial^2\varepsilon_s^*(\phi, M)/\partial M\partial p < 0$. Therefore, more productive (resp., volatile) firms sort into denser cities if they are more volatile (resp., productive). This second result ensures that the share of churning firms increases with density, even though average productivity also increases with density. Then, again under Assumption 3, we can also show that $\forall \varepsilon \in [0, 1], \partial\phi_B^*(M)/\partial M > \partial\phi_W^*(\varepsilon, M)/\partial M > \partial\phi_C^*(\varepsilon, M)/\partial M$. The relationship between productivity and density is stronger when firms choose strategy B , followed by strategy W and then C . Therefore, the productivity-density gradient decreases with volatility, which proves result 4.3.

A.3.3. Additional figures

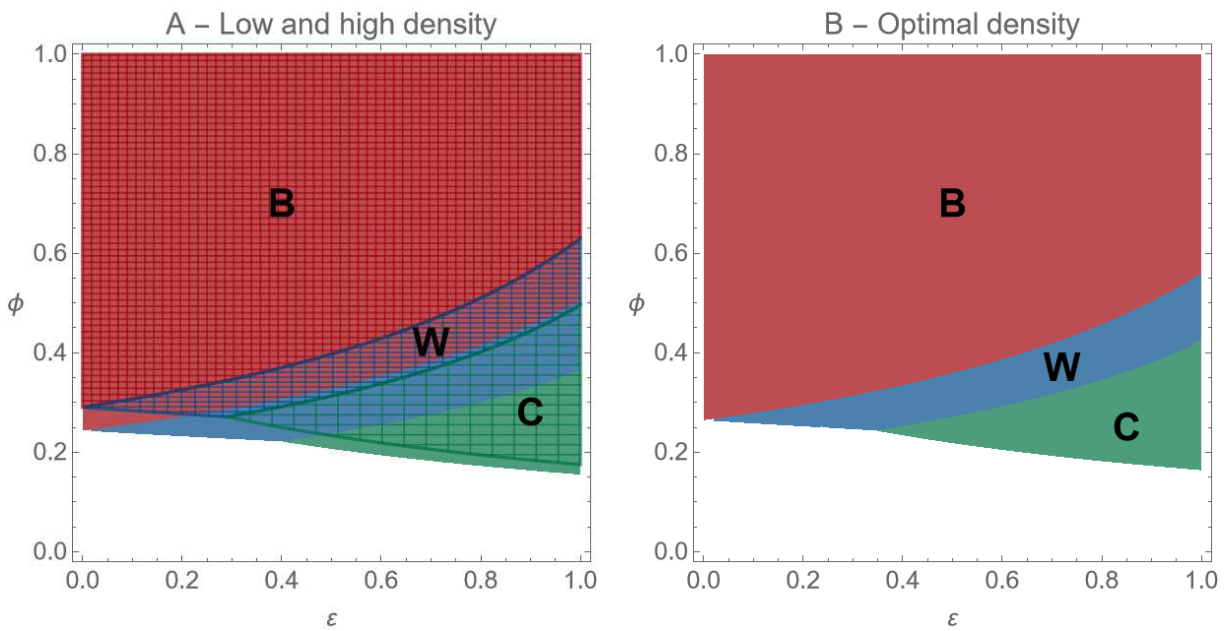


Figure A.6 — Strategy choice and city choice

Notes: Calibration: $\xi = 0.2$, $r = 0.03$, $\delta = 0.1$, $c = 0.1$, $\mu(M) = 0.3M^{0.05}$, $R(M) = 0.2M^{0.1}$, $\bar{p} = 1$. Panel A: The figure represents the set of (ε, ϕ) combinations associated with each adopted strategy, for a low density ($M = 1$, plain colors), and for a high density ($M = 10$, mesh lines). The letters locate adopted strategies when $M = 10$. Panel B: The figure represents the set of (ε, ϕ) combinations associated with each adopted strategy, for the optimal level of density, found by numerical search.

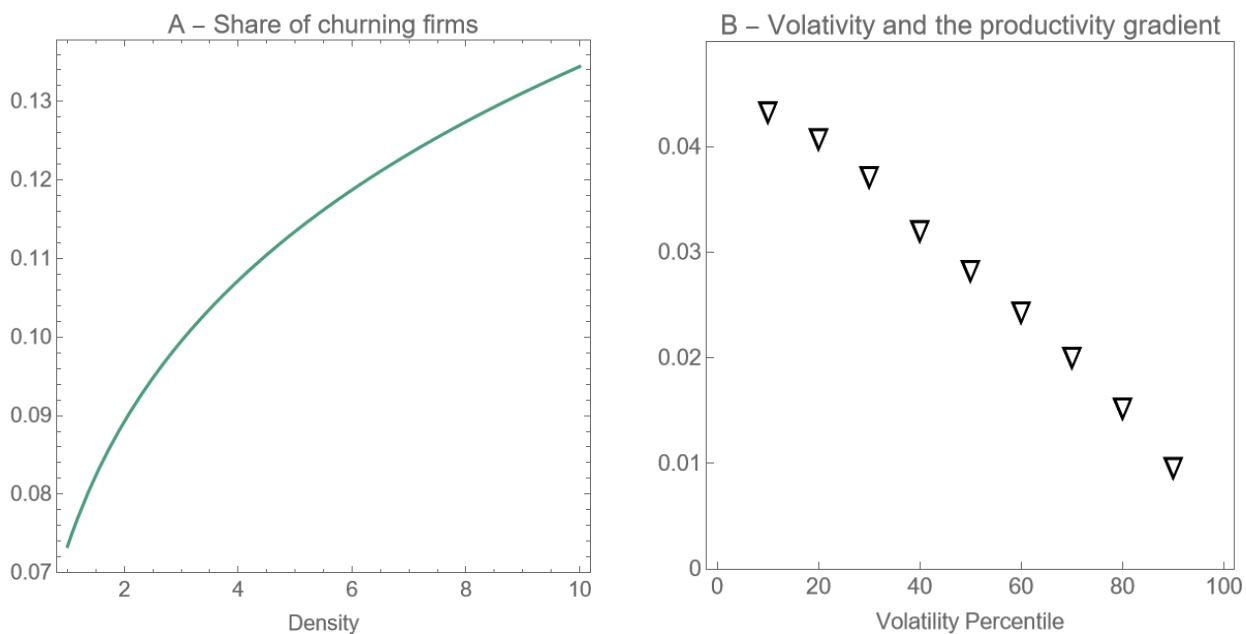


Figure A.7 — Volatility and the sorting of firms

Calibration: $c = 0.1$, $\xi = 0.2$, $r = 0.03$, $\delta = 0.1$, $\mu(M) = 0.3M^{0.05}$, $R(M) = 0.2M^{0.1}$, $\bar{p} = 1$. We assume that ϕ and ε are independent and uniformly distributed over $[0, 1]$. The optimum is found by a numerical search. Panel A: The share of churning firms is given by $\mathcal{C}(M) = \frac{\int \int \mathbf{1}_{C=s^*(\phi, \varepsilon, M)} h(\phi, \varepsilon) d\phi d\varepsilon}{\sum_s \int \int \mathbf{1}_{s=s^*(\phi, \varepsilon, M)} h(\phi, \varepsilon) d\phi d\varepsilon}$. Panel B: For each decile d of the distribution of ε , let $d_\varepsilon = [d - 5\%, d + 5\%]$. The average productivity of firms in this decile of volatility is given by $\phi^*(M, d) = \frac{\sum_s \int \int_{d_\varepsilon} \mathbf{1}_{s=s^*(\phi, \varepsilon, M)} \phi_s^*(\varepsilon, M) h(\phi, \varepsilon) d\phi d\varepsilon}{\sum_s \int \int_{d_\varepsilon} \mathbf{1}_{s=s^*(\phi, \varepsilon, M)} h(\phi, \varepsilon) d\phi d\varepsilon}$. To approximate the productivity-density gradient, we compute the slope of this function at the median density, such that $\nabla(d) = \phi^*(6, d) - \phi^*(5, d)$.

A.3.4. Endogenous matching rate

In order to recover $\mathbf{B}(M)$, $\mathbf{W}(M)$ and $\mathbf{C}(M)$ in equation (A.11)-(A.13), we introduce an entry process: in each city, we assume that there is a continuum of latent firms with known distribution $h(\phi, \varepsilon)$. Those firms pay a cost f_E to draw (ϕ, ε) . Free entry means that

$$\forall M, f_E = \int \int \max \{0, \max_s \{E_s(\phi, \varepsilon, M)\}\} h(\phi, \varepsilon) d\phi d\varepsilon. \quad (\text{A.14})$$

Market tightness is given by the ratio of the number of vacancies to that of unemployed workers

$$\theta(M) = \frac{\sum_s \mathbf{V}_s(M)}{M - \sum_s \mathbf{A}_s(M)}. \quad (\text{A.15})$$

If we denote $\alpha(M) = M/F(M)$ with $F(M) = \mathbf{B}(M) + \mathbf{W}(M) + \mathbf{C}(M)$ the density of firms and use a parametric assumption $\mathcal{M}(\cdot)$ on the matching technology, we can solve numerically for the fixed point given by equation (A.14)-(A.15) and recover the values of $\alpha(M)$ and $\mu(M)$.

We illustrate this method with the same calibration as in the main text.³⁸ The results are shown in Figure A.8, for two different scenarios: either the full model, where firms may follow any of the three strategies, or a restricted model where firms may only follow the B strategy. Under this parametrization, the impact of density on the worker finding rate is higher under the restricted model: the indirect negative effect of churning through firm entry trumps the direct positive effect. Conversely, even if the number of firms does not adjust, firms that are allowed to churn still benefit from better matching conditions in larger cities. To substantiate this point, we simulate two counterfactual situations where α is set to its value for the lowest density ($M = 1$): matching economies, as measured by $d\mu(M)/dM$, drop by half when firms are allowed to churn, and drop to zero when they are not.

38. We use a Cobb-Douglas matching function $\mathcal{M}(V, U) = \sqrt{VU}$ and we set $f_E = 7.5$ so that the resulting unemployment rate lies between 5% and 10%.




Figure A.8 — Endogenous worker finding rate

Calibration: $c = 0.1$, $\xi = 0.2$, $r = 0.03$, $\delta = 0.1$, $R(M) = 0.2M^{0.1}$, $\bar{p} = 1$. We assume that ϕ and ε are independent and uniformly distributed over $[0, 1]$. The value of $\mu(M)$ is obtained as the numerical solution to the fixed point problem described in the text, with a Cobb-Douglas matching technology $\mathcal{M}(V, U) = \sqrt{VU}$ and $f_E = 7.5$. All: $\mu(M)$ is obtained by solving the full model; B: $\mu(M)$ is obtained by solving the restricted model. Continuous lines mean that firm density adjusts to meet the free entry condition. Dashed lines mean that firm density is set to its equilibrium value when $M = 1$.

A.3.5. Model extension

In this Section, we provide an extension to our base model in the spirit of Melitz (2003), introducing monopolistically competitive firms facing a CES demand system that draw heterogeneous productivity and demand volatility upon entry. This allows us to (i) explicitly model demand shocks, (ii) introduce firm size, (iii) clarify the link between productivity thresholds and the share of volatile firms active in a given market, and (iv) exhibit that the same forces as used in our parsimonious model are needed to generate the main results of the paper. We use the model to perform comparative statics. All proofs and derivations are available upon request.

Assumptions — We analyze a differentiated goods sector in one city of size M . There are M workers-consumers inelastically providing one unit of labor that spend a fraction γ of their income I on the sector's differentiated goods over which they have CES preferences.

There is a unit mass of potential entrants in each city. Firms pay a fixed entry cost $f_E > 0$ after which they learn their productivity ϕ and demand volatility. Productivity is drawn from a Pareto distribution of shape ν and scale $\phi_{min} > 0$. A proportion χ of firms face constant demand, while $(1 - \chi)$ are exposed to volatile demand flows. More specifically, the firm observes the demand for its variety at each period

$$q(p, \varepsilon) = \varepsilon^\eta \bar{Q} p^{-\eta},$$

where \bar{Q} is aggregate real consumption and ε is the firm's idiosyncratic demand.³⁹ The corresponding indirect demand is $p(q, \varepsilon) = \varepsilon (\bar{Q})^{\frac{1}{\eta}} (q)^{-\frac{1}{\eta}}$ which allows us to link this extension with the base model in the main paper. For the firms with no demand volatility, we normalize $\varepsilon = 1$. Volatile firms alternate between ε_l and ε_h (respectively low and high demand) at a rate ξ . We consider the case where $(\varepsilon_l + \varepsilon_h)/2 = 1$, i.e. volatile and non-volatile firms only differ by the second-moment of their indirect demand process, conditional on their productivity. In our exposition, we further assume that $\varepsilon_l = 0$ and $\varepsilon_h = 2$. These are extreme values that, however, permit analytical characterization of the problem without resorting to simulations.

In the rest of the model, we assume monopolistic competition, i.e. firms are input price takers and view any aggregate parameters (e.g. labor market tightness, or real consumption) as being exogenous. Once productivity and volatility are revealed, the firm decides whether to pay a set-up cost $f_p > 0$ to produce, then chooses a hiring strategy.

In comparison with the stylized model in Section 1.3, we concentrate on the two extreme strategies, namely Business-as-usual and Churning, which corresponds to the case where c is small. While the baseline model considers one-job firms, we now consider the firm's decision on a measure of job openings, given an exogenous filling rate $\mu(M)$ and separation rate δ for each of these positions. If a worker fills a position, she produces ϕ units of the differentiated good. The firm assesses the present discounted value (PDV) of each created position given operational and hiring costs, as well as the job filling and separation rates. When the firm enters, it doesn't know the state of its demand ε , but it immediately chooses the type of each position it creates. A position is permanent if the firm will not fire the employed worker

39. We assume that $\nu > \eta - 1$ to ensure finite aggregate productivity levels.

under any circumstances and only an exogenous separation can destroy it. A position will be a churning one if the firm may also fire workers when demand is low.

Business as usual strategy — Firms using this strategy keep employment constant and hence also production \bar{q} and the associated cost. It is straightforward to show that non-volatile firms always choose this strategy. Volatile firms instead choose between Business-as-usual and Churning. Under Business-as-usual, the firm produces at full capacity and adjusts its price to demand shocks.

The employment level \bar{L}_B chosen by a Business-as-usual firm upon entry, assuming it does not know a priori the state of the demand, maximizes expected profits

$$\Pi_B(\phi, \varepsilon, M) = \left[\frac{\mu(M) \left\{ \left[\frac{p_t + p_h}{2} \right] \phi - R(M) \right\} - c(\delta + r)}{r[r + \delta + \mu(M)]} \right] \bar{L}_B \quad (\text{A.16})$$

subject to $p_t = \varepsilon_t (\bar{Q})^{\frac{1}{\eta}} (\bar{q})^{-\frac{1}{\eta}}$, where t corresponds to the low or high state. $\bar{q} = \phi \kappa_B \bar{L}_B$ is the expected (constant) output given the adopted strategy, where $\kappa_B = \mu(M)/[\delta + \mu(M)]$ is the fraction of time a position is filled. The term under brackets is simply the present discounted value of a permanent position and is as in equation (A.3) in the baseline model (using the fact that the demand shocks average to one across states), except that now we explicitly allow the firm to optimize on its price p .

Letting $\Gamma_1 = \mu(M)/\{r[r + \delta + \mu(M)]\}$ and $\Gamma_2 = \{\mu(M)R(M) + c(\delta + r)\}/\{r[r + \delta + \mu(M)]\}$, we can rewrite the problem of the firm as

$$\max_{\bar{L}_B} \left[\Gamma_1 \left(\frac{\phi \kappa_B \bar{L}_B}{\bar{Q}} \right)^{-1/\eta} \phi - \Gamma_2 \right] \bar{L}_B.$$

The optimal solution gives

$$\begin{aligned} \bar{L}_B &= \left(\frac{\eta}{\eta - 1} \frac{\Gamma_2}{\Gamma_1} \right)^{-\eta} \left(\frac{\bar{Q}}{\kappa_B} \right) \phi^{\eta-1} \\ p_t &= \varepsilon_t \frac{\eta}{\eta - 1} \frac{\Gamma_2}{\phi \Gamma_1}. \end{aligned}$$

Plugging the equilibrium strategies into equation (A.16) implies

$$\Pi_B(\phi, \varepsilon, M) = \frac{\Gamma_2}{\eta - 1} \left(\frac{\eta}{\eta - 1} \frac{\Gamma_2}{\Gamma_1} \right)^{-\eta} \frac{\bar{Q}}{\kappa_B} \phi^{\eta-1},$$

which does not depend on the firm's volatility.

Churning strategy — Consider now a firm that lays off some workers in the low state of demand. Given the firm faces zero demand in the low state ($\varepsilon_l = 0$), it can be shown that if firing some workers is preferred in the low demand state, the firm would want to fire its entire workforce.⁴⁰ Following the same steps as before, the problem of the firm consists in choosing \bar{L}_C to maximize expected discounted profits

$$\Pi_C(\phi, \varepsilon, M) = \frac{1}{2} \left[\frac{\mu(M) \{p_h \phi - R(M)\} - c(r + \delta + \xi)}{r[r + \delta + \xi + \mu(M)]} \right] \bar{L}_C \quad (\text{A.17})$$

$$= \frac{1}{2} \left[\Phi_1 \varepsilon_h \left(\frac{\phi \kappa_C \bar{L}_C}{\bar{Q}} \right)^{-\frac{1}{\eta}} \phi - \Phi_2 \right] \bar{L}_C, \quad (\text{A.18})$$

where $\Phi_1 = \mu(M)/\{r[r + \delta + \xi + \mu(M)]\}$, $\Phi_2 = \{\mu(M)R(M) + c(r + \delta + \xi)\}/\{r[r + \delta + \xi + \mu(M)]\}$ and $\kappa_C = \mu(M)/\{(\delta + \xi + \mu(M))\}$ is the fraction of time a position is filled when the firm is in the high state. In the low state, the firm is idle and receives zero profits. Again, the term under brackets corresponds to the present discounted value of a position, as in equation (A.5). We thus obtain that

$$p_h = \frac{\eta}{\eta - 1} \frac{\Phi_2}{\phi \Phi_1}$$

$$\bar{L}_C = \left(\frac{1}{\varepsilon_h} \frac{\eta}{\eta - 1} \frac{\Phi_2}{\Phi_1} \right)^{-\eta} \frac{\bar{Q}}{\kappa_C} \phi^{\eta-1}$$

40. It should be noted here that handling non-zero demand in the low state is straightforward conceptually. In such a scenario, the firm churns over \bar{L}_C and hires \bar{L}_C^P permanent workers to ensure sufficient production to serve demand in the low state. However, given the maximization problem of the firm is then not separable in \bar{L}_C and \bar{L}_C^P , it is impossible to get a full characterization for the optimal measures of positions, prices nor the value function. This is why we set $\varepsilon_l = 0$ to obtain an explicit solution to the problem.

and the expected profit

$$\Pi_C(\phi, \varepsilon, M) = \frac{1}{2} \varepsilon_h^\eta \frac{\Phi_2}{\eta - 1} \left(\frac{\eta}{\eta - 1} \frac{\Phi_2}{\Phi_1} \right)^{-\eta} \frac{\bar{Q}}{\kappa_C} \phi^{\eta-1}. \quad (\text{A.19})$$

In this model, the impact of demand shocks on firms' outcomes depends on their strategy. Under Business-as-usual, the firm's employment is constant and prices are adjusted to shocks. Instead, the churning strategy allows firms to adjust to demand shocks through quantities, and thus prices are independent of demand shocks.⁴¹

The firm compares the expected gains from pursuing the Business-as-usual strategy and Churning. The latter is preferred if the following holds

$$\Gamma_2 \left(\frac{\Gamma_1}{\Gamma_2} \right)^\eta \kappa_C < \frac{1}{2} \varepsilon_h^\eta \Phi_2 \left(\frac{\Phi_1}{\Phi_2} \right)^\eta \kappa_B. \quad (\text{A.20})$$

The condition depends on structural parameters and location through $R(M)$ and $\mu(M)$ but not on firm productivity ϕ . A sufficient condition is that $R(M) > c$ and $\mu(M) > \xi$, exactly the conditions needed for churning in the base model (see Assumption 1).

Moreover, as $\xi \rightarrow 0$ firms always prefer to churn. It is only if demand levels change frequently, that firms may prefer to keep workers instead of constantly readjusting the labor force and saving on operating and hiring costs in low states of demand at the expense of waiting to hire when its demand turns high. In the calibration used in Section 1.3, condition (A.20) is met. In what follows, we will thus discern between two types of firms: non-volatile ones (that follow the Business-as-usual strategy) and volatile that churn.

Equilibrium distribution of volatile and non-volatile firms — After firms learn their productivity and demand volatility, they decide whether to pay a fixed cost f_p to produce. This defines the cutoff productivity values for firms that break even, conditional on their

41. Note that this is true in the extreme case in which production is zero in the low state. But if churning firms had to combine permanent and churning positions, their price in the high state would be more complicated to determine as the high state would combine permanent and adjustable positions.

volatility denoted $\phi^*(0)$ and $\phi^*(\varepsilon)$ respectively

$$\Pi_B(\phi^*(0), 0, M) = \Pi_C(\phi^*(\varepsilon), \varepsilon, M) = f_p. \quad (\text{A.21})$$

This implies a linear relationship between the productivity cutoffs for volatile and non-volatile firms $\phi^*(\varepsilon) = \Omega\phi^*(0)$ where $\Omega^{\eta-1} \equiv 2\varepsilon_h^{-\eta} \frac{\Gamma_2}{\Phi_2} \left(\frac{\Gamma_2}{\Phi_2} \frac{\Phi_1}{\Gamma_1} \right)^{-\eta} \frac{\kappa_C}{\kappa_B}$. If equation (A.20) holds, $\Omega < 1$ and $\phi^*(\varepsilon) < \phi^*(0)$ i.e. the productivity cut-off is lower for volatile than non-volatile firms, as in the base model of Section 1.3.

For a firm considering paying entry cost f_E , the expected profit depends on

$$\left[\begin{aligned} & \chi \left[(1 - G(\phi^*(0))) \int_{\phi^*(0)}^{\infty} (\Pi_B(\phi, 0, M) - f_p) \frac{g(\phi)}{(1 - G(\phi^*(0)))} d\phi \right] \\ & + (1 - \chi) \left[(1 - G(\phi^*(\varepsilon))) \int_{\phi^*(\varepsilon)}^{\infty} (\Pi_C(\phi, \varepsilon, M) - f_p) \frac{g(\phi)}{(1 - G(\phi^*(\varepsilon)))} d\phi \right] \end{aligned} \right] = f_E.$$

Letting $\tilde{\phi}(\cdot) := \left(\frac{1}{1 - G(\phi^*(\cdot))} \int_{\phi^*(\cdot)}^{\infty} \phi^{\eta-1} g(\phi) d\phi \right)^{\frac{1}{\eta-1}}$, this simplifies to

$$\left[\chi (1 - G(\phi^*(0))) \left(\left(\frac{\tilde{\phi}(0)}{\phi^*(0)} \right)^{\eta-1} - 1 \right) + (1 - \chi) (1 - G(\phi^*(\varepsilon))) \left(\left(\frac{\tilde{\phi}(\varepsilon)}{\phi^*(\varepsilon)} \right)^{\eta-1} - 1 \right) \right] = \frac{f_E}{f_p} \quad (\text{A.22})$$

Since $\phi^*(\varepsilon) = \Omega\phi^*(0)$ we can obtain the solution for $\phi^*(0)$ from equation (A.22). With the Pareto productivity distribution, condition (A.22) becomes

$$\chi (\phi^*(0))^{-\nu} + (1 - \chi) (\phi^*(\varepsilon))^{-\nu} = \frac{f_E (\nu - \eta + 1)}{f_p (\eta - 1)} (\varphi_{\min})^{-\nu}.$$

The measures of non-volatile firms N_B and volatile firms N_C that respectively use Business-as-usual and Churning strategies are given by $N_B = \chi[1 - G(\phi^*(0))]N$ and $N_C = (1 - \chi)[1 - G(\phi^*(\varepsilon))]N$ where N is the measure of entering firms. The share of non-volatile to volatile firms is determined by the productivity cutoffs $\phi^*(0)$ and $\phi^*(\varepsilon)$. With the Pareto productivity

distribution, this is given by

$$\frac{N_B}{N_C} = \frac{\chi}{(1-\chi)} \left(\frac{\phi^*(\varepsilon)}{\phi^*(0)} \right)^\nu = \frac{\chi}{(1-\chi)} (\Omega)^\nu.$$

The share of firms of each type co-moves with the relative productivity thresholds of non-volatile and volatile firms. We can show that for η high enough, a faster filling rate $\mu(M)$ and higher operating costs $R(M)$ ceteris paribus (holding the labor market tightness constant) increase the share of volatile churning firms as those benefit more from faster hiring and savings on operating in low demand periods. If we further assume that f_p is increasing with $R(M)$, we can show that for ν high enough, $\partial(\phi^*(\varepsilon))/\partial R(M) > 0$ or selection on productivity increases with city size. Given that $\partial\Omega/\partial R(M) < 0$, the productivity-density gradient is thus flatter for volatile firms.

Chapter 2

Moving to opportunity? Networks, information and skill-biased migration

Abstract: This paper studies the role of information frictions in workers' regional migration. I exploit a strategy that isolates changes in workers' information on migration opportunities based on their network of former coworkers. Using matched employer-employee data from France, I find that increases in information on migration opportunities lead to higher labor market transitions to other cities, but this effect is significantly weaker for low-skilled workers, who also tend to be less mobile and have more spatially concentrated networks. I then estimate a quantitative model of migration to assess the role of information frictions on the skill mobility gap. Results suggest that this channel contributes substantially to differences in migration propensities between low- and high-skilled workers and to the resulting wage inequality between these two skill groups. Moreover, information frictions magnify the misallocation across space of low-skilled workers, who, due to a weaker network and scarce information on their outside options, are less mobile and less likely to move to high-wage cities.¹

1. This chapter benefited from the guidance of my PhD supervisors Benoit Schmutz and Isabelle Mejean, and from the feedback of many conference and seminar participants, in particular at CREST.

2.1 — Introduction

High economic inequalities between local labor markets are a striking feature of many developed countries (Kline and Moretti, 2013; Bilal, 2023; Kuhn et al., 2021), yet labor mobility across cities and regions is relatively limited. In addition, data shows that it is skill-biased: high-skilled workers are more likely to move to another local labor market than low-skilled workers (Amior, 2024). The literature has emphasised high barriers to mobility limiting migration (Kennan and Walker, 2011), especially among low-skilled workers (Diamond, 2016; Balgova, 2018; Caldwell and Danieli, 2024). One example of such barriers are information frictions, namely the costs of gathering and processing information on a range of location attributes that can shape migration decisions.

The information required to weight the value of alternative locations is substantial, as migration decisions affect many aspects of workers' lives. One key aspect is the kind of job opportunities that each city can offer, and this information can be heavily influenced by the spatial distribution of social networks. Evidence shows that social networks are crucial for obtaining jobs (Granovetter, 1973; Topa, 2011). However, personal connections tend to be concentrated in the city where workers reside, while they are much less developed in other locations. As such, the spatial distribution of social networks may represent an important type of information friction. Since personal connections are key to access to information on outside options, their smaller presence outside the city of residence implies higher search frictions in these more distant local labor markets.

What is the role of social networks for accessing information on job opportunities in other cities? Is there inequality in networks and does this affect migration opportunities? To address these questions, I exploit a strategy that isolates changes in a worker's information set on migration opportunities from other factors that may drive migration. This strategy relies on the fact that individuals often learn about jobs through social contacts, including former coworkers (Caldwell and Harmon, 2019). I measure a worker's information about job opportunities by weighting plant-specific changes in labor demand by each worker's coworker network in that plant. These networks consist of former coworkers who have changed job and

allow to identify which new positions a worker is more likely to hear about in each city, where cities are identified by commuting zones (CZs).² Networks vary across workers within the same plant and occupation, and within workers over time, since they depend on individual mobility histories and on tenure in the current job. The data come from French matched employer-employee databases between 2010 and 2015, which are combined following Babet et al. (2022) to obtain a full panel of quarterly job histories for private-sector employees.

I start by reporting a set of stylized facts on how workers' mobility decisions between cities correlate with individual and time-specific measures of information on local job opportunities. My baseline indicator weights the number of new positions at each connected plant and location by an individual's exposure to that plant through her coworker network, as well as by a proxy of the quality of the worker's connection to the plant.³ The identifying assumption is that, conditional on the included covariates, unobserved determinants of individual mobility are uncorrelated with time-varying labor demand in potential destination locations. The main threat to identification is that the network structure and the variation in demand at connected plants and locations could proxy for other characteristics of the individual and location that also influence mobility decisions. I address this issue by adding a restrictive set of fixed effects (FEs) that control non-parametrically for possible omitted variables. In my baseline specification, I include sector by time FEs and worker FEs. The first control for unobserved changes in demand for worker skills that might be both correlated with exposure to local opportunities, and affect mobility decisions. Individual FEs instead rule out endogeneity of the network, namely the possibility that more mobile workers sort into more mobile coworker networks.

As a first stylized fact, I show that increases in labor demand at former coworkers' current workplaces in different cities raise the likelihood of migration decisions. In my baseline specification, a one standard deviation increase in the exposure to migration opportunities through networks raises the probability of migration in the next quarter by 1.8% of mean migration. This effect is entirely driven by increases in moves to coworker-connected locations

2. In what follows, I use interchangeably the terms commuting zone (CZ), city, and location.

3. I weight each coworker in the network by the amount of time that the worker spent together with her former colleague in the workplace where they originally met. This amount of time is divided by the worker's tenure in her former workplace, and thus expressed as a share.

(plants and CZs where the worker knows at least one coworker): a one standard deviation increase in information on migration opportunities raises mobility to connected CZs by 8.6% relative to mean migration. I examine the robustness of these results to a large set of alternative specifications, for example including firm FEs to control for unobserved shocks that are specific to workers of the same origin firm.

I then study the role of networks in providing information on migration opportunities among two groups of workers: high-skilled workers, defined as white-collar managers, and low-skilled workers, defined as blue-collar manual workers. The distinction between these two skill groups is motivated by the observation that networks are highly segregated by skill: networks of high-skilled workers are mostly composed of high-skilled workers, and similarly for low-skilled workers. If high-skilled workers are more mobile than low-skilled workers, their networks are likely to be more mobile. As a result, the first will have networks that are more spread out in multiple cities, through which they will be able to obtain more information on opportunities in other locations. This suggests that skill-biased networks, and the resulting differences in information frictions between skill groups, may in part explain skill-biased migration patterns.

Comparing the network characteristics of high- and low-skilled workers reveals that the latter have much smaller and spatially concentrated coworker networks. This pattern comes on top of a smaller diversification of local opportunities across CZs for low-skilled workers. By contrast, among high-skilled workers a few large cities represent booming centers with growing job opportunities. Networks exacerbate these dynamics: while spatially dispersed networks inform high-skilled workers about distant opportunities in growing cities, for low-skilled workers the similarity of job opportunities across cities combines with a network that is more concentrated in the city of residence. On top of these patterns, the elasticity of migration to information on local opportunities is also skill-biased: it is three times larger among high-skilled than among low-skilled workers.

These stylized facts not only suggest that the distribution of networks and of local job opportunities display significant differences across the skill distribution, but also that their role in migration decisions of high- and low-skilled workers may differ. At the same time, it is difficult to precisely disentangle the impact of information frictions. The larger effect of

information on the mobility decisions of high-skilled workers could simply be the result of having more to gain from migrating, due to the concentration of their job opportunities in few booming cities.

Building on these stylized facts, I thus develop a quantitative model of job mobility that enables to disentangle the effect of information frictions from other barriers to migration. I specify my model as a two-step nested logit: an upper model explains the decision to change job, and a lower model illustrates the job-location decision, including migration. In this framework, networks transmit information on job opportunities in each location. Information frictions arise from the fact that networks are mainly concentrated in one's city of residence, and smaller networks in other locations reduce the amount of information on job offers arriving from other cities. The expected value of living in each location also depends on publicly available information on the quality of the local labor market, on local amenities, and on migration costs. The latter include a fixed cost of moving and a home-bias component, reflecting the cost of moving to a city that is in a different region from the one of birth.

Estimating this model enables to compare the effect of private information on local job opportunities to other push and pull factors of migration. It also helps to quantify the impact of each channel on the migration gap between high- and low-skilled workers, and on the resulting wage inequality. In line with previous evidence in the literature (Kennan, 2015; Diamond, 2016; Balgova, 2018; Caldwell and Danieli, 2024), I first show that the variation in local amenities and in migration costs explains a large share of migration decisions, while the role of private information on local job opportunities is smaller but not negligible. In particular, private information reduces the migration probability because the majority of an individual's network is located in the city of residence, which limits knowledge on job opportunities in other local labor markets. The negative effect on migration propensity is stronger for high-skilled workers, who are more reactive to this channel, thus narrowing the overall migration gap between the two skill groups by 5.6%.

At the same time, by placing more weight on locations with better opportunities, this channel increases expected wages⁴ for both groups, but especially for high-skilled workers.

4. As detailed in Section 2.4.4, the expected wage is a weighted average of the mean skill-specific wage in each CZ, weighted by the model's predicted mobility rate to each CZ.

This results in higher wage inequality between the two skill groups, especially among high-wage destinations (cities that pay wages in the top quartile of the wage distribution), where the migration gap is larger. Low-skilled workers struggle to gain information on local job opportunities in cities that would offer them better wages, which makes it less likely for them to move there.

While the responsiveness of high-skilled workers to information on local job opportunities is substantially higher than that of the low-skilled, it could simply reveal that the latter group has less to gain from migrating to different cities because the distribution of their local opportunities is more spatially homogeneous. To disentangle the effect of networks from that of opportunities, I simulate a counterfactual network distribution at the skill-group, origin and destination cell, and examine what would happen if low-skilled workers had the same spatial distribution of coworker networks as high-skilled workers, holding constant the distribution of local opportunities. This exercise thus yields a counterfactual migration probability for low-skilled workers that is due to a larger and more spatially dispersed network, and thus to lower information frictions.

Results show that equating the spatial distribution of coworker networks in the two groups would reduce the skill mobility gap by 4.8%. This effect is reflected in expected wages, with a reduction of 12.2% in wage inequality between high- and low-skilled workers. This is a sizeable effect and suggests that information frictions matter substantially for the misallocation of workers across space and for their resulting access to local opportunities and to better wages. The lower response of low-skilled workers to private information on local job opportunities is not only caused by a more spatially homogeneous distribution of job opportunities, which makes migration less salient. Instead, it also reflects their weaker professional network, which co-locates less well with local job opportunities across space, thus giving rise to information frictions that limit the potential wage gains of low-skilled workers.

This paper is related to several existing literatures. There is a large quantitative economic geography literature studying mobility between local labor markets (Schmutz and Sidibé, 2019; Zerecero, 2021; Monras, 2020; Bilal, 2023). A strand of this literature has highlighted the high barriers to mobility that limit internal migration (Kennan and Walker, 2011), in particular among low-skilled workers (Topel, 1986; Bound and Holzer, 2000; Wozniak, 2010;

Moretti, 2011; Kennan, 2015; Diamond, 2016; Balgova, 2018; Schmutz et al., 2021; Caldwell and Danieli, 2024). However, most studies of internal migration do not consider an explicit information channel, or, when they do, they tend to focus on specific settings, based on natural experiments, such as the impact of news about opportunities related to the US fracking boom (Wilson, 2021), or randomized information treatments (Gharad et al., 2014). Nevertheless, there is growing evidence that information shocks affect job application behavior, especially for those workers who live further away from posted vacancies (Skandalis, 2019).

Other recent work has highlighted the fact that information transmission often takes place through social networks, and in particular that the role of coworkers is important in providing information on job opportunities. For example, Caldwell and Harmon (2019) show that increases in labor demand at former coworkers' current firms lead to job-to-job mobility and wage growth. These studies have so far not considered the role of information networks as a friction affecting mobility between cities. Conversely, networks have been thoroughly studied in the international migration literature, as a source of information on jobs and conditions in the destination country (see Munshi, 2020, for a detailed review). One close paper that combines mobility between local labor markets with information frictions is Porcher (2022), who however exploits information acquired through weak ties, assuming that workers in same municipality share information on job opportunities in previously visited locations. My paper aims to bridge the literatures on regional migration, information frictions and the role of social networks, by shedding light on how inequality in access to information affects migration decisions, and how this differs across the skill distribution of workers.

The remainder of the paper is organized as follows. In Section 2.2, I describe the data, provide descriptive statistics, and detail my measure of information on migration opportunities. In Section 2.3, I outline a reduced-form estimation strategy and describe the stylized facts. I characterize the heterogeneity in local opportunities, in the distribution of coworker networks and in migration responses across the skill distribution. I then detail my quantitative framework in Section 2.4 and present results of a counterfactual exercise. Section 2.5 concludes.

2.2 — Data and measurement

In this Section, I describe the data, how coworker networks are identified, and how information on local job opportunities is measured.

2.2.1. Data

Following Babet et al. (2022), I use 2005-2019 data from the matched French employer-employee cross-sectional datasets (INSEE-DADS) to construct a quarterly exhaustive panel (see Appendix B.1.1 for additional details). DADS contains rich information describing job posts and workers, such as age and gender of the worker, occupation and sector.

In the main analysis, I focus on the 2010-2015 period, thus on 24 consecutive quarters. The choice of quarterly frequency is motivated by the need to obtain a precise description of work histories, while exploiting a sufficiently long period of time over which migration decisions can mature. I use the 280 commuting zones (CZs) in continental France as main geographical units to describe urban areas. CZs cover continental France exhaustively and capture areas where the majority of workers live and work.

2.2.2. Descriptive statistics and sample restrictions

Workers — The first column of Table 2.1 provides descriptive information on the set of workers who appear at least once in DADS data over the period 2010-2015. I restrict the sample to males with the objective of better identifying voluntary moves, since evidence shows that females are more likely to move to follow their male partner (Patnaik et al., 2022). The average worker in this sample is around 43 years old and earns a mean hourly wage of €22.3 before taxes.

I distinguish between high- and low-skilled workers, which I identify by managers and manual workers respectively. The advantage of focusing on these two groups of workers in the skill distribution is that they rarely intersect⁵ and therefore operate on two separate

5. In the 2010-2015 estimation sample described below, only 0.10% of individuals switch between these two broad skill classes.

labor markets. Almost 57% of individuals in the sample are low-skilled workers, a majority has had a permanent contract⁶ and has worked full-time at least once over this period. Moreover, around 5% of workers has made at least one migration over the 2010-2015 period, defined as a between-firm job-to-job transition to a new CZ. Panel B of Table 2.1 describes some characteristics of the employment history of these workers over the 2010-2015 period, including the mean number of plants (1.4) and firms (1.3) at which they work.

Table 2.1 — Descriptives: workers

	All workers	Clean sample
Number of worker-quarter observations	122,074,757	81,984,630
Number of workers	8,493,404	5,531,862
<i>A. Demographics</i>		
Age	42.5 (10.1)	42.5 (10.2)
% low-skill workers	56.8%	59.8%
% ever permanent contract	83.4%	91.4%
% ever full-time work	90.8%	92.5%
% ever migrated	5.0%	4.6%
<i>B. Employment</i>		
Hourly wage (€)	22.3 (57.3)	19.6 (39.1)
Number of plants	1.4 (0.7)	1.4 (0.8)
Number of firms	1.3 (0.6)	1.3 (0.6)
Number of occupations	1.6 (1.0)	1.7 (1.1)
Number of sectors	1.1 (0.4)	1.1 (0.4)
Number of quarters employed	9.5 (5.8)	10.5 (6.8)

Notes: The first column describes the entire INSEE-DADS dataset of male workers who appear at least once in the period 2010-2015. The “clean” sample is a sub-sample of male workers that appear at least once in the period 2010-2015, obtained after cleaning the data as described in the main text. In both columns, summary statistics are calculated over the period 2010-2015. The first entry in each row is the mean. Standard deviations are reported in parentheses. The number of sectors is measured at the plant level. In the cleaned panel, the hourly wage is winsorized at the bottom and top 5th percentiles.

Column 2 of Table 2.1 reports the same summary statistics for the clean sample of workers

6. In France, permanent contracts (known as “contrat à durée indéterminée,” or CDI) are defined as contracts with no time limit, fulfilled on a full-time or part-time basis between employer and employee.

that is the basis for my empirical analysis. Compared to column 1, I first restrict the panel to each worker’s main job in each quarter and I focus on workers between the ages of 25 and 60 that are born in France, to obtain a sample of workers for which work histories are most complete. Moreover, I limit the sample to employees that have always worked in the private sector, thus excluding civil servants and self-employed workers. Finally, I focus on workers that have always worked in plants with size of 2-500 employees, since in large workplaces it is unlikely that a worker knows or shares information with all of her former coworkers (Caldwell and Harmon, 2019).⁷ This cleaned sample is described in the second column of Table 2.1 and includes 5.5 million male workers. Appendix B.1.2 provides additional details on the data cleaning procedure. Table B.1 in Appendix B.1.3 describes the firms in the cleaned sample.

Migration patterns and wage inequality — Migration patterns in the data are reported in Table 2.2. Over 2010-2015, 5.7% of individuals move at least once between CZs and 2.4% of migration events involve a return to the region of birth.⁸ Workers who move make on average 1.1 migrations over the period. Table 2.2 also reports the share of individuals that move at least once between municipalities, and between larger aggregations than CZs such as regions.

There are noticeable differences between the migration patterns of low- and high-skilled workers. The latter are twice more likely to move than low-skilled workers, they move on average 1.4 times further away, and they tend to reside in larger cities. At the same time, the two skill groups are relatively similar in terms of the share of return migrations, the average age of migration and the number of migration events among movers. Tables B.2 and B.3 in Appendix B.1.3 describe migration flows by skill group between some of the largest CZs in France for the period 2010-2015.

High- and low-skilled workers are also characterized by different labor market returns of job mobility. Table B.4 in Appendix B.1.3 highlights that, among high-skilled workers,

7. This type of restriction is common in the network literature. For instance, Hensvik and Skans (2016) study firms with less than 500 employees. Caldwell and Harmon (2019) use connections formed in firms of 2-1000 employees, while Saygin et al. (2021) include connections formed in all firms with less than 3000 workers, but only within firms involved in mass layoffs.

8. In INSEE-DADS data, region of birth is available at the department (“département”) level. The department is one of the three levels of the French government under the national level, corresponding to NUTS 3 of the European territorial units for statistics.

Table 2.2 — Migration patterns

	All	Low-skilled	High-skilled
<i>A. % of workers moving at least once</i>			
between municipalities	11.1%	10.0%	16.6%
between CZs	5.7%	5.2%	9.5%
between department	5.6%	4.4%	11.0%
between regions	2.5%	1.9%	4.9%
% of return migrations	2.4%	3.2%	1.2%
<i>B. Characteristics of CZ migrations</i>			
Age of first migration	35.5	36.9	35.3
	(9.5)	(9.9)	(9.1)
Age of migration	35.6	36.9	35.3
	(9.5)	(9.9)	(9.0)
Number of migrations	1.1	1.1	1.2
	(0.4)	(0.4)	(0.4)
Migration distance (km)	173.8	144.8	204.9
	(198.5)	(182.3)	(208.0)
Density of origin CZ	1294.3	983.4	1695.3
	(2541.1)	(2197.0)	(2865.0)
Density of destination CZ	1256.9	978.4	1670.0
	(2506.2)	(2192.2)	(2850.8)

Notes: Summary statistics are based on the “clean” sample of male workers over the period 2010-2015 described in column 2 of Table 2.1. The first entry in each row is the mean. Standard deviations are reported in parentheses. The % of return migrations is calculated as the number of migrations to a CZ in the same department of birth, divided by the number of total migrations between departments. Density of origin or destination CZs is defined as the number of working-age individuals in each CZ, divided by CZ area in squared meters (see Appendix B.1.3 for additional details).

migration transitions yield higher wage gains than job-to-job transitions within the same CZ, while this is not the case on average for the low-skilled. At the same time, expected wages for the two groups vary substantially over space, resulting in large differences in wage inequality across cities, as shown in Figure B.1 of Appendix B.1.3. The highest wage gaps are observed in the largest cities such as Paris and its metropolitan area.

Estimation sample — After having cleaned the data as detailed above, I transform the panel into a single “spell” of zeros and ones to estimate a discrete-time duration model of job mobility and location choice over the period 2010-2015. A spell consists of zeros between the first quarter of 2010 and the last quarter before the first job-to-job transition. If a transition takes place between period t and period $t + 1$, period $t + 1$ is coded as 1. I drop

individual-period observations after recording a gap in employment history in the 2010-2015 sample. These gaps could be due to unemployment or non-participation spells, to a switch to self-employment, or to a move abroad, all of which are not documented in my data. To reflect the idea of a duration model as in Buggle et al. (2023), I drop observations after the first job-to-job transition has taken place. Finally, I keep the subset of individuals for which a job-to-job transition is observed over the 2010-2015 period to examine the timing of job mobility decisions and specifically the role of information on outside options on this decision.⁹ This selection identifies over 700 thousand male workers who represent my main estimation sample. I further obtain two additional subsamples for each skill group, one of all low-skilled male workers, and one of all high-skilled male workers. The job transitions of these samples are described in Table B.5 of Appendix B.1.3.

2.2.3. Measuring information transmission

My empirical strategy exploits the fact that individuals often learn about new job opportunities through their former coworkers. Moreover, they should be more likely to hear about job opportunities when their former coworkers' firms are expanding more relative to other periods. While this takes inspiration from Caldwell and Harmon (2019), I am specifically interested in *migration* opportunities and thus in the information that spatially dispersed networks can provide on job opportunities in distant cities.

Coworker networks

I use coworker networks to measure information transmission on local job opportunities.¹⁰ To precisely estimate tight connections, I identify each worker's former colleagues as those that have worked in the same *plant*. I then select ex-colleagues that have changed job, where a change in job is defined as a transition to a new firm between consecutive quarters. When such job-to-job transitions are also characterized by a change in CZ, they are defined as

9. In robustness checks not reported here, I reproduce these results using an alternative sample where individuals are not dropped after making their first job-to-job transitions.

10. In Appendix B.1.4 I compare the role of professional networks relative to other types of social networks when it comes to finding jobs, based on the French Labor Force Survey.

migrations. These job-to-job transitions connect the worker to (the plants of) new firms and potentially to new CZs. The key idea is that a worker is more likely to know about job opportunities in plants and locations where she has connections. Crucially, I keep the plant dimension when identifying where former coworkers have moved to, because coworkers are more likely to hear about job opportunities in their current plant rather than in other plants of the same firm.

There are two additional restrictions that I impose, following Caldwell and Harmon (2019). First, I exclude connections that were formed more than three years ago, both because older connections are likely to be less informative, and because, without a fixed window, network size or quality would vary mechanically over the sample window.¹¹ Second, I remove connections who are working at plants or firms where the individual worked in the past three years so that the network does not vary mechanically with mobility.¹²

As illustrated in Figure B.2 of Appendix B.1.4, for two workers that met in the same plant, networks may differ for two key reasons. Firstly, networks may vary between workers who previously worked at different firms (or plants within the same firm) and thus developed different networks. Secondly, networks may vary between workers who joined the same plant at different times, due to when precisely other former colleagues joined or left the plant.

Table 2.3 provides some summary statistics on coworker networks. The first column shows that, on average and in each quarter, workers know 8.3 coworkers outside of their own firm. These coworkers link them to 6.0 distinct plants of other firms. However, the distribution is very skewed (the median worker has only 1 former coworker), and the standard deviation of the number of coworkers is high (33.0), compared to the mean. The second and third columns compare the networks of low- and high-skilled workers and show that low-skilled workers have significantly weaker networks by most metrics: the number of connections and connected plants, the number of connected sectors and CZs, or the average size of a connection's plant. In other words, coworker networks are skill-biased: low-skilled workers tend to have smaller,

11. I use the full 2005-2019 dataset to construct the network and ensure that individuals at entry in the estimation sample are observed for at least three years preceding entry.

12. If a worker moves from plant A to plant B, I do not include her former coworkers at A in the network, unless they move to another plant or firm. If I did not do this, workers who moved more would mechanically see a large increase in network size. In addition, the network would be heavily weighted towards the plant that they just left.

less diverse, and more concentrated networks. Moreover, networks are highly segregated: low-skilled workers are more likely to be connected to firms that mostly hire low-skilled workers and that pay lower wages, while the opposite is true for high-skilled workers. Lastly, most of the variation in coworker networks comes from that between-workers within-time.¹³

Table 2.3 — Network characteristics

	All	Low-skilled	High-skilled
<i>A. Characteristics of coworkers</i>			
Number of coworkers	8.3 (33.0)	4.3 (15.2)	14.5 (53.7)
% in non-home CZ	31.5% (38.2)	33.0% (40.1)	31.5% (37.4)
<i>B. Characteristics of connected plants</i>			
Plant size	359.6 (726.4)	252.4 (589.5)	387.9 (700.1)
Hourly wage (€)	19.1 (9.4)	17.4 (8.8)	23.4 (10.8)
% low-skilled workers	32.1% (26.8)	49.9% (26.5)	22.4% (22.2)
% high-skilled workers	26.4% (22.2)	18.6% (16.3)	42.5% (24.2)
Number of sectors	4.3 (5.0)	3.4 (3.6)	5.7 (6.6)
Number of plants	6.0 (20.5)	3.3 (8.5)	10.1 (33.7)
Number of CZs	2.2 (5.2)	1.6 (3.4)	2.8 (6.5)
Weighted CZ distance	95.8 (117.2)	98.9 (121.0)	101.0 (120.5)

Notes: Summary statistics are based on the “clean” sample of male workers over the period 2010-2015 described in column 2 of Table 2.1. The first entry in each row is the mean. Standard deviations are reported in parentheses. Plant size is measured by the number of employees. The number of sectors refers to sectors of connected plants. Weighted CZ distance uses the number of coworkers in each plant to weight the distance between CZs (measured in km) where each plant is located.

13. Focusing on the 2010-2015 estimation sample of all male workers, the between-worker standard deviation in the number of coworkers represents 85% of the overall standard deviation.

Measuring information transmission

Coworker networks expose workers to information about job opportunities in other firms and locations. Information on job opportunities can vary over time, as connected firms expand and contract. The interaction of exposure through the network and differential growth rates at connected firms generates an information shock that resembles a shift-share instrument.¹⁴

To study the transmission of information through networks, I build a variable that captures changes to local opportunities at plants to which the worker is connected through her network. For each worker i , location j , and quarter t , this variable is defined as the sum (among the set F_{ijt} of connected plants in location j) of changes in demand at each connected plant f , multiplied by the share of the worker's connections in that plant

$$\text{Info}_{ijt} = \sum_{f \in F_{ijt}} \left(\mathbf{s}_{ifjt} \mathbf{v}_{fjt} \mathbf{q}_{ifjt} \right), \quad (2.1)$$

where \mathbf{s}_{ifjt} is the share of the individual's network in each plant f and \mathbf{v}_{fjt} indicates changes in demand at plant f corresponding to the number of new positions relative to the previous quarter: $\mathbf{v}_{fjt} = (\text{Empl}_{fjt} - \text{Empl}_{fj,t-1})^+$.¹⁵ Finally, \mathbf{q}_{ifjt} is a weight which aims to capture quality of information pass-through. In my baseline specification, I weight each coworker in the network by the amount of time that the worker spent together with her former colleague in the workplace where they originally met. This amount of time is divided by the worker's tenure in her former workplace, and thus expressed as a share.

In the data, variation in Info_{ijt} within an individual is driven by changes in plant demand, not changes in network composition. Table B.7 in Appendix B.1.4 shows that the number of connections in a worker's coworker network are highly autocorrelated, with autocorrelations above 0.8 after one year. Other characteristics of networks are also relatively stable over time. By contrast, the number of new positions at connected plants is significantly less autocorrelated. This suggests that the timing and magnitude of expansions at connected

14. Figure B.3 in Appendix B.1.4 shows the timeline of coworker network formation, the arrival of the information shock, and the potential decision to change job.

15. Following Caldwell and Harmon (2019), I focus on the number of new positions, rather than the overall number of hires, which reflects changes in both labor demand and churn.

firms is likely to be random to the individual worker, thus helping to identify its causal effect on job mobility decisions.

2.3 — Stylized facts

In this Section, I present a set of stylized facts on the role of exposure to information on job opportunities and its impact on job mobility and migration.

2.3.1. Empirical strategy

A first stylized fact that the data enables to uncover is whether individuals are more likely to change job after having been exposed to information on job opportunities through their former coworkers. I estimate a discrete-time duration model of the job mobility decisions. Section 2.2.2 explains how I transform individual-level information into spells, where each person can change job once. My outcome of interest is thus a dummy that takes the value one in the quarter t in which individual i changes job. In the year $t + 1$ after the job transition occurs, the individual exits the sample. I focus on two different types of transitions: all job-to-job transitions, and the subset of those that involve a change in CZ, which I define as migrations. This leads to the following specification

$$Prob[\text{Transition}_{it} = 1] = \Phi(\alpha_i + \alpha_{kt} + \beta \text{Info}_{it} + \mathbf{X}'_{it} \delta + \gamma \mathbf{n}_{it} + \epsilon_{it}). \quad (2.2)$$

In this equation, $\text{Info}_{it} = \sum_j \text{Info}_{ijt}$ measures exposure to information on job opportunities. When studying the impact on any job-to-job transition, the relevant information set is the sum of information on job opportunities in the home CZ ($j = h$) and in other potential destinations ($j \neq h$). When the dependent variable is migration transitions, the focus is on information on job opportunities in destinations that differ from the CZ of residence ($j \neq h$). The variable $\mathbf{n}_{it} = \sum_j \sum_{f \in F_{ijt}} \mathbf{n}_{ifjt}$ measures the size of the individual's network in each period t , summed over all connected plants and cities. A comprehensive set of individual characteristics is included in \mathbf{X}_{it} : age and its square, whether the contract is fixed-term,

whether the job is part-time, tenure in the firm, and log wage.

The identifying assumption to estimate equation (2.2) is that, conditional on the included covariates, the distribution of coworker networks and changes in labor demand at connected plants are uncorrelated with unobserved determinants of mobility. However, the network structure and the variation in demand at connected plants and locations could proxy for other characteristics of the individual and location that also influence mobility decisions. To address the threat of omitted variables, I include individual FEs (α_i) and sector \times quarter FEs (α_{kt}). Adding this restrictive set of FEs aims to tackle two types of unobserved factors.

The first key challenge is posed by network endogeneity (Manski, 2000; Moffitt, 2001). This is the possibility that mobile workers sort into mobile coworker networks, leading to an endogenous choice of network. For example, workers with higher mobility preferences may be more likely to work with colleagues that have similar preferences. The mobility decision of a worker and her network may be determined by unobserved preferences or factors such as ability and ambition that are common across workers that have collaborated in the past. To deal with worker sorting, I include individual FEs (as well as individual-clustered standard errors), which account for differences in mobility preferences.¹⁶

Secondly, unobserved changes in demand for worker skills could be correlated with information on job opportunities and also affect mobility decisions. I include sector \times quarter (or occupation \times quarter or sector \times occupation \times quarter) FEs to absorb aggregate variation in demand for specific skills. In Section 2.3.2, I discuss a set of robustness checks, which provide further evidence that results are driven by changes in workers' information about their migration opportunities.

The functional form Φ in equation (2.2) depends on the estimation procedure. Although nonlinear estimators are preferable given my data structure, the linear probability model (LPM) offers the advantage of dealing with high-dimensional FEs, allowing for a rich clustering structure. My baseline regressions are estimated using the LPM, while in Appendix B.2.1 I present results using logit models.

16. Unobserved mobility preferences are controlled for by individual FEs as long as they are fixed over the sample period. While workers may change their unobserved propensity to move over time, the fact that I use a relatively limited sample period in my analysis (2010-2015) reduces omitted variable bias originating from changes in workers' "mobility" type over time.

2.3.2. Impact of information on job transitions

In what follows, I focus on the estimation sample of all male workers described in Table B.5 and I study their work history over the period 2010-2015. I exploit within-individual variation (across time) in exposure to information on mobility opportunities. I thus examine whether workers are more likely to change job in periods in which they are more exposed (through their network) to information on new job opportunities, after controlling for changes in demand at the sector level.

Results for all workers

Table 2.4 displays estimates from equation (2.2) using the linear probability model (LPM). In the first three columns the outcome variable is an indicator for whether the individual makes a job-to-job transition in the next period. Info_{it} thus combines information on local ($j = h$) and migration ($j \neq h$) job opportunities. Model (1) includes sector \times quarter FEs and individual FEs. Model (2) replaces sector \times quarter FEs with occupation \times quarter FEs, while model (3) uses sector \times occupation \times quarter FEs. The three models also control for the worker's total number of connections (\mathbf{n}_{it}), as well as for individual characteristics. The last three columns focus on job-to-job transitions involving a change in CZ and replicate models (1)-(3). In this case, Info_{it} includes only the relevant information set, namely migration opportunities in CZs other than the one of residence.

Table 2.4 shows that workers are more likely to make a job-to-job transition after quarters in which they are more exposed to information on outside options. In model (1), a one standard deviation increase in Info_{it} leads to a 0.4 percentage points higher probability that an individual will change job in the following quarter (3.2% of the sample mean). These results are consistent with the idea that individuals stay in contact with and discuss job opportunities with their former coworkers. When new vacancies arise at former coworkers' plants, the coworker is more likely to hear about these increased opportunities than the average worker, and may communicate this to her former colleagues. The remaining coefficients for models (2) and (3) indicate that estimates remain stable when including different non-parametric controls for changes in demand for certain skill groups (at the sector and occupation level),

Table 2.4— Baseline estimates

Dependent variable	Job-to-job transition			Migration transition		
	(1)	(2)	(3)	(4)	(5)	(6)
Mean		11.7			3.6	
Info_{it}	0.373 (0.032)	0.380 (0.032)	0.350 (0.030)	0.063 (0.010)	0.061 (0.010)	0.060 (0.010)
n_{it}	1.881 (0.041)	1.946 (0.042)	1.865 (0.043)	0.764 (0.022)	0.762 (0.022)	0.768 (0.023)
Individual controls	✓	✓	✓	✓	✓	✓
Individual FEs	✓	✓	✓	✓	✓	✓
Sector× <i>t</i> FEs	✓			✓		
Occ× <i>t</i> FEs		✓			✓	
Sector×Occ× <i>t</i> FEs			✓			✓
Adjusted <i>R</i> ²	0.2066	0.1964	0.2303	0.1320	0.1327	0.1442
Observations	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762

Notes: LPM regressions based on the 2010-2015 estimation sample of all male workers. In the first three columns **Info_{it}** includes information on local ($j = h$) and migration ($j \neq h$) job opportunities, while in the last three columns it comprises only information on migration opportunities. Estimates are expressed in percentage points. Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

suggesting the results are driven by information transmission through networks and not by unobserved demand shocks. In addition, the scale of the network (**n_{it}**) has a positive effect on the transition probability. A larger network scale implies a larger number of coworkers that have moved to new firms, and thus a larger information set.

Similar results are observed when focusing on migration decisions: higher information on migration opportunities increases mobility across CZs. In model (4) a one standard deviation increase in **Info_{it}** increases the probability that an individual will change job and CZ in the following quarter by 0.06 percentage points (1.8% of the sample mean). While this effect is contained, information on local job opportunities has substantial influence on the direction of migration, increasing moves to coworker-connected locations (CZs where the worker knows at least one coworker) as shown in Table 2.5. In model (1), a one standard deviation increase in **Info_{it}** raises mobility to connected CZs by 5.6% relative to mean migration. The effect is even stronger if we consider migrations to connected plants: in model (1), a one standard deviation increase in **Info_{it}^M** raises mobility to connected plants by 8.6% relative to mean migration.

Table 2.5 — Connected and unconnected migration decisions

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Migration to connected and unconnected CZs</i>	Connected CZ mean = 0.6			Unconnected CZ mean = 3.0		
Info_{it}^M	0.032 (0.005)	0.031 (0.005)	0.031 (0.005)	0.032 (0.007)	0.030 (0.007)	0.029 (0.008)
Adjusted R^2	0.0619	0.0615	0.0754	0.1386	0.1397	0.1495
Observations	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762
<i>Panel B. Migration to connected and unconnected plants</i>	Connected plant mean = 0.2			Unconnected plant mean = 3.4		
Info_{it}^M	0.014 (0.002)	0.014 (0.002)	0.013 (0.002)	0.049 (0.009)	0.048 (0.009)	0.047 (0.009)
Adjusted R^2	0.0495	0.0482	0.0686	0.1335	0.1344	0.1448
Observations	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762

Notes: LPM regressions based on the 2010-2015 estimation sample of all male workers. Estimates are expressed in percentage points. Info_{it}^M denotes information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardised and individual-clustered standard errors are reported in parentheses. All models include individual controls (age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage), individual FEs, and network size \mathbf{n}_{it} . Model (1) adds sector $\times t$ FEs, model (2) instead includes occupation $\times t$ FEs, while model (3) controls for sector \times occupation $\times t$ FEs.

The LPM is my reference specification, since it can be estimated quickly on a very large dataset, even with computationally demanding FEs. However, I also present results of a discrete choice (logit) model that are consistent with the results obtained by LPM (see Table B.9 in Appendix B.2.1).¹⁷

In Section B.2.2 of the Appendix, I report a large set of robustness specifications, providing further evidence that results are driven by changes in workers' information about their migration opportunities. A first robustness test is to examine how information on *local* job opportunities affects migration decisions, as this should play a minor role relative to information on migration opportunities (Table B.12). In Table B.13 I exploit within-firm variation to control for additional sources of unobserved changes in demand. I also compare the baseline model, where the weight \mathbf{q}_{fjt} captures the quality of information pass-through, with specifications where it is set to 1 (Table B.14). In Table B.15 I proxy the firm's quality by the wage that it pays relative to mean retributions in the same CZ, or in the same CZ and

17. To speed up computation, results are estimated on a random subsample of 10,000 individuals taken from the main estimation sample. Logit models do not allow to control for high-dimensional FEs, hence these models do not include individual FEs. For comparison, Table B.10 in Appendix B.2.1 reports results of the same models without individual FEs estimated by LPM.

sector. Table B.16 checks the robustness of my results to using occupation-specific coworker networks. Table B.17 examines how the strength of information pass-through varies by the size of plants where former coworkers met. Lastly, Table B.18 shows that changes in labor demand at plants of more recent former coworkers matter more than changes in plants of coworkers with whom individuals worked with in the more distant past. As a placebo test, it also highlights that changes in demand at workers' future coworkers' firms have a negligible effect on mobility decisions.

Heterogeneity by skill group

Coworker networks play an important role in connecting workers to migration opportunities, but their distribution differs substantially between skill groups, as seen in Table 2.3. In this Section, I report four additional empirical regularities, which further highlight that the transmission of information through coworker networks varies across the skill distribution.

A first consideration is that the variation in migration opportunities across cities is significantly larger for high-skilled than for low-skilled workers. A more homogeneous distribution of local opportunities among low-skilled workers reduces their incentives to move, since no alternative location offers a significant improvement compared to the current residence. Figure 2.1 maps, for each skill-specific distribution, mean migration opportunities in each CZ in shaded colors, normalized by the sample mean. A value of 1 implies that local opportunities correspond to the mean in the CZ distribution, while a value greater (or smaller) than 1 indicates a larger (respectively smaller) number of opportunities in a given city. These maps suggest that, compared to low-skilled workers, migration opportunities of high-skilled workers display higher spatial variation and a stronger concentration of offers in specific high-growth cities.¹⁸

Figure 2.1 also maps the (average) spatial distribution of the coworker network for individuals living in Clermont-Ferrand, a medium-sized city in central France (and the largest red node in each map). The size of each node represents the scale of the network in each CZ,

18. Nevertheless, the spatial variation in offers is still significant for low-skilled workers: across CZs, the standard deviation of offers for low-skilled workers is 63% that of high-skilled workers (see panel A of Table B.8 in Appendix B.1.4, which provides additional details on the skills-specific distributions of coworker networks and local opportunities).

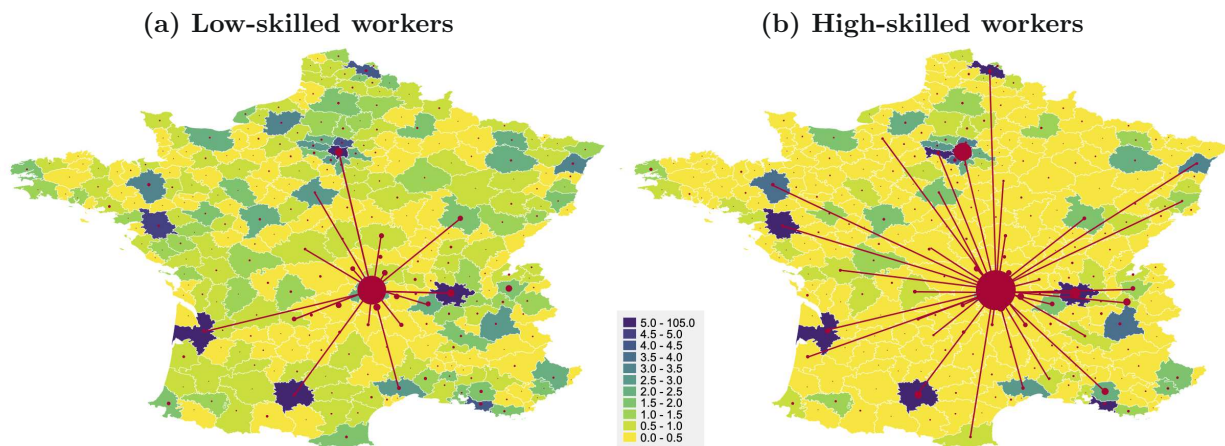


Figure 2.1 — Migration opportunities and networks for the average worker living in Clermont-Ferrand

Notes: The shaded colors depict total job creation in each CZ over the period 2010-2015 by skill group. Local opportunities are measured as in Section 2.2.3 by the number of quarterly new positions within each firm: local opportunities in city j and quarter t are defined as $v_{jt} = (E_{jt} - E_{jk,t-1})^+$. I distinguish positions by the type of occupation hired, and sum all local opportunities within each CZ-quarter-skill cell. I then average all local opportunities over the period 2010-2015, for each CZ-skill cell. Finally, total opportunities in each CZ are normalized by the mean across CZs (within each skill group distribution) and the resulting ratio is plotted. The red nodes represent the average size of the coworker network in each CZ in the range $[0, 3]$, based on the skill-specific estimation samples of male workers over the period 2010-2015 described in Table B.5. I focus on the subset of workers living in Clermont-Ferrand (the largest node in the map). I then take the average size of the coworker network in each CZ over the period 2010-2015.

in the range $[0, 3]$. A comparison of low- and high-skilled workers highlights a second stylized fact: the smaller and less geographically dispersed networks of low-skilled workers imply that they receive less information about migration opportunities.

A third element is that the networks of high-skilled workers tend to better co-locate with local opportunities. By contrast, those of low-skilled are less well distributed, which makes it more difficult to hear about migration opportunities. Figure 2.1 overlays the spatial distribution of local opportunities and that of the network for the average high- and low-skilled worker in Clermont-Ferrand, showing that networks better co-locate with opportunities among high-skilled workers (correlation of 0.20) compared to the low-skilled (correlation of 0.10).¹⁹ Table B.8 in Appendix B.1.4 puts together the previously discussed intuitions by reporting summary statistics on local opportunities, networks, and network-weighted opportunities across CZs.

19. Correlations are computed separately by skill, using data of panels (a) and (b) of Figure 2.1.

Overall, these stylized facts suggest that coworker networks may be more valuable for high-skilled workers, as they provide access to spatially dispersed opportunities that are further away from their current residence. I thus analyze the response to information on migration opportunities separately by skill group, replicating the analysis in Section 2.3.2 for high- and low-skilled workers. The results, described in Table 2.6, highlight that the effect of information on migration opportunities is skill-biased: low-skilled workers are significantly less responsive to information on migration opportunities, compared to high-skilled workers.

For high-skilled workers in model (1), a one standard deviation increase in Info_{it}^M leads to a 0.18 percentage points increase in the migration probability (a 4.9% increase relative to mean migration). For low-skilled workers, this value is 0.06 percentage points (1.9% of mean migration). Hence, for high-skilled workers, the effect of Info_{it}^M is three times larger than that of low-skilled workers. Table B.11 in Appendix B.2 reports results estimated using a discrete choice (logit) model, which are consistent with those estimated with the LPM.

Table 2.6 — Migration transitions by skill group

Dependent variable	Migration transition					
	Low-skilled			High-skilled		
Mean	2.9			3.7		
	(1)	(2)	(3)	(1)	(2)	(3)
Info_{it}^M	0.056 (0.013)	0.057 (0.013)	0.054 (0.013)	0.181 (0.033)	0.179 (0.033)	0.174 (0.032)
\mathbf{n}_{it}	1.482 (0.036)	1.481 (0.036)	1.567 (0.038)	1.960 (0.039)	1.959 (0.039)	1.952 (0.039)
Individual controls	✓	✓	✓	✓	✓	✓
Individual FEs	✓	✓	✓	✓	✓	✓
Sector-t FEs	✓			✓		
Occ-t FEs		✓			✓	
Sector-Occ-t FEs			✓			✓
Adjusted R^2	0.0970	0.0959	0.1052	0.0733	0.0730	0.0822
Observations	3,405,396	3,405,396	3,405,396	2,225,897	2,225,897	2,225,897

Notes: LPM regressions estimated on the skill-specific 2010-2015 estimation samples of male workers, whose transitions are described in Table B.5. Estimates are expressed in percentage points. Info_{it}^M denotes information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

Thus, high-skilled workers not only see a larger variance in their migration opportunities across space, with jobs being highly concentrated in a few high-growth cities, but they also

have a better knowledge of migration opportunities, thanks to a larger and more spatially dispersed network. The presence of former coworkers in these more dynamic cities can provide access to key information on job opportunities, to which high-skilled workers are ultimately more responsive.

The natural question that arises is whether this evidence points to the presence of information frictions or whether it simply reveals that low-skilled workers have less to gain from migrating to different cities. If information frictions were null, then the observed differences in migration patterns between high- and low-skilled workers would reflect true differences in local opportunities, mobility costs and location preferences. By contrast, the presence of information frictions is likely to distort mobility decisions and the spatial distribution of workers, generating misallocation patterns.

To understand the role of information frictions for the misallocation of workers, I build a quantitative model that can be taken to the data. This enables to assess the role of private information on migration opportunities by comparing it to other push and pull factors of migration, such as mobility costs and public information on the quality of the local labor market. It also enables to disentangle the effect of three key mechanisms: the spatial distribution of local opportunities, that of networks, and the co-location of the two.

2.4 — A quantitative model of migration with networks

In this Section, I build a random utility model of location choice and estimate it in a nested logit setup. The nested logit model is a natural framework for the data at hand: I observe job mobility decisions as a binary choice repeated in every quarter between 2010 and 2015. Conditional on deciding to change job, I observe the discrete choice of location.²⁰

My analysis considers a repeated static choice, in the sense that individual job mobility decisions do not factor in the expected future realizations of relevant job mobility determinants. Moreover, my setup does not model the general equilibrium feedback effect of job mobility on economic activity, neither in the origin city nor in the destination local labor markets. I

20. A similar approach has been used by Bugge et al. (2023) to study migration decisions.

focus on the worker side only, assuming that workers do not bargain over wages, and I do not model unemployment.²¹

In my setup, the expected value of each location depends on the quality of the local labor market, as well as on moving costs. Workers can also learn about local job opportunities through their individual-specific networks. Information frictions arise from workers having an incomplete knowledge of job opportunities. The biased geography of their private information on job opportunities implies that information frictions systematically distort the arrival of job opportunities across individuals.

2.4.1. Model setup

In the model, the first decision faced by each worker is whether to change job, which can be summarized as a choice between two nests. The first nest, denoted by S , represents the choice of staying in the worker's current job and, implicitly, in her current home location $j = h$. The alternative nest, denoted by M , identifies the choice of moving to a new job and thus comprises jobs in both the home location $j = h$ and in other destinations $j \neq h$. In a second step, if the worker has decided to change job, she will choose a specific destination $j \in M$. The decision to stay in her current job or to change job is referred to as the upper-level model, while the destination choice represents the lower-level model.

Workers choose the location that maximises their expected utility U_{ijt} from living in city j in period t , defined as

$$U_{ijt} = W_{iNt} + V_{ijNt} + \varepsilon_{ijt}. \quad (2.3)$$

There are two types of observable components in the utility function. The upper part component W_{iNt} varies across nests with $N \in \{S, M\}$. The second component V_{ijNt} varies across alternatives j within nest N , and captures city-specific characteristics that are valued by the worker, such as the quality of the local labor market. Finally, ε_{ijt} can be interpreted as the unobserved costs and benefits of choosing destination city j for individual i in a specific period t .²²

21. I drop individual-period observations after a gap in a worker's employment history, as explained in Section 2.2.2.

22. ε_{ijt} is an idiosyncratic shock that is specific to individual i , city j and time t , but it is not captured by

If the worker decides to change job (nest $N = M$), the lower-level observed part of utility is assumed to take the following form

$$V_{ijMt} = \pi_{jt} - \tau_{hj} - \eta_{ij} + \beta \mathbf{Info}_{ijt}. \quad (2.4)$$

In this equation, the term π_{jt} captures public information on each local labor market j at time t . This can be thought of as the sum of publicly accessible information on job opportunities such as online job advertisements, as well as housing costs and other local amenities and disamenities.²³ Since this information is public, it does not vary across workers. In practice, I measure π_{jt} using information on worker i 's expected local wages and on local housing prices. The latter are assumed to capture the value of local amenities as in hedonic pricing models (see Appendix B.3.1 for more details). I thus define $\pi_{jt} = \phi_1 \mathbf{Wage}_{jt} + \phi_2 \mathbf{Amenities}_j$.

The term τ_{hj} represents spatial mobility costs, which are origin-destination specific. Mobility costs capture factors such as physical and cultural distance and are defined as

$$\tau_{hj} = \delta_0 \mathbf{1}\{j \neq h\} + \delta_1 \mathbf{Dist}_{hj}, \quad (2.5)$$

where $\mathbf{1}\{j \neq h\}$ is a dummy for whether the potential destination differs from the current residence, while \mathbf{Dist}_{hj} measures the bilateral distance between worker i 's current location h and an alternative city j . The third element in equation (2.4) is η_{ij} , which identifies the home bias and is defined as

$$\eta_{ij} = \lambda \mathbf{1}\{j \notin B_i\}, \quad (2.6)$$

where $\mathbf{1}\{j \notin B_i\}$ is a dummy equal to 1 if location j is not in the same region where individual i was born B_i .

Lastly, $\beta \mathbf{Info}_{ijt}$ measures the value of private information in the worker's utility. Asymme-

the other components in U_{ijt} . It is assumed to be independent and identically distributed across locations, and independent of all other terms in U_{ijt} .

23. I implicitly control for sector and quarter FEs which are constant within an individual because each worker makes a unique location decision. Similarly, firm FEs cannot be added as the model is already saturated.

tries in information on job opportunities between different individuals, or information frictions, are captured in this term. Info_{ijt} is defined as in equation (2.1), where the probability that a worker hears about one of the v_{fjt} offers at firm f in city j scales with the share of people that they know at that firm, namely s_{ifjt} . In the context of a random search model, the parameter $\beta \in [0, 1]$ in equation (2.4) can be thought of as the joint probability of learning about an opening through professional ties and receiving an offer.

Turning to the upper level of the model, given that the staying nest $N = S$ is a singleton (it only comprises one possible outcome, namely staying in the current job in city h), the upper-level decision is a binary choice, and only relative levels of job change determinants matter. This means that, when the worker does not change job, V_{ijSt} can be normalized to zero without loss of generality. The utility of a worker i who decides to stay in her current job (and city) at time t therefore depends on W_{iSt} , which is specified as

$$W_{iSt} = \alpha_i + \alpha_{kt} + \mathbf{X}'_{it}\beta. \quad (2.7)$$

The utility of staying includes an individual component α_i , which captures time-invariant individual heterogeneity, addressing the fact that some workers are inherently more likely to change jobs than others. The second element α_{kt} measures variation in demand for specific skills (as earlier, it is measured by sector \times quarter, occupation \times quarter, or sector \times occupation \times quarter FE). Lastly, the vector \mathbf{X}'_{it} includes a set of observable individual-level characteristics that influence the utility of staying in the current job: age and its square, whether the current contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

As derived in Section B.3.2 the upper model can then be estimated as a binary logit of the job mobility decision

$$\mathbb{P}_{iMt} = \Lambda\left(\rho_2 I_{iMt} + \alpha_i + \alpha_{kt} + \mathbf{X}'_{it}\beta\right), \quad (2.8)$$

where $\Lambda(x) = \frac{1}{1+\exp(x)}$. The model is estimated separately by skill group, using each skill-specific estimation sample as in Section 2.3.2. In this setup, workers making a job-to-job

transition can choose among 280 possible CZs covering continental France.

2.4.2. Model estimation

Table 2.7 reports the estimated (standardized) coefficients of the location decision in the lower model, separately by skill group. These results show that location-specific information on job opportunities has a positive effect on the probability of choosing a destination. In model (1), a one standard deviation increase in Info_{ijt} leads to a 2.3% higher odds of choosing destination j among low-skilled workers, and to a 2.2% higher odds of choosing destination j among high-skilled workers.

Table 2.7 — Lower model: location choice

Dependent variable	$\mathbb{P}_{ijt M}$					
	Low-skilled			High-skilled		
	(1)	(2)	(3)	(1)	(2)	(3)
Info_{ijt}	0.023 (0.005)		0.048 (0.007)	0.022 (0.006)		0.087 (0.014)
$\mathbb{1}\{j \neq h\}$	-3.456 (0.050)			-3.792 (0.040)		
$\text{Info}_{ijt} \times \mathbb{1}\{j = h\}$		0.114 (0.028)			0.184 (0.043)	
$\text{Info}_{ijt} \times \mathbb{1}\{j \neq h\}$		0.016 (0.007)			0.032 (0.005)	
$\mathbb{1}\{j \in B_i\}$	-1.681 (0.056)	-1.640 (0.047)	-1.638 (0.047)	-1.531 (0.051)	-1.911 (0.039)	-1.910 (0.039)
Dist_{hj}	-2.182 (0.068)	-6.409 (0.158)	-6.496 (0.157)	-1.241 (0.032)	-3.199 (0.071)	-3.266 (0.071)
Wage_{jt}	0.001 (0.008)	0.006 (0.006)	0.007 (0.006)	0.013 (0.004)	0.024 (0.002)	0.025 (0.003)
Amenities_j	0.177 (0.005)	0.207 (0.007)	0.208 (0.007)	0.218 (0.002)	0.270 (0.004)	0.273 (0.003)
Pseudo R^2	0.5966	0.5094	0.5083	0.4788	0.3550	0.3526
Observations	97,868,400	97,868,400	97,868,400	61,689,040	61,689,040	61,689,040

Notes: Conditional logit regressions based on the skill-specific estimation samples. Estimates are expressed in log odds and individual-clustered standard errors are reported in parentheses. In this table, the variables Info_{ijt} , Wage_{jt} , Dist_{hj} , and Amenities_j are standardized to ease comparisons between their effects.

In model (1) the estimated coefficients of Info_{ijt} for low- and high-skilled workers are not statistically different, which may appear to contradict the stylized facts discussed in Section

2.3.2. Models (2) and (3) help to understand this, by studying the interaction between \mathbf{Info}_{ijt} and the migration cost ($\mathbb{1}\{j \neq h\}$), which in model (1) has a large negative effect on location decisions. The key driver of this apparent inconsistency is thus the decision to stay in the same CZ. In column (2), I estimate the coefficient on \mathbf{Info}_{ijt} separately for destinations corresponding to the current residence and for those involving mobility to a different city. This highlights that information on local job opportunities is especially important for the decision to stay in the current residence versus that of moving, and especially among high-skilled workers. By contrast, low-skilled workers stay in their CZ for reasons that are relatively less correlated with economic opportunities compared to high-skilled workers. This is consistent with the reduced-form results reported in Section 2.3.2. A similar intuition emerges when eliminating the migration cost indicator variable in model (3), resulting in an almost doubling of the coefficient on \mathbf{Info}_{ijt} among high-skilled workers, compared to the low-skilled.²⁴

Looking at other mobility costs, distance negatively impacts location choice,²⁵ in line with the broad literature that has estimated migration gravity regressions (Beine et al., 2015).²⁶ In addition, the home bias coefficient in Table 2.7 highlights very large fixed costs of moving to a different city, especially if that city is in a different region from the one where the worker was born. It is interesting to note that, while in model (1) the home-bias seems to play a larger role in location decisions of low-skilled workers, the migration cost dummy appears more relevant for high-skilled workers. This reflects the fact that high-skilled workers are generally much more mobile: in this sample, only 19.6% of high-skilled workers still live in the same region where they were born, compared to a share of 37.9% among low-skilled workers.

Examining public information on local labor markets, expected wages have no effect on the location decisions of low-skilled workers, while they positively impact those of the

24. In additional results reported in Table B.21 of the Appendix, I interact \mathbf{Info}_{ijt} with distance. This highlights that information on local job opportunities is less relevant in distant markets where it is more costly to move. The value of information decays over space and the rate of decay is comparable across skill groups.

25. The coefficient implies that a one standard deviation increase in the distance from the city of residence (corresponding to 200km) is associated with an 88.7% decrease in the odds of choosing that location for low-skilled workers, who tend to do a more concentrated job search, and a 71.1% decrease for high-skilled workers. In the rest of the analysis, distance is expressed in km and is defined to be zero when $j = h$. Only results reported in Table 2.7 are based on standardized variables to ease comparisons between their effects.

26. Table B.21 of the Appendix also reveals that moving costs increase with age, by interacting distance (or fixed migration costs) with age-group dummies, a similar result as in Schluter and Willems (2019).

high-skilled. At the same time, among high-skilled workers, the role of expected wages is smaller than that of private information on local job opportunities. Lastly, local amenities positively impact location decisions, especially among high-skilled workers.

I use the results of model (1) to calculate the inclusive value for each individual and period and estimate the probability of making a job-to-job transition in the upper model. Table 2.8 displays the results, using alternative models that mirror those seen in Section 2.3.2, now taking into account the lower-level destination choice through inclusive utility. Models (1) and (2) are based on the binomial logit estimator, first adding CZ and sector \times quarter FEs, and then replacing the latter by occupation \times quarter FEs. Models (3) and (4) are estimated by LPM, adding individual FEs.

Table 2.8 — Upper model: job-to-job transitions

Dependent variable	P_{iMt}							
	Low-skilled				High-skilled			
	Logit		LPM		Logit		LPM	
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Inclusive value	0.234 (0.022)	0.185 (0.026)	0.038 (0.004)	0.039 (0.004)	0.329 (0.027)	0.337 (0.030)	0.048 (0.003)	0.049 (0.003)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Individual FEs			✓	✓			✓	✓
CZ FEs	✓	✓			✓	✓		
Sector \times t FEs	✓		✓		✓		✓	
Occ \times t FEs		✓		✓		✓		✓
R^2	0.1207	0.1128	0.1004	0.0847	0.1069	0.1013	0.0770	0.0649
Observations	3,405,396	3,405,396	3,405,396	3,405,396	2,225,897	2,225,897	2,225,897	2,225,897

Notes: Regression results based on the skill-specific estimation samples. The inclusive value is calculated based on model (1) in Table 2.7. Estimates of the Logit models are expressed as log odds and allow for a structural interpretation of the coefficient ρ_2 from equation (2.8). Individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage. The R^2 reports the adjusted R^2 for LPM models and the Pseudo R^2 for Logit models.

The structural interpretation of the coefficient ρ_2 on inclusive utility from equation (2.8) is confined to models (1) and (2) based on the binomial logit estimator. The theoretical requirement that $\rho_2 \leq 1$ is therefore satisfied, ensuring that the assumed tree structure of the location choice is consistent with utility maximization. At the same time, Table 2.8 highlights that the estimated ρ_2 is higher among high-skilled workers compared to low-skilled individuals. This suggests that information on the possibility of choosing other jobs in one of the destination cities is more relevant for the job mobility decisions of high-skilled workers.

2.4.3. Quantification exercises

The model enables to assess the determinants of the skill mobility gap and the resulting wage inequality between high-skilled and low-skilled workers. To identify the migration gap, I start from the conditional probability of choosing a given destination in the job mobility nest ($\mathbb{P}_{ij|M}$, see Appendix B.3.2 for more details), which is estimated in the lower-level model separately by skill group. I aggregate the probabilities of choosing non-local destinations, to obtain the individual migration probability

$$\tilde{\mathbb{P}}_{i|M} = \sum_{j \neq h} \mathbb{P}_{ij|M}. \quad (2.9)$$

The migration gap is then obtained by averaging $\tilde{\mathbb{P}}_{i|M}$ across workers within each skill-group, and obtaining the difference between high- and low-skilled workers.²⁷ $\tilde{\mathbb{P}}_{i|M}$ depends on all location-specific $\mathbb{P}_{ij|M}$, which are a function of the underlying components of utility in each city j . Changing these underlying components impacts the probability of choosing each destination, shifting the migration probability and the resulting migration gap. In the following quantification exercises, I compare the relative value of each component of local utility by separately shutting down each channel and examining its effect on the expected migration rate relative to the baseline. The gap between this counterfactual migration rate and the baseline reflects the contribution of each component of utility to the migration decision.

The second moment of interest is wage inequality. I first compute a worker's expected wage, corresponding to the sum of average wages paid in each location weighted by the predicted probability of locating in each city

$$\text{Wage}_i = \mathbb{P}_{ij|M} * \text{Wage}_j, \quad (2.10)$$

where average wages in each destination (Wage_j) are skill-group-specific. A worker's expected

27. The t subscript is suppressed since workers make a single job-mobility decision in this setup.

wage (Wage_i) depends on location probabilities ($\mathbb{P}_{ij|M}$), which are a function of the underlying components of utility in each city j , and therefore react to changes in their relative value. I average workers' expected wages (Wage_i) within each skill group and calculate the wage gap as the difference in average expected wages between the two skill groups.

Quantifying the migration gap

Table 2.9 displays in the first two columns the percentage change in the predicted migration rate induced by each component of utility for either low-skilled or high-skilled workers. The third column reports the migration gap under each counterfactual, which can be compared to the baseline migration gap of 8.4 percentage points. The last column summarizes the percentage change in the migration gap that is induced by each component of the utility function.

Table 2.9 — Effect of each component of utility on the migration gap

	% change in migration relative to baseline		Implied migration gap (percentage points)	% change in migration gap relative to baseline
	Low-skilled	High-skilled		
Info_{ijt}	-1.7%	-2.7%	8.9	-5.6%
Wage_{jt}	0.0%	-0.2%	8.4	-0.7%
Amenities_j	-2.9%	-10.7%	11.9	-29.2%
Dist_{hj}	-66.9%	-52.9%	-6.9	220.9%
$\mathbb{1}\{j \neq h\}$	-70.3%	-62.9%	4.3	93.5%
$\mathbb{1}\{j \notin B_i\}$	-22.4%	-14.9%	6.8	24.0%

Notes: Results are based on estimates from model (1) in Table 2.7. To quantify the contributions of each element of utility to the variation in the migration probability, I sequentially set each element of equation (2.4) to zero, and compare the implied migration probability to the baseline migration probability estimated by model (1) in Table 2.7. The gap between the two reflects the role of each component of utility in the migration decision. The migration gap is expressed in percentage points. The baseline migration gap is 8.4.

These results show that, compared to a counterfactual where Info_{ijt} is set to zero across all cities, heterogeneous individual information on local job opportunities reduces the probability of migration. This results from the fact that the majority of an individual's network is located in the city of residence, which biases location decisions towards staying. The negative effect on migration propensity is stronger for high-skilled workers, who are more reactive to this channel, thus narrowing the migration gap between the two skill groups by 5.6%.

By contrast, wage differentials across cities have no effect on migration decisions of low-skilled workers, since wages tend to be relatively homogeneous across locations and concentrated around the minimum wage. The distribution of expected wages has also a negligible (negative) impact on migration decisions of high-skilled workers, thus leaving the migration gap almost unchanged.

In addition, the spatial variation in local amenities negatively affects the propensity to move, especially among high-skilled workers. Without differentiation in amenities, cities would be more homogeneous, making workers more indifferent between their choice. By contrast, the diversity of amenities across cities implies that workers are only going to consider moving to locations that offer a relative improvement in amenities (Gagné et al., 2022). This reduces the pool of potential destinations, which negatively impacts workers' propensity to move. This effect is especially strong among high-skilled workers because they tend to already live in cities with better than average amenities. As such, the set of alternatives that offer them relative improvements in amenities is smaller. This mechanism thus reduces the migration gap relative to a baseline scenario with no spatial differences in local amenities.

Regarding the role of migration costs, Table 2.9 shows that, without differences in distance between cities, low-skilled workers would move more than high-skilled workers, hence the implied migration gap under this scenario would be negative. Distance makes both groups move less, but especially low-skilled workers, which significantly increases the mobility gap. This is in line with previous evidence in the literature, which finds that migration costs are more binding for low-skilled workers.²⁸ Fixed migration costs also negatively impact migration, especially among low-skilled workers, resulting in a steep increase in the migration gap. A similar but weaker effect is observed for the home bias.

Quantifying wage inequality

The previous results suggest that private information on local job opportunities influences migration decisions. Ultimately, however, the question that matters most is whether this

28. See for example Topel (1986), Bound and Holzer (2000), Wozniak (2010), Moretti (2011), Kennan (2015), Diamond (2016), Balgova (2018), Caldwell and Danieli (2024). Other evidence, such as Schmutz et al. (2021), has shown that migration costs relative to wages are higher for low-skilled workers.

increases wage inequality by distorting migration decisions. This could be driven by the fact that, among low-skilled workers, low private information on local job opportunities reduces their incentives to move to cities that would offer better opportunities and wages.

A first answer to this question is provided by the first row of Table 2.10, which displays the overall effect of Info_{ijt} on the expected wage of low- and high-skilled workers, on the resulting wage inequality, and on the migration gap between the two skill groups. Summing up across all job-to-job transitions, information on local job opportunities increases expected wages for both skill groups by enabling workers to place more weight on locations with better job opportunities. However, since the effect is stronger for high-skilled workers, this raises wage inequality between the two skill groups.

Table 2.10 — Effect of Info_{ijt} on wage inequality

	% change in wage relative to baseline		% change in wage inequality relative to baseline	% change in migration gap relative to baseline
	Low-skilled	High-skilled		
All job-to-job transitions	2.1%	3.3%	4.9%	-5.6%
<i>Transitions by type of destination</i>				
To high-wage cities	3.0%	4.8%	6.3%	17.7%
To low-wage cities	1.0%	0.6%	-3.6%	-0.1%
<i>Transitions by type of origin</i>				
From high-wage cities	3.0%	4.3%	4.7%	-8.2%
From low-wage cities	1.1%	1.0%	0.8%	-1.0%

Notes: Results are based on estimates from model (1) in Table 2.7. To measure the impact of Info_{ijt} on wage inequalities, I compute each worker's expected wage as the weighted sum of average wages paid in each location, and the predicted probability of locating in each city (including the current residence). I compare expected wages when location probabilities are based on the baseline estimation of model (1) in Table 2.7, and when information frictions are set to zero. The gap between the counterfactual expected wage and the baseline expected wage reflects the contribution of Info_{ijt} . Wage inequality is then calculated as the difference in mean wages between high- and low-skilled workers. The first row of this table calculates the effect of Info_{ijt} on wage inequality considering all job mobility decisions, from any origin city to any destination city. The following rows focus on high-wage (destination or origin) cities, defined as cities where the average wage is in the top quartile of the skill-specific city-wage distribution, and on low-wage cities, defined as those where the average wage is in the bottom quartile. The last column of this table summarizes the Info_{ijt} -induced migration gap by type of destination and type of origin.

The second and third rows quantify the effect of private information on wage inequality by the type of destination of job transitions, distinguishing between cities that pay relatively high average wages (above the 75th percentile, calculated separately for each skill group) and

those that offer lower wages (below the 25th percentile). This shows that private information on local job opportunities increases expected wage inequality in high-wage destinations by 6.3%, by more strongly attracting high-skilled workers (the migration gap to high-wage cities increases by 17.7%). Information networks thus allow the high-skilled to better reap the benefits of moving to growing cities. By contrast, low-skilled workers struggle to gain information on local job opportunities in cities that would offer them better wages, which reduces their incentives to move there.

Table 2.10 also highlights that private information has almost no effect on the migration gap to low-wage destinations. Moreover, by enabling workers to target cities that offer relatively better opportunities, even among low-wage destinations, this mechanism leads to higher expected wages. The effect is weaker among high-skilled workers, since they have relatively less to gain from residing in low-wage cities, thus decreasing wage inequality.

The last two rows of Table 2.10 examine the effect of private information on expected wages depending on the origin of job transitions. Info_{ijt} makes both types of workers more likely to remain in high-wage cities, which raises expected wages. Since the effect is stronger for high-skilled workers, this results in higher wage inequality between the two skill groups. Lastly, Info_{ijt} has a limited impact on wage inequality for job transitions originating from low-wage cities, and a similarly small effect on the migration gap. In other words, private information on local job opportunities does not significantly affect the propensity of low-skilled workers to live in low-wage cities, but it reduces their probability of moving to high-wage destinations, as seen above.

To conclude this quantification exercise, I use results of the lower model to obtain the inclusive value for each individual-time observation, and examine the impact of the private information channel on overall job-to-job transitions in the model's upper level. Results reported in Table B.22 of Appendix B.3.3 show that this mechanism has a negligible effect on job-to-job transition probabilities, with a change relative to the baseline that is rounded to 0%. Thus, the private information channel only affects the direction of job transitions, conditional on changing job.

2.4.4. The role of information frictions

The above decomposition has highlighted that private information is a more important determinant in the migration decision of high-skilled workers, and that it more strongly attracts this skill group to high-wage cities. However, this could simply reveal that low-skilled workers have less to gain from moving to different cities, because the distribution of their local opportunities is more spatially homogeneous. In this case, observed differences in migration patterns between high- and low-skilled workers would reflect true discrepancies in local opportunities, mobility costs and location preferences, and would not be caused by information frictions. An alternative explanation is that this different response to private information is also at least partly driven by a weaker network among low-skilled workers that co-locates less well with the spatial distribution of local opportunities.

To more precisely disentangle the contributions of networks and that of local opportunities, I thus examine a counterfactual scenario in which the network of low-skilled workers is simulated to resemble the one observed among high-skilled workers. An intuitive way to understand this exercise is to analyze the expected value of Info_{ihjt}^s over the distribution of workers within each skill group and origin-destination-period cell. I include a skill-group superscript s , and the h subscript to highlight the origin city of each worker i . The expected value of Info_{ihjt}^s is given by

$$\text{Info}_{hjt}^s = \mathbb{E}(\text{Info}_{ihjt}^s) = \text{Cov}(\tilde{\mathbf{s}}_{ihjt}^s; \tilde{\mathbf{v}}_{hjt}^s) + \mathbb{E}(\tilde{\mathbf{s}}_{ihjt}^s)\mathbb{E}(\tilde{\mathbf{v}}_{hjt}^s), \quad (2.11)$$

where I aggregate the key moments over the (individual-specific) set of connected firms in city j (F_{ihjt}^s), and define $\tilde{\mathbf{s}}_{ihjt}^s = \sum_{f \in F_{ihjt}^s} \mathbf{s}_{ifhjt}^s \mathbf{q}_{ifhjt}^s$ and $\tilde{\mathbf{v}}_{hjt}^s = \sum_{f \in F_{ihjt}^s} \mathbf{v}_{fht}^s$. This equation summarizes how the skill-biased nature of local opportunities and networks affects the kind of information on mobility opportunities that the worker receives. The covariance $\text{Cov}(\tilde{\mathbf{s}}_{ihjt}^s; \tilde{\mathbf{v}}_{hjt}^s)$ points to the fact that private information depends on the overlap of local networks and opportunities in each location. The terms $\mathbb{E}(\tilde{\mathbf{s}}_{ihjt}^s)$ and $\mathbb{E}(\tilde{\mathbf{v}}_{hjt}^s)$ describe the role of network size and of the level of opportunities in each location, which both increase Info_{hjt}^s . Their multiplicative form further indicates that networks and opportunities need to co-locate in

space to have a positive impact on Info_{hjt}^s . The distribution of information on local job opportunities can then be collapsed to a degenerate one by calculating Info_{hjt}^s for each skill-group and origin-destination-period cell. I similarly obtain the terms $\text{Cov}(\tilde{\mathbf{s}}_{ihjt}^s; \tilde{\mathbf{v}}_{hjt}^s)$, $\mathbb{E}(\tilde{\mathbf{s}}_{ihjt}^s)$, and $\mathbb{E}(\tilde{\mathbf{v}}_{hjt}^s)$.

To understand the role of information frictions, the relevant counterfactual is one in which the moments $\text{Cov}(\tilde{\mathbf{s}}_{ihjt}^s; \tilde{\mathbf{v}}_{hjt}^s)$ and $\mathbb{E}(\tilde{\mathbf{s}}_{ihjt}^s)$ of low-skilled workers match those of high-skilled workers. This enables to study the location decision of low-skilled workers under a counterfactual network distribution that is better aligned with the spatial distribution of local opportunities, and thus under reduced information frictions. This exercise can be intuitively understood by examining the maps in Figure 2.1 of Section 2.3.2. It corresponds to taking the average low-skilled worker in a given CZ and her average distribution of local opportunities, but examining what would happen if she had a network resembling that of the average high-skilled worker in the same origin CZ, holding the distribution of local opportunities constant.

Results are reported in Table 2.11, which displays the effect of information frictions for low-skilled workers, relative to the counterfactual network spatial distribution. The joint effect of $\text{Cov}(\tilde{\mathbf{s}}_{ihjt}^s; \tilde{\mathbf{v}}_{hjt}^s)$ and $\mathbb{E}(\tilde{\mathbf{s}}_{ihjt}^s)$ is the most important driver of the lower migration probability among low-skilled workers, increasing the migration gap with high-skilled workers by 4.8% relative to the counterfactual. This is entirely driven by the mismatch between the network's spatial distribution and that of local opportunities, and not by the size of the local network itself. City-specific opportunities $\mathbb{E}(\tilde{\mathbf{v}}_{hjt}^s)$ are still important for raising the migration gap with high-skilled workers by 2.0%, but to a lesser extent than $\text{Cov}(\tilde{\mathbf{s}}_{ihjt}^s; \tilde{\mathbf{v}}_{hjt}^s)$.

The last column of Table 2.11 depicts the effect of each component of Info_{hjt}^s on wage inequality. The spatial distribution of local opportunities drives the largest increase in wage inequality, which is 12.7% higher than under the scenario where low-skilled workers faced similar local opportunities as high-skilled workers. The effect of this channel on wage inequality represents a useful benchmark for understanding the role of information frictions.²⁹ Information frictions, and in particular the joint effect of $\text{Cov}(\tilde{\mathbf{s}}_{ihjt}^s; \tilde{\mathbf{v}}_{hjt}^s)$ and $\mathbb{E}(\tilde{\mathbf{s}}_{ihjt}^s)$, increase

29. By contrast, since this is a partial equilibrium model, it is less relevant to assess changes in local job opportunities.

Table 2.11 — Baseline and counterfactual Info_{hjt}^s distribution — low-skilled

Moments of Info_{hjt}^s	% change in migration probability	% change in migration gap	% change in wage inequality
$Cov(\tilde{s}_{ihjt}^s; \tilde{v}_{hjt}^s) \ \& \ \mathbb{E}(\tilde{s}_{ihjt}^s)$	-1.6%	4.8%	12.2%
$Cov(\tilde{s}_{ihjt}^s; \tilde{v}_{hjt}^s)$	-1.4%	4.2%	11.0%
$\mathbb{E}(\tilde{s}_{ihjt}^s)$	0.0%	0.1%	6.3%
$\mathbb{E}(\tilde{v}_{hjt}^s)$	-0.7%	2.0%	12.7%

Notes: The table compares the baseline network distribution for low-skilled workers (i.e. the actual distribution in the data) to the counterfactual distribution, where the underlying moments of private information on local job opportunities are simulated to match those of high-skilled workers. I thus calculate the effect on the migration probability, the migration gap and wage inequality induced by the baseline network distribution, relative to each separate counterfactual. This shows how the distribution of coworker networks puts low-skilled workers at a disadvantage relative to the high-skilled. The baseline migration gap is measured in percentage points and corresponds to 8.8. The baseline wage gap is €12.2.

wage inequality by 12.2% relative to the counterfactual, almost the same size as the impact implied by the spatial distribution of local opportunities.³⁰ This is a sizeable effect and suggests that information frictions matter significantly for the misallocation of workers across space, and for their resulting access to local opportunities and to better wages. The lower response of low-skilled workers to private information on local job opportunities is not only caused by a more spatially homogeneous distribution of job opportunities, which makes migration less salient for these workers. It also reflects their weaker network, which co-locates less well with their spatial distribution of job opportunities, thus giving rise to information frictions that limit the potential wage gains of low-skilled workers.

2.5 — Conclusion

This paper uses a novel empirical strategy to study skill-biased regional migration and the channel of information frictions using matched employer-employee data for France. Exploiting the fact that individuals often learn about jobs through networks, including former coworkers, I am able to isolate changes in a worker’s information about migration opportunities from

30. It is important to note that these results represent an upper bound, since moments are not adjusted for composition effects: part of these effects are related to worker sorting on unobservables, which I do not control in this setting.

other factors driving migration. Results suggest that information on migration opportunities plays an important role in determining location decisions for workers changing job. At the same time, this effect varies considerably across the skill distribution and suggests that information frictions may be harder to overcome for low-skilled workers, who have smaller, less mobile networks, and thus limited access to information on migration opportunities.

Recent evidence has shown that low out-migration responses to local shocks can be explained by the strength of ties, especially in depressed areas (Zabek, 2023). Individuals who live in such areas are more likely to have been born there than to have migrated from other cities. As a result, they are more likely to possess stronger connections in their birth city, and a limited network in other locations. My paper sheds light on the other side of this story, since the spatially concentrated networks of low-skilled workers will be less successful in providing information on job opportunities in other cities. This results in higher information frictions for low-skilled workers, and thus lower out-mobility.

My study also provides an additional piece of evidence to understand labor market polarisation, or the great divergence, whereby initially skilled and typically larger cities have become even more skilled relative to initially less skilled and typically smaller cities (Davis et al., 2020). While networks help high-skilled workers obtain information on job opportunities in distant cities and support their location decisions, low-skilled workers are less likely to gain access to migration opportunities through their network, resulting in higher exposure to negative local labor market shocks. The increased unemployment risk among low-skilled workers may end up exacerbating political, as well as labor market, polarization.

One key limitation of this paper is the static framework used to quantitatively assess the role of the information channel relative to other push and pull factors of migration. Future research could incorporate this mechanism into a dynamic model and also explore how coworker networks are built over time.

B.1 — Data appendix

B.1.1. Data description

I use 2005-2019 data from the matched French employer-employee cross-sectional datasets (INSEE-DADS) to construct an exhaustive quarterly panel. The usual DADS panel is a sample of one 12th of all workers (those born in October of each year). DADS cross sections are instead exhaustive but do not allow to follow workers over time. However, they provide variable values both for the current and previous year for all individuals, hence it is possible to match individuals across years based on these variables. I follow Babet et al. (2022) and construct a full panel of quarterly employee job histories. As in their work, there is a single match for 98% of individuals. DADS contains rich information describing job posts and workers, such as age and gender of the worker, occupation and sector.³¹

One additional advantage of DADS is that it provides detailed information on each job position that the worker has held in each year, including beginning and end dates. This enables to rearrange the data to obtain any desired time frequency. I build my panel to have a quarterly frequency. This choice is motivated by the fact that a quarterly frequency is significantly more precise than a yearly frequency to describe work histories. More fine-grained frequencies, such as a monthly frequency, are possible, but they are not as relevant for describing migration decisions, which I assume to be matured over a longer period of time (3 months in my set up). Moreover, the quarterly frequency of the panel aims to better account for the time lag between receiving an information signal and the final recruitment of the worker. In the rest of the paper, I use a measure of employment equal to the full-time equivalent, based on the number of days worked in each quarter. I consider a full-time worker as an employee that works 90 days in each quarter.

DADS contains information on both municipality of residence and municipality of work, which enables to study mobility patterns precisely. Throughout the rest of my analysis, I

31. Occupations are defined based on a 4-digit INSEE nomenclature of professions (variable PCS). The definition of sector is based on the INSEE French activity nomenclature (NAF rev. 2, 2008). Throughout the analysis, I mainly use 2-digit sector identifiers, also named A88 in INSEE-DADS.

will focus on CZs as main urban-area geographical units. The advantage of using CZs in my setting is that they cover continental France exhaustively and that they capture regions where the majority of workers live and work. There are 280 CZs in continental France (excluding Corsica and oversea territories).³² I exploit the entire 2005-2019 DADS panel to identify transitions precisely and to build my coworker network variables.

B.1.2. Cleaning worker histories

To obtain the quarterly panel, I identify the main job of each worker in each quarter. This is defined as the one with most days or hours worked in the quarter.³³ Since DADS is based on firms' payroll reports, it does not contain information on unemployment. However, for the purpose of my analysis I am not interested in unemployment spells, but only in job-to-job transitions. I define these transitions as changes in employer (firm) between two consecutive quarters. When such job-to-job transitions are also characterized by a change in CZ, transitions are defined as migrations.³⁴ Thus, new jobs that follow a gap in DADS longer than one quarter are not considered in my analysis: workers with such gaps are dropped from my sample, since it is not possible to identify the cause of this gap in the data.

I further restrict the sample by focusing on workers between the ages of 25 and 60. Very young (under 25) workers will not have had enough time in the labor market to develop a network; older (over 60) workers are likely close to retirement. Moreover, I focus on workers that are born in France and in my main analysis I restrict the sample to males. The first restriction aims to obtain a sample of workers for which work histories are most complete.

32. The disadvantage of using CZs is that, in some cases, they do not correspond to cities, but instead cover vast geographical areas that are very sparsely populated. Nevertheless, this geographical unit has been used extensively in the literature studying internal migration patterns, such as by Bilal (2023).

33. For ties, I keep the new job that was started in a given quarter. I drop second-jobs, only keeping the main job in terms of duration in each quarter, and clean the data to ensure continuity of main job during the year. For example, a worker could have two jobs that have the same number of days and hours per quarter. However one of the two jobs is only worked for two quarters in the year and intermittently, while the second is carried out for 3 quarters in the year and without breaks between quarters worked. In this case, the latter will be chosen as the main job.

34. Workers also move to a new CZ through a transition between different plants of the same firm. In the 2010-2015 estimation sample these transitions represent 0.2% of all individual-quarter observations. By contrast, migration events identified as transitions between different firms constitute 3.6% of all individual-quarter observations. I discard within-firm transitions because it is more difficult to consider them as voluntary, since they could be due to internal firm strategies.

For foreign-born workers that move to France at a given point in time, it is not possible to observe their work history prior to moving to France. In addition, foreign-born workers are more likely to move back to their country of origin, thus displaying shorter work histories in DADS.³⁵ Focusing on males has instead the aim of better identifying voluntary moves, since evidence shows that females are more likely move to follow their male partner (Patnaik et al., 2022).

Moreover, I limit the sample to employees that have always worked in the private sector. I thus exclude civil servants and self-employed workers.³⁶ Self-employed workers are not covered by DADS. Civil servants are dropped from my sample because they are much less mobile than private-sector workers. Following Schmutz and Sidibé (2019), I drop workers that have ever had a non-standard type of employment contract, such as apprenticeships or interim contracts, in the period 2010-2015.³⁷ I further focus on workers that have always worked in plants with size of 2-500 employees in the full period of my data (2005-2019, and in particular in the network building phase). This is because in large workplaces it is unlikely that a worker knows, or shares information with, all of their former coworkers (Caldwell and Harmon, 2019).

Finally, I drop workers if they ever have missing sector or hourly wage, likely due to misreporting of their data. INSEE-DADS contains detailed information on the total number of hours worked in each year and at each plant, which can be used to obtain hours worked in each quarter. In addition, the dataset contains information on the total salary before tax earned in each position, as well as the total relative number of days worked, which can be used to calculate the total salary earned in each quarter. The hourly wage is then constructed from information on the quarterly salary and the respective number of hours worked, after winsorizing the quarterly salary at the bottom and top 5th percentiles.

In part of the analysis, I distinguish between high- and low-skilled workers, which I

35. French-born workers may also go abroad, and in that period their work history will not be observed in DADS data. Workers with gaps are thus dropped from my cleaned sample.

36. In 2020, civil servants represented 20% of employees in France, while self-employed workers represented 12% of the French workforce, according to data from the French national statistical agency (INSEE).

37. I also exclude workers who have ever worked from home. This is identified by the variable “CPFD” in INSEE-DADS data, which describes the type of contract, for example distinguishing between full-time and part-time work, as well as working from home (in versions of the data up to 2018). Jobs that entail working from home are dropped as I am interested in workers who go to the office and interact with their colleagues.

identify by managers and manual workers respectively. In INSEE-DADS, manual workers are identified through a 4-digit occupation code (known as PCS code in the dataset) starting with 6, while managers are defined as workers with PCS codes starting with 3. While alternatives to occupation, such as education level, could provide a more precise definition of skill, this data is not available in INSEE-DADS.

B.1.3. Additional sample descriptives

Table B.1 describes the firms and respective plants in INSEE-DADS data (column 1). Workers are spread across more than 6.4 million firms and 8.3 million plants in the period 2010-2015. The average plant has 6.3 employees, though there is substantial variance (the standard deviation is 58.0). On average, each firm is composed of 23.3 plants, and 12.1% of plants are located in the Paris CZ. The second column of Table B.1 shows that roughly a quarter of plants fall within the “network” sample: they have between 2 and 500 employees, on average, throughout my sample period. Most (99%) of the plants that are excluded from my network sample are single-employee plants. There is also less variation in plant size within the network sample: the average plant has around 15.2 employees and the standard deviation is 36.3.

In Table 2.2 of Section 2.2.2 and in Table B.1, density is based on working-age population (aged 15-64) taken from the census at the municipality level, where the breakdown of the population by age and municipality is available at a 5-year frequency. For years in which the population is not available, we use data from the previous non-missing year. The area of CZs is based on INSEE 2020 shapefiles (*“base des zones d’emploi”*).

Tables B.2 and B.3 report 2010-2015 migration flows between some of the largest CZs in France, for low- and high-skilled workers respectively. Paris attracts by far the largest migration flows, but there are also large migration flows between other smaller cities of France (identified in the table by “Rest of France”).

Table B.4 highlights that, on average, migration transitions yield higher wage gains compared to local transitions among high-skilled workers, but not among the low-skilled. At the same time, expected wages for the two groups vary substantially over space, resulting in large differences in wage inequality across cities, as shown in Figure B.1. The largest wage gaps are observed in biggest cities, such as Paris and its metropolitan area.

Table B.1 — Descriptive statistics: firms

	All	Network
Number of firms	6,430,707	1,121,533
Number of plants	8,280,499	1,739,163
Size firm	618.8 (7,295.2)	1,888.4 (12,590.5)
Size plant	6.3 (58.0)	15.2 (36.3)
Number of plants per firm	23.3 (199.3)	67.8 (344.8)
Number of plants per CZ	3.1 (42.9)	7.3 (76.2)
<i>A. Characteristics of workers in plants</i>		
% low-skill	59.8% (34.0)	52.3% (31.8)
% high-skill	35.2% (31.1)	29.1% (25.6)
% ever migrated	36.7% (39.0)	15.6% (18.0)
Hourly wage (€)	11.6 (11.4)	15.6 (14.3)
<i>B. Characteristics of plant locations</i>		
% of plants in Paris	12,1%	14,6%
Density of CZ	1,174.3 (2,501.2)	1,240.6 (2564.7)

Notes: This table presents descriptive statistics on firms and their plants in INSEE-DADS. The first column describes all plants included in INSEE-DADS over the period 2010-2015. The second column restricts the data to the set of plants in the “network” sample, i.e. plants with 2-500 employees. The average number of employees over the sample period is used to define the network sample. The first entry in each row is the mean. Standard deviations are reported in parentheses. Workers are defined as having made a migration if they have changed their employer (firm) and moved to a new CZ between two non-missing consecutive quarters.

Table B.5 describes the transitions of the estimation samples of male workers. In the sample of all workers, 11.7% of all observations represent job-to-job transitions involving a change in firm (this is the sum of 8.7% local transitions and 3.8% migration transitions). For low-skilled workers, 2.9% of job-to-job transitions involve a change in firm and CZ, while this share is 3.7% among high-skilled workers.

Table B.2 — Migration patterns by city: low-skilled workers


Origin CZ	Destination CZ					
	Paris	Lyon	Marseille	Toulouse	Lille	Rest of France
Paris	-	1.2%	1.0%	1.0%	0.8%	95.6%
Lyon	5.3%	-	2.1%	1.2%	0.2%	91.2%
Marseille	4.2%	4.2%	-	1.4%	0.2%	90.0%
Toulouse	5.2%	1.3%	1.7%	-	0.3%	91.5%
Lille	4.0%	1.4%	0.4%	0.6%	-	93.7%
Rest of France	8.8%	2.7%	1.6%	1.4%	1.9%	83.6%

Notes: Summary statistics are based on migration events for the estimation sample of low-skilled male workers over the period 2010-2015 described in Table B.5. Reading of table: among the low-skilled migrations originating from the city of Paris, 1.2% went to Lyon, 1.0% went to Marseille, 1.0% went to Toulouse, etc.

Table B.3 — Migration patterns by city: high-skilled workers

Origin CZ	Destination CZ					
	Paris	Lyon	Marseille	Toulouse	Lille	Rest of France
Paris	-	6.5%	2.5%	2.5%	2.4%	86.1%
Lyon	18.4%	-	2.9%	1.2%	1.0%	76.1%
Marseille	13.2%	6.0%	-	2.8%	0.7%	77.6%
Toulouse	18.2%	4.6%	2.5%	-	0.5%	74.1%
Lille	17.1%	2.7%	0.6%	0.7%	-	78.8%
Rest of France	16.2%	4.8%	2.1%	2.0%	2.5%	72.3%

Notes: Summary statistics are based on migration events for the estimation sample of high-skilled male workers over the period 2010-2015 described in Table B.5. Reading of table: among the high-skilled migrations originating from the city of Paris, 6.5% went to Lyon, 2.5% went to Marseille, 2.5% went to Toulouse, etc.

**Figure B.1 — Wage inequality between high- and low-skilled workers**

Notes: The shaded colors depict the difference in mean wages (expressed in €) between high- and low-skilled workers in each CZ over the period 2010-2015. The data is binned in ten deciles. Data comes from the “clean” sample of male workers described in column 2 of Table 2.1.

Table B.4 — Wage effects of job mobility

	All	Low-skilled	High-skilled
% change in wage after any job transition	6.2%	4.5%	7.6%
	(58.1)	(47.1)	(57.5)
% change in wage after migration transition	5.8%	4.0%	8.3%
	(50.7)	(53.5)	(64.9)
% change in wage after local transition	6.4%	4.7%	7.3%
	(47.1)	(44.2)	(53.1)

Notes: Summary statistics are based on the “clean” sample of male workers over the period 2010-2015 described in column 2 of Table 2.1. The first entry in each row is the mean. Standard deviations are reported in parentheses. I use quarterly wages and measure job-to-job transitions between consecutive quarters. Quarterly wages are winsorized at the bottom and top 5th percentiles.

Table B.5 — Descriptives: transitions in the estimation sample

	All workers		Low-skilled		High-skilled	
	Same CZ	Change in CZ	Same CZ	Change in CZ	Same CZ	Change in CZ
<i>A. Number of observations</i>						
Job-to-job transition	537,967		276,713		162,346	
Change of firm	502,177	219,402	259,449	100,299	142,769	82,923
Same firm	35,790	14,182	17,264	7,275	19,577	7,932
No transition	53,84,211		3,021,109		1,972,696	
Total observations	6,155,762		3,405,396		2,225,897	
<i>B. as % of Number of observations</i>						
Job-to-job transition	8.7%	3.8%	8.1%	3.2%	7.3%	4.1%
Change of firm	8.2%	3.6%	7.6%	2.9%	6.4%	3.7%
Same firm	0.6%	0.2%	0.5%	0.2%	0.9%	0.4%
Number of individuals	721,579		349,530		220,318	

Notes: Transitions are calculated for the three main estimation samples described in Section 2.2.2, using individuals with full quarterly job histories, and dropping observations after the first job-to-job transition.

B.1.4. Additional descriptives of coworker networks

There are good reasons to focus on coworkers to measure social networks. Firstly, these networks are precisely identified in administrative datasets and can be observed dynamically.³⁸ Moreover, survey evidence suggests that work-related contacts are critical for finding jobs. For example, according to the French Labor Force Survey (Table B.6), 56.6% of workers

38. Mobility costs and frictions could also depend on non-professional social networks, for example friends and family (see Munshi, 2020 for a review). However, these types of social networks are not measured in my data.

actively looking for a job through their social network³⁹ have used their professional contacts, and this share is very similar among low-skilled workers (58.9%) and among the high-skilled (59.9%).⁴⁰ An equally important role of social networks is reflected in the arrival rate of offers: 67.7% of respondents that received offers had looked for jobs through their social networks, and out of the latter, 56.8% used their professional network. Professional contacts are also more likely to possess good knowledge of the specific abilities of a given worker and to be aware of potential job openings than, for example, neighbours, friends or family members who may lack the attachment to the relevant labor market segment (Antoninis, 2006). These properties should make coworkers particularly valuable social contacts when looking for a new job.

These descriptives reveal a very similar role of networks for the two skill groups and only a slightly stronger relevance of professional networks for low-skilled workers when it comes to receiving offers. This is important for my identification strategy (detailed in Section 2.3.1), since I compare the role of information on job offers obtained through coworker networks among low- and high-skilled workers.

Table B.6 — Social networks and job search

	All	Low-skilled	High-skilled
Search through social networks	66.9%	68.2%	68.3%
of which professional network	56.6%	58.9%	59.9%
Offer through social networks	67.7%	67.9%	71.7%
of which professional network	56.8%	61.5%	53.9%

Notes: Summary statistics are based on the French Labour Force Survey for the year 2015, using the appropriate survey weights.

Figure B.2 describes in a stylized example how two workers in the same firm may have unique coworker networks, due to different mobility and tenure patterns.⁴¹ Firstly, networks

39. In this survey, social networks comprise family, friends and professional contacts. More specifically, research through family and friends is identified through the variable MRC, while research through professional contacts is identified through the variable MRD.

40. In the French labor Force Survey, the skill group of individuals is identified by the variables QPRCR and QPRCU. Low-skilled individuals are those with value 1 (manual or specialized worker), while high-skilled workers are identified by values 6 and 7 (managers).

41. I focus on first-degree connections to identify tighter links and therefore a stronger information pass-through, and discard second-degree and higher-order network links. A recent paper focusing on higher-order networks and mobility in a developing country setting is Blumenstock et al. (2023).

may vary between workers who previously worked at different firms (or plants within the same firm) and thus developed different networks. Secondly, networks may vary between workers who joined the same plant at different times, due to when precisely *other* former colleagues joined or left the plant. Figure B.3 shows the timeline of coworker network formation, the arrival of the information shock, and the potential decision to change job. The coworker network is formed before time $t = 0$, the information shock arrives between period $t = 0$ and $t = 1$, and may be followed by a job transition.⁴²

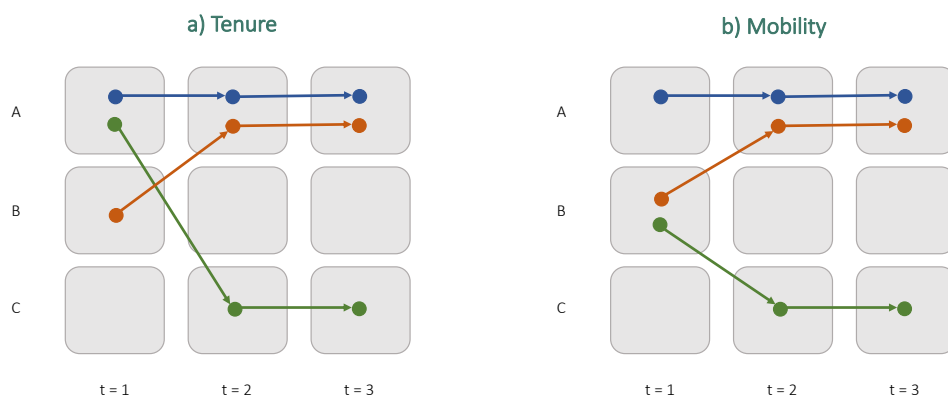


Figure B.2— Variation in networks: tenure and mobility

Notes: Coworker networks can vary between workers in the same plant both due to their history at other firms and due to differences in tenure at their current workplace. Panel A shows how networks vary between workers in the same workplace due to differences in tenure. It shows an example where the incumbent (blue) worker has better information than a new worker (orange). In the first period, the blue worker works with the green worker at firm A; the orange worker is alone at firm B. In period two, the orange worker moves to firm A and the green worker moves to firm C. In the third period the blue worker's coworker network will include the green worker (firm C), but the orange worker's network will not. Panel B shows how networks vary between workers in the same plant due to differences in mobility histories. If we compare the blue and the orange worker in the third period, we see that the latter has a larger network (composed of the green worker), since she was more mobile in the past.

Table B.7 describes autocorrelation patterns of network characteristics and information on local job opportunities. It shows that the number of connections in a worker's coworker network are highly autocorrelated, with autocorrelations above 0.8 after one year. Other characteristics of networks are also relatively stable over time. By contrast, the number of new positions at connected plants is significantly less autocorrelated.

42. The quarterly frequency of the panel enables to account for the time lag between receiving an information signal and the final recruitment of the worker.




Figure B.3 — Timeline: coworker network formation and job transitions

Notes: For each quarter, I use data from the previous 12 quarters (3 years) to construct the coworker network. I use changes in employment from last quarter (period 0) to the current quarter (period 1) to construct the information shock. I look at job transitions from period 0 to period 1.

Table B.7 — Autocorrelation of network characteristics

	Lag						
	1	2	3	4	5	6	7
Number of connections	0.9567	0.9132	0.8694	0.8268	0.7859	0.7454	0.7060
Non-local connections	0.9597	0.9167	0.8747	0.8331	0.7915	0.7519	0.7163
Change in demand at connected plants	0.2143	0.1235	0.1475	0.2022	0.1380	0.0883	0.1105
<i>A. Characteristics of network</i>							
Number of plants	0.9704	0.9369	0.9016	0.8664	0.8322	0.7999	0.7708
Number of CZs	0.9545	0.9117	0.8719	0.8350	0.7997	0.7663	0.7357
Number of sectors	0.9806	0.9548	0.9304	0.9082	0.8871	0.8681	0.8513
<i>B. Characteristics of connected plants</i>							
Size	0.8848	0.7794	0.6925	0.6191	0.5579	0.5038	0.4480
Hourly earnings	0.9173	0.8653	0.8205	0.7854	0.7559	0.7289	0.6945
% blue collar	0.9459	0.8972	0.8552	0.8192	0.7899	0.7640	0.7400
% in non-home CZ	0.9218	0.8496	0.7852	0.7272	0.6744	0.6248	0.5778

Notes: Change in demand at connected firms (v_{fjt}) is calculated as a mean across all connected firms (both local and based in other CZs) in each quarter. Panel (A) and (B) are also based on both local and non-local network connections. Summary statistics are calculated from the estimation sample of all-skilled workers, whose transitions are described in Table B.5.

Table B.8 reports skill-specific summary statistics on local opportunities, networks, and network-weighted opportunities across CZs, calculated in the estimation samples of high- and low-skilled workers. Network-weighted opportunities capture the subset of migration opportunities that workers are more likely to hear about through their network measured in equation (2.1). As seen in Figure 2.1, the standard deviation of local opportunities across CZs is significantly larger for high-skilled workers (Panel A of Table B.8). A similar story is also captured by Moran's I, a measure of spatial autocorrelation, which suggests that

local opportunities are more than twice less spatially concentrated for high-skilled workers, compared to low-skilled workers.⁴³

Table B.8 — Summary statistics of networks by CZ

	Low-skilled	High-skilled
<i>A. Local opportunities by CZ</i>		
Mean	54,629.1	25,023.1
Median	31,286.5	5,314.5
Std. Dev.	97,992.0	154,723.0
Moran's I (weight = x^2)	0.07	0.03
Moran's I (weight = \exp^2)	0.23	0.05
<i>B. Network by CZ</i>		
Mean	0.01	0.02
Median	0.01	0.01
Std. Dev.	0.02	0.11
Moran's I (weight = x^2)	0.07	0.03
Moran's I (weight = \exp^2)	0.24	0.05
<i>C. Network-weighted local opportunities by CZ</i>		
Mean	0.47	0.75
Median	0.14	0.11
Std. Dev.	1.56	5.18
Moran's I (weight = x^2)	0.08	0.03
Moran's I (weight = \exp^2)	0.29	0.04

Notes: Local opportunities in Panel A are the sum of job opportunities in each CZ-skill-quarter cell, averaged over the period 2010-2015. Summary statistics in panel B and C are calculated for the two skill-specific samples of male workers over the period 2010-2015, described in Table B.5. I take the average of their network size or of network-weighted local opportunities in each CZ over the period 2010-2015. In these summary statistics, the network includes coworkers in the home CZ.

At the same time, the coworker network tends to be more spatially concentrated for low-skilled workers, and this is mostly driven by its large presence in the city of residence (Panel B of Table B.8). By contrast, high-skilled workers display a significantly larger standard deviation across CZs in their network and a much smaller Moran's I. Panel C of Table B.8 combines the spatial distribution of networks and that of opportunities by examining characteristics of the distribution across CZs of network-weighted local opportunities, as measured in equation (2.1). This reveals that the network distribution across cities enhances the spatial autocorrelation of job opportunities for low-skilled workers. By contrast, the

43. Results are robust to using alternative weights to calculate the Moran's I, such as the second power and the second exponential.

spatial autocorrelation of network-weighted opportunities is lower for high-skilled workers since they have a more diverse network and a more varied set of opportunities.

B.1.5. Future networks

In Section B.2.2 we describe robustness tests based on future coworker networks, to show that they have a limited impact on workers' job mobility decisions.

An individual's future coworker network is the subset of individuals with whom she will work in either (i) the next year (from tomorrow up until 365 from now), (ii) or in the next two-three years (from 366 days from now), but who are not currently working at her firm, and who therefore she has not yet met. Thus, a worker's future coworker network consists of workers who will join her at her current (or her future) firm, but have, themselves, not moved yet.

Following Caldwell and Harmon (2019), I leave out connections at plants and firms where the worker will be employed in the next three years, otherwise the size of the future coworker network would mechanically vary with job mobility decisions. I also exclude individuals who are in her former coworker network.

B.2 — Empirical appendix

B.2.1. Non-linear models

Table B.9 — Baseline estimates by logit

Dependent variable	Job-to-job transition			Migration transition		
	(1)	(2)	(3)	(4)	(5)	(6)
Mean		11.7			3.7	
Info_{it}	0.043 (0.006)	0.040 (0.005)	0.041 (0.005)	0.018 (0.008)	0.026 (0.008)	0.027 (0.010)
n_{it}	0.095 (0.009)	0.101 (0.008)	0.098 (0.009)	0.109 (0.015)	0.128 (0.015)	0.111 (0.016)
Individual controls	✓	✓	✓	✓	✓	✓
CZ× <i>t</i> FEs	✓			✓		
CZ FEs		✓	✓		✓	✓
Sector× <i>t</i> FEs		✓			✓	
Occ× <i>t</i> FEs			✓			✓
Pseudo <i>R</i> ²	0.1468	0.1498	0.1476	0.1955	0.1685	0.1823
Observations	85,336	85,336	85,336	85,336	85,336	85,336

Notes: Logit regressions estimated on subset of 10,000 workers from the 2010-2015 estimation sample of all male workers. Estimates are expressed in log odds. In the first three columns **Info_{it}** includes information on local ($j = h$) and migration ($j \neq h$) job opportunities, while in the last three columns it comprises only information on migration opportunities. Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage. For comparison, Table B.10 reports the results of the same models estimated by LPM, thus excluding individual FEs.

Table B.10 — Baseline estimates by LPM - no individual FEs

Dependent variable	Job-to-job transition			Migration transition		
	Mean	11.7		3.7		
	(1)	(2)	(3)	(4)	(5)	(6)
Info_{it}	0.981 (0.161)	0.892 (0.152)	0.952 (0.156)	0.074 (0.065)	0.104 (0.059)	0.117 (0.059)
n_{it}	1.185 (0.139)	1.284 (0.138)	1.237 (0.140)	0.421 (0.076)	0.449 (0.075)	0.375 (0.075)
Individual controls	✓	✓	✓	✓	✓	✓
CZ×t FEs	✓			✓		
CZ FEs		✓	✓		✓	✓
Sector×t FEs		✓			✓	
Occ×t FEs			✓			✓
Pseudo R^2	0.2138	0.1647	0.1558	0.0725	0.0611	0.0659
Observations	85,336	85,336	85,336	85,336	85,336	85,336

Notes: LPM regressions estimated on subset of 10,000 workers from the 2010-2015 estimation sample of all male workers. Estimates are expressed in percentage points. In the first three columns **Info_{it}** includes information on local ($j = h$) and migration ($j \neq h$) job opportunities, while in the last three columns it comprises only information on migration opportunities. Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

B.2.2. Robustness tests

This Section summarizes results of robustness specifications, providing further evidence that results are driven by changes in workers' information about their migration opportunities.

Local information — One robustness test is to examine the role of information on *local* job opportunities on migration decisions, which should play a minor role relative to the more relevant information set on migration opportunities. I define local information on job opportunities as $\mathbf{Info}_{it}^H = \sum_{j=h} \mathbf{Info}_{ijt}$, thus comprising connected firms in the home location of worker i . The first three columns of Table B.12 show that adding \mathbf{Info}_{it}^H to the previously estimated models has a weak effect on the migration probability.⁴⁴ This small but positive effect can be explained by the fact that the decision to change job is driven by overall information on job opportunities (both local and non-local). Thus, information on local job

44. The mirror pattern is observed when analysing local job-to-job transitions, which are mainly driven by information on local jobs, although information on jobs in other cities still has a minor effect (in the last three columns of Table B.12).

Table B.11 — Migration decisions by skill: Logit regressions

Dependent variable	Migration transition					
	Low-skilled			High-skilled		
Mean	2.9			3.7		
	(1)	(2)	(3)	(1)	(2)	(3)
Info_{it}^M	0.018 (0.007)	0.026 (0.007)	0.033 (0.009)	0.030 (0.015)	0.025 (0.007)	0.029 (0.005)
\mathbf{n}_{it}	0.129 (0.018)	0.141 (0.016)	0.120 (0.015)	0.131 (0.014)	0.152 (0.015)	0.153 (0.015)
Individual controls	✓	✓	✓	✓	✓	✓
CZ-t FEs	✓			✓		
CZ FEs		✓	✓		✓	✓
Sector-t FEs		✓			✓	
Occ-t FEs			✓			✓
Pseudo R-squared	0.1658	0.1427	0.1494	0.1576	0.1244	0.1309
Observations	97,180	97,180	97,180	102,043	102,043	102,043

Notes: Logit regressions estimated on two subsets of 10,000 workers obtained from the estimation sample for each skill group. Estimates are expressed as log odds. Info_{it}^M identifies information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardized and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

opportunities has a complementary effect, rather than a substitution effect, on migration decisions. The quantitative model developed in Section 2.4 accounts for this effect explicitly, by positing a nested logit structure where total information on job opportunities affects the decision to change job, while the location decision takes place as a second stage of the model.

Within-firm variation — One worry is that sector \times quarter or occupation \times quarter (or even sector \times occupation \times quarter) FEs may not suffice to control for unobserved changes in labor demand. In particular, there could be unobserved shocks that are specific to the workers of a given firm. For example, firm policies that decrease job benefits could increase both the mobility of a coworker and that of her network. One solution is to exploit within-firm variation by adding firm FEs to the baseline model.

Firm FEs may also provide an alternative to control for unobserved changes in demand, as there may be substantial unobserved variation in workers' skills even within detailed occupation groups. For instance, software engineers may differ in their knowledge of Python, and their varying skills may be differently valued by firms. In some cases, skill variation

Table B.12 — Local and non-local information

Dependent variable	Migration transition			Local transition		
	(1)	(2)	(3)	(4)	(5)	(6)
Mean		3.6			8.2	
\mathbf{Info}_{it}^M	0.063 (0.010)	0.061 (0.010)	0.060 (0.010)	0.042 (0.010)	0.041 (0.010)	0.039 (0.010)
\mathbf{Info}_{it}^H	0.018 (0.008)	0.017 (0.008)	0.025 (0.009)	0.563 (0.064)	0.583 (0.066)	0.519 (0.059)
\mathbf{n}_{it}	0.762 (0.022)	0.761 (0.022)	0.766 (0.023)	1.988 (0.032)	1.154 (0.032)	1.071 (0.033)
Individual controls	✓	✓	✓	✓	✓	✓
Individual FEs	✓	✓	✓	✓	✓	✓
Sector \times t FEs	✓			✓		
Occ \times t FEs		✓			✓	
Sector \times Occ \times t FEs			✓			✓
Adjusted R^2	0.1320	0.1327	0.1442	0.1973	0.1831	0.2278
Observations	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762

Notes: LPM regressions based on the 2010-2015 estimation sample of all male workers. \mathbf{Info}_{it}^M identifies information on migration opportunities (i.e. in CZs such that $j \neq h$), while \mathbf{Info}_{it}^H identifies information on local ($j = h$) opportunities. Estimates are expressed in percentage points. Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

within sectors or occupations may be the result of on-the-job training received in specific firms. As a result, individuals with a shared work history may have skills that are similar in ways that are not observed.

Moreover, adding firm FEs can limit endogeneity issues caused by the fact that higher ability workers tend to be hired by better firms, especially in larger cities (Dauth et al., 2022), and that better firms could be hiring more on average. The combination of these two factors could increase the potential mobility opportunities for high-ability workers, making the hiring patterns of connected plants non-random to (unobserved) individual ability. While this may be an important concern, the precise timing of connected firms' hiring decisions and the specific magnitude of these expansions are likely to be random to the individual receiving this information through her network. This is corroborated by the descriptive statistics provided in Table B.7 of Appendix B.1.4, which indicate a low quarterly autocorrelation in the number of new positions at connected plants.

In Table B.13 I exploit within-firm variation to control for these factors. Compared to

model (1) of Table 2.4, I sequentially add firm \times quarter FEs in model (1), firm FEs and sector \times quarter FEs in model (2) or occupation \times quarter FEs in model (3), firm \times sector \times quarter FEs in model (4), and firm \times occupation \times quarter FEs in model (5). Results are consistent with the baseline estimates in Table 2.4: exposure to information on migration opportunities has a positive and statistically significant effect on the probability to migrate in a given quarter, even when exploiting within-firm variation.

Table B.13 — Within-firm variation

Dependent variable Mean	Migration transition				
	(1)	(2)	(3)	(4)	(5)
Info_{it}^M	0.135 (0.043)	0.063 (0.019)	0.098 (0.038)	0.061 (0.019)	0.129 (0.042)
n_{it}	0.500 (0.073)	0.764 (0.067)	0.440 (0.082)	0.762 (0.068)	0.562 (0.081)
Individual controls	✓	✓	✓	✓	✓
Individual FEs	✓	✓	✓	✓	✓
Firm \times t FEs	✓				
Firm FEs		✓	✓		
Sector \times t FEs		✓			
Occ \times t FEs			✓		
Firm \times Sect \times t FEs				✓	
Firm \times Occ \times t FEs					✓
Adjusted R^2	0.3459	0.0979	0.3746	0.0987	0.3525
Observations	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762

Notes: LPM regressions based on the 2010-2015 estimation sample of all male workers. Estimates are expressed in percentage points. Info_{it}^M identifies information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

Quality of coworker connection — My baseline specification weights information by the quality of the coworker connection, proxied by the time worked together at the same plant. I compare these results with estimates where the quality weight q_{fjt} is set to 1. The results, summarised in Table B.14, display a substantially smaller effect of exposure to information on the migration probability relative to baseline estimates, suggesting that the quality of connections is an important element in determining the strength of information pass-through.

Table B.14 — Unweighted measure of information

Dependent variable	Migration transition		
	3.6		
Mean	(1)	(2)	(3)
Info_{it}^M	0.044 (0.007)	0.042 (0.007)	0.042 (0.008)
\mathbf{n}_{it}	0.765 (0.022)	0.763 (0.022)	0.769 (0.023)
Individual controls	✓	✓	✓
Individual FEs	✓	✓	✓
Sector \times t FEs	✓		
Occ \times t FEs		✓	
Sector \times Occ \times t FEs			✓
Adjusted R^2	0.1633	0.1327	0.1442
Observations	6,155,762	6,155,762	6,155,762

Notes: LPM regressions based on the 2010-2015 estimation sample of all male workers. Estimates are expressed in percentage points. Info_{it}^M identifies information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

Quality of the connected firm — The transition probability could depend on the probability that an offer comes from a sufficiently good firm. One proxy of the firm’s quality is the wage that it pays relative to mean retributions in the same CZ, or in the same CZ and sector.⁴⁵ Using these alternative weights results in a smaller effect of exposure to information on the migration probability (Table B.15), compared to both the baseline results and the results with $\mathbf{q}_{fjt} = 1$ (Table B.14). This suggests that plant quality is a weaker prediction of the strength of information pass-through than the quality of the network connection.

Occupation-specific networks — The effective strength of ties between former colleagues is difficult to measure in my data. However, workers in more similar occupations are more likely to have interacted during their work experience in the same plant, compared to workers in different occupations. As an additional test, I check the robustness of my results to using occupation-specific coworker networks, reducing the coworker network to former colleagues

45. Determining which workplaces are likely to be more attractive is in practice difficult. For example, workers may have preferences over non-wage characteristics (Sorkin, 2018), such as in-work benefits. However, these are not measured in my data.

Table B.15 — Firm quality weights

Dependent variable Mean	Migration transition 3.6		
	Baseline weight (1)	CZ wage weight (2)	CZ-sector wage weight (3)
Info_{it}^M	0.063 (0.010)	0.038 (0.007)	0.039 (0.007)
n_{it}	0.764 (0.022)	0.765 (0.022)	0.765 (0.022)
Individual controls	✓	✓	✓
Individual FEs	✓	✓	✓
Sector \times t FEs	✓	✓	✓
Adjusted R^2	0.1320	0.1320	0.1320
Observations	6,155,762	6,155,762	6,155,762

Notes: LPM regressions based on the 2010-2015 estimation sample of all male workers. Estimates are expressed in percentage points. Info_{it}^M identifies information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

who worked in the same 1-digit occupation group. Results of this analysis, displayed in Table B.16, show a stronger effect of information on the migration probability when this narrower network definition is used. This provides suggestive evidence that former coworkers transmit information on job opportunities, especially when their ties are stronger.

Connections formed at smaller and larger plants — Connections formed at smaller plants could be expected to be tighter than those formed at larger plants, leading to more effective information pass-through when workers met in smaller workplaces. On the other hand, working in larger plants results in a wider and more diverse set of professional contacts, which can provide more accurate information on job motility opportunities. Table B.17 compares the baseline results, estimated on a subset of workers who worked in plants of 2-500 employees, with those that worked in smaller plants (2-200 employees), and with those that worked in larger plants (500 or more employees). Results highlight that information on migration opportunities is more relevant for mobility decisions when it comes from coworkers who met in a larger workplace.

Table B.16 — Occupation-specific networks

Dependent variable	Migration transition		
		3.6	
Mean	(1)	(2)	(3)
Info_{it}^M	0.071 (0.011)	0.067 (0.010)	0.067 (0.010)
n_{it}	0.568 (0.022)	0.572 (0.022)	0.545 (0.022)
Individual controls	✓	✓	✓
Individual FEs	✓	✓	✓
Sector× t FEs	✓		
Occ× t FEs		✓	
Sector×Occ× t FEs			✓
Adjusted R^2	0.1318	0.1326	0.1440
Observations	6,155,762	6,155,762	6,155,762

Notes: LPM regressions based on the 2010-2015 estimation sample of all male workers. Estimates are expressed in percentage points. Info_{it}^M identifies information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

Table B.17 — Plant size

Dependent variable	Migration transition		
	2-200	2-500	500+
Mean	3.6	3.5	3.2
	(1)	(2)	(3)
Info_{it}^M	0.044 (0.010)	0.063 (0.010)	0.235 (0.042)
n_{it}	0.823 (0.023)	0.0.764 (0.022)	1.091 (0.043)
Individual controls	✓	✓	✓
Individual FEs	✓	✓	✓
Sector× t FEs	✓	✓	✓
Adjusted R^2	0.1332	0.1308	0.1465
Observations	4,680,588	6,155,762	2,105,399
N. workers	556,049	721,579	225,949

Notes: LPM regressions based on the 2010-2015 estimation sample of all male workers. Estimates are expressed in percentage points. Info_{it}^M identifies information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

Past and future coworkers — Depending on when they worked together, some coworkers are more or less likely to provide the worker with information. Since workers may lose contact with their former colleagues over time, coworkers that met in the more distant past are likely to be less valuable sources of information. Future coworkers are also less likely to be a source of information in the current period because they have not yet worked together. I thus construct distinct networks comprising individuals the worker worked with (1) 4-5 years ago, (2) 2-3 years ago, and (3) 1 year ago. Similarly, I construct networks of coworkers that the worker will work with in (4) 1 year, and (5) 2-3 years (see Section B.1.5 for more details). Table B.18 presents these results and shows that changes in labor demand in plants of more recent former coworkers matter more than changes in plants of coworkers with whom individuals worked with in the more distant past. In addition, changes in demand at a workers' future coworkers' firms have a negligible effect on mobility decisions.

Table B.18 — Past and future coworkers

Dependent variable Mean	Migration transition				
	3.6				
	Past coworkers			Future coworkers	
	1 year	2-3 years	4-5 years	1 year	2-3 years
	(1)	(2)	(3)	(4)	(5)
Info_{it}^M	0.086 (0.011)	0.054 (0.009)	0.021 (0.006)	0.014 (0.014)	0.012 (0.027)
n_{it}	0.546 (0.017)	0.368 (0.016)	0.099 (0.014)	0.038 (0.024)	0.033 (0.040)
Individual controls	✓	✓	✓	✓	✓
Individual FEs	✓	✓	✓	✓	✓
Sector \times t FEs	✓	✓	✓	✓	✓
Adjusted R^2	0.1321	0.1317	0.1316	0.1331	0.1316
Observations	6,155,762	6,155,762	6,155,762	6,155,762	6,155,762

Notes: LPM regressions based on the 2010-2015 estimation sample of all male workers. Estimates are expressed in percentage points. Info_{it}^M identifies information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

Lagged information — Older information on migration opportunities should be less relevant than newer information for migration decisions. In Table B.19, I present alternative

models based on 1-4 lags of information on migration opportunities, confirming that older information has a smaller impact on migration transitions.

Table B.19 — Lagged information

Dependent variable Mean	Migration transition			
	1 Lag (1)	2 Lags (2)	3 Lags (3)	4 Lags (4)
Info_{it}^M	0,042 0,010	0,019 0,006	0,018 0,007	0,026 0,012
\mathbf{n}_{it}	0,751 0,023	0,699 0,024	0,668 0,026	0,629 0,028
Individual controls	✓	✓	✓	✓
Individual FEs	✓	✓	✓	✓
Sector× <i>t</i> FEs	✓	✓	✓	✓
Adjusted R^2	0.1292	0.1246	0.1137	0.1153
Observations	6,155,762	6,155,762	6,155,762	6,155,762

Notes: LPM regression results based on the 2010-2015 estimation sample of all male workers. Estimates are expressed in percentage points. Info_{it}^M identifies information on migration opportunities (i.e. in CZs such that $j \neq h$). Variables are standardised and individual-clustered standard errors are reported in parentheses. Individual controls include age and its square, whether the contract is fixed-term, whether the job is part-time, tenure in the firm, and log wage.

B.3 — Model appendix

B.3.1. Data construction

Expected local wages — For each CZ and quarter, I calculate worker i 's expected wage as the mean wage in her 2-digit sector, 2-digit occupation and age group, to reflect work experience. I identify 4 age groups: 20-34, 35-44, 45-54, 55-65. Before calculating these averages, wages are winsorized at the bottom and top 5th percentiles.

Local amenities — I quantify the value of local amenities using data on local house prices obtained from French housing transaction data (*Demande de Valeur Foncière*). Hedonic pricing models treat house prices as a sum of individual amenities that cannot be sold separately in the market and estimate the contribution of amenities to the house price. I use house prices to measure (non-labor market) local amenities rather than population density, although the two measures are highly correlated. House prices are more orthogonal than city density to the quality of local labor markets as captured by expected wages (Table B.20). I use housing sales data for the period 2014-2015 for the sample of detached homes and apartments. For each sale, house price per squared meter is obtained by dividing the sale price by the surface of the housing unit (excluding that of other terrains). Data is geo-coded at the municipality level and then averaged within each CZ. I do not have sufficient data points to exploit the time dimension of this dataset, hence I average house prices over 2014-2015.

Table B.20 — Measures of local amenities

	House prices	CZ density	Average wages
House prices	1.000		
CZ density	0.219	1.000	
Average wages	0.381	0.409	1.000

Notes: Correlation coefficients between CZ-level house prices, density, and expected wages. House prices are averaged over the period 2014-2015, density and expected wages reported in this table refer to the year 2015. All correlations are statistically significant at the 1% level.

Mobility costs — To calculate distances between CZs, I use municipality-level population data to calculate the population-weighted centroid of each CZ, and calculate bilateral CZ distances using the resulting centroids as coordinates. The fixed migration cost is a dummy equal to 1 if the potential destination differs from the current residence. The home bias corresponds instead to a dummy equal to 1 if a CZ is in a region that differs from that of birth. As discussed in Section 2.2.2 region of birth is available at the department (“département”) level in INSEE-DADS data.

B.3.2. Empirical implementation of the model

The objective is to quantify the probability P_{ijt} of choosing a destination $j \in M$ between periods t and periods $t + 1$, for a worker i currently living in city h . As in Anderson et al. (1992), Train (2001), and Cameron and Trivedi (2005), I characterize the nested choice through two logit equations. Assuming that ε_{ijt} follows a generalized extreme value distribution, the probability of choosing a destination $j \in M$ can be decomposed into the product of conditional and marginal probabilities: the probability P_{iMt} that nest M is chosen (in the upper model), and the conditional probability $P_{ijt|M}$ that location j is chosen among the potential locations in nest M (the lower model)

$$P_{ijt} = P_{ijt|M}P_{iMt}. \quad (\text{B.1})$$

The conditional probability of choosing a given destination (lower-level model) can be written as

$$P_{ijt|M} = \exp\left(\frac{V_{ijMt}}{\rho_2} - I_{iMt}\right), \quad (\text{B.2})$$

where the inclusive value I_{iMt} links the upper and lower models by bringing information from the lower model into the upper model and is given by

$$I_{iMt} = \ln \sum_{j \in M} \exp(V_{ijMt}/\rho_2). \quad (\text{B.3})$$

The marginal probability of choosing nest M and moving to a new job (upper-level model) takes the following logit form⁴⁶

$$P_{iM} = \frac{1}{1 + \exp\left(\frac{W_{iSt} - \rho_2 I_{iMt}}{\rho_1}\right)}. \quad (\text{B.4})$$

Following Train (2001) and Cameron and Trivedi (2005), I normalise ρ_1 to 1, since ρ_1 and ρ_2 cannot be identified separately.⁴⁷ The model can be estimated by maximum likelihood techniques using information on the variables entering the utility function. In the following, I adopt a backward estimation procedure. I first estimate the lower model with the conditional logit estimator and using information contained in equation (B.2). I then use results of the lower model to compute the inclusive value I_{iMt} for each individual and period in the panel based on equation (B.3). This inclusive value serves to estimate the upper model in equation (B.4) as a binary logit of the job mobility decision, combining the latter with the observed utility from equation (2.7) to estimate

$$P_{iMt} = \Lambda\left(\rho_2 I_{iMt} + \alpha_i + \alpha_{kt} + \mathbf{X}'_{it}\beta\right), \quad (\text{B.5})$$

where $\Lambda(x) = \frac{1}{1 + \exp(x)}$. The model is estimated separately by skill group.

46. In this equation, W_{iSt} makes the current job more attractive, thus reducing the probability of changing job. By contrast, the inclusive utility term $\rho_2 I_{iMt}$ summarizes all the relevant information coming from the possibility of choosing other jobs in one of the destination cities (including the home location h), and thus increases the probability of changing job.

47. An important theoretical requirement is then that $\rho_2 \leq 1$ for the model to be consistent with utility maximization for all possible values of the explanatory variables (Anderson et al., 1992). With $\rho_2 = 1$, the shocks are uncorrelated within a nest, and the model collapses to the standard multinomial logit.

B.3.3. Additional quantification exercises

Table B.21 — Lower model: robustness

Dependent variable	$\mathbb{P}_{ijt M}$					
	Low-skilled			High-skilled		
	(1)	(2)	(3)	(1)	(2)	(3)
Info_{ijt}		0.194 (0.004)	0.020 (0.004)		0.025 (0.008)	0.024 (0.008)
Dist_{hj}	-2.139 (0.066)		-2.131 (0.065)	-1.157 (0.031)		-1.156 (0.031)
Info_{ijt} × Dist_{hj}	-0.010 (0.003)			-0.011 (0.004)		
Dist_{hj} × 1{Age_{it} ∈ [20, 34]}		-1.864 (0.067)			-1.004 (0.034)	
Dist_{hj} × 1{Age_{it} ∈ [35, 44]}		-2.264 (0.096)			-1.238 (0.044)	
Dist_{hj} × 1{Age_{it} ∈ [45, 54]}		-2.488 (0.109)			-1.382 (0.057)	
Dist_{hj} × 1{Age_{it} ∈ [55, 65]}		-2.737 (0.187)			-1.488 (0.102)	
1{j ≠ h}	-3.463 (0.049)	-3.436 (0.050)		-3.960 (0.041)	-3.939 (0.041)	
1{j ≠ h} × 1{Age_{it} ∈ [20, 34]}			-3.126 (0.056)			-3.692 (0.048)
1{j ≠ h} × 1{Age_{it} ∈ [35, 44]}			-3.596 (0.064)			-4.015 (0.054)
1{j ≠ h} × 1{Age_{it} ∈ [45, 54]}			-3.772 (0.069)			-4.249 (0.063)
1{j ≠ h} × 1{Age_{it} ∈ [55, 65]}			-3.931 (0.104)			-4.668 (0.104)
1{j ∈ B_i}	-1.731 (0.057)	-1.743 (0.057)	-1.749 (0.057)	-1.520 (0.056)	-1.523 (0.056)	-1.532 (0.056)
Wage_{jt}	0.013 (0.004)	0.014 (0.005)	0.014 (0.004)	0.021 (0.008)	0.020 (0.009)	0.020 (0.009)
Amenities_j	0.183 (0.004)	0.183 (0.004)	0.183 (0.004)	0.232 (0.002)	0.233 (0.002)	0.234 (0.002)
Pseudo R ²	0.5924	0.5934	0.5936	0.4939	0.4949	0.4952
Observations	97,868,400	97,868,400	97,868,400	61,689,040	61,689,040	61,689,040

Notes: Conditional logit regressions based on the skill-specific estimation samples. Estimates are expressed in log odds and individual-clustered standard errors are reported in parentheses. In this table, the variables **Info_{ijt}**, **Wage_{jt}**, **Dist_{hj}**, and **Amenities_j** are standardized to ease comparisons between their effects.

Table B.22 — Upper model: effect of information frictions on job-to-job mobility

	Low-skilled				High-skilled			
	Logit		LPM		Logit		LPM	
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Baseline	0.115	0.115	0.110	0.110	0.109	0.111	0.105	0.104
$\mathbf{Info}_{ijt} = 0$	0.115	0.115	0.110	0.110	0.109	0.111	0.105	0.104
% change compared to baseline	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

Notes: Results based on estimates of model (1) in Table 2.7. The first row shows the predicted job-to-job mobility under the baseline model estimation for each alternative specification and skill group. The second row shows the predicted job-to-job mobility when the \mathbf{Info}_{ijt} channel is shut down. The % change in job-to-job mobility compared to the baseline is very small and rounded to 0% in all specifications.

Chapter 3

Affordable housing and the labor market: evidence from the city of Copenhagen

With Ismir Mulalic and Jos van Ommeren

This paper investigates the impact of affordable housing policies on long-run labor market outcomes in a quasi-natural experimental setting for Copenhagen. We compare two key affordable housing policies – home purchase subsidies and public rental housing – both directed towards low-income households. We find that subsidized home-buying has a negative effect on labor market participation, compared to subsidized rental, as home-buyers experience large wealth increases, which reduce work incentives. Comparing children of home-buyers and renters in public housing, we find limited effects on their adult labor market participation. However, we show that public housing privatisation can in the long run increase housing inequalities, due to the high inter-generational persistence of home-ownership.

3.1 — Introduction

Increasing geographic income segregation and growing housing costs have put housing support at the forefront of current policy debates (Saiz, 2023). Across the world, government expenditure on affordable housing policies is becoming a significant item in public budgets, with public housing constituting a large share of this expenditure. According to the OECD Affordable Housing Database, public housing accounted for 7% of the total housing stock in OECD countries in 2020, ranging from over 20% in Austria, Denmark, and the Netherlands, to 10% in the US, and 3% in Germany. However, the impact of housing support policies, in particular on long-run labor supply, is still not well understood. Analyzing the effect of these interventions is challenging, since households eligible for housing support tend to be a selected non-random sample with distinct socioeconomic characteristics.

Overcoming these challenges, this paper exploits a quasi-natural experiment – the 1996 privatization of subsidized public housing owned by the city of Copenhagen – to explore the impact of housing support on labor supply. This intervention offered residents receiving rental subsidies in municipality-owned public housing the opportunity to buy their home at below-market prices. While municipality-owned public housing was privatized, public housing associations continued to provide subsidized public housing. This setup thus enables us to compare two key forms of housing support. The first policy, which will be called the treatment, corresponds to a large housing purchase subsidy for residents of municipality-owned public housing. At the same time, individuals renting from public housing associations could not buy their home and thus represent an ideal comparison group.

Understanding the impact of housing support policies is crucial because they can have a long-lasting impact on how individuals behave. Of particular interest is the effect that these different policies have on labor supply, and therefore productivity, and thus also on the ability of recipients to potentially move beyond the need for housing support by becoming financially independent. Empirical evidence is however scarce and ambiguous, and the policy debate on this topic is still ongoing.¹

1. For example, in Denmark at the end of the 1990s the Danish Economic Council advocated for the abolition of rental subsidies. In their latest 2023 report they however shifted their stance and recommended

The first type of housing support that we study are rental subsidies in the form of subsidized rental within public housing.² Despite representing a growing share of total housing stock across many OECD countries, evidence on their impact on labor supply is scarce. One notable exception is van Dijk (2020), who investigates labor market effects of admission to public housing through a lottery scheme in Amsterdam in the short run (one to two years after admission). Her results show that moving into public housing has heterogeneous effects on labor market outcomes, and, in particular, a positive impact for households that move from economically worse to economically better neighborhoods.³

The second type of housing support policy that we study are home purchase subsidies, which are widespread across many countries, for example as tax-deductible mortgage interest payments, or through purchase subsidies for first-time buyers (Schwartz et al., 2006; Sodini et al., 2023). These policies are supported by the idea that homeownership generates positive externalities, such as household and neighborhood stability, greater civic engagement, and increases in social capital (Green and White, 1997; Di Pasquale and Glaeser, 1999; Dietz and Haurin, 2003), especially in low-income areas (Shlay, 2006). In addition, recent evidence looking at the privatization of public housing suggests that home-buyers increase their labor supply in the short run, and that this effect can be driven by higher returns to labor market earnings (Hausman et al., 2022).⁴

One additional justification for public involvement in the home-ownership market is that it provides one of the few ways for low-income households to accumulate wealth (Wainer maintaining public housing (Danish Economic Council, 2023).

2. In Denmark a second form of rental subsidies consists in monetary assistance to low-income households through housing benefits.

3. The literature on housing benefits is instead larger, especially in the US, where, for example Jacob and Ludwig (2012) find that housing assistance vouchers reduce labor force participation. Rental subsidies have also been shown to have negative consequences for the housing market, reducing quantity and quality of housing supply (Moon and Stotsky, 1993; Eriksen and Rosenthal, 2010). Moreover, there is evidence that rent control causes housing misallocation, as tenants and housing are not efficiently matched (Glaeser and Luttmer, 2003).

4. There also exists an older literature which argues that reduced residential mobility among homeowners has a negative effect on employment because it restricts labor mobility (Oswald, 1996), although evidence is mixed. Munch et al. (2006), for example, find that that home-ownership hampers the propensity to move for job reasons but improves the chances of finding local jobs. In the Danish context it is also important to note that public housing residents are much less mobile than those in unregulated rental markets, since waiting lists are known to be extremely long, exceeding 20 years for the most popular public housing estates in Copenhagen (Tsenkova, 2021). Such lengthy waiting list, and the resulting inertia of the public housing market, may end up restricting individual mobility among public housing residents compared to home-buyers.

and Zabel, 2020). Since households do not fully internalize the value of home-ownership as a means for asset accumulation, there is a role for government involvement. These wealth effects can be substantial (Wainer and Zabel, 2020), especially in the last 25 years which saw extraordinary house price increases in most major cities, including Copenhagen.

Rising housing wealth is likely to reduce labor supply, as highlighted in the literature. For example, Henley (2004) reports a reduction in working hours in response to housing wealth gains in Britain over 1992-2001. Using variations in local UK housing prices for identification, Disney and Gathergood (2017) also find a negative relation between housing wealth increases and labor supply.⁵ These results match those of a broader established literature on the wealth effect and labor supply (Joulfaian and Wilhelm, 1994; Inbens et al., 2001; Algan et al., 2003; Brown et al., 2010; Cesarini et al., 2017).

Besides the effect of housing support policies on adults, a related question is their impact on children who grew up under differential subsidized housing arrangements. On this matter, the literature is mostly focused on the US and examines housing vouchers that enabled families to move out of public housing located in high-crime low-employment neighbourhoods, in exchange for housing benefits to settle down in better areas. (Chetty and Hendren, 2018; Chyn, 2018; Haltiwanger et al., 2020; Pollakowski et al., 2022). These studies generally find a positive effect of moving out of public housing projects on adult labor market outcomes, but they cannot disentangle the location effect from that of public housing. Another strand of the literature has examined the effect of home-ownership on children more generally. While some evidence points to increased welfare and economic success among children growing up in owned homes (Green and White, 1997; Haurin et al., 2002; Boehm and Schlottmann, 1999), other results are ambiguous (Mohanty and Raut, 2009; Bourassa et al., 2015).

A typical challenge in studying the effect of subsidized housing policies is that it is hard to find settings with plausible exogenous variation. Households residing in subsidized rental housing, or those eligible for housing purchase subsidies, tend to be a selected sample, with distinct socioeconomic characteristics that impact their labor market outcomes. This is particularly problematic for evaluating long-term labor supply, as the decision to buy a house

5. In addition, Li et al. (2020) employ a regression discontinuity design in house size due to Chinese housing policies to identify the effect of housing wealth on labor supply, showing that higher housing wealth reduces labor market participation.

is typically related to expectations over long-run labor market participation. This implies that homeownership is likely to be endogenous to labor supply decisions, even when home-buyers and renters are observationally identical.

Our unique quasi-natural setting enables to overcome the latter issue. In a difference-in-difference framework with two-way fixed effects, we compare individuals living in public housing owned by the municipality at the time of the policy announcement (the treatment group), with individuals living in public housing owned by public housing associations, who were not offered the possibility to purchase their homes (the control group). This resembles an ideal experiment, in which a large random sample of public housing tenants are suddenly offered to buy their home at subsidized prices, and almost all tenants accept this offer. Crucially, residents could not anticipate the sales policy when they were assigned to municipality housing versus housing offered by public housing associations.

One complication in our setting is that not all residents renting from public housing associations are comparable to those living in municipality housing. While municipality housing is only assigned based on priority lists to individuals at high housing risk (single mothers, individuals with health issues or that have been long-term unemployed, and refugees), public housing associations allocate apartments based on both priority and waiting lists. Waiting lists differ from priority lists by allowing any individual or household, without income limit, to sign up and wait until a subsidized housing unit becomes available. Due to high congestion, however, waiting lists are extremely long, exceeding 20 years for the most popular public housing estates in Copenhagen (Tsenkova, 2021).

Since we do not have direct access to priority lists, the challenge is therefore to identify, among public housing residents, those that were more likely to be on priority lists due to their specific demographic characteristics, and who are observationally identical to individuals living in municipality public housing at the start of the policy. We obtain this control group through one-to-one matching on *ex ante* (before treatment) characteristics that relate to eligibility for priority lists, in particular whether individuals are single parents, their health status, and their employment history. As a result, our identifying assumption is that selection into the type of public housing (municipality versus public housing associations) is uncorrelated with future labour supply decisions, conditional on observed household characteristics.

We study both the intensive and extensive impact of home-ownership on labor supply with a long-term focus, since the large wealth effect implied by this policy is likely to play out in long-run labor supply decisions. We find that home purchase has a negative impact on labor market participation. Compared to renting, becoming a homeowner reduces the probability of working in a given year by 2.8%, while yearly labor market earnings decrease by over 4% after buying the home. This is driven by a reduction in the probability of working full-time and not by a fall in hourly wages.

A key explanation for falling labor supply may be the wealth effect experienced by home-buyers, who, compared to their peers in subsidized rental, could be described as having won a lottery. Between 1996 and 2016, house prices in Copenhagen rose by more than 300% on average. On top of this, when buying their home the treatment group received a significant discount of 33% to 72% relative to 1996 house prices. Such a large wealth effect could by itself more than explain the reduced incentives to work. While we cannot track precisely whether home-buyers sold their apartments to cash in this wealth, we find that they were more likely to move to other apartments later on, compared to public housing residents who are subject to high mobility frictions. These effects are mostly driven by males, while females appear to react very little to the treatment. This can be explained by the fact that females in our sample are more than 10 times more likely to be single parents, making it very complicated for them to reduce labor supply.⁶

We then examine the intergenerational effect of subsidized home purchase versus subsidized rental, by comparing children of the treatment and the control group. Children were on average 7-8 years old at the time of the policy, and we compare outcomes once they reach the age of 30. We find no effect on labor supply, both at the extensive and intensive margins. By contrast, we observe a large impact on the intergenerational transmission of home status. Children of home-buyers are 46% more likely to reside in the home of their parents, who in

6. In addition, home purchase and long-run labor supply decisions interact with the take-up of public benefits. As home-buyers reduce long-run labor supply, they are no longer eligible for unemployment benefits, which in the Danish “flex-security” system require availability to work with a one-day notice. Thus, despite reducing their labor supply, home-buyers are less likely to receive unemployment benefits, while they increase their dependence on other social benefits such as cash assistance. By contrast, due to the high Danish retirement age, there is almost no effect on anticipated retirement. Overall, relative to renters, home-buyers see an increase in total public benefits by almost 4% per year on average, which could also explain falling labor supply by reducing work incentives.

the meantime are more likely than renters to move out. At the age of 30, they have not yet inherited the home of their parents, but this is likely to take place in later years (which we cannot yet observe in our data). This translates into a higher probability of living closer to the city centre and in better neighbourhoods in the long run, while children of renters are more likely to move to Copenhagen's suburbs.

Overall, our results suggest that public housing privatisation can in the long run increase housing inequalities due to the high inter-generational persistence of home-ownership. Moreover, the reduction in the public housing stock through privatisation seems at odds with the need to respond to growing house prices and income segregation in urban centres. At the same time, current regulations of the public housing system could be improved. An important step would be the enforcement of an income ceiling to access these services. This could reduce congestion and increase mobility across public housing units for renters, improving matching of housing to renters, with possible spillovers on labor supply decisions.

The remainder of the paper is organized as follows. Section 3.2 discusses our empirical setting. We first explain the quasi-natural experiment and then describe the data and provide descriptive statistics. In Section 3.3, we outline our estimation strategy and present the empirical results. In Section 3.4 we examine the intergenerational effect of the housing support policies. Section 3.5 concludes.

3.2 — Empirical setup and data

This Section outlines our empirical setup and the data used for analysis. Section 3.2.1 describes the quasi-natural experiment, while Section 3.2.2 provides a description of the data.

3.2.1. The quasi-natural experiment

In the early 1990s, the municipality of Copenhagen faced a large budget deficit. The city had at the time a higher concentration of pensioners, unemployed individuals, singles, and socially challenged individuals than any other municipality in the country. This strained the municipality's resources, while tax revenues failed to grow at the same rate as the rest of

Denmark.

In 1996, the municipality of Copenhagen decided to unexpectedly sell all its public housing units, which at the time were more than 15,000. These apartments represented 22% of overall public housing, while the remaining 78% of the public housing stock was provided by non-profit housing associations. Residents were offered the option to buy these municipality housing units, and by 2001, almost all (90%) of them were sold to incumbent residents as cooperatives, a form of shared private ownership (Figure 3.1).⁷

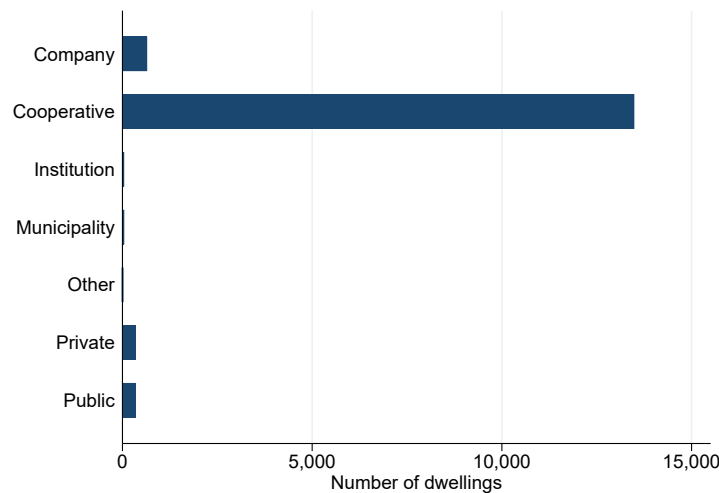


Figure 3.1 — Municipality housing units by ownership type in 2001

Notes: Number of municipality housing units by type of ownership in 2001. As seen in the Figure, 88% of the dwellings were sold to cooperatives by 2001.

The purchase of these dwellings was highly subsidized and a number of policies were put in place to help residents secure loans at very low cost. To ensure that the municipality received fair value for the properties and that residents did not overpay, homes were sold based on public assessment values. These valuations were significantly below market prices for similar housing units. When comparing 1996 assessment prices (per squared meter) to the market price index at the neighbourhood level, estimates of the discount range between 33% and 72%, with a mean of 61% (column 1 of Table 3.1). Relative to 1996 earnings, these

7. In Denmark, cooperative housing is called *andelsbolig*. It is possible to acquire an *andel* (a share in a cooperative, tied to a dwelling) that gives the exclusive right to use the respective housing unit. The full ownership of a housing complex belongs to a cooperative called *andelsboligforening* (housing cooperative association). Appendix C.1 describes in more detail the sale procedure.

discounts were considerable, with a median of 228% and a large variance, mostly driven by heterogeneity in earnings, with some home-buyers having very small incomes in 1996 (column 2 of Table 3.1).

Table 3.1 — Housing discounts and house price appreciation

	(1) Housing discount in 1996	(2) Housing discount relative to 1996 earnings	(3) House price increase 1996-2016	(4) House price increase relative to 1996 earnings
5th percentile	33.2%	96.1%	242.8%	198.9%
25th percentile	57.5%	155.9%	270.8%	326.2%
50th percentile	62.6%	227.5%	375.1%	517.2%
75th percentile	67.3%	489.0%	472.4%	1,164.0%
95th percentile	71.7%	7,017.6%	480.6%	15,509.9%
Mean	61.4%	2,373.7%	365.5%	5,364.4%
Standard deviation	92.11	11,803.45	85.23	26,544.33

Notes: Summary statistics are based on the sample of 8,287 homebuyers described in Section 2.2.1. The housing discount is obtained by comparing 1996 assessment prices per squared meter at the ownership identifier level with a neighbourhood-level house price index per squared meter in 1996. Earnings in 1996 include income from self-employment. The housing discount relative to 1996 earnings is obtained after multiplying the discount price by the size of the apartment, and dividing the result by 1996 earnings. The house price increase between 1996 and 2016 is calculated using the neighbourhood-level house price index. The house price increase relative to 1996 earnings is obtained by multiplying 1996 assessment prices by the house price increase over the period 1996-2016, multiplying this by the size of the apartment, and dividing the result by 1996 earnings. The second and fourth columns are winsorized at the 1st and 99th percentiles.

These substantial discounts allowed almost all residents to purchase their home, despite home-buyers belonged to the poorest share of the population. As a result, only 10% of the municipality's housing stock was not sold to incumbent residents and was instead bought mainly by companies (7%), transferred to public housing associations (2%) or to other public institutions (1%). In many cases, these remaining apartments were small properties in poor locations with limited development opportunities, or they were relatively expensive because they had been recently renovated.

In addition to the home price discounts, the period under study was characterized by rapid house price increases, with a median increase of 375% between 1996 and 2016 across all neighbourhoods in Copenhagen (column 3 of Table 3.1). Comparing this to 1996 earnings (column 4 of Table 3.1), the overall wealth effect was substantial, with a median of 517%

and a large variance.⁸ As a result, there was almost no selection effect and practically all residents bought the municipality housing units. This is different from other settings studied in the literature, for example the privatization of public housing in Israel (Hausman et al., 2022), where only 34% of eligible households bought the apartment.

3.2.2. Data and descriptive statistics

We combine various administrative datasets from Denmark Statistics, covering the entire population over the period 1986-2018. The data contains detailed information on individual income and tax records, labor market participation, household characteristics, individual demographics and health status, residence and descriptives of housing units at yearly frequency.

We identify housing units that were owned by the municipality in 1993, before the policy was announced. The privatization of these housing units took place between 1997 and 2001, with a concentration of sales in 1998 (Figure C.1 in Appendix C.1). Both types of public housing units were spread around the city of Copenhagen in similar neighbourhoods (Figure C.2 in Appendix C.1), which allows us to assess the effect of these policies independently of neighbourhood effects.

The treatment group corresponds to the subset of individuals that lived in municipality-owned public housing in Copenhagen continuously between 1996 and its privatization.⁹ The control group, by contrast, includes individuals who rented from public housing associations continuously between 1996 and 2001.

Municipality housing units are assigned using priority lists to individuals with high housing risk, such as single mothers, individuals with health issues or that have been long-term unemployed, and refugees. A key difference with public housing associations is that the latter assign housing units based on both priority lists and waiting lists. While the first methodology mirrors the procedure used by municipality housing units, waiting lists are not subject to an income ceiling and anyone can sign up.¹⁰ Our objective is therefore to identify

8. Some of these very large wealth effects were again driven by the fact that a portion of home-buyers had very small incomes in 1996.

9. The treatment sample excludes Frederiksberg, a neighbourhood within the city of Copenhagen that was not affected by the policy because it is formally a different municipality from the rest of Copenhagen.

10. Waiting lists are also extremely congested, with an average of 20 years for the most popular housing

the subset of renters in public housing associations that are on priority lists and that are thus more directly comparable to those living in municipality housing.

We start by focusing on working age individuals, thus excluding individual-period observations when individuals had less than 25 or more than 60 years of age, and impose the condition that individuals must have been at least 25 years old in 1996, the year preceding the start of the sales period. Moreover, we exclude student residences and individuals that were students in 1996.¹¹ We further drop housing units with size above the top 1st percentile or with more than 10 residents in 1996. Furthermore, we keep individuals in the sample if they appear for at least 5 years between 1986 and 2018, including the year 1996.

This cleaning process leaves us with a total of 8,287 individuals in the treatment group, who lived in municipality public housing before privatization. Home-buyers can be compared to renters from public housing associations, who constitute a larger and more heterogeneous sample of 29,306 individuals. Their greater heterogeneity stems from the fact that only a portion of these individuals are on priority lists, while the rest is on waiting lists. This implies that renters from public housing associations are on average not directly comparable with treated individuals, as seen in Panel A of Table 3.2.

To identify a control group that is observationally identical to the treated group, we use a one-to-one matching procedure on ex ante (1996) characteristics. More specifically, we use exact matching on gender, parent and single status, and matching based on propensity score for other demographic characteristics (number of children, size of the household, age, civil status, highest completed education and whether they were born in Denmark), indicators of labor market participation (total income from work, wage income, hourly wage, number of days employed in the year, whether the contract is full-time, total social assistance benefits, an indicator describing the main source of income or employment in the income year, number of years in unemployment before 1996), measures of geographic mobility (total number of apartment moves between 1986 and 1996, and a dummy for whether the individual changed apartment between 1995 and 1996), and descriptives of health (number of hospital visits and days for non-pregnancy related illnesses). This second subsample is described in Panel B estates (Tsenkova, 2021).

11. Student residences are identified as dwellings where total residents are more than 10 and the share of students is greater than 70%.

Table 3.2— Descriptives of buyers and renters in full and matched samples in 1996

	Home-buyers	Renters	Difference	T-test of means
<i>Panel A: full sample</i>				
N. individuals	8,287	29,306		
N. households	6,996	23,021		
Age	41.5	41.9	0.4	2.8
% female	46.4%	54.3%	7.9	12.7
Number of children	0.40	0.53	0.1	12.1
Size of household	2.31	2.60	0.3	15.9
% single	37.9%	30.4%	-7.5	-12.6
% single parent	6.0%	7.4%	1.4	4.6
% danish born	73.6%	72.3%	-1.3	-2.5
% college graduate	10.5%	15.9%	5.4	13.5
N. years unemployed pre 1996	1.6	1.2	-0.4	-15.3
% receiving UI	17.8%	14.3%	-3.5	-7.5
% days unemployed	27.5%	21.0%	-6.5	-9.7
% full-time	52.6%	58.5%	5.9	9.5
Hourly wage (DKK)	106	117.1	11.1	9.7
Yearly earnings (DKK)	76,061.2	106,896.0	30,834.8	23.7
% receiving social benefits	15.1%	13.9%	-1.2	-2.5
Total cash assistance (DKK)	39,984.5	35,358.3	-4,626.2	-7.5
% retired	27.2%	19.7%	-7.5	-13.9
N. of apartment changes pre 1996	1.5	1.3	-0.2	-9.7
% changed apartment in 1996	14.3%	12.9%	-1.4	-3.2
% ever hospitalized	11.0%	8.9%	-2.1	-5.5
% ever hospitalized in family	17.2%	18.3%	1.1	2.3
N. hospital days	0.86	0.69	-0.2	-2.3
<i>Panel B: 1-1 matched sample</i>				
N. individuals	8,287	8,287		
N. households	6,996	7,009		
Age	41.5	41.5	0.0	-0.4
% female	46.4%	46.4%	0.0	0.0
Number of children	0.40	0.40	0.0	0.2
Size of household	2.31	2.31	0.0	0.2
% single	37.9%	37.9%	0.0	0.00
% single parent	6.0%	6.0%	0.0	0.0
% danish born	73.6%	73.1%	-0.5	-0.8
% college graduate	10.5%	9.7%	-0.8	-1.7
N. years unemployed pre 1996	1.6	1.6	0.0	-1.0
% receiving UI	17.8%	17.9%	0.1	0.2
% days unemployed	27.5%	28.0%	0.5	0.7
% full-time	52.6%	52.2%	-0.4	-0.6
Hourly wage (DKK)	106.0	105.9	-0.1	-0.1
Yearly earnings (DKK)	76,061.2	75,547.1	-514.1	-0.3
% receiving social benefits	15.1%	14.5%	-0.6	-0.9
Total cash assistance (DKK)	39,984.5	39,853.2	-131.3	-0.2
% retired	27.2%	27.9%	0.7	1.1
N. of apartment changes pre 1996	1.5	1.4	-0.1	-1.7
% changed apartment in 1996	14.3%	13.7%	-0.6	-1.2
% ever ill	11.0%	10.1%	-0.9	-1.9
% ever ill in family	17.2%	17.4%	0.2	0.2
N. hospital days	0.86	0.87	0.0	0.1

Notes: Descriptive statistics for treatment and control groups. Summary statistics are calculated based on individual characteristics in 1996 (the last year before the sales period). The % of days employed per year is defined as days during which the individual had a job contract out of a total of 365 in the year. Health descriptives are related to non-pregnancy hospital stays. Total cash assistance includes social benefits and unemployment insurance, but excludes retirement benefits.

of Table 3.2. In our main analysis, we observe 8,287 home-buyers and 8,287 public housing renters up to 5 years preceding the year of purchase and up to 10 years following it, with over 200 thousand observations.

3.3 — Effect on labor market participation

3.3.1. Empirical strategy

We study the intensive and extensive effect of subsidized home purchase (the treatment) versus subsidized rental (the control) on labor market participation, with a focus on understanding the longer-term effects of these different housing support policies. Our approach is based on comparing outcomes of individuals who bought the public housing unit to those who did not, after treatment relative to before. We estimate the average effect of the policy by comparing labor market participation up to 5 years preceding the year of purchase with up to 10 years following it.

To study the impact of subsidized purchase versus subsidized rental, we implement a difference-in-differences specification with the following elements. First we include a dummy that equals 1 after the year of purchase for buyers, whose effect is captured by β^{Buyer} . Second, we include a dummy that turns to 1 after the year of purchase for each buyer and her matched control renter, whose impact is measured by β^{All} . This is equivalent to an interaction term between a dummy that identifies buyers and a dummy that turns to 1 after the purchase year. In addition, we include individual fixed effects (FEs) α_i , year FEs α_t , and a set of individual and household time-varying controls captured by \mathbf{X}'_{it} . The resulting OLS difference-in-differences specification is the following

$$y_{it} = \alpha_i + \alpha_t + \beta^{Buyer} I_{it}^{Buyer} + \beta^{All} I_{it}^{All} + \mathbf{X}'_{it} \gamma + \varepsilon_{it}. \quad (3.1)$$

Individual FEs control for permanent individual characteristics, while \mathbf{X}'_{iht} summarizes a set of individual and household time-varying controls in period t . In our baseline regressions

we include a set of important interaction terms: the second and third power of age, the interaction of age and gender, and the interaction of age and highest completed education. While individual and year FEs control for age, age is likely to interact with other key individual characteristics in its effect on labor market participation. In addition, higher polynomials of age control for the non-linear impact of this variable.¹² Finally, ε_{iht} in equation (3.1) is an idiosyncratic error term and standard errors are clustered at the level of each match between treated individual and control.¹³

The main outcomes of interest measured by y_{it} are a set of variables describing labor market participation. We examine both the extensive and the intensive margins of labor supply, the first captured by a dummy equal to 1 if the main source of income in the year was labor income, the second measured by yearly labor income. The coefficient of interest in equation (3.1) is β^{Buyer} which denotes the differential labor market participation of individuals affected by the subsidized home purchase policy, compared to those living in subsidized rental apartments.

A key identifying assumption is that of parallel trends, namely that trends in labor market participation would be the same for both buyers and renters in the absence of treatment. We can test this assumption by examining the evolution of the dependent variables over time before the home sales period.

The difference-in-differences approach described in equation (3.1) measures the average change in labor supply between the period preceding the home purchase and the period following it, for the treatment relative to the control group. To study the effect of labor market participation for home-buyers relative to renters distinctively by year, we implement an OLS specification that corresponds to the dynamic version of equation (3.1). We include a set of leads and lags around the home purchase year for buyers (L_{it}^{Buyer}) associated with the coefficients $\{\beta_{\tau}^{Buyer}\}_{\tau=-5}^{10}$, where τ denotes time relative to year of purchase. We also

12. We obtain very similar results when additionally controlling for individual time-varying characteristics that are also used in the matching procedure: civil status, highest completed education, number of children, a dummy indicating non-pregnancy related hospitalization in a given year, and the relative number of days in the hospital.

13. The dataset also contains detailed geographical level information. The location of each apartment can be defined at the parish level, where parishes identify small neighbourhoods that divide the municipality of Copenhagen in 55 areas. However, we do not include parish FEs in our regression model, otherwise identification would be based on individuals who move to a different neighbourhood over the sample period.

add a set of leads and lags around the time of purchase that is common to both buyers and their respective placebo control (L_{it}^{All}) associated with the coefficients $\{\beta_{\tau}^{All}\}_{\tau=-5}^{10}$. As in equation (3.1), we augment this with individual FEs α_i , year FEs α_t , and a set of individual and household time-varying controls captured by \mathbf{X}'_{it} . This results in the following OLS specification

$$y_{it} = \sum_{\tau=-5}^{10} \beta_{\tau}^{Buyer} \mathbf{1}[L_{it}^{Buyer} = \tau] + \sum_{\tau=-5}^{10} \beta_{\tau}^{All} \mathbf{1}[L_{it}^{All} = \tau] + \alpha_i + \alpha_t + \mathbf{X}'_{it} \gamma + \varepsilon_{it}. \quad (3.2)$$

The year FEs control for time trends, while the lags and leads common to all individuals (L_{it}^{All}) control for joint dynamics around the time of home purchase. Figure 3.2 in Section 3.3.2 depicts the dynamic effect and shows that the parallel trends assumption is satisfied.¹⁴

3.3.2. Main results

Labor supply

Table 3.3 reports a first set of results, estimating the home-ownership effect on the extensive margin of labor supply (columns 1 and 2) and on the intensive margin of labor supply (columns 3 and 4) up to ten years after treatment. More specifically, the first column uses as dependent variable a dummy equal to one if the worker has received positive income from working in a given year, including from self-employment. The dependent variable in the second column is a dummy equal to one if the worker has received positive income from working as an employee in a given year, thus excluding self-employment. For both variables, home purchase has a negative effect on labor market participation: becoming a homeowner decreases the probability of working in a given year by 2.8% relative to the sample mean.

Looking at the intensive margin of labor supply, the third column of Table 3.3 examines the effect of home-ownership on total yearly labor income (including income from self-employment, expressed in DKK), while the fourth column uses as dependent variable yearly income from employment, excluding self-employment. According to both metrics, we observe a reduction

14. However, the limited number of observations per year results in relatively large confidence intervals.

Table 3.3 — Estimates of home-ownership on labor supply

Dependent variable	(1) Working	(2) Employed	(3) Labor income	(4) Wage income
Home-buyer	-0.012 (0.005)	-0.012 (0.005)	-4,669.0 (1,257.0)	-4,303.0 (1,267.0)
Individual FEs	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
Controls	✓	✓	✓	✓
N. obs	222,916	222,916	222,916	222,916
Adjusted R^2	0.634	0.626	0.739	0.737
Mean dep. var.	0.422	0.416	107,239.1	102,130.1
Coeff / mean	-2.8%	-2.8%	-4.4%	-4.2%

Notes: OLS regressions estimated up to 5 years preceding the year of purchase and up to 10 years following it, based on equation (3.1). The means of dependent variables are calculated over this time window. The sample is described in Panel B of Table 3.2. Standard errors in parentheses are clustered at the match identifier level. Labor and wage income are expressed in DKK.

in labor income among home-buyers relative to renters by over 4% of the sample mean. Overall, Table 3.3 suggests that home-buyers reduce their labor market participation relative to renters in the long run.

Table 3.4 presents additional results on labor market participation, to yield a clearer picture of the margins of adjustments. Column (1) examines the impact of home-ownership on the hourly wage, which is only available for a subset of workers that are employees, and highlights the absence of any effect on hourly wages. The probability of working full-time reported in column (2) instead decreases by 2.5% (relative to the sample mean) and the coefficient is statistically significant. These first two columns thus suggest that the decrease in earnings is due to individuals working fewer hours or days in the year, and not due to lower wages: individuals who continue to work reduce their working time to part-time. Moreover, there is no impact on the probability of being self-employed (column 3).

These results suggest that home-buyers tend to reduce labor supply, both at the extensive and at the intensive margins, for example by switching to part-time. This negative response of labor supply is likely driven by the large wealth effect experienced by home-buyers, as shown in Table 3.1. These outcomes contrast with those described by Hausman et al. (2022), who exploit a public housing privatisation event in Israel and find that home-buyers increase their labor supply. However, Hausman et al. (2022) focus on the very short-run effects of

Table 3.4 — Estimates of home-ownership on labor supply

Dependent variable	(1) Hourly wage	(2) Working full time	(3) Self- employment
Home-buyer	-0.009 (1.016)	-0.013 (0.004)	-0.002 (0.002)
Individual FEs	✓	✓	✓
Year FEs	✓	✓	✓
Controls	✓	✓	✓
N. obs	115,273	222,916	222,916
Adjusted R^2	0.471	0.460	0.503
Mean dep. var.	126.77	0.508	0.028
Coeff / mean	0.0%	-2.5%	-7.4%

Notes: OLS regressions estimated up to 5 years preceding the year of purchase and up to 10 years following it, based on equation (3.1). The means of dependent variables are calculated over this time window. The sample is described in Panel B of Table 3.2. Standard errors in parentheses are clustered at the match identifier level. The hourly wage is expressed in DKK.

public housing privatization (up to 6 months after the start of the sales period), while they are silent on the long-run effect. The large wealth effect implied by the home purchase subsidies is unlikely to appear in their results, especially since this effect amplifies over time, pushed by house price appreciation. By contrast, our long panel of administrative register data enables us to follow individuals over time and over space.

Housing and mobility

As described in Section 3.3.2, the steep increase in wealth among treated individuals may be large enough to explain falling labor supply. This is especially the case if we consider that, among our sample of low-income households, buying a home would be by far the single largest investment made over their lifetime. At an age of about 40 years old, treated individuals found themselves in the position of having reached this important life milestone at a very affordable price.

One way to cash in this wealth could be to sell the apartment and move to a different home, or even back into public housing since residents could have retained their place on waiting lists. Column (1) of Table 3.5 however suggests that home-buyers are less likely to move into public housing following their home purchase, compared to renters. The dependent

variable in this case is a dummy equal to 1 in the year in which individuals move into public housing.

Table 3.5 — Estimates of home-ownership on housing decisions

Dependent variable	(1) Moving into public housing	(2) Selling apartment	(3) Renting	(4) Living in owned home
Home-buyer	-0.006 (0.003)	0.005 (0.001)	-0.024 (0.003)	0.025 (0.003)
Individual FEs	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
Controls	✓	✓	✓	✓
N. obs	222,916	222,916	222,916	222,916
Adjusted R^2	0.865	0.017	0.267	0.307
Mean dep. var.	0.539	0.007	0.966	0.029
Coeff / mean	-1.1%	63.7%	-2.5%	86.1%

Notes: OLS regressions estimated up to 5 years preceding the year of purchase and up to 10 years following it, based on equation (3.1). The means of dependent variables are calculated over this time window. The sample is described in Panel B of Table 3.2. Standard errors in parentheses are clustered at the match identifier level.

At the same time, it is difficult to investigate whether home-buyers sell their home after having bought it. Home transaction data cannot track precisely the sale of these apartments, as they were part of cooperatives.¹⁵ In column (2), we proxy the sale of an apartment by a dummy equal to 1 if individuals change apartment and in the same year move to a new apartment where they are classified as owners. Looking at this variable suggests that home-buyers see a relative increase in the propensity to sell their home. Columns (3) and (4) confirm that the treatment group is less likely to rent and more likely to own their home throughout the 10 years following the home sales policy.

The data further allow us to observe mobility patterns, and in particular whether individuals changed apartment, neighbourhood, or even city. Column (1) of Table 3.6 uses as dependent variables a dummy that turns to 1 when individuals change apartment, showing that home-buyers are relatively more likely to move compared to renters. Similar results hold for changing neighbourhood (column 2), or even municipality (column 3). This increased mobility means that home-buyers are slightly less likely to live in Copenhagen's centre on

15. Home transaction data is identified by a property identifier, which in the case of cooperatives is the same for the entire cooperative. Therefore, it is not linked to the share of the cooperative owned by each individual.

average (column 4), and more likely to move outside of the metropolitan area (column 6), although they don't appear to move to Copenhagen's suburbs (column 5).

Table 3.6 — Estimates of home-ownership on mobility probability

Dependent variable	(1) Changing apartment	(2) Changing neighbour- hood	(3) Changing municipality	(4) Living in CPH centre	(5) Living in CPH suburbs	(6) Living outside CPH
Home-buyer	0.047 (0.006)	0.058 (0.005)	0.047 (0.004)	-0.029 (0.004)	0.002 (0.003)	0.003 (0.001)
Individual FEs	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
N. obs	222,916	222,916	222,916	222,916	222,916	222,916
Adjusted R^2	0.547	0.513	0.460	0.394	0.397	0.004
Mean dep. var.	0.162	0.120	0.068	0.935	0.038	0.007
Coeff / mean	29.0%	48.3%	69.6%	-3.1%	4.9%	37.1%

Notes: OLS regressions estimated up to 5 years preceding the year of purchase and up to 10 years following it, based on equation (3.1). The means of dependent variables are calculated over this time window. The sample is described in Panel B of Table 3.2. Standard errors in parentheses are clustered at the match identifier level. Columns (4)-(6) describe three alternative scenarios: the probability of staying in Copenhagen's (CPH) city centre (the municipality of Copenhagen and that of Frederiksberg), the probability of moving outside of the centre (in the suburbs) but within Copenhagen's metropolitan area, and the probability of moving outside of the Copenhagen metropolitan area.

These results confute the Oswald hypothesis, according to which home-ownership restricts residential mobility (Oswald, 1996). Buying a home makes individuals more mobile because they can sell it or rent it and move somewhere else. This result is likely not specific to our context: in many European cities there are high frictions to mobility for residents of public housing due to long waiting lists.

Public benefits

Purchasing a home and the resulting long-run labor supply decisions interact in important ways with the kind of public benefits accessed by individuals. Table 3.7 analyzes the effect of home-buying on take-up of public benefits, relative to individuals that continued to rent their apartment. These results show that home-buyers are less likely to receive unemployment benefits (column 1), although the effect is not statistically significant. As home-buyers reduce their long-run labor supply, they are no longer eligible for unemployment benefits,

which in the Danish “flex-security” system require availability to take up a new job with a one-day notice. To compensate, home-buyers are more likely to receive other other types of social benefits after their home purchase (column 2).¹⁶ By contrast, there is almost no effect on anticipated retirement (column 3), since retirement age for the Danish state pension is relatively high (between 65 and 68, depending on the year of birth).

Table 3.7 — Estimates of home-ownership on up-take of public benefits

	(1)	(2)	(3)	(4)
Dependent variable	Unemployment	Social benefits	Retired	Total public benefits
Home-buyer	-0.004 (0.004)	0.012 (0.005)	0.005 (0.004)	1,403.0 (603.4)
Individual FEs	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
Controls	✓	✓	✓	✓
N. obs	222,916	222,916	222,916	221,964
Adjusted R^2	0.244	0.432	0.780	0.508
Mean dep. var.	0.105	0.158	0.274	34,591.17
Coeff / mean	-3.7%	7.7%	1.9%	4.1%

Notes: OLS regressions estimated up to 5 years preceding the year of purchase and up to 10 years following it based on (3.1). The means of dependent variables are calculated over this time window. The sample is described in Panel B of Table 3.2. Standard errors in parentheses are clustered at the match identifier level. Social benefits include the daily allowance (activation and similar, sickness, maternity and leave) and cash assistance. Total public benefits are expressed in DKK and include social assistance and unemployment benefits, but exclude retirement allowances.

The end result is actually an increase in the amount of yearly public benefits (the sum of unemployment and social assistance benefits, excluding early retirement) received by home-buyers relative to renters by almost 4% of the sample mean (column 4). On top of the wealth effect, this increase in public benefits may also partly explain the falling labor supply among home-buyers, through a reduction in work incentives.

Overall, home-buyers see lower labor market participation than renters in the long run and

16. These social benefits are the sum of the daily allowance and cash assistance. The daily allowance includes different types of public benefits: sickness allowance, maternity benefits, activation and rehabilitation benefits, education and labor market benefits for unemployed people who have used up their unemployment benefits. By contrast, cash assistance is an additional public benefit that is given to adults who would otherwise not be able to support themselves or their family. Eligibility criteria for cash assistance require registration with the job center and accepting the job or activation offers (such as training and inclusion programs) that individuals receive. Cash assistance also differs from the daily allowance because it is paid from the general social security funds, while other transfers are paid by specific social security funds to which individuals contribute, for example the health insurance fund when it comes to sickness allowance or maternity benefits.

continue to depend on the social security system. We observe similar results when reducing the horizon of observations to up to 5 years after purchasing the home (Table C.1 of the Appendix), or when considering only 5-10 years after purchase (Table C.2 of the Appendix).

Dynamic effects

Figure 3.2 details the timing of these effects. The figures each present the evolution of the treatment effect by year, before and after purchasing the home. Overall, treatment effects for both extensive and intensive margins are relatively flat around zero prior to the home purchase. After home purchase, however, homeowners' participation in the labor market relative to renters decreases and stabilizes at this lower level (panel a), while labor income falls gradually after purchasing the home (panel b). The fall in labor market income is driven by a lower probability of working full-time, and not by a fall in the hourly wage.

While these dynamic effects suggest that the parallel trends assumption is satisfied, the limited number of observations per year results in relatively large confidence intervals. Figures C.3 and C.4 in Appendix C.2 display the dynamic effects of variables describing housing, mobility, and the take-up of social benefits.

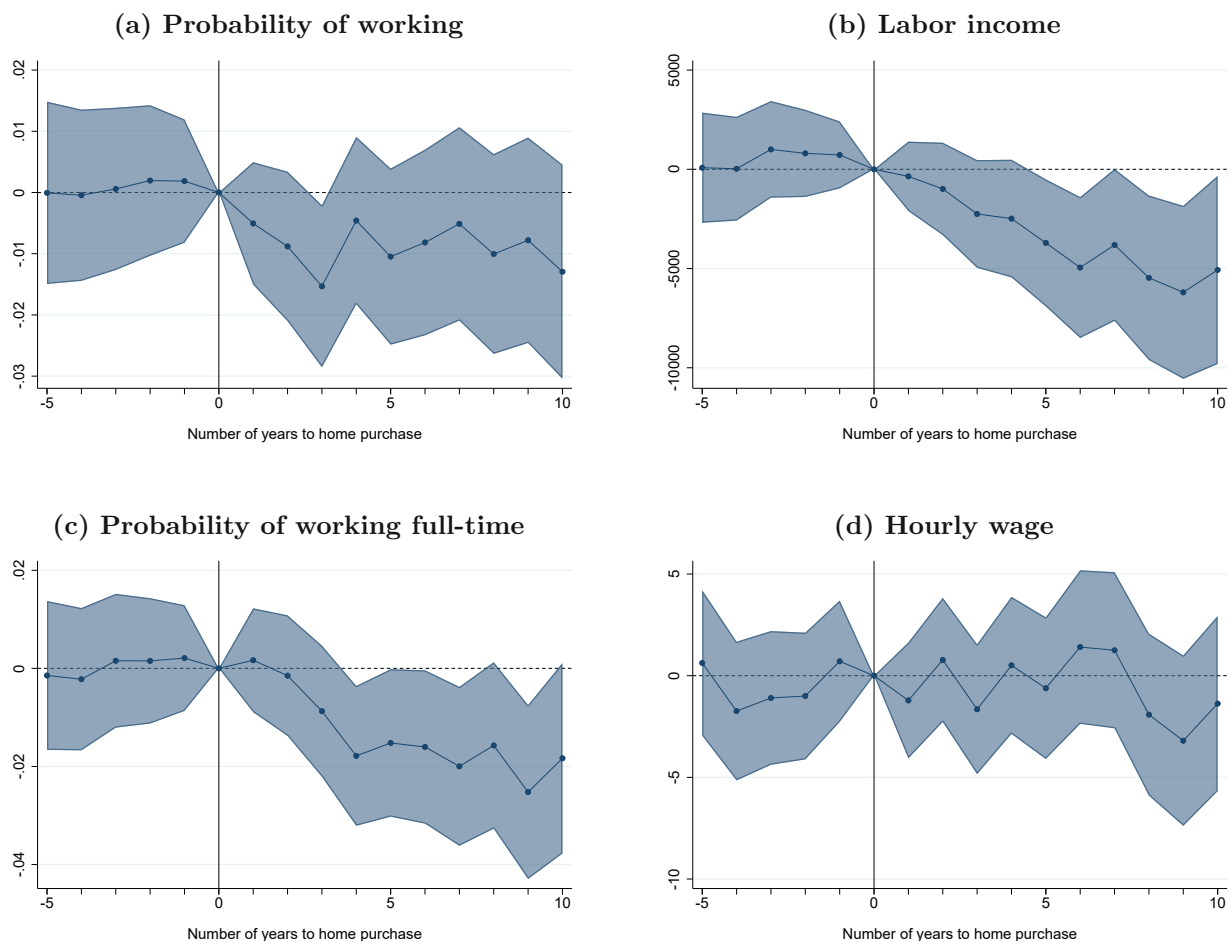


Figure 3.2 — Pre-trends coefficients and treatment effects: labor supply

Notes: Points on the graphs represent coefficients on $\text{treatment} \times \text{year}$ interactions, where treatment is defined as purchasing the home during the sales period (1997-2001). The sample is described in Panel B of Table 3.2. The first row describes the extensive and intensive margins, namely (a) the probability of working in the year (including self-employment), and (b) yearly labor income. The second row focuses on the margins of labor market adjustment: (c) the probability of working full-time, and (d) the hourly wage. We include individual and year FEs, and we control for the interaction of age and gender, the interaction of age and highest completed education, the second and third powers of age.

Heterogeneity by gender

Table 3.8 summarises the characteristics of buyers and renters in the matched sample, distinguishing by gender.

Table 3.8 — 1996 descriptives of buyers and renters in matched sample by gender

	Females	Males	Difference	T-test of means
N. individuals	7,694	8,880		
Age	41.2	41.8	0.6	3.8
Number of children	0.82	0.03	-0.8	-61.2
Size of household	2.49	2.15	-0.3	-14.9
% single	29.8%	44.9%	15.1	20.3
% single parent	12.5%	0.0%	-12.5	-31.5
% danish born	73.8%	73.1%	-0.7	-1.0
% college graduate	12.6%	7.9%	-4.7	-10.1
N. years unemployed pre 1996	1.4	1.9	0.5	17.8
% receiving UI	14.7%	20.6%	5.9	10.1
% days unemployed	24.7%	30.3%	5.6	6.5
% full-time	48.6%	55.7%	7.1	9.1
Hourly wage (DKK)	103.9	107.7	3.8	2.7
Yearly earnings (DKK)	69,368.6	81,414.2	12,045.6	7.7
% receiving social benefits	18.1%	11.9%	-6.2	-11.2
Total cash assistance (DKK)	40,633.4	39,299.7	-1,333.7	-1.7
% retired	28.0%	27.2%	-0.8	-1.2
N. of apartment changes pre 1996	1.2	1.7	0.5	21.4
% changed apartment in 1996	11.8%	15.9%	4.1	7.6
% ever hospitalized	11.0%	10.1%	-0.9	-1.8
% ever hospitalized in family	19.3%	15.6%	-3.7	-6.3
N. hospital days	0.74	0.98	0.24	2.5

Notes: Summary statistics by gender for buyers and renters in the matched sample described in Panel B of Table 3.2. Summary statistics are calculated based on individual characteristics in 1996 (the last year before the sales period). The % of days employed per year is defined as days during which the individual had a job contract out of a total of 365 in the year. Health descriptives are related to non-pregnancy hospital stays.

It is interesting to note that males and females living in public housing (either owned by the municipality or by public housing associations) differ substantially along several dimensions. In particular, females are much more likely to be single mothers, and to therefore live in households with kids but without a partner. By contrast, there are few male parents in our sample, and males tend to live with their partner only. Females are also on average significantly more educated than males and less likely to be unemployed. This is however

not reflected in higher hourly wages relative to males and especially in yearly earnings, as females are less likely to work full-time.

We repeat the analysis of Section 3.3.2 separately by gender in Table 3.9, which shows that the decrease in labor supply is mostly driven by males. Compared to females, males see larger declines in labor market participation after buying their home, both at the extensive and intensive margins. Among females, we observe a smaller and not statistically significant effect of home-ownership on labor supply. The almost null effect of home-ownership on female labor supply can be interpreted in light of the descriptives seen in Table 3.8. Since females are much more likely to be single mothers, decreasing labor supply after purchasing the home is less feasible for them.

Table 3.9 — Estimates of home-ownership on labor supply by gender

Dependent variable	Males			Females		
	(1) Working	(2) Labor income	(3) Working full time	(4) Working	(5) Labor income	(6) Working full time
Home-buyer	-0.020 (0.007)	-7,830.0 (1,801.0)	-0.015 (0.007)	-0.003 (0.007)	-1,068.0 (1,736.0)	-0.010 (0.008)
Individual FEs	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
N. obs	118,408	118,408	118,408	104,508	104,508	104,508
Adjusted R^2	0.628	0.740	0.616	0.640	0.736	0.601
Mean dep. var.	0.412	113,786.3	0.521	0.434	99,836.1	0.493
Coeff / mean	-4.8%	-6.9%	-2.9%	-0.7%	-1.1%	-2.0%

Notes: OLS regressions estimated up to 5 years preceding the year of purchase and up to 10 years following it, based on equation (3.1). The means of dependent variables are calculated over this time window. The sample is described in Panel B of Table 3.2. Standard errors in parentheses are clustered at the match identifier level. Labor income is expressed in DKK.

When it comes to the effect on take-up of public benefits, female home-buyers see a larger increase in the probability of unemployment, although the coefficient is not statistically significant for both gender groups (Table 3.10). By contrast, among males we observe larger effects of home-ownership versus rental on increasing take-up of social benefits. The overall size of total public benefits increases by 8% of the sample mean among male home-buyers relative to renters, while there is no effect among females. These responses can be linked to the larger decline in labor supply among males following the home purchase, implying that

males increasingly rely on public benefits.

Table 3.10 — Estimates of home-ownership on public benefits by gender

Dependent variable	Males			Females		
	(1) Unem- ployment	(2) Social benefits	(3) Total public benefits	(4) Unem- ployed	(5) Social benefits	(6) Total public benefits
Home-buyer	-0.002 (0.006)	0.015 (0.006)	2,674.0 (799.3)	-0.007 (0.005)	0.009 (0.007)	-38.6 (914.6)
Individual FEs	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
N. obs	118,408	118,408	118,408	104,508	104,508	104,508
Adjusted R^2	0.263	0.413	0.504	0.213	0.445	0.512
Mean dep. var.	0.120	0.137	33,480.5	0.088	0.181	35,847.8
Coeff / mean	-1.3%	10.6%	8.0%	-7.4%	5.1%	-0.1%

Notes: OLS regressions estimated up to 5 years preceding the year of purchase and up to 10 years following it, based on equation (3.1). The means of dependent variables are calculated over this time window. The sample is described in Panel B of Table 3.2. Standard errors in parentheses are clustered at the match identifier level. Social benefits include the daily allowance (activation and similar, sickness, maternity and leave) and cash assistance. Total public benefits are expressed in DKK and include social assistance and unemployment benefits, but exclude retirement allowances.

3.4 — Intergenerational effect

In Section 3.3.2, we have seen that home-buying reduces labor market participation, especially among male workers. One question that arises is whether observed labor supply response may have been larger if individuals in our sample had not already been on an established career path. At the start of the sales period, public housing residents were on average around 40 years old. They had therefore already completed their education and accumulated experience in the labor market. This could have limited the impact of the policy on their labor supply. A different story could be true for their children, since, at the time, their education decisions still had to be finalized and they had not yet made the first steps in the labor market. Put differently, we may be able to observe a delayed response of home-buying choices on future generations. The opposite hypothesis is instead that labor supply effects of home purchase

policies dilute through generations, especially as the Danish public education system equalizes opportunities among children with different backgrounds.

Buying a house is a long-term project to build family wealth, which can have important inter-generational effects on children (Dietz and Haurin, 2003). For many households, this is the biggest investment in their lifetime, a key source of wealth that is crucial for the household's financial security. Some evidence points to home-ownership being associated with positive impacts on children welfare and economic success, including higher education status (Green and White, 1997), well-being (Haurin et al., 2002), and children's own probability of becoming home-owners as adults (Boehm and Schlottmann, 1999). However, the effect of home-ownership on inter-generational mobility has been mostly studied descriptively using (administrative) panels and controlling for as many observables as possible, since it is difficult to find plausible sources of exogenous variation (Aaronson, 2000). Another approach has been that of using house price shocks that increase the value of housing for households that already own a house (Daysal et al., 2022).

At the same time, it is difficult to find settings that allow to study the impact of subsidized home-ownership on children from low-income households, since affordable housing policies tend to be focused on rental. In this context, there is some evidence that living in public housing negatively affects children's labor market outcomes as adults. A large part of this evidence is focused on the US and studies housing vouchers that allowed households to rent at subsidized prices from the private rental market, instead of living in public housing projects (Chetty and Hendren, 2018; Chyn, 2018; Pollakowski et al., 2022).¹⁷ These results are however very specific to the US context, where public housing projects are often located in deprived neighbourhoods that are more distant from job centers, and they do not allow to disentangle the public housing effect from that of location.

17. Pollakowski et al. (2022), for example, compare households living in public housing projects in the US to those whose public housing projects were demolished in the context of the HOPE VI program, in exchange for housing vouchers. They find that subsidized rental positively impacted adult labor market participation of children affected by the HOPE VI program, mainly driven by improved job accessibility relative to those children who grew up in public housing projects. Similar evidence by Chyn (2018) and Chetty and Hendren (2018) finds that housing vouchers improved outcomes of children, as households had the opportunity to move to better neighbourhoods. In contrast to these studies, Jacob et al. (2014) explore the effect of housing assistance on child outcomes, comparing families that received housing vouchers through a lottery to those that do not (but who did not live in public housing projects), and find negligible effects on a range of child outcomes.

Another aspect to consider is that, while home-ownership is increasingly becoming a wealth maker and a measure of economic status, it is also a rising source of housing inequalities. Across the world, we observe increasing inter-generational persistence of home-ownership and a fall in home-ownership rates among those whose parents did not own their own home (Boehm and Schlottmann, 1999; Daysal et al., 2022).

3.4.1. Empirical strategy and summary statistics

Our unique setting enables to study the impact of home-ownership on inter-generational mobility by exploiting the quasi-natural experiment of the home sales policy. We compare the children of home-buyers to those of renters at a given age, in a cross-sectional framework, following Haltiwanger et al. (2020). We thus run the following OLS regression

$$y_{ia} = \beta^{PBuyer} I_i^{PBuyer} + \alpha_a + \alpha_{ih} + \mathbf{X}'_{ia} \gamma + \varepsilon_{ia}, \quad (3.3)$$

where i identifies the child, a indicates the specific age at which we compare the children of buyers and renters, and h denotes the household identifier in 1996. In the main results we focus on children that were on average 7-8 in 1996 and compare their outcomes at age $a = 30$ years old, the last year observed in our data.

The dummy I_i^{PBuyer} identifies children of parents who bought the house, versus children of renters, α_a controls for the year when individual i is 30, while α_{ih} corresponds to 1996 household FEs. The coefficient of interest is thus β^{PBuyer} , which captures the difference in the outcome variable between children of home-buyers and of public housing renters at the age of 30. We also include a set of additional controls \mathbf{X}'_{ia} that are either fixed for each individual (gender and whether the child was born in Denmark) or taken at age a (civil status, highest-completed education, number of children, family size, whether the individual went to the hospital in the year due to non-pregnancy related illness, and the number of days in the hospital). Standard errors are clustered at the 1996 household level.

Table 3.11 confronts the children of buyers (1,621 children) to those of renters (1,835 children) at age 30. The sample is relatively homogeneous when comparing adult outcomes

Table 3.11 — Descriptives: children of home-buyers and of public housing renters

	Children of home buyers	Children of renters	Difference	T-test of means
N. individuals	1,621	1,835		
N. households at age 30	1,614	1,825		
Age in 1996	8.1	7.1	-1.0	-8.0
<i>A. Demographic characteristics at age 30</i>				
% female	46.9%	48.2%	1.3	1.0
Number of children	0.40	0.40	0.0	0.3
Size of household	2.09	2.12	0.0	0.7
% single	62.3%	63.9%	1.6	0.9
% single parent	10.7%	9.8%	-0.9	-0.8
% danish born	54.0%	51.5%	-2.5	-1.5
% college graduate	24.4%	28.4%	4.0	2.7
<i>B. Labor supply at age 30</i>				
% receiving UI	4.5%	4.0%	-0.5	-0.7
% days unemployed	21.7%	20.3%	-1.4	-1.1
% full-time	76.0%	78.1%	2.1	1.2
Hourly wage (DKK)	144.6	147.0	2.4	0.8
Yearly earnings (DKK)	197,966.2	209,892.6	11,926.4	2.2
% receiving social benefits	16.2%	14.2%	-2.0	-1.7
Total cash assistance (DKK)	39,445.2	37,912.9	-1,532.3	-0.7
<i>C. Health at age 30</i>				
% ever hospitalized	8.8%	7.4%	-1.4	-1.5
% ever hospitalized in family	27.6%	26.7%	-0.9	-0.6
N. hospital days	0.39	0.26	-0.1	-1.2
<i>D. Housing characteristics at age 30</i>				
% renters	74.8%	75.3%	0.5	0.3
% home owners	13.8%	14.6%	0.8	0.7
% public housing	30.4%	50.0%	19.6	10.8
% private housing	19.4%	20.4%	1.0	0.7
% cooperative housing	41.6%	20.0%	-21.6	-12.7
% living in CPH centre	65.5%	65.7%	0.2	0.2
% living in CPH suburbs	22.5%	25.3%	2.8	2.0
% living outside of CPH metropolitan area	12.1%	9.0%	-3.1	-3.0
% living in 1996 apartment of parents	9.9%	6.2%	-3.7	-4.0

Notes: Summary statistics for children of treatment and control groups. Summary statistics are calculated based on individual characteristics in 1996 (the last year before the sales period) or at age 30. The % of days employed per year is defined as days during which the individual had a job contract out of a total of 365 in the year. Health descriptives are related to non-pregnancy hospital stays. The last row describes individuals that live in the same house inhabited by their parents in 1996. For children of home-buyers, this corresponds to the home that was eventually purchased by their parents during the sales period. Copenhagen (CPH) centre includes the municipality of Copenhagen and the one of Frederiksberg.

at this age, with some differences especially in terms of education status and yearly earnings: at 30, children of renters appear to have a more advanced education and to earn more on average. Housing characteristics are also different among children of home-buyers and those of renters, who seem to mirror their parents when it comes to living in cooperatives versus in public housing. Children of home-buyers are also twice more likely to live in the same house that was purchased by their parents during the sales policy.

3.4.2. Results for children

Table 3.12 summarizes regression results on labor supply and take-up of public benefits estimated based on equation (3.3). Both the extensive margin of labor supply (column 1) and the intensive margin (column 2) are not statistically different between the two groups, although the negative coefficients suggest that children of home-buyers have reduced labor supply relative to children of renters. The effect of parents' home-buying on children's take-up of public benefits (columns 3-4) is also not statistically significant, although the positive coefficient indicate increased reliance of children of home-buyers on public benefits.

Table 3.12 — Estimates of home-ownership on children's labour supply

	(1)	(2)	(3)	(4)	(5)
	Working	Labor income	Unemployment	Social assistance	College graduate
Children of home-buyers	-0.023 (0.015)	-3,719 (4,985)	0.003 (0.007)	0.011 (0.012)	-0.032 (0.015)
1996 household FEs	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
N. obs	3,456	3,456	3,456	3,456	3,456
Adjusted R^2	0.154	0.311	0.013	0.088	0.080
Mean dep. var.	0.665	204,305.2	0.043	0.151	0.265
Coeff / mean	-3.5%	-1.8%	7.5%	7.5%	-12.0%

Notes: OLS regressions estimated at age 30 based on equation (3.3). The sample is described in Table 3.11. Standard errors in parentheses are clustered at the family identifier level. Labor income is expressed in DKK.

At the same time, we find that the probability of pursuing higher education is lower among children of homeowners relative to those of renters (column 5 of Table 3.12). In the literature, the probability of pursuing higher education is usually associated with higher parents' income

and may also rise when parents own a home (Green and White, 1997), although evidence is mixed (Mohanty and Raut, 2009; Bourassa et al., 2015). Since parents who are homeowners earn less relative to parents who are renters (Section 3.3.2),¹⁸ the income effect on children education appears to be stronger than the wealth effect due to the family owning a home. Nevertheless, lower education among children of home-buyers does not have a statistically significant impact on their earnings at age 30, although this effect may become stronger as they get older.

Another channel is the impact that home-ownership of parents may have on the residential outcomes of children when they become young adults, which could itself interact with labor supply decisions. Columns (1)-(3) of Table 3.13 examine whether children of homeowners are more likely to live in public housing, in cooperatives or in private housing, relative to children of renters.

Table 3.13 — Estimates of home-ownership on children’s housing type

	(1)	(2)	(3)	(4)	(5)	(6)
	Public housing	Cooperative	Private housing	Renting	Owning	Parents’ house
Children of home-buyers	-0.203 (0.018)	0.219 (0.017)	-0.002 (0.015)	-0.001 (0.015)	-0.004 (0.012)	0.036 (0.009)
1996 household FEs	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
N. obs	3,456	3,456	3,456	3,456	3,456	3,456
Adjusted R^2	0.087	0.086	0.056	0.058	0.058	0.048
Mean. dep. var	0.406	0.302	0.199	0.751	0.142	0.080
Coef / mean	-50.0%	72.4%	-1.0%	-0.1%	-3.0%	45.6%

Notes: OLS regressions estimated at age 30 based on equation (3.3). The sample is described in Table 3.11. Standard errors in parentheses are clustered at the family identifier level. Column (1) uses as dependent variable a dummy equal to 1 if individuals reside in public housing, while columns (2) and (3) denote individuals living in cooperative housing or in private housing. Columns (4) and (5) use respectively as dependent variables dummies that turn to 1 if individuals rent their home, or if they live in a home that they own. Column (6) uses as dependent variable a dummy equal to 1 for individuals who live in the same house inhabited by their parents in 1996. For children of home-buyers, this corresponds to the home that was eventually purchased by their parents during the sales period.

Results show that children of homeowners are significantly more likely to live in cooperatives (column 2) as adults rather than renting subsidized housing (column 1). However, this

18. Note that education status is matched between parent treatment and control groups in Section 3.3.2.

does not translate into increased home-ownership at age 30 (column 5) because most children living in cooperatives in the sample actually rent the apartments from their parents. By contrast, the probability of living in private housing (column 3) is not statistically different among the two sets of children. Overall, children of home-buyers are more likely to live in the house that their parents bought in 1996, compared to children of renters (column 6), and this effect is large (46% of the sample mean) and statistically significant.¹⁹

A related aspect is whether children of home-buyers end up living closer to the city centre and in richer neighborhoods, and the resulting impact on commuting time. Our results, displayed in Table 3.14, suggest that children of home-buyers are slightly less likely to live closer to the city centre (although the coefficient is not statistically significant) and that they have a 20% longer commuting time (relative to the sample mean) compared to children of renters.

Table 3.14 — Estimates of home-ownership on children’s housing location

	(1)	(2)	(3)	(4)	(5)	(6)
	Distance from city centre (km)	Commuting distance (km)	Living in rich neighbourhood	Living in CPH centre	Living in CPH suburbs	Living outside CPH
Children of home-buyers	0.595 (0.376)	2.46 (0.090)	0.071 (0.014)	0.000 (0.016)	-0.027 (0.014)	0.026 (0.010)
1996 household FEs	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
N. obs	3,456	3,456	3,456	3,456	3,456	3,456
Adjusted R^2	0.038	0.011	0.040	0.065	0.065	0.044
Mean dep. var.	7.547	12.105	0.212	0.656	0.240	0.104
Coeff / mean	7.9%	20.3%	33.7%	0.1%	-11.1%	25.1%

Notes: OLS regressions estimated at age 30 based on equation (3.3). The sample is described in Table 3.11. Standard errors in parentheses are clustered at the family identifier level. Distance from the city centre is expressed in kilometers. Commuting distance is calculated as the shortest road distance from the employee’s residential address to the workplace address and expressed in kilometers. Rich neighbourhoods are identified as the top 25 (out of 55 in Copenhagen) by house price appreciation between 1996 and 2016, calculated based on the house price index. Columns (4)-(6) describe three alternative scenarios: the probability of staying in Copenhagen’s (CPH) city centre (the municipality of Copenhagen and that of Frederiksberg), the probability of moving outside of the centre but within Copenhagen’s metropolitan area, and the probability of moving outside of the Copenhagen region.

19. Children of public housing renters cannot inherit the home of their parents. This means they don’t have priority in retaining the apartment if their parents relocate.

At the same time, relative to children of renters, children of home-buyers are more likely to live in richer neighbourhoods (column 3) and this effect is both large (34% of the sample mean) and statistically significant. The last three columns suggest that the probability of living in Copenhagen's metropolitan area does not differ substantially between the two groups (column 4), although children of renters are relatively more likely to move to Copenhagen's suburbs (column 5), while children of home-buyers have a higher chance of moving outside of the metropolitan area to other cities (column 6).

These results thus highlight that the intergenerational transmission of home status is an important mechanism. Children of home-buyers are 46% more likely to reside in the home that their parents bought, after the latter move out. While it appears that at the age of 30, these children have not yet inherited the home of their parents, this is likely to take place in later years. This translates into a higher probability of living in better neighbourhoods in the long run, while children of renters are more likely to move to Copenhagen's suburbs.

3.5 — Conclusion

Rising geographic income segregation in large cities, combined with increasing housing costs, has been a long-standing phenomenon. To respond to these challenges, local governments have implemented various housing support policies. Empirically, we know very little about the long-run effects of these policies, especially on labor market participation. This paper investigates the impact of housing support policies on long-run labor market outcomes.

Utilizing a quasi-natural experiment – the 1996 privatization of all public housing owned by the city of Copenhagen – and using panel data that covers the entire population from 1986 to 2018, we find that home purchase assistance had a negative impact on home-buyers' labor market participation. In the long run, individuals who continued to rent subsidized public housing performed better in the labor market, in particular when it comes to earnings, which on average decreased by over 4% per year among home-buyers. This is mostly driven by a reduction in full-time work among home-buyers, and by falling labor market participation incentives. Home-buyers experienced a substantial wealth gain and in the long run increased

their dependency from public benefits.

Studying the effect of home-ownership versus subsidized rental on individuals that were children at the time of the home purchase policy, we find limited effects on labor market participation of future generations. However, evidence suggests that public housing privatization can in the long run increase housing inequalities, due to the high inter-generational persistence of home-ownership.

These results may be of interest to policymakers and academics focusing on housing support policies. In urban areas, providing affordable housing for low-income households remains a key political and economic challenge. Our analysis shows that housing provision through subsidized rental continues to be one of the most important pillars of low-income housing policy. By contrast, public housing privatization has no positive effect in the long run, neither on individuals nor on their children, and could in the end lead to increased urban inequalities.

C.1 — Data appendix

In 1995, Copenhagen’s budget deficit amounted to DKK 2.9 (USD 0.42) billion, which led to the decision of privatizing public housing. The apartments owned by the municipality were first sold in 1996 to the municipality-owned real estate company TOR for a public assessed value of DKK 3.235 billion (USD 0.464 billion). Shortly after acquiring the properties, TOR sent out flyers to tenants, to inform them of the opportunity of purchasing their apartments through a cooperative scheme. Alongside this information, the leaflets included a guide describing the process of apartment acquisition starting in 1997. Residents were encouraged to engage in discussions, establish an association, and seek legal counsel. They were also advised to explore financing options, maintain open communication with TOR, and be prepared to act quickly. Figure C.1 shows the timeline of housing sales, while Figure C.2 displays the neighbourhood spatial distribution of municipality housing and of dwellings owned by public housing associations.




Figure C.1 — Municipality housing units by sale year

Notes: Number of municipality housing units sold in each year.

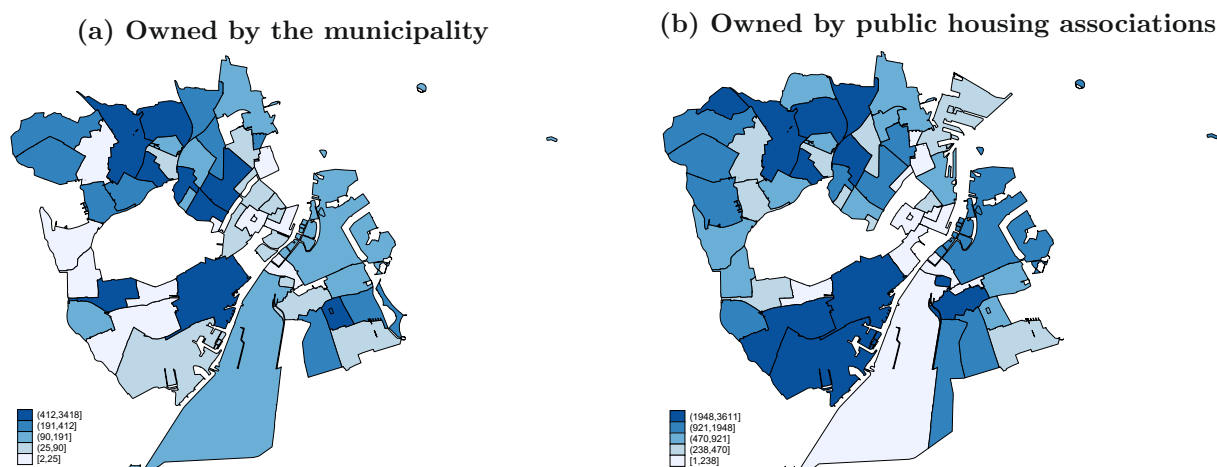


Figure C.2—Public housing by type of ownership

Notes: The maps display the number of housing units owned by the municipality of Copenhagen (panel a) and by public housing associations (panel b) in each neighbourhood. Colors identify 5 percentiles in each distribution. We exclude Frederiksberg, a neighbourhood in the centre the city of Copenhagen that was not affected by the policy because it is formally a different municipality from the rest of Copenhagen.

C.2 — Empirical appendix

Table C.1 — Estimates of home ownership on labor supply: 5-year window around purchase year

	(1)	(2)	(3)	(4)	(5)
Dependent variable	Working	Employed	Labor income	Wage income	Working full time
Home-buyer	-0.012 (0.005)	-0.012 (0.005)	-3,421.0 (1,100.0)	-3,212.0 (1,107.0)	-0.008 (0.005)
Individual FEs	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
N. obs	167,050	167,050	166,301	167,044	166,978
Adjusted R^2	0.646	0.639	0.756	0.757	0.630
Mean dep. var.	0.401	0.396	93,759.2	89,622.2	0.504
Coeff / mean	-3.0%	-3.0%	-3.6%	-3.6%	-1.6%

Notes: OLS regressions estimated up to 5 years preceding the year of purchase and up to 5 years following it based on equation 3.1. The sample is described in Panel B of Table 3.2. Standard errors in parentheses are clustered at the match identifier level. Labor and wage income are expressed in DKK.

Table C.2—Estimates of home ownership on labor supply: 5-10 years after purchase year

	(1)	(2)	(3)	(4)	(5)
Dependent variable	Working	Employed	Labor income	Wage income	Working full time
Home-buyer	-0.010 (0.006)	-0.010 (0.006)	-5,798.0 (1,724.0)	-52,39.0 (1,739.0)	-0.018 (0.007)
Individual FEs	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
N. obs	164,168	164,168	163,440	164,163	161,931
Adjusted R^2	0.274	0.610	0.725	0.721	0.589
Mean dep. var.	0.417	0.411	106,242.3	101,180.3	0.517
Coeff / mean	-2.5%	-2.3%	-5.5%	-5.2%	-3.4%

Notes: OLS regressions estimated up to 5 years preceding the year of purchase and 5-10 years following it based on equation 3.1. The sample is described in Panel B of Table 3.2. Standard errors in parentheses are clustered at the match identifier level. Labor and wage income are expressed in DKK.

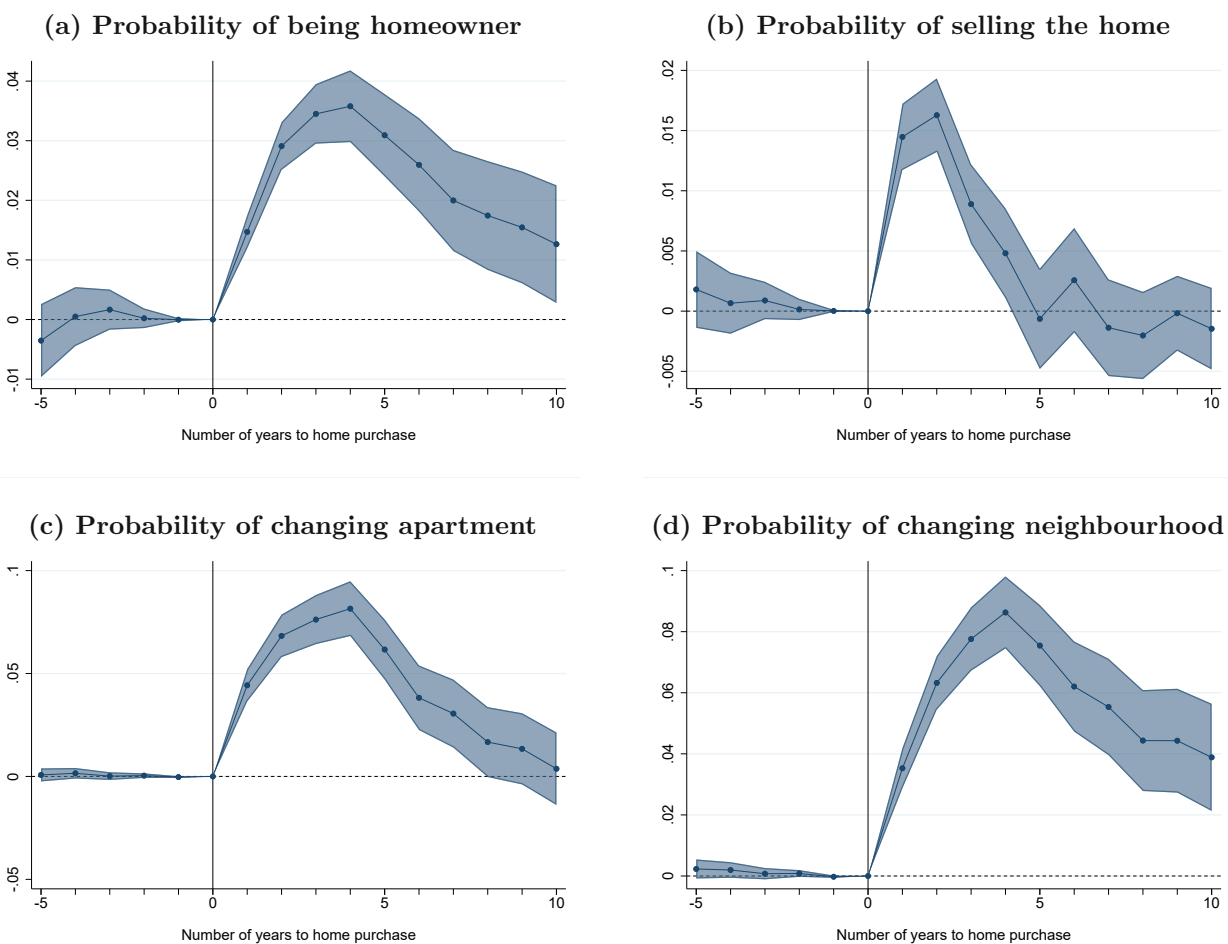


Figure C.3— Pre-trends coefficients and treatment effects: housing

Notes: Points on the graphs represent coefficients on $\text{treatment} \times \text{year}$ interactions, where treatment is defined as purchasing the home during the sales period (1997-2001). The sample is described in Panel B of Table 3.2. Column (a) measures the probability of being a homeowner, and (b) the probability of selling the home (proxied by a dummy equal to 1 if individuals change apartment and in the same year move to a new apartment where they are classified as owners). Column (c) captures the probability of changing apartment. Column (d) describes the probability of changing neighbourhood. We include individual and year FEs, and we control for the second and third power of age, the interaction of age and gender, and the interaction of age and highest completed education.

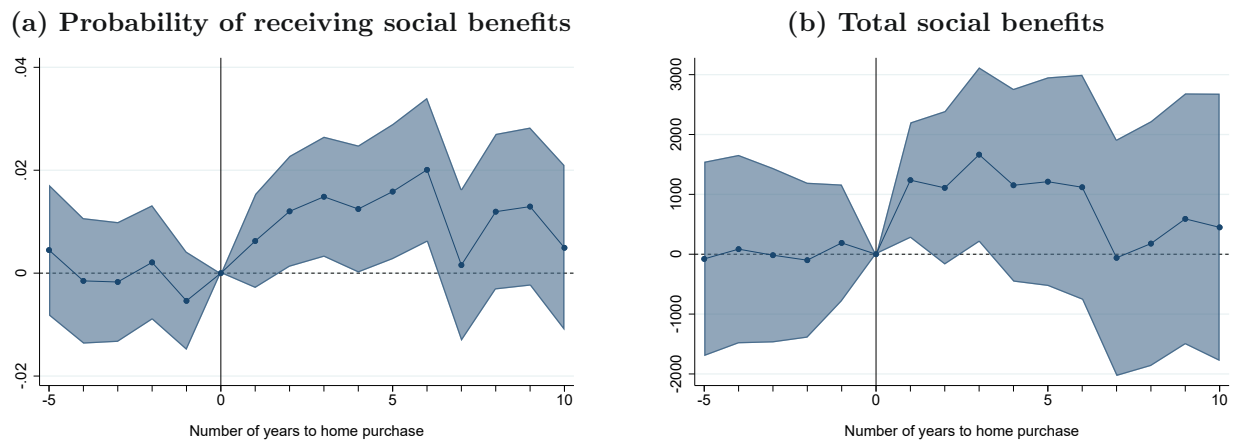


Figure C.4 — Pre-trends coefficients and treatment effects: public benefits

Notes: Points on the graphs represent coefficients on treatment \times year interactions, where treatment is defined as purchasing the home during the sales period (1997-2001). The sample is described in Panel B of Table 3.2. Column (a) measures the probability of receiving social benefits, and (b) the amount of social benefits received, measured in DKK. We include individual and year FEs, and we control for the second and third power of age, the interaction of age and gender, and the interaction of age and highest completed education.

Bibliography

- Aaronson, Daniel.** 2000. “A Note on the Benefits of Homeownership.” *Journal of Urban Economics* 47 (3): 356–369.
- Ackerberg, Daniel A., Kevin Caves, and Garth Frazer.** 2015. “Identification Properties of Recent Production Function Estimators.” *Econometrica* 83 (6): 2411–2451.
- Ahlfeldt, Gabriel M., and Elisabetta Pietrostefani.** 2019. “The economic effects of density: A synthesis.” *Journal of Urban Economics* 111 93–107.
- Algan, Yann, Arnaud Chéron, Jean-Olivier Hairault, and François Langot.** 2003. “Wealth effect on labor market transitions.” *Review of Economic Dynamics* 6 (1): 156–178.
- Amior.** 2024. “Education and Geographical Mobility: The Role of the Job Surplus.” *American Economic Journal: Economic Policy (Forthcoming)*.
- Amior, Michael, and Alan Manning.** 2018. “The Persistence of Local Joblessness.” *American Economic Review* 108 (7): 1942–1970.
- Anderson, Simon P., Andre de Palma, and Jacques-Francois Thisse.** 1992. *Discrete Choice Theory of Product Differentiation*. The MIT Press.
- Antoninis, Manos.** 2006. “The wage effects from the use of personal contacts as hiring channels.” *Journal of Economic Behavior and Organization* 59 (1): 133–146.
- Babet, Damien, Olivier Godechot, and Marco G. Palladino.** 2022. “Explaining the Dynamics of Wage Inequality in France.”

- Balgova, Maria.** 2018. “Why don’t less educated workers move? The role of job search in migration decisions.”
- Beine, Michel, Simone Bertoli, and Jesús Fernández-Huertas Moraga.** 2015. “A Practitioners’ Guide to Gravity Models of International Migration.” *The World Economy* 39 (4): 496–512.
- Bilal, Adrien.** 2023. “The Geography of Unemployment.” *The Quarterly Journal of Economics* 138 (3): 1507–1576.
- Blanchard, and Katz.** 1992. “Regional Evolutions.” *Brookings Papers on Economic Activity* 1.
- Bleakley, Hoyt, and Jeffrey Lin.** 2012. “Thick-market effects and churning in the labor market: Evidence from US cities.” *Journal of Urban Economics* 72 (2-3): 87–103.
- Blumenstock, Joshua E, Guanghua Chi, and Xu Tan.** 2023. “Migration and the Value of Social Networks.” *The Review of Economic Studies* 1–32.
- Boehm, Thomas P., and Alan M. Schlottmann.** 1999. “Does Home Ownership by Parents Have an Economic Impact on Their Children?” *Journal of Housing Economics* 8 (3): 217–232.
- Bound, John, and Harry J. Holzer.** 2000. “Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s.” *Journal of Labor Economics* 18 (1): 20–54.
- Bourassa, Steven C., Donald R. Haurin, and Martin Hoesli.** 2015. “What affects children’s outcomes: house characteristics or homeownership?” *Housing Studies* 31 (4): 427–444.
- Brown, Jeffrey R., Courtney C. Coile, and Scott J Weisbenner.** 2010. “The Effect of Inheritance Receipt on Retirement.” *Review of Economics and Statistics* 92 (2): 425–434.
- Buggle, Johannes, Thierry Mayer, Seyhun Orcan Sakalli, and Mathias Thoenig.** 2023. “The Refugee’s Dilemma: Evidence from Jewish Migration out of Nazi Germany.” *The Quarterly Journal of Economics* 138 (2): 1273–1345.

- Burgess, Simon, Julia Lane, and David Stevens.** 2000. "Job Flows, Worker Flows, and Churning." *Journal of Labor Economics* 18 (3): 473–502.
- Caldwell, and Harmon.** 2019. "Outside Options, Bargaining, and Wages: Evidence from Coworker Networks."
- Caldwell, Sydnee, and Oren Danieli.** 2024. "Outside Options in the Labour Market." *Review of Economic Studies* 1–41.
- Cameron, A. Colin, and Pravin K. Trivedi.** 2005. *Microeconometrics: Methods and Applications*. Cambridge University Press.
- Cesarini, David, Erik Lindqvist, Matthew J. Notowidigdo, and Robert Östling.** 2017. "The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries." *American Economic Review* 107 (12): 3917–3946.
- Chetty, Raj, and Nathaniel Hendren.** 2018. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." *The Quarterly Journal of Economics* 133 (3): 1107–1162.
- Chyn, Eric.** 2018. "Moved to Opportunity: The Long-Run Effects of Public Housing Demolition on Children." *American Economic Review* 108 (10): 3028–3056.
- Combes, Pierre-Philippe.** 2000. "Economic Structure and Local Growth: France, 1984 - 1993." *Journal of Urban Economics* 47 (3): 329–355.
- Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon.** 2019. "The Costs of Agglomeration: House and Land Prices in French Cities." *The Review of Economic Studies* 86 (4): 1556–1589.
- Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sebastien Roux.** 2012. "The Productivity Advantages of Large Cities: Distinguishing Agglomeration From Firm Selection." *Econometrica* 80 (6): 2543–2594.
- Combes, Pierre-Philippe, and Laurent Gobillon.** 2015. "Chapter 5 - The Empirics of Agglomeration Economies." In *Handbook of Regional and Urban Economics*, edited by

- Duranton, Gilles, J. Vernon Henderson, and William C. Strange Volume 5. of Handbook of Regional and Urban Economics 247–348, Elsevier.
- Comin, Diego, and Sunil Mulani.** 2006. “Diverging Trends in Aggregate and Firm Volatility.” *The Review of Economics and Statistics* 88 (2): 374–383.
- Comin, Diego, and Thomas Philippon.** 2006. “The Rise in Firm-Level Volatility: Causes and Consequences.” In *NBER Macroeconomics Annual 2005, Volume 20*, 167–228, National Bureau of Economic Research.
- Cuñat, Alejandro, and Marc Melitz.** 2012. “Volatility, labor market flexibility, and the pattern of comparative advantage.” *Journal of the European Economic Association* 10 (2): 225–254.
- Danish Economic Council.** 2023. “Danish Economy, Autumn 2023.”
- Dauth, Wolfgang, Sebastian Findeisen, Enrico Moretti, and Jens Suedekum.** 2022. “Matching in Cities.” *Journal of the European Economic Association* 20 (4): 1478–1521.
- Davis, Donald, Eric Mengus, and Tomasz Michalski.** 2020. “Labor Market Polarization and The Great Urban Divergence.” *NBER Working Paper Series* 26955.
- Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda.** 2007. “Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms.” In *NBER Macroeconomics Annual 2006, Volume 21*, 107–180, National Bureau of Economic Research.
- Davis, Steven J, John Haltiwanger, Ron Jarmin, Javier Miranda, Christopher Foote, and Eva Nagypal.** 2006. “Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms.” *NBER Macroeconomics Annual* 21.
- Daysal, N. Meltem, Michael Lovenheim, and David Wasser.** 2022. “The Intergenerational Transmission of Housing Wealth.” *NBER Working Paper Series* 31669.
- Di Pasquale, Denise, and Edward L. Glaeser.** 1999. “Incentives and Social Capital: Are Homeowners Better Citizens?” *Journal of Urban Economics* 45 (2): 354–384.

- Diamond, Rebecca.** 2016. “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000.” *American Economic Review* 106 (3): 479–524.
- Diamond, Rebecca, and Enrico Moretti.** 2021. “Where is Standard of Living the Highest? Local Prices and the Geography of Consumption.” *NBER Working Paper Series* 29533.
- Dietz, Robert D., and Donald R. Haurin.** 2003. “The social and private micro-level consequences of homeownership.” *Journal of Urban Economics* 54 (3): 401–450.
- van Dijk, Winnie.** 2020. “The socio-economic consequences of housing assistance.”
- Disney, Richard, and John Gathergood.** 2017. “House Prices, Wealth Effects and Labour Supply.” *Economica* 85 (339): 449–478.
- Duranton, Gilles.** 2007. “Urban Evolutions: The Fast, the Slow, and the Still.” *American Economic Review* 97 (1): 197–221.
- Duranton, Gilles, and Diego Puga.** 2004. “Micro-foundations of urban agglomeration economies.” In *Handbook of Regional and Urban Economics*, edited by Henderson, J. V., and J. F. Thisse Volume 4. of *Handbook of Regional and Urban Economics*, Chap. 48 2063–2117, Elsevier.
- Eriksen, Michael D., and Stuart S. Rosenthal.** 2010. “Crowd out effects of place-based subsidized rental housing: New evidence from the LIHTC program.” *Journal of Public Economics* 94 (11): 953–966.
- Findeisen, Sebastian, and Jens Südekum.** 2008. “Industry churning and the evolution of cities: Evidence for Germany.” *Journal of Urban Economics* 64 (2): 326–339.
- Gaigne, Carl, and Mathieu Sanch-Maritan.** 2019. “City size and the risk of being unemployed. Job pooling vs. job competition.” *Regional Science and Urban Economics* 77 222–238.

- Gagné, Carl, Hans R.A. Koster, Fabien Moizeau, and Jacques-François Thisse.** 2022. “Who lives where in the city? Amenities, commuting and income sorting.” *Journal of Urban Economics* 128 2–27.
- Gan, Li, and Qinghua Zhang.** 2006. “The thick market effect on local unemployment rate fluctuations.” *Journal of Econometrics* 133 (1): 127–152.
- Gaubert, Cecile.** 2018. “Firm sorting and agglomeration.” *American Economic Review* 108 (11): 3117–3153.
- Gharad, Bryan, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak.** 2014. “Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh.” *Econometrica* 82 1671–1748.
- di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Mejean.** 2014. “Firms, Destinations, and Aggregate Fluctuations.” *Econometrica* 82 (4): 1303–1340.
- Glaeser, Edward L., Hedi D. Kallal, José A. Scheinkman, and Andrei Shleifer.** 1992. “Growth in Cities.” *Journal of Political Economy* 100 (6): 1126–1152.
- Glaeser, Edward L., and Erzo F.P. Luttmer.** 2003. “The misallocation of housing under rent control.” *American Economic Review* 93 (4): 1027 – 1046.
- Granovetter, Mark S.** 1973. “The Strength of Weak Ties.” *American Journal of Sociology* 78 (6): 1360–1380.
- Green, Richard K., and Michelle J. White.** 1997. “Measuring the Benefits of Home-owning: Effects on Children.” *Journal of Urban Economics* 41 (3): 441–461.
- Gregg, Paul, Stephen Machin, and Alan Manning.** 2004. *Mobility and Joblessness*. 371–410, University of Chicago Press.
- Haltiwanger, John C., Mark J. Kutzbach, Giordano E. Palloni, Henry Pollakowski, Matthew Staiger, and Daniel Weinberg.** 2020. “The Children of HOPE VI Demolitions: National Evidence on Labor Market Outcomes.” *NBER Working Paper Series* 28157.

- Haurin, Donald R., Toby L. Parcel, and R. Jean Haurin.** 2002. “Does Homeownership Affect Child Outcomes?” *Real Estate Economics* 30 (4): 635–666.
- Hausman, Naomi, Tamar Ramot-Nyska, and Noam Zussman.** 2022. “Homeownership, Labor Supply, and Neighborhood Quality.” *American Economic Journal: Economic Policy* 14 (2): 193–230.
- Heise, Sebastian, and Tommaso Porzio.** 2023. “Labor Misallocation Across Firms and Regions.” NBER Working Papers, updated 28792, National Bureau of Economic Research.
- Helsley, Robert W., and William C. Strange.** 1990. “Matching and agglomeration economies in a system of cities.” *Regional Science and Urban Economics* 20 (2): 189–212.
- Henderson, V, A Kuncoro, and M Turner.** 1995. “Industrial Development in Cities.” *Journal of Political Economy* 103 (5): 1067–1090.
- Henley, Andrew.** 2004. “House Price Shocks, Windfall Gains and Hours of Work: British Evidence.” *Oxford Bulletin of Economics and Statistics* 66 (4): 439–456.
- Hensvik, Lena, and Oskar Nordström Skans.** 2016. “Social Networks, Employee Selection, and Labor Market Outcomes.” *Journal of Labor Economics* 34 (4): 825–867.
- Hoynes, Hilary.** 2000. *Finding Jobs: Work and Welfare Reform*. Chap. The Employment and Earnings of Less Skilled Workers Over the Business Cycle 23–71, Russell Sage Foundation: New York.
- Inbens, Guido W., Donald B. Rubin, and Bruce I. Sacerdote.** 2001. “Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players.” *American Economic Review* 91 (4): 778–794.
- Jacob, Brian A., Max Kapustin, and Jens Ludwig.** 2014. “The Impact of Housing Assistance on Child Outcomes: Evidence from a Randomized Housing Lottery.” *The Quarterly Journal of Economics* 130 (1): 465–506.
- Jacob, Brian A., and Jens Ludwig.** 2012. “The Effects of Housing Assistance on Labor Supply: Evidence from a Voucher Lottery.” *American Economic Review* 102 (1): 272–304.

- Joulfaian, David, and Mark O. Wilhelm.** 1994. "Inheritance and Labor Supply." *The Journal of Human Resources* 29 (4): 1205–1234.
- Kennan, John.** 2015. "Spatial Variation in Higher Education Financing and the Supply of College Graduates." *NBER Working Paper Series* 21065.
- Kennan, John, and James R. Walker.** 2011. "The Effect of Expected Income on Individual Migration Decisions." *Econometrica* 79 (1): 211–251.
- Kline, Patrick, and Enrico Moretti.** 2013. "Place Based Policies with Unemployment." *American Economic Review* 103 (3): 238–243.
- Krugman, Paul.** 1992. *Geography and Trade*. Volume 1. of MIT Press Books, The MIT Press.
- Kuhn, Moritz, Iourii Manovskii, and Xincheng Qiu.** 2021. "The Geography of Job Creation and Job Destruction." *NBER Working Paper Series* 29399.
- Li, Han, Jiangyi Li, Yi Lu, and Huihua Xie.** 2020. "Housing wealth and labor supply: Evidence from a regression discontinuity design." *Journal of Public Economics* 183 1–28.
- Manski, Charles F.** 2000. "Economic Analysis of Social Interactions." *Journal of Economic Perspectives* 14 (3): 115–136.
- Mayer, Thierry, Isabelle Mejean, and Benjamin Nefussi.** 2010. "The location of domestic and foreign production affiliates by French multinational firms." *Journal of Urban Economics* 68 (2): 115–128.
- Melitz, Marc J.** 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–1725.
- Moffitt, Robert A.** 2001. *Policy Interventions, Low-Level Equilibria, and Social Interactions*. 45–82, The MIT Press.
- Mohanty, Lisa L., and Lakshmi K. Raut.** 2009. "Home Ownership and School Outcomes of Children: Evidence from the PSID Child Development Supplement." *The American Journal of Economics and Sociology* 68 (2): 465–489.

- Monras.** 2020. “Economic Shocks and Internal Migration.” *CEPR Discussion Papers* 12977.
- Moon, Choon-Geol, and Janet G Stotsky.** 1993. “The effect of rent control on housing quality change: a longitudinal analysis.” *Journal of Political Economy* 101 (6): 1114–1148.
- Moretti, Enrico.** 2011. “Chapter 14 - Local Labor Markets.” In *Handbook of Labor Economics*, edited by Card, David, and Orley Ashenfelter Volume 4. 1237–1313, Elsevier.
- Mortensen, Dale T., and Christopher A. Pissarides.** 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *The Review of Economic Studies* 61 (3): 397–415.
- Munch, Jakob R., Michael Rosholm, and Michael Svarer.** 2006. “Are Homeowners Really More Unemployed?.” *The Economic Journal* 116 (514): 991–1013.
- Munshi, Kaivan.** 2020. “Social Networks and Migration.” *Annual Review of Economics* 12 (1): 503–524.
- Nekoei, Arash, and Andrea Weber.** 2020. “Seven Facts about Temporary Layoffs.” cepr discussion papers, C.E.P.R. Discussion Papers.
- OECD.** 2005. “How Persistent Are Regional Disparities in Employment? The Role of Geographic Mobility.” Technical report, OECD Employment Outlook.
- OECD.** 2021. *Brick by Brick: Building Better Housing Policies*. OECD.
- Oswald, Andrew J.** 1996. “A conjecture on the explanation for high unemployment in the industrialized nations: Part I.” *The Warwick Economics Research Paper Series* 475.
- Overman, Henry G., and Diego Puga.** 2010. “Labor Pooling as a Source of Agglomeration: An Empirical Investigation.” In *Agglomeration Economics*, 133–150, National Bureau of Economic Research.
- Papageorgiou, Theodore.** 2022. “Occupational Matching and Cities.” *American Economic Journal: Macroeconomics* 14 (3): 82–132.

- Patnaik, Arpita, Joanna Venator, Matthew Wiswall, and Basit Zafar.** 2022. “The role of heterogeneous risk preferences, discount rates, and earnings expectations in college major choice.” *Journal of Econometrics* 231 (1): 98–122.
- Pollakowski, Henry O., Daniel H. Weinberg, Fredrik Andersson, John C. Haltiwanger, Giordano Palloni, and Mark J. Kutzbach.** 2022. “Childhood Housing and Adult Outcomes: A Between-Siblings Analysis of Housing Vouchers and Public Housing.” *American Economic Journal: Economic Policy* 14 (3): 235–272.
- Porcher, Charly.** 2022. “Migration with Costly Information.”
- Saiz, Albert.** 2023. “The Global Housing Affordability Crisis: Policy Options and Strategies.” *IZA Policy Paper* 203.
- Saygin, Perihan Ozge, Andrea Weber, and Michèle A. Weynandt.** 2021. “Coworkers, Networks, and Job-Search Outcomes among Displaced Workers.” *ILR Review* 74 95–130.
- Schluter, Christian, and Guillaume Willems.** 2019. “A Dynamic Empirical Model of Frictional Spatial Job Search.”
- Schmutz, Benoit, and Modibo Sidibé.** 2019. “Frictional labour mobility.” *Review of Economic Studies* 86 (4): 1779–1826.
- Schmutz, Benoit, Modibo Sidibé, and Elie Vidal-Naquet.** 2021. “Why are low-skilled workers less mobile? The role of mobility costs and spatial frictions.” *Annals of Economics and Statistics* 142 283–303.
- Schwartz, Amy E., Ingrid G. Ellen, Michael H. Schill, and Ioan Voicu.** 2006. “The external effects of place-based subsidized housing.” *Regional Science and Urban Economics* 36 (6): 679–707.
- Shlay, Anne B.** 2006. “Low-income Homeownership: American Dream or Delusion?” *Urban Studies* 43 (3): 511–531.
- Skandalis, Daphne.** 2019. “Breaking News: Information About Firms’ Hiring Needs Affects the Direction of Job Search.”

- Sodini, Paolo, Stijn Van Nieuwerburgh, and Roine Vestman.** 2023. "Identifying the Benefits from Homeownership: A Swedish Experiment." *The American Economic Review* 113 (12): 3173–3212.
- Sorkin, Isaac.** 2018. "Ranking Firms Using Revealed Preference." *The Quarterly Journal of Economics* 133 (3): 1331–1393.
- Topa, Giorgio.** 2001. "Social interactions, local spillovers and unemployment." *Review of Economic Studies* 68 (2): 261–295.
- Topa, Giorgio.** 2011. "Chapter 22 - Labor Markets and Referrals." Volume 1. of Handbook of Social Economics 1193–1221, North-Holland.
- Topel, Robert H.** 1986. "Local Labor Markets." *Journal of Political Economy* 94 (3, Part 2): S111–S143.
- Train, Kenneth E.** 2001. *Discrete Choice Methods with Simulation*. Cambridge University Press.
- Tsenkova, Sasha.** 2021. "Resilience of social housing systems in Vienna, Amsterdam, and Copenhagen." *Cities and Affordable Housing: Planning, Design and Policy Nexus* 171–193.
- Wainer, Allison, and Jeffrey Zabel.** 2020. "Homeownership and wealth accumulation for low-income households." *Journal of Housing Economics* 47 1–20.
- Weingarden, Alison.** 2020. "Worker Churn at Establishments over the Business Cycle." notes, Federal Reserve.
- Wilson, Riley.** 2021. "Moving to Jobs: The Role of Information in Migration Decisions." *Journal of Labor Economics* 39 (4): 1083–1128.
- Wozniak, Abigail.** 2010. "Are College Graduates More Responsive to Distant Labor Market Opportunities?" *Journal of Human Resources* 45 (4): 944–970.
- Zabek, Mike.** 2023. "Local Ties in Spatial Equilibrium." *American Economic Journal: Macroeconomics* 16 (2): 287–317.

Zerecero, Miguel. 2021. "The birthplace premium."

Titre : Essais en géographie économique et économie urbaine

Mots clés : marchés locaux du travail, économies d'agglomération, localisation des entreprises, mobilité des travailleurs, coûts de mobilité, accessibilité du logement

Résumé : Cette thèse étudie les décisions de localisation des entreprises et des travailleurs, ainsi que la manière dont celles-ci interagissent avec les caractéristiques locales des marchés du travail. Le premier chapitre se concentre sur les entreprises et s'intéresse à un nouveau mécanisme qui les incite à s'implanter dans les villes les plus denses : le rôle de la volatilité de la demande et de son interaction avec la productivité des entreprises. Ce processus tient à des conditions d'embauche plus fluides qui attirent les entreprises productives pouvant plus rapidement réduire ou augmenter leurs effectifs dans les villes

les plus densément peuplées. Le deuxième chapitre s'intéresse aux choix de localisation des travailleurs et à la manière dont les coûts de mobilité, en particulier les frictions informationnelles, affectent les migrations régionales. Cela permet de mettre en lumière les mécanismes à l'origine du biais migratoire lié aux compétences, à savoir la régularité empirique selon laquelle les travailleurs hautement qualifiés sont significativement plus mobiles que les travailleurs peu qualifiés. Le troisième chapitre porte sur l'interaction des politiques de logement abordable avec les incitations à participer au marché du travail.

Title : Essays in economic geography and urban economics

Keywords : local labor markets, agglomeration economies, firm location, worker mobility, mobility costs, housing affordability

Abstract : This thesis studies spatial location decisions of firms and workers, and how these interact with local labor market characteristics. The first chapter focuses on firms and explores a novel mechanism that incentivizes firms to locate in denser cities: the role of volatile demand and its interaction with firm productivity. This channel arises since faster hiring conditions in thicker labor markets attract productive firms that can more swiftly downsize or expand in denser cities. The second chapter explores the location decisions of workers and how regional migration is affected by mobility costs, in particular information fric-

tions. This helps shed light on the mechanisms driving skill-biased migration, namely the empirical regularity that high-skilled workers are substantially more mobile than low-skilled workers. The third chapter studies the interaction of affordable housing policies with incentives for labor market participation. A quasi-natural experimental setting enables to analyze a large public housing privatization event in the city of Copenhagen directed towards low-income households, and to compare the impact of subsidized home purchase versus subsidized rental on long-run labor market outcomes.