UNIVERSITY OF NAPLES FEDERICO II

DOCTORAL THESIS

Labor Market Flexibility and Life-Cycle Decisions: Three Essays on Work, Retirement, and Family Behavior

Author
Nicol BARBIERI

Supervisor
Prof. Monica LANGELLA

PhD Director
Prof. Marco PAGNOZZI

This thesis is submitted in fulfillment of the requirements for the degree of Doctor of Philosophy in Economics - XXXVII Cycle

Naples School of Economics
Department of Economics and Statistical Sciences

Declaration of Authorship

I declare that this thesis and the work presented within it are entirely my own, generated as a result of my own research. I submit this thesis in fulfillment of the requirements for the degree of Doctor of Philosophy in Economics, XXXVII Cycle, at the Naples School of Economics, Department of Economics and Statistical Sciences, University of Naples Federico II.

To the best of my knowledge, this thesis contains no material previously published by any other individual, except where proper acknowledgment has been made. Furthermore, this thesis does not include any material that has been accepted for any other academic degree or non-degree program, in English or any other language.

I also confirm that I have clearly attributed all published works consulted during my research. Whenever I have quoted from the work of others, the source has been properly cited.

In cases where this thesis is based on work done collaboratively with other people, I have made exactly what was done by others and what I have contributed myself.

To my sister Irma,

for walking beside me and lighting my way with love.

Acknowledgements

During my PhD, I have accumulated a deep debt of gratitude to many people whose help and encouragement have been essential in reaching this stage. These few lines are a small way to acknowledge, at least in words, a part of that debt.

First and foremost, I am deeply grateful to Monica Langella, my advisor, for her constant guidance, generosity, and support throughout these years. Her insight and patience have been a constant source of inspiration and have guided me closely in my journey toward becoming a researcher, profoundly shaping both this dissertation and the way I see and pursue my work. Each of our discussions felt like fitting together pieces of a larger puzzle that she helped me arrange with clarity and purpose, allowing me to grow both intellectually and personally.

I am also grateful to all the faculty members at the Centre for Studies in Economics and Finance and at the Department of Economics and Statistics of the University of Naples Federico II for creating such a stimulating academic environment, where their research, seminars, and teaching have greatly enriched my academic journey. I am especially thankful to Tullio Jappelli, Marco Pagnozzi, and Marco Pagano for their guidance and constant presence.

Among them, a special thanks goes to Annalisa Scognamiglio for her invaluable comments, enthusiasm, and mentorship, which have supported me since well before the beginning of my PhD. Her continuous and genuine interest in both my work and personal growth has been truly motivating and has deeply shaped my experience over these years.

I am sincerely thankful to my PhD colleagues and research assistants for sharing ideas, challenges, and laughter both inside and outside academia. Their friendship and collaboration have made this experience not only stimulating but also truly meaningful. My heartfelt gratitude goes to Antonio, who has been there since the very beginning. His friendship has made these years brighter and more enjoyable. I am also thankful to Matteo for his encouragement and many engaging conversations we have shared along the way. I owe special thanks to Nello for being both a brilliant coauthor and an unfailing source of thoughtful guidance and genuine friendship. I feel lucky to have shared this experience with him, and grateful that part of this thesis was shaped together.

I am sincerely grateful to all my friends. Among them, a very special thanks goes to Simona, who has always been by my side in difficult times, even while facing her own, and who, as has long been the case, continues to share every moment of joy with me.

Finally, I owe my deepest gratitude to my parents, Michele and Maria, for supporting every dream I have ever had and for always making them possible. I am endlessly grateful to my sister, Irma, to whom I dedicate this thesis, for being my biggest fan and for never stopping believing in me and in my abilities. Her love and trust have always given me, and continue to give me, the strength to keep going.

Abstract

This dissertation studies how labor market flexibility influences individual and family decisions across the life cycle. Through three self-contained essays, it analyzes the effects of remote work, retirement, and family choices on labor supply, financial behavior, and fertility.

The first essay, "Work from Home, Quit Behavior, and Monopsony Power: Evidence from the Italian Labor Market" examines how the rapid expansion of remote work during the COVID-19 pandemic transformed job mobility and wage dynamics. Using rich administrative employer-employee data, the analysis shows that occupations with higher potential for remote work experienced a marked decline in voluntary quits after 2020 and a moderation in wage growth. Estimates from the quit elasticity further indicate that the wage elasticity of quits increased, suggesting that remote work altered the trade-off between wage compensation and flexibility, with potential implications for employers' monopsony power.

The second essay, "The Impact of Retirement on Portfolio Allocation: Evidence from Italian Households" - a joint work with Nello Esposito (Ph.D. student, University of Naples Federico II) - studies how retirement affects individuals' financial behavior using panel data from the Bank of Italy's Survey on Household Income and Wealth. Exploiting exogenous variation in pension eligibility rules as an instrument for retirement, the study finds that retirement entails a considerable reallocation of individuals' portfolios toward a significantly higher share of financial assets allocated to stocks. These effects are more substantial among individuals retiring through seniority schemes, consistent with the wealth and income-stability mechanisms predicted by life-cycle and portfolio-choice models.

The third essay, "Making Room for a Baby: The Impact of Working from Home on Fertility Decisions" explores whether the diffusion of remote work affected women's fertility behavior after the pandemic. Combining survey data with occupation-based teleworkability measures, the analysis shows that women in teleworkable occupations were significantly more likely to be on maternity leave after 2020. A two-stage least squares framework confirms that this relationship is not driven by selection into remote work but by the exogenous feasibility of performing tasks from home. The results suggest that increased flexibility in work arrangements reduced the opportunity cost of childbearing, particularly for highly educated and high-income women.

Overall, this dissertation highlights the role of labor market flexibility in shaping economic behavior across key life stages - from employment stability to retirement planning and family formation. By linking labor market institutions with household and individual outcomes, it contributes to an understanding of how the changing nature of work interacts with economic decision-making.

Contents

List of Figures

List of Tables

Chapter 1

Work from Home, Quit Behavior, and Monopsony Power: Evidence from the Italian Labor Market

Nicol Barbieri *

This paper studies the impact of Working from Home (WFH) on voluntary quits in the context of the Italian labor market. It investigates how and to what extent the adoption of WFH during the COVID-19 pandemic shaped workers' quitting behavior and how this translated into changes in employers' market power. Using very rich administrative employer-employee data, I implement a Difference-in-Differences (DiD) analysis that exploits occupational differences in WFH feasibility to identify the causal effect of WFH on workers' likelihood of voluntarily quitting and their wages. Then, I investigate whether this shift translates into changes in workers' voluntary quit elasticity. My results show that workers in teleworkable occupations experienced a decline in voluntary quits and relatively lower wages, suggesting that WFH can act as a non-wage amenity that supports job attachment. Moreover, while workers' quit elasticity declined in the post-pandemic period, this reduction was less pronounced for those in teleworkable occupations. These findings suggest that WFH mitigated the loss of wage responsiveness that characterized the post-pandemic labor market, partly protecting workers' outside options while reinforcing attachment through non-monetary amenities.

* I thank my PhD advisor, Monica Langella, for her guidance, insightful feedback, and continuous support throughout this project. This work has also significantly benefited from the invaluable discussions with Annalisa Scognamiglio. I am also grateful to Chiara Lacava and Valeria Zurla for their helpful comments and encouragement. I learned enormously from informal conversations with Nello Esposito, Martina Peluso, Matteo Moracci, Matteo Vasca, and Antonio Genovese. Special thanks to all the participants of the 20th Doctorissimes at PSE, especially David Margolis, for his priceless suggestions, the 2nd ECWE at the University of Verona, the HEC Paris Economics PhD Conference, and the Law, Markets and Public Policy Workshop for their helpful comments. Many thanks also to all participants at the Royal Economic Society (RES) 2025 Annual Conference, the 40th AIEL Annual Conference, and the Remote Work Conference at Stanford University, for their inestimable advice. I am also indebted to the PhD CSEF Seminar members at the University of Naples Federico II for their sustained guidance, and to those who participated in the NSFE 2024 Summer School. I obtained the data from the Italian Ministry of Labor and Social Policies. This work has also been funded under grant Prin PNRR, P2022KHP8L, "Inequalities and the labor markets".

1.1 Introduction

The COVID-19 pandemic led to a significant and unexpected global shift toward Work from Home (WFH) arrangements (see, e.g., Zarate et al., 2024; Aksoy et al., 2022; Ramani and Bloom, 2021; Barrero et al., 2020). Surprisingly, this shift has continued to persist even in the years following the pandemic (see, for instance, Aksoy et al., 2025; Ketter et al., 2025; Buckman et al., 2025; Bick et al., 2023; Barrero et al., 2023; Barrero et al., 2021b), suggesting that the pandemic-induced changes may have had long-lasting effects on labor market dynamics, both from workers' and employers' perspectives (see, for example, Akan et al., 2025; Autor et al., 2023; De Fraja et al., 2022).

On the workers' side, WFH has the potential to make labor markets less local than before and to expand job opportunities beyond geographical constraints, thereby increasing worker mobility and competition and shaping how firms recruit and retain talent. At the same time, WFH can be considered as a non-wage amenity, as workers may highly value the flexibility and reduced commuting constraints it provides, which may strengthen their attachment to the current employer (see, e.g., Mas, 2025; Cullen et al., 2025; Sockin, 2022; Mas and Pallais, 2017). Depending on how it is implemented - fully or in hybrid form - WFH is also often associated with productivity gains through time savings and better allocation of working hours (see Choudhury et al., 2024; Angelici and Profeta, 2024; Bloom et al., 2024, Bloom et al., 2022; Battiston et al., 2021; Barrero et al., 2021a; Emanuel and Harrington, 2021, Bloom et al., 2015).

However, these benefits are not universal. For some workers, WFH can act as a disamenity, blurring the boundaries between work and personal life and limiting interpersonal interactions and knowledge spillovers (see, e.g., Emanuel and Harrington, 2024; Gibbs et al., 2023; Bao et al., 2022; Atkin et al., 2022), suggesting that its net effect on labor mobility is not straightforward.

On the employers' side, the prospect of productivity gains may incentivize the adoption of WFH arrangements, positioning them as a potential long-term organizational strategy (see, e.g., Thorstensson, 2020). Firms may increasingly enable employees to work remotely, motivated not only by potential efficiency gains but also by lower office costs, higher employee satisfaction, and greater engagement. In this sense, WFH can also serve as a retention tool, helping firms strengthen workers' attachment by offering valued non-wage amenities.

Yet, many firms remain reluctant to fully embrace WFH, partly due to concerns about monitoring, collaboration, and organizational culture (see, for instance, Barrero et al., 2025; Flynn et al., 2024; Bloom et al., 2023).

This heterogeneity in responses suggests that WFH may have far-reaching implications for employer-employee relationships and for broader labor market dynamics, as differences in preferences, mobility costs, and the valuation of job amenities ultimately shape how wages are set and how easily workers move across jobs (see, for instance, Autor et al., 2023; Barrero et al., 2023; Kahn and Tracy, 2023; De Fraja et al., 2022; Brinatti et al., 2021).

In frictional labor markets, heterogeneous preferences, commuting time, and job-search constraints generate distortions that can grant market power to employers (see, e.g., Azar and Marinescu, 2024; Manning, 2021; Webber, 2015; Manning, 2003). In other words, workers' differing valuations of job amenities (e.g., short commuting or flexible working conditions) are not fully priced into wages, allowing employers to set pay below the competitive benchmark. These mobility frictions make labor supply to firms not perfectly elastic, as workers face costs and constraints in changing jobs. This, in turn, creates wedges between workers' marginal product and their earnings, providing firms with scope to exercise monopsony power.

Within this framework, WFH defines a unique workplace arrangement, as it has the potential to relax labor market frictions by reducing mobility constraints and expanding workers' outside options. In particular, WFH may remove geographical barriers, enabling workers to consider jobs beyond their local labor markets. By lowering these mobility costs, WFH can increase the elasticity of labor supply to the firm, as workers become more responsive to wage differentials across employers. These extended outside options could therefore weaken monopsony power, especially in sectors and occupations where WFH is more prevalent.

Nonetheless, the flexibility associated with WFH could reduce turnover, as workers may be less inclined to quit jobs that offer non-wage benefits as a substitute for monetary compensation (see, for example, Mas, 2025; Babet and Chabaud, 2024; Lavetti, 2023; Mas and Pallais, 2017). In this sense, WFH can simultaneously strengthen worker attachment to employers while also expanding outside options, thereby reshaping the balance of employer-employee relationships and bargaining power.

This dual mechanism highlights two different theoretical predictions: while increased outside options could raise the elasticity of quitting to the firm, greater flexibility may lower voluntary separation rates, increasing workers' attachment to their employer.

Though increasing evidence points to positive effects of WFH on workers' well-being and productivity (see, among others, Bloom et al., 2024; Angelici and Profeta, 2024; Choudhury et al., 2024), little is known about how WFH influences workers' decisions to quit - a key indicator of mobility and competition in labor markets. Moreover, the impact of WFH on wages and its implications for employers' market power remain only partially understood (see, on this point, Autor et al., 2023). This question has become even more relevant in light of the phenomenon known as the *Great Resignation* or *Big Quit* (see, for example, Lee et al., 2023; Forsythe et al., 2022; Gittleman, 2022), which signals a shift in workers' preferences and priorities, with WFH opportunities playing a critical role. In the Italian context, Casarico and Lattanzio, 2022 shows that immediately after the pandemic, job separations decreased due to job-retention schemes and mobility restrictions that helped prevent layoffs. However, as restrictions were lifted, job mobility resumed, and reallocation disproportionately occurred toward sectors with higher WFH feasibility, consistent with a sectoral reallocation effect (see, e.g., Basso et al., 2023). This process was not uniform across industries, as sectors differed in the extent to which tasks could be adapted to remote settings. Moreover, workers who voluntarily left their jobs exhibited significantly higher re-employment probabilities than before the pandemic, with an increasing share of transitions directed toward industries better able to accommodate remote tasks (see, for instance, Gómez and Lattanzio, 2024).

While these studies have primarily examined sectoral reallocation and workers' re-employment probabilities, they have not analyzed how WFH explicitly interacts with workers' quitting behavior, wage dynamics, and voluntary quit elasticity.

This paper addresses these gaps by evaluating how WFH shapes both workers' voluntary separation decisions and their wages in the Italian labor market, providing the first systematic analysis of quits and wages jointly, while making predictions on how these adjustments translate into changes in employers' market power. By leveraging a unique combination of administrative data and task-based classifications of "teleworkability", i.e., the degree to which each occupation can be done from home, as commonly measured in the literature (see, for example, Barrero et al., 2023; Basso et al., 2022; Barbieri et al., 2021; Dingel, 2020; Hensvik et al., 2020), I quantify the relationship between WFH and workers' quitting behavior and wages. Then, I investigate whether WFH mitigates employers' market power, studying how WFH shapes workers' quitting elasticity in response to wage variations.

To do so, I implement a Difference-in-Differences (DiD) analysis, where I compare workers' quitting probability and wages in occupations defined as "teleworkable", i.e., occupations more suitable for remote options, with those less conducive to working remotely. The key advantage of this approach is that the teleworkability measure is based on pre-pandemic job-task characteristics, which reflect the intrinsic feasibility of WFH rather than subsequent adjustments in work practices following the COVID-19 shock. This ensures that the variation I exploit - i.e., the occupational differences in WFH feasibility - is predetermined and plausibly exogenous to the COVID-19 event, under the assumption that, absent the pandemic, quits and wages would have evolved in parallel across teleworkable and non-teleworkable jobs.

My findings suggest that working remotely has a significant impact on workers' quitting probability, with teleworkable occupations showing a lower likelihood of voluntary quitting post-pandemic compared to non-teleworkable occupations. In particular, a 1 percent increase in WFH is associated with a 1 percentage point decrease in the probability of voluntary quitting. Given the average quit rate of 10 percent for the entire sample, WFH leads to a 10 percent reduction in voluntary quitting relative to the mean. This result points to WFH as a possible attachment channel, reducing separations by supplying flexibility and non-wage amenities that workers may highly value.

To better explore this potential mechanism, I first estimate the impact of WFH on wages. My results show that after the pandemic, occupations that are more easily switched to WFH experienced a reduction in daily wages of 1.5 percent. Then, I estimate the wage elasticity of quits. I find that quit elasticity becomes less negative after the pandemic, but the reduction is smaller for teleworkable occupations. This suggests that workers in more WFH-friendly jobs remained relatively more responsive to wage differences compared to those in non-teleworkable jobs, consistent with an expansion of their outside options.

Overall, these results suggest that WFH may serve as a non-monetary job amenity that partially substitutes for wage compensation in teleworkable occupations, while at the same time mitigating the post-pandemic decline in wage responsiveness and preserving workers' ability to respond to outside options.

I also explore heterogeneity in how WFH shaped quits and wages across workers' socio-economic characteristics. Men experienced lower reductions in quits as opposed to women, potentially implying that women workers benefited more strongly from the attachment effects of WFH. This could reflect the fact that women were disproportionately employed in occupations where remote work was more widely implemented, or that they placed a higher value on the job security and stability associated with flexible work arrangements during the pandemic. Moreover, women's quitting behavior may have been shaped by competing factors, such as increased caregiving responsibilities or greater constraints on mobility, which partially complemented the flexibility feature of WFH. By contrast, men faced a sharper wage decline than women, suggesting that in male-dominated teleworkable occupations, more substantial attachment effects came at the cost of more pronounced wage adjustments. One possible explanation can be rooted in compensating differentials: if men perceived the flexibility of WFH as less valuable than women, employers could moderate wages without triggering higher quits, since attachment remained strong. In other words, the higher the value attached to an amenity like WFH, the greater the scope for wage moderation. Another explanation relates to career concerns. If male workers in teleworkable occupations viewed WFH as compatible with long-term career advancement within the firm, the combination of WFH and perceived internal career opportunities may have further reduced their willingness to switch jobs. This reduced mobility, in turn, may have strengthened employers' ability to compress wages.

Age differences in quits are more pronounced for younger workers (16-35 years old), indicating that WFH increased mobility for older workers while reinforcing attachment among younger ones. The younger group benefited the most from WFH, attaching a higher value to its flexibility, which substantially reduced their voluntary mobility compared to the pre-pandemic period. This stronger attachment likely allowed employers to compress wages for younger cohorts, who typically face more outside options in regular times. By contrast, older workers experienced flatter effects, suggesting that WFH altered the behavior for younger cohorts, who are usually the most mobile in the labor market.

Education plays a central role: during the pandemic, voluntary quits declined more sharply among highly educated workers in teleworkable occupations, indicating that this group placed particular value on the flexibility offered by WFH. For lower-educated workers, quits also declined but to a smaller extent, while medium-educated workers exhibited only modest and transitory effects. When examining wages, however, the picture shifts: lower and medium educated workers in teleworkable jobs experienced a pronounced wage penalty by 2023, whereas higher-educated workers not only avoided this decline but also registered modest wage gains. This asymmetry potentially suggests that rents were disproportionately extracted from lower-educated workers, who had fewer outside options relative to their higher-educated.

Contract type also matters: temporary and full-time workers saw more substantial reductions in quits, reflecting the value they placed on security and flexibility, whereas open-ended and part-time workers remained more mobile. Employers could thus leverage WFH as an attachment tool, especially for more precarious workers.

Finally, regional patterns indicate that the stabilizing effect of WFH on quits was strongest in Northern Italy, where its adoption was more feasible, suggesting structural disparities across labor markets. This regional heterogeneity implies that the implications of WFH are context-dependent, amplifying attachment where remote adoption is higher.

Overall, my results suggest that the massive shock in WFH implementation induced by COVID-19 had a large effect on wages and the probability of voluntary quitting. In occupations where WFH is feasible, workers may prioritize the non-monetary benefits associated with remote work, such as flexibility and reduced commuting time, over wage increases. Employers may leverage this preference shift as attachment tool. Moreover, WFH can expand workers' outside options by relaxing geographic constraints.

My contribution to the existing literature unfolds along two main dimensions. First, I shift the focus from aggregate reallocation to individual quitting behavior, offering the first causal evidence on how WFH shaped voluntary separations in the Italian labor market. In contrast to Casarico and Lattanzio, 2022, who focuses on the immediate decline in separations driven by job-attachment schemes, I examine the longer-run adjustment in quits once restrictions were lifted, documenting how WFH affected separations beyond policy-driven dynamics. Unlike Basso et al., 2023, who emphasizes aggregate sectoral reallocation, I link quits directly to wage dynamics, uncovering how WFH influences both mobility and compensation. Moreover, while Gómez and Lattanzio, 2024 analyzes employment flows and sectoral transitions both in Italy and Spain, my analysis explicitly evaluates the causal impact of WFH on quitting behavior and wages in the Italian labor market, complementing their descriptive approach. Similarly, while they document cross-sector transitions and rising re-employment probabilities, I investigate how WFH reshapes the underlying decision to quit and its implications for earnings. Finally, I also extend beyond Faberman et al., 2022, who studies job search and wage dynamics through job ladders but does not consider how non-wage amenities alter quits or wage setting.

Second, by extending the analysis to wages and the elasticity of quits to the firm, thereby connecting WFH to employers' market power, I focus on a dimension not yet addressed in the literature. Particularly, I contribute to the huge debate on monopsony power, where Manning, 2021 and Manning, 2003 provide the main theoretical foundation for imperfect labor market competition, showing how mobility frictions and heterogeneous preferences grant firms wage-setting power. I draw on this literature (see, for instance, Langella and Manning, 2021; Webber, 2015, among others), analyzing how and to what extent WFH can create the scope for monopsony power. Other studies emphasize frictions arising from labor market concentration (see, for example, Azar and Marinescu, 2024; Azar et al., 2020; Dube et al., 2020) or commuting constraints (see, for instance, Caldwell and Danieli, 2024; Datta, 2023; Datta, 2022; Naidu and Carr, 2022). My contribution is to shift the focus to the quit elasticities only and to examine how they vary with WFH. In doing so, I provide the first evidence on how WFH reshapes wage outcomes, thereby uncovering a new channel through which non-wage amenities influence quit elasticities and, ultimately, employers' market power.

The rest of the paper is structured as follows. Section 1.2 provides an overview of the institutional context, while Section 1.3 explains the data and the methodology used to define teleworkable occupations. It also discusses the integration of teleworkability measures with the administrative dataset and provides descriptive statistics. Section 1.4 outlines the empirical strategy, focusing on the DiD event study design. I will present the main results, starting with the workers' quitting behavior, the wage effects of WFH, and then exploring how the results translate into changes in quit elasticity. I will also show the robustness and the placebo test. Section 1.5 includes a detailed discussion of heterogeneity in the results. Finally, Section 1.6 concludes.

1.2 Institutional Background

The COVID-19 pandemic reached Italy in early 2020, with the first wave spreading rapidly across Northern regions. In response, the government imposed a nationwide lockdown starting on 10 March 2020, distinguishing between "essential" and "non-essential" activities. Non-essential sectors, such as most manufacturing, retail trade, and entertainment, were forced to stop in-person activities, while essential services like healthcare and some public administration continued to operate. To mitigate job losses, the government introduced emergency labor market measures. Two key policies were introduced in March 2020:

1. A unique subsidy program known as the COVID-related STW (short-time work) compensation scheme (i.e., *Cassa Integrazione Guadagni*), providing crucial financial support for nine weeks, retroactively effective from February 23. This innovative scheme was designed to help preserve employment relationships by offering a government-backed subsidy for both partial and full-time reductions in working hours. It replaced a significant 80 percent of lost earnings due to these reductions, up to a certain threshold, allowing employees to face the diminished income. What set this initiative apart was its broadened scope. It extended the protective embrace of the usual STW provisions to smaller firms with fewer than 15 employees, generally excluded from such support, alongside those already benefiting from the extraordinary STW, a specialized type of aid governed by Italian employment protection laws. Typically, these two forms of assistance could not

be combined, but the COVID-related scheme created an exception, enabling businesses to optimize their support[1].

2. A wide prohibition on layoffs, designed to shield employees from the threat of job loss for a crucial period of 60 days, starting on March 17. This mandate not only applied to future dismissals but also had the power to retroactively safeguard those whose layoffs were still in the approval pipeline, dating back to February 23[2].

These interventions primarily benefited firms in non-essential sectors (see, for instance, Di Nicola et al., 2021). As such, COVID-related policies likely helped reduce layoffs, independently of WFH feasibility. Moreover, these emergency measures were not temporary stopgaps: both the special STW scheme and the layoff ban were extended multiple times throughout 2020, reflecting the prolonged impact of the pandemic on the labor market. The pandemic STW was gradually phased out only in the second half of 2021, as economic conditions improved and restrictions were lifted. Similarly, the layoff ban was lifted in a staggered fashion, with different timelines based on sector and firm characteristics: for most firms in the industrial and construction sectors, the ban ended in July 2021, while for the hardest-hit service sectors, including tourism and retail, protections were extended until October 2021. These staggered exits from emergency support imply that the intensity of protection varied both over time and across sectors, with teleworkable occupations often remaining under protection for a longer period.

These policies may have introduced compositional biases in observed separation behavior. By reducing involuntary separations also in WFH sectors, pandemic measures could mechanically inflate the share of separations in these industries. Then, the lower separation rate in WFH-friendly sectors may reflect, in part, the effect of these policies, rather than a behavioral response to WFH per se. The key problem is whether these emergency measures undermine my identification strategy by affecting separations independently of WFH feasibility. Importantly, both the STW scheme and the layoff ban applied broadly across sectors, including those where WFH was possible. In this sense, the protective measures were not systematically targeted at teleworkable versus non-teleworkable occupations, but rather at the essential/non-essential classification and firm size. This implies that the variation in teleworkability I exploit is orthogonal to the presence of these policies. While the layoff ban and STW scheme did reduce overall separation rates, they are unlikely to bias the relative comparison between teleworkable and non-teleworkable occupations, which is at the core of my empirical strategy. Consistent with this interpretation, Figure A3 shows that the share of layoffs dropped

[1] Firms participating in this program were granted greater flexibility; they could renew temporary contracts while bypassing typical regulatory constraints. This financial lifeline allowed businesses to strategically reduce labor costs during the lockdown, ensuring survival and the preservation of jobs. In Appendix A, I show that the share of part-time and open-ended contracts, both in levels (Figure A1) and relative to 2019 (Figure A2), does not display any systematic divergence between teleworkable and non-teleworkable occupations, reassuring that changes in contract composition do not drive results.

[2] This policy aimed to provide a buffer of security and reassurance amid uncertainty. The nationwide dismissal ban should have led to a sudden and uniform decline in layoffs across occupations.

sharply in 2020 for both teleworkable and non-teleworkable occupations, with parallel dynamics thereafter. When expressed relative to 2019 levels, the changes likewise display no systematic divergence across groups. This evidence suggests that COVID-related emergency policies reduced layoffs broadly, but did not differentially distort separation patterns across teleworkability groups.

I exclusively focus on voluntary quits and control for detailed individual, sectoral, and occupational fixed effects, thereby ruling out compositional differences in separations driven by sector- or occupation-specific exposure to emergency policies. Since COVID-19 related interventions, such as the layoff ban and the STW scheme, were applied broadly at the sectoral or firm-size level, but were not targeted based on teleworkability, these fixed effects absorb persistent differences in policy intensity across groups. What remains is the variation in quitting behavior between teleworkable and non-teleworkable occupations within the same policy environment, making it unlikely that the estimated effects are mechanically driven by policy-induced separation bias.

The focus exclusively on voluntary quits, that is, separations explicitly labeled as resignations, is a crucial choice in this institutional setting: during the pandemic, government interventions, such as the STW scheme and the layoff ban, significantly reduced involuntary separations (such as dismissals or end-of-contract terminations). In addition, under the STW scheme, workers on reduced hours were eligible to receive up to 80 percent of their forgone earnings, subject to a cap. While this compensation helped preserve employment relationships, it may have also created disincentives for firms to dismiss workers and may have led to voluntary resignations that were indirectly induced by reduced income or limited career prospects during the subsidy period. However, this level of support was relatively generous, especially for lower-wage workers, and likely discouraged quitting during the height of the pandemic[3]. In regular times, quits often respond to wage losses, dissatisfaction, or career concerns. By contrast, during the subsidy period, the income cushion muted these usual drivers of mobility, so that observed voluntary quits are less likely to reflect routine job-to-job transitions and more likely to capture stronger preferences or constraints. Workers had little financial incentive to leave their current jobs voluntarily unless they faced significant dissatisfaction or external pressure. In this context, observed voluntary quits are more likely to reflect strong preferences or constraints, not noise caused by policy. By restricting the analysis to voluntary quits, I aim to better isolate worker-initiated mobility decisions, net of policy-driven job attachment effects. To address this, I compare only voluntary exits across teleworkable and non-teleworkable occupations, thereby isolating worker-driven mobility choices. Moreover, by including individual fixed effects, I control for time-invariant unobserved heterogeneity across workers, such as

[3] Workers lost, on average, 90% of their usual working hours, but suffered only a 27.3% reduction in gross monthly earnings (see Di Nicola et al., 2021). In practice, while working activity was almost entirely suspended, income losses were substantially cushioned. Given that the SWT covered about 80% of forgone wages up to a cap (around €1,000 - €1,200 gross per month), the income protection was particularly strong for low- and middle-wage workers. Thus, although some voluntary quits may have occurred, the generous level of income replacement likely discouraged quitting for the majority of affected workers during the peak of the pandemic.

preferences for job stability, risk attitudes, or family constraints, that could otherwise bias the comparison between teleworkable and non-teleworkable occupations. This ensures that the estimated effects are identified from within-worker changes over time, rather than reflecting systematic differences in the workforce composition.

1.3 Data

1.3.1 Labor Force Survey

The Italian Labor Force Survey (LFS) is a cross-sectional dataset provided by ISTAT, the Italian National Institute of Statistics, with quarterly information on the labor market status and other socio-economic characteristics of a representative sample of the Italian population. The dataset contains about 95,000 observations for each quarter. To assess the frequency of remote work opportunities for employees, I analyzed the period spanning from the first quarter of 2009 to the second quarter of 2023[4]. Since my analysis focuses on employer-employee relationships, I restrict the sample to dependent employees aged 16-65, excluding self-employed workers. After these sample selection criteria, I am left with 2,318,536 observations. The dataset provides detailed information on industries and occupations categories up to the 4-digit level, classified according to the ATECO 2007[5] and the CP2011[6] system. The survey includes questions about the possibility of working from home, which allows me to build a variable that measures whether workers in each industry and occupation can perform their jobs from home. Importantly, this question allows me to build a measure of realized WFH. The question asks how often employees worked from home in the last four weeks before the interview, and it is structured as follows: "In the last four weeks, did you work at home?". Employees can specify the frequency at which they carry out their work duties remotely. This can be two or more times per week, at least once per week, or not at all. Using this information, I built the variable WFH equal to one for workers who declared to have worked from home twice or more than twice and less than twice per week (and zero otherwise). This variable measures realized WFH, rather than merely capturing the potential for teleworking based on occupational characteristics. This realized measure of WFH is particularly valuable, as it enables me to test whether ex-ante teleworkability indices - based on occupational characteristics - actually predict ex-post WFH adoption in the data. In other words, the LFS provides a unique opportunity to validate the predictive content of potential WFH measures against observed remote work behavior. For this dataset, descriptive statistics are presented in Appendix A.

[4] I have excluded observations from 2023 when evaluating the prevalence of WFH within industries and occupations at a very narrow level since I lack detailed information for that year. Nonetheless, I will still provide evidence on the prevalence of WFH in 2023 at a more general industry level (1-digit ATECO 2007 classification) in Figure A5 and Figure A6.

[5] The ATECO 2007 classification is the Italian classification of economic activities and coincides with the Nomenclature of Economic Activities (NACE Rev. 2) used in the European Union.

[6] The CP2011 codes refer to the Italian system for classifying occupations. This system is designed to align with the International Standard Classification of Occupations (ISCO-08).

Overall, Table A1 shows that from 2009 to 2023, only 4,13% of workers were remote before the interview, while most were not teleworking. Table A2 shows that in the years before the pandemic, the share of people employed in WFH was, on average, about 1.7. However, looking at the years immediately following the pandemic (meaning 2021 and 2022), the share of people working from home increased, reaching approximately 14%.

Figure 1.1 displays that the shift towards WFH persisted even after the pandemic.

FIGURE 1.1: Evolution of WFH over time

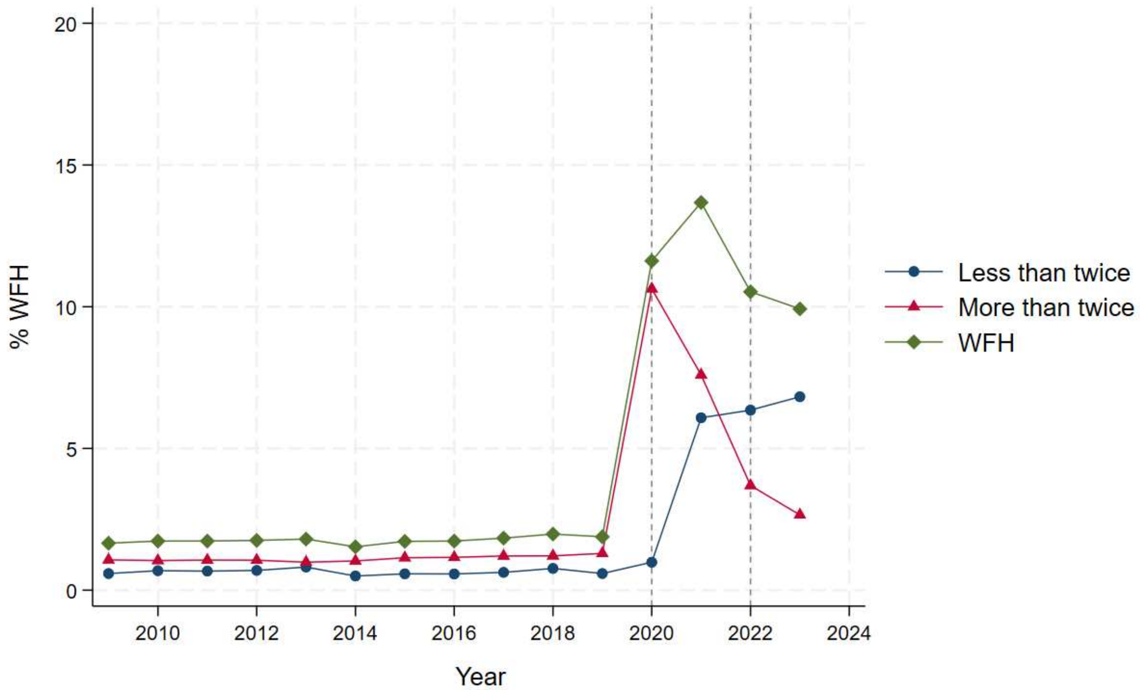

Notes. This figure illustrates the evolution of the percentage of workers engaging in WFH from 2009 to 2023. The green line represents the overall share of workers who reported working from home, obtained by summing those working from home less than twice per week (blue line) and those working from home more than twice per week (red line). Hence, the blue and red lines are disaggregated components of the green line. The two vertical dashed lines indicate significant periods, namely the onset of COVID-19 and the subsequent years.

This sustained increase highlights the potential lasting effect of the pandemic on work habits, particularly among workers whose roles enable them to work from home regularly. In fact, the enduring shift towards remote work is largely driven by those who work from home less than twice a week, suggesting a stable change in how work is organized.

Figure A4 reveals the industries in which these workers are mainly employed. The Information and Communication sector leads with approximately 60%, followed by the Financial and Insurance Activities sector at around 30%. Professional, Scientific, and Technical Activities sectors, along with Public Administration and Education, also show substantial increases, with WFH percentages ranging between 15% and 25%, in line with Dingel, 2020 for the U.S. labor market and Barbieri et al., 2022 for Italy. Other sectors, such as Real Estate Activities, Administrative Services, and various Service Activities, show moderate increases in remote options. Looking at what happened before the COVID-19 pandemic, the share of people teleworking in these industries was surprisingly low. The pandemic has forced many companies to adopt remote arrangements, where WFH was feasible even before but not used, and this has led to a persistent and significant shift in attitudes towards remote work.

Table 1.1 provides a comprehensive overview of both the top and bottom five occupations that have moved to WFH arrangements in the aftermath of the pandemic.

TABLE 1.1: Top and Bottom 5 Occupations by WFH after 2020

Occupation	**WFH**
Top 5 Occupations	
Mathematical, physical, natural sciences specialists	0.642
Engineers, architects	0.438
Social, management specialists	0.418
Entrepreneurs, administrators, directors of large private companies	0.381
Education, research specialists	0.352
Bottom 5 Occupations	
Agricultural machinery operators	0.003
Skilled construction workers	0.004
Semi-skilled production workers	0.004
Industrial plant operators	0.006
Vehicle operators	0.006

Notes. This table highlights the top and bottom five occupations at the 2-digit ISCO-08 level based on the share of workers performing their jobs from home (WFH) after 2020 from the LFS.

Topping the list are Mathematical, Physical, and Natural Sciences Specialists, closely followed by Engineers and Architects (in line with Hensvik et al., 2020 and Adams-Prassl et al., 2022), suggesting that occupations with digital infrastructure and less reliance on physical tasks are better suited for WFH adoption.

1.3.2 Task-Based WFH Potential Measure

This shift toward remote work, particularly in industries and occupations that were previously under-utilizing WFH, underscores the role of task characteristics in determining WFH feasibility. While the Italian Labor Force Survey offers information on WFH prevalence, relying solely on survey responses may not fully capture the underlying potential for remote work. Self-reported data can be influenced by workers' and firms' current preferences or constraints, which might not align with the intrinsic characteristics of the tasks performed in each occupation. To address this limitation, I incorporate a task-based measure of WFH feasibility for the analysis. This approach evaluates the potential for remote work based on the intrinsic characteristics of occupational tasks, such as the degree of physical presence required or the reliance on digital tools. This decision is also driven by the key endogeneity challenge that arises when estimating the effect of WFH on quits. Workers who prefer flexibility may self-select into remote jobs, and firms may offer WFH to retain employees prone to quitting. For this reason, I focus on the *potential* for WFH rather than actual adoption. Unlike observed WFH, which reflects individual and firm choices, task-based feasibility is exogenous to quitting decisions, allowing for a more credible identification strategy.

Following Basso et al., 2022, I define the "teleworkability" of each occupation, i.e., the extent to which an occupation can be done from home, relying on their task-based classification system. This approach constructs WFH measures by leveraging detailed information from the O*NET database, which documents the nature of tasks and required skills for occupations in the U.S. labor market[7] (see, for a similar approach, Alipour et al., 2023; Mongey et al., 2021; Gottlieb et al., 2021; Dingel, 2020; Boeri et al., 2020). Specifically, this classification assesses the extent to which tasks associated with each occupation can be performed remotely, taking into account factors such as the need for face-to-face interaction, the physical presence required for task execution, and the reliance on digital tools and technology[8]. Importantly, this methodology maps these O*NET-based characteristics to international occupational classifications, including ISCO-08, which is crucial for my analysis as it links occupations to the Italian CP2011 system. This mapping enables me to determine whether specific occupations in the Italian labor market are teleworkable based on the intrinsic nature of their tasks. Crucially, since this teleworkability measure is based on pre-pandemic task requirements, it captures the intrinsic feasibility of remote work and is not affected by subsequent changes in work practices introduced after the COVID-19 pandemic.

Table 1.2 categorizes occupations as teleworkable or non-teleworkable based on their task requirements, following the task-based approach developed by Basso et al., 2022.

[7] Although the WFH classification is based on the U.S. O*NET database, it can be reliably applied to Italy by leveraging a crosswalk from U.S. occupational codes (SOC) to international classifications (ISCO), which are harmonized with the Italian Labour Force Survey (EU-LFS). Following Basso et al., 2022, O*NET-based measures are mapped to ISCO codes through a weighted procedure that adjusts for differences in occupational structures between the U.S. and Europe. Their empirical validation confirms that the resulting teleworkability index predicts COVID-19-related labor market outcomes in Italy, supporting its applicability in this context.

[8] Basso et al., 2022 introduces a novel classification for the "teleworkability" of 3-digit ISCO occupations into four categories based on their exposure to infection risk and feasibility for remote work, using data from the CSP, EU-LFS, and O*NET. The categories are the following: (1) jobs that can be done from home, (2) jobs that cannot be done from home but have low infection risk, (3) unsafe jobs that cannot be done from home and present high infection risk, and (4) essential occupations carried out during lockdowns. For the main analysis, I use the first category.

TABLE 1.2: Top & Bottom 5 Teleworkable and Non-Teleworkable Occupations

Teleworkable Occupations	**Non-Teleworkable Occupations**
Legislators and Senior Officials	Production Managers in Agriculture, Forestry
Managing Directors, Chief Executives	Medical Doctors
Business Services and Administration Managers	Nursing, Midwifery Professionals
Sales, Marketing and Development Managers	Social and Religious Professionals
Physical, Earth Science Professionals	Physical, Engineering Science Technicians
...	...
Numerical Clerks	Subsistence Fishers, Hunters, Trappers
Client Information Workers	Building Finishers and Related Trades Workers
Secretaries (general)	Machinery Mechanics and Repairers
General Office Clerks	Electrical Equipment Installers and Repairers
Legal, Social and Religious Associate Professionals	Mining and Mineral Processing Plant Operators

Notes. This table highlights the top and bottom five teleworkable and non-teleworkable occupations based on their task requirements as classified in Basso et al., 2022 at the 3-digit ISCO-08 level.

The teleworkable occupations include roles such as Legislators, Managers, and Professionals that involve tasks compatible with remote work, such as those relying on digital tools or requiring minimal physical presence. Non-teleworkable occupations, by contrast, include roles such as construction workers, Cooks, and Farmers, where tasks require physical presence or manual interaction.

To strengthen the validity of this task-based index, I cross-check it with survey evidence from the LFS, which directly asks respondents about their actual use of WFH. The results in Figure 1.2 confirm that the teleworkability measure strongly predicts realized WFH adoption during the pandemic, providing empirical support for the accuracy of this classification in capturing remote work feasibility in the Italian context.

FIGURE 1.2: Validation of the Teleworkability Index with Actual WFH Adoption

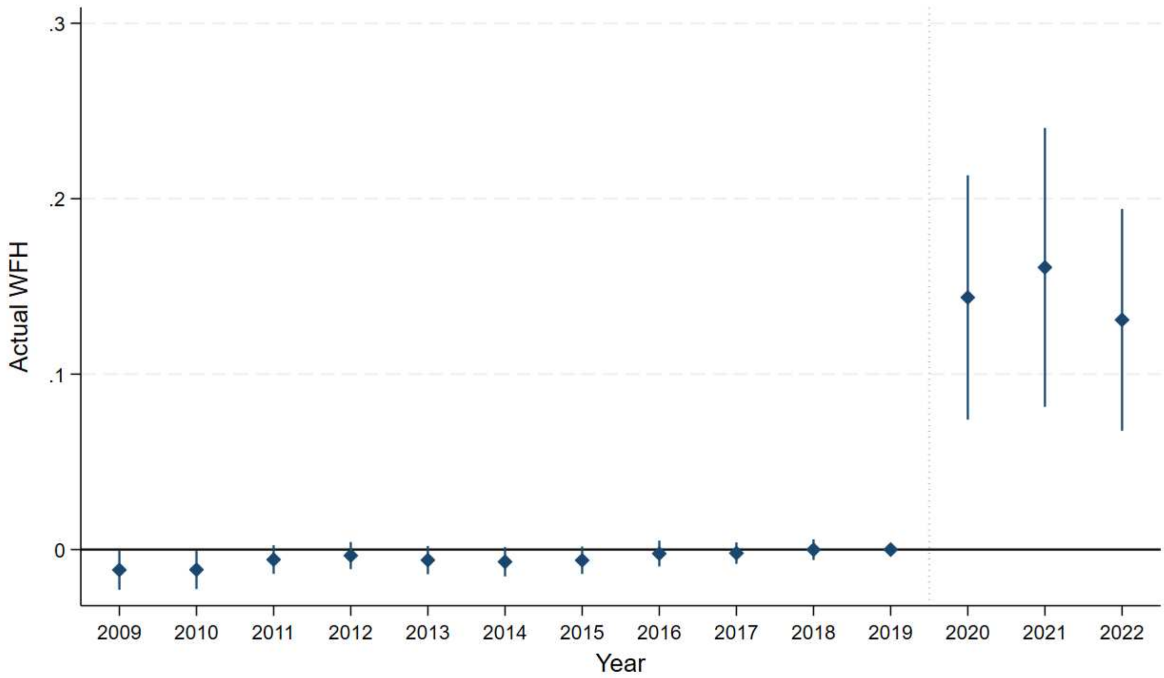

Notes. This figure plots the relationship between the teleworkability index proposed by Basso et al., 2022, and actual WFH adoption from 2009 to 2022, based on survey evidence from the LFS. Estimates are obtained by regressing actual WFH on the teleworkability measure, while controlling for time, occupations (at the 3-digit level), and sector (at the 2-digit level) fixed effects. Standard errors are clustered at the 3-digit occupation level.

In addition to this task-based classification, I complement the analysis by leveraging estimates of the share of hours worked at home within different occupations from Hensvik et al., 2020. This provides an additional measure of the intensity of teleworkability, reflecting how frequently tasks associated with a given occupation can be teleworkable.

Figure 1.3 shows the distribution of the hours worked from home for teleworkable and non-teleworkable occupations. Teleworkable occupations exhibit a distribution skewed toward higher shares of hours worked from home, with a notable concentration around 20%, meaning that many teleworkable roles consistently incorporate remote work into their routines. In contrast, non-teleworkable occupations have a distribution concentrated near zero, reflecting the minimal feasibility of remote work. The overlap between the two distributions in the lower range of the WFH share highlights that even some teleworkable occupations may not utilize remote work extensively, suggesting a consistent variability in teleworking adoption across job types. The stark differences in the distributions between the two groups emphasize the role of occupational characteristics in determining the intensity of WFH adoption.

FIGURE 1.3: Distribution of WFH Teleworkable vs Non-Teleworkable

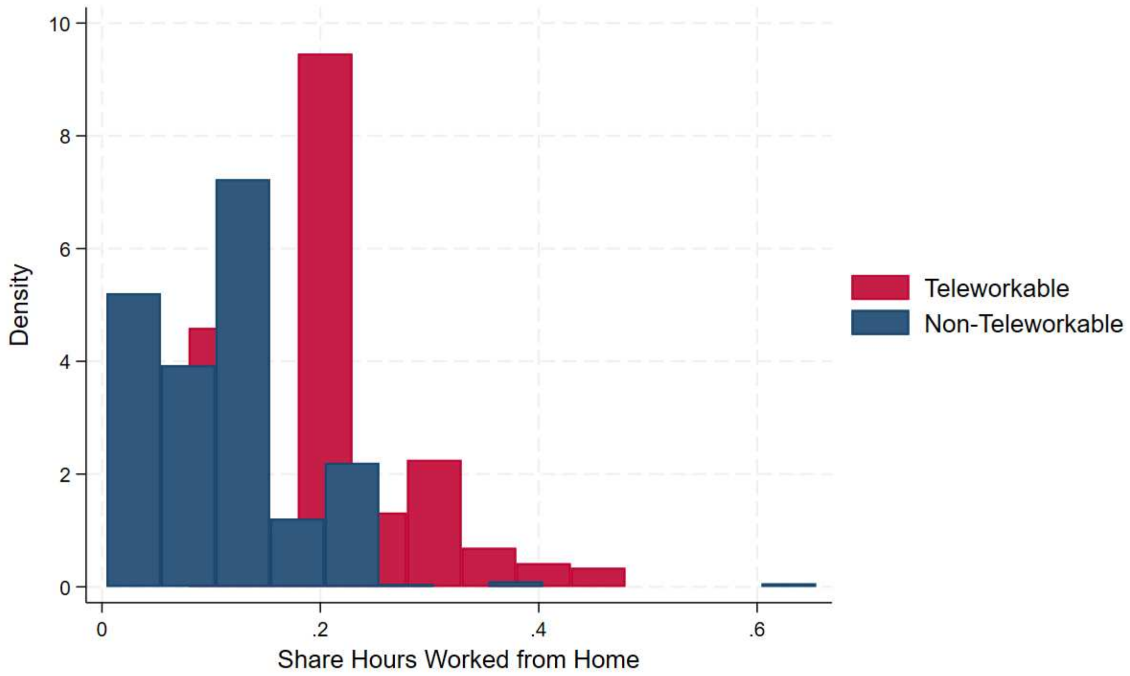

Notes. The figure displays the distribution of the share of hours worked from home from Hensvik et al., 2020 for teleworkable and non-teleworkable occupations. Teleworkable occupations are those classified as capable of being performed remotely based on task characteristics, while non-teleworkable occupations lack this feasibility, according to the category from Basso et al., 2022.

Building on this distinction, I use this classification to define as treated those occupations classified as "teleworkable", i.e., those with the unique capability of being performed from a remote setting, allowing for a clear identification of roles where remote work is a feasible and prevalent option. I primarily use the task-based teleworkability classification from Basso et al., 2022 to define teleworkable versus non-teleworkable occupations for the main causal analysis. Additionally, I leverage the continuous measure developed by Hensvik et al., 2020, which estimates the share of tasks that can be performed remotely, to complement the analysis and assess the intensity of remote work feasibility across occupations. Specifically, I use this continuous measure in robustness checks, replacing the binary indicator with this continuous WFH measure.

This dual approach ensures that the results are not relying on a single classification, but instead hold across alternative definitions of WFH feasibility. Moreover, the combination of a binary and a continuous measure allows me to distinguish between the extensive margin (whether an occupation can be performed remotely at all) and the intensive margin (the extent to which tasks within an occupation are suitable for remote work). In doing so, I can better capture the heterogeneity in remote work potential across occupations, strengthening the credibility of the empirical findings.

1.3.3 Administrative data

To evaluate the impact of WFH on workers' likelihood of quitting their jobs and to assess their quit elasticity, I use data from *Campione Integrato delle Comunicazioni Obbligatorie* (CICO), a 13 per cent sample of workers obtained from the administrative system that collects mandatory notifications that employers submit to the Italian Ministry of Labor when they activate or terminate a contract. In particular, the sampling process includes mandatory notifications that employers submit when they activate or terminate a contract between 2010 and 2023 for workers born on the 1st, 9th, 10th, and 11th of each month, so that the raw observation in the data is a job relationship with an employer within a calendar year. The dataset includes 23,615,462 contracts. This unique sampling feature allows me to track employees' career histories, including job changes and exits from the labor market, while also preserving representativeness at the national level. This feature is crucial for studying job mobility and labor supply elasticity, as it provides a longitudinal view of workers' employment trajectories rather than a static cross-section. Moreover, the data records the type of contract (full-time or part-time, temporary, open-ended, or apprenticeship), the activation and end date for each contract, and the reason for the termination[9]. The granularity of this information allows me to precisely distinguish voluntary quits from other separations, which is essential for my empirical strategy. The data provides detailed information on the type of workers' occupation[10], employers' industries[11] and the gross monthly wage at the start of each contract. Moreover, I also have information on workers' demographic and personal characteristics - such as gender, year of birth, and education level - which enables me to control for observable sources of heterogeneity in job mobility and wage outcomes. The richness of the dataset is particularly valuable in this context, as it allows me to combine contract-level and worker-level dimensions to study how WFH reshaped both separation behavior and wage dynamics.

For the analysis, I excluded contracts for individuals working abroad and only considered workers aged 16-65. I only considered temporary or open-ended contracts, as well as full-time and part-time workers. In addition, I excluded workers in the public sector, since employment relationships in the Italian labor market follow distinct contractual rules, are less exposed to market competition, and are differently affected by WFH policies. These restrictions ensure a focus on the private labor market, where workers' quitting behavior and wage adjustments

[9] Reasons for contract termination are recorded by employers through mandatory administrative reporting. A potential concern is that employers might misreport the reason for termination by recording a dismissal as a voluntary resignation. This risk is substantially mitigated by major institutional reforms implemented in Italy. Since 2016, following the Jobs Act, voluntary resignations must be submitted directly by workers through a mandatory electronic procedure, independent of the employer. This system requires workers to confirm their intention to resign via a dedicated online platform managed by the Ministry of Labor. Employers cannot unilaterally declare a resignation without the worker's explicit confirmation.

[10] The CICO dataset adopts the Italian occupational classification CP2011, which is directly aligned with the international ISCO-08 standard at the 3-digit level. To ensure comparability with the teleworkability measures from Basso et al., 2022 and Hensvik et al., 2020, which are based on ISCO-08, I map CP2011 codes into their corresponding 3-digit ISCO-08 categories.

[11] The CICO dataset provides a very detailed ATECO 2007 industry classification compared to the LFS. I focus the analysis on the 2-digit industry level, ensuring sufficient sample size within each cell, avoiding excessive fragmentation that could bias estimates due to very small groups.

are most directly linked to employers' market power. The final dataset has 15,608,483 observations.

Table 1.3 reports descriptive statistics about the dataset at the contract level. Workers in the dataset are predominantly male (61%), with the majority holding at most lower secondary education (63%). Most contracts are full-time (71%) and open-ended (71%), while temporary contracts account for 23% of the sample. The average worker is 44 years old, and the mean daily wage is about €51.

TABLE 1.3: Descriptive Statistics CICO

	Mean	**SD**	**Min**	**Max**
Male	0.61	0.49	0	1
Female	0.39	0.49	0	1
Lower Secondary School	0.63	0.48	0	1
Upper Secondary School	0.29	0.45	0	1
College or More	0.08	0.28	0	1
Part-Time	0.29	0.45	0	1
Full-Time	0.71	0.45	0	1
Open-Ended Contracts	0.71	0.45	0	1
Temporary Contracts	0.23	0.42	0	1
Age	43.60	11.43	16	65
Daily Wage	50.67	25.93	12	289
Observations	15608483	15608483	15608483	15608483

Notes. This table presents the mean, standard deviation, min, and max from CICO of workers' socioeconomic characteristics. Data are at the contract level from 2010 to 2023.

I converted this dataset into an individual panel annual data from 2016 to 2023, using the year of starting and ending contracts for each employee in the dataset, where the unit of analysis is the worker observed at a yearly frequency[12]. Having information about the reasons for the contract termination, I focus on voluntary resignations by constructing a dummy variable that equals one if the separation is due to a voluntary quit and zero otherwise[13]. As a result, I can analyze different types of job-to-job separations over consecutive years.

Figure 1.4 shows the share of quits over total contract terminations, while other descriptive statistics about the quit variable are detailed in Appendix A.

[12] For workers holding multiple contracts in overlapping periods, I assume that the main contract is the one that has the longest duration. In case two or more contracts have the exact same duration and start in the same year, I conventionally select the one with the highest wage.

[13] Quits are constructed following Table A3: (1) Resignation, (2) Resignation for Just Cause, (3) Resignation During the Probation Period, (4) Mutual Termination, (5) Dismissal for Justified Reason During Training, (6) Resignation of Mother Worker, (7) Mutual Termination Ex Art. 14, DL 104/2020.

FIGURE 1.4: Voluntary Quit over Years

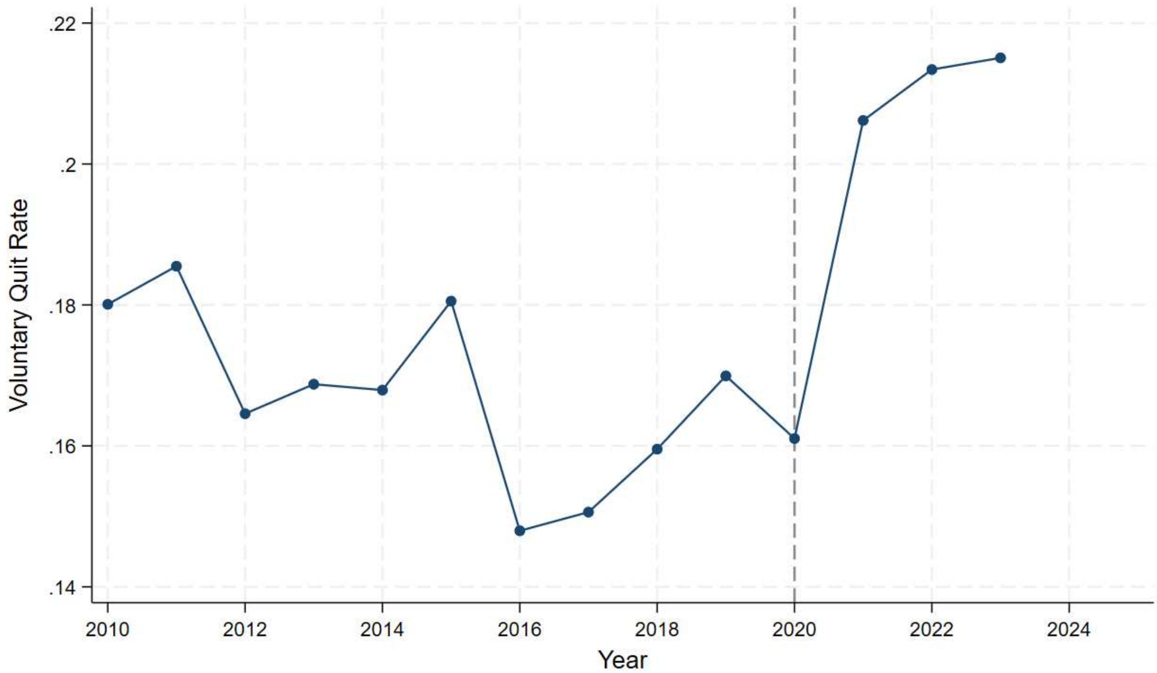

Notes. This figure shows the voluntary quit rate over time. The voluntary quit rate is computed as the number of contracts that end in voluntary quits in a given year (numerator) divided by the total number of contracts that terminate in that year. The quit variable is computed from the CICO dataset at the contract level.

After a period of relative stability in the 2010s, the quit rate exhibits a marked increase starting in 2021, with levels reaching their historical peak in 2022-2023. This sharp rise is consistent with the phenomenon widely documented in the literature as the "Great Resignation", whereby voluntary separations surged in the aftermath of the COVID-19 pandemic (see, for instance, Basso et al., 2023). I merge the information about the teleworkability of occupations, based on the classification by Basso et al., 2022, with this dataset containing information about workers' quitting behavior. This approach enables me to study workers' decisions to leave their jobs precisely and examine how these decisions relate to teleworkability, as well as other job and worker socio-economic characteristics.

1.4 Empirical Strategy and Results

Estimating the causal impact of WFH on voluntary quitting behavior raises significant challenges due to the endogenous nature of remote work. A key concern is self-selection of workers into WFH jobs: individuals who place a high value on flexibility may be both more likely to choose occupations or firms that offer remote work opportunities and, independently, more prone to quit jobs more frequently. In this case, the observed association between WFH and quits may simply capture differences in underlying worker preferences rather than the causal effect of WFH itself. In addition, there is a risk of reverse causality: firms may strategically offer WFH options to employees or occupations with higher turnover risk as a attachment strategy, making it unclear whether WFH reduces quitting or is instead targeted to groups already predisposed to leave. Together, these issues highlight the difficulty of disentangling the causal effect of WFH from both worker-driven selection and firm-driven responses.

For all the reasons stated above, simply comparing the voluntary quitting probability between individuals who opted for remote work and those who did not is likely to produce a biased estimate of the effect of WFH on voluntary quitting. Therefore, instead of relying on the actual implementation of remote work, I focus on the *potential* use of WFH in each occupation.

This approach is crucial for my identification strategy, as the potential for remote work within each occupation should remain uncorrelated with any workers' pre-existing preferences for WFH. The potential for WFH is determined by the intrinsic characteristics of the tasks associated with each occupation, such as the degree of interaction required, the need for physical presence, or the reliance on digital tools rather than on personal traits or workers' preferences.

Importantly, the teleworkability indices I use were constructed prior to the pandemic and are based on detailed occupational task content (e.g., reliance on ICT tools, physical/manual requirements, interpersonal interaction). This timing helps ensure that the measure reflects the inherent technological and organizational feasibility of remote work rather than ex-post adjustments to COVID-19. A potential concern is that the pandemic may have altered the effective feasibility of remote work for some tasks, for instance, through accelerated digitalization or organizational investments in remote infrastructures. However, this would primarily affect the adoption margin (whether firms actually implement WFH) rather than the underlying potential margin (whether tasks can, in principle, be performed remotely). In this sense, my approach identifies how differences in pre-determined teleworkability interacted with the COVID-19 shock, while abstracting from within-occupation technological upgrading that may have occurred in response to the pandemic. A task-based approach highlights that occupations vary greatly in their adaptability to remote work: jobs relying on digital tools and tasks such as information processing or written communication are well-suited to WFH, while those requiring manual labor, in-person collaboration, or the use of specialized equipment are not. While some within-occupation variation in WFH may exist - for example, workers in the same occupation might face different employers, industries, or personal circumstances - I account for this in two ways. First, I control for individual characteristics that absorb systematic differences in workers' propensity to work remotely within the same occupation. Second, I include detailed occupation fixed effects, which capture all time-invariant task characteristics that define teleworkability at the occupational level. This combination ensures that any residual within-occupation heterogeneity is unlikely to bias the estimates, as the identifying variation arises from differences in teleworkability across occupations interacting with the pandemic shock, rather than from unobserved sorting within occupations.

Although this strategy addresses within-occupation selection into WFH, it does not eliminate all concerns about occupational sorting: individuals may self-select into occupations with higher WFH potential based on unobserved characteristics, such as risk aversion, job stability preferences, or career orientation. To mitigate this issue, I include an extensive set of controls. Most importantly, individual fixed effects absorb all time-invariant worker characteristics, whether observed (e.g., education, contract types) or unobserved (e.g., preferences for flexibility, career orientation, or baseline attachment to the labor market). This means that identification comes from changes over time within the same individual, thereby ruling out the possibility that stable worker traits are driving the results. In addition, occupation, sector, and region fixed effects capture structural differences in labor market conditions and institutional features, while worker-level controls such as age and contract type account for time-varying characteristics. By combining individual fixed effects with these additional controls, my estimates exploit within-individual variation in exposure to WFH across time, making the interpretation of my results conditional on occupational choice but robust to unobserved heterogeneity across workers.

By leveraging the exogenous task-based WFH adaptability in each occupation, I implement a Difference-in-Differences (DiD) event study analysis, comparing changes in voluntary quitting probability across occupations with varying potential for remote work before and after the onset of the COVID-19 pandemic. Occupations are classified into treatment and control groups based on their potential to adapt to WFH, as determined by the share of tasks that can easily be performed remotely. This classification builds on the methodology used by Basso et al., 2022, which identifies "teleworkable" jobs before the COVID-19 pandemic as those that can be performed entirely from home since they do not require workers to leave their homes, interact with coworkers in person, or engage in tasks that necessitate physical presence at a specific location. I take this classification to define those more exposed to WFH during COVID-19, i.e., teleworkable occupations, as the treated group. I compare the likelihood of leaving the job between teleworkable and non-teleworkable occupations within a DiD framework. Then, I run the following regression analysis:

$$Q_{i,t} = \alpha_i + \lambda_o + \eta_r + \mu_s + \gamma_t + \delta\, WFH_o + \beta\, WFH_o \times Post_t + \varphi\, \mathbf{X'}_{i,t} + \varepsilon_{i,t} \tag{1.1}$$

where $Q_{i,t}$ indicates whether the worker i voluntarily quits his current job at time t. WFH_o is the teleworkability indicator defined by Basso et al., 2022, which indicates if the occupation is defined as teleworkable. $X'_{i,t}$ is a set of workers' socio-economics characteristics (including wages), while α_i, λ_o, η_r, μ_s, γ_t are individual, occupation, region, industry and time fixed effects, respectively. In other words, I exploit the interaction between pre-determined occupational teleworkability and the exogenous time shock introduced by the pandemic. The identifying variation comes from comparing changes in voluntary quit rates before and after 2020 across teleworkable and non-teleworkable occupations, while controlling for individual, occupational, sectoral, and regional fixed effects. Then, my main parameter of interest, i.e., β, captures the differential change in voluntary quitting probability between teleworkable and non-teleworkable occupations over time. It should be interpreted as the effect of being in an occupation suitable for remote work on quitting behavior during and after COVID-19[14].

Table 1.4 shows the estimated coefficient of equation (1.1).

TABLE 1.4: Linear Probability Model of WFH on Voluntary Quitting

	(1)	(2)	(3)	(4)
WFH × Post	-0.01***	-0.01***	-0.01***	-0.01***
	(0.00)	(0.00)	(0.00)	(0.00)
WFH	-0.02***	-0.02***	-0.02***	
	(0.00)	(0.00)	(0.00)	
Controls	✓	✓	✓	✓
Individuals FEs	·	✓	✓	✓
Regions FEs	·	✓	✓	✓
Sectors FEs	·	·	✓	✓
Occupations FEs	·	·	·	✓
Observations	7195022	7195022	7194360	7194360
Dep. Var. Mean	0.10	0.10	0.10	0.10

Notes. This table presents the estimated coefficients for the probability of voluntary quitting from equation (1.1). This first model includes controls for age, education, type of contract, wages, and individual fixed effects. The second model adds region fixed effects. The third one includes 2-digit industries fixed effects. Finally, the last model also includes fixed effects for 3-digit occupations. *Post* coefficient is omitted due to collinearity with year fixed effects, while *WFH* coefficient in model (4) is omitted due to collinearity with occupations fixed effects. Standard errors (in parentheses) are clustered at the worker level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The results indicate that individuals in occupations classified as more suitable for WFH are less likely to quit their jobs. The coefficient on WFH is negative and significant, indicating that, even before the pandemic, teleworkable occupations exhibited lower quit rates compared to non-teleworkable ones. This baseline difference reflects persistent occupational characteristics, such as job stability or skill requirements, that make quits structurally less frequent in teleworkable jobs, and reinforces the interpretation that the post-2020 effect captures an additional shift due to the expansion of remote work. In fact, across all specifications, the effect of teleworkability is consistently negative and highly significant, indicating that workers in teleworkable occupations are about one percentage point less likely to voluntarily quit their jobs after the pandemic compared to those in non-teleworkable occupations. Given that the mean quit rate is 10%, this represents a reduction of approximately 10% relative to the baseline

[14] Formally, the empirical design corresponds to an Intention-to-Treat (ITT) analysis, where treatment is defined by pre-pandemic occupational teleworkability. Since teleworkability is fixed at the occupation level and does not vary over time or across individuals, the estimates identify differential trends in voluntary quitting between treated (teleworkable) and control (non-teleworkable) groups. Thus, the findings should be interpreted as the impact of being in a job that aligns better with remote work on employees' decisions to leave their positions.

probability of quitting. The stability of the coefficient across models (from column 1 to column 4) suggests that the result is not driven by individual, regional, sectoral, or occupational composition effects.

To check whether treated and control groups, i.e., teleworkable and non-teleworkable occupations, are comparable, the trends in voluntary quitting probability for teleworkable and non-teleworkable occupations should have followed a similar trajectory without the WFH feasibility due to the pandemic. To test this assumption, I implement an event-study analysis that examines the pre-treatment differences in quitting behavior across the two groups. Specifically, I implement the following event-study regression:

$$Q_{i,t} = \alpha_i + \lambda_o + \eta_r + \mu_s + \gamma_t + \sum_{t=2016, t\neq 2019}^{2023} \beta_t (WFH_o \times D_t) + \varphi \mathbf{X'}_{i,t} + \varepsilon_{i,t} \tag{1.2}$$

where $Q_{i,t}$ indicates whether the worker voluntarily quits. WFH_o is a dummy variable indicating whether the occupation is defined as teleworkable, interacted with an indicator D_t for time. $X'_{i,t}$ are a set of workers' socio-economics characteristics, while α_i, λ_o, η_r, μ_s, γ_t are individual, occupation, region, industry and time fixed effects, respectively as in equation (1.1). The event-study coefficients β_t measure the difference in the likelihood of voluntary quitting between teleworkable and non-teleworkable occupations in year t, relative to the baseline year 2019. These coefficients capture the dynamic evolution of the effect of teleworkability on quitting behavior before and after the onset of the pandemic. Standard errors are clustered at the worker level to account for within-worker correlation over time.

Figure 1.5 shows the estimated evolution of the voluntary quitting probability of equation (1.2) for the teleworkable occupations.

The assumption of parallel trends is supported since the pre-2019 coefficients are not significantly different from zero. Before the pandemic, the two groups followed a similar trend, with no significant differences in quitting probabilities observed between potentially teleworkable and non-teleworkable occupations.

The plot in Figure 1.5 reveals a sharp divergence in quitting probabilities between potentially teleworkable and non-teleworkable occupations after 2019, coinciding with the onset of the COVID-19 pandemic. The negative coefficients observed for 2020 and 2021 indicate that individuals in teleworkable occupations were significantly less likely to quit compared to their non-teleworkable counterparts during the pandemic. The probability of voluntary quitting for teleworkable occupations decreased by approximately 0.5 percentage points in 2020 and 1 percentage points in 2021 relative to non-teleworkable occupations when no controls are included. These reductions correspond to a relative decrease of 5% to 10% in quitting probabilities during the pandemic, and the effects are particularly striking, reflecting the increased importance of teleworkability in retaining workers during the pandemic.

FIGURE 1.5: Difference in the Probability of Quitting Teleworkable vs Non-Teleworkable

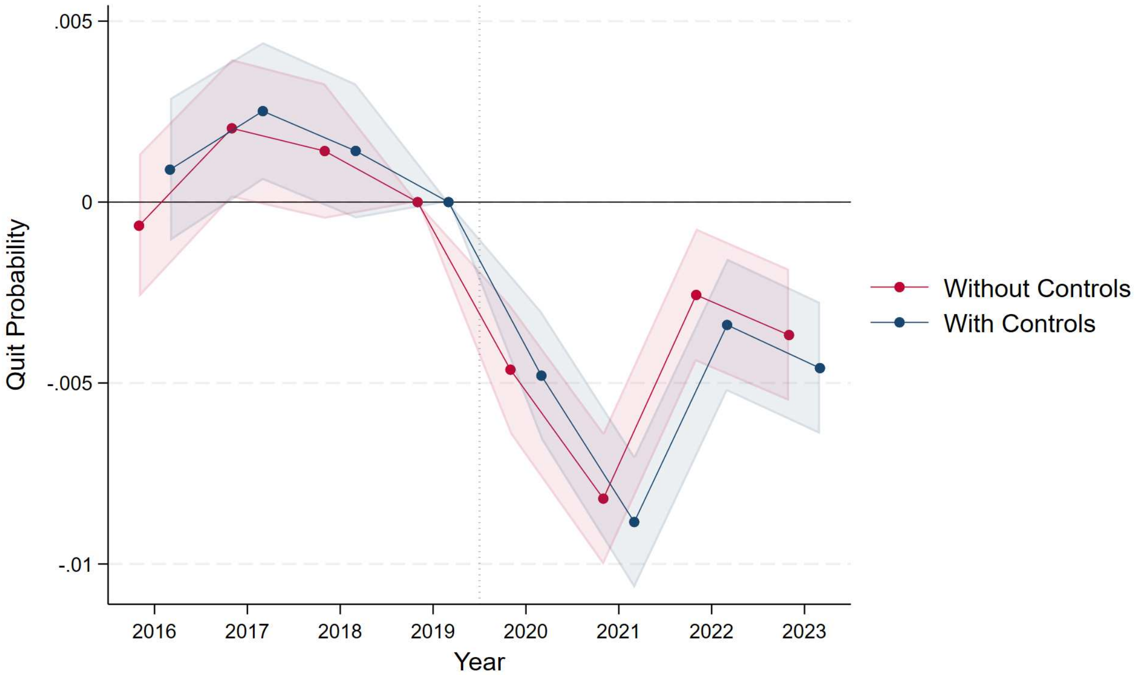

Notes. This figure shows the estimated event-study coefficients from (1.2) for the probability of voluntary quitting among teleworkable and non-teleworkable occupations over time. The red line represents estimates without controls, while the blue line includes controls for workers' socio-economic characteristics, as well as individual, occupation, region, industry, and year fixed effects. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of COVID-19.

In the post-pandemic years (2022-2023), the gap in quitting probabilities narrows, making the effects less pronounced and significant. During this period, the reduction in the likelihood of quitting for teleworkable occupations is approximately 0.2 to 0.3 percentage points, corresponding to a relative decrease of 2% to 3% compared to the average quit rate. These results suggest that the potential for WFH provided a particularly strong protective effect against voluntary quitting during the pandemic, with its influence diminishing after the pandemic.

To see whether this negative effect of WFH on voluntary quitting probabilities translates into changes in employers' market power, I first study whether WFH has the potential to affect wages. My aim is to test whether the WFH could act as non-monetary compensation. Then, I run the following event-study specification:

$$log\left(w_{i,t}\right) = \lambda_o + \eta_r + \mu_s + \gamma_t + \sum_{t=2016, t\neq 2019}^{2023} \beta_t \left(WFH_o \times D_t\right) + \varphi \mathbf{X'}_{i,t} + \varepsilon_{i,t} \tag{1.3}$$

where $log\ (w_{i,t})$ indicates the log of the workers' daily wages. WFH_o is a dummy variable indicating if the occupation is defined as teleworkable. $X'_{i,t}$ are a set of workers' socio-economic characteristics, λ_o, η_r, μ_s, γ_t are occupation, region, industry, and time fixed effects, respectively. The event-study coefficients β_t measure the difference in the log daily wages between teleworkable and non-teleworkable occupations in year t, relative to the baseline year 2019. Standard errors $\varepsilon_{i,t}$ are clustered at the worker level.

Figure 1.6 shows the estimated coefficient for the log daily wages of equation (1.3). Before the COVID-19 pandemic, teleworkable and non-teleworkable occupations displayed comparable wage trends. Interestingly, following the pandemic, occupations better suited for remote work experienced a gradual decline in daily wages. Starting in 2020, the wage differential for teleworkable occupations begins to decline significantly, reaching a reduction of approximately 0.01 to 0.015 in log daily wages by 2023. This indicates that workers in teleworkable occupations experienced a relative decline in daily wages compared to their non-teleworkable counterparts during the pandemic and its aftermath. The overall trend suggests a persistent wage penalty for teleworkable occupations compared to those that are non-teleworkable.

FIGURE 1.6: Difference in Log Daily Wages Teleworkable vs Non-Teleworkable

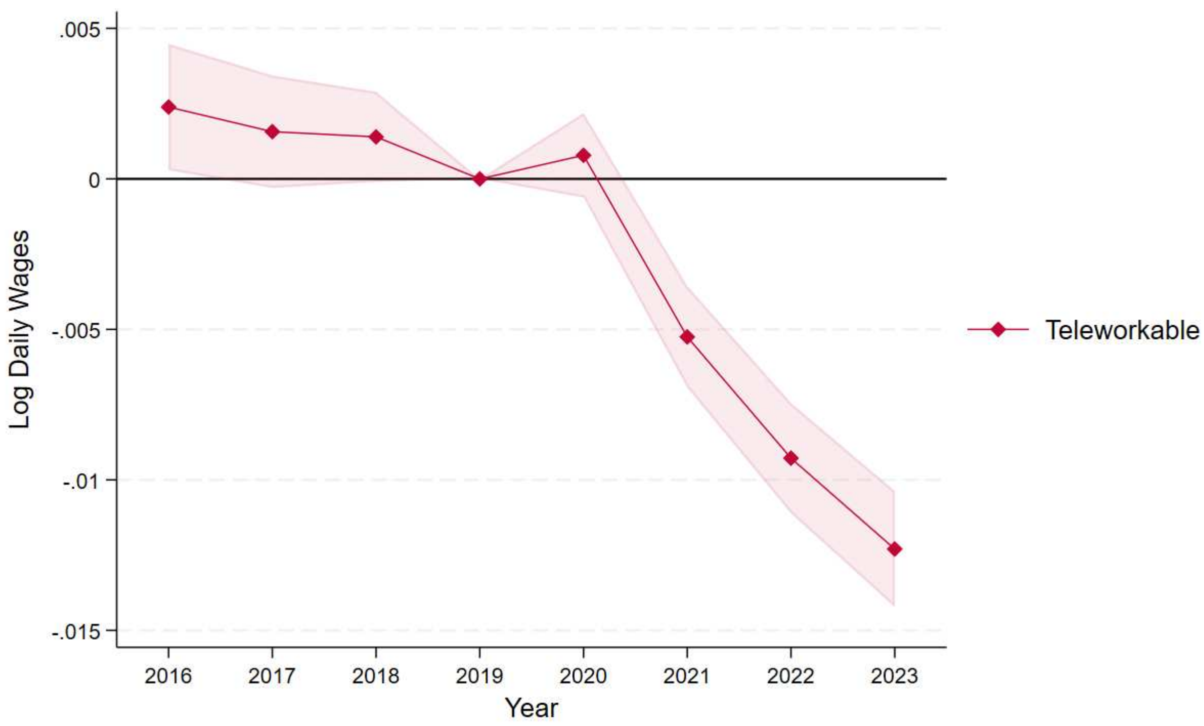

Notes. This figure shows the estimated event-study coefficients from (1.3) for the log daily wages among teleworkable and non-teleworkable occupations over time. The red line represents estimates, including controls for workers' socio-economic characteristics, region, years, and industry effects. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of COVID-19.

The relative decline in wages for teleworkable occupations may be connected to underlying productivity and profit dynamics in the firms and sectors where these jobs are concentrated. Recent evidence suggests that while WFH improves employee satisfaction and attachment, the productivity effects are heterogeneous and often weaker in the long run (see, for instance, Bloom et al., 2015; Barrero et al., 2021a). At the same time, firms in digital-intensive industries, where teleworkable occupations are prevalent, often experienced rising revenues during the pandemic, yet these gains did not necessarily translate into higher wages (see, e.g., Autor et al., 2023; Morikawa, 2022).

This pattern is consistent with a scenario where WFH adoption generated only limited productivity spillovers while increasing firms' bargaining power, leading to a decoupling between firm performance and wages. Although my dataset does not allow me to observe productivity or profits directly, the wage penalty I document aligns with these broader trends.

I provide some interesting heterogeneity analysis for this effect in Appendix B. Figure B1 compares the log daily wages between teleworkable and non-teleworkable occupations by gender. While both men and women experienced wage declines in teleworkable occupations, men showed stronger wage pressures than women. A priori, one might expect women to bear a more substantial wage penalty from working in teleworkable occupations. This is because women are disproportionately represented in lower-paid teleworkable jobs (e.g., administrative and teaching roles). If WFH is perceived as a non-wage amenity, employers may adjust wages downward more for female workers, thereby amplifying the gender wage gap.

However, the estimates reveal the opposite pattern: men in teleworkable occupations experience a sharper decline in wages compared to women. This result is consistent with sectoral composition effects. Male workers are more concentrated in high-paying and competitive teleworkable sectors, such as finance, ICT, and professional services, where the shift to remote work may have intensified wage compression or altered bargaining dynamics. By contrast, many female teleworkable jobs are located in public or highly regulated sectors, where wages are less responsive to labor market adjustments. Additionally, differences in bargaining behavior may contribute to the observed pattern. Men are generally more mobile across jobs and more likely to negotiate wage increases actively (see, for instance, Card et al., 2016). In the context of teleworkable occupations, especially in competitive, high-paying sectors such as finance, ICT, and consulting, this higher propensity to search and bargain could have intensified competition among male workers. Firms, facing greater outside options and renegotiation attempts, may have responded by compressing wages within these male-dominated occupations. By contrast, women are more concentrated in teleworkable jobs with lower mobility and stronger wage rigidity, where such competitive pressures are weaker.

Figure B2 compares the evolution of log daily wages across teleworkable occupations by educational attainment. For individuals with lower levels of education (i.e., those with primary and lower secondary school) and medium levels of education (i.e., those with upper secondary school), the decline in wages may reflect their concentration in teleworkable roles that are less specialized or involve routine tasks, where competition for remote positions intensified during the pandemic. In theory, one might expect less-educated workers to be less affected by wage adjustments in teleworkable occupations, since their tasks are often routine and wages already compressed. However, employers in these roles likely faced greater flexibility to reduce wages as the pool of available workers expanded and bargaining power for less-educated workers weakened. In contrast, higher-educated workers often engage in teleworkable occupations requiring specialized skills, such as professional or managerial roles, where demand remains strong and wage competition is less severe. The observed wage growth for highly educated workers may also reflect increased productivity in high-skilled remote roles and their stronger negotiation power in a tight labor market.

These results raise important questions about the interplay between wages and quit behavior, particularly for occupations better suited for remote work. The wage penalty observed for teleworkable jobs may alter the incentives for workers to remain in their positions, potentially affecting the elasticity of quits with respect to wages. In standard labor economics theory, a wage decline is typically associated with an increase in quit rates, as workers seek better-paying opportunities elsewhere. However, in the case of teleworkable occupations, the flexibility and non-monetary benefits associated with remote work may counteract this effect, making workers less sensitive to wage declines.

Opposite effects may arise: the wage penalty, combined with the widespread adoption of WFH across firms, could lead to higher quit elasticity for teleworkable occupations. When WFH becomes a common feature offered by most employers, its value as a differentiating non-wage amenity diminishes, and workers may become more responsive to wage differences when deciding whether to quit. The geographic mobility and broader job market access facilitated by remote work could increase workers' ability to switch jobs in response to wage changes, amplifying their responsiveness to wage. As remote work becomes normalized, its non-monetary benefits may no longer fully offset declining wages, potentially leading workers to prioritize financial stability over flexibility.

Thus, from a general and theoretical perspective, remote work can be interpreted as a job amenity, consistent with the theory of compensating differentials (see, for instance, Rosen, 1986; Roback, 1982). If WFH is highly valued as an amenity, lower wages may be accepted by workers as part of the compensation bundle, which explains the observed penalty. This reasoning clarifies the wage-quit link: when the amenity dominates, wage reductions translate into fewer quits than expected, as flexibility acts as a attachment mechanism. However, when the amenity becomes widespread and no longer differentiates jobs, the compensating differential story weakens, and quits should respond more strongly to wage variation. In other words, the same wage penalty can underpin both lower and higher quit elasticities, depending on how much value workers attach to remote work relative to wages.

To study whether and in what direction remote work affects workers' quit elasticity, I model the voluntary quitting probability as an instantaneous separation rate according to the following model (see Langella and Manning, 2021):

$$Pr(Q = 1|x) = e^{e^{\beta x_t}} \tag{1.4}$$

where the probability of quitting a job Q, in a given interval of time t, is a function of suitable controls x, including wages. This type of model can be estimated through a complementary log-log model of having a separation[15]. This model, which belongs to the family of the logit model, becomes especially useful when outcomes are heavily skewed and the estimated coefficient can be interpreted as the relevant quit elasticity.

Table 1.5 shows the result from the complementary log-log model of equation (1.4).

TABLE 1.5: Quit Elasticity

	(1)	(2)
Log Daily Wages	-0.14***	-0.17***
	(0.00)	(0.00)
Log Daily Wages × Post		0.05***
		(0.00)
Controls	✓	✓
Year FEs	✓	✓
Region FEs	✓	✓
Sector FEs	✓	✓
Occupation FEs	✓	✓
Observations	7601453	7601453

Notes. This table presents the estimated coefficients from the complementary log-log model on the probability of voluntary quitting from (1.4). This model includes controls for age, gender, education, working region, occupation (3-digit level), and industry (2-digit level). Standard errors (in parentheses) are clustered at the worker level * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

15 This model can be also expressed as $log(-log(1 - Pr(Q = 1|x))) = \beta_1 x_1 + \beta_2 x_2 + ... + \beta_k x_k$. Each coefficient in the model represents the change in the log of the hazard (or odds) of the event occurring for a one-unit increase in the corresponding predictor variable, holding all other variables constant. To interpret the coefficients in terms of odds ratios, they must be exponentiated. The exponentiated coefficient provides the multiplicative change in the odds of the event occurring for a one-unit increase in the predictor variable.

The estimated coefficients can be interpreted as quit elasticities, measuring how sensitive quitting behavior is to wage changes. Before COVID-19, the elasticity was negative (-0.172), indicating that higher wages reduced the probability of quitting. After the pandemic, the elasticity became less negative (-0.122), implying that workers were less responsive to wage differences. Overall, the elasticity weakened after COVID-19, suggesting that quitting decisions became less sensitive to wage changes.

Several factors can account for this result. The pandemic fundamentally reshaped labor market dynamics, with many workers reevaluating their job preferences and priorities. Increased uncertainty and changes in work organization may have led workers to value stability and non-wage aspects of their jobs - such as flexibility or security - more than before. As a result, wage differences became a less decisive factor in quitting decisions, consistent with the observed decline in the wage elasticity of quits. However, this aggregate pattern may hide important heterogeneity. For teleworkable occupations, the expansion of remote opportunities may have expanded workers' effective choice set and mobility, thereby mitigating the overall decline in wage responsiveness.

Given these results, I estimate the exact specification of equation (1.4), interacting log daily wages with the measure of WFH to explore whether the variation in quit elasticity observed in the post-pandemic period can be directly linked to higher WFH exposure. Table 1.6 reports how the quit elasticity interacts with wages and the measure of WFH.

Results indicate that the overall reduction in quit elasticity after the pandemic is less pronounced for occupations with higher WFH feasibility. Although workers in teleworkable jobs remain less elastic on average, their wage sensitivity increased relative to non-teleworkable occupations in the post-pandemic period.

This pattern is consistent with the idea that remote work expanded workers' outside options by relaxing geographical constraints and facilitating job mobility across regions and employers. As a result, teleworkable occupations experienced a relative increase in wage responsiveness. In this sense, WFH appears to partly mitigate employers' market power by enhancing workers' ability to respond to wage differentials, while continuing to provide valuable non-wage amenities that strengthen job attachment.

This outcome highlights an important trade-off: while WFH increases job satisfaction and attachment through flexibility and non-wage amenities, its wider adoption across firms can also expand workers' outside options, making them more responsive to wage differences as similar amenities become available elsewhere. In this sense, the spread of remote work may both strengthen worker retention and enhance wage competition, thereby reshaping employer market power in more teleworkable occupations.

TABLE 1.6: Quit Elasticity with WFH

	(1)	(2)
Log Daily Wages × Post	0.06***	0.08***
	(0.00)	(0.00)
WFH × Post	-0.05***	0.21***
	(0.01)	(0.04)
Log Daily Wages × WFH		0.06***
		(0.01)
Log Daily Wages × WFH × Post		-0.07***
		(0.01)
Log Daily Wages	-0.18***	-0.19***
	(0.00)	(0.00)
Controls	✓	✓
Year FEs	✓	✓
Region FEs	✓	✓
Sector FEs	✓	✓
Occupation FEs	✓	✓
Observations	7601453	7601453

Notes. This table presents the estimated coefficients from the complementary log-log model on the probability of voluntary quitting as in (1.4), interacting the wages with the measure of WFH. This model includes controls for age, gender, education, working region, occupation, and industry (at the 3-digit level). Standard errors (in parentheses) are clustered at the worker level * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Appendix B presents some descriptive facts on how the quit elasticity varies across the wage distribution. Figure B3 shows that after the pandemic, the quit elasticity is less negative at most wage levels compared to the pre-COVID period. This suggests that workers across the wage distribution became less sensitive to wage changes when quitting after the pandemic. Notably, this pattern is particularly pronounced at the higher end of the wage distribution.

1.4.1 Robustness Checks and Placebo Test

One possible concern of my analysis is the distinction between the potential for WFH, as defined by task-based measures, and its actual implementation during the pandemic. These are conceptually different: potential WFH captures the feasibility of performing tasks remotely, while actual WFH reflects the realized adoption of remote work arrangements. While both concepts can be measured with some imprecision, relying solely on potential WFH might overlook variation in actual take-up across occupations and industries. To address this, I test the robustness of my results using alternative measures of WFH exposure. Specifically, I adopt the continuous classification of Hensvik et al., 2020, which uses the observed share of hours worked from home within occupations, and compare it to the binary teleworkability indicator of Basso et al., 2022. Results are reported in Appendix A.

Figure A10 shows that results are consistent across both approaches, indicating that my conclusions are not sensitive to the measure employed and reflect differences in WFH exposure.

The placebo test is designed to validate my identification strategy by testing whether treatment effects appear in groups that should not logically experience the impact of remote work. This test evaluates whether WFH drives the effect by randomly assigning treatment status to a group that is inherently non-teleworkable and re-estimating the model. For example, suppose assigning treatment status to occupations such as construction workers, which do not involve tasks suited for remote work. If the results show significant reductions in quitting probabilities for this placebo group, it would suggest that the observed effects in the actual treated group may be driven by unrelated factors or model misspecification.

Reassuringly, Table A4 shows that the coefficient for the placebo treatment group interacted with the post-treatment period is close to zero and not statistically significant. The lack of a significant effect in the placebo group reinforces the causal interpretation of my results.

1.5 Heterogeneity Analysis

Looking at heterogeneity effects is crucial for understanding how WFH interacts with quitting behavior across different groups of workers. It comes naturally to guess that adopting remote work influences quitting behavior differently depending on factors such as workers' demographic characteristics (e.g., gender, age, or education), geographical areas, or types of contracts (e.g., open-ended or part-time contracts). Then, I conducted several heterogeneity analyses to identify subgroups for whom the effects of WFH on quitting behavior are particularly pronounced. Results are reported in Appendix B.

Figure B5 presents the findings from the event-study analysis on the probability of quitting across teleworkable and non-teleworkable occupations, differentiated by gender. It shows that although both women and men experienced similar negative trends in voluntary quits, the effects are particularly more pronounced for men in teleworkable occupations two years following the pandemic. Over time, these negative effects disappear after the pandemic's peak. The larger reduction in quitting probability for men in teleworkable occupations could reflect higher male participation in jobs where remote work was widely implemented or stronger preferences among men for stability during economic uncertainty. For women, the reductions in quits were also significant and more pronounced, potentially due to competing factors such as greater caregiving responsibilities.

Figure B6 shows the analysis by age. The probability of voluntarily leaving a job in those occupations considered more teleworkable does not significantly differ between adults (ages 36-64) and young workers (ages 16-35), especially during the pandemic and post-pandemic periods.

Both groups exhibit similar reductions in quitting probabilities for teleworkable occupations relative to non-teleworkable ones during the pandemic years, and these effects diminish in the post-pandemic period. The convergence of quitting probabilities could suggest that teleworkable occupations mitigate age-related differences in labor market mobility. For younger workers, who typically exhibit higher mobility, the flexibility of remote work may have reduced the need to seek new opportunities. For older workers, who are generally less mobile, the benefits of remote work may have further incentivised job attachment.

Figure B7 plots the event study on the probability of quitting across various education levels. Workers with lower secondary education exhibit the most pronounced reductions in quitting probabilities for teleworkable occupations during the pandemic (2020-2021). Post-pandemic, the quit probabilities for all groups begin to converge, with workers of higher education experiencing a slight increase in quitting probabilities for teleworkable occupations, though not statistically significant. Workers with lower education seem to benefit most from the stability provided by teleworkable occupations during the pandemic. This group likely faced fewer alternative employment opportunities in the uncertain labor market, making them more likely to remain in teleworkable jobs when such opportunities were available.

Figure B8 shows the event-study estimates of the probability of quitting teleworkable versus non-teleworkable occupations by type of employment contract (open-ended vs. temporary). Interestingly, the trends reveal significant differences in quitting behavior between workers with open-ended and temporary contracts, particularly during and after the pandemic. Quitting probabilities decrease significantly for workers with open-ended contracts in teleworkable occupations, reflecting a strong preference for job stability and security during the heightened uncertainty of the pandemic. Workers with temporary contracts also experience a reduction in quitting probabilities. Still, the decline is less pronounced, likely due to the inherently higher mobility associated with temporary work and limited opportunities for job transitions during the pandemic. In the post-pandemic years (2022-2023), workers with open-ended contracts show a gradual rise in quitting probabilities in teleworkable occupations, converging toward pre-pandemic levels as labor markets stabilize and workers regain confidence in exploring new opportunities.

Figure B9 compares teleworkable and non-teleworkable occupations for workers with part-time and full-time contracts. The feasibility of WFH due to COVID-19 does not affect the likelihood of quitting for those workers with part-time contracts. Interestingly, WFH has stronger negative effects for workers with a full-time contract. In particular, full-time workers in remote options experienced a decrease in the probability of voluntarily leaving their jobs. The increased probability of quitting for full-time workers in teleworkable occupations after the pandemic may reflect several post-pandemic labor market dynamics. During the pandemic, full-time workers likely valued the stability and flexibility offered by teleworkable roles, leading to reduced quitting probabilities. However, as remote work became more normalized in the post-pandemic period, its non-monetary benefits, such as flexibility and reduced commuting, may have diminished in perceived value.

Figure B10 illustrates the estimated event-study coefficients for the probability of voluntary quitting in teleworkable versus non-teleworkable occupations across Italian regions. Before 2019, the trends for the North, Central, and South of Italy show some variation but generally remain close to zero, indicating limited regional differences in quitting behavior between teleworkable and non-teleworkable occupations. After the onset of the COVID-19 pandemic, notable regional divergence emerged. Workers in the North exhibit a marked increase in quit probability for teleworkable occupations relative to non-teleworkable ones, while the Central and Southern regions display less pronounced changes. This suggests that adopting WFH and its associated impact on labor mobility and quitting behavior were more significant in the North, possibly reflecting regional disparities in WFH feasibility and labor market dynamics.

1.6 Conclusions

This paper provides novel evidence about the adoption of Work from Home (WFH) arrangements and its potential to reshape labor market dynamics in Italy in the post-pandemic period. Using a combination of administrative employer-employee matched data and task-based measures of teleworkability, I show that WFH significantly affects both workers' quitting behavior and employers' market power. First, I find that occupations more suited for remote work experienced a notable reduction in voluntary quits during and after the COVID-19 pandemic. This suggests that the non-monetary benefits of WFH, such as flexibility and reduced commuting costs, act as attachment mechanisms, enhancing job stability even in a period marked by heightened labor mobility. Second, my results reveal a significant relationship between WFH and employers' market power. The less pronounced decline in the wage sensitivity of quitting behavior in teleworkable occupations indicates that WFH acted as a non-wage amenity that strengthened attachment, while at the same time mitigating the overall reduction in workers' responsiveness to wage differences observed after the pandemic. Overall, these results highlight an important trade-off: WFH increases job satisfaction and attachment, yet preserves comparatively greater wage responsiveness by expanding workers' outside options in teleworkable occupations, relative to non-teleworkable ones.

1.7 Appendix A: Additional Summary Statistics

TABLE A1: Work from Home (WFH) Frequency and Percent Distribution

WFH	**Freq.**	**Percent**	**Cum.**
No	2,215,695	95.56	95.56
Yes	95,808	4.13	99.70
Missing	7,033	0.30	100.00
Total	2,318,536	100.00	

Notes. This table displays the frequency and percentage distribution of remote workers from 2009 to 2023. WFH is constructed to be equal to 1 for workers who declared to WFH two or more times per week and less than twice per week before the interview (and 0 otherwise). *Missing* category represents the respondents who did not answer the survey question.

TABLE A2: WFH Yearly Mean Values

Year	**Mean**
2009	0.017
2010	0.017
2011	0.017
2012	0.018
2013	0.018
2014	0.015
2015	0.017
2016	0.017
2017	0.018
2018	0.020
2019	0.019
2020	0.117
2021	0.137
2022	0.105
2023	0.099

Notes. This table presents the year's mean of actual WFH (constructed as in Table A1).

TABLE A3: Reason for Contract Termination

Reason Termination	**Freq.**	**%**	**Cum.**
Individual Dismissal	424	0.00	0.00
Collective Dismissal	97,574	0.63	0.63
Resignation	2,065,901	13.24	13.86
Resignation for Just Cause	47,843	0.31	14.17
Resignation During the Probation Period	110,196	0.71	14.88
Failure to Pass the Probation Period	248,395	1.59	16.47
Change of Initially Set Term	614,024	3.93	20.40
Death	19,457	0.12	20.53
Retirement	44,745	0.29	20.81
Other	488,393	3.13	23.94
Dismissal for Just Cause	169,531	1.09	25.03
Mutual Termination	98,146	0.63	25.66
Dismissal for Objective Justified Reason	811,947	5.20	30.86
Dismissal for Subjective Justified Reason	56,193	0.36	31.22
Service Termination	3,906	0.03	31.24
Activity Cessation	116,272	0.74	31.99
Dismissal for Just Cause During Training	5	0.00	31.99
Dismissal for Justified Reason During Training	6	0.00	31.99
Resignation for Just Cause or Justified Reason	3	0.00	31.99
Termination at the End of the Probation Period	8	0.00	31.99
Resignation of Mother Worker	10,022	0.06	32.05
Termination with Retirement Requirements	149	0.00	32.05
Mutual Termination Ex Art. 14, DL 104/2020	3,317	0.02	32.08
Termination at the End of Transformed Contract	1,810,023	11.60	43.67
Termination at the End	7,086,851	45.40	89.08
Not Defined	1	0.00	89.08
Still Active	1,705,151	10.92	100.00
Total	15,608,483	100.00	

Notes. This table provides the frequency (Freq.), percentage (%), and cumulative percentage (Cum.) for each category of reasons for contract termination.

FIGURE A1: Share of Contract Types by Year

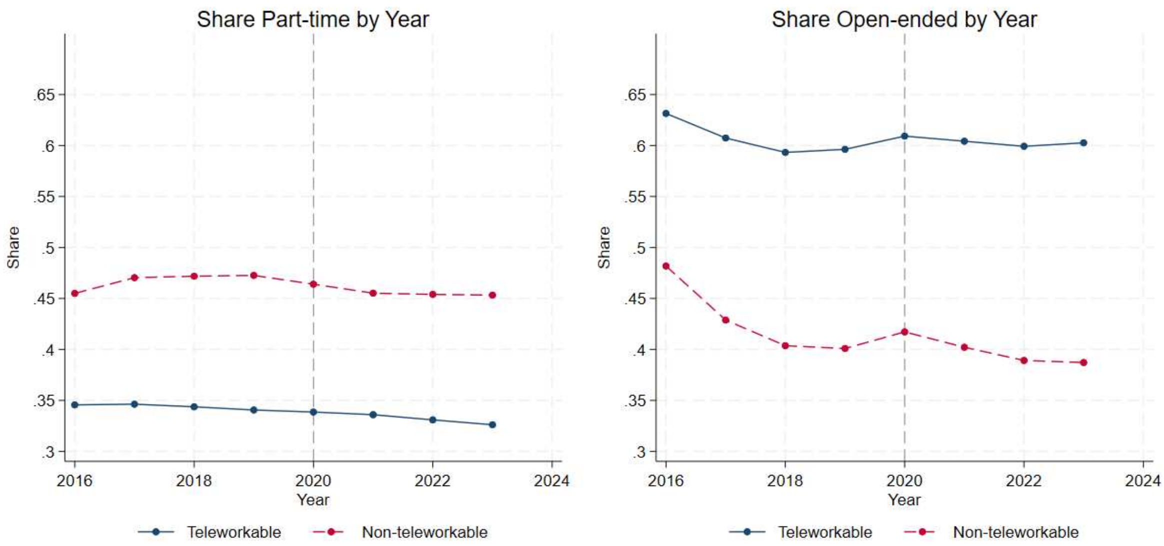

Notes. The figures plot the share of part-time contracts (left panel) and open-ended contracts (right panel) by year, separately for teleworkable and non-teleworkable occupations. Shares are computed as the number of contracts of a given type divided by the total number of contracts in the corresponding group and year. Trends are broadly parallel and do not exhibit systematic divergence after the COVID-19 pandemic.

FIGURE A2: Change in Contract Shares Relative to 2019

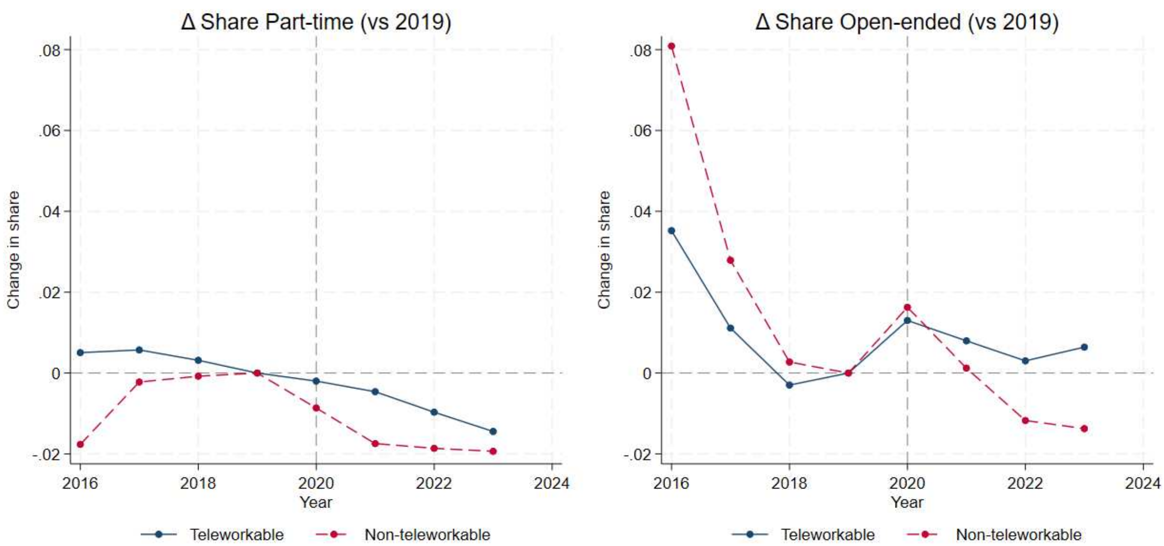

Notes. The figures plot the change in the share of part-time contracts (left panel) and open-ended contracts (right panel), measured relative to 2019, for teleworkable and non-teleworkable occupations. Shares are defined as the ratio of the number of contracts of a given type (numerator) to the total number of contracts in the same group and year (denominator). The plotted values represent the difference between this share in year t and the corresponding share in 2019. The absence of differential shifts suggests that COVID-related emergency policies did not alter contract composition across groups in a way that could confound the main results.

FIGURE A3: Layoff Dynamics before and after the COVID-19 Shock

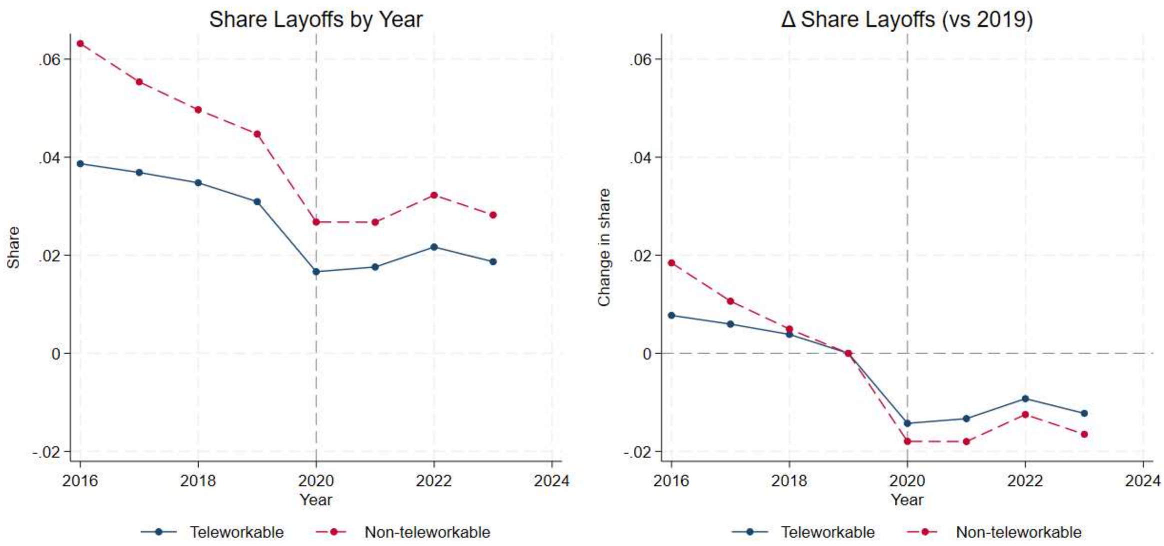

Notes. This figure shows the share of layoffs before and after the pandemic. The left panel plots the share of contracts ending in layoffs by year, while the right panel displays changes relative to 2019, separately for teleworkable and non-teleworkable occupations. The measure is constructed as the number of contracts that terminate with a layoff (numerator) divided by the total number of contracts observed in the corresponding group and year (denominator). Both groups experienced a sharp and parallel drop in layoffs in 2020, coinciding with the nationwide dismissal ban introduced in March. This pattern indicates that the policy acted as a common shock, uniformly reducing layoffs across occupations.

FIGURE A4: Average WFH across sectors pre and post COVID-19 pandemic

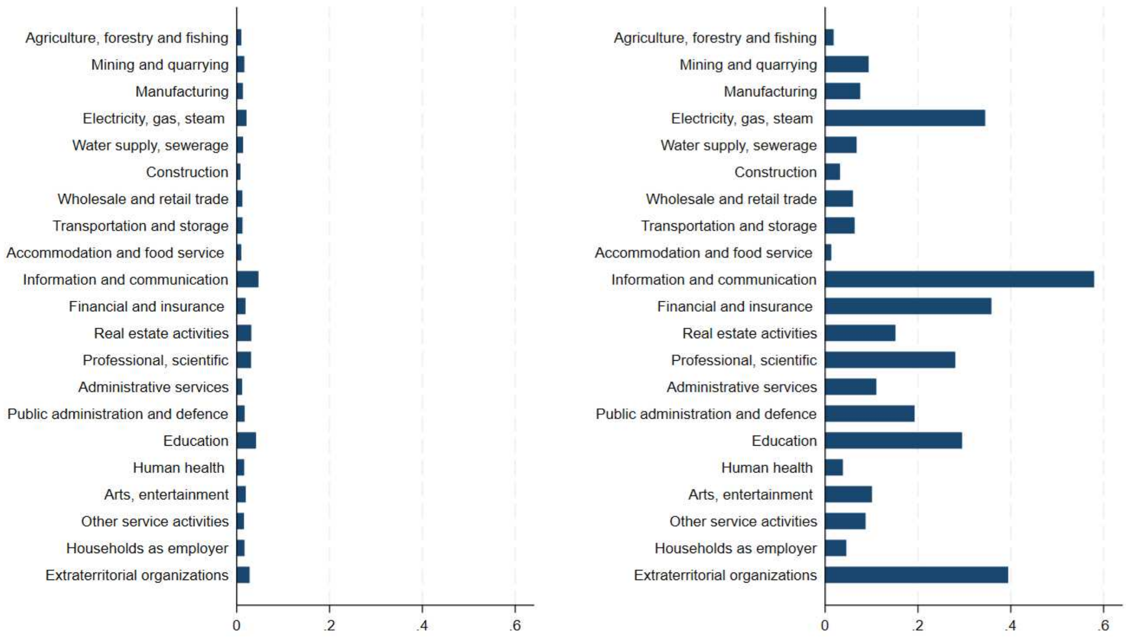

Notes. This graph shows a comparative snapshot of WFH percentage distribution across the 2-digit ATECO 2007 classification. The left panel displays the period 2009-2019. The right panel presents the corresponding figures for the period 2021 to 2022.

FIGURE A5: Incidence of WFH in each sector in 2023

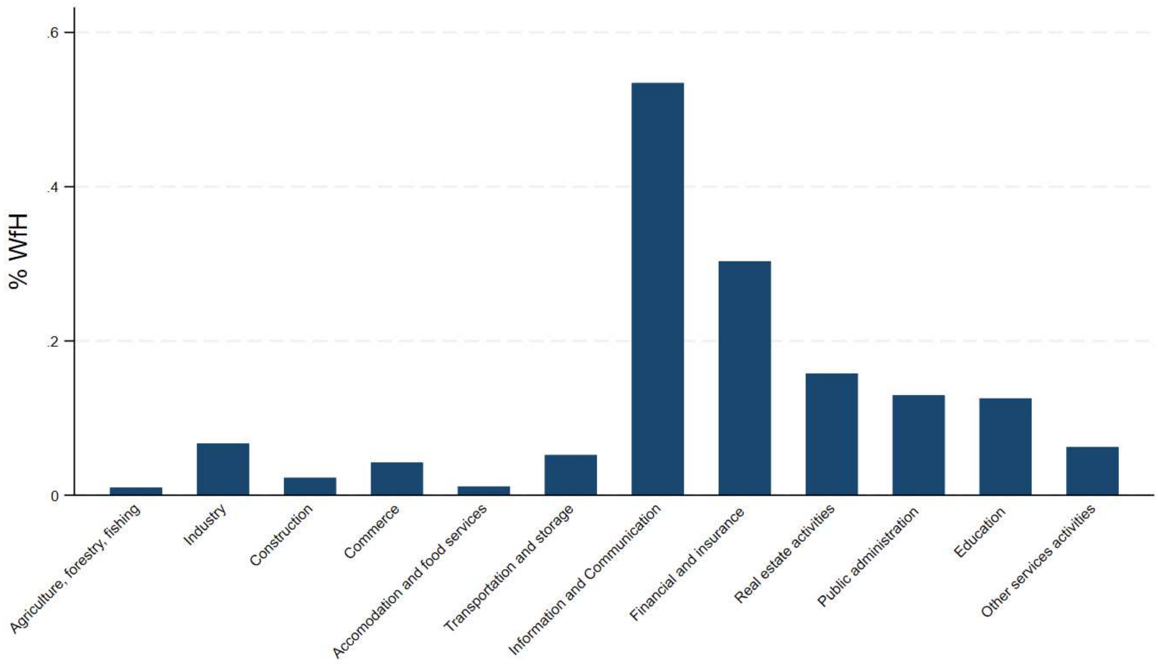

Notes. Bar chart of the % of workers in 2023 in WFH at the broadest ATECO 2007 classification.

FIGURE A6: Incidence of WFH in each sector in 2023

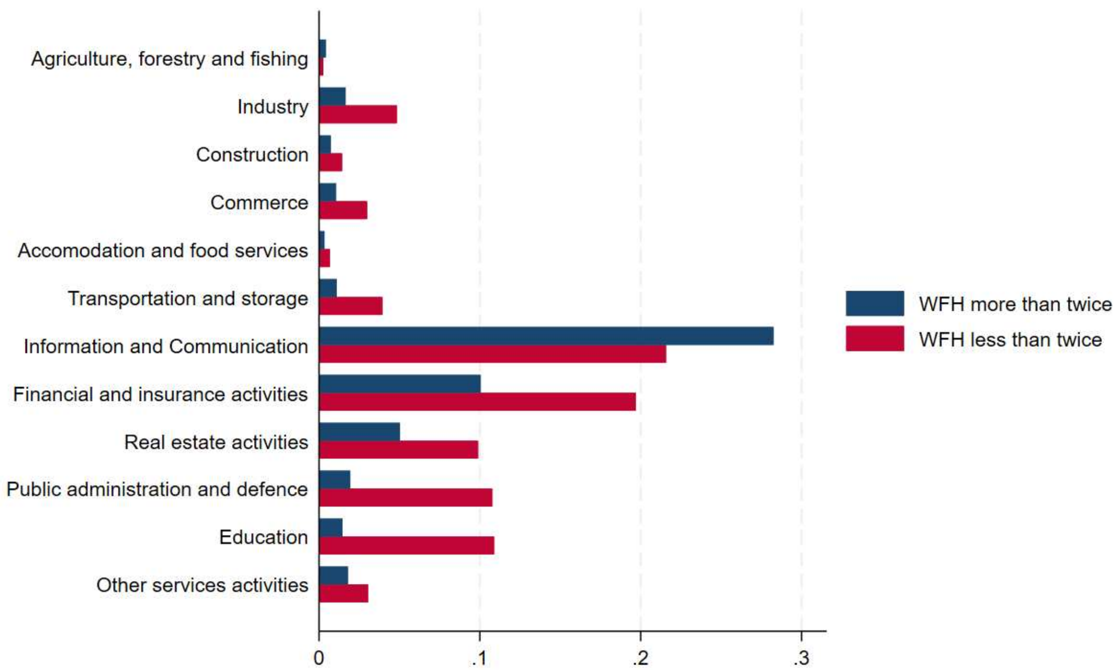

Notes. Bar chart of the % of workers in 2023 in WFH at the broadest ATECO 2007 classification, distinguishing between those who WFH more than twice a week and those who WFH less than twice a week.

FIGURE A7: Voluntary Quit over Years by Gender

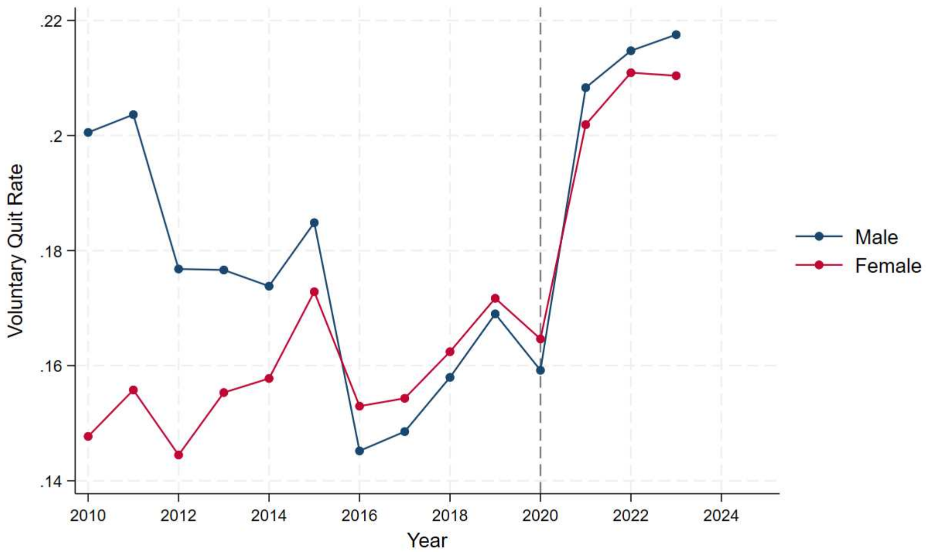

Notes. The figure displays the average quit rate over the years by gender. The blue line shows the average quit for men, while the red line displays the same statistics for women. The quit variable is computed from the CICO dataset.

FIGURE A8: Voluntary Quit over Years by Age

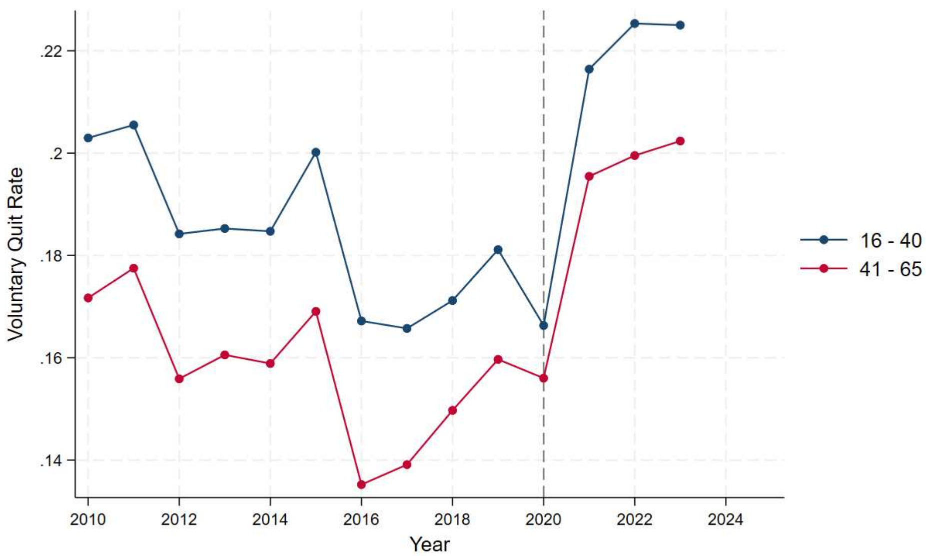

Notes. The figure displays the average quit rate over the years by age. The blue line shows the average quit for younger workers, while the red line displays the same statistics for older workers. The quit variable is computed from the CICO dataset.

FIGURE A9: Voluntary Quit over Years by Occupations

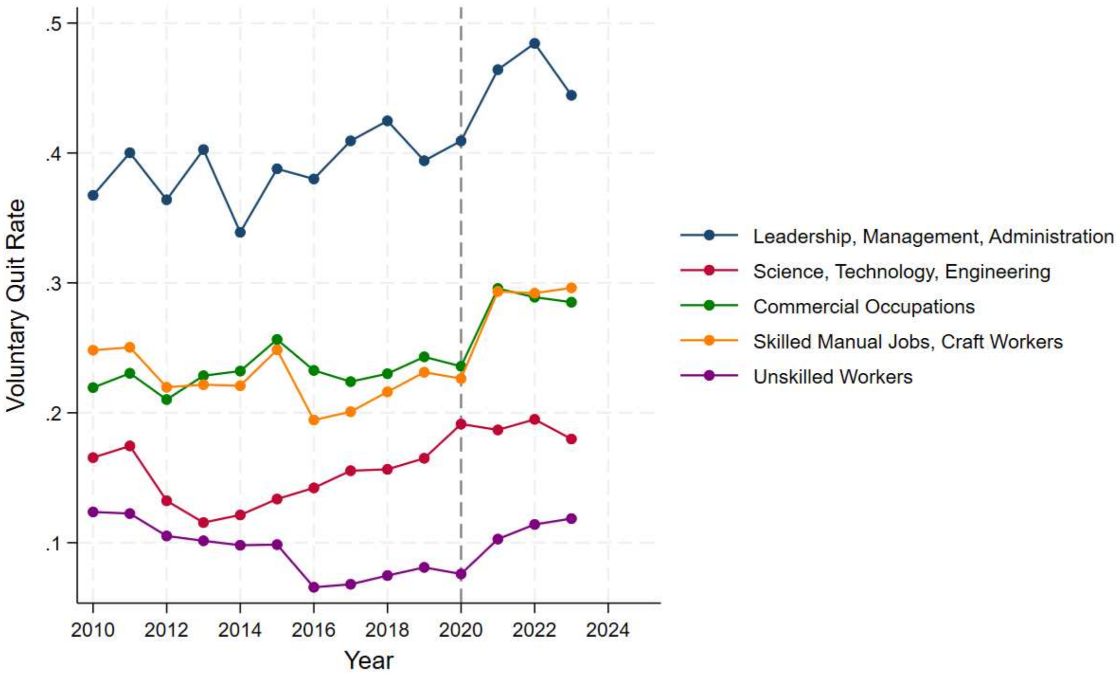

Notes. The figure displays the average quit rate over the years by occupation. The occupation classification is provided here at a very broad level. The quit variable is computed from the CICO.

FIGURE A10: Robustness Check using the Share of Hours Worked from Home

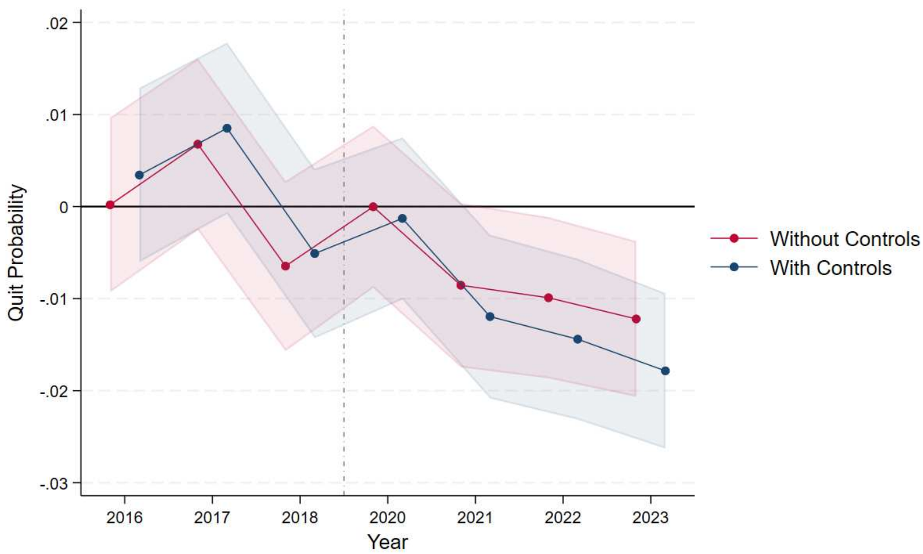

Notes. This figure shows the estimated event-study coefficients for the probability of voluntary quitting among teleworkable and non-teleworkable using the definition of teleworkability from Hensvik et al., 2020. The red line represents estimates without controls, while the blue line includes controls for workers' socio-economic characteristics, region, and industry effects. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

TABLE A4: Placebo Test

	(1)
	Voluntary Quits
Placebo Treatment × Post	0.0000769
	(0.000269)
Controls	✓
Year FEs	✓
Region FEs	✓
Sectors FEs	✓
Occupation FEs	✓
Observations	7601453

Notes. This table presents the results of the placebo test, where the treatment group is randomly assigned among occupations that should not be affected by WFH exposure. Controls for workers' characteristics, year fixed effects, region fixed effects, and sector fixed effects are included. Standard errors (in parentheses) are clustered at the worker level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

1.8 Appendix B: Heterogeneity Analysis

Differences in wages

FIGURE B1: Impact of WFH on Wages by Gender

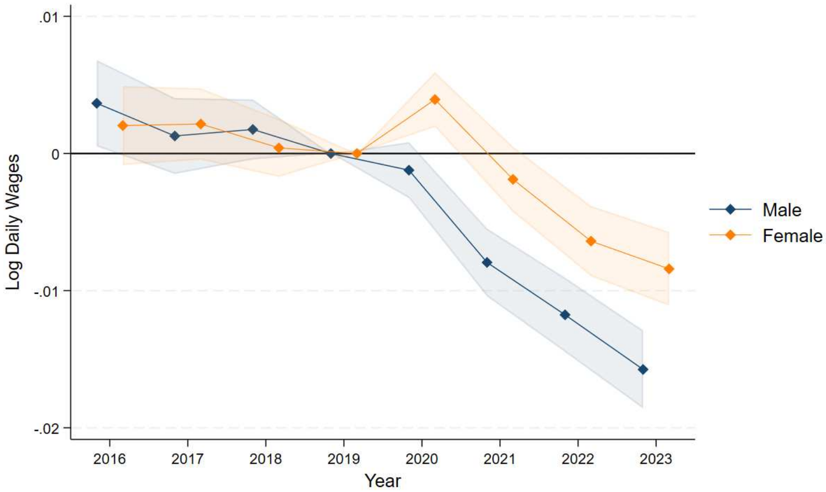

Notes. This figure shows the estimated event-study coefficients from (1.3) for the log daily wages among teleworkable and non-teleworkable occupations over time by gender. The orange line represents estimates for females in teleworkable occupations, while the blue line shows the estimates for men in teleworkable occupations. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

FIGURE B2: Impact of WFH on Wages by Education

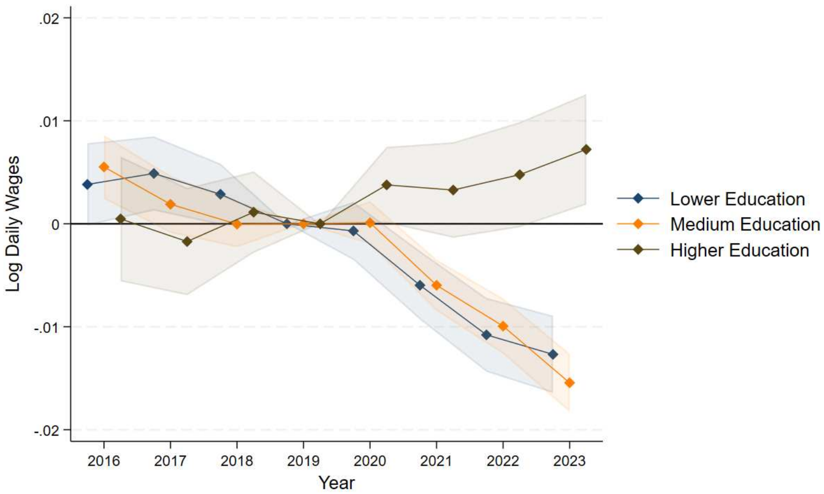

Notes. This figure shows the estimated coefficients from (1.3) for the log daily wages among teleworkable and non-teleworkable occupations over time by education. The blue line represents estimates for lower-educated workers in teleworkable occupations; the orange line represents estimates for those with a medium level of education, while the green line shows the estimation for higher-educated workers. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

FIGURE B3: Voluntary Quit Elasticity Pre and Post Covid-19

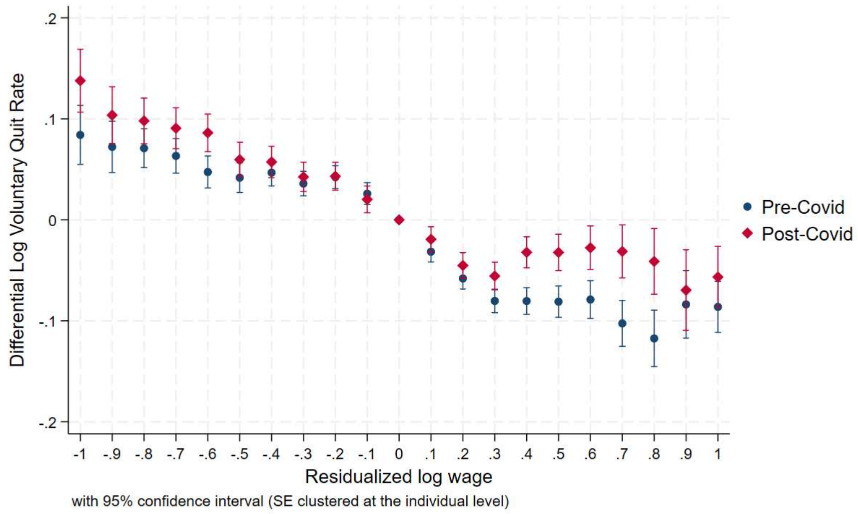

Notes. This figure shows the voluntary quit elasticity across the residualized log daily wage in the years before and after COVID-19. The *blue* elasticity is related to the year before the pandemic. The *red* elasticity refers to years just after. The year 2020 is excluded from the computation, so the Post-Covid quit elasticity regards only the period spanning 2021-2023.

FIGURE B4: Voluntary Quit Elasticity Teleworkable vs Non-Teleworkable

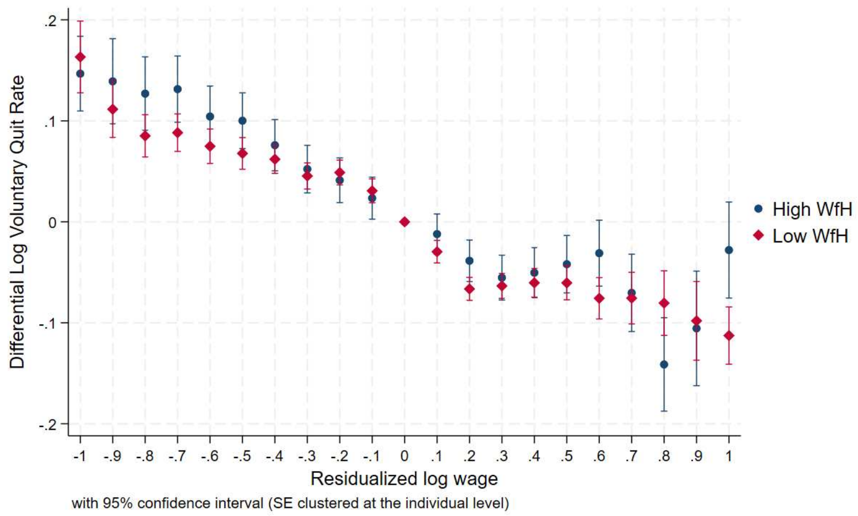

Notes. This figure shows the voluntary quit elasticity across the residualized log daily wage for teleworkable and not teleworkable. The *blue* elasticity is related to teleworkable. The *red* elasticity refers to non-teleworkable. The year 2020 is excluded from the computation, so the Post-Covid quit elasticity regards only the period spanning 2021-2023.

Differences in quitting

FIGURE B5: Probability of Quitting Teleworkable vs Non-Teleworkable by Gender

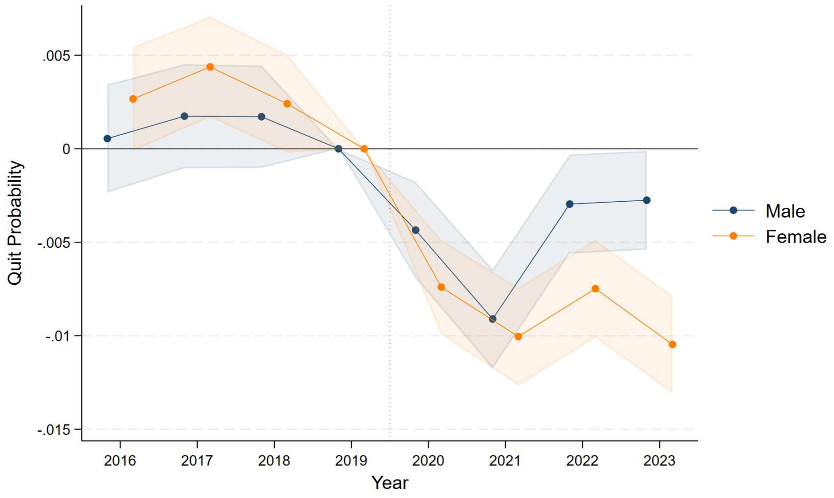

Notes. This figure shows the estimated event-study coefficients from (1.2) for the probability of voluntary quitting among teleworkable and non-teleworkable occupations by gender. The orange line represents estimates for females in teleworkable occupations, while the blue line shows the estimation for men in teleworkable occupations. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

FIGURE B6: Probability of Quitting Teleworkable vs Non-Teleworkable by Age

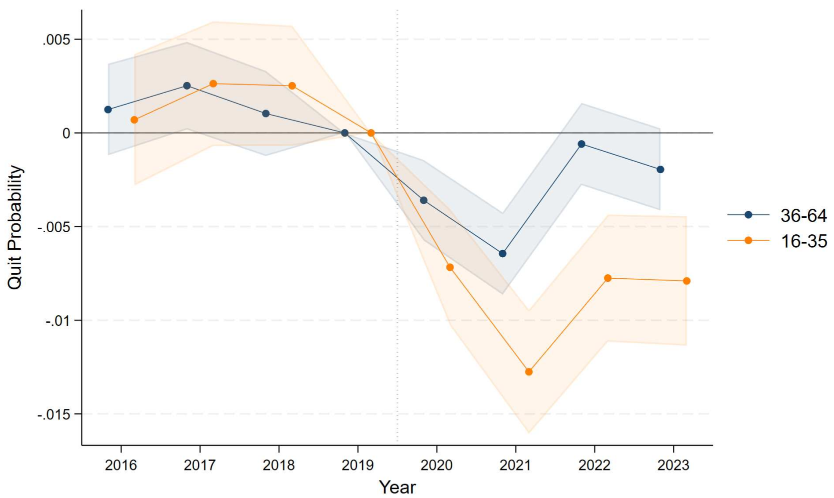

Notes. This figure shows the estimated event-study coefficients for the probability of voluntary quitting from (1.2) among teleworkable and non-teleworkable occupations by age. The orange line represents estimates for those 16-35 in teleworkable occupations, while the blue line shows the estimation for those older than 35. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

FIGURE B7: Probability of Quitting Teleworkable vs Non-Teleworkable by Education

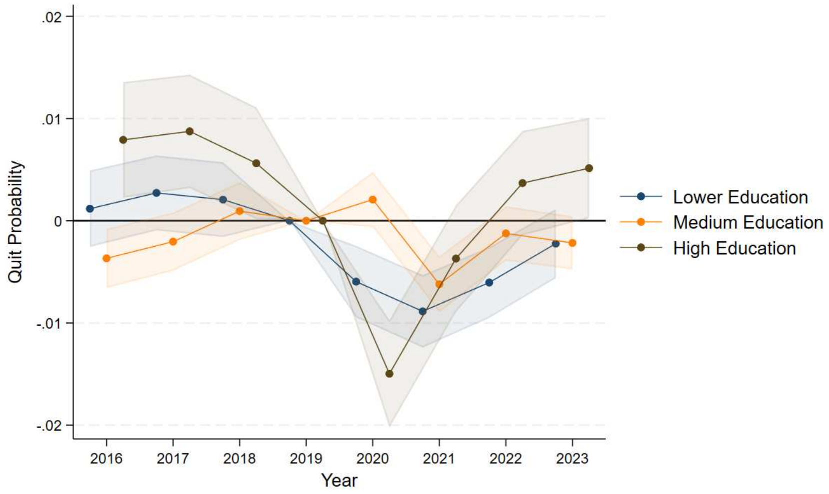

Notes. This figure shows the estimated event-study coefficients for the probability of voluntary quitting from (1.2) among teleworkable and non-teleworkable occupations by education. The orange line represents estimates for those with lower secondary education; the blue line shows the estimation for workers with upper secondary education, while the green line shows the estimation for those with college or more education. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

FIGURE B8: Probability of Quitting Teleworkable vs Non-Teleworkable by Type of Contract

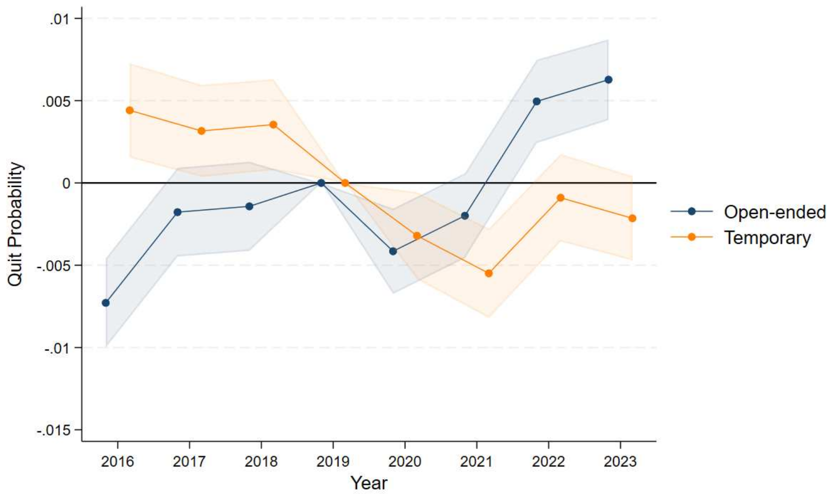

Notes. This figure shows the estimated event-study coefficients for the probability of voluntary quitting from (1.2) among teleworkable and non-teleworkable occupations by type of contract. The orange line represents estimates for workers with an open-ended contract in teleworkable occupations, while the blue line shows the estimation for those workers with temporary contracts in teleworkable occupations. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

FIGURE B9: Probability of Quitting Teleworkable vs Non-Teleworkable by Type of Contract (Part-time and Full-time)

Notes. This figure shows the estimated event-study coefficients for the probability of voluntary quitting from (1.2) among teleworkable and non-teleworkable occupations by type of contract, distinguishing between part-time and full-time. The orange line represents estimates for workers with a full-time contract in teleworkable occupations, while the blue line shows the estimation for those workers with a part-time contract in teleworkable occupations. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

FIGURE B10: Probability of Quitting Teleworkable vs Non-Teleworkable by Italian Regions

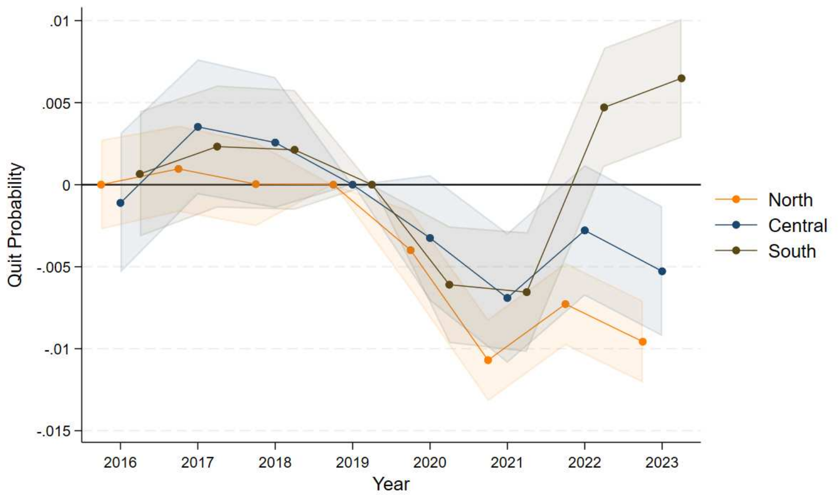

Notes. This figure shows the estimated event-study coefficients for the probability of voluntary quitting from (1.2) among teleworkable and non-teleworkable occupations by Italia regions. The orange line represents estimates for workers in the North, the blu line for those in the central, and the green line for those in the South of Italy. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

Chapter 2

The Impact of Retirement on Portfolio Allocation: Evidence from Italian Households

Nicol Barbieri and Nello Esposito *

This paper studies how retirement affects household portfolio allocation using rich panel data from the Italian Survey on Household Income and Wealth (SHIW) between 2012 and 2022. Exploiting exogenous variation in pension eligibility rules, we estimate a two-stage least squares with individual and time fixed effects to identify the causal impact of retirement on financial investment behavior. We find that retirement leads to a significant increase in stock holdings: the share of financial wealth invested in stocks rises by 8.2 percentage points, and the probability of participating in the stock market increases by 23 percentage points. This reallocation is mostly driven by publicly traded stocks. We interpret these findings through the lens of precautionary saving: as retirement reduces income risk exposure, the incentive to hold liquid, low-return assets weakens, and people shift toward higher-yielding investments. This effect is particularly pronounced for seniority-eligible individuals - i.e., those with longer, more stable career paths and higher pension benefits - while it is not statistically significant among old-age retirees, who are more likely to have experienced long unemployment spells and remain liquidity-constrained.

* *University of Naples Federico II, Department of Economics and Statistics, Italy*

We thank Tullio Jappelli and Monica Langella for their guidance, insightful feedback, and continuous support throughout this project. Special thanks to all PhD CSEF Seminar participants at the University of Naples Federico II for their inestimable contributions and suggestions. Nello Esposito gratefully acknowledges financial support from the Italian Fulbright Commission.

2.1 Introduction

Retirement is a turning point in a person's life: it entails not only changes in daily routines but also a substantial reorganization of financial resources. Although retirement is largely anticipated (i.e., people know when they are approaching pension eligibility), it still constitutes a significant income loss, given that pension benefits are typically lower than earnings during working life - especially for those, the majority in Italy, who do not hold a private pension plan on top of the public scheme (see, e.g., OECD, 2023). At the same time, retirement also entails a substantial reduction in income risk, as individuals stop facing employment uncertainty and their pension income becomes largely predetermined and stable over time - unlike labor earnings, which are more exposed to idiosyncratic shocks. Understanding how investment decisions adjust at retirement is therefore relevant from a policy perspective, as households' wealth reallocation at this stage of their life cycle can influence the demand for different asset classes, thereby affecting asset prices and macroeconomic fluctuations. This is even more salient in a country like Italy, where the share of retirees relative to the working population is expected to increase in the coming decades, and where household portfolios remain conservative, with strong preferences for real estate and low-risk investments (see, e.g., INPS, 2025).

From a theoretical perspective, retirement may be considered as an anticipated and permanent reduction in income. According to the Permanent Income Hypothesis (i.e., Friedman, 1957) and standard portfolio choice theory (i.e., Merton, 1971), such predictable changes should not alter either consumption or portfolio allocation. However, retirement also entails a significant decline in income-risk exposure, as people no longer face the possibility of job loss or unemployment spells, and pension benefits are generally more stable and less exposed to idiosyncratic shocks than labor earnings. This reduction in income risk is particularly relevant when markets are incomplete: income fluctuations play a critical role in shaping the strength of precautionary saving, i.e., the need to rely on self-insurance to smooth consumption in response to income shocks (see, e.g., Jappelli and Pistaferri, 2017). Therefore, as households transition from uncertain labor income to predictable pension benefits, the precautionary demand for safer assets may decline, potentially triggering rebalancing decisions toward higher-yielding investments. At the same time, other factors associated with aging - such as shorter investment horizons, increased liquidity needs, and possibly greater risk aversion - may work in the opposite direction, inducing more conservative portfolio choices.

As a result, how people rebalance their portfolios at retirement is still unclear *ex-ante* and remains an open empirical question. In this paper, we fill this gap, analyzing how retirement affects individuals' portfolio choices in the context of the Italian pension system.

To address the endogeneity of the retirement decision, we implement a two-stage least squares (2SLS) analysis in which pension eligibility - determined by the retirement schemes available in Italy between 2012 and 2022 - is used as an instrument for actual retirement. We also include individual and time fixed effects (TWFEs), allowing us to compare the same individual while working and after retirement, controlling for common time shocks and individual-specific unobserved heterogeneity that could affect portfolio choices over time.

Between 2012 and 2022, Italian workers had access to three main retirement pathways. The first, which we will refer to as *seniority*, allowed individuals to retire after accumulating 42 years of contributions for men and 41 for women. The second option, that we will call *old age*, allows retirement at age 67 with at least 20 years of contributions, for both men and women. Finally, between 2019 and 2022, individuals could retire with 38 years of contributions and a minimum age of 62 (raised to 64 in 2022); we will refer to this option as *early retirement*.

These eligibility rules will be used to construct our instrumental variable. Specifically, we define an individual as eligible for retirement if, in a given year, she meets the requirements for at least one of the retirement schemes described above. We then use eligibility as an instrument for actual retirement in our 2SLS estimation.

Our identification strategy relies on the fact that eligibility is determined by institutional rules and thus exogenous to individual portfolio decisions, while strongly predicting the probability of retiring, once conditioning on time-invariant unobserved heterogeneity. Indeed, a potential concern is that eligibility rules are partly determined by accumulated years of contributions, which, in turn, reflect individuals' career histories. Since career stability and earnings trajectories may be correlated with portfolio decisions, this could generate endogeneity in our setting. In our 2SLS estimation, we address this issue in two ways: first, we rely on a TWFEs strategy, which allows us to compare the same individual across working and retirement, thereby absorbing time-invariant unobserved heterogeneity related to career choices. Second, in the robustness analysis, we explicitly control for long unemployment spells, which capture the main observable dimension of career instability.

We draw data from the Bank of Italy's *Survey of Household, Income, and Wealth* (SHIW). This is a biannual survey[2] that collects detailed information on the wealth composition of a representative panel of Italian households. The survey also asks respondents to report their years of contribution - i.e., the number of years they have worked under a regular contract - which allows us to construct retirement eligibility as defined above.

Our results show that retirement leads to a rebalancing of household portfolio allocation: retirees distribute a larger share of their financial assets to stocks and are more likely to enter the stock market. This shift suggests that households are taking on greater financial risk, moving away from safe, liquid assets such as deposits and bonds. Importantly, these results reflect a reallocation of existing financial wealth rather than an increase in overall saving: our specifications control for total household wealth, ensuring that the estimated effects capture changes in portfolio composition rather than wealth accumulation. Particularly, our findings suggest that retirement increases the share of financial assets invested in stocks by 8.2 percentage points.

[2] The SHIW is typically conducted every two years, with the exception of the 2016 and 2020 waves.

It also raises the probability of stock market participation by 23 percentage points. This effect holds for publicly traded stocks and managed savings funds, but not for private company stocks. Retirement increases the share of financial assets held in publicly traded stocks by 7 percentage points and the likelihood of holding such assets by 23 percentage points. Moreover, the share of stocks plus managed savings - i.e., mutual funds, pension funds, and other professionally managed investment products - increases by 12 points, while the likelihood of holding them increases by 28 percentage points. By contrast, we do not find significant effects on the share of financial assets held in private equity.

Taken together, our results show that the decline in income risk at retirement potentially outweighs other opposing forces, such as a decline in income levels or shorter investment horizons: retirees reallocate their wealth away from safe, liquid assets toward higher-yielding yet accessible instruments, including publicly traded equities and managed savings. No significant effect is observed for private equity, which remains an illiquid and less accessible investment option.

These patterns observed in our findings can be rationalized by a precautionary saving mechanism. While working, households need to rely on precautionary savings to ensure against income fluctuations (see, for instance, Kimball, 1990 and Dynan, 1993). As argued by Auclert et al., 2024, the financial instrument better suited for this purpose is a low-return, safe asset with no transaction costs, like bonds or bank accounts. The transaction costs associated with stocks, mostly related to market liquidity (see on this point Foucault et al., 2023), make them a relatively bad insurance instrument. Yet, when people stop working and their pension benefits are computed, there is no more uncertainty about earnings and the probability of losing their jobs. Therefore, the precautionary motive weakens, and the incentive to get higher financial returns strengthens. This is consistent with existing evidence on the role of income risk and liquidity in shaping saving behavior. When labor income is uncertain, households rely on liquid and low-return assets to self-insure against income shocks. As argued by Kaplan et al., 2014, transaction costs associated with illiquid assets such as housing give rise to the so-called "wealthy hand-to-mouth" consumers - i.e., individuals who hold substantial illiquid wealth but very little liquidity, and therefore do not smooth transitory income shocks because the adjustment costs exceed the size of the shocks themselves. In the same spirit, Graves, 2025 shows that households use liquid assets to absorb income shocks and that the consumption response to unemployment is much smaller among those holding more liquid than illiquid wealth.

These studies reinforce the idea that income risk is a key determinant of portfolio liquidity at retirement: as income risk declines, households have less need for precautionary liquidity and can reallocate wealth toward higher-return and illiquid assets. Then, lower exposure to income risk and higher liquidity weaken the precautionary motive for saving and strengthen the desire for high-return investments.

To validate this interpretation, we perform a reduced-form analysis in which we regress the outcomes of interest on the different retirement options. We find that the effects are statistically significant only for individuals eligible for *seniority* retirement and, occasionally, for *early retirement*. This pattern can be explained by the fact that "regular" retirees (i.e., *seniority* eligible) have presumably had more stable career paths and received higher severance payments[3]. Therefore, they are further away from the borrowing constraint and need to rely less on precautionary savings to begin with. By contrast, *old-age* retirees tend to have more fragmented career trajectories and enter retirement with lower levels of liquid wealth and smaller severance payments. The key mechanism here relates to liquidity rather than income: even if their pension income is stable, their limited cash on hand keeps them closer to the borrowing constraint. We provide supporting evidence for this mechanism by showing that *old-age* retirees are more likely to have experienced long unemployment spells (more than 12 months) and to hold lower levels of liquid wealth - namely, real estate, financial assets, and bank deposits. This pattern is consistent with the interpretation that a weaker liquidity position sustains a stronger precautionary saving motive among this retirees' group.

To further explore possible heterogeneity in our results, we run separate 2SLS regressions across different subgroups.

We first split the sample into two groups based on net wealth distribution and find that the effect of retirement on portfolio rebalancing is significant only for individuals above the wealth distribution median. This result is consistent with Fagereng et al., 2018, who show that higher income risk reduces investment in stocks mainly for those below the median of the wealth distribution. These results potentially suggest that households with fewer liquidity constraints and greater financial flexibility are more responsive to the reduction in income risk that accompanies retirement. By contrast, individuals in the lower half of the wealth distribution are more likely to remain close to the borrowing constraint even after retirement and therefore continue to prioritize liquid, safe assets over higher returns.

We also find a stronger response among individuals residing in the North of Italy, for married households, and for women. These groups may enjoy better access to financial markets or additional sources of insurance, such as spousal income, which can make them more able or willing to bear financial risk. Interestingly, the slightly larger effect observed for women at the extensive margin contrasts with the existing literature, which typically finds that men are more likely to participate in the stock market and hold risky assets (see, for instance, Eckel and Grossman, 2008; Croson and Gneezy, 2009; Baghai, Soares, et al., 2023).

Overall, these observed patterns point to a consistent mechanism: the portfolio rebalancing at retirement is typically stronger among groups facing fewer financial frictions and greater opportunities to adjust their investment behavior.

[3] Workers in Italy receive a severance payment upon retirement, which could further relax liquidity constraints. This one-time transfer increases households' cash holdings (see, for instance, Battistin et al., 2009), making them better able to smooth consumption and self-insure against shocks without holding excessive amounts of low-yield liquid assets.

Finally, we perform a series of robustness tests to address potential identification threats and validate the credibility of our 2SLS strategy. We design these tests to ensure that the estimated effect of retirement on portfolio allocation is not driven by unobserved heterogeneity in career paths, wealth accumulation, or age composition. We aim to rule out that our instrument captures systematic differences in financial behavior unrelated to actual retirement.

After confirming the instrument's relevance, we introduce a large set of controls that might explain portfolio choices, such as past long unemployment spells and other sources of income, including financial and real estate income.

We also perform an event study analysis to verify that portfolio choices follow parallel trends prior to eligibility, ruling out pre-existing differences in investment behavior. Moreover, we re-estimate the model excluding individuals younger than 55, those with positive liabilities, and those with negative net wealth, to test whether the results are driven by outliers or by individuals close to financial distress. Finally, we restrict the sample to observations closer to the eligibility threshold, both before and after reaching it, to ensure that our estimates are not influenced by individuals far from retirement age.

Across all these tests, although the sample size shrinks considerably - and hence statistical power decreases - the 2SLS coefficients remain very close to the baseline estimates, both in magnitude and statistical significance. This stability reinforces the credibility of our identification strategy and suggests that our findings are not sensitive to alternative specifications or sample definitions.

Our paper is connected to different strands of literature and contributes to them along several dimensions. First, it mirrors the literature on the *retirement consumption puzzle* (see Banks et al., 1998). Recent empirical work by Kolsrud et al., 2024 and Olafsson and Pagel, 2024 in Scandinavian countries has shown an economically and statistically significant decline in consumption upon retirement. Results like these, observed in Italy as well by Battistin et al., 2009, are at odds with standard economic theory, as consumers are supposed not to vary consumption in response to predicted income variations. Aguiar and Hurst, 2005 rationalizes this puzzle by noting that there is a difference between consumption and expenditure, as the former also includes home production and the depreciation of durables (like TVs, laptops, cars, etc.), evidently not as precisely measured as expenditure by survey data. Another explanation for the puzzle is the presence of complementarity between labor and some consumption items. Commuting costs are a perfect example: as people retire, they no longer need to commute to work, and thus, transportation expenditure decreases. We speak to this literature by shifting the focus from consumption to investment behavior[4]. While much of the existing literature has examined how retirement affects spending patterns, less attention has been paid to how households adjust the composition of their wealth when income risk declines. Understanding

[4] In Appendix D, we provide evidence that the drop in consumption concerns only old-age retirees, who, for the reasons outlined above, are more likely to be liquidity constrained. There is, on the other hand, no convincing evidence of a drop in consumption for seniority and early retirees, who are on average further away from the borrowing limit and are more likely to behave as Permanent Income Hypothesis (PIH) maximizers.

this adjustment is crucial, as portfolio reallocation not only determines long-run household financial security but also affects aggregate asset demand, prices, and macroeconomic fluctuations. Rather than studying how households spend in retirement, we examine how they reallocate their portfolios in response to changes in income level and risk. Our findings are not puzzling in light of standard theory: life-cycle portfolio choice models with idiosyncratic income risk and incomplete markets predict increased stock market participation once individuals exit the labor force and receive stable, predictable pension income (see, e.g., Cocco et al., 2005). In this sense, our contribution complements the consumption literature by providing new evidence on how households adjust their financial choices in response to reduced income uncertainty.

Second, we contribute to the extensive literature in consumption and household finance on the impact of income risk on consumption and portfolio choice. The papers in this strand, either by estimating structural models or relying on natural experiments, test the implications of the precautionary saving model and the buffer stock model[5]. As mentioned in the very first lines, retirement also represents a reduction in income risk exposure, and therefore, our results might be seen as the reduced-form impact of income risk on portfolio choice. In this regard, our paper is deeply connected to Fagereng et al., 2018, who look at the impact of uninsured wage risk on participation in the stock market in Norway. Using employee-employer matched and tax record data, and applying the methodology of Guiso et al., 2005 to extract the component of wage volatility not insured by the employer, they show that risk reduces stock holdings, and that the effect is greater for households belonging to the bottom of the wealth distribution. Guiso et al., 1996 find, in the 1989 SHIW cross section, a negative correlation between holdings of risky assets and subjective income variance, elicited through survey questions. Angerer and Lam, 2009 show that portfolio choices are mostly affected by permanent income volatility, while they do not respond to transitory income risk. Clark et al., 2022, leveraging on the 2012 labor market reform in Italy, shows that lower employment protection causes households to invest more in safer assets and reduce their exposure to risk assets (e.g., stocks). Silva et al., 2025 show that there is a negative effect of going public (i.e., employers' IPOs) on wage volatility and a positive impact on the holdings of risky assets. We differentiate from these papers in that we focus on the variation in risk exposure due to retirement: in other words, we are comparing the same individuals before and after retirement, isolating the effect of income risk reduction within the life cycle.

Third, we contribute to the literature on households' participation in the stock market and how this changes in response to wealth shocks. It is a well-known fact that the stock market participation among households is low. Van Rooij et al., 2011 point to low financial literacy (measured by "The Big Three" as in Lusardi, 2008) as the main reason in the United States. Moreover, Lusardi et al., 2017 shows that differences in financial literacy are among the main drivers of cross-sectional wealth inequality in the United States. Fagereng et al., 2017 and

[5] See Jappelli and Pistaferri, 2017 for a review of the main theoretical and empirical contribution and Jappelli and Pistaferri, 2025 for a recent test of the buffer-stock model by Carroll and Samwick, 1997 and Deaton, 1991.

Fagereng et al., 2020 also shows limited stock market participation in Norway. Christelis et al., 2024 estimates the causal effect of wealth shocks on stock holdings by assigning randomly hypothetical lottery wins to the respondents to the Consumer Expectation Survey (CES). They find that the increase in stock market participation, even though statistically significant, especially for lottery gains as high as 50.000 €, is contained (around 8.4 percentage points). Also, the share of wealth invested in stocks (the intensive margin) increases only slightly (by 1.7 percentage points for a hypothetical wealth shock of 50.000 €). This result indicates that the assumption of Constant Relative Risk Aversion (CRRA) preferences is a reasonable one for most investors, as heterogeneous stock market participation is mostly explained by risk aversion heterogeneity[6]. We contribute to this literature first by showing that retirement increases stock market participation in Italy and then that the estimated causal effect of retirement on stock market participation remains stable even when controlling for household wealth, suggesting that our results are not driven by mechanical wealth changes, but rather by shifts in risk exposure associated with the transition to retirement.

Deeply connected to our paper are also the works that study the impact of pension reforms in the 1990s implemented in Italy on retirement expectations (see Bottazzi et al., 2006), savings and portfolio composition (see, for instance, Attanasio and Brugiavini, 2003, Bottazzi et al., 2011, and Daminato and Padula, 2024). Unlike these studies, however, we do not look at the impact of pension reforms affecting people who are currently working, but at how stock holdings change in the transition from work to retirement.

Finally, we also speak to the literature that looks at the macro-finance consequences of population aging. Auclert et al., 2021 argues that the aging in the population will increase the wealth-to-income ratio, and thus lower interest rates. Our results suggest that a larger fraction of the population will participate in the stock market and increase their holdings on the intensive margin as well. Therefore, stock market returns will decline due to higher demand from retail investors, pushing prices up.

The paper is structured as follows. Section 2.2 describes the data, the retirement options we use to build the instrument, and the sample selection. Section 2.3 depicts the empirical strategy and the baseline results. We also discuss the validity of the instrument, while Section 2.4 shows the main mechanisms that rationalize our results and some heterogeneity analysis. Finally, Section 2.5 concludes.

2.2 Data and Retirement Options

We draw on the Bank of Italy's Survey of Household Income and Wealth (SHIW), a comprehensive dataset that provides insights into the financial lives of Italian households. SHIW

[6] In the standard portfolio choice as in Merton, 1969, the share of wealth invested in risky assets is not a function of wealth itself. It is a decreasing function of risk aversion and the variance of the asset return.

contains detailed information on demographics, income, consumption, and wealth. In particular, respondents are asked to self-report in detail their holdings of stocks (both public and private), bonds, bank accounts, and other forms of financial investments. They are also asked to report their liabilities, as well as the value of the real estate properties they own. Moreover, SHIW also contains information on individuals' years of contribution, i.e., how many years the respondents have worked under a regular contract, which is crucial in order to construct our instrumental variable.

We select the sample following the approach by Jappelli and Pistaferri, 2020; Jappelli and Pistaferri, 2025. For starters, we only keep household heads[7] who were employed in 2012. We define an individual as retired if she is not working and receives a retirement pension[8]. Then, to ensure we are tracking the same individuals over time, we drop observations with inconsistent demographics. Therefore, we drop households whose head changes sex over time and with impossible ages[9]. For instance, if a person is 30 years old in 2012, she must be 32 in 2014, 34 in 2016, 38 in 2020, and 40 in 2022. To minimize the concerns related to the measurement error of (self-reported) years of contribution, we only keep those individuals for whom the increase in the years of contribution is lower than or equal to the time difference between two consecutive waves. We drop individuals who report fewer years of contribution than the preceding wave. After this sample selection criteria, we are left with a sample size of around 2000 observations, pooled from the waves of 2012, 2014, 2016, 2020, and 2022.

In the period under analysis, people could leave the labor market following three different paths to retirement. According to the first, that we will refer to as *seniority*, workers might retire, regardless of their age, as soon as they accumulate 42 years of contribution if they are men and 41 if they are women. The second path, that we will call *old age*, requires workers, both women and men, to be at least 67 years old and to have 20 years of contribution. Finally, between 2019 and 2021, a third option known as *Quota 100* (Q100) was introduced. According to Q100, people as old as 62 and with 38 years of contribution could decide to retire. In 2022, Q100 was replaced by *Quota 102* (Q102), meaning 64 years of age and 38 years of contribution. We will refer to these two last options together as *early retirement*, as in fact they represented, in general, a milder eligibility criterion, at the cost of lower pension wealth upon retirement.

People could choose which option to use and, typically, their choice is dictated by their careers, as well as their earnings trajectories during their working history. For example, seniority

[7] Restricting the analysis to heads prevents complications arising from changes in household composition or from the identity of the head switching across survey waves.

[8] Our dataset contains detailed information about the nature of the pension income that can come from several sources, including old-age, seniority, disability, and survivor pensions. In this analysis, we exclude individuals receiving disability or survivor pensions and retain only those receiving old-age or seniority pensions, as these are directly linked to retirement decisions and eligibility.

[9] These cases likely reflect changes in the identity of the household head across waves or reporting errors, rather than actual demographic variation. We exclude them to ensure that the same individual is consistently tracked over time.

represents the typical option used to retire, and is the one that entails the biggest pension benefit[10]. This is because, if a person has a continuous working life, not characterized by long unemployment spells, she reaches the eligibility criteria for seniority in her early 60s, even earlier if she started working in her teenage years. By contrast, old age entails a much lower pension wealth, and it is typically "chosen" by individuals who have experienced large unemployment spells during their lives. This potentially reflects the fact that irregular career paths often prevent individuals from accumulating the required contribution years for more generous schemes. In the Italian public pension system, pension benefits are commonly calculated based on both the number of contribution years and the worker's past earnings, according to a contribution-based or mixed formula, depending on the year of entry into the labor market. Under the contribution-based regime, which applies to workers who started contributing after 1996, the benefit equals the sum of all contributions paid over the working life, capitalized at an annual rate linked to GDP growth, multiplied by a conversion coefficient that depends on the retirement age. Workers with earlier contribution histories fall under a mixed regime that combines earnings-related and contribution-based components, resulting in higher replacement rates.

Age and years of contribution are what we need to define eligibility for the three retirement options mentioned above: seniority, old age, and early retirement. In particular, we define:

$$
\begin{aligned}
seniority_{it} &= \begin{cases} 1 \quad \text{if} & yc_{it} \geq 42 \text{ and } female_i = 0 \\ 1 \quad \text{if} & yc_{it} \geq 41 \text{ and } female_i = 1 \\ 0 \quad \text{otherwise} & \end{cases} \\
old\ age_{it} &= \begin{cases} 1 \quad \text{if} & age_{it} \geq 67 \text{ and } yc_{it} \geq 20 \\ 0 \quad \text{otherwise} & \end{cases} \\
early_{it} &= \begin{cases} 1 \quad \text{if} & yc_{it} \geq 38 \text{ and } age_{it} \geq 62 \text{ and } t = 2020 \\ 1 \quad \text{if} & yc_{it} \geq 38 \text{ and } age_{it} \geq 64 \text{ and } t = 2022 \\ 0 \quad \text{otherwise} & \end{cases}
\end{aligned}
\tag{2.1}
$$

where yc_{it} is the years worked by individuals i at time t. Following the eligibility rules defined above, we construct our instrumental variable for actual retirement as follows:

[10] According to INPS, 2025, in 2024, roughly 30% of new pensions in Italy were seniority/early retirement pensions, confirming the widespread relevance of this option among workers with continuous careers.

$$eligible_{it} = \begin{cases} 1 & \text{if} \quad seniority_{it} = 1 \\ 1 & \text{if} \quad old\ age_{it} = 1 \\ 1 & \text{if} \quad early_{it} = 1 \\ 0 & \text{otherwise} \end{cases} \tag{2.2}$$

In this setting, we implement a 2SLS regression where retirement status is instrumented by eligibility as defined above, controlling for both individual time shocks. Our identification strategy relies on the assumption that, conditional on individual and time fixed effects, eligibility for retirement affects portfolio decisions only through actual retirement.

The use of an instrumental variable is motivated by the potential endogeneity of the retirement decision in the portfolio composition. Retirement status may be correlated with unobserved factors that also influence portfolio choices, such as unmeasured expectations about future income, health shocks, or individual preferences toward risk and saving. For instance, risk-tolerant workers might both retire earlier and invest more in stocks, leading to an upward bias in OLS estimates. Using pension eligibility as an instrument isolates the exogenous variation in retirement induced by institutional rules, allowing us to recover the causal effect of retirement on portfolio allocation.

Figure 2.1 shows the take-up ratios, i.e., the ratio between the eligible individuals and the actual retirees, for the three retirement options.

FIGURE 2.1: Retirement Options Take Up

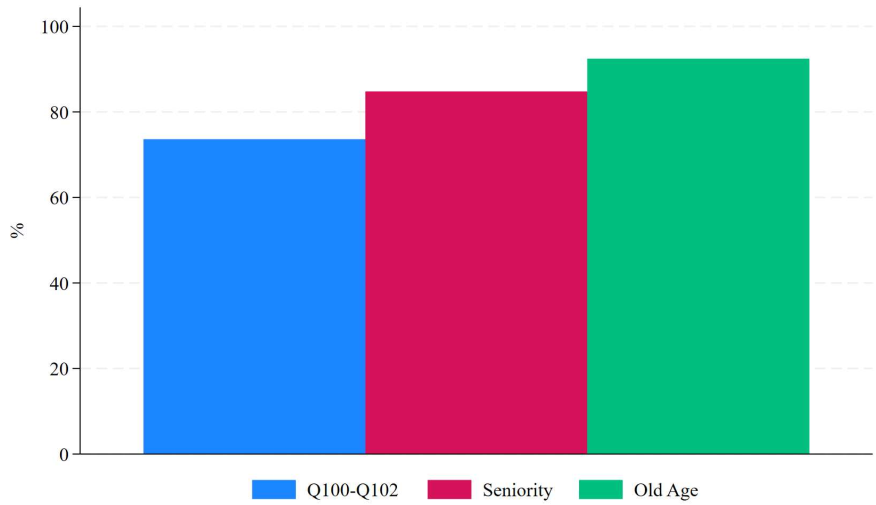

Notes. This figure displays the take-up rates for the three retirement pathways: early retirement (blue bars), seniority-based retirement (red bars), and old-age retirement (green bars). Note that early retirement was available only between 2019 and 2022, corresponding to the Quota 100 and Quota 102 schemes (Q100-Q102).

As we can see, almost 88% of old age-eligible individuals are retired, while the take-up ratio goes down to slightly above 80% for seniority. The share of early eligible individuals who are retired is below 80%.

The main message of Figure 2.1 is that the majority of people who are eligible for retirement do choose to retire, consistently with Battistin et al., 2009.

To define eligibility, we assume the following assignment rule (or preference order) among the three schemes: (1) seniority, (2) early, and (3) old age. If a person satisfies the requirements for more than one option in a given year, we assign her to the highest-ranking scheme for which she is eligible. This rule is not meant to reflect individual preferences, but rather the institutional structure and observed behavior in the data. In practice, workers tend to prioritize seniority retirement whenever possible, as it provides a more generous pension benefit and does not restrict post-retirement employment, unlike early retirement. By contrast, old-age pensions are typically accessed by individuals who cannot meet the contribution requirements for the other schemes. This institutional background helps explain the observed take-up patterns: the take-up rate for early retirement is lower than for seniority. At the same time, old-age pensions represent a residual option for workers with fragmented careers[11].

Figure 2.2 displays the relationship between actual retirement status and the distance to retirement eligibility.

FIGURE 2.2: Actual Retirement vs Eligibility

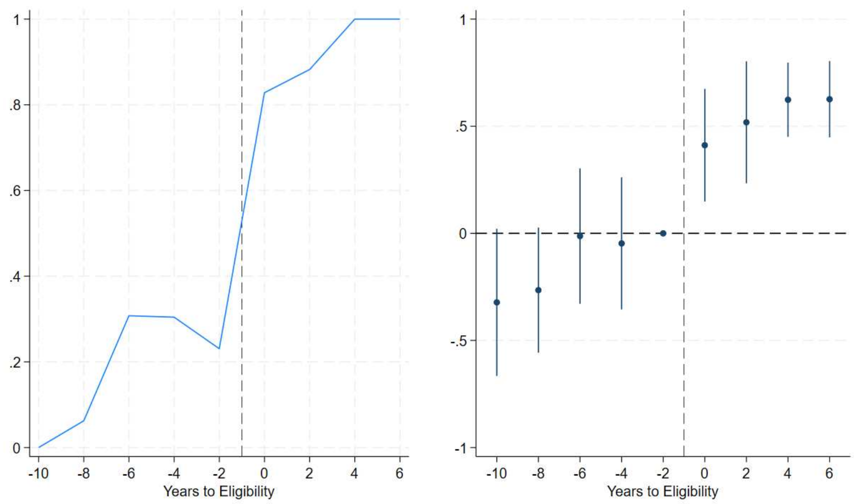

Notes. The left panel of this Figure plots the relationship between actual retirement and the distance to eligibility. The vertical dashed line indicates the retirement threshold (distance = -1). The sample includes individuals within a window of ±10 years around eligibility. The right panel shows the event study version of the left panel.

[11] As a robustness in Appendix A, we perform the main analysis by dropping this assumption and allowing individuals to be contemporaneously eligible for more options. Results are practically unchanged. The main reason is that there are very few people in the sample eligible for more than one option at the same time.

The figure shows that there is an overall positive relationship between distance to retirement eligibility and actual retirement. In particular, the share of retirement-eligible individuals who are in retirement is 100% six years later. This pattern confirms the predictive power of eligibility for actual retirement behavior in a way that validates the relevance of our instrument.

Figure 2.2 also speaks to the plausibility of the monotonicity assumption: it is unlikely that anyone in the sample would have retired if they had not been eligible, while not retiring after becoming eligible remains a rare behavior.

In Table 2.1, we report the descriptive statistics for non-eligible, early retirement-age eligible, seniority-eligible, and old-age eligible.

TABLE 2.1: Summary Statistics from SHIW data by Eligibility Group

	Not Eligible		Early		Seniority		Old Age	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Age	50.52	8.51	64.27	1.36	64.67	3.48	68.62	2.90
Married	0.72	0.45	0.83	0.38	0.62	0.49	0.73	0.45
Family Size	2.95	1.30	2.13	0.73	1.98	0.87	2.10	0.82
Centre	0.21	0.41	0.20	0.41	0.26	0.44	0.27	0.45
North	0.49	0.50	0.53	0.51	0.66	0.48	0.27	0.45
Medium Education	0.77	0.42	0.83	0.38	0.82	0.39	0.63	0.48
High Education	0.20	0.40	0.17	0.38	0.15	0.36	0.25	0.43
Years of Contribution	25.42	10.18	40.10	1.45	42.98	1.15	36.21	5.95
Retired	0.04	0.20	0.73	0.45	0.79	0.41	0.89	0.32
Nondurable Consumption	28.26	13.79	23.64	10.44	28.57	13.30	24.75	12.41
Cash on Hand (1000 €)	81.57	149.60	103.05	104.23	109.79	85.88	76.26	80.68
Experienced long unemployment	0.21	0.41	0.03	0.18	0.03	0.18	0.15	0.36
Net Wealth (1000 €)	274.68	313.97	315.77	266.32	405.81	375.39	261.17	264.05
Financial Assets (1000 €)	39.72	138.13	59.79	89.39	68.79	77.47	40.71	69.00
Share of Stocks (x 100)	1.69	7.96	2.71	10.37	4.70	11.84	0.92	5.55
Participate in the Stock Market	0.08	0.28	0.10	0.31	0.21	0.41	0.04	0.18

Notes. This table reports descriptive statistics for individuals classified by their eligibility for different retirement schemes: not eligible, early retirement, seniority, and old age. Variables include different socio-economic and demographic characteristics, labor market experience, and financial indicators. Long unemployment refers to unemployment spells lasting more than 12 months. Education is divided into three groups: low education (primary school or less), medium education (lower or upper secondary school), and high education (tertiary education or above).

As the first column shows, 4% of the sub-sample of the non-eligible is retired. This is because, sometimes, people benefit from early retirement options offered in case of mass layoffs or business restructuring. As a robustness check, we also perform the analysis dropping these observations. What we get is a much stronger first stage, but no significant changes in the results. Results are reported in Appendix A.

Seniority-eligible individuals hold more cash on hand (income + financial assets) than old age- and early-eligible, and there are fewer of them who declare they have experienced long unemployment spells (more than 12 consecutive months). Moreover, seniority-eligible individuals hold a higher share of stocks and are more likely to participate in the financial market than other retirement-eligible individuals. In Appendix B we provide additional descriptive evidence that complements our analysis, including information on retirement behavior and portfolio composition of our sample.

2.3 Empirical Strategy and Baseline Results

Our empirical strategy consists of a 2SLS, in which we instrument actual retirement with retirement eligibility defined as in the above section. Our main equation is:

$$y_{it} = \alpha_i + \lambda_t + \beta retired_{it} + \varphi \mathbf{X}_{it} + \varepsilon_{it} \tag{2.3}$$

In (2.3), y_{it} is either the share of stocks (listed and unlisted) held in i's portfolio at time t as a share of gross financial wealth, or a dummy equal to 1 if i holds stocks at time t. $\mathbf{X}_{it}$ is a vector of controls including age bins (to control for the life cycle), education, family size, civil status, and, in the most demanding specification, net wealth, i.e,. the sum of financial assets, liabilities, and real estate assets. The variable $retired_{it}$ is a dummy equal to 1 if i is retired at time t. α_i and λ_t are individual and time fixed effects. The former are meant to purge the analysis from time-invariant characteristics affecting retirement and stock holdings at the same time. The first example that came to our mind is whether an individual suffered from financial losses in the past that scared her and forced her to postpone retirement. The latter, on the other hand, capture common shocks that affect individuals in a presumable homogeneous way, like, for instance, bad news about the sustainability of the National Social Security Institute. Finally, ε_{it} is the error term clustered at the individual level, to account for serial correlation in unobserved shocks over time for each individual.

Estimating (2.3) by OLS is going to lead to an inconsistent estimate of β because people choose when (and how) to retire, and this choice is likely to be correlated with time-variant unobservable characteristics affecting portfolio decisions as well. One potential source of endogeneity arises from unobservable preferences for work. Individuals who enjoy working or derive nonmonetary satisfaction from their jobs are likely to postpone retirement. At the same time, because they work longer, they accumulate higher lifetime wealth, which may translate into a greater ability to invest in riskier assets. Then, the observed association between retirement and portfolio reallocation would partly reflect differences in lifetime earnings driven by unobserved preferences rather than the causal effect of retirement itself.

Another problem is reverse causality: it might be that an individual gets a very high return from the financial market, which convinces her to stop working. Georgarakos et al., 2023, randomly providing individuals with hypothetical lottery gains, show that people would respond by reducing labor supply both on the intensive and the extensive margin. Cesarini et al., 2017 and Fagereng et al., 2021, in Sweden and Norway, respectively, show that people winning lotteries persistently reduce their gross and net earnings.

To tackle these issues, we instrument actual retirement with eligibility. Pension eligibility is determined exclusively by institutional rules that depend on age and contribution history, and not on individual preferences, expectations, or financial behavior. Conditional on individual and time-fixed effects, these institutional thresholds are exogenous to unobservable factors such as a taste for work or expectations about future income risk, as well as to shocks to financial returns that might directly influence portfolio choices. In other words, eligibility affects portfolio decisions only through its impact on the probability of retiring, satisfying the exclusion restriction.

The first stage we are going to estimate is then the following:

$$retired_{it} = \alpha_i + \lambda_t + \gamma eligible_{it} + \rho \mathbf{X}_{it} + v_{it} \tag{2.4}$$

In (2.4), $eligible_{it}$ is a dummy defined as in the section above. Our coefficient of interest is β, that is, the LATE of retirement on the portfolio, i.e., the effect on the compliers. In this context, the compliers are defined as those individuals who retired *because* they became eligible, so they satisfy the age and years of contribution requirements for at least one of the three retirement options.

Table 2.2 shows the baseline results for all stocks, both on the intensive (first three columns) and the extensive (last three columns) margins.

The effect of retirement on the share of stocks is around 10 percentage points (portfolio shares are multiplied by 100), while the effect on stock holdings is around 22 percentage points. Adding controls makes little difference for the intensive margin, while it makes estimates much more precise for the extensive margin. Finally, the addition of wealth doesn't change practically anything: the coefficient attached to wealth is both imprecisely estimated and economically insignificant (a 1000 € increase in net wealth increases the share of stocks by 0.01 percentage points).

This result speaks in favor of Merton, 1971 and is consistent with Christelis et al., 2024, suggesting that the driver of households' heterogeneity in portfolio composition is mainly driven by heterogeneity in risk aversion as opposed to differences in wealth levels.

TABLE 2.2: The Effect of Retirement on Stock Holdings - All stocks

	Intensive Margin			Extensive Margin		
	(1)	(2)	(3)	(4)	(5)	(6)
Retired	5.17***	8.04***	8.04***	0.15*	0.23**	0.23**
	(1.90)	(2.65)	(2.65)	(0.08)	(0.11)	(0.11)
Age 30-39		0.66	0.65		0.00	0.01
		(1.04)	(1.05)		(0.03)	(0.03)
Age 40-49		0.86	0.86		0.06	0.06
		(1.78)	(1.78)		(0.06)	(0.06)
Age 50-60		0.51	0.51		0.06	0.06
		(2.60)	(2.60)		(0.09)	(0.09)
Over 60		-2.61	-2.61		-0.04	-0.04
		(2.89)	(2.89)		(0.10)	(0.10)
Family Size		1.78**	1.78**		0.02	0.02
		(0.84)	(0.84)		(0.02)	(0.02)
Married		-1.81*	-1.80*		-0.03	-0.03
		(1.04)	(1.05)		(0.06)	(0.06)
Medium Education		0.88	0.89		-0.02	-0.03
		(1.24)	(1.26)		(0.05)	(0.05)
High Education		2.64	2.64		0.06	0.06
		(1.98)	(1.99)		(0.09)	(0.09)
Net Wealth (1000 €)			-0.00			0.00*
			(0.00)			(0.00)
Observations	1086	1086	1086	1214	1214	1214
F-stat 1st stage	74.54	36.71	36.69	79.88	40.92	40.88

Notes: This table reports the 2SLS estimates of equation (2.3). In columns (1)–(3), the dependent variable is the share of total financial assets held in stocks (intensive margin). In columns (4)–(6), the dependent variable is a dummy equal to 1 if the household holds any stocks (extensive margin). Retired is a dummy equal to 1 if the individual is retired and 0 otherwise. Because retirement status may be endogenous to portfolio decisions, it is instrumented using retirement eligibility, defined in equation (2.2) based on age and years of contribution. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.3 and Table 2.4 decompose the result for public stocks and private stocks.

TABLE 2.3: The Effect of Retirement on Stock Holdings - Public stocks

	Intensive Margin			Extensive Margin		
	(1)	(2)	(3)	(4)	(5)	(6)
Retired	3.95**	6.66***	6.67***	0.13*	0.22**	0.22**
	(1.89)	(2.55)	(2.54)	(0.08)	(0.10)	(0.10)
Age 30-39		0.95	0.77		0.00	0.00
		(1.12)	(1.08)		(0.03)	(0.03)
Age 40-49		2.45	2.45		0.06	0.06
		(1.93)	(2.12)		(0.06)	(0.06)
Age 50-60		3.29	3.27		0.06	0.06
		(2.83)	(2.98)		(0.08)	(0.08)
Over 60		1.78	1.79		-0.03	-0.03
		(3.15)	(3.29)		(0.09)	(0.09)
Family Size		1.23*	1.24*		0.02	0.02
		(0.73)	(0.72)		(0.02)	(0.02)
Married		2.45	2.73		0.00	-0.00
		(2.65)	(2.93)		(0.06)	(0.06)
Medium Education		0.59	0.77		-0.02	-0.02
		(1.66)	(1.79)		(0.05)	(0.05)
High Education		1.96	2.10		0.06	0.06
		(2.39)	(2.48)		(0.09)	(0.09)
Net Wealth (1000 €)			-0.00			0.00
			(0.00)			(0.00)
F-stat 1st stage	74.54	36.71	36.69	79.88	40.92	40.88
Observations	1086	1086	1086	1214	1214	1214

Notes: This table reports the 2SLS estimates of equation (2.3) using retirement eligibility as an instrument for retirement status. The dependent variable is the share of public stocks held in total financial assets (columns 1–3, intensive margin) and a dummy for holding any public stocks (columns 4–6, extensive margin). Retired equals 1 if the individual is retired and 0 otherwise. and it is instrumented with retirement eligibility, defined as in equation (2.2). Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 2.4: The Effect of Retirement on Stock Holdings - Private stocks

	Intensive Margin			Extensive Margin		
	(1)	(2)	(3)	(4)	(5)	(6)
Retired	1.21	1.39	1.38	0.02	0.01	0.02
	(0.97)	(1.60)	(1.60)	(0.03)	(0.04)	(0.04)
Age 30-39		-0.29	-0.12		0.00	0.01
		(0.78)	(0.72)		(0.01)	(0.01)
Age 40-49		-1.59	-1.59		0.00	0.00
		(1.48)	(1.72)		(0.02)	(0.02)
Age 50-60		-2.78	-2.77		-0.01	-0.00
		(1.90)	(2.13)		(0.02)	(0.02)
Over 60		-4.39**	-4.40*		-0.02	-0.01
		(2.04)	(2.30)		(0.03)	(0.03)
Family Size		0.55	0.54		-0.00	-0.00
		(1.04)	(1.05)		(0.01)	(0.01)
Married		-4.26	-4.54		-0.03	-0.03
		(3.11)	(3.42)		(0.03)	(0.03)
Medium Education		0.29	0.12		-0.00	-0.00
		(0.71)	(0.83)		(0.01)	(0.01)
High Education		0.68	0.54		0.00	0.00
		(0.91)	(1.04)		(0.01)	(0.01)
Net Wealth (1000 €)			0.00			0.00
			(0.00)			(0.00)
Observations	1086	1086	1086	1214	1214	1214

Notes: This table reports the 2SLS estimates of equation (2.3) using retirement eligibility as an instrument for retirement status. The dependent variable is the share of private stocks held in total financial assets (columns 1–3, intensive margin) and a dummy equal to 1 if the household holds any private stocks (columns 4–6, extensive margin). Retired equals 1 if the individual is retired and 0 otherwise and it is instrumented with retirement eligibility, defined as in equation (2.2). Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.5 looks at stocks and managed savings (i.e, mutual and pension funds) together.

TABLE 2.5: Stocks + Managed Savings

	Intensive Margin			Extensive Margin		
	(1)	(2)	(3)	(4)	(5)	(6)
Retired	7.24***	10.70***	10.70***	0.17*	0.25**	0.25**
	(2.60)	(3.70)	(3.70)	(0.10)	(0.12)	(0.12)
Age 30-39		1.06	0.95		0.01	0.01
		(1.33)	(1.37)		(0.04)	(0.04)
Age 40-49		1.60	1.60		0.11	0.11
		(2.25)	(2.28)		(0.08)	(0.08)
Age 50-60		1.65	1.64		0.11	0.11
		(3.25)	(3.26)		(0.10)	(0.10)
Over 60		-2.02	-2.01		0.02	0.02
		(4.01)	(4.02)		(0.11)	(0.11)
Family Size		2.52***	2.52***		0.05**	0.05**
		(0.93)	(0.93)		(0.02)	(0.02)
Married		-4.01**	-3.85**		-0.12	-0.12
		(1.64)	(1.66)		(0.09)	(0.09)
Medium Education		0.81	0.91		-0.02	-0.02
		(1.95)	(2.02)		(0.06)	(0.06)
High Education		-3.10	-3.01		-0.01	-0.02
		(4.43)	(4.48)		(0.06)	(0.06)
Net Wealth (1000 €)			-0.00			0.00
			(0.00)			(0.00)
Observations	1086	1086	1086	1214	1214	1214

Notes: This Table shows the 2SLS estimates of equation (2.3). From columns 1 to 3, the dependent variable is the ratio between the sum of stocks and managed savings held by household i at time t over financial assets (intensive margin). From columns 4 to 6, the dependent variable is a dummy equal to 1 if i owns either stocks or shares of various forms of managed savings (extensive margin). *Retired* is a dummy equal to 1 if i is retired 0 otherwise, and is instrumented by *Eligible*, that is a dummy defined as in (2.2). Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Looking at the effect in Table 2.3 and Table 2.4, the effect of retirement is present only for the former, both on the intensive and the extensive margin. Specifically, retirement raises the share of public stock holdings by around 7.12 percentage points and increases the probability of owning public stocks by 23 percentage points. The result on the extensive margin gives further credibility to our empirical setting. Indeed, the causal effect on the portfolio share might be contaminated by heterogeneous variation in the prices of these stocks. It could still be, nevertheless, that people decide to enter (exit) the stock market because of its good (bad) performance. However, the presence of time FEs nets out any possible confounder related to overall time trend. In addition, the overall performance of the stock market is surely not correlated with the variables that define retirement eligibility: age and years of contribution.

Regarding the effect on stocks and managed savings, Table 2.5 shows a positive impact on both the intensive and the extensive margins, amounting to 11.7 percentage points and 28 percent, respectively.

To interpret these estimates as causal effects, we rely on an instrumental variable strategy that exploits retirement eligibility as an exogenous source of variation in retirement status. For the instrument to really estimate the Local Average Treatment Effect (LATE) of retirement on stock holding, it has to satisfy two requirements. The first is relevance: we need retirement eligibility to predict actual retirement. The second requirement, the exclusion restriction, is that the instrument affects stock holdings only through its effect on retirement. If the instrument were not relevant, we would expect no strong statistical relationship between eligibility and actual retirement - that is, becoming eligible would not significantly predict the probability of retiring. To test this, it suffices to look at the statistical relationship between eligibility and actual retirement. Then, we estimate the first stage in equation (2.4) by OLS and report the results in Table 2.6.

TABLE 2.6: First Stage Estimation

	(1)	(2)	(3)
Eligible for Retirement	0.54^{***}	0.42^{***}	0.42^{***}
	(0.06)	(0.07)	(0.07)
Socio-Economics Controls	·	✓	✓
Net Wealth (1000 €) Controls	·	·	✓
Observations	1214	1214	1214

Notes. This table reports the first-stage estimates from equation (2.4), where actual retirement status is regressed on retirement eligibility. Column (1) includes only individual and time fixed effects. Column (2) adds demographic controls. Column (3) further includes net wealth. Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The instrument is strongly correlated with actual retirement: being eligible for retirement makes actual retirement more likely by 50 percentage points, controlling for individual and time FEs, as well as for the same variables we control for in equation (2.3). The F-stat of the excluded instrument test ranges between 68 and 75, i.e., well above the benchmark recommended by the literature for the relevance of an instrument. In line with Battistin et al., 2009, Figure 2.2 - left panel - shows that the probability of being retired increases as people get closer to retirement. The share of retired 10 years before eligibility is 0, while it is 1 4 years after retirement. Moreover, increases by 60 percentage points from -2 to 0 years to eligibility, confirming that most people retire as soon as they are eligible. The right panel, showing the event study version of the left panel, conveys the same message. It also certifies that the negative different at -4 and -2 are not statistically distinguishable from the coefficient at -6, since their confidence intervals overlap.

The first argument we provide is that pension eligibility is built upon variables that are very difficult to manipulate: age and years of contribution. Regarding the latter, the Italian law allows converting the years spent in higher education (e.g., University) and in the compulsory military service as years of contribution, in exchange for a front-loaded *una tantum* payment to the Social Security Institute[12]. In principle, people might select into this option according to unobservable characteristics that also influence investment decisions. If this is the case, the 2SLS coefficient might still yield an inconsistent estimate. We address this concern by including in the analysis only workers whose years of contribution increased by no more than 2 years, that is, the frequency of SHIW[13]. This also allows us to deal with potential measurement error in the years of contribution, a problem also acknowledged by Battistin et al., 2009.

The second concern is the presence of unobservable heterogeneity that affects both eligibility and portfolio choices. It could indeed be the case that individuals retire at different times because of career choices and paths we cannot observe in the available data. We respond in two ways. First, we always include individual fixed effects, which purge the analysis of time-invariant heterogeneity that affects both careers and consumption/investment decisions. In addition, as a further robustness check (see Appendix A), we control for a dummy variable indicating whether the individual experienced an unemployment period of more than 12 months. Results are essentially unchanged.

We then check for pre-trends with respect to eligibility. We want to test whether stock holdings differed systematically before people became eligible for retirement. We therefore estimate the following event study equation both for the intensive and the extensive margin of portfolio choice, where D_{-k} and D_k represent the distance, before and after, to retirement eligibility:

[12] This option is called in Italian *riscatto della Laurea*. Importantly, during the 2012-2022 period, no major reform affected the cost or the accessibility of the "risccatto della Laurea" scheme, implying that the possibility to convert education years into contributory years did not vary systematically across individuals.

[13] Between 2016 and 2020, we drop observations whose years of contribution either went down or increased by more than 4 years.

$$y_{it} = \alpha_i + \lambda_t + \sum_{k=2}^{10} \gamma_{-k} D_{-k} + \sum_{k=2}^{6} \gamma_k D_k + \rho \mathbf{X}_{it} + u_{it} \tag{2.5}$$

Figure 2.3 shows the estimated coefficients $\hat{\gamma}_{-k}$ and $\hat{\gamma}_k$, alongside 95% confidence intervals.

FIGURE 2.3: Pre-Trends in Stock Market Outcomes

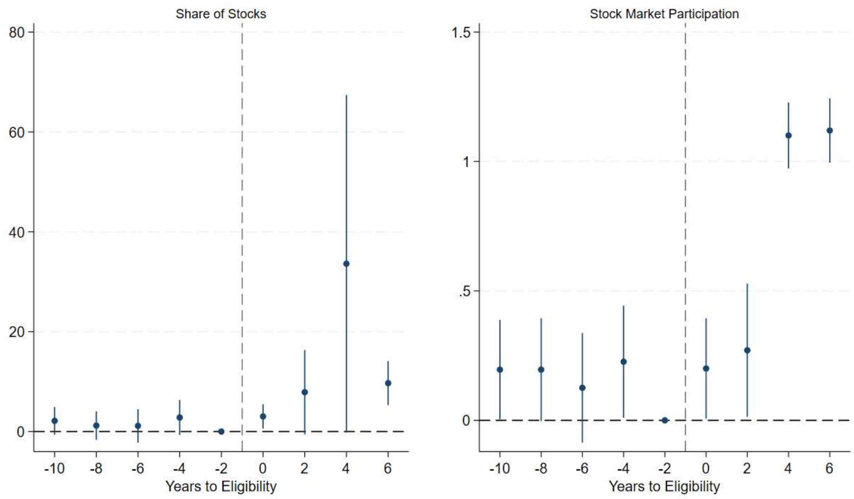

Notes. This figure shows the OLS estimated coefficients of equation (2.4). The left panel plots the evolution of the share of financial wealth held in stocks, while the right panel shows the probability of participating in the stock market. Each point represents the average difference relative to the year of eligibility (year -2), controlling for individual and time fixed effects. The vertical lines represent 95% confidence intervals.

The figure plots the evolution of the share of stocks held (left panel) and the probability of participating in the stock market (right panel) as a function of the number of years before and after eligibility (ranging from 10 years before to 6 years after becoming eligible for retirement). The p-value of the test of all pre-trend coefficients being jointly equal to 0 is 0.12 for the intensive margin and 1.61 for the extensive margin. In other words, the pre-eligibility period, the OLS coefficients are close to zero and statistically insignificant, suggesting the absence of pre-trends in portfolio allocation. Immediately after eligibility, however, there is a sharp increase: the stock share jumps as well as the stock market participation. Therefore, we cannot reject the null hypothesis of no effect in the years before individuals become eligible for retirement.

As a final test, we look at returns from financial and real estate assets before actual retirement and before eligibility. The concern behind this test is that systematic differences in asset returns, due to particularly good (or bad) investments, might be the true responsible of the changes in the portfolio composition. For instance, if individuals who retire happen to experience unusually high returns on their financial assets, this could mechanically inflate their stock holdings, confounding our estimates.

To address this concern, we estimate (2.4) for returns from real estate (e.g., rents) and returns from financial investments, and report the results in Figure 2.4.

FIGURE 2.4: Pre-trends in Real Estate and Financial Income

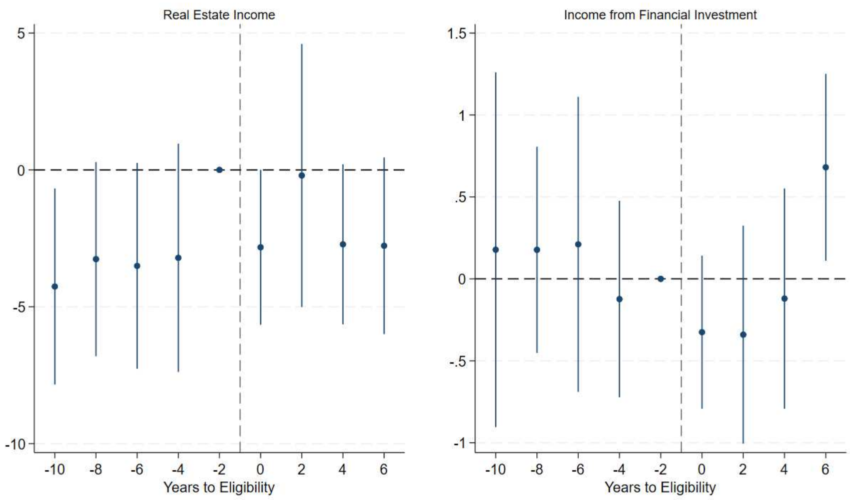

Notes. This figure shows the OLS estimated coefficients of equation (2.4). The left panel plots the evolution of the real estate income, while the right panel shows the income from financial investment. Each point represents the average difference relative to the year of eligibility (year -2), controlling for individual and time fixed effects. The vertical lines represent 95% confidence intervals.

Overall, there is no clear trend before eligibility. In Appendix A, we also control for both types of income, and the results do not change. This exercise helps rule out differential asset returns as a confounding factor and strengthens the interpretation that the observed reallocation toward stocks is a behavioral response to retirement-induced changes in income risk. As a further check, in Appendix A, we perform the main analysis dropping households with negative net wealth, people with financial liabilities, and people younger than 55 years old. Results are practically unchanged in their economic and statistical significance.

In Appendix C, we show the results obtained using different identification strategies. In particular, we estimate, via the Regression Discontinuity Design (RDD) approach, an equation in which we interact the retirement dummy with years distant to eligibility. We allow for both a linear and a quadratic function of the distance to eligibility.

Results in Table C1 and Table C3 display the results for the intensive margin using a linear and a quadratic function, respectively, for the distance to eligibility. The estimated coefficients are close to the baseline estimates. Results concerning the extensive margin (see Table C2 and Table C4 with linear and quadratic polynomials, respectively) are not statistically different from 0. Then, Figure C1 performs a McCrary test (see McCrary, 2008) of the distribution of the years of contribution around 38 (the threshold for *early*) and 42 (the threshold for *seniority*).

Both tests cannot reject the null that there is no discrete jump in the density around the cut-offs, therefore suggesting that there is no manipulation of the eligibility criteria on the side of the years of contribution. Moreover, Figure C2 shows the estimated coefficients of the TWFE and 2SLS, including observations within close retirement eligibility bandwidths: between -4 and 4 years, -6 and 6 years, -8 and 8 years. Results preserve both economic and statistical significance, suggesting that they are not influenced by the behavior of people far away from eligibility.

2.4 Driving Mechanism and Heterogeneity

In this section, we dig deeper into the baseline result and investigate which eligibility category is driving it. To this end, we perform a reduced-form analysis, in which we regress the outcome variable against the three eligibility dummies defined in Section 2.2. In particular, we estimate by OLS the following equation and report the estimated coefficients in Table 2.7.

$$y_{it} = \alpha_i + \lambda_t + \theta early_{it} + \delta seniority_{it} + \omega old\ age_{it} + \kappa \mathbf{X}_{it} + v_{it} \tag{2.6}$$

The three eligibility dummies are mutually exclusive. If an individual qualifies for more than one option in a given year, we assign her to the highest-ranking scheme according to the preference order defined in Section 2.2, namely seniority, early retirement, and old-age. This ensures that each person-year observation is classified under one, and only one, eligibility rule.

Looking at the intensive margin (Panel A) of Table 2.7, only the coefficient attached to $seniority_{it}$ is positive and statistically different from zero. In particular, being eligible for seniority retirement is associated with a share of stocks (listed and unlisted) over financial wealth, higher by around 10 percentage points. On the other hand, the coefficient attached to *early*, even though sometimes close to the one of *seniority*, is seldom statistically different from zero. To conclude, the coefficient attached to *old age* is never statistically and economically different from zero.

We interpret this result as follows. Seniority-eligible retirees and, in part, early-eligible retirees have had, on average, a more stable career path. A more stable career path thus translates into a lower realized earning volatility over the life cycle, a higher level of cash-on-hand upon retirement, and then a weaker precautionary saving motive. For this reason, when they retire, they are more willing to invest in stocks to earn higher returns. Additionally, their financial situation upon retirement is more predictable, allowing them to treat retirement as a planned, rather than reactive, transition. This forward-looking behavior likely reinforces their willingness to re-optimize their portfolio allocation, taking advantage of their improved ability to bear financial risk.

TABLE 2.7: Reduced-Form Analysis

	Stocks - Total	Stocks - Listed	Stocks + Managed
Panel A: Intensive Margin			
Seniority Eligible	6.46***	6.21***	7.82***
	(1.86)	(1.94)	(2.48)
Q100 Eligible	4.53	3.54	8.25**
	(3.01)	(2.87)	(3.30)
Old Age Eligible	1.03	0.59	2.11
	(1.00)	(1.00)	(2.04)
Observations	1086	1086	1086
Panel B: Extensive Margin			
Seniority Eligible	0.24***	0.24***	0.22**
	(0.08)	(0.08)	(0.09)
Q100 Eligible	0.09	0.09	0.20**
	(0.10)	(0.10)	(0.10)
Old Age Eligible	0.01	0.00	0.02
	(0.05)	(0.05)	(0.07)
Observations	1214	1214	1214

Notes: This Table shows OLS estimates of equation (6). In Panel A, the dependent variable is the share of stocks over financial assets. In Panel B, the dependent variable is a dummy equal to 1 if i holds stocks at time t. $Senioriy_{it}$, $Early_{it}$, $Old\ Age_{it}$ are dummies defined as in (2.1). Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

On the other hand, old age-eligible individuals have had, on average, a highly unstable career path, perhaps with irregular contracts and/or long unemployment spells. Therefore, upon retirement, they have much lower cash-on-hand levels and still have a strong precautionary motive for saving. As a result, these individuals tend to maintain a more conservative asset allocation, prioritizing liquidity and capital preservation over risky investments.

Table 2.8 shows the correlation between eligibility for old-age retirement and wealth, conditioning on eligibility for at least one retirement option.

TABLE 2.8: Old Age Eligibility and Wealth

	Long Term Unemployment	Total Wealth	Real Estate	Financial Wealth	Bank Account
Old Age Eligible	0.12***	-114.95***	-93.43**	-25.12**	-8.83**
	(0.04)	(42.50)	(36.93)	(10.49)	(4.43)
Observations	205	205	205	205	205

Notes: This table shows the correlation, among those who are eligible for retirement, between old age eligibility and, the probability of being in unemployment for more than 12 months (first column), total wealth (second column), real estate wealth (third column), financial wealth and bank account (columns 4 and 5 respectively). Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The first column shows that being eligible for old-age retirement increases the probability of having experienced long-term unemployment spells (more than 12 months) by 13 percentage points. Moreover, old age-eligible individuals hold less total wealth, less real estate wealth, fewer financial assets, and hold less money in their bank accounts. These correlations suggest that individuals eligible for old age retirement have a lower level of cash on hand and, therefore, are more sensitive to income fluctuations.

In Appendix C we show that consumption drops, upon retirement, only for old-age retirees. This result is driven by the exact mechanism proposed for stock holdings, but in reverse. Old age retirees are, on average, on lower levels of cash on hand and more likely to be liquidity constrained, and thus sensitive to predicted income variations.

The baseline results mask noticeable heterogeneity along different dimensions, as shown by Figure A1 to Figure A2. The first of them we look at is wealth. We split the sample into two, according to whether wealth is above or below the median. The effect is statistically significant only for households above the median. This is true both for the intensive and the extensive margin. This result is in line with Fagereng et al., 2018, who show that wealthy people do not reduce their investment in risky assets in response to a higher exposure to uninsured wage risk.

The second dimension we look at is geographical location: North (the richest and developed Region of Italy) versus Center-South (relatively underdeveloped). The effect is statistically different from 0 only in Northern Italy. This happens because retirees in the North are, on average, richer and have experienced smoother career paths. Therefore, as the intuition outlined above commands, they should increase their investments in stocks.

In addition, we look at the differential effect of whether the individual is married or not. The rationale for this heterogeneity test is that married couples provide additional insurance against income fluctuations and therefore weaken the precautionary saving motive. In line with this conjecture, we find that the effect is only significant for married couples, while it's insignificant for singles. Finally, we look at how the estimated effects differ between males and females. We find that the effect is stronger for women than for men.

2.5 Conclusion

In this paper, we analyze how households' portfolios change upon retirement. We address the endogeneity of retirement by instrumenting it with eligibility to three retirement options available in Italy between 2012 and 2022: seniority, early, and old age retirement. We find that the share of stocks over financial assets (i.e., the intensive margin of stock holdings) increases by 8.4 percentage points after retirement, and that the probability of participating in the stock market (i.e., the extensive margin) increases by 23 percentage points. We interpret this result as follows: when people retire, they cease to be exposed to income risk, since their pension is known with certainty and insured from macroeconomic fluctuations. Moreover, they receive from their employer the severance payment, increasing their liquidity and increasing their ability to smooth consumption in the occurrence of shocks. Therefore, their need for precautionary savings decreases, increasing their financial risk-bearing capacity.

To informally argue on the goodness of exclusion restriction, we show that stock market participation exhibits no pre-trends before retirement. We then control for financial returns, as well as for real estate income. Results are essentially unchanged.

To gauge insights on the driving mechanism and understand who drives the main results, we perform a reduced-form analysis in which we regress outcomes against the single retirement eligibility options. We find that the effect is driven mainly by seniority and early retirees who have had a more stable career path, are on higher levels of cash on hand, and are therefore better-positioned to absorb financial shocks. On the other hand, old-age retirees are more likely to have experienced long unemployment spells, hold less cash and financial assets. For this reason, they still need to self-insure against the borrowing constraints.

This paper offers an important policy message for a country like Italy, where the population of retirees relative to the working population is expected to rise in the forthcoming decades. A more active stock market participation by the public, driven by retirees, will probably mean a more efficient and more dynamic financial market. Our findings highlight the role of retirement design in influencing household financial behavior. As public pension systems evolve, understanding how retirement affects asset allocation is crucial for enhancing individual financial resilience and market stability.

2.6 Appendix A: Additional Figures and Tables

Heterogeneity Analysis

FIGURE A1: Heterogeneity by Wealth Halves

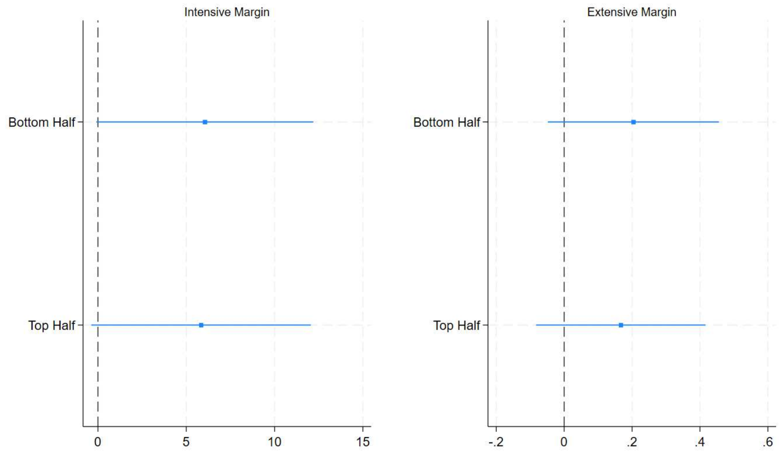

Notes. This figure shows the estimated effect of retirement on stock holdings separately for individuals in the bottom and top halves of the net wealth distribution. The left panel refers to the intensive margin (share of wealth invested in stocks), while the right panel refers to the extensive margin (stock market participation). Each point represents the 2SLS estimate of the retirement effect, with horizontal lines indicating 95% confidence intervals.

FIGURE A2: Heterogeneity by Geographical Area

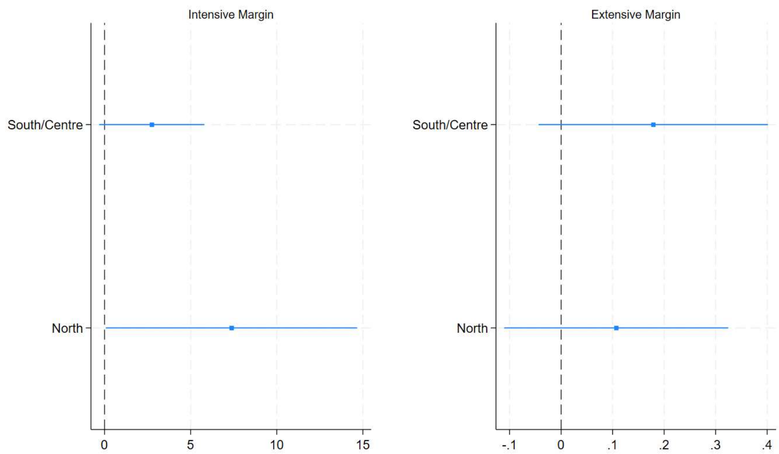

Notes. This figure shows the estimated effect of retirement on stock holdings separately for individuals in the North and South of Italy. The left panel refers to the intensive margin (share of wealth invested in stocks), while the right panel refers to the extensive margin (stock market participation). Each point represents the 2SLS estimate of the retirement effect, with horizontal lines indicating 95% confidence intervals.

FIGURE A3: Heterogeneity by Marital Status

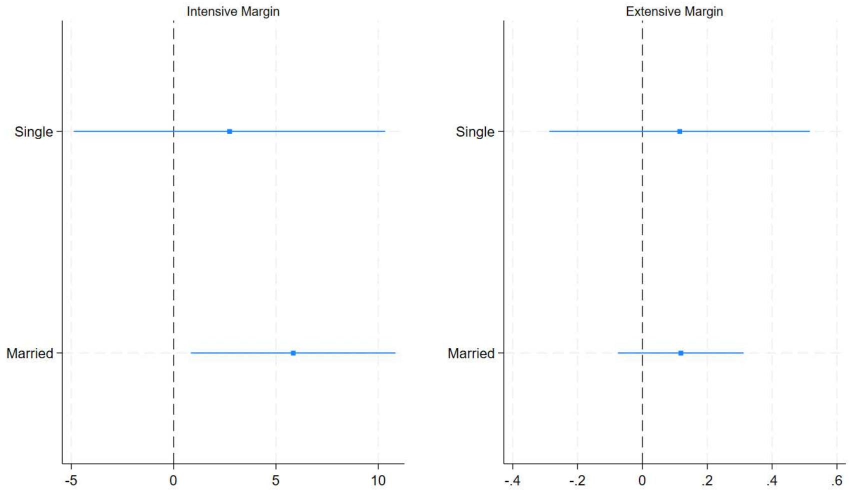

Notes. This figure shows the estimated effect of retirement on stock holdings separately for singles and married people. The left panel refers to the intensive margin (share of wealth invested in stocks), while the right panel refers to the extensive margin (stock market participation). Each point represents the 2SLS estimate of the retirement effect, with horizontal lines indicating 95% confidence intervals.

FIGURE A4: Heterogeneity by Gender

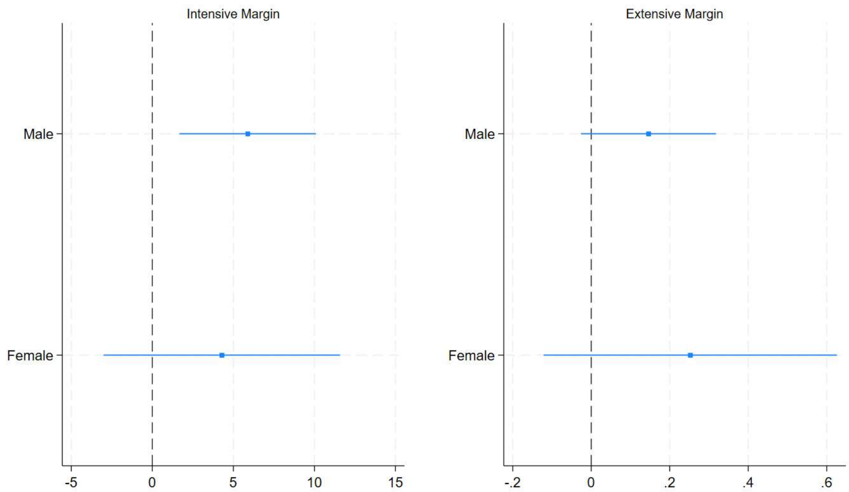

Notes. This figure shows the estimated effect of retirement on stock holdings separately for both male and female. The left panel refers to the intensive margin (share of wealth invested in stocks), while the right panel refers to the extensive margin (stock market participation). Each point represents the 2SLS estimate of the retirement effect, with horizontal lines indicating 95% confidence intervals.

FIGURE A5: Financial and Real Estate Annual Incomes Before Retirement Eligibility

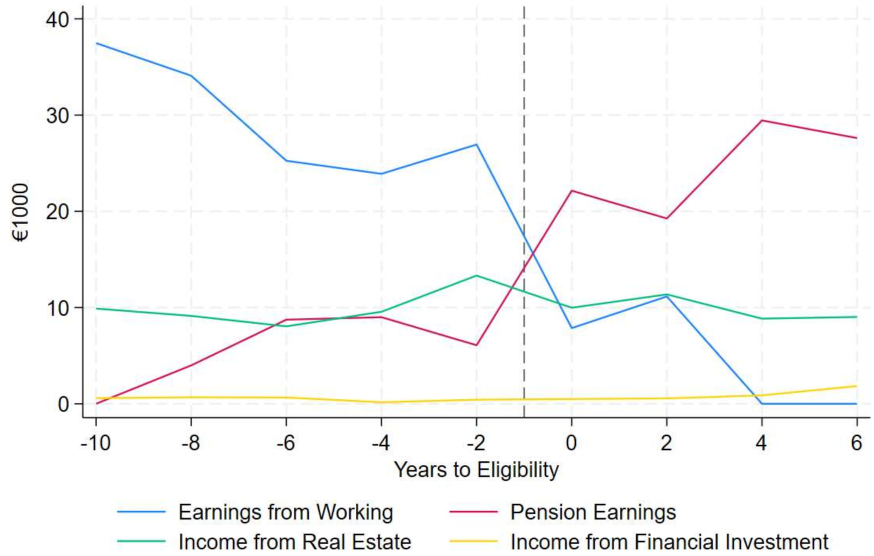

Notes. This figure plots the average income levels from different sources, i.e., earnings from work, retirement benefits, income from real estate, and income from financial investments, over the distance to retirement eligibility. The x-axis measures years to eligibility, while the y-axis shows average income levels.

Additional Controls

TABLE A1: Controlling for Unemployment Spells and other Incomes - Intensive margin

	Stocks-All	Stocks-Listed	Stocks + Managed
Retired	7.81^{***}	7.26^{***}	10.52^{***}
	(2.59)	(2.65)	(3.62)
Years of Unemployment	0.06	-0.04	-0.20
	(0.26)	(0.20)	(0.41)
Return from Financial Investments (1000 €)	-0.15	0.74	-0.26
	(0.30)	(0.72)	(0.35)
Real Estate Income (1000 €)	0.04	0.02	0.03
	(0.05)	(0.05)	(0.10)
Other Controls	✓	✓	✓
F-stat 1st stage	36.94	36.94	36.94
Observations	1086	1086	1086

Notes. This table reports the 2SLS estimates of the effect of retirement on the share of financial wealth invested in stocks, also controlling for the number of years spent in unemployment, income from financial investments, and real estate assets. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A2: Controlling for Unemployment Spells and other Incomes - Extensive margin

	Stocks-All	Stocks-Listed	Stocks + Managed
Retired	0.23**	0.22**	0.25**
	(0.11)	(0.10)	(0.13)
Years of Unemployment	-0.00	-0.00	-0.01
	(0.00)	(0.00)	(0.01)
Return from Financial Investments (1000 €)	-0.00	0.01	0.00
	(0.01)	(0.01)	(0.01)
Real Estate Income (1000 €)	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)
Other Controls	✓	✓	✓
F-stat 1st stage	41.20	41.20	41.19
Observations	1214	1214	1214

Notes. This table reports the 2SLS estimates of the effect of retirement on the share of financial wealth invested in stocks, controlling for the number of years spent in unemployment, income from financial investments, and real estate assets. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Excluding households with debts

TABLE A3: Excluding people with debt - Intensive Margin

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		8.38***	7.89***	7.96*
		(2.90)	(2.85)	(4.10)
Eligible for Retirement	0.43***			
	(0.09)			
Controls	✓	✓	✓	✓
F-stat 1st stage		23.67	23.67	23.67
Observations	611	611	611	611

Notes: This table presents results excluding people holding positive financial liabilities. The first column reports the first stage of a 2SLS regression where retirement status is instrumented with eligibility for retirement. The remaining columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A4: Excluding people with debt - Extensive Margin

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		0.28**	0.30***	0.19
		(0.12)	(0.12)	(0.14)
Eligible for Retirement	0.42***			
	(0.08)			
Controls	✓	✓	✓	✓
F-stat 1st stage		25.33	25.33	25.33
Observations	670	670	670	670

Notes: This table presents results excluding people holding positive financial liabilities. The first column reports the first stage of a 2SLS regression where retirement status is instrumented with eligibility for retirement. The remaining columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Excluding households with negative net wealth

TABLE A5: Excluding Households with negative net wealth - Intensive Margin

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		7.59***	7.25***	10.22***
		(2.60)	(2.70)	(3.62)
Eligible for Retirement	0.42***			
	(0.07)			
Controls	✓	✓	✓	✓
F-stat 1st stage		36.69	36.69	36.69
Observations	1046	1046	1046	1046

Notes: This table presents results excluding people with negative net wealth. The first column reports the first stage of a 2SLS regression where retirement status is instrumented with eligibility for retirement. The remaining columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A6: Excluding Households with negative net wealth - Extensive Margin

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		0.20*	0.23**	0.22*
		(0.11)	(0.10)	(0.12)
Eligible for Retirement	0.42***			
	(0.07)			
Controls	✓	✓	✓	✓
F-stat 1st stage		40.15	40.15	40.15
Observations	1046	1156	1156	1156

Notes: This table presents results excluding people with negative net wealth. The first column reports the first stage of a 2SLS regression where retirement status is instrumented with eligibility for retirement. The remaining columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Excluding household heads younger than 55

TABLE A7: Excluding household heads younger than 55 - Intensive Margin

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		11.05**	9.50**	13.05**
		(4.91)	(4.67)	(6.46)
Eligible for Retirement	0.30***			
	(0.09)			
Controls	✓	✓	✓	✓
F-stat 1st stage		9.85	9.85	9.85
Observations	448	448	448	448

Notes: This table presents results excluding people younger than 55. The first column reports the first stage of a 2SLS regression where retirement status is instrumented with eligibility for retirement. The remaining columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A8: Excluding household heads younger than 55 - Extensive Margin

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		0.44^{*}	0.41^{*}	0.36
		(0.23)	(0.22)	(0.24)
Eligible for Retirement	0.32^{***}			
	(0.09)			
Controls	✓	✓	✓	✓
F-stat 1st stage		10.68	10.68	10.68
Observations	483	483	483	483

Notes: This table presents results excluding people younger than 55. The first column reports the first stage of a 2SLS regression where retirement status is instrumented with eligibility for retirement. The remaining columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Excluding non eligible retired

TABLE A9: Excluding People in Retirement but Not Eligible - Intensive Margin

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		5.05^{***}	4.46^{***}	6.94^{***}
		(1.67)	(1.65)	(2.06)
Eligible for Retirement	0.80^{***}			
	(0.06)			
Controls	✓	✓	✓	✓
F-stat 1st stage		207.72	207.72	207.72
Observations	1028	1028	1028	1028

Notes: This table presents results excluding people in retirement but not eligible for any of the considered options. The first column reports the first stage of a 2SLS regression where retirement status is instrumented with eligibility for retirement. The remaining columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A10: Excluding People in Retirement but Not Eligible - Extensive Margin

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		0.14*	0.16**	0.17**
		(0.07)	(0.07)	(0.08)
Eligible for Retirement	0.81***			
	(0.05)			
Controls	✓	✓	✓	✓
F-stat 1st stage		232.34	232.34	232.34
Observations	1152	1152	1152	1152

Notes: This table presents results excluding people in retirement but not eligible for any of the considered options. The first column reports the first stage of a 2SLS regression where retirement status is instrumented with eligibility for retirement. The remaining columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

FIGURE A6: First Stage Excluding Not Eligible But Retired

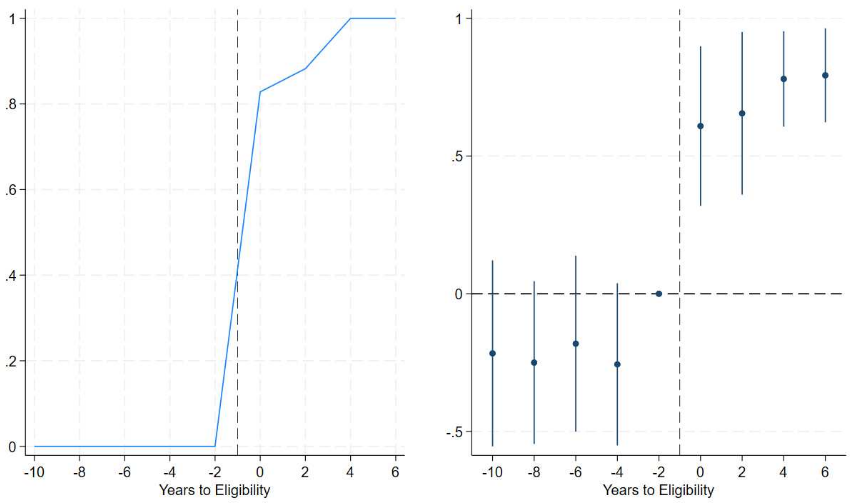

Notes: This Figure shows the probability of being retired as a function of distance to eligibility, dropping individuals who are not eligible for any of the schemes but are retired.

FIGURE A7: Annual Incomes Before Retirement Eligibility Excluding Not Eligible But Retired

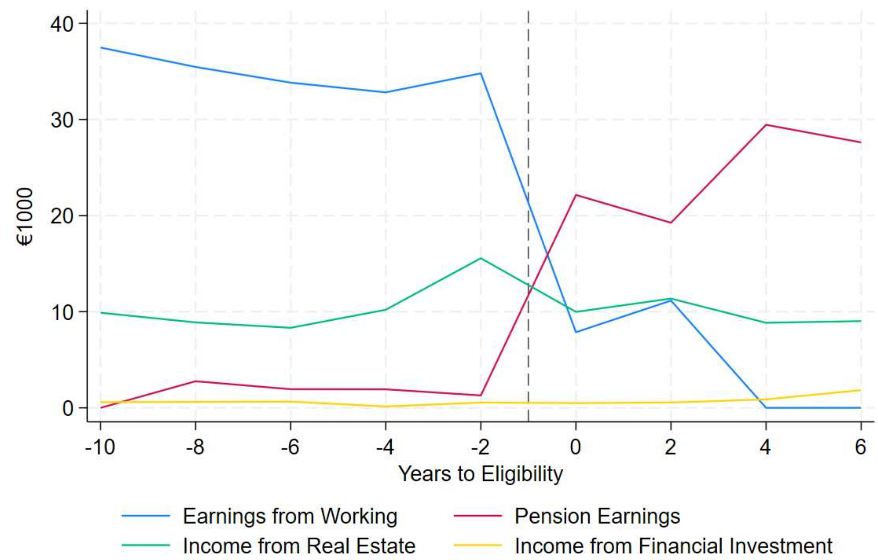

Notes. This figure displays the evolution of different income sources around the year of retirement eligibility, excluding individuals who are not eligible but already retired.

Alternative Eligibility Definition

TABLE A11: Alternative Specification of the Instrument. Intensive Margin Results

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		7.81***	7.26***	10.52***
		(2.59)	(2.65)	(3.62)
eligible	0.42***			
	(0.07)			
Controls	✓	✓	✓	✓
F-stat 1st stage		36.94	36.94	36.94
Observations	1086	1086	1086	1086

Notes: This table presents results using an alternative specification of the instrumental variable strategy. The first column reports the first stage of a 2SLS regression where retirement status is instrumented with eligibility for retirement. The remaining columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A12: Alternative Specification of the Instrument. Extensive Margin Results

	First Stage	Stocks-All	Stocks-Listed	Stocks + Managed
Retired		0.23**	0.22**	0.25**
		(0.11)	(0.10)	(0.13)
eligible	0.42***			
	(0.07)			
Controls	✓	✓	✓	✓
F-stat 1st stage		41.20	41.20	41.20
Observations	1214	1214	1214	1214

Notes: This table presents results using an alternative specification of the instrumental variable strategy. The two columns report the second-stage estimates of the effect of retirement on the share of financial wealth invested in stocks, separately for total stocks, publicly listed stocks, and privately held (unlisted) stocks. Standard errors are clustered at the individual level. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

FIGURE A8: First Stage with Alternative Specification of the Instrument

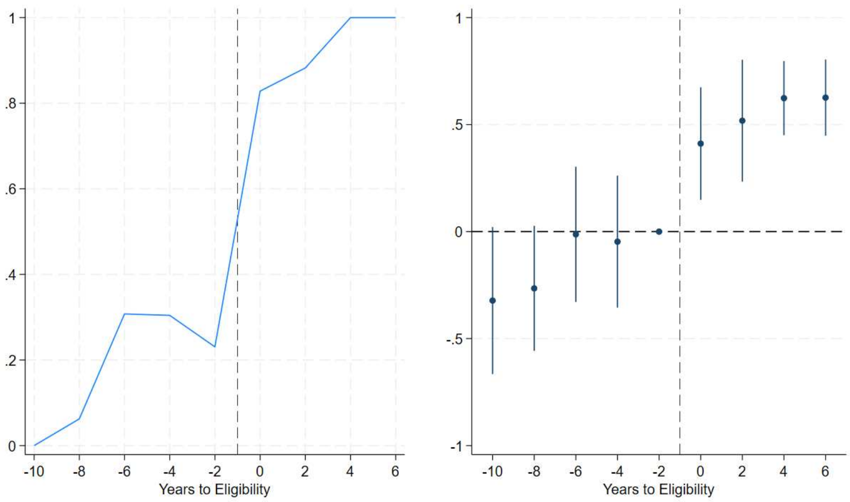

Notes: This Figure shows the probability of being retired as a function of distance to eligibility, allowing for contemporaneous eligibility.

2.7 Appendix B: Additional Descriptive Statistics

FIGURE B1: Years of Contribution in 2020 and 2022 by Region

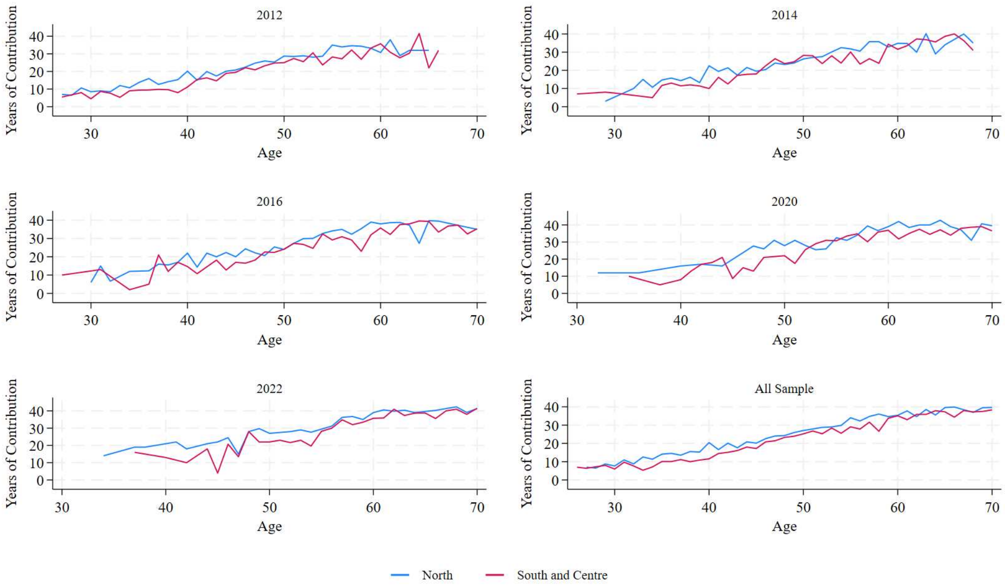

Notes. This figure shows the average years of contribution by age, separately for individuals in the North (blue line) and in the South/Centre (red line) of Italy. The left panel refers to 2020 and the right panel to 2022. The vertical dashed line marks the minimum age requirement for early retirement (62 in 2020, 64 in 2022), while the horizontal dashed line indicates the contribution requirement (38 years).

FIGURE B2: Mean Stock Holdings as a Share of Financial Assets by Retirement over Time

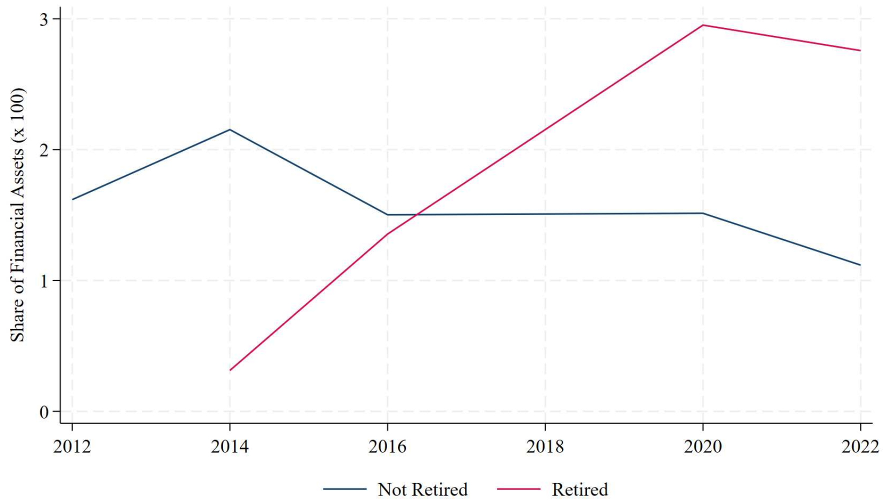

Notes. This figure displays the average share of stock holdings over total financial assets by retirement status between 2012 and 2022. The red line represents retired individuals, while the blue line shows those not yet retired. The year 2018 is missing.

FIGURE B3: Mean Stock Market Participation by Retirement over Time

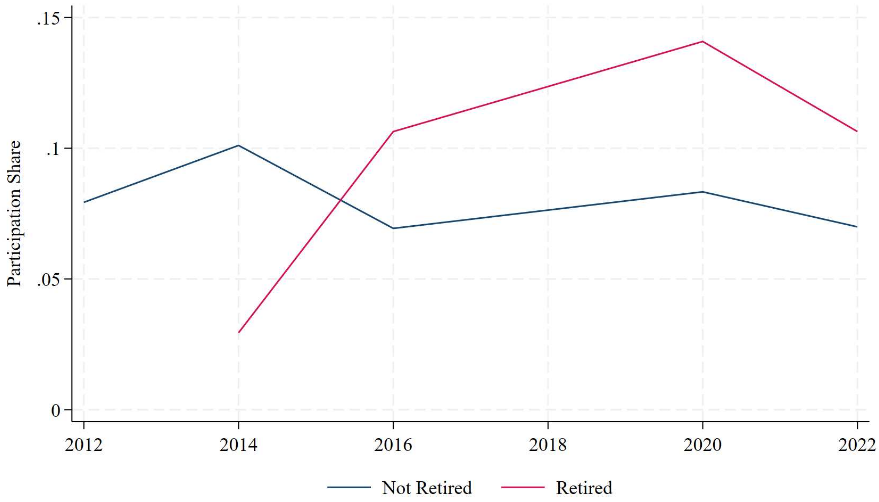

Notes. This figure shows the average stock market participation rate by retirement status between 2012 and 2022. The red line refers to retired individuals, while the blue line represents those not retired. Participation is defined as holding a positive amount of stocks. The year 2018 is missing.

FIGURE B4: Net Wealth and Real Estate Wealth

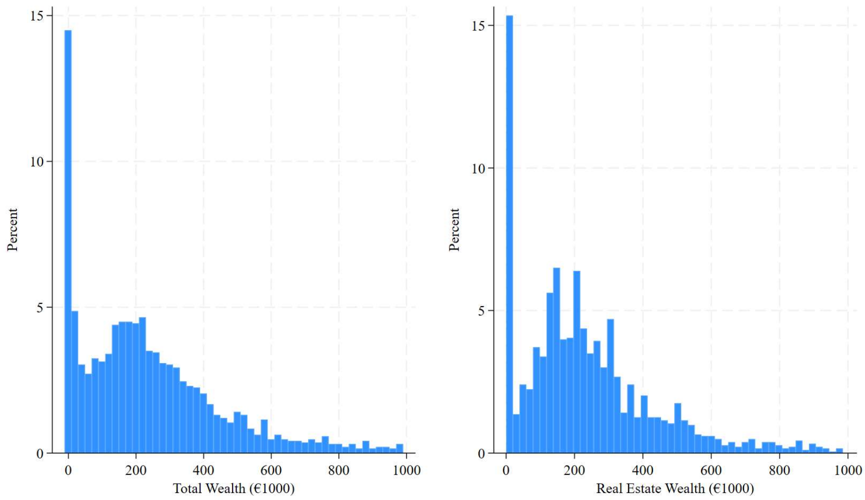

Notes: This Figure shows the pooled distribution of net wealth (left panel) and real estate wealth. Net wealth is the sum of financial wealth, real estate wealth and financial liabilities.

FIGURE B5: Financial Wealth

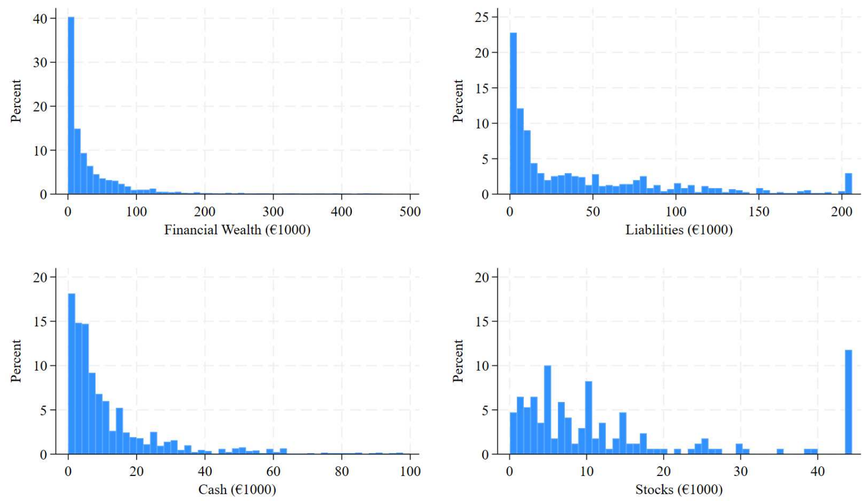

Notes: The upper left panel shows the pooled distribution of financial assets, the upper right panel the distribution of financial liabilities. The bottom left panel shows the distribution of cash, while the bottom right panel the distribution of stocks.

FIGURE B6: Pooled Distribution of Stock Holdings as a Share of Financial Assets by Retirement Status

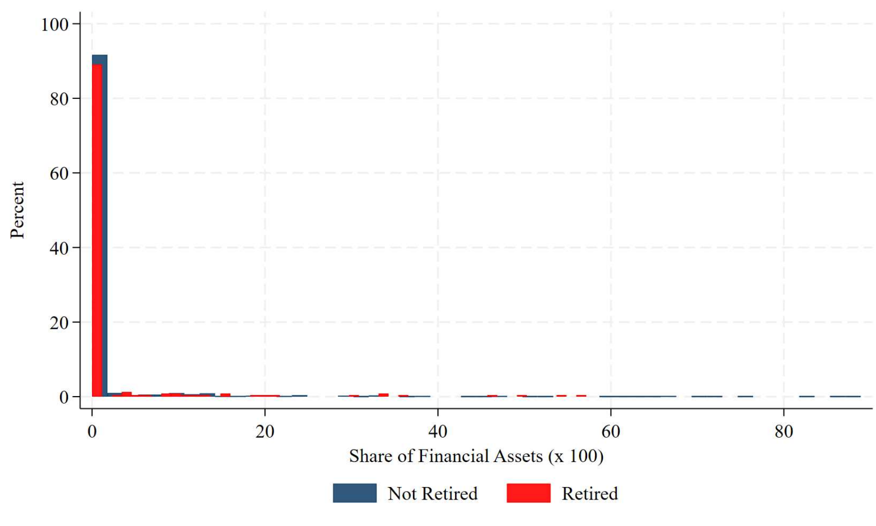

Notes: This figure shows the pooled cross-sectional distribution of stock holdings as a share of total financial assets by retirement status. The red bars represent retired individuals, while the blue bars represent those not retired. The distribution is heavily skewed to the left for both groups, with most individuals holding either no stocks or a very small share of stocks relative to their financial assets.

FIGURE B7: Life Cycle Profiles of Financial Assets and Stock Holding

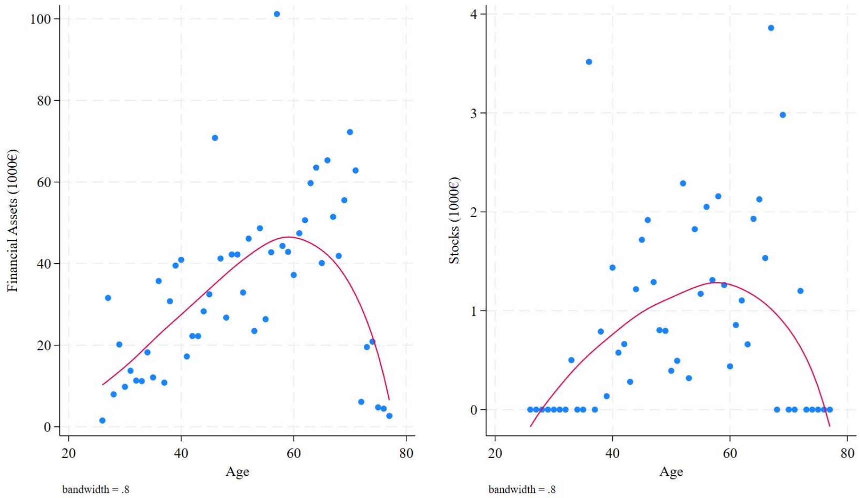

Notes: This Figure shows the li cycle profile of the value of financial assets (left panel) and the value of stocks (right panel).

FIGURE B8: Life Cycle Profiles of Stock Holding - Shares

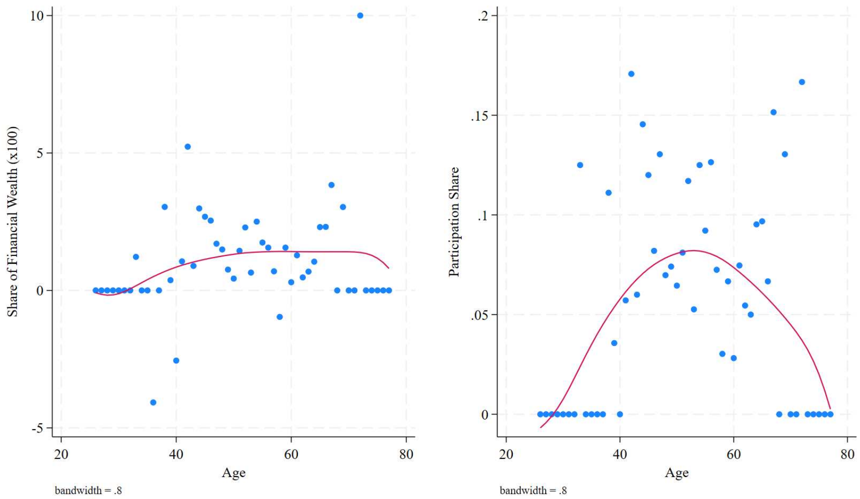

Notes: This Figure shows the life cycle profile of the share of stocks over financial assets (left panel) and the stock market participation (right panel).

2.8 Appendix C: RDD Results

We estimate the following equations, only considering observation distant -4 and 4 years from retirement eligibility:

$$Y_{it} = \beta retired_{it} + \gamma retired_{it} \times distance_{it} + \theta distance_{it} + \varphi \mathbf{X}_{it} + \epsilon_{it} \tag{C1}$$

$$\begin{aligned} Y_{it} &= \beta retired_{it} + \gamma_1 retired_{it} \times distance_{it} + \theta_1 distance_{it} \\ &+ \gamma_2 retired_{it} \times distance^2_{it} + \theta_2 distance^2_{it} + \varphi \mathbf{X}_{it} + \epsilon_{it} \end{aligned} \tag{C2}$$

TABLE C1: Linear RDD - Intensive Margin

	Stocks-All		Stocks-Listed		Stocks + Managed	
	(1)	(2)	(3)	(4)	(5)	(6)
Retired	5.04***	4.78***	4.91***	4.69***	5.87**	6.66***
	(1.57)	(1.55)	(1.57)	(1.55)	(2.87)	(2.46)
Years to Eligibility x Retired	1.34***	1.28***	1.33***	1.26***	1.47**	1.06
	(0.45)	(0.44)	(0.45)	(0.45)	(0.57)	(0.83)
Years to Eligibility	-0.18	-0.15	-0.17	-0.14	0.16	0.57
	(0.15)	(0.16)	(0.15)	(0.16)	(0.37)	(0.75)
Controls		✓		✓		✓
Observations	177	177	177	177	177	177

Notes: This Table shows the estimated coefficients of (C1), considering a window of the eligibility between -4 and 4. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE C2: RDD Results - Extensive Margin

	Stocks-All		Stocks-Listed		Stocks + Managed	
	(1)	(2)	(3)	(4)	(5)	(6)
Retired	0.13*	0.12*	0.12*	0.12*	0.16**	0.16**
	(0.07)	(0.07)	(0.07)	(0.06)	(0.07)	(0.06)
Years to Eligibility x Retired	0.04**	0.04**	0.04**	0.04*	0.04**	0.04*
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Years to Eligibility	-0.00	-0.00	-0.00	0.00	0.00	0.00
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Controls		✓		✓		✓
Observations	182	182	182	182	182	182

Notes: This Table shows the estimated coefficients of (C1), considering a window of the eligibility between -4 and 4. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE C3: Quadratic RDD - Intensive Margin

	Stocks-All		Stocks-Listed		Stocks + Managed	
	(1)	(2)	(3)	(4)	(5)	(6)
Retired	4.25***	3.85**	4.09**	3.74**	5.38**	6.58**
	(1.59)	(1.51)	(1.59)	(1.51)	(2.57)	(2.55)
Years to Eligibility x Retired	1.81***	1.71***	1.82***	1.72***	2.43***	2.47***
	(0.56)	(0.60)	(0.56)	(0.60)	(0.51)	(0.71)
Years to Eligibility Squared x Retired	0.14*	0.15*	0.15*	0.15*	0.21***	0.27***
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
Years to Eligibility	-0.22	-0.12	-0.22	-0.11	-0.30	-0.29
	(0.19)	(0.24)	(0.19)	(0.24)	(0.27)	(0.66)
Years to Eligibility Squared	-0.01	0.00	-0.01	0.00	-0.05	-0.09
	(0.01)	(0.03)	(0.01)	(0.03)	(0.05)	(0.07)
Controls		✓		✓		✓
Observations	177	177	177	177	177	177

Notes: This Table shows the estimated coefficients of (C2), considering a window of the eligibility between -4 and 4. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE C4: Quadratic RDD - Extensive Margin

	Stocks-All		Stocks-Listed		Stocks + Managed	
	(1)	(2)	(3)	(4)	(5)	(6)
Retired	0.07	0.06	0.06	0.06	0.10	0.11
	(0.07)	(0.06)	(0.07)	(0.06)	(0.07)	(0.07)
Years to Eligibility x Retired	0.08***	0.07***	0.07***	0.07***	0.08***	0.08***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Years to Eligibility Squared x Retired	0.01***	0.01***	0.01***	0.01***	0.01***	0.01***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Years to Eligibility	-0.00	-0.00	-0.00	0.00	-0.01	-0.01
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Years to Eligibility Squared	-0.00	0.00	-0.00	0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Controls		✓		✓		✓
Observations	182	182	182	182	182	182

Notes: This Table shows the estimated coefficients of (C2), considering a window of the eligibility between -4 and 4. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

FIGURE C1: McCary Test around 38 and 42 years of contribution

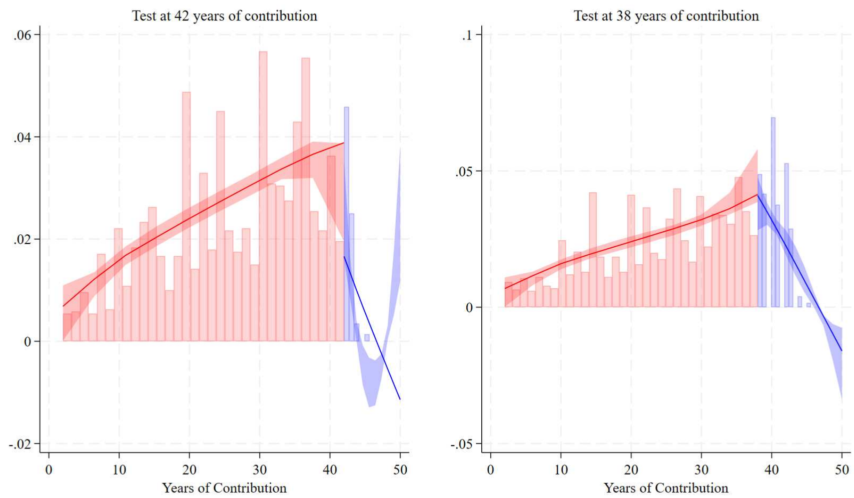

Notes: This Figure shows the outcome of a McCrary test performed on the distribution of the years of contribution. The left panel shows the test around 42, the threshold for seniority eligibility. The righ panel shows the same test for 38, the threshold of Q100.

FIGURE C2: 2SLS Results using observations in different eligibility bandwidths

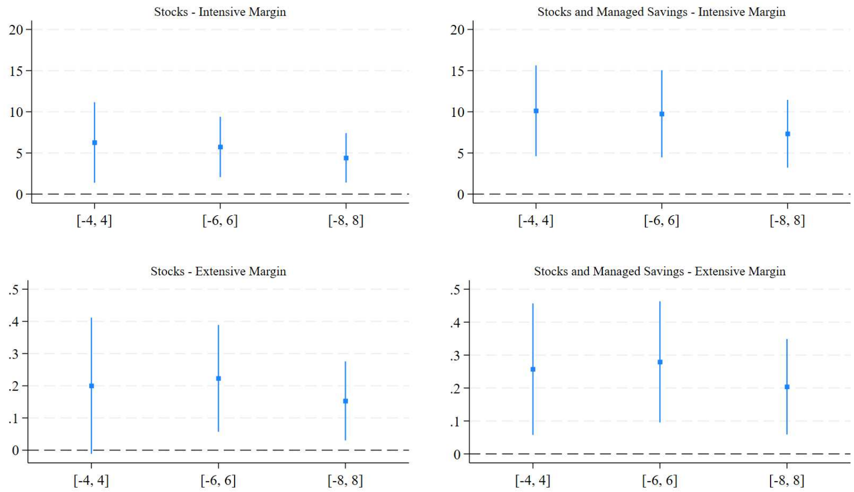

Notes: This Figure shows the 2SLS results using only individuals whose distance to retirement eligibility is within the range indicated on the horizontal axis. The first column reports results for the stocks (intensive and extensive margin). The second column reports the results for stocks and managed savings.

2.9 Appendix D: (Quasi) Retirement Consumption Puzzle

Battistin et al., 2009, using SHIW data from 1989 to 2004, shows that households reduce consumption when the head retires. A similar pattern in found more recently by Kolsrud et al., 2024 and Olafsson and Pagel, 2024 in Sweden and Denmark, respectively. This result is at odds with standard macro theory, since consumers are supposed to smooth consumption and not be affected in their decisions by expected income changes.

In this Appendix, we show two interesting results. First, in line with Battistin et al., 2009, we document a drop in nondurable consumption upon retirement. We perform the same TWFE 2SLS estimation as in the main analysis. However, and this is the second result, this is not properly a *retirement consumption puzzle* because the drop in consumption concerns only the old age eligible, while there is no significant effect for seniority and early eligibility. The mechanism behind this result is analogous to the one proposed in the main analysis on stock holdings: old age retirees are more likely to be liquidity constrained and thus to be sensitive to expected income changes (Jappelli and Pistaferri 2017, Deaton 1991). In other words, seniority retirees and early retirees behave, in this sense, as PIH consumers do not change their consumption upon retirement.

TABLE D1: 2SLS Estimates for Nondurable Consumption

	Nondurable Consumption (1000 €)			Log of Consumption ($\times$ 100)		
	(1)	(2)	(3)	(4)	(5)	(6)
Retired	-8.38***	-6.58*	-8.38***	-28.89**	-28.89**	-28.67**
	(3.01)	(3.95)	(3.01)	(12.53)	(12.53)	(12.41)
Net Wealth (1000 €)						0.04***
						(0.01)
Observations	1214	1214	1214	1214	1214	1214
F-1st Stage	79.88	40.92	79.88	40.92	40.92	40.88
Controls		✓	✓		✓	✓

Notes: This Table shows the 2SLS estimates of equation (3). From columns 1 to 3, the dependent variable is the level of nondurable consumption (in 1000 €) of household i at time t over financial assets (intensive margin). From columns 4 to 6, the dependent variable is the log multiplied by 100. *Retired* equals 1 if i is retired 0 otherwise, and is instrumented by *Eligible*, defined as in equation (2). Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE D2: Reduced Form for Consumption

	Nondurable Consumption (1000 €)			Log of Consumption (× 100)		
	(1)	(2)	(3)	(4)	(5)	(6)
Seniority Eligible	-5.39*	-3.16	-3.79	-17.55**	-10.86	-12.57
	(2.81)	(2.74)	(2.75)	(8.18)	(8.08)	(8.09)
Q100 Eligible	-3.04	-1.14	-1.07	-12.46	-7.25	-7.07
	(2.89)	(3.01)	(2.95)	(10.11)	(10.25)	(10.02)
Old Age Eligible	-4.47**	-3.02	-2.62	-19.40***	-14.37**	-13.27**
	(1.99)	(1.98)	(1.94)	(6.42)	(6.36)	(6.23)
Net Wealth (1000 €)			0.01***			0.04***
			(0.00)			(0.01)
Observations	1214	1214	1214	1214	1214	1214
Controls		✓	✓		✓	✓

Notes: This Table shows OLS estimates of equation (6). In Panel A, either the level of nondurable consumption or the log. In Panel B, the dependent variable is a dummy equal to 1 if i holds stocks at time t. $Senioriy_{it}$, $Early_{it}$, $Old\ Age_{it}$ are dummies defined as in (1). Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.10 Appendix E: SHIW Data and Dataset Construction

The Survey of Household Income and Wealth (SHIW) started in the 1960s and collects information on saving and investment behavior of Italian Households. In the last years, SHIW asks retirement-specific questions, like the years of contribution, the expected retirement age and the expected pension benefit, as a share of the last working earnings. In particular, we use the first piece of information, together with age, to define eligibility to the three retirement options explained in the main text.

We use the waves of 2012, 2014, 2016, 2020 and 2022, and keep track of the same household head. This sample selection is follows the same criterion as Jappelli and Pistaferri, 2000; Jappelli and Pistaferri, 2020; Jappelli and Pistaferri, 2025. Since there is not a individual identifier, but just a family identifier, we focus on the household head and make sure that she does not change from one wave to another. We therefore drop households whose head changed sex, report inconsistent ages and years of contribution over time. For instance, if i is 30 in 2012 and has worked for 5 years, she must be 32 and have worked for maximum 7 years in 2014, 34 in 2014 and so on. This procedure yields a sample or around 2000 household-year observations.

SHIW has two types of datasets: historical and historical. The historical dataset is just a unique dataset with all waves and containing variables in aggregated form at the household level. The yearly datasets contain more detailed information as years of contribution in-depth demographics, wealth and employment status. We use both the historical and the yearly datasets, merging them for every year using the family unique identifier *nquest*. The table below describes the variables we use, and how we derive them.

TABLE E1: List of Used Variables

Variable	Units	Description and Source Variables
nquest	number	cross-sectional identifier
year	year	time
age	number of years	
yc	number of years	years of contribution: acontrib
disann	number of years	year of unemployment
dislav	dummy	experienced long term unemployment
female	dummy	gender: sesso == 2
low_ed	dummy	low education: studio $<$ 3
mid_ed	dummy	middle education (high school): studio == 4 or 5
high_ed	dummy	high education (college decree or more): studio $>$ 5
south	dummy	living in the South: area5 == 4 or 5
centre	dummy	living in the Centre: area5 == 3
north	dummy	living in the North: area5 == 1 or 2
family_size	number	household components: ncomp

Continued on next page

Variable	**Units**	**Description and Source Variables**
retired	dummy	retired: qual > 7
cn	1000 €	nondurable consumption
cd	1000 €	durable consumption
c	1000 €	total consumption: cn + cd
lcn	log ×100	log of nondurable consumption: $\log(cn)$
lc	log ×100	log of total consumption: $\log(c)$
ytp	1000 €	family pension earnings
yca	1000 €	real estate income
ycf	1000 €	financial income
ar	1000 €	real estate wealth
af	1000 €	total financial wealth
pf	1000 €	total liabilities
w	1000 €	net wealth: ar + af - pf
bank_account	1000 €	amount of cash: ld
public_debt	1000 €	government securities, af2
bonds_private	1000 €	private bonds (companies and banks): lobb
assvita	1000 €	amount paid annually for life insurance
pensint	1000 €	amount paid annually for private pensions
stocks_listed	1000 €	stocks traded on the stock market (market value at year-end): lazq
stocks_unlisted	1000 €	stocks of private businesses: lazi + lsrl + lper - stocks listed until 2016; lazi2 - stocks listed from 2020
stocks	1000 €	total stocks: *stocks_listed* + *stocks_unlisted*
managed_savings	1000 €	managed savings: lgp
others	1000 €	other assets: lat
foreign	1000 €	foreign assets: lte
share_bank_account	%	share of cash: bank account / af
share_public_debt	%	share of government securities, public debt / af
share_bonds_private	%	share of private bonds: bonds private / af
share_stocks_listed	%	share of listed stocks: stocks listed / af
share_stocks_unlisted	%	share of stocks of private businesses: stocks unlisted / af
share_stocks	%	share of total stocks: stocks / af
share_managed	%	share of managed savings: managed / af
share_others	%	share of other assets: others/ af
share_foreign	%	share of foreign assets: foreign / af
p_bank_account	dummy	hods cash: *bank_account* > 0
p_public_debt	dummy	holds government securities: *public_debt* > 0
p_bonds_private	dummy	holds private bonds: *bonds_private* > 0
p_stocks_listed	dummy	holds listed stocks: *stocks_listed* > 0
p_share_stocks_unlisted	dummy	holds stocks of private businesses: *stocks_unlisted* > 0
p_share_stocks	dummy	holds total stocks: *stocks* > 0
p_share_others	dummy	holds other assets: *others* > 0

Continued on next page

Variable	Units	Description and Source Variables
$p\_share\_foreign$	dummy	holds foreign assets: $foreign > 0$

Chapter 3

Making Room for a Baby: The Impact of Working from Home on Fertility Decisions

Nicol Barbieri *

This paper studies how the expansion of Working from Home (WFH) during the pandemic affected women's fertility decisions in Italy. Using data from the Italian Labor Force Survey (2015-2022), I exploit variation in occupations' "teleworkability", i.e, the extent to which each occupation can be performed from home based on pre-pandemic task-based characteristics, to study maternity leave take-up as a proxy for childbirth. I implement a Difference-in-Differences (DiD) design and a Two-Stage Least Squares strategy (2SLS), instrumenting actual WFH with pre-pandemic occupation-level teleworkability interacted with the post-2020 period. My results show that women in teleworkable occupations became significantly more likely to take maternity leave after the pandemic. The effects are concentrated among highly educated, high-income, and full-time women, those with larger families, as well as in service-sector, white-collar jobs, highlighting that remote work reduces the opportunity cost of childbearing where career-family trade-offs are most binding.

* I thank my advisor, Monica Langella, for her guidance throughout this project. Many thanks to Annalisa Scognamiglio for her insightful feedback. This work has been funded under grant Prin PNRR, P2022KHP8L, "Inequalities and the labor markets".

3.1 Introduction

The COVID-19 pandemic triggered an expansion of Working from Home (WFH) practices worldwide, reshaping the organization of labor and family life. While the economic and productivity implications of WFH have been extensively debated (see, e.g., Choudhury et al., 2024; Bloom et al., 2024; Gibbs et al., 2023; Bloom et al., 2015), less attention has been devoted to its consequences for demographic behaviors, in particular, fertility decisions. Understanding this relationship is crucial, as declining fertility rates are a pressing concern for many advanced economies, where the ability to reconcile work and family responsibilities remains a key determinant of childbearing choices (see, e.g., Doepke et al., 2023; Goldin, 2014; Adsera, 2004; Becker, 1992).

In this context, WFH may represent a fundamental shift in the cost-benefit analysis of fertility, as it expands workers' ability to combine employment and family responsibilities. Flexible working arrangements have long been recognized as a valued job amenity, especially for women with children, and are often capitalized into wages (see, e.g., Mas, 2025; Mas and Pallais, 2017). Recent evidence suggests that access to job flexibility increases labor force attachment among mothers and supports fertility intentions by reducing the opportunity costs of childbearing (see, e.g., Hupkau and Petrongolo, 2020; Farré and González, 2019). Therefore, by lowering commuting time and allowing a reallocation of time within the household, WFH can further reduce the tension between work and caregiving, potentially making childbearing more compatible with sustained labor market participation (see, for instance, Goldin, 2021; Del Boca and Sauer, 2009).

In this sense, WFH could be viewed as a novel form of workplace flexibility, one that directly alters the time and opportunity costs of fertility decisions. Moreover, the implications of WFH for fertility may differ substantially between parents and non-parents. For women who are already mothers, WFH can ease time constraints, reduce childcare costs, and facilitate the decision to have additional children. For non-parents, the effects may operate through different channels. On the one hand, increased flexibility and reduced commuting time may encourage earlier family formation. Yet, the blurred boundaries between work and private life, as well as potentially reduced social interactions, could delay entry into parenthood. These contrasting mechanisms suggest that the relationship between WFH and fertility decisions is likely heterogeneous and context-specific.

In this paper, I study how and to what extent the adoption of WFH, driven by the COVID-19 pandemic, has affected women's fertility decisions in Italy. The Italian context provides an ideal setting for this type of analysis: fertility rates have been persistently low, labor market rigidities are pronounced, and family policies remain relatively limited as opposed to other European countries (see, for studies on this topic, Carta and De Philippis, 2024; Zurla, 2023; Del Boca, 2002; Del Boca et al., 2004; Billari and Kohler, 2004). These structural features make fertility choices in Italy particularly sensitive to changes in work arrangements.

The expansion of WFH following the pandemic created a quasi-natural experiment to study whether the flexibility in the workplace translated into higher take-up of maternity leave and, ultimately, into changes in fertility decisions. To identify the effect of WFH on fertility behavior, I exploit heterogeneity in occupations' WFH feasibility, following the teleworkability index proposed by Basso et al., 2022. I implement a Difference-in-Differences (DiD) and a Two-Stage Least Squares (2SLS) strategy, comparing women in "teleworkable" versus "non-teleworkable" occupations, i.e, jobs whose task content makes them feasible or not for WFH according to a pre-pandemic task-based classification, before and after the pandemic, with a focus on the probability of being on maternity leave as a proxy for fertility decisions.

My results show that WFH significantly increases women's likelihood of taking maternity leave, suggesting that WFH can ease the work-family trade-off and positively affect fertility decisions.

To rationalize this result, I conduct some heterogeneous analysis. Two main pathways arise. First, the effect is primarily concentrated among highly educated and high-income women, who are more likely to be employed in teleworkable occupations and face higher opportunity costs associated with childbearing. This suggests that the additional flexibility introduced by WFH can help relax career-family constraints precisely where these trade-offs are most binding. Highly educated and high-income women tend to have stronger career attachment, steeper wage profiles, and greater advancement opportunities, which amplify the opportunity costs of childbearing. For them, interruptions related to motherhood imply larger foregone earnings and higher risks of career setbacks, making the decision to have children particularly sensitive to changes in workplace flexibility. Second, the effect is more substantial for women with larger families and for full-time workers, pointing to the role of WFH in reducing time and resource pressures for women with greater household responsibilities or stronger labor market attachment. The results are also more pronounced in the service sector and among white-collar occupations, where WFH is more formally institutionalized.

Taken together, my results suggest that the expansion of WFH acted as a novel source of flexibility in the labor market, lowering the opportunity costs of motherhood for certain groups of women. Importantly, however, the benefits appear unevenly distributed.

My analysis builds on a growing literature investigating the impact of labor market institutions and work arrangements on family outcomes. Previous studies have documented how childcare availability, parental leave policies, and job security influence fertility decisions (see, for example, Lalive and Zweimüller, 2009; Del Boca and Sauer, 2009; Goldstein et al., 2003). More recently, research has started exploring the interplay between WFH and gender roles (see, for instance, Hensvik et al., 2021; Hupkau and Petrongolo, 2020; Alon et al., 2020), but evidence on fertility remains scarce. I contribute to this debate by providing one of the first causal estimates of the effect of WFH on women's fertility.

I also speak to the literature on heterogeneity in fertility responses across socioeconomic groups.

For instance, Adsera, 2004 shows that unstable employment conditions and rigid labor markets disproportionately discourage fertility among women with lower education and weaker labor market attachment. Building on this work, I show that WFH reshapes these dynamics: while its fertility-enhancing effects are concentrated among highly educated and high-income women, the mechanism differs from labor market insecurity, as it operates through increased flexibility and reduced opportunity costs of childbearing. In this way, I extend the heterogeneity perspective by showing how institutional changes in work arrangements can alter traditional socioeconomic slopes in fertility behavior.

My findings also empirically speak to classic theories of fertility and the opportunity cost of childbearing (see, on this topic, Becker, 1992; Hotz et al., 1997). The papers emphasize how higher female wages and stronger labor market attachment raise the opportunity cost of children, leading to delayed or reduced fertility. Consistent with this view, I provide recent empirical evidence showing that the flexibility introduced by WFH mitigates this trade-off, particularly for highly educated and high-income women, who typically face the steepest career-family conflicts.

Finally, by exploiting the sudden expansion of WFH induced by COVID-19, I contribute to the emerging literature that uses the pandemic as a natural experiment to estimate causal effects of novel shocks on social and demographic outcomes. This includes, for example, studies quantifying lockdown impacts on fertility intentions (see, for instance, Mooi-Reci et al. 2023) and analyses investigating how teleworking reshaped work-family dynamics in sharing caregiving responsibilities (i.e., Angelici and Profeta, 2024; Angelici and Savio, 2024; Del Boca et al., 2022; Del Boca et al., 2021). My focus complements these works by uncovering the fertility-related consequences of remote work adoption - specifically, maternity leave behavior - in a causal framework.

The remainder of the paper is structured as follows. Section 3.2 provides background on fertility policies and WFH adoption in Italy. Section 3.3 presents the data and shows some descriptive statistics, while the empirical strategy and the main results are reported in Section 3.4. In Section 3.5 I discuss potential mechanisms, while in Section 3.6 I present robustness checks and placebo test. Finally, Section 3.7 concludes.

3.2 Institutional Context

In Italy, social insurance for working mothers is centered on maternity and parental leave provisions, which constitute the core of the country's family policy framework. Compulsory Maternity Leave (ML) covers a period of five months, usually beginning two months before the expected date of birth and continuing until three months after. During this period, mothers are legally prohibited from working and receive a government transfer equal to 80 percent of their previous earnings (see, on this topic, INPS, 2025). After compulsory leave, parents are entitled to up to ten months of Parental Leave (PL), which can be taken until the child turns eight.

The first six months are paid at a 30 percent replacement rate if used before the child's sixth birthday, while the remaining months are unpaid[2]. In practice, PL is mostly taken by mothers, who account for more than 90 percent of recipients, and it is concentrated in the first year after childbirth (see, e.g., OECD, 2025). In addition, Italian law grants mothers job protection from 300 days before the expected birth until the child's first birthday, ensuring reinstatement under the same conditions as before birth (see, for instance, Carta and De Philippis, 2024).

Despite this relatively generous system of statutory leave, Italy's family policy framework remains weak in its capacity to support female employment and fertility, particularly when compared to other European countries (see, e.g., Del Boca et al., 2004; Billari and Kohler, 2004). Female labor force participation rates are among the lowest in the EU, and access to affordable childcare services is limited (see, for instance, Del Boca and Vuri, 2007; Brilli et al., 2017). Consequently, women continue to bear the majority of childcare responsibilities, and gender norms regarding household work division remain highly traditional (see, e.g., Barbieri et al., 2015; Arpino et al., 2014). These structural features make the reconciliation between work and family particularly challenging, especially for women in full-time or career-oriented jobs, who face higher opportunity costs of childbearing (see, e.g., Doepke et al., 2023; Adsera, 2004).

These institutional and cultural factors make fertility decisions in Italy particularly sensitive to changes in work organization and time flexibility. In this context, the diffusion of remote work represents a potentially important margin of adjustment: by altering how and where work is performed, it can directly affect the compatibility between employment and family life.

Against this backdrop, WFH has historically been rare in Italy. Before 2020, only a limited set of firms-mostly large and high-productivity enterprises-had adopted formal "smart-working" arrangements, and take-up was concentrated among high-skilled employees in the service sector (see Boeri et al., 2020; Angelici and Profeta, 2024). Estimates from Eurostat, 2024 indicate that only around 5 percent of employees reported working from home regularly before the pandemic, one of the lowest shares in Europe.

The COVID-19 pandemic radically changed this scenario. Starting in March 2020, the government introduced a series of emergency decrees that made remote work the default arrangement whenever tasks could be performed from home (Law Decree 18/2020, "Cura Italia"). Employers were allowed to implement WFH without prior individual agreements, and coverage was extended across both public and private sectors. This exceptional framework remained in place through 2021 and was progressively phased out thereafter, with many firms maintaining hybrid or partial WFH arrangements (see, for instance, Basso et al., 2023). As a result, Italy witnessed a dramatic and sudden expansion of remote work: the share of employees

[2] The legislative decree Law 151/2001, "*Testo unico delle disposizioni legislative in materia di tutela e sostegno della maternità e della paternità*" reorganized and harmonized all previous provisions on maternity and paternity protection, consolidating a fragmented legal framework into a single national code. It remains the cornerstone of Italian family and labor law, periodically updated through subsequent budget laws and decrees to extend eligibility and adapt replacement rates.

working from home rose from 5 percent before the pandemic to over 15 percent in 2020-2021 (see Eurostat, 2024), with particularly high adoption in the service sector and among white-collar professionals.

These institutional changes - coupled with a longstanding structure of gendered labor markets - shape both women's labor supply and fertility-related choices. In particular, they define the extent to which remote work can complement existing family policies. By lowering commuting costs, reducing coordination frictions, and providing greater flexibility in time allocation, WFH has the potential to mitigate the opportunity costs associated with childbearing, but only for women with stable employment and formal access to statutory leave entitlements. Conversely, for those in precarious or informal jobs, remote work opportunities remained largely unavailable, perpetuating inequality in work–family reconciliation. This makes the Italian setting particularly well-suited to study how workplace flexibility interacts with fertility decisions, especially in a context where fertility rates are persistently low, labor market dualism is strong, and policy support for working parents remains limited.

3.3 Data and Teleworkability Measure

The Italian Labor Force Survey (LFS) is a dataset collected by ISTAT, the Italian National Institute of Statistics, that provides quarterly information on labor market status and other socioeconomic characteristics of a representative sample of the Italian population (about 95,000 observations per quarter). I use data from the first quarter of 2015 to the third quarter of 2022, totaling 1,504,690 observations. I restrict the sample to exclude self-employed workers. This also allows me to avoid self-selection into WFH adoption: self-employed workers often have the flexibility to create their work environments, making them more inclined to adopt remote work practices[3]. Due to the features of my dataset, which provides information on maternity leaves, I focus exclusively on women aged 16-46 years. For each woman, I can observe whether she has open-ended or part-time contracts, whether she works in the private or public sectors, and in which occupation[4] and sector[5] she is employed. Importantly, I can observe if she has taken maternity leave, which is crucial information to have a proxy for fertility decisions.

To build my dependent variable, I use two LFS questions about the reasons for respondents having no working time or reduced working time during the reference week. The two questions were formulated as follows: “What is the main reason why you did not work last week?", and “What is the main reason why you worked less than usual?". For both questions, one possible answer is “Compulsory Maternity Leave", which in Italy typically covers the 2 months before and 3 months after childbirth. Using this information, I build my dependent variable

[3] Self-employed workers are also not entitled to compulsory maternity leave in Italy. Since my outcome variable is based on maternity leave take-up, they could never appear as treated cases in the data. Restricting to dependent employees ensures that the measure of maternity leave reflects a comparable institutional entitlement.

[4] The LFS contains detailed information about occupations, up to the 4-digit ISCO-88 classification.

[5] The classification refers to the ATECO 2008, i.e., the Italian classification aligning the ISIC one.

Maternity Leave equal to 1 for women who declare "Compulsory maternity leave"[6]. Because each respondent can report only one main reason for non-work in the reference week, this variable is mutually exclusive with all other leave categories and therefore provides a proxy for the moment of childbirth[7].

The survey includes questions that measure the actual use of WFH. The question asks how often employees worked from home in the last four weeks before the interview: "In the last four weeks, did you work at home?". Employees can then specify the frequency. This can either be two or more times per week, less than twice per week, or not at all. Using this information, I build the *Actual WFH* variable equal to 1 for workers who declare to work from home twice or more than twice and less than twice per week (and 0 otherwise).

It is important to acknowledge a potential contemporaneity issue between the two measures. *Maternity Leave* is defined with respect to the reference week, while *Actual WFH* refers to remote work during the four weeks preceding the interview, which include the reference week. This overlap implies that a woman on maternity leave for the entire four-week period is mechanically recorded as not working from home. In other words, the two variables are not entirely independent by construction. However, a woman on maternity leave in the reference week may still report remote work if she was working earlier in the four-week period. This mechanical feature can only bias the estimates downward, since it reduces the probability of observing WFH among women on leave, rather than generating spurious positive correlations[8].

In Table 3.1, I report the summary statistics of the main variables involved in the analysis.

Approximately 28% of women have reduced working hours due to maternity leave. Only 5% of women actually worked from home during the reference week, reflecting both institutional constraints and the limited adoption of WFH practices in the Italian labor market during the observed period.

[6] From 2021 onward, the LFS introduced the same answer category for male respondents labeled "Compulsory Paternity Leave". The 10-day compulsory paternity leave was first introduced on an experimental basis in 2013 (Law 92/2012) and became permanent with Budget Law 234/2021. Nonetheless, the number of male respondents reporting paternity leave is very limited, preventing any significant analysis of fathers' fertility-related behavior.

[7] The *Maternity Leave* indicator used in the analysis captures only women who are formally employed and explicitly report maternity leave as the main reason for absence from work in the reference week. This implies that the analysis focuses on the intensive margin of labor supply among employed women and does not capture childbirth events among the unemployed. However, this approach is consistent with previous studies using labor force or administrative data to proxy fertility events through maternity leave take-up among employed women, such as Hiriscau, 2024; De Paola et al., 2021; Lalive and Zweimüller, 2009, among others.

[8] As a robustness check, I exclude observations in which women are simultaneously on maternity leave and working from home, to rule out potential mechanical overlap between the two conditions. Results, reported in the Appendix B, remain unchanged.

TABLE 3.1: Summary Statistics of Labor Force Survey (LFS)

	Mean	**Sd**
Young	0.73	0.45
Medium Education	0.19	0.40
High Education	0.80	0.40
Italian	0.91	0.29
North	0.62	0.49
South	0.56	1.17
Blue Collars	0.30	0.46
Permanent Contracts	0.85	0.36
Full-time Contracts	0.66	0.47
Big Firms	0.38	0.49
Agriculture	0.01	0.11
Industry	0.15	0.36
Construction	0.01	0.10
Services	0.69	0.46
Actual WFH	0.03	0.18
Maternity Leave	0.28	0.45
Observations	27,048	

Notes. This table shows the Mean and the Standard deviation (Sd) of the main variables used in the analysis.

Table 3.2 shows summary statistics separately for the pre- and post-pandemic periods, only for the two main variables of interest, i.e., maternity leave and actual WFH.

TABLE 3.2: Summary Statistics by Pre/Post Pandemic

	Mean	**Sd**	**Obs.**
Panel A: Pre-pandemic			
Maternity Leave	0.31	0.46	17008
Actual WFH	0.01	0.12	17008
Panel B: Post-pandemic			
Maternity Leave	0.24	0.43	9926
Actual WFH	0.07	0.25	9926

Notes. This table reports means, standard deviations, and number of observations for the key variables of interest, separately for the pre-pandemic period (Panel A) and the post-pandemic period (Panel B).

Before the pandemic (Panel A), only 2% of women reported WFH. The share of women on maternity leave was 31%. By contrast, Panel B shows that after the pandemic, the share of women working from home increased to 12%, while the incidence of maternity slightly declined to 24%.

I complement these data with an occupation-based indicator that captures whether the respondent's job can, in principle, be performed remotely. This distinction is very important: while the Actual WFH varies over time, this indicator is time-invariant and based on pre-determined job-task characteristics, as it is constructed solely from the task content of occupations and thus fixed before the pandemic and independent of individual fertility decisions.

I define the “teleworkability" of each occupation, i.e., the extent to which an occupation can be done from home, relying on the task-based classification system proposed by Basso et al., 2022. This approach constructs WFH measures by leveraging detailed information from the O*NET database, which documents the nature of tasks and required skills for occupations in the U.S. labor market[9] (see, for a similar approach, Dingel, 2020; Boeri et al., 2020; Mongey et al., 2021; Gottlieb et al., 2021; Alipour et al., 2023). Specifically, this classification evaluates the extent to which tasks associated with each occupation can be performed remotely, considering factors such as the need for face-to-face interaction[10].

Table 3.3 displays the summary statistics for the teleworkability measure before and after the pandemic.

TABLE 3.3: Summary Statistics for Teleworkability by Pre/Post Pandemic

	Mean	**Sd**	**Obs.**
Panel A: Pre-pandemic			
Teleworkability	0.38	0.48	17008
Panel B: Post-pandemic			
Teleworkability	0.37	0.48	9926

Notes. This table reports the means, standard deviations, and number of observations of the *Teleworkability* measure for the pre-pandemic (Panel A) and post-pandemic (Panel B).

[9] Although the WFH classification is based on the U.S. O*NET database, it can be reliably applied to Italy by leveraging a crosswalk from U.S. occupational codes (SOC) to international classifications (ISCO), which are harmonized with the Italian Labour Force Survey (EU-LFS). Following Basso et al., 2022, O*NET-based measures are mapped to ISCO codes through a weighted procedure that adjusts for differences in occupational structures between the U.S. and Europe.

[10] Basso et al., 2022 introduce a novel classification for the “teleworkability" of 3-digit ISCO occupations into four categories based on their exposure to infection risk and feasibility for remote work, using data from the CSP, EU-LFS, and O*NET. The categories are the following: (1) jobs that can be done from home, (2) jobs that cannot be done from home but have low infection risk, (3) unsafe jobs that cannot be done from home and present high infection risk, and (4) essential occupations carried out during lockdowns. For the main analysis, I use the first classification. In Appendix B, I present the main effect using the third classification and the one proposed by Hensvik et al., 2020 as a robustness checks.

On average, 38% of women are in teleworkable occupations before the pandemic.

In Appendix A, I report some statistics on the teleworkable occupations. Table A1 displays examples of occupations that are fully teleworkable and non-teleworkable among women[11] based on the teleworkability score within each occupation. Among women, roles such as keyboard operators, legal associate professionals, and tellers are entirely teleworkable, whereas occupations involving manual or service-oriented tasks, such as assemblers or textile workers, show no potential for remote work.

I rely on this occupational-level teleworkability in two different ways: first, I use it to define the treated and control groups in the Difference-in-Differences design, and then as an exogenous source of variation to instrument actual WFH adoption. Teleworkability is a time-invariant measure based solely on the pre-determined task content of occupations. To generate identifying variation, I interact this index with a post-pandemic dummy, which captures the exogenous expansion of remote work opportunities after 2020. In the two-stage least squares framework, the first stage relates actual WFH adoption to this interaction, while the second stage estimates the effect of predicted WFH on the take-up of maternity leave. I adopt this instrumental variable strategy not only to address potential endogeneity in actual WFH adoption, but also to have a more detailed estimate on compliers.

3.4 Empirical Strategy and Baseline Results

To identify the impact of WFH on fertility, I adopt a two-step empirical strategy. First, I estimate a reduced-form Difference-in-Differences (DiD) model that relates the probability of being on maternity leave to the differential exposure of occupations to remote work before and after 2020. This captures the Intention-to-Treat (ITT) effect of being employed in jobs that became more compatible with working from home following the pandemic.

Second, to address the potential endogeneity arising from self-selection into remote work, I implement a two-stage least squares (2SLS) specification. In this setting, actual WFH is instrumented using the teleworkability index proposed by Basso et al., 2022, interacted with the post-pandemic expansion of remote work opportunities. The index measures, for each occupation, the extent to which tasks can be performed from home based on their content and skill requirements.

The underlying idea is that the COVID-19 shock made this occupational feature suddenly relevant for the adoption of remote work: jobs that were more teleworkable by design became disproportionately more likely to be performed from home, independently of workers' fertility plans or employer-specific policies. This exogenous variation isolates the component of

[11] In Table A2, I report examples of occupations that are fully teleworkable also for men. In general, male workers in professional and client-facing roles, such as teaching associates and business service agents, are fully teleworkable, while those in agriculture, cleaning, or craft-related occupations exhibit zero teleworkability.

remote work adoption driven by the intrinsic feasibility of the occupation, satisfying both the relevance and exclusion conditions for instrument validity.

Within this framework, the DiD specification corresponds to the ITT estimate, while the 2SLS recovers the Local Average Treatment Effect (LATE) of actual remote work on fertility among women whose probability of working from home increased because of the pandemic-induced shift in teleworkable occupations.

Therefore, I run the following specification:

$$Maternity\ Leave_{i,t} = \alpha_t + \gamma_r + \theta_s + \lambda_o + \beta\ Teleworkability_o \times Post_t + \varphi\, \mathbf{X}'_{i,t} + \varepsilon_{i,t} \tag{3.1}$$

where $Maternity\ Leave_{i,t}$ is a binary indicator equal to 1 if woman i is on maternity leave during the reference week at time t. The variable $Teleworkability_o \times Post_t$ is an interaction between a dummy indicating whether occupation o can accommodate remote working, based on the classification by Basso et al., 2022, and a post-treatment dummy equal to 1 for years after 2020. The vector $\mathbf{X}'_{i,t}$ includes individual-level controls such as education, family size, and civil status. This specification also includes fixed effects for year (α_t), region (γ_r), industry (θ_s), and occupation (λ_o). Standard errors are clustered at the individual level.

My identification assumption is that, in the absence of the pandemic, fertility trends would have evolved similarly across occupational groups. Under this assumption, the coefficient on the interaction term β can be interpreted as the causal effect of increased remote work feasibility on fertility decisions.

Table 3.4 reports the results from the specification (3.1).

TABLE 3.4: DiD Estimates of Teleworkability on Women's Fertility Decisions

	Maternity Leave	**Maternity Leave**
Teleworkable × Post	0.03^{***}	0.03^{***}
	(0.01)	(0.01)
Observations	27048	27048
Controls	No	Yes

Notes. This table reports the estimated coefficient on the interaction between Teleworkability and Post from (3.1). Column (1) includes year and occupation fixed effects. Column (2) adds individual-level controls (age, education, type of contract, family size, blue-collar status, and firm size), along with year, region, industry, and occupation fixed effects. The main effects of Post and Teleworkability are not reported due to perfect collinearity with the included fixed effects. Standard errors (in parentheses) are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The coefficient β is positive and statistically significant, also when controlling for year, region, industry, and occupation fixed effects. This suggests that, after the pandemic, women in teleworkable occupations were significantly more likely to be on maternity leave relative to those in non-teleworkable occupations. This effect is about 3 percentage points.

This result could reveal that access to occupations that accommodate remote work may have played a role in women's fertility behavior immediately after the pandemic. One plausible interpretation is that the increased flexibility associated with teleworkable occupations lowered the perceived opportunity cost of childbearing, making it easier for women to combine work and family responsibilities. This interpretation is consistent with prior literature highlighting how work-family reconciliation policies and job flexibility can influence fertility decisions (see, for instance, Vignoli et al., 2022; Harknett et al., 2014).

Importantly, this effect emerges despite the fact that not all women in teleworkable occupations actually worked remotely, underscoring the relevance of occupational feasibility itself as a driver of behavioral responses. In other words, the option value of remote work - associated with the occupation - may have altered women's fertility decisions by relaxing time constraints and lowering the expected opportunity cost of childbearing, even in the absence of actual WFH adoption.

To interpret this estimate causally, the identifying assumption of parallel trend must hold - namely, that in the absence of the pandemic shock, fertility trends for women in teleworkable and non-teleworkable occupations would have evolved in parallel. Formally, the model can be expressed as:

$$Maternity\ Leave_{i,t} = \alpha_t + \gamma_r + \theta_s + \lambda_o + \sum_{t=2015, t\neq 2019}^{2022} \beta_t (Teleworkability_o \times Year_t) + \varphi\, \mathbf{X}'_{i,t} + \varepsilon_{i,t} \tag{3.2}$$

Figure 3.1 that fertility trajectories were parallel across the two groups before 2020, supporting the validity of this identifying assumption. Moreover, the effect seems to be more pronounced in the first two years after the pandemic.

The estimation from (3.1) would reflect the ITT effect, as it measures the average impact of exposure to occupations that are more likely to adopt remote work, without conditioning on actual WFH adoption.

However, this does not identify the causal effect of remote work itself, since not all individuals in teleworkable occupations necessarily transitioned to WFH, and selection into remote work may be correlated with unobserved preferences for fertility.

FIGURE 3.1: Parallel Trends

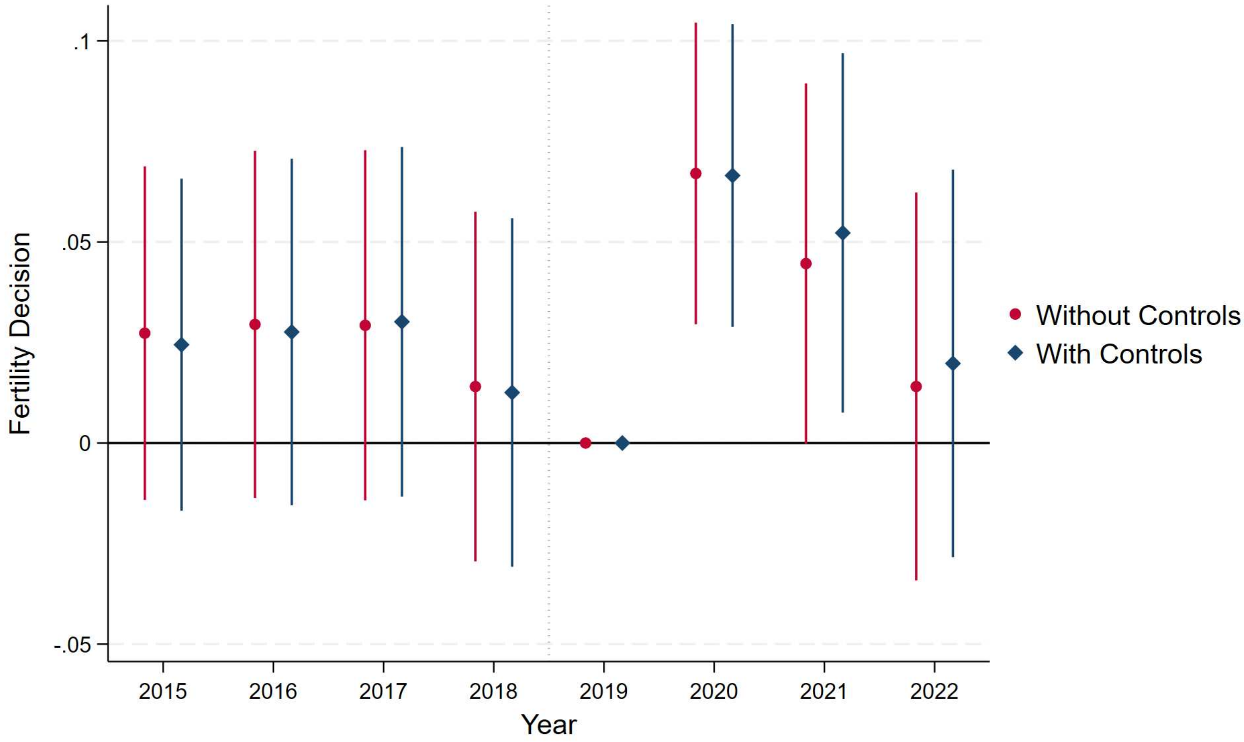

Notes. This figure plots the estimated coefficients from the event-study specification associated with Equation (3.2), showing differences in the probability of being on maternity leave between teleworkable and non-teleworkable occupations over time. The vertical line marks the onset of the COVID-19 pandemic (2020). Estimates are shown both without controls (red circles) and with individual, regional, sectoral, and occupational controls (blue diamonds).

To assess whether the observed relationship reflects a causal effect of actual WFH adoption, I implement a 2SLS strategy that instruments individual-level use of remote work with the interaction between occupational teleworkability by a post-2020 indicator.

The underlying assumption is that, conditional on fixed effects and controls, the variation in WFH adoption driven by teleworkability after the pandemic is exogenous to unobserved determinants of fertility. Then, my main equation is:

$$Maternity\ Leave_{i,t} = \alpha_t + \gamma_r + \theta_s + \lambda_o + \beta\, WFH_{i,t} + \varphi\, \mathbf{X'}_{i,t} + \varepsilon_{i,t} \tag{3.3}$$

As stated above, estimating (3.3) by OLS is likely to yield a biased and inconsistent estimate of β, since the individual decision or opportunity to work from home is endogenous. Women who are able or willing to work remotely may systematically differ from those who do not, in ways that are correlated with fertility preferences and unobserved job or family characteristics. For instance, individuals with stronger family orientation or more flexible employers may both be more likely to adopt remote work and to plan childbirth, leading to an upward bias in the OLS estimate. Conversely, women facing greater career constraints may self-select out of remote positions, generating a downward bias.

To address this endogeneity concern, I instrument the actual use of WFH with the teleworkability index proposed by Basso et al., 2022, which captures the exogenous feasibility of performing job tasks from home.

Formally, my first-stage equation is specified as:

$$WFH_{i,t} = \alpha_t + \gamma_r + \theta_s + \lambda_o + \delta\, Teleworkability_o \times Post_t + \varphi\, \mathbf{X'}_{i,t} + \varepsilon_{i,t} \tag{3.4}$$

where $WFH_{i,t}$ is a binary indicator equal to 1 if individual i was working from home at time t. The instrument, $Teleworkability_o \times Post_t$, captures whether the occupation is classified as teleworkable (based on Basso et al., 2022), interacted with an indicator for the post-2020 period. $\mathbf{X}'_{i,t}$ includes the same individual-level controls as in the DiD specification. The model also includes fixed effects for year (α_t), region (γ_r), industry (θ_s), and occupation (λ_o).

A strong and statistically significant first-stage coefficient on the interaction term δ would confirm that occupational teleworkability, interacted with the post-2020 period, is a relevant predictor of individual WFH adoption and therefore a valid instrument.

Table 3.5 presents the results from the first-stage and reduced-form regressions used to implement the 2SLS strategy.

TABLE 3.5: First Stage and Reduced Form (2SLS Components)

	First Stage		**Reduced Form**	
	WFH		**Maternity Leave**	
Teleworkable × Post	0.19***	0.19***	0.03***	0.03***
	(0.02)	(0.02)	(0.01)	(0.01)
Observations	27,048	27,048	27,048	27,048
F-statistic	**60.65**	**60.61**		
Controls	No	Yes	No	Yes

Notes. This table reports the first-stage from (3.4) and reduced-form regressions from (3.1) underlying the 2SLS estimation for women in the sample. Columns (1)–(2) report the first stage, where the dependent variable is the probability of working from home (WFH). Columns (3)–(4) correspond to the reduced form, where the dependent variable is being on maternity leave. Columns (2) and (4) include individual-level controls and fixed effects for year, region, industry, and occupation. Standard errors (in parentheses) are clustered at the individual level * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The first-stage estimates confirm that the interaction between occupational teleworkability and the post-pandemic period is a strong and significant predictor of actual WFH adoption. The first-stage coefficient is 0.19, with F-statistics well above the conventional threshold ($\geq$ 60), alleviating concerns about weak instruments.

The 2SLS estimator isolates variation in WFH adoption that is plausibly exogenous, allowing me to recover a LATE for a specific subpopulation of compliers - i.e., individuals whose likelihood of WFH changed due to the increased teleworkability of their occupation in the post-pandemic period.

The key insight from Table 3.5 is that the ratio of the reduced-form coefficient (β), estimated from the equation (3.1), to the first-stage coefficient (δ), estimated from the equation (3.4), yields the 2SLS estimate:

$$\theta_{2SLS} = \frac{\beta}{\delta} \tag{3.5}$$

The Wald-type estimator identifies the LATE of actual WFH on fertility, under the assumption that the instrument satisfies the exclusion restriction. Dividing the reduced-form (0.03) by the first-stage coefficient (0.19), implies that actual WFH adoption increases the probability of being on maternity leave by approximately 15 percentage points.

To interpret this coefficient as the causal effect of working from home on fertility decision, two key assumptions must be satisfied: first, the instrument I adopt must be relevant, meaning that it must predict the probability of real adoption of working from home. The strong and statistically significant first-stage coefficient on the interaction between occupational teleworkability and the post-pandemic period, reported in Table 3.5, confirms that teleworkability is a powerful predictor of actual WFH adoption. A one-unit increase in teleworkability - moving from a job that is entirely non-teleworkable to one that is fully teleworkable - is associated with a 19 percentage point increase in the probability of working from home. The first-stage F-statistic of approximately 60 far exceeds the conventional threshold of 10, indicating that weak instrument concerns are negligible.

I also run the following event-study specification:

$$WFH_{i,t} = \alpha_t + \gamma_r + \theta_s + \lambda_o + \sum_{t=2015, t\neq 2019}^{2022} \beta_t (Teleworkability_o \times Year_t) + \varphi\, \mathbf{X}'_{i,t} + \varepsilon_{i,t} \tag{3.6}$$

Figure 3.2 shows the results from (3.6). After 2020, the same event-study analysis shows a sharp and persistent increase in the probability of WFH among women employed in teleworkable occupations, with the effect peaking in 2021 and remaining significant thereafter.

FIGURE 3.2: Event-Study Teleworkability Index on Actual WFH

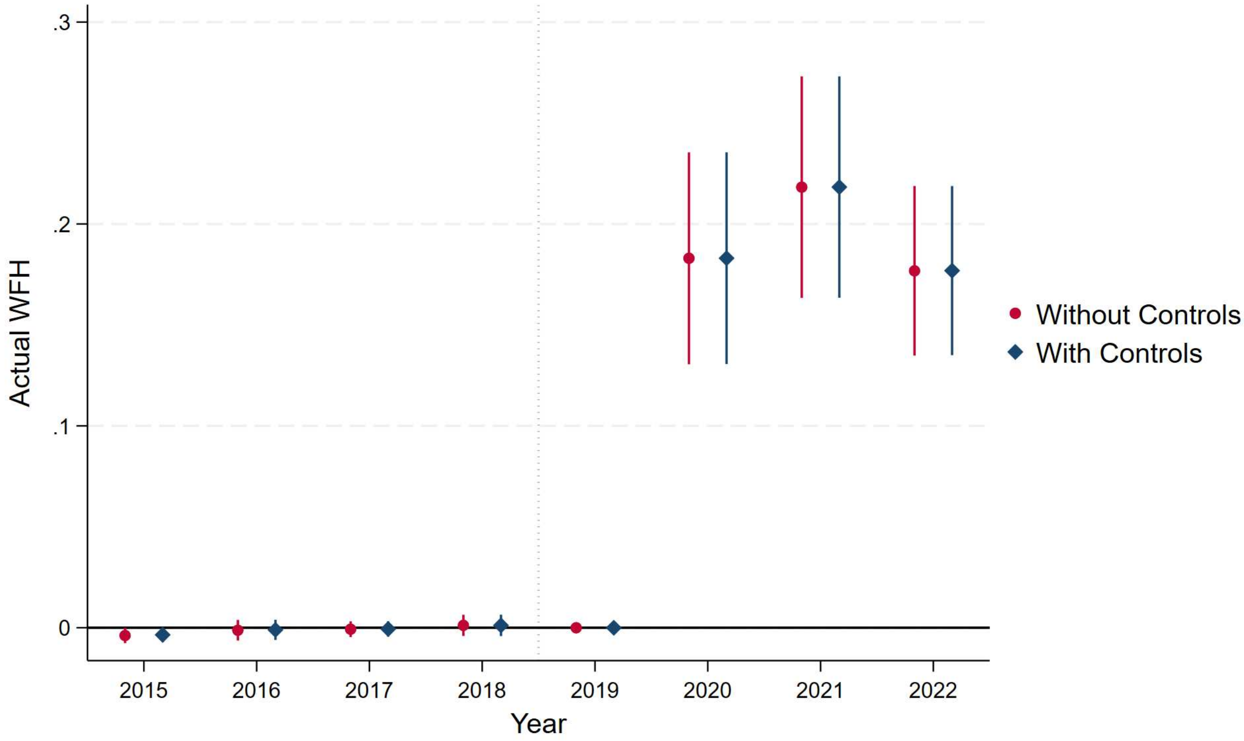

Notes. This figure shows the estimated event-study coefficients from (3.6)for the probability of actual use of WFH on the teleworkability measure. The red line represents estimates without controls, while the blue line includes controls for workers' socio-economic characteristics, region, industry, and year fixed effects. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic.

Overall, these post-pandemic dynamics confirm that the instrument captures the exogenous expansion of remote work opportunities induced by the COVID-19 shock, rather than endogenous occupational or family-related choices.

Second, the exclusion restriction hypothesis must hold: To assess whether the instrument satisfies the exclusion restriction, I verify that the teleworkability index, constructed entirely from pre-pandemic task characteristics, is not correlated with fertility outcomes prior to 2020. Specifically, I estimate an event-study specification where the teleworkability index is interacted with year dummies. The resulting coefficients, plotted in Figure 3.1, show no significant differential trends in maternity leave between teleworkable and non-teleworkable occupations before the onset of the pandemic, supporting the assumption of parallel pre-trends. This absence of anticipatory effects indicates that teleworkability is predetermined with respect to fertility behavior and thus a valid source of exogenous variation.

The combination of strong predictive power, lack of pre-trends, and the exogenous nature of the teleworkability measure provides robust evidence that the instrument satisfies both the relevance and exclusion conditions required for the 2SLS estimates to identify a causal effect of WFH on fertility decisions.

3.5 Heterogeneity Analysis

To understand the mechanisms underlying the relationship between remote work and fertility, I explore in Appendix B if the effects vary across different subgroups of women.

3.5.1 Main Heterogeneity Patterns

Figure B1 shows the reduced-form estimates of the effect of teleworkable occupations on the probability of being on maternity leave, separately for women with high and low levels of education. The positive and statistically significant effect for highly educated women suggests that the expansion of remote work opportunities following the pandemic increased their likelihood of taking maternity leave .

This effect may reflect that highly educated women are more likely to be employed in occupations where remote work is feasible and institutionally supported. The additional flexibility provided by these jobs might reduce the opportunity cost of childbearing, allowing women to reconcile work and family life more easily. In contrast, the effect for low-educated women is not statistically significant and more imprecisely estimated, which could be due to limited access to teleworkable occupations or greater exposure to jobs with lower autonomy and fewer formal leave entitlements.

Figure B2 displays reduced-form estimates by occupational category, distinguishing between blue-collar and white-collar occupations. The positive association between teleworkable occupations and the probability of being on maternity leave is more pronounced for white-collar jobs. In contrast, the effect for blue-collar women is highly imprecise and not statistically distinguishable from zero. These findings reinforce the idea that remote work affects fertility primarily in jobs where it is technologically and institutionally feasible.

Taken together, the evidence from education and occupation suggests that the fertility response to WFH is stronger among women who are more career-attached and face higher opportunity costs of childbearing. This motivates a closer look at whether the magnitude of the effect also varies by earnings capacity and labor-market intensity.

Figure B3 shows that the effect of post-COVID remote work exposure on the probability of being on maternity leave is significantly more pronounced among high-income women compared to their lower-income counterparts. While the predicted probability of taking maternity leave remains relatively stable or even decreases slightly for low-income women after the pandemic, it increases notably for women in the top tercile of the income distribution. This divergence suggests that the flexibility introduced by teleworkable occupations may have alleviated some of the opportunity costs traditionally associated with childbearing - particularly for women with higher earnings, who face steeper trade-offs between career advancement and family formation.

In other words, the ability to work remotely may have enabled high-income women to better reconcile their professional responsibilities with motherhood, thereby increasing the likelihood of taking birth-related leave. These findings are consistent with the interpretation that WFH acts as a mechanism that reduces barriers to fertility for economically advantaged women, reinforcing the importance of considering heterogeneity in economic constraints and labor market attachment when evaluating the broader effects of WFH.

In line with this interpretation, Figure B4 shows that the positive association between remote work exposure and the probability of taking maternity leave is not limited to high-income women. It is also more pronounced among full-time workers compared to their part-time counterparts. While full-time employment is generally associated with greater labor market attachment and higher opportunity costs of childbearing, the introduction of remote work may have alleviated these constraints by offering more flexible working arrangements. This suggests that the benefits of remote work in supporting fertility decisions are particularly salient for women with stronger labor force commitment, who might otherwise have faced greater trade-offs between career and family.

Overall, these patterns support the interpretation of WFH as a mechanism that lowers the opportunity cost of motherhood for women with higher attachment to the labor market.

3.5.2 Additional Consistency Checks

I also examine heterogeneity along dimensions that primarily serve as a consistency check, rather than distinct mechanisms.

Figure B5 reports results by age group. As expected, the effect is concentrated among younger women - those of fertile age - while it is smaller and statistically insignificant among older women. This pattern confirms that the estimated relationship aligns with biological and life-cycle fertility patterns.

Figure B6 explores differences by family size. The association between teleworkable occupations and maternity leave appears slightly stronger for women with larger families, consistent with the notion that remote work is more relevant for those facing greater household responsibilities. However, given the potential endogeneity of family size, these estimates should be interpreted primarily as a robustness check.

Figure B7 and Figure B8 show results by sector and firm size. The effect is more pronounced in the service sector - where remote work is more institutionalized - and in smaller firms, although differences are modest.

Overall, these additional analyses confirm that the main findings are not driven by compositional factors or by specific subsets of the sample. The heterogeneous effects suggest that remote work primarily benefits women for whom work-family trade-offs are most binding.

3.6 Robustness Checks and Placebo Tests

A potential concern with the main analysis relates to how the feasibility of remote work is measured. The baseline specification relies on a binary indicator of *teleworkability*, which classifies occupations as either teleworkable or non-teleworkable based on the nature of the tasks performed. While intuitive and widely used, this dichotomous approach abstracts from important within-occupation variation and may introduce measurement error. In particular, it captures the *potential* ability to work remotely - derived from occupational characteristics - rather than the actual extent to which remote work is implemented in practice. As a result, the binary measure might fail to reflect the true heterogeneity in WFH exposure across occupations and over time.

To address this concern, I use the continuous teleworkability index proposed by Hensvik et al., 2020, which is constructed from the average share of hours worked from home by occupation. This measure is based on observed working patterns rather than potential feasibility, thus providing a more accurate proxy for actual WFH exposure. I re-estimate the main specification, replacing the binary treatment variable with this continuous indicator.

The results in Figure B9 show that the continuous measure closely predicts the observed distribution of remote work. Moreover, Figure B10 confirms that the estimated effects remain stable when this measure is used, suggesting that the findings are not driven by how the teleworkability index is constructed but rather reflect genuine differences in exposure to remote work opportunities across occupations.

A second robustness check focuses on the nature of occupations that cannot be performed remotely but instead require high levels of physical proximity. The aim here is to verify that the estimated effects are not merely capturing occupational differences in exposure to in-person work during the pandemic. To this end, I use the occupational classification developed by Basso et al., 2022, which distinguishes between *safe* and *unsafe* jobs. In their framework, unsafe jobs are those that cannot be carried out from home and that involve frequent close contact with colleagues or customers—such as teachers, health professionals, and waiters - thus being particularly constrained by lockdowns and social distancing measures during COVID-19.

The results, reported in Figure B11 and Figure B12, show that the main findings are robust to this alternative classification. Occupations identified as unsafe - those most affected by restrictions on in-person work - exhibit patterns similar to those found using the teleworkability measure, strengthening the interpretation that the observed effects reflect genuine differences in remote work feasibility rather than sector-specific shocks.

A potential source of simultaneity arises because the LFS records both maternity leave and WFH status for the same reference week. In principle, women on maternity leave should not be recorded as working from home at the same time. To address this, I re-estimate specification in (3.1) after excluding the small number of observations (73 cases) where both conditions hold (*i.e.*, women reported being on maternity leave and working from home in the same reference week). Table Table B1 reports the results: they remain virtually unchanged, confirming that the main estimates are not driven by contemporaneous reporting or measurement overlap between WFH and fertility status.

To further validate the identification strategy, I perform a series of placebo tests. First, I test for pre-trends in fertility decisions by estimating an event-study model where the interaction between teleworkability and year dummies is included for the pre-pandemic period (2015–2018). Figure B13 shows that the estimated coefficients are close to zero and statistically insignificant, and the joint F-test of pre-trends cannot be rejected ($p = 0.61$), supporting the validity of the parallel trends assumption.

Second, I conduct placebo DiD exercises by shifting the treatment period to pre-COVID years (2019, 2018, 2017, and 2016). The results in Table B2 show that all placebo coefficients are small and statistically insignificant, providing no evidence of spurious treatment effects prior to the pandemic. Taken together, these checks confirm that the estimated effects are not driven by pre-existing trends or measurement artifacts, but by the post-pandemic expansion of remote work opportunities.

3.7 Concluding Remarks

This chapter investigates how the expansion of remote work following the COVID-19 pandemic affected women's fertility decisions in Italy. Using microdata from the Italian Labour Force Survey (2015–2022) and a task-based measure of occupational teleworkability, I exploit the sudden and uneven diffusion of remote work across occupations to estimate its impact on the probability of being on maternity leave, used as a proxy for childbirth.

The empirical analysis combines a Difference-in-Differences design with an instrumental variable strategy. The first step estimates the Intention-to-Treat effect by comparing fertility behavior among women employed in more versus less teleworkable occupations before and after 2020. The second step instruments actual WFH adoption with the interaction between occupational teleworkability and the post-pandemic period, thus isolating the exogenous variation in remote work opportunities driven by the pandemic shock. The results are consistent across both approaches: after 2020, women in teleworkable occupations became significantly more likely - by around three percentage points - to take maternity leave relative to those in non-teleworkable jobs. The first-stage estimates confirm that teleworkability strongly predicts actual WFH adoption, and the event-study analysis validates the absence of pre-trends in fertility, supporting the credibility of the identification strategy.

The findings suggest that access to occupations compatible with remote work facilitated the reconciliation of work and family responsibilities during and after the pandemic. The estimated effect likely reflects a reduction in the opportunity cost of childbearing: women in teleworkable jobs may have perceived greater flexibility and security in managing employment and motherhood, even when not directly working from home. This interpretation is further supported by the heterogeneity analysis, which shows that the effect is stronger for highly educated, high-income, and full-time women - groups for whom time constraints and career - family trade-offs are typically more binding.

These results provide new evidence on how workplace flexibility can influence fertility behavior, particularly in labor markets characterized by rigid working arrangements and persistent gender inequalities.

By lowering the costs of maintaining labor force attachment around childbirth, remote work may have acted as an informal family policy instrument, supporting fertility without direct public intervention. However, the benefits appear to be concentrated among more advantaged women, suggesting that the expansion of WFH could inadvertently widen inequalities in the ability to combine work and family life.

While the COVID-19 pandemic created an exceptional context, its legacy in shaping work organization and family decisions is likely to persist. Understanding how the structural diffusion of remote work interacts with women's labor supply and fertility choices remains crucial for designing policies that promote both gender equality and demographic sustainability.

3.8 Appendix A: Additional Descriptive Statistics

TABLE A1: Examples of Teleworkable and Non-Teleworkable Occupations (Women)

Occupation	Teleworkability	Obs.
Teleworkable Occupations		
Street Vendors and Related Workers	✓	46
Legal, Social and Religious Associate Professionals	✓	1,726
Mathematicians, Actuaries and Statisticians	✓	48
Tellers, Money Collectors and Related Clerks	✓	2,881
Keyboard Operators	✓	766
Non-Teleworkable Occupations		
Assemblers	✗	1,492
Sewing, Textile and Related Trades Workers	✗	50
Material-Recording and Transport Clerks	✗	4,111
Protective Services Workers	✗	143
Artistic, Cultural and Culinary Associate Professionals	✗	638

Notes. This table reports five occupations with an average teleworkability score of 1 (i.e., fully teleworkable) and five with a score of 0 (i.e., fully non-teleworkable) among female workers.

TABLE A2: Examples of Teleworkable and Non-Teleworkable Occupations (Men)

Occupation	Teleworkability	Obs.
Teleworkable Occupations		
Professional Teaching Associates	✓	3,070
Life Science Professionals	✓	270
Social Work and Counseling Professionals	✓	366
Business Services Agents	✓	781
Client Information Workers	✓	3,665
Non-Teleworkable Occupations		
Handicraft and Printing Workers	✗	2,815
Market Gardeners and Crop Growers	✗	2,966
Domestic, Hotel and Office Cleaners and Helpers	✗	4,584
ICT Operations Technicians	✗	464
Advertising and Public Relations Professionals	✗	743

Notes. This table reports five occupations with an average teleworkability score of 1 (i.e., fully teleworkable) and five with a score of 0 (i.e., fully non-teleworkable) among male workers.

3.9 Appendix B: Additional Tables and Figures

TABLE B1: DiD Estimates of Teleworkability on Women's Fertility Decisions

	Maternity Leave	**Maternity Leave**
Teleworkable × Post	0.03**	0.03***
	(0.01)	(0.01)
Observations	26,975	26,975
Controls	No	Yes

Notes. This table reports Difference-in-Differences estimates of the effect of occupational *Teleworkability* on the probability that a woman is on maternity leave, as specified in equation (3.1), while excluding contemporaneous cases. The coefficient on *Teleworkable* × *Post* captures the change in the likelihood of taking maternity leave after 2020 for women employed in occupations more suitable for remote work relative to those in non-teleworkable jobs. Column (1) includes year and occupation fixed effects. Column (2) additionally controls for individual characteristics (age, education, type of contract, family size, blue-collar status, and firm size) and includes year, region, industry, and occupation fixed effects. Standard errors, reported in parentheses, are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE B2: Placebo DiD Estimates

	(1)	(2)	(3)	(4)
Placebo 2019	0.02			
	(0.01)			
Placebo 2018		0.01		
		(0.01)		
Placebo 2017			0.01	
			(0.01)	
Placebo 2016				0.01
				(0.02)
Observations	27048	27048	27048	27048
Controls	Yes	Yes	Yes	Yes

Notes. This table reports placebo DiD estimates where the treatment period is artificially shifted to pre-pandemic years (2016-2019). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

FIGURE B1: Reduce Form by Education

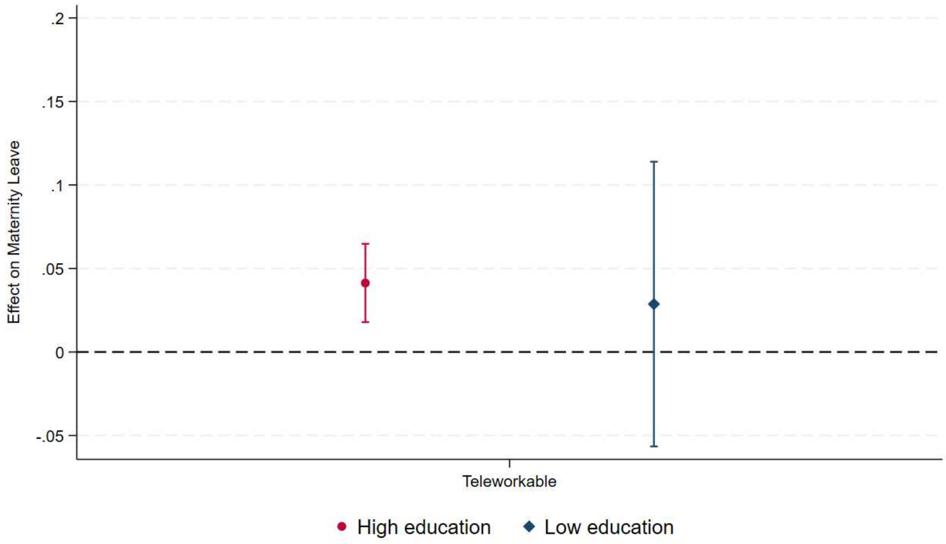

Notes. This figure presents reduced-form estimates from (3.1) of the effect of teleworkable occupations on the probability of being on maternity leave , separately for women with high (red estimate) and low (blue estimate) levels of education.

FIGURE B2: Reduced Form by Occupations

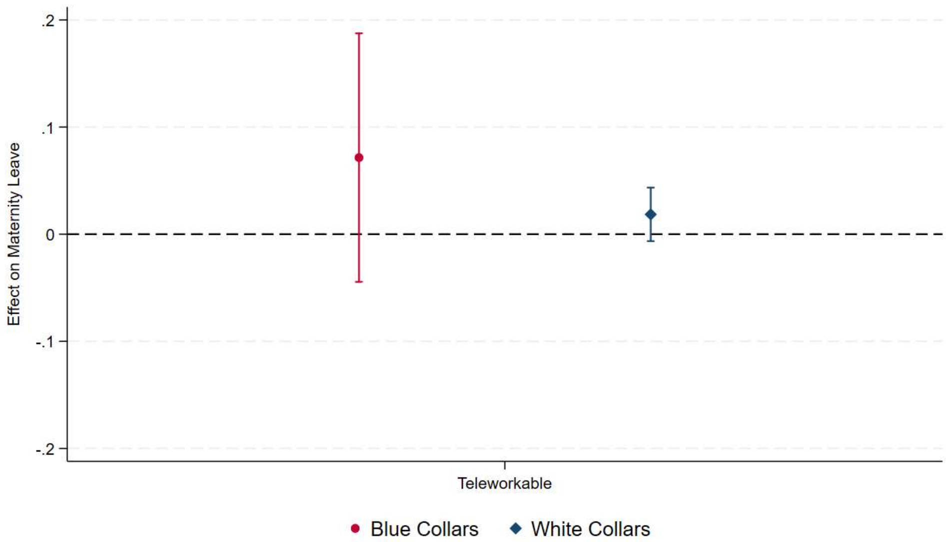

Notes. This figure presents reduced-form estimates from (3.1) of the effect of teleworkable occupations on the probability of being on maternity leave , separately for low skills (red estimate) and high skills (blue estimate) jobs.

FIGURE B3: High vs Low Labor Income

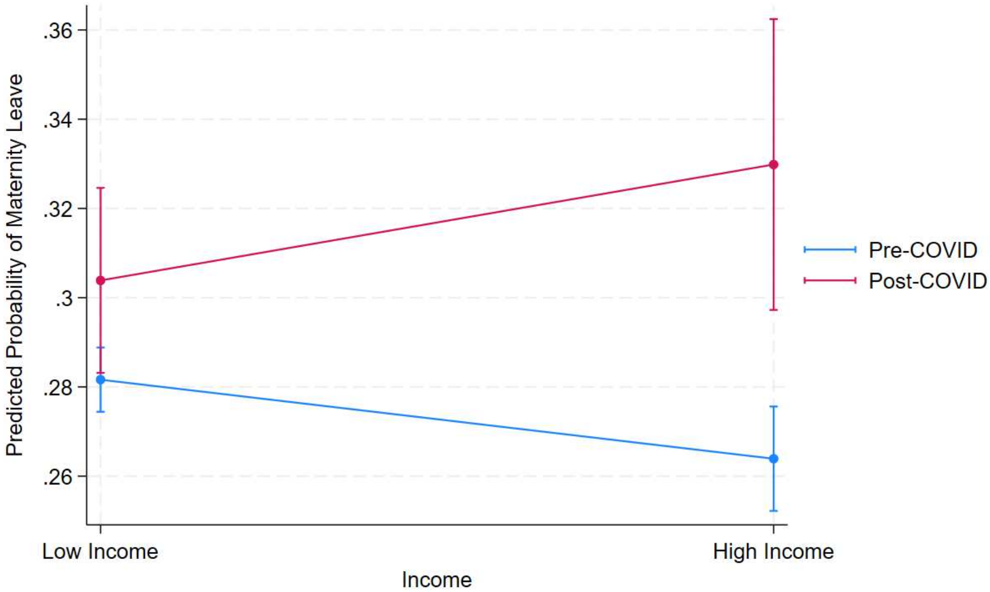

Notes. This figure shows the predicted probability of being on maternity leave by income group (low vs high), separately for the pre- and post-COVID periods from (3.1). The estimates are based on the interaction between post-COVID teleworkability exposure and a dummy for high labor income (top tercile of the distribution).

FIGURE B4: Full-time vs Part-time

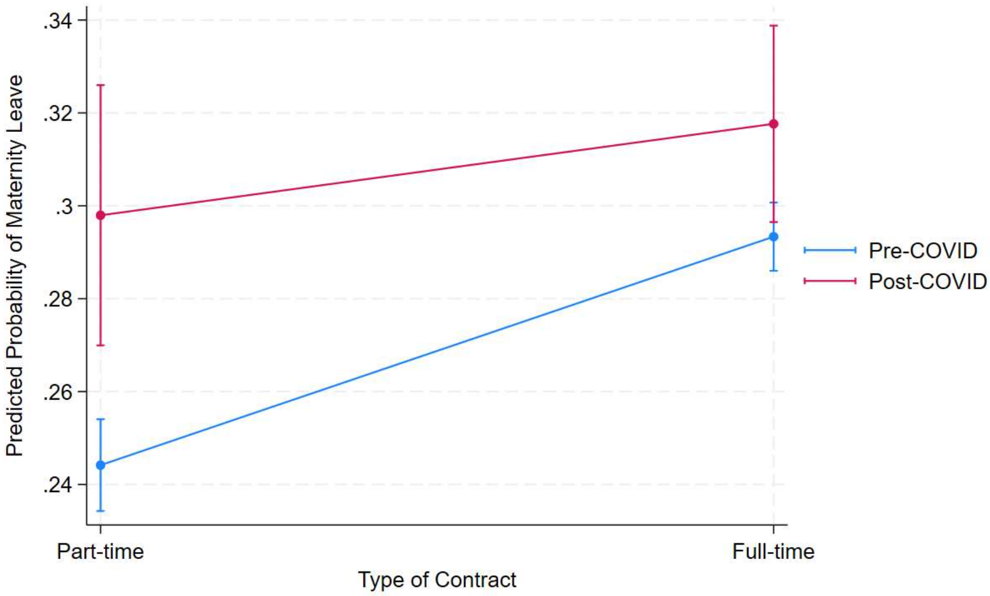

Notes. This figure shows the predicted probability of being on maternity leave by contracts, separately for the pre- and post-COVID periods from (3.1). The estimates are based on the interaction between post-COVID teleworkability exposure and a dummy for full-time.

FIGURE B5: Reduced Form by Age

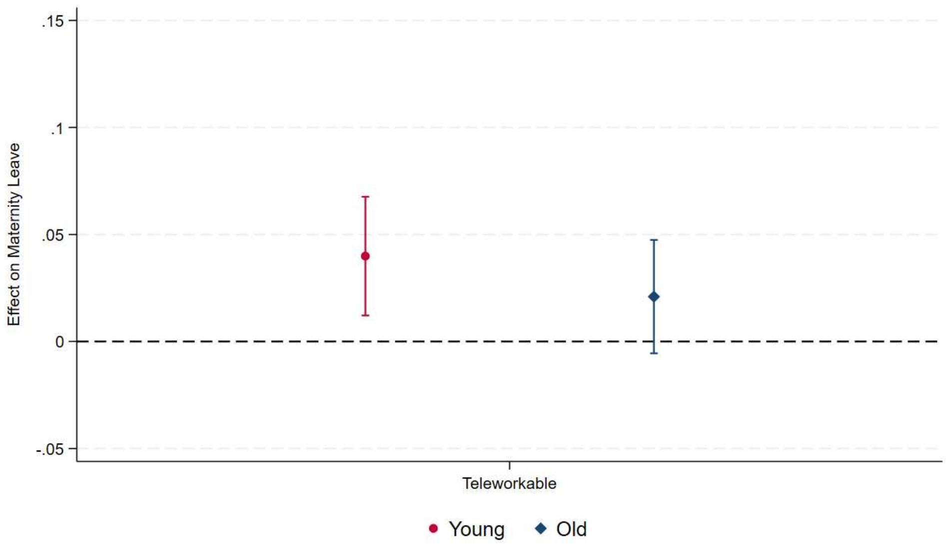

Notes. This figure presents reduced-form estimates of the effect from (3.1) of teleworkable occupations on the probability of being on maternity leave , separately for young women (red estimate) and older (blue estimate) ones.

FIGURE B6: Reduced Form by Family Size

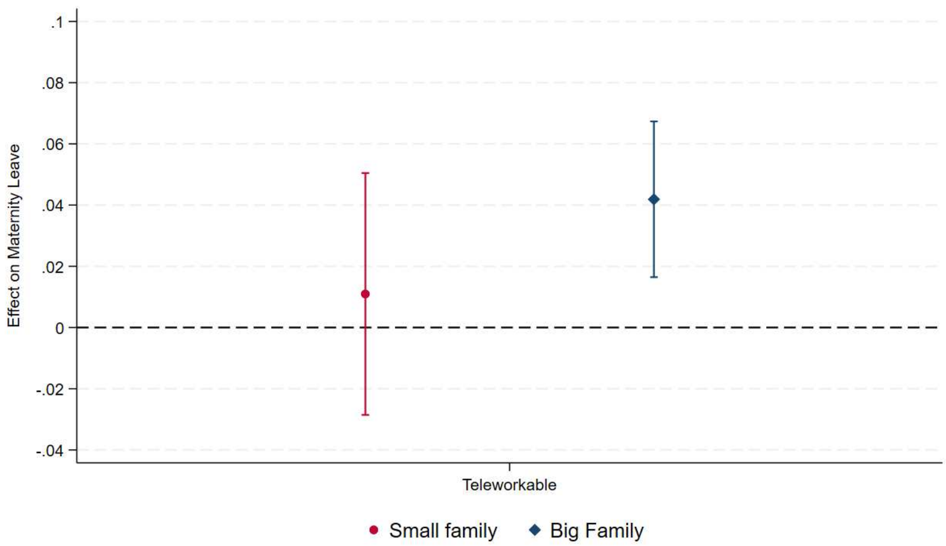

Notes. This figure presents reduced-form estimates from (3.1) of the effect of teleworkable occupations on the probability of being on maternity leave , separately for women with small families (red estimate) and those with larger families (blue estimate).

FIGURE B7: Reduced Form by Main Sectors

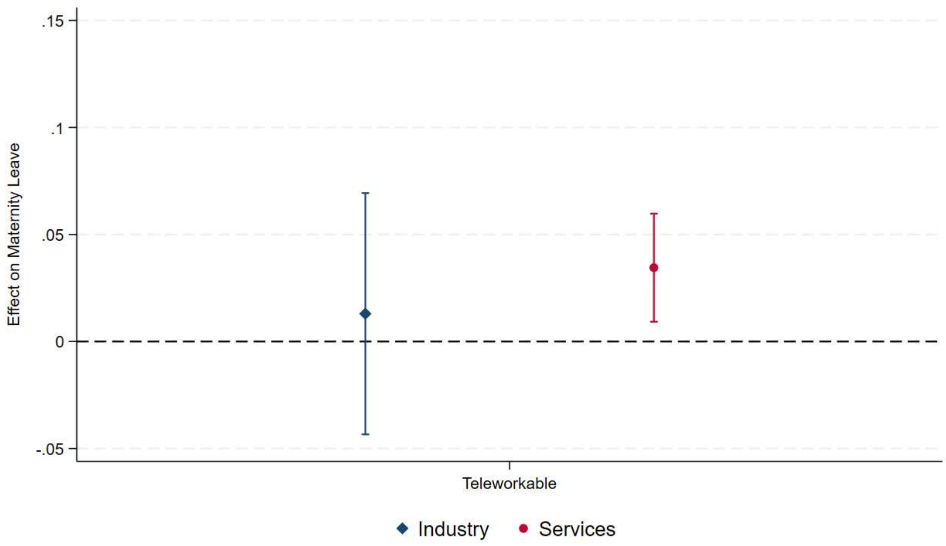

Notes. This figure presents reduced-form estimates of the effect of teleworkable occupations on the probability of being on maternity leave from (3.1), separately for women in industry (red estimate) and those in services (blue estimate).

FIGURE B8: Reduced Form by Firm Size

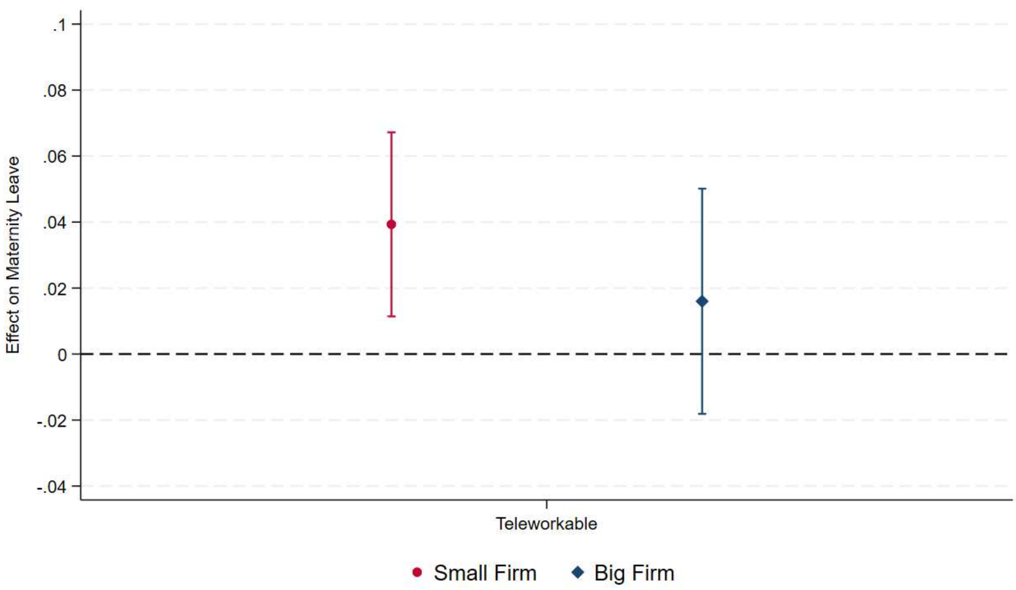

Notes. This figure presents reduced-form estimates of the effect of teleworkable occupations on the probability of being on maternity leave from (3.1), separately for women in small firms (red estimate) and those in big firms (blue estimate).

FIGURE B9: Robustness Check - Hours of WFH as Teleworkability Measure

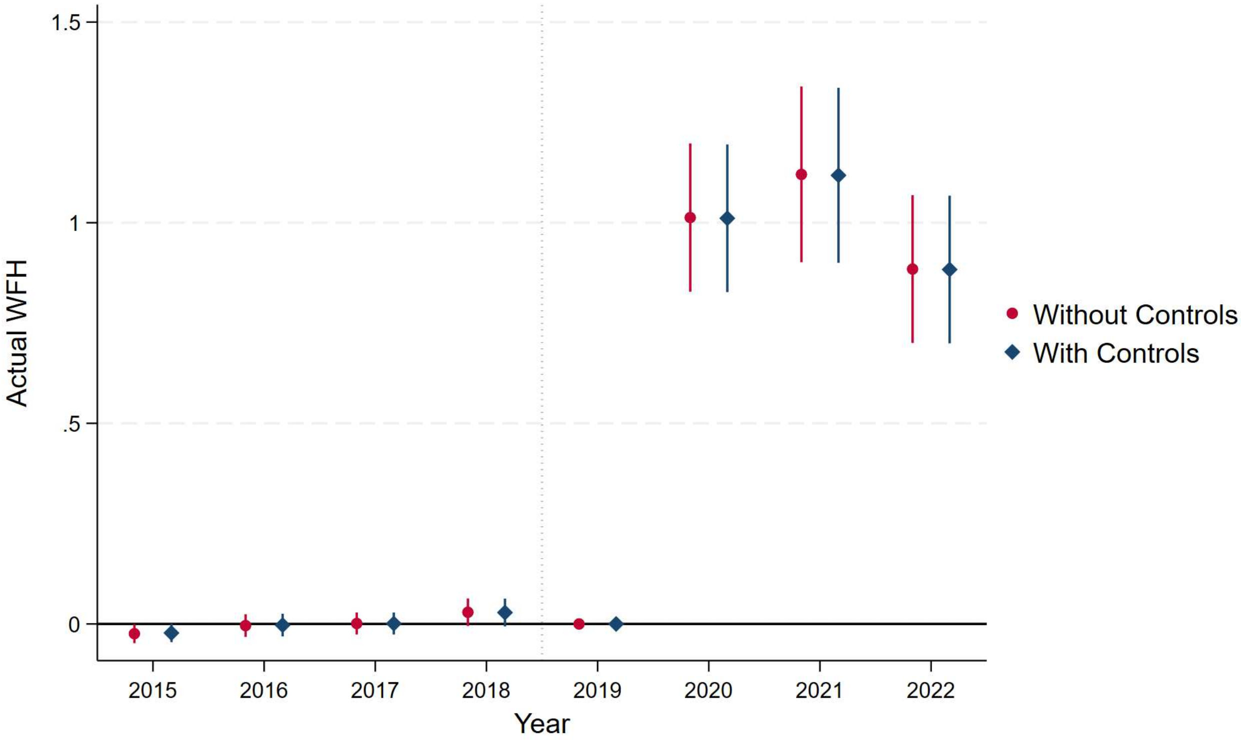

Notes. This figure shows the estimated event-study coefficients for the probability of actual use of WFH from (3.6) on the teleworkability measure based on the hours worked at home. The red line represents estimates without controls, while the blue line includes controls for workers' socio-economic characteristics, region, industry, and year fixed effects. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic

FIGURE B10: Robustness Check - Fertility vs Average Hours of WFH

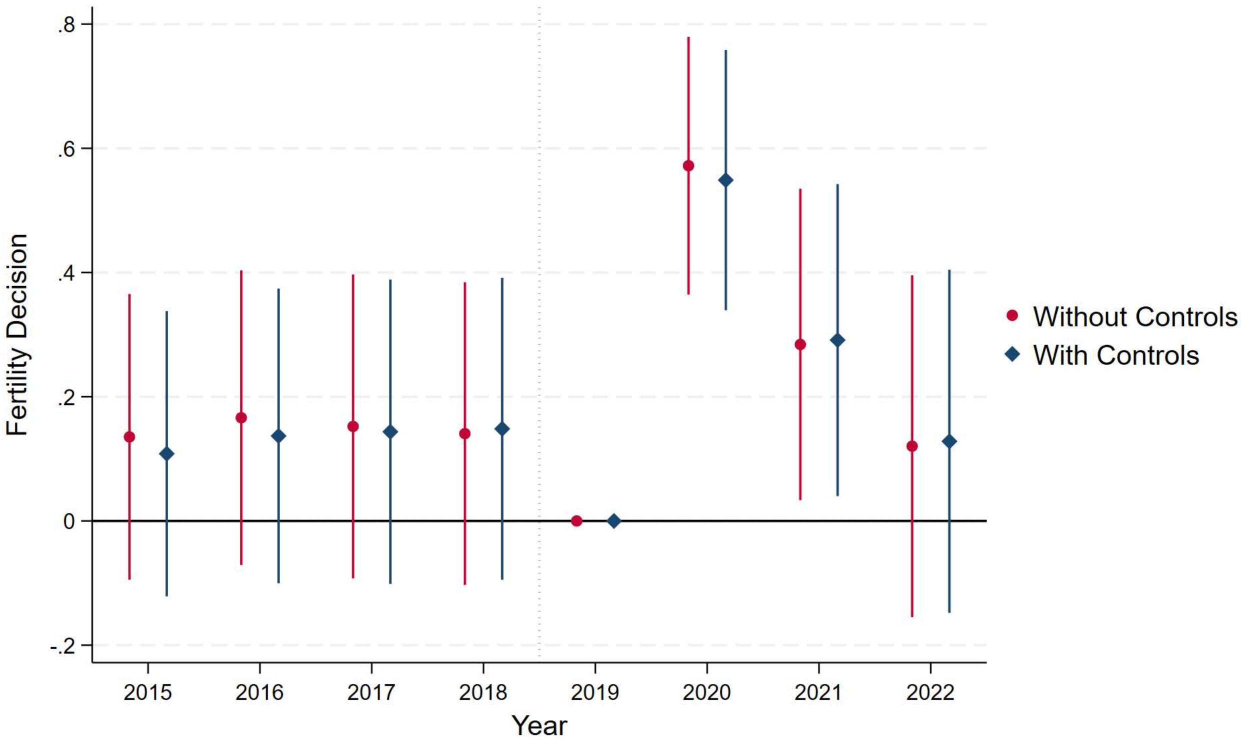

Notes. This figure shows the estimated event-study coefficients for the probability of actual use of WFH on the average hours of WFH as a proxy for teleworkability from (3.1). The red line represents estimates without controls, while the blue line includes controls for workers' socio-economic characteristics, region, industry, and year fixed effects. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic

FIGURE B11: Robustness Check - Hours of WFH as Teleworkability Measure

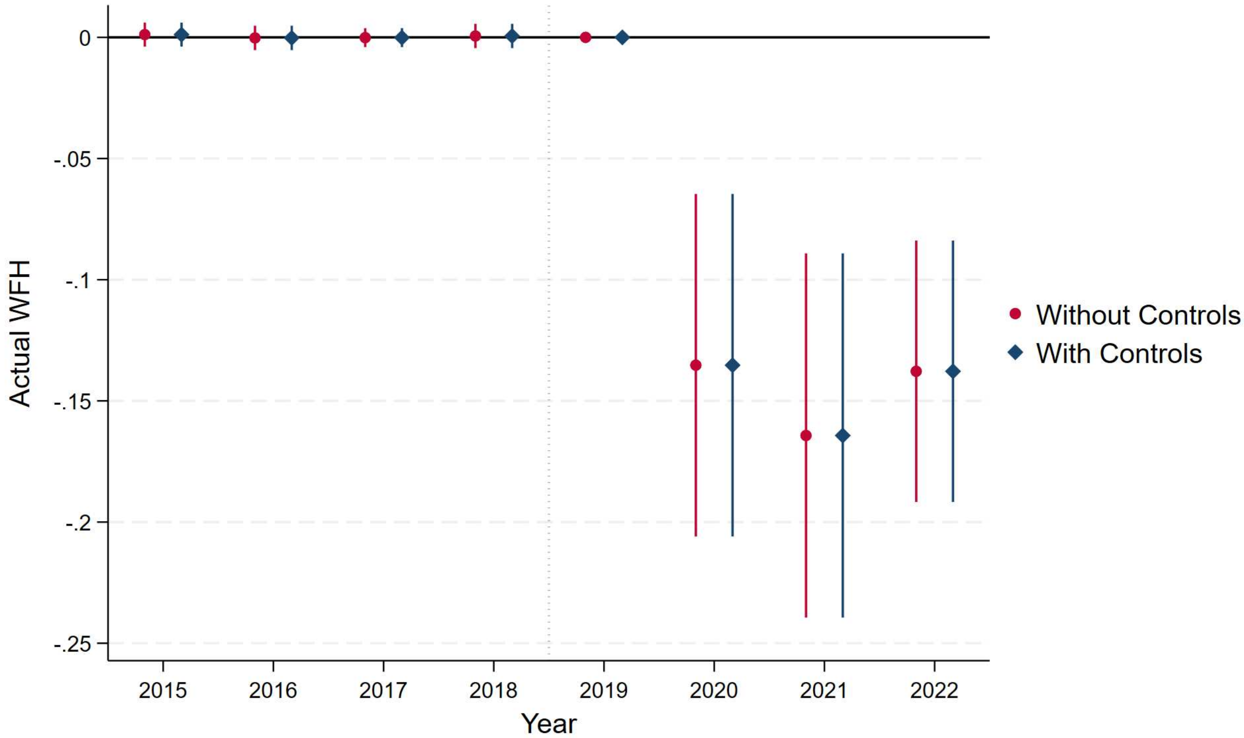

Notes. This figure shows the estimated event-study coefficients for the probability of actual use of WFH from (3.6) on the unsafe job index. The red line represents estimates without controls, while the blue line includes controls for workers' socio-economic characteristics, region, industry, and year fixed effects. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic

FIGURE B12: Robustness Check - Fertility vs Unsafe Job

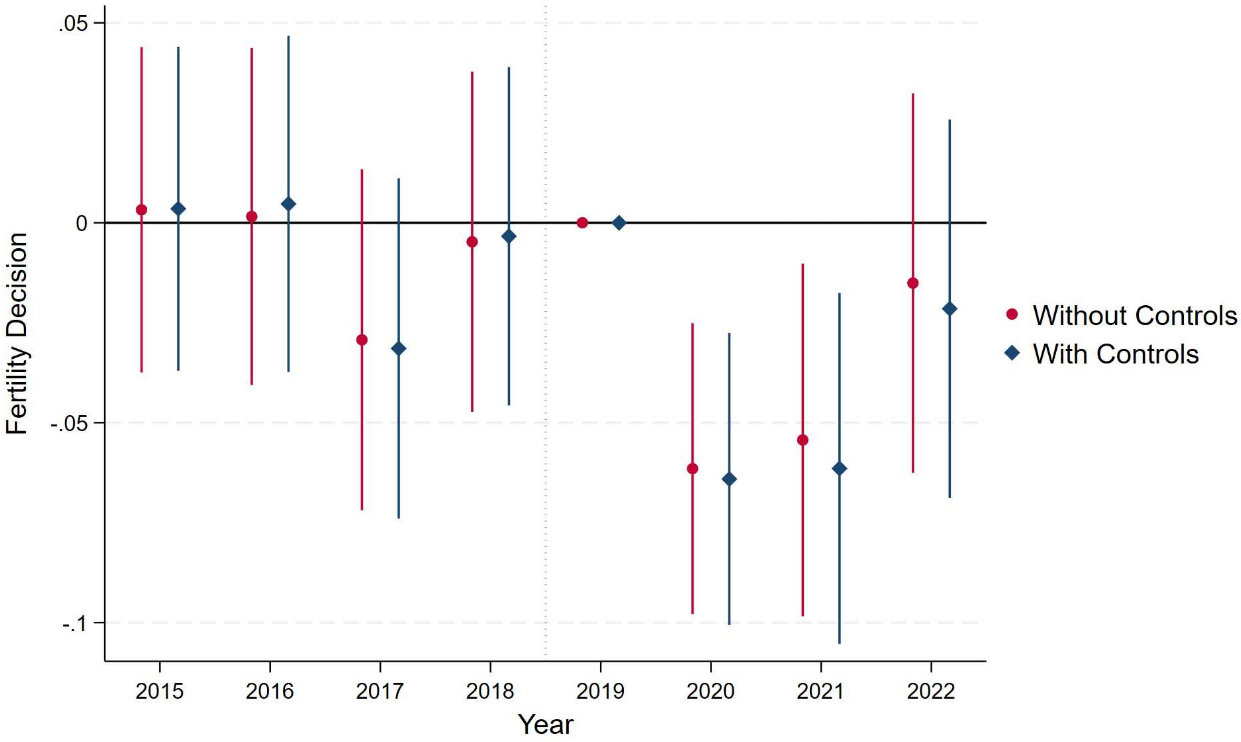

Notes. This figure shows the estimated event-study coefficients for the probability of actual use of WFH from (3.1) on the unsafe jobs index. The red line represents estimates without controls, while the blue line includes controls for workers' socio-economic characteristics, region, industry, and year fixed effects. The shaded areas indicate 95% confidence intervals. The coefficients are normalized relative to the baseline year 2019, allowing for a comparison of trends before and after the onset of the COVID-19 pandemic

FIGURE B13: Test for joint parallel trends

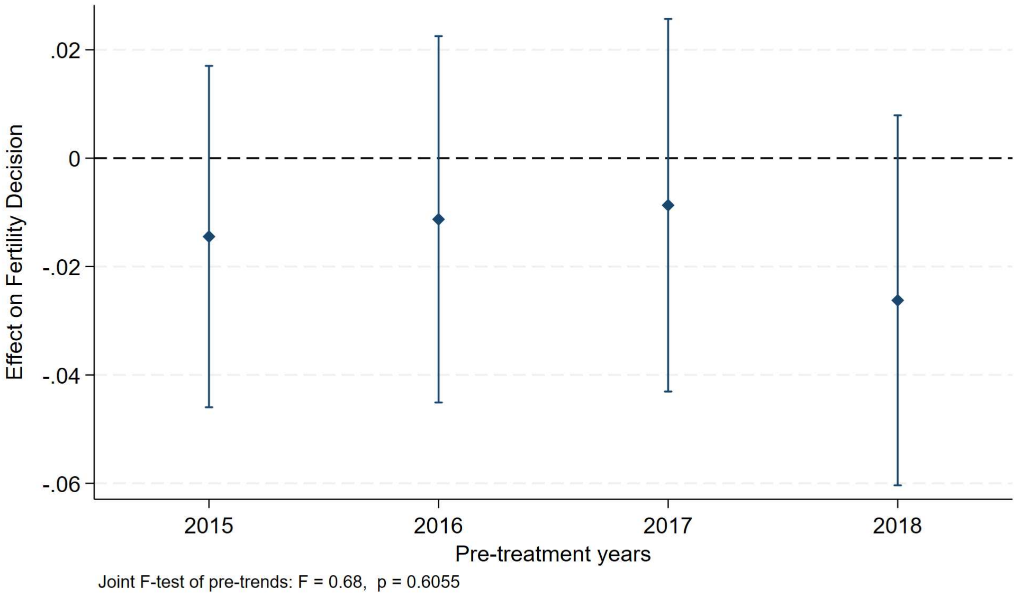

Notes. This figure shows the predicted probability of being on maternity leave by income group (low vs high), separately for the pre- and post-COVID periods. The estimates are based on the interaction between post-COVID teleworkability exposure and a dummy for high labor income (top tercile of the distribution).

Bibliography

Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh (2022). "Work that can be done from home: Evidence on variation within and across occupations and industries". In: *Labour Economics* 74, p. 102083.

Adsera, Alicia (2004). "Changing fertility rates in developed countries. The impact of labor market institutions". In: *Journal of population economics* 17.1, pp. 17–43.

Aguiar, Mark and Erik Hurst (2005). "Consumption versus Expenditure". In: *Journal of Political Economy* 113.5, pp. 919–948. eprint: https://doi.org/10.1086/491590.

Akan, Mert, Jose Maria Barrero, Nicholas Bloom, Thomas Bowen, Shelby R Buckman, Steven J Davis, and Hyoseul Kim (2025). *The new geography of labor markets*. Tech. rep. National Bureau of Economic Research.

Aksoy, Cevat Giray, Jose Maria Barrero, Nicholas Bloom, Steven J Davis, Mathias Dolls, and Pablo Zarate (2022). "Working from home around the world". In: *Brookings Papers on Economic Activity* 2022.2, pp. 281–360.

— (2025). "The global persistence of work from home". In: *Proceedings of the National Academy of Sciences* 122.27, e2509892122.

Alipour, Jean-Victor, Oliver Falck, and Simone Schüller (2023). "Germany's capacity to work from home". In: *European Economic Review* 151, p. 104354.

Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey, and Michèle Tertilt (2020). *The impact of COVID-19 on gender equality*. Tech. rep. National Bureau of economic research.

Angelici, Marta and Paola Profeta (2024). "Smart working: work flexibility without constraints". In: *Management Science* 70.3, pp. 1680–1705.

Angelici, Marta and Giulia Savio (2024). "Gender Gaps in Housework Activities in Europe before and after COVID-19". In: *CESifo Economic Studies* 70.3, pp. 306–319.

Angerer, Xiaohong and Pok-Sang Lam (2009). "Income Risk and Portfolio Choice: An Empirical Study". In: *The Journal of Finance* 64.2, pp. 1037–1055.

Arpino, Bruno, Chiara D Pronzato, and Lara P Tavares (2014). "The effect of grandparental support on mothers' labour market participation: An instrumental variable approach". In: *European Journal of Population* 30.4, pp. 369–390.

Atkin, David, M Keith Chen, and Anton Popov (2022). *The returns to face-to-face interactions: Knowledge spillovers in Silicon Valley*. Tech. rep. National Bureau of Economic Research.

Attanasio, Orazio P. and Agar Brugiavini (Aug. 2003). "Social Security and Households' Saving*". In: *The Quarterly Journal of Economics* 118.3, pp. 1075–1119. eprint: https://academic.oup.com/qje/article-pdf/118/3/1075/5472817/118-3-1075.pdf.

Auclert, Adrien, Hannes Malmberg, Frederic Martenet, and Matthew Rognlie (2021). *Demographics, wealth, and global imbalances in the twenty-first century*. Tech. rep. National Bureau of Economic Research.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2024). "The intertemporal keynesian cross". In: *Journal of Political Economy* 132.12, pp. 4068–4121.

Autor, David, Arindrajit Dube, and Annie McGrew (2023). *The unexpected compression: Competition at work in the low wage labor market*. Tech. rep. National Bureau of Economic Research.

Azar, José and Ioana Marinescu (2024). "Monopsony power in the labor market". In: *Handbook of Labor Economics*. Vol. 5. Elsevier, pp. 761–827.

Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska (2020). "Concentration in US labor markets: Evidence from online vacancy data". In: *Labour Economics* 66, p. 101886.

Babet, Damien and Mathieu Chabaud (2024). *Follow the Money?: Workers' Mobility, Wages and Amenities*. Insee, Institut national de la statistique et des études économiques.

Baghai, Ramin, Margarida Soares, et al. (2023). "Are women more exposed to firm shocks?" In: *Rui Alexandre Colaço and Soares, Margarida, Are Women More Exposed to Firm Shocks*.

Banks, James, Richard Blundell, and Sarah Tanner (1998). "Is there a retirement-savings puzzle?" In: *American Economic Review*, pp. 769–788.

Bao, Lingfeng, Tao Li, Xin Xia, Kaiyu Zhu, Hui Li, and Xiaohu Yang (2022). "How does working from home affect developer productivity?—A case study of Baidu during the COVID-19 pandemic". In: *Science China Information Sciences* 65.4, p. 142102.

Barbieri, Barbara, Silvia Balia, Isabella Sulis, Ester Cois, Cristina Cabras, Sara Atzara, and Silvia De Simone (2021). "Don't call it smart: Working from home during the pandemic crisis". In: *Frontiers in psychology* 12, p. 741585.

Barbieri, Paolo, Rossella Bozzon, Stefani Scherer, Raffaele Grotti, and Michele Lugo (2015). "The rise of a Latin model? Family and fertility consequences of employment instability in Italy and Spain". In: *European societies* 17.4, pp. 423–446.

Barbieri, Teresa, Gaetano Basso, and Sergio Scicchitano (2022). "Italian workers at risk during the Covid-19 epidemic". In: *Italian Economic Journal* 8.1, pp. 175–195.

Barrero, Jose Maria, Nicholas Bloom, Kathryn Bonney, Cory L Breaux, Cathy Buffington, Steven J Davis, Lucia S Foster, Brian McKenzie, Keith Savage, and Cristina Tello-Trillo (2025). *Tapping Business and Household Surveys to Sharpen Our View of Work from Home*. Tech. rep. National Bureau of Economic Research.

Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis (2020). *COVID-19 is also a reallocation shock*. Tech. rep. National Bureau of Economic Research.

Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis (2021a). *Internet access and its implications for productivity, inequality, and resilience*. Tech. rep. National Bureau of Economic Research.

— (2021b). *Why working from home will stick*. Tech. rep. National Bureau of Economic Research.

Barrero, José María, Nicholas Bloom, and Steven J Davis (2023). "The evolution of work from home". In: *Journal of Economic Perspectives* 37.4, pp. 23–49.

Basso, Gaetano, Tito Boeri, Alessandro Caiumi, and Marco Paccagnella (2022). "Unsafe jobs, labour market risk and social protection". In: *Economic Policy* 37.110, pp. 229–267.

Basso, Gaetano, Domenico Depalo, and Salvatore Lattanzio (2023). "Worker flows and reallocation during the recovery". In: *Labour Economics* 83, p. 102399.

Battistin, Erich, Agar Brugiavini, Enrico Rettore, and Guglielmo Weber (2009). "The retirement consumption puzzle: evidence from a regression discontinuity approach". In: *American Economic Review* 99.5, pp. 2209–2226.

Battiston, Diego, Jordi Blanes i Vidal, and Tom Kirchmaier (2021). "Face-to-face communication in organizations". In: *The Review of Economic Studies* 88.2, pp. 574–609.

Becker, Gary S (1992). "Fertility and the economy". In: *Journal of Population Economics* 5.3, pp. 185–201.

Bick, Alexander, Adam Blandin, and Karel Mertens (2023). "Work from Home before and after the COVID-19 Outbreak". In: *American Economic Journal: Macroeconomics* 15.4, 1–39.

Billari, Francesco and Hans-Peter Kohler (2004). "Patterns of low and lowest-low fertility in Europe". In: *Population studies* 58.2, pp. 161–176.

Bloom, Nicholas, Jose Maria Barrero, Steven J. Davis, Benjamin Meyer, and Emil Mihaylov (2023). "Where Managers and Employees Disagree About Remote Work". In: *Harvard Business Review*. Online article.

Bloom, Nicholas, Ruobing Han, and James Liang (2022). *How hybrid working from home works out*. Tech. rep. National Bureau of economic research.

— (2024). "Hybrid working from home improves retention without damaging performance". In: *Nature*, pp. 1–6.

Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying (2015). "Does working from home work? Evidence from a Chinese experiment". In: *The Quarterly journal of economics* 130.1, pp. 165–218.

Boeri, Tito, Alessandro Caiumi, and Marco Paccagnella (2020). "Mitigating the work-safety trade-off". In: *Covid Economics* 2.8, pp. 60–66.

Bottazzi, Renata, Tullio Jappelli, and Mario Padula (2006). "Retirement expectations, pension reforms, and their impact on private wealth accumulation". In: *Journal of Public Economics* 90.12, pp. 2187–2212.

— (2011). "The portfolio effect of pension reforms: evidence from Italy". In: *Journal of Pension Economics & Finance* 10.1, pp. 75–97.

Brilli, Ylenia, Nevena Kulic, and Moris Triventi (2017). "Who cares for the children? Family social position and childcare arrangements in Italy, 2002–12". In: *Childcare, Early Education and Social Inequality*. Edward Elgar Publishing, pp. 31–48.

Brinatti, Agostina, Alberto Cavallo, Javier Cravino, and Andres Drenik (2021). *The international price of remote work*. Tech. rep. National Bureau of Economic Research.

Buckman, Shelby R, Jose Maria Barrero, Nicholas Bloom, and Steven J Davis (2025). *Measuring work from home*. Tech. rep. National Bureau of Economic Research.

Caldwell, Sydnee and Oren Danieli (2024). "Outside Options in the Labour Market". In: *Review of Economic Studies* 91.6, pp. 3286–3315.

Card, David, Ana Rute Cardoso, and Patrick Kline (2016). “Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women”. In: *The Quarterly journal of economics* 131.2, pp. 633–686.

Carroll, Christopher D and Andrew A Samwick (1997). “The nature of precautionary wealth”. In: *Journal of monetary Economics* 40.1, pp. 41–71.

Carta, Francesca and Marta De Philippis (2024). “The forward-looking effect of increasing the full retirement age”. In: *The Economic Journal* 134.657, pp. 165–192.

Casarico, Alessandra and Salvatore Lattanzio (2022). “The heterogeneous effects of COVID-19 on labor market flows: evidence from administrative data”. In: *The Journal of Economic Inequality* 20.3, pp. 537–558.

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”. In: *American Economic Review* 107.12, pp. 3917–3946.

Choudhury, Prithwiraj, Tarun Khanna, Christos A Makridis, and Kyle Schirmann (2024). “Is hybrid work the best of both worlds? Evidence from a field experiment”. In: *Review of Economics and Statistics*, pp. 1–24.

Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, and Geoff Kenny (2024). “Wealth shocks and portfolio choice”. In: *Journal of Monetary Economics*, p. 103632.

Clark, ANDREW E, CONCHITA D’AMBROSIO, and ANTHONY Lepinteur (2022). *Job insecurity, savings and consumption: An Italian experiment*.

Cocco, Joao F, Francisco J Gomes, and Pascal J Maenhout (2005). “Consumption and portfolio choice over the life cycle”. In: *The Review of Financial Studies* 18.2, pp. 491–533.

Croson, Rachel and Uri Gneezy (2009). “Gender differences in preferences”. In: *Journal of Economic literature* 47.2, pp. 448–474.

Cullen, Zoë, Bobak Pakzad-Hurson, and Ricardo Perez-Truglia (2025). “Home Sweet Home: How Much Do Employees Value Remote Work?” In: *AEA Papers and Proceedings*. Vol. 115. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 276–281.

Daminato, Claudio and Mario Padula (2024). “The life-cycle effects of pension reforms: a structural approach”. In: *Journal of the European Economic Association* 22.1, pp. 355–392.

Datta, Nikhil (2022). “Essays on Monopsony Power, Wage Floors and Atypical Work Arrangements”. PhD thesis. UCL (University College London).

— (2023). “The measure of monopsony: the labour supply elasticity to the firm and its constituents”. In.

De Fraja, Gianni, Jesse Matheson, Paul Mizen, James Rockey, Shivani Taneja, and Gregory Thwaites (2022). “Remote Work and Compensation Inequality”. In: *Available at SSRN 4962603*.

De Paola, Maria, Roberto Nisticò, and Vincenzo Scoppa (2021). “Employment protection and fertility decisions: The unintended consequences of the Italian Jobs Act”. In: *Economic Policy* 36.108, pp. 735–773.

Deaton, Angus (1991). “Saving and liquidity constraints”. In: *Econometrica* 59.5, pp. 1221–1248.

Del Boca, Daniela (2002). "The effect of child care and part time opportunities on participation and fertility decisions in Italy". In: *Journal of population economics* 15.3, pp. 549–573.

Del Boca, Daniela, Noemi Oggero, Paola Profeta, and Mariacristina Rossi (2021). *Did COVID-19 Affect the Division of Labor Within the Household?: Evidence from Two Waves of the Pandemic in Italy*. Tech. rep. JSTOR.

Del Boca, Daniela, Noemi Oggero, Paola Profeta, Mariacristina Rossi, et al. (2022). "The impact of COVID-19 on the gender division of housework and childcare: Evidence from two waves of the pandemic in Italy". In: *IZA Journal of Labor Economics* 11.1, pp. 1–20.

Del Boca, Daniela, Silvia Pasqua, and Chiara Pronzato (2004). "Why are fertility and women's employment rates so low in Italy? Lessons from France and the UK". In: *Lessons from France and the UK (August 2004)*.

Del Boca, Daniela and Robert M Sauer (2009). "Life cycle employment and fertility across institutional environments". In: *European Economic Review* 53.3, pp. 274–292.

Del Boca, Daniela and Daniela Vuri (2007). "The mismatch between employment and child care in Italy: the impact of rationing". In: *Journal of Population Economics* 20.4, pp. 805–832.

Di Nicola, Fernando, Edoardo Di Porto, Eliana Franzese, Joanna Kopinska, Stefania Lucchini, et al. (2021). "RAPPORTO ANNUALE INPS, vol XIX". In: *RAPPORTO ANNUALE.... INPS*.

Dingel, JI (2020). *How Many Jobs Can be Done at Home?*

Doepke, Matthias, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt (2023). "The economics of fertility: A new era". In: *Handbook of the Economics of the Family*. Vol. 1. 1. Elsevier, pp. 151–254.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri (2020). "Monopsony in online labor markets". In: *American Economic Review: Insights* 2.1, pp. 33–46.

Dynan, Karen E. (1993). "How Prudent are Consumers?" In: *Journal of Political Economy* 101.6, pp. 1104–1113.

Eckel, Catherine C and Philip J Grossman (2008). "Men, women and risk aversion: Experimental evidence". In: *Handbook of experimental economics results* 1, pp. 1061–1073.

Emanuel, Natalia and Emma Harrington (2021). "Working remotely". In: *Selection, Treatment, and Market Provision of Remote Work, Harvard Job Market Paper* 2.

— (2024). "Working remotely? Selection, treatment, and the market for remote work". In: *American Economic Journal: Applied Economics* 16.4, pp. 528–559.

Eurostat (2024). *Employed persons working from home by professional status — % of total employment (LFSA_EHOMP)*. URL: https://ec.europa.eu/eurostat/databrowser/product/page/LFSA_EHOMP (visited on 10/16/2025).

Faberman, R Jason, Andreas I Mueller, and Ayşegül Şahin (2022). "Has the willingness to work fallen during the covid pandemic?" In: *Labour Economics* 79, p. 102275.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso (2017). "Asset market participation and portfolio choice over the life-cycle". In: *The Journal of Finance* 72.2, pp. 705–750.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri (2020). "Heterogeneity and persistence in returns to wealth". In: *Econometrica* 88.1, pp. 115–170.

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri (2018). "Portfolio choices, firm shocks, and uninsurable wage risk". In: *The Review of Economic Studies* 85.1, pp. 437–474.

Fagereng, Andreas, Martin B Holm, and Gisle J Natvik (2021). "MPC heterogeneity and household balance sheets". In: *American Economic Journal: Macroeconomics* 13.4, pp. 1–54.

Farré, Lídia and Libertad González (2019). "Does paternity leave reduce fertility?" In: *Journal of Public Economics* 172, pp. 52–66.

Flynn, Sean, Andra C Ghent, and Vasudha Nair (2024). "Determinants and consequences of return to office policies". In: *Available at SSRN 4757876*.

Forsythe, Eliza, Lisa B Kahn, Fabian Lange, and David Wiczer (2022). "Where have all the workers gone? Recalls, retirements, and reallocation in the COVID recovery". In: *Labour Economics* 78, p. 102251.

Foucault, Thierry, Marco Pagano, and Ailsa Röell (2023). *Market liquidity: theory, evidence, and policy*. Oxford University Press.

Friedman, Milton (1957). "The permanent income hypothesis". In: *A theory of the consumption function*. Princeton University Press, pp. 20–37.

Georgarakos, Dimitris, Tullio Jappelli, Geoff Kenny, and Luigi Pistaferri (Sept. 2023). *Labor Supply Response to Windfall Gains*. CSEF Working Papers 682. Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy.

Gibbs, Michael, Friederike Mengel, and Christoph Siemroth (2023). "Work from home and productivity: Evidence from personnel and analytics data on information technology professionals". In: *Journal of Political Economy Microeconomics* 1.1, pp. 7–41.

Gittleman, Maury (2022). "The "great resignation" in perspective". In: *Monthly Labor Review*.

Goldin, Claudia (2014). "A grand gender convergence: Its last chapter". In: *American economic review* 104.4, pp. 1091–1119.

— (2021). "Career and family: Women's century-long journey toward equity". In.

Goldstein, Joshua, Wolfgang Lutz, and Maria Rita Testa (2003). "The emergence of sub-replacement family size ideals in Europe". In: *Population research and policy review* 22.5, pp. 479–496.

Gómez, Ángel Luis and Salvatore Lattanzio (2024). "Worker reallocation in Italy and Spain after the COVID-19 pandemic". In: *Latin American Journal of Central Banking* 5.1, p. 100105.

Gottlieb, Charles, Jan Grobovšek, Markus Poschke, and Fernando Saltiel (2021). "Working from home in developing countries". In: *European Economic Review* 133, p. 103679.

Graves, Sebastian (2025). "Does Unemployment Risk Affect Business Cycle Dynamics?" In: *American Economic Journal: Macroeconomics* 17.2, 65–100.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese (1996). "Income Risk, Borrowing Constraints, and Portfolio Choice". In: *The American Economic Review* 86.1, pp. 158–172.

Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi (2005). "Insurance within the firm". In: *Journal of Political Economy* 113.5, pp. 1054–1087.

Harknett, Kristen, Francesco C Billari, and Carla Medalia (2014). "Do family support environments influence fertility? Evidence from 20 European countries". In: *European Journal of Population* 30.1, pp. 1–33.

Hensvik, Lena, Thomas Le Barbanchon, and Roland Rathelot (2020). "Which jobs are done from home? Evidence from the American Time Use Survey". In.

— (2021). "Job search during the COVID-19 crisis". In: *Journal of Public Economics* 194, p. 104349.

Hiriscau, Andra (2024). "The effect of paid maternity leave on fertility and mothers' labor force participation". In: *Journal of Labor Research* 45.3, pp. 350–384.

Hotz, V Joseph, Jacob Alex Klerman, and Robert J Willis (1997). "The economics of fertility in developed countries". In: *Handbook of population and family economics* 1, pp. 275–347.

Hupkau, Claudia and Barbara Petrongolo (2020). "Work, care and gender during the Covid-19 crisis". In: *Fiscal studies* 41.3, pp. 623–651.

INPS (2025). *XXIV Rapporto annuale*. Accesso 6 ottobre 2025. URL: https://www.inps.it/it/it/dati-e-bilanci/rapporti-annuali/xxiv-rapporto-annuale.html (visited on 10/06/2025).

Jappelli, Tullio and Luigi Pistaferri (2000). "Using subjective income expectations to test for excess sensitivity of consumption to predicted income growth". In: *European Economic Review* 44.2, pp. 337–358.

— (Nov. 2017). *The Economics of Consumption: Theory and Evidence*. Oxford University Press.

— (2020). "Reported MPC and unobserved heterogeneity". In: *American Economic Journal: Economic Policy* 12.4, pp. 275–297.

— (2025). "Permanent income shocks, target wealth, and the wealth gap". In: *American Economic Journal: Macroeconomics* 17.1, pp. 102–125.

Kahn, Matthew E and Joseph Tracy (2023). *A Human Capital Theory of Who Escapes the Grasp of the Local Monopsonist*. Tech. rep. National Bureau of Economic Research.

Kaplan, Greg, Giovanni L Violante, and Justin Weidner (2014). *The wealthy hand-to-mouth*. Tech. rep. National Bureau of Economic Research.

Ketter, Laura, Todd Morris, and Lizi Yu (2025). "A new equilibrium: COVID-19 lockdowns and WFH persistence". In: *arXiv preprint arXiv:2506.16671*.

Kimball, Miles S. (1990). "Precautionary Saving in the Small and in the Large". In: *Econometrica* 58.1, pp. 53–73.

Kolsrud, Jonas, Camille Landais, Daniel Reck, and Johannes Spinnewijn (2024). "Retirement consumption and pension design". In: *American Economic Review* 114.1, pp. 89–133.

Lalive, Rafael and Josef Zweimüller (2009). "How does parental leave affect fertility and return to work? Evidence from two natural experiments". In: *The Quarterly Journal of Economics* 124.3, pp. 1363–1402.

Langella, Monica and Alan Manning (2021). "Marshall lecture 2020: The measure of monopsony". In: *Journal of the European Economic Association* 19.6, pp. 2929–2957.

Lavetti, Kurt (2023). "Compensating wage differentials in labor markets: Empirical challenges and applications". In: *Journal of Economic Perspectives* 37.3, pp. 189–212.

Lee, Dain, Jinhyeok Park, and Yongseok Shin (2023). *Where are the workers? From great resignation to quiet quitting*. Tech. rep. National Bureau of Economic Research.

Lusardi, Annamaria (2008). *Financial literacy: an essential tool for informed consumer choice?* Tech. rep. National Bureau of Economic Research.

Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S. Mitchell (2017). "Optimal Financial Knowledge and Wealth Inequality". In: *Journal of Political Economy* 125.2, pp. 431–477. eprint: https://doi.org/10.1086/690950.

Manning, Alan (2003). "The real thin theory: monopsony in modern labour markets". In: *Labour economics* 10.2, pp. 105–131.

— (2021). "Monopsony in labor markets: A review". In: *ILR Review* 74.1, pp. 3–26.

Mas, Alexandre (2025). "Non-Wage Amenities". In.

Mas, Alexandre and Amanda Pallais (2017). "Valuing alternative work arrangements". In: *American Economic Review* 107.12, pp. 3722–3759.

McCrary, Justin (2008). "Manipulation of the running variable in the regression discontinuity design: A density test". In: *Journal of Econometrics* 142.2. The regression discontinuity design: Theory and applications, pp. 698–714.

Merton, Robert C (1969). "Lifetime portfolio selection under uncertainty: The continuous-time case". In: *The review of Economics and Statistics*, pp. 247–257.

— (1971). "Optimum consumption and portfolio rules in a continuous-time model". In: *Journal of Economic Theory* 3.4, pp. 373–413.

Mongey, Simon, Laura Pilossoph, and Alexander Weinberg (2021). "Which workers bear the burden of social distancing?" In: *The Journal of Economic Inequality* 19.3, pp. 509–526.

Mooi-Reci, Irma, Trong-Anh Trinh, Esperanza Vera-Toscano, and Mark Wooden (2023). "The impact of lockdowns during the COVID-19 pandemic on fertility intentions". In: *Economics & Human Biology* 48, p. 101214.

Morikawa, Masayuki (2022). "Work-from-home productivity during the COVID-19 pandemic: Evidence from Japan". In: *Economic inquiry* 60.2, pp. 508–527.

Naidu, Suresh and Michael Carr (2022). "If You Don't Like Your Job, Can You Always Quit? Pervasive Monopsony Power and Freedom". In: *Journal of Law and Political Economy, 3 (1)*.

OECD (2023). "Pensions at a Glance 2023: OECD and G20 Indicators, OECD Publishing, Paris". In.

OECD (2025). *OECD Family Database*. Accesso 14 ottobre 2025. URL: https://www.oecd.org/en/data/datasets/oecd-family-database.html (visited on 10/14/2025).

Olafsson, Arna and Michaela Pagel (2024). "Retirement puzzles: New evidence from personal finances". In: *Journal of Public Economics* 234, p. 105103.

Ramani, Arjun and Nicholas Bloom (2021). *The Donut effect of COVID-19 on cities*. Tech. rep. National Bureau of Economic Research.

Roback, Jennifer (1982). "Wages, rents, and the quality of life". In: *Journal of political Economy* 90.6, pp. 1257–1278.

Rosen, Sherwin (1986). "The theory of equalizing differences". In: *Handbook of labor economics* 1, pp. 641–692.

Silva, Rui, Ramin Baghai, and Margarida Soares (2025). "The Impact of IPOs on Workers". In: *Available at SSRN*.

Sockin, Jason (2022). "Show me the amenity: Are higher-paying firms better all around?" In.

Thorstensson, Esra (2020). *The Influence of Working from Home on Employees' Productivity: Comparative document analysis between the years 2000 and 2019-2020*.

Van Rooij, Maarten, Annamaria Lusardi, and Rob Alessie (2011). "Financial literacy and stock market participation". In: *Journal of Financial Economics* 101.2, pp. 449–472.

Vignoli, Daniele, Alessandra Minello, Giacomo Bazzani, Camilla Matera, and Chiara Rapallini (2022). "Narratives of the future affect fertility: Evidence from a laboratory experiment". In: *European Journal of Population* 38.1, pp. 93–124.

Webber, Douglas A (2015). "Firm market power and the earnings distribution". In: *Labour Economics* 35, pp. 123–134.

Zarate, Pablo, Mathias Dolls, Steven J Davis, Nicholas Bloom, Jose Maria Barrero, and Cevat Giray Aksoy (2024). *Why Does Working from Home Vary Across Countries and People?* Tech. rep. National Bureau of Economic Research.

Zurla, Valeria (2023). "How Should We Design Parental Leave Policies? Evidence from Two Reforms in Italy". In: *Evidence from Two Reforms in Italy (May 15, 2023)*.