SAPIENZA
UNIVERSITÀ DI ROMA

And yet it moves?
Insights into income and geographical mobility in Italy

Department of Economics and Law
PhD in Economics (XXXV cycle)

Francesca Subioli
ID number 1658114

Advisor
Prof. Michele Raitano

Academic Year 2022/2023

***And yet it moves?* Insights into income and geographical mobility in Italy**

PhD thesis. Sapienza University of Rome

This thesis has been typeset by LaTeX and the Sapthesis class.

Author's email: francesca.subioli@uniroma1.it; francesca.subioli@gmail.com

To my darling,
tireless supporter

Our desire is not that others might be relieved while you are hard pressed, but that there might be equality.
At the present time your plenty will supply what they need, so that in turn their plenty will supply what you need.
The goal is equality, as it is written: «The one who gathered much did not have too much,
and the one who gathered little did not have too little».

2 Corinthians 8, 13-15

Non si tratta infatti di mettere in difficoltà voi per sollevare gli altri, ma che vi sia uguaglianza.
Per il momento la vostra abbondanza supplisca alla loro indigenza,
perché anche la loro abbondanza supplisca alla vostra indigenza, e vi sia uguaglianza, come sta scritto:
«Colui che raccolse molto non abbondò e colui che raccolse poco non ebbe di meno».

2 Corinzi 8, 13-15

Contents

Introduction

This thesis stems from research work I carried out in various universities and institutions to provide new perspectives on the analysis of mobility in different research contexts. The first two chapters deal with *earnings* mobility and its association with inequality (Chapter 1) and polarization (Chapter 2), while the third chapter is concerned with *geographical* mobility in response to labour market shocks (Chapter 3). In all three cases, the focus of the analysis is on the Italian labour market. However, while the first and last works are applied in nature, the second one is a theoretical paper whose application is instrumental to understanding the theory and demonstrating its empirical relevance. I briefly outline below the contents of the three chapters.

Chapter 1 – "Differences set in stone: evidence on the inequality-mobility trade off in Italy" – is designed to give emphasis to the concept of earnings mobility in the assessment of economic inequality. The rationale is that observing individual income dynamics is crucial to assess the characteristics of the process shaping income inequality and its consequences on individual and social well-being. Given the level of cross-sectional inequality, a mobile society faces different challenges than one where people are stuck in their income positions for their whole life or see their income stagnate. Moreover, people prefer a stable income stream to a fluctuating one, and the policy concern should deal with the level of income as with its dynamics. Using Italy as our case study, we characterise the long-run evolution of intragenerational mobility in the last forty years and find evidence of a trade-off between income inequality and 'good' mobility, and complementarity with the worst notions of mobility related to income instability. Exploiting individual-level estimates of good and bad mobility, we also uncover patterns of *unequal mobility* (OECD, 2018) – the concentration of low upward mobility and frequent fluctuations among the most vulnerable groups.

Chapter 2 – "Inter-temporal income polarization" – proposes an extension of Esteban and Ray (1994) income polarization index to incorporate the time dimension. Income polarization captures the extent to which an income distribution concentrates around two or more income levels. Polarization measurement is typically rationalized as measuring potential conflict in a society when people feel alienated from one another when distant in income but feel identified with other people of similar income levels. We introduce time in this model following the idea that the two key ingredients of polarization – *alienation* and *identification* – may have fewer implications for potential conflict if individual incomes vary over time and feelings of alienation or identification have therefore limited time to form and consolidate.

Accordingly, the second chapter proposes an *inter-temporal* income polarization measure using panel data, in which memory parameters allow past income differences to determine the degree of alienation and identification in a society's income distribution. This leads to measures of income polarization that are sensitive to the history of interpersonal income proximity and distances in income trajectories. The empirical relevance of this longitudinal perspective is demonstrated through an application to Italian data.

Chapter 3 – "Labour market dynamics and geographical reallocation" – deals with a different notion of mobility, namely migration. Understanding the responsiveness of the geographical allocation of workers to local labour market dynamics is a first-order issue in the economic literature, being migration a major mechanism to absorb labour demand variations through people moving across regions (Blanchard and Katz, 1992) or changing their commuting behaviour (Monte et al., 2018) in response to employment opportunities. The work is based on a unique source of administrative data on the universe of labour market flows for Italy and exploits it to study how local labour demand shocks affect internal migration through an instrumental variable approach.

Besides providing a new and comprehensive picture of job and migration flows in Italy from 2010 to 2018, the estimates reveal that job creation has a strong effect on the in-migration rate, whereas job destruction has a much milder effect on the out-migration rate, the latter being a less responsive adjustment margin. Crucially, it seems that the large responsiveness of in-migration does not work through an increase in the number of relocating workers, but rather through changes in their destination alternatives. Moreover, the effects of labour market shocks on geographical mobility vary by distance: the positive effect of job creation on in-migration flows has a much larger geographical reach than that of job destruction on out-migration, which instead creates out-migration flows that are locally concentrated.

As a final note on this thesis, the first and second chapters are enriched with details of (the preliminary version of) the programs developed in Stata to make the proposed methodologies easily accessible to other researchers. The intragenerational mobility indices and graphical tools described in Chapter 1 can be reproduced for panel data through the program `intramob` (Appendix A.2), while the program `itempolar` (Appendix B.2) can be used to measure inter-temporal income polarization as in Chapter 2.

I would like to thank the two referees Elena Bárcena-Martín and Conchita D'Ambrosio for their careful reading and valuable suggestions on the first version of this thesis.

Chapter 1

Differences set in stone: evidence on the inequality-mobility trade off in italy

JEL Codes: D31, D63.
Keywords: Earnings inequality, Great Gatsby curve, Intragenerational mobility, Earnings dynamics, Unequal mobility, Italy.

Notes and Acknowledgements

This chapter is an extension of joint work with my advisor Michele Raitano (Sapienza University of Rome), whom I warmly thank for his supportive and inspiring supervision. I also thank the Giacomo Brodolini Foundation for allowing us to use the AD-SILC dataset for the project. This chapter benefited from comments and suggestions on previous versions from the members of the INEQ group at Sapienza University of Rome and of the PRIN RISIng 'Rising inequality and Social Insecurity of the middle classes', and from participants at the 9th Meeting of the Society for the Study of Economic Inequality (ECINEQ), the XXXIII Annual Conference of the Italian Society of Public Economics (SIEP), the 62nd Annual Conference of the Italian Economic Association (SIE), and the Workshop on Wealth Inequality and Intergenerational Mobility organized by Vienna University of Economics and Business. I am particularly grateful to Philippe Van Kerm for some comments that greatly improved this work. Appendix A.2 contains details (ado and help files) of a Stata program I developed to perform the analyses in this chapter and which I would like to share with other researchers interested in intragenerational mobility.

1.1 Introduction

In recent years, the economic literature and the policy debate have been increasingly concerned with the rise in income inequality experienced since the last decades of the 20th century in most high-income countries (see, among others, OECD (2008) and OECD (2011)). Despite inequality is shaped by complex processes acting through various mechanisms and is influenced by several possible determinants (Atkinson, 2016), shared wisdom argues that these trends have been mainly due to processes acting in the markets and, specifically, the bulk of the increase in inequality seems attributable to the rise in labour income dispersion (Salverda et al., 2014; Hoffmann et al., 2020). However, also due to the scarcity of accurate and long longitudinal data, most analyses on trends in earnings inequality provide pictures of what happened at various points in time (typically years), focusing on 'snapshots' of the income distribution. The usual focus is on cross-sectional inequality – across people at a point in time –, neglecting what happens to individuals from one period to the next (Burkhauser and Couch, 2009).

Whatever the magnitude of period inequality, observing individual income dynamics is crucial to assess the characteristics of the process shaping inequality and its consequences on individual and social well-being for mainly two reasons. The first one is related to social welfare: as pointed out by Jenkins (2011) and OECD (2018), a society with a certain level of income inequality where individuals change their positions in the income ladder faces different challenges than one with the same (or a lower) level of inequality where individuals are stuck in their income positions during their whole life. The second reason for tracking individual careers concerns people themselves: in general, individuals are concerned not only with the average income they receive over a certain period, but also with its *pattern* over time. Since people prefer a stable income stream to a fluctuating one, having a stable stream may be considered welfare-enhancing per se (Shorrocks, 1978). Moreover, there is well-known evidence from psychology and behavioural economics that people are much more averse to losing what they already have than what they could potentially gain (Kahneman and Tversky, 1984; Kahneman et al., 1991). If this is the case, the policy concern should deal with the dynamics of income as well as its level. Therefore, we believe that the assessment of income inequality from a welfare perspective should be complemented with information on the underlying mobility processes.

While the empirical association between income inequality and intergenerational persistence – i.e. absence of mobility across subsequent generations – has been widely studied starting from the work of Corak (2013), there has been less attention to the *intragenerational* persistence – absence of mobility within the same generation –, possibly also due to data requirements. We are aware of some works investigating through longitudinal data whether a high level of inequality is mitigated by a similarly high level of income mobility through cross-country comparisons.[1] Taking as reference the United States as a high-inequality country, comparisons between the levels of mobility in the US and in Europe reveal that the differences in income mobility are not so pronounced, not even with respect to the Nordic European countries (Gangl, 2005). A very recent work (Guvenen et al., 2022) covering a wide range of countries all over the world, finds a positive and weak correlation (0.35) between country-level inequality and persistence in income positions after five years.[2]

Besides the aggregate correlation between inequality and mobility, a further major issue in the evaluation of the income movements underlying inequality is related to the assessment of who are the *winners* and *losers* of income mobility:

> «[...] 'unequal mobility' can occur when unpredictable income changes combine with low levels of long-term (upward) income mobility and when this concerns mostly the most vulnerable population groups.» (OECD (2018), p. 65).

Indeed, it may be the case that only part of the population benefits from a desirable notion of mobility – upward and smooth income growth –, while another part suffers its more negative aspects which take the form of income instability.

With this framework in mind, we use Italy as our case study – a country characterised by a steep rise in labour income inequality in the last decades – and characterise long-run patterns of inequality and mobility across several cohorts of workers with a twofold aim. First, to understand whether the well-proved increase in earnings inequality has been compensated by higher mobility between workers, or has been due to widening persistent differences. Second, we go into details of income changes and distinguish 'good' mobility – i.e., upward and predictable changes – from mere volatility – i.e. frequent and unpredictable fluctuations – to assess *who* is concerned and whether there is a vulnerability problem related to income dynamics that policymakers should be concerned about.

In measuring the dynamics underlying income inequality changes between two points in time, the empirical literature has encountered some substantial challenges. First, no univocal methods and measures have emerged, given the complex and multifaceted nature of the concept of income mobility itself. As reviewed in Fields and Ok (1999), Jenkins (2011) and Jäntti and Jenkins (2015), the conceptualization of income mobility depends on the reference period – mobility from when to when? –, the reference group – mobility relative to whom? –, and the reference concept of income – mobility of what? –. Such

[1] See Burkhauser and Couch (2009) for a review of these works. More recent works are Alves and Martins (2012) comparing the US and Nordic countries, Aaberge and Mogstad (2014) for European countries, and OECD (2018) for OECD countries.

[2] We report in Figure A.1 two intragenerational Great Gatsby curves from Gangl (2005) and Guvenen et al. (2022). The expression 'Great Gatsby curve' is due to a speech by the economist Alan Krueger – "The Rise and Consequences of Inequality" – on January 12th, 2012 at the Center for American Progress. It is the graphical representation of the positive relationship between cross-sectional inequality (measured by the Gini index) and intergenerational earnings persistence (measured by the intergenerational income elasticity) across countries.

complexity naturally led to a proliferation of conceptualizations, measurement tools and indices, each of which is useful for isolating a specific facet of income mobility. We believe that the best approach to this complexity is to take into account as many different aspects of income mobility as possible, rather than choosing only one. This can return a comprehensive picture of the dynamics underlying inequality, not tied to the type of measure chosen. In this respect, our approach is in the same spirit as Jenkins (2011). We may also be surprised to find different pictures depending on the specific aspect we look at.

As said, most of the mobility indices, as well as the inequality ones, require setting a *reference group* to compare income values and positions. With longitudinal data, two strategies are possible to compare different generations: a *time approach*, comparing people's income at any age in a given calendar year, and a *cohort approach*, fixing age and comparing people belonging to the same cohort of birth regardless the calendar year. Inequality and mobility measures are heavily influenced by the life-cycle features of earnings: even when considering only individual income from labour, leaving aside the impact of demographic events like a marriage or the birth of a child, a typical income trajectory should rise up to a certain age due to the accumulation of experience and then decrease with retirement. Therefore, measures of inequality based on a calendar year approach can be affected by changes in the demographic structure: inequality may increase from one year to the next either because there is more dispersion in earnings at a given age, or because the age composition of the population has changed.

For the purpose of this work, we believe that a cohort approach is more suited: we compare people within their own generation, assuming that their reference group are those having a similar age in the same years – their *peers*. This means that within their group people share the macroeconomic conditions that are specific to their generation at a given life-cycle phase. Our goal with this setting is to compare the inequality and mobility prospects of different generations of workers: as employment and earnings prospects may change across cohorts, comparing the within-generation inequality and mobility values is informative with respect to *intergenerational fairness* concerns (Raitano et al., 2021). In fact, we already know from previous studies for Italy that the progressive 'dualization' of the labour market that started in the mid-80s – imposing worse contractual arrangements to new entrants while maintaining secure conditions for incumbent workers – has led to a serious gap in the economic well-being of different generations of workers, especially in terms of career prospects.[3]

A second crucial challenge for mobility measurement is related to data. By its very nature, mobility depends on time; therefore, the choice of the concept of mobility adopted is also

[3] For empirical evidence on the consequences of labour market flexibilization for new entrants in Italy see, among others, Rosolia and Torrini (2007), Barbieri and Scherer (2009), Naticchioni et al. (2016), Rosolia and Torrini (2016), Raitano and Fana (2019), Hoffmann et al. (2022). For a detailed discussion of the reforms that shaped this 'dual' labour market, see Boeri and Garibaldi (2007) and Hoffmann et al. (2022).

driven by the time coverage of the available data, and their capacity to follow individuals over time. A typical problem from this point of view is panel attrition, often characterizing survey data, but also the simple fact of observing individuals at a distance of time. To address this issue, we rely on a matched survey-administrative dataset for the Italian private sector covering a long time span (1975-2018) and following the entire careers of workers born in very different economic contexts. Some peculiar characteristics of this dataset, detailed in Section 1.3, make it particularly suited for the purpose of this work with respect to other survey and administrative sources available for Italy.

We contribute to the literature on income inequality and income mobility by providing the first cross-cohort intragenerational Great Gatsby curves for a single country, and providing a strategy to study the individual-level vulnerability due to income dynamics. In Section 1.2, we detail how we measure intragenerational inequality and mobility at the individual and aggregate level, distinguishing notions of 'good' and 'bad' mobility. Then, we provide information about the data in Section 1.3. Section 1.4 presents and discusses the results of the analysis: first, we describe the evolution across subsequent cohorts of several indices of intragenerational inequality and mobility, also focusing on non-linearities along the income distribution (Section 1.4.1 and 1.4.2). Then, we discuss in Section 1.4.3 our estimates of the correlation between inequality and mobility levels, and present the underlying intragenerational Great Gatsby curves. In Section 1.4.4, we show and discuss our findings on the phenomenon of unequal mobility. Finally, Section 1.5 provides a heterogeneity analysis for the main results by gender, level of education and macro area of work, and Section 1.6 concludes.

1.2 Methodology

1.2.1 General setting

To get empirical estimates of the inequality-mobility trade off in the intragenerational context, our strategy consists of measuring for separate cohorts of Italian workers the aggregate inequality and mobility levels, and then simply estimating their correlation. A positive correlation would be a sign of complementarity between inequality and mobility and, therefore, of a possible compensation between the two. A negative correlation, on the contrary, would signal a trade-off: greater inequality would also come with the burden of less income mobility, rather than signalling a more dynamic society. Taking inspiration from the intergenerational Great Gatsby curve, we also employ a scatter plot to visualize the relationship; however, our units are not different countries at the same point in time but rather different birth cohorts of the same country.

Then, to answer our second research question related to the existence – and the pattern over time – of 'unequal mobility', we move to the micro level and exploit the individual-level estimates of earnings mobility: by measuring individual mobility in a way that separates unpredictable income changes from long-term upward mobility following Nichols (2008) and Nichols and Rehm (2014), we study the combination between the two and their relation to lifetime income. This approach allows for a transparent and intuitive detection of vulnerabilities related to wage dynamics.

Following a cohort approach, we fix a common age window for all workers; we need some assumptions about which is the best moment for observing one's career and getting the best proxy of the lifetime earnings experience.[4] We fix the age at 35-45, a long and central phase of the career when we assume formal education is completed and retirement is still a long way off. However, we are aware that this age group can have very different implications for men and women, as it is a fertile period when maternity leaves and childcare may affect women's careers more than men's.

Inequality and mobility measurement We include in the baseline analysis zero earnings to take into account periods of non-employment that may have a strong impact on income dynamics. Most of the results are compared with the case of only positive earnings to infer how much periods of non-employment affect the inequality and mobility estimates. Moreover, we adopt a personal-level perspective rather than a household-level one not simply because of data limitations, but also because we want to track personal income experiences gross of behavioural choices related to family formation. Importantly, our analysis includes both women and men.

For measuring inequality, we use the Generalized Entropy index of degree two for the reasons we will explain in Section 1.2.2 and distinguish *overall* – across people and time –, *permanent* – based on long-term income experience –, and *average cross-sectional* inequality – the mean of period-by-period snapshot inequality. The more mobility is in place, the more permanent inequality departs from the other two measures.

Regarding mobility, we rely on a vast set of indices for two reasons. First, mobility is a multifaceted phenomenon, and different measures of it are not alternatives but complementary. Second, the direction of the association between a specific notion of mobility and inequality is not a priori determined. Therefore, we let it vary according to the concept of mobility used in each case.

The indices of mobility are presented in Section 1.2.2 and 1.2.2 divided into *bi-periodical* –

[4]This issue is usually a concern in the literature on intergenerational mobility because, when analysing the effect of parents' characteristics on children's outcome, it is crucial not to disregard at which stage of life parents and children are observed. For example, Haider and Solon (2006) and Böhlmark and Lindquist (2006), respectively for the US and Sweden, find evidence that the difference between current and lifetime earnings for men is minimized around age 35. Conversely, a simple rule does not emerge for women, who display more variety in their life-cycle income patterns especially because of maternity periods. Nybom and Stuhler (2016) warn that age-earnings profiles may be worker, country or cohort-specific even for male workers, so the choice of the same point in age for every worker may be misleading.

based on a comparison between an origin and a destination income – and *dynamics* measures – based on the income movements in each period between origin and destination points (Jenkins, 2011). We also explain for some indices the graphical tools we employ to visualize mobility patterns. Importantly, we attempt to classify each of the measures as 'good' or 'bad' for the society and the individual, without using formal welfare evaluation methods but simply through reasonable arguments.[5] This classification is crucial for interpreting the inequality-mobility trade off and drawing conclusions in terms of intergenerational fairness.

Measuring 'unequal mobility' The aim of this part of the work is to see whether, and how much, «[...] unpredictable income changes combine with low levels of long-term (upward) income mobility and [...] this concerns mostly the most vulnerable population groups.»[6] We follow two steps: first, we choose the three measures of, respectively, *unpredictable income changes*, *low levels of long-term (upward) income mobility*, and *vulnerability*. To separate unpredictable income changes and long-term (upward) income mobility, we rely on the 'income risk decomposition' proposed by Nichols (2008) and described in detail in Section 1.2.2. It proxies long-term predictable mobility through the steepness of an individual linear trend, and unpredictable income changes through the intensity of deviations from that trend. The third element of the framework is vulnerability, defined in Calvo and Dercon (2013) as the extent of a threat of poverty. The concept refers to the expectation of dropping below a poverty threshold without being able to ensure this risk (Ceriani, 2018; Calvo, 2018). An individual with no or restricted access to social insurance and credit is defined vulnerable if her current consumption is close to the poverty line. In our setting, being absent measures of consumption and possible buffers, we simplify the framework by measuring vulnerability as low 'permanent income', where permanent income is average earnings in the time window observed. If the credit market is perfect – i.e. consumption smoothing is possible at any time –, the permanent income is a measure of the long-term economic well-being of people: it summarises their economic possibilities and accounts for the fact that the ability to save and borrow to address income shocks is strictly linked to the overall income potential. A low permanent income status is therefore assumed to be an indicator of a long-lasting proximity to the poverty line.

Once we have the estimates of good mobility, bad mobility, and permanent income, we employ a *heat map* graphical tool to study their correlation. The hypothesis of unequal mobility would be verified if we find good mobility to be negatively correlated with bad mobility but positively related to permanent income: people enjoying better overall economic conditions would also benefit from smooth and positive income growth and be

[5]For a social-welfare evaluation approach incorporating the insecurity aversion of individuals, see Gottschalk and Spolaore (2002) and Jäntti et al. (2014).

[6]OECD (2018), p. 65.

protected by unexpected income shocks, while the reverse would be true at the bottom of the permanent income distribution. The heat map will allow us to visualise the permanent income distribution as a function of combinations of good and bad mobility levels.

1.2.2 Measuring intragenerational income mobility

Bi-periodical indices

Bi-periodical mobility indices are based on the comparison between an *origin* and a *destination* income distribution computed in two different periods, the second being later than the first. As stated above, in our setting each worker is observed in the age window 35-45: therefore, we choose as origin income the earnings averaged from age 35 to 37, and as destination income the earnings averaged from age 43 to 45.[7] Once origin and destination incomes are defined, *relative* or *positional* mobility indices compare the two income distributions measuring changes in relative positions, while *absolute* mobility indices compare one's own income value at destination with that at origin, regardless of relative position, and then aggregate such changes through a simple average.[8]

Positional mobility As a first bi-periodical index, we use a modification of the Hart (1976) mobility index employing a Spearman's rank correlation coefficient ρ with ranks normalized in the interval $[0, 1]$:

$$\text{Rank mobility} = 1 - \rho_n = 1 - cov(r_o, r_d) \tag{1.1}$$

r_o and r_d being the origin and destination normalized ranks.[9] With this procedure, the origin and destination income distributions are forced to be standard uniforms; therefore, the beta coefficient from a linear regression of the rank of destination on the rank of origin is the simple covariance between the two. Being based on income ranks rather than on income values, ρ_n measures how much the rank of destination increases with the rank of origin: a correlation of -1 indicates perfect rank reversal, one of 0 indicates no monotonic relation between the two distributions – i.e. origin independence –, and a correlation of 1 indicates complete dependence, that is no rank mobility.

To look graphically at this notion of mobility, we plot the line fitted through the scatter plot of r_d on r_o together with the 45 degrees line that is the place of complete immobility for comparison. Moreover, we look at non-linearities by plotting the average rank of origin and

[7] Averaging income in a short interval to slightly smooth it is a standard procedure to build mobility measures mitigating the effect of year or age-specific shocks.

[8] For a detailed explanation of income mobility indices, see Jäntti and Jenkins (2015).

[9] The rank is obtained by ordering people from the lowest to the highest level of income and normalized using the formula $\frac{rank-1}{max(rank)-1}$. We order the zeros and equal values of income by adding random numbers from a uniform distribution.

destination inside 10 equal-sized bins for both variables; it may be the case that different parts of the origin income distribution are more mobile than others, so that the average hides important differences depending on the starting point.

Again based on the comparison between the normalized rank at origin and that at destination, we use a measure of 'average jump' following the idea of Bartholomew (1968). Separating rank movements *to the right* and *to the left*, we define the average jump up as the mean rank difference for those improving their position ($\sum_{i:r_d>r_o}(r_d - r_o)$) and the average jump down as the same measure but for those who end up in a lower rank ($\sum_{i:r_d<r_o}(r_d - r_o)$). Since normalized ranks lie in the interval [0,1], the jump is the average fraction of the income distribution climbed up or passed when falling down, giving a proxy of the 'distance' covered in the process of positional mobility and allowing to inspect any asymmetry in it.

Another possibility of measuring relative mobility is by comparing the two positions in terms of income *quantiles* and computing the aggregate probability to change quantile through a *transition matrix* – i.e. looking at the share of people reaching a certain destination quantile given the origin position. Let $i = 1, ..., q$ be the quantile of origin income and $j = 1, ..., q$ be the quantile of destination income; then, n_{ij} is the number of people moving from quantile i to quintile j, and $n_{i.}$ is the number of people starting from quantile i whatever their destination quantile. We compute for each cohort the probability of reaching a higher quantile as $\sum_{j>i} \frac{n_{ij}}{n_{i.}}$, of falling into a lower quantile as $\sum_{j<i} \frac{n_{ij}}{n_{i.}}$, of exit from the bottom quantile as $\sum_{j\neq 1} \frac{n_{1j}}{n_{1.}}$, and of falling from the top quantile as $\sum_{j\neq q} \frac{n_{qj}}{n_{q.}}$.

Absolute mobility Measures of absolute mobility do not consider income positions, but rather income value changes from origin to destination. The typical index of absolute mobility is the average income growth in the population (Fields and Ok, 1999). Let y_o be the origin income and y_d be the destination income. While Fields and Ok (1999) use the log difference to measure individual income growth, we use $(y_d/y_o) - 1$ to include zero earnings.[10] When measured directly on income, the growth rate may assume very high values and some outliers may heavily influence the index if the aggregation rule is the simple average. Therefore, we adopt two alternative strategies to address this issue: as a first solution, we compute the median rather than the mean income growth across workers. Second, we keep the average aggregation rule but using a bounded underlying growth rate proposed in Davis and Haltiwanger (1992), $g_{DH} = (y_d - y_o)/(\frac{y_d+y_o}{2})$. This growth rate is symmetric around zero and lies in the interval [-2; 2]. [11]

[10]There are cases in which the growth rate is not defined being the denominator $y_o = 0$. We assign a growth rate of 0 if $y_o = y_d = 0$, and a growth rate of 1 if $y_o = 0$ and $y_d > 0$.

[11]It is monotonically related to the traditional growth rate g, and the relation is $g = \frac{2g_{DH}}{2-g_{DH}}$ (Davis and Haltiwanger, 1992). The two measures are approximately equal for small values.

Indices of dynamics

Income risk decomposition Among the possible mobility indices of dynamics, summarizing individual income movements in a population, we choose the method proposed in Nichols (2008) and applied in Nichols (2010), Nichols and Rehm (2014), Latner (2018), and OECD (2018). Called 'income risk decomposition', this method separates permanent inequality, mobility and volatility through the decomposition of an inequality index with longitudinal data. Overall inequality – across people and time – is measured through a subgroup decomposable index (Shorrocks, 1984), and individuals themselves are the population subgroups.

Nichols (2008) uses the Generalized Entropy (GE from now on) index with parameter $\alpha = 2$ because it has some desirable properties: (i) the family of the GE indices share with the more classical Gini coefficient the Lorenz consistency property; (ii) GE indices also allow additive subgroup decomposability; (ii) the GE index of degree 2 does not require log transformation of income, allowing the inclusion of zeros.[12]. Being Lorenz consistent, the GE2 index is scale invariant, so it allows comparison of different countries or the same country in different periods by removing the effect of the overall income level from the measure of inequality.

Let $i = 1, ..., L$ workers be followed for $t = 1, ..., T$ periods, for a total of $N = LT$ observations. Applying a decomposition by 'people subgroups', the *between-group* inequality component measures permanent inequality across workers — i.e. inequality in average incomes over the observed time window –, while the *within-group* inequality component measures average personal inequality over time, which is a combination of mobility risk and volatility. Formally, let y_{it} be the annual real gross earnings of worker i at time t, $\bar{y}$ be the average annual earnings among all $N = LT$ observations in the time window T for the L workers in the sample, and $\bar{y}_i$ be the average earnings of worker i in the window T – i.e. her permanent earnings. Then, the overall inequality across people and time in the window T can be decomposed as in Equation (1.2):

$$GE(2) = \frac{1}{2\bar{y}^2}\left[\frac{1}{L}\sum_{i=1}^{L}\frac{1}{T}\sum_{t=1}^{T}(y_{it}-\bar{y})^2\right] = \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{L}\sum_{i=1}^{L}\frac{1}{T}\sum_{t=1}^{T}(\bar{y}_i-\bar{y})^2\right]}_{\text{Between-workers inequality (B)}} + \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{L}\sum_{i=1}^{L}\frac{1}{T}\sum_{t=1}^{T}(y_{it}-\bar{y}_i)^2\right]}_{\text{Within-worker inequality (W)}} \tag{1.2}$$

The between-workers inequality is the variance of individual-level average income $\bar{y}_i$, divided by twice squared average income $\bar{y}$. It corresponds to the definition of long-term inequality as the dispersion in permanent incomes. On the other hand, the within-worker inequality is the average across workers of the individual-level variance of income over

[12] When the parameter α is neither 0 (Mean Log Deviation) nor 1 (Theil Index), all the GE indices can be computed using income without log transformation.

time, again divided by twice the squared mean income. We do not need to weigh the personal variances since all the individuals are observed for the same number of years in this formulation.

As a further and crucial step, the numerator of the within-worker inequality component can be further decomposed into what Nichols (2008) calls 'mobility risk' and 'volatility'. In practice, the individual income process is seen as made of three components: (i) the average, permanent, income; (ii) a linear trend summarizing smooth and directional income growth; (iii) volatility around the income trend. Equation (1.3) models this process:

$$y_{it} = \alpha_i + \beta_i t + \epsilon_{it} \tag{1.3}$$

If time t is centred at zero, α_i coincides with the permanent income $\bar{y}_i$, and the income trend $\beta_i t$ is demeaned – i.e. has mean zero. The choice of a linear trend, which may be controversial when considering the entire life-cycle income pattern that is usually modelled as convex, can be considered particularly suitable when looking at incomes in a medium-short age window sufficiently far from retirement. Moreover, a linear trend is theoretically preferable because of its smooth pattern: if we believe that 'good' mobility for the individual is a predictable income path, directional and not affected by relevant and frequent fluctuations, a linear pattern seems to be the most reasonable and transparent choice.

As discussed in Nichols (2008), the length of the period T must be at least three (two observations to estimate a linear trend, and the third to allow deviation from it). However, the variance of the idiosyncratic error term used to characterize volatility will tend to be dramatically understated for small lengths. We decide to adopt here a wide range T=11 since our data allow us to follow the workers continuously for many years.

Going on with the decomposition, substituting the income process described in Equation (1.3) in the within-worker component of inequality, we obtain:

$$W = \frac{1}{2\bar{y}^2}\left[\frac{1}{L}\sum_{i=1}^{L}\frac{1}{T}\sum_{t=1}^{T}(\alpha_i + \beta_i t + \epsilon_{it} - \bar{y}_i)^2\right] \tag{1.4}$$

Since $\alpha_i = \bar{y}_i$ by construction, we end up with

$$W = \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{L}\sum_{i=1}^{L}\frac{1}{T}\sum_{t=1}^{T}(\beta_i t)^2\right]}_{\text{Mobility risk}} + \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{L}\sum_{i=1}^{L}\frac{1}{T}\sum_{t=1}^{T}\epsilon_{it}^2\right]}_{\text{Volatility}} + \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{L}\sum_{i=1}^{L}\frac{1}{T}\sum_{t=1}^{T}(2\beta_i t\epsilon_{it})\right]}_{\text{Residual component}} \tag{1.5}$$

According to this further decomposition, aggregate mobility risk is the mean-normalized average across people of the individual variance of the income trend, while volatility is the mean-normalized average across people of the individual mean squared residual from the personal trend.[13] There is a residual component of covariance which has a very small order of magnitude and is negligible in the computations.

Why the income risk decomposition Figure 1.1 shows four examples of income trajectories taken from our data to look at very different income experiences and see the motivation under our choice of the method from Nichols (2008). For each worker, the

Figure 1.1. Very different earnings trajectories

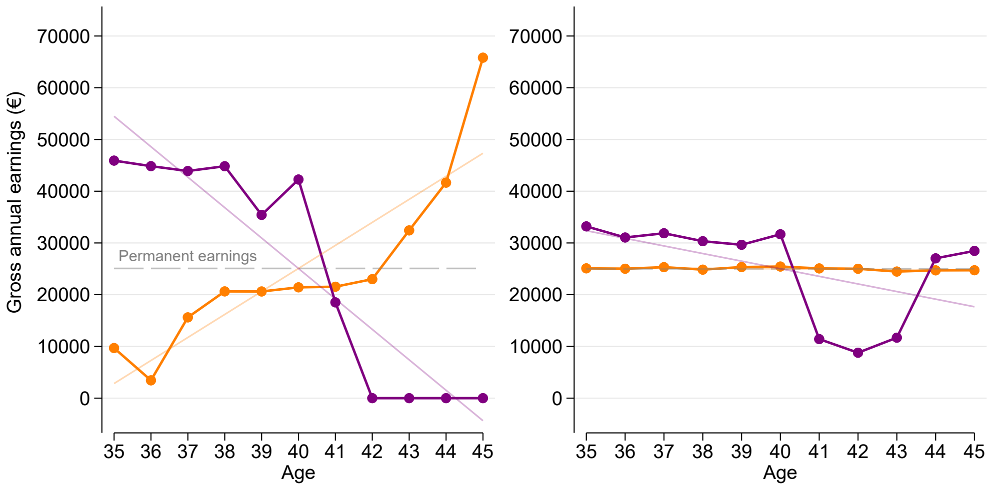

Note: The figure plots four representative career paths taken from real data. The four workers have approximately the same 'permanent earnings' (average earnings in the age window) but very different economic experiences in terms of direction and steepness of income trend (the solid line). Annual earnings are real (2015 price level) and gross of personal income taxes and social contributions and include income from any source. *Source*: AD-SILC data 1975-2018.

figure shows the permanent earnings in the age window from 35 to 45 (dashed grey line), a linear trend (solid lines), and the actual earnings records. We selected four workers with approximately the same permanent earnings: if average income is taken to proxy their economic well-being, we can say that there is no (permanent) inequality and the four workers enjoy the same level of well-being. However, they have completely different patterns over time: in the left panel, we see a worker with an exceptional career progression, ending up at age 45 with an income more than 6 times higher than the level at age 35, and enjoying quite smooth growth over time. In contrast, the other worker in the left panel experiences downward mobility and loses his job at the age of 42 after one year of halving his previous income.

The careers described in the left panel are good examples to understand why a simple measure of volatility measuring the dispersion of income deviation from the mean misses accounting for the existence of a 'good', desirable variability of income. In fact, attributing

[13]To see why $\left[\sum_{i=1}^{L} \frac{1}{T} \sum_{t=1}^{T} (\beta_i t^2)\right]$ is the individual-level variance of the points on the linear trend $\{\beta_i t\}_{t=1}^{T}$, remember that time t is centred at zero and the trend is demeaned, so that its average is zero by construction.

every deviation from the mean of the worker with a steep career progression to volatility means assuming that positive income growth is actually perceived as instability.

In the right panel of Figure 1.1, we compare a worker with a completely flat income in the window, and one experiencing a large drop (more than 2/3) at age 41 with a recovery thereafter. A flat income trajectory in the middle of one's career is not a good sign, since the accumulation of experience is not rewarded. On the other hand, the large and persistent (for two periods) income drop suffered by the person in purple in the right panel needs to be ensured through savings accumulated before or through borrowing relying on future earnings.

The framework described enables to look at the income experience as a three-dimensional phenomenon: the permanent component reflects overall experience, the result of variations in various directions that may offset each other; the mobility component reflects the 'smoothness' of the career progression; the volatility component reflects its instability.

Good and Bad mobility To reinforce this framework developed by Nichols (2008), we introduce a novelty to ease the interpretation of the results from a welfare point of view: the mobility risk component in Equation (1.4) is neutral with respect to the direction of the income trend; it measures the intensity, the speed of linear mobility, regardless of its direction. This is certainly a shortcoming for the interpretation, since we may consider desirable a rise in mobility risk which may actually come from an acceleration of 'linear falls'. To rule out this possibility, we further decompose the mobility risk: we divide the L workers into two types according to the direction of the income trend. $u = 1, ..., U$ are those with an upward linear trend, and $d = 1, ..., D$ those with a downward one. Mobility risk can be expressed as the sum of *upward* and *downward* mobility risk as follows:

$$\text{Mobility risk} = \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{L}\sum_{u=1}^{U}\frac{1}{T}\sum_{t=1}^{T}(\beta_u t)^2\right]}_{\text{Upward mobility risk}} + \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{L}\sum_{d=1}^{D}\frac{1}{T}\sum_{t=1}^{T}(\beta_d t)^2\right]}_{\text{Downward mobility risk}} \tag{1.6}$$

Therefore, we end up with a conceptual framework according to which overall inequality is the sum of permanent inequality, upward and downward mobility risk, volatility, and a residual component. The permanent inequality component is the part of inequality that is not smoothed out over time by mobility; it is due to differences in permanent income across workers, so it reflects inequality across people in terms of their lifetime economic possibilities. On the other hand, the upward mobility component is the expression for 'good' mobility, because it measures the intensity of smooth and linear income growth, which is the kind of absolute mobility that we consider more desirable for people and for the society. Finally, we include the sum of the downward mobility and the volatility components in a concept of 'bad' mobility: income changes that follow a linear progression but go down are equivalent to fluctuations, since they are neither desirable nor predictable during mid-career.

$$GE(2) = \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{LT}\sum_{i=1}^{L}\sum_{t=1}^{T}(\bar{y}_i - \bar{y})^2\right]}_{\text{Permanent inequality}} + \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{LT}\sum_{u=1}^{U}\sum_{t=1}^{T}(\beta_u t)^2\right]}_{\text{Good mobility}} + \tag{1.7}$$
$$+ \underbrace{\frac{1}{2\bar{y}^2}\left[\frac{1}{LT}\sum_{d=1}^{D}\sum_{t=1}^{T}(\beta_d t)^2\right] + \frac{1}{2\bar{y}^2}\left[\frac{1}{LT}\sum_{i=1}^{L}\sum_{t=1}^{T}{e_{it}}^2\right]}_{\text{Bad mobility}}$$

1.3 Data

Data source We need for our analysis longitudinal data covering a long part of individuals' careers. For this purpose, we use a selection of the Administrative-SILC (AD-SILC) dataset developed by merging through fiscal codes the waves from 2004 to 2017 of the IT-SILC survey (the Italian component of the European Union Statistics on Income and Living Conditions, EU-SILC) with social security records collected by the Italian National Social Security Institute (INPS). The INPS archives record employment and earnings histories of all individuals working in Italy from the moment they enter the formal labour market. Reliable earnings data are available from 1974 for employees in the private sector and later on for other types of employment. In the version of the dataset employed in this work, the latest year of observation is 2018.

In addition to the demographic characteristics, the administrative component allows to have detailed information on the gross annual earnings, allowances, the weeks worked in the year and the type of employment contract, while not suffering from attrition problems. On the other side, the survey component provides information on the level of education, which is always a great absentee in micro-level analyses using administrative data while being an important determinant of income.

This dataset is particularly suited for our analysis because of two characteristics that are crucial and rare in the existing literature on income mobility: (i) workers are followed for a large part of their career, allowing us to distinguish between short and long-term mobility; (ii) they are followed *continuously* as long as they participate in the formal labour market – without memory biases and, mostly, the gaps from attrition characterizing panel data from surveys. This latter feature largely improves the analysis: volatility is traditionally considered to be a short-term issue and requires observations very close in time, while mobility can be studied both as a short and long-term phenomenon. Having a long span of income records without 'holes' allows us to study mobility and volatility at same time looking also at their interaction.

The final sample records the income history of workers born between 1940 and 1973. We divide the sample into 30 five-year-long cohorts of birth, each of which overlaps with the preceding one for every year but the last one, from 1940-1944 to 1969-1973. Therefore, we observe earnings patterns from 1975 (when those born in 1940 are 35 years old) to 2018 (when those born in 1973 are 45 years old), and each cohort covers a calendar period of 15 years (for example, the first cohort 1940-1944 covers the period 1975-1989). All the analyses are performed *within* each cohort to allow comparison of intragenerational inequality and mobility over time.

Sample selection The sample is restricted excluding individuals without Italian citizenship, since the retrospective panel under-represents them in older cohorts. We focus on those working as employees in the private sector, which is the only category covering a very long-time span in INPS archives.[14] We use as a measure of economic well-being real (2015 price level) annual earnings from any job, also including allowances for sickness, maternity, unemployment and CIG, and gross of personal income taxes and social contributions.[15] Our aim is to capture through this measure of income the overall economic experience of workers before redistribution. The choice of annual earnings reflects our interest in economic well-being that includes the intensity of work during the year – in terms of weeks worked in the year and hours worked in the week -, as well as the hourly wage. The bottom and top 0.1% of the earnings distribution in each year are dropped to minimize measurement errors that may occur at the tails and to get rid of serious outliers.

It is possible that some workers, especially women in older cohorts, are out of the sample if they don't have any job for which social contributions are due to the INPS in the year. Those must be cases - without even sickness, maternity, unemployment, and CIG allowances, which we observe in the administrative archives -, spent either in non-employment, in inactivity, or in undeclared work. We assume that in the out-of-archives years the income from work is zero, so as to take into account periods of non-employment. We believe the treatment of the zeros to be a major issue in the mobility analysis: if the interest is in the overall economic well-being of a person, ignoring periods of non-employment and focusing on positive incomes naturally leads to a biased picture of reality.

As a final restriction, we select workers observed *continuously* for eleven years (from age 35 to age 45) with either positive or zero income from labour. Those workers with periods spent in jobs other than private employment or with missing information when 35-45 are excluded, since we do not want to impute zero earnings to people who are actually working

[14]On average, the dependent sector (public and private) represented about 69% of total employment at the end of the 70s, 71% at the end of the last century, about 75% in 2010 and 77% in 2018 (source ISTAT).

[15]The Cassa Integrazione Guadagni (CIG) is a short-work scheme for supporting the wages of employees for which firms going through specific crisis events request a reduction or a suspension of the employment relationship. It is limited in time and subject to specific requirements for both the employer's nature, the type of crisis, and the employment contract.

in a different form. Unfortunately, we are not able to distinguish periods of non-employment from informal work.[16] To avoid the inclusion of people mostly out of the labour market, we restrict the sample to workers with at least six years of positive earnings when 35-45. Importantly, while this will be our baseline sample, we will also check the differences in results when using a reduced sample from which zero earners are excluded for comparison.

Summary statistics Table A.1 for the baseline sample, and Table A.2 for the restricted sample excluding zero earners report in the Appendix for each cohort summary statistics on annual gross earnings and the composition of the sample in terms of gender, education and geographical area.[17] The sample includes 26,645 workers including those with at most five periods of non-employment when aged 35-45, and 21,849 workers when including only positive earnings. We can clearly observe in our sample that the Italian labour market has faced relevant structural changes linked to increasing women participation (women were 29.9% of workers in the first cohort, 45.2% in the last) and to the educational upgrading (workers with tertiary education were 2.7% in the first cohort, 14.8% in the last one). As regards the level and variability of earnings, we confirm with our data the well-known stagnation in average income from labour from the 90s coupled with increasing standard deviation.

1.4 Results

As a first set of results, we briefly look at the trend across subsequent cohorts of the indices of intragenerational inequality and mobility described in Section 1.2 and reported for three representative cohorts (the first, the last, and one in the middle) in the Appendix in Table A.3 with the percentage variation from the first to the last cohort.[18]

1.4.1 Bi-periodical mobility patterns

Positional mobility The left and centre panels in Figure 1.2 plot the indices of positional mobility. Starting from quintile mobility, we see that between 20 and 25% of workers within each cohort move to a different quintile, and that the probability of moving to a higher or a lower quintile is almost symmetric: more than 50% of workers remain in their

[16]According to ISTAT estimates, undeclared work involved 14.5% of employment in 1995, 12.4% in 2005, and 12.8% in 2018, with slightly lower percentages if excluding self-employed. Informal work can impact our measure of earnings through two channels: by 'hiding' workers who are actually employed, and by distorting downwards the earnings of workers who are employed in the formal labour market but also receive unrecorded pay.

[17]In this work, we include the two main islands of the country (Sicily and Sardinia) in the macro area 'South'.

[18]Table A.5 and A.6 in the Appendix report all the indices for the 30 cohorts with standard errors obtained through 100 bootstrap repetitions. The normal-based confidence intervals at 95% confidence level in Figure 1.2 and 1.4 are based on those standard errors.

origin quintile when they reach age 43-45 and, among those who move, half improve their position, and half worsen it. Looking at the tails of the income distribution, we find that between 30 and 35% of workers starting in the bottom quintile at age 35-37 manage to get out of it after 10 years, while between 15 and 25% of those who start from the top quintile end up in a lower position.

Figure 1.2. Earnings mobility patterns

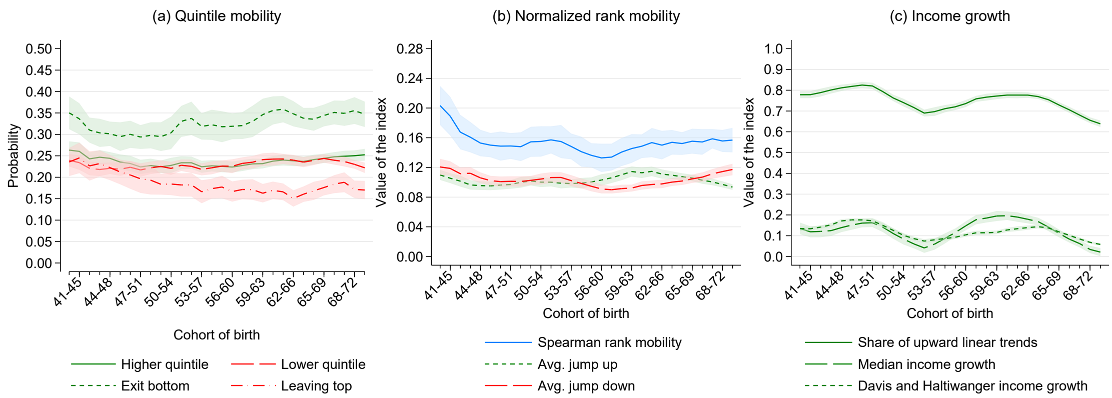

Note: The figure plots several intragenerational mobility indices for 30 five-year-long rolling cohorts of birth of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. IT-SILC sample weights are used to compute the indices and normal-based confidence intervals (95%) are obtained through 100 bootstrap repetitions. *Source*: AD-SILC data 1975-2018.

These indices uncover a relevant persistence in income positions in 10 years in middle-career, and the existence of 'sticky floors' – low positions hard to escape from – and 'sticky ceilings' – high positions that are unlikely to be left. The ceiling seems to be 'stickier' than the floor: less than one-fifth of the top earners change position after 10 years, meaning that being a top earner is a persistent status. Moreover, if we look at trends across cohorts, the probability of leaving the top decreases by 28%: from 0.24 for cohort 1940-1944, to 0.17 for cohort 1969-1973.

Measuring mobility through changes in normalized ranks, we see that the correlation between origin and destination positions is high (between 0.80 and 0.85) making the Spearman's mobility index lie between 0.15 and 0.20. The mobility index decreased by 23% from the first to the last cohort, but the drop occurred mainly for the first cohorts and then mobility remained stable at lower levels. Looking at the direction of rank changes after 10 years, also in this case we find symmetry in the movements: workers climbing the income ladder, as well as those who fall down, cross about 10% of the income distribution.

Absolute mobility The right panel in Figure 1.2 plots the median income growth for each cohort, the average growth à la Davis and Haltiwanger, and the share of workers enjoying an upward linear trend at age 35-45.[19] Interestingly, the three indices tell the same story in terms of patterns across cohorts: with some cyclicalities, the long-run trend of income growth across generations is markedly decreasing: 56% less median income growth from the first to the last cohort, and 84% less average growth à la Davis and Haltiwanger. We move from a picture of 13% income growth after 10 years for workers born in 1940-1944, to one between 2 and 6% for those born between 1969 and 1973.[20] Finally, the share of upward trends is consistent with this picture: the probability of experiencing smooth upward growth is 18% lower for the last cohort (0.64) than for the first one (0.78). A value of 0.64 means that almost 40% of the workers belonging to the last cohort do not benefit from 'good' mobility in the central phase of their careers.

Non-linearities What we have seen so far provides an average picture of intragenerational mobility as measured by bi-periodical indices. However, average values may hide very different behaviours along the income distribution. To inspect such non-linearities in the association between origin and destination income, we rely on the graphical tool in Figure 1.3: for three cohorts of birth (again the first, the last, and one in the middle), we plot destination income against origin income using the *value* in Euros in the left panel, and the *normalized rank* in the right one. Since a full scatterplot would be unreadable, we average income (left panel) and income ranks (right panel) in 10 bins, and also plot the linear fit of the full scatter. The reference to read the graph is the 45-degree line, which is the place of perfect immobility.

We see that in the left panels the fitted line is always above the 45-degree line, meaning that, in general, destination income is higher in value than origin income, and confirmed by the positive average growth we measured (Figure 1.2 and Table A.3). However, the slope changes: for cohort 1940-1944, the bottom of the origin distribution experiences on average greater income growth than the top, while the reverse is true for the subsequent cohorts.[21] Looking at the scatter points, clear non-linearities emerge: the middle-class of origin distribution has almost stagnating income, even falling for the last cohort, while the

[19]We do not show in the graph the average income growth because the standard errors for some cohorts are too wide to make the estimates credible (see Table A.5 in the Appendix). As explained in Section 1.2.2, the average growth rate is very sensitive to outliers; including zero earners and measuring origin income as an individual mean at age 35-37, we may have very small values of origin income that result in outstanding levels of growth. We see this also by comparing the level of mobility including zero earners in Table A.3 (61% growth for the first cohort!) with that for positive earners only in Table A.4 (19%) in the Appendix. For this reason, we rely on the more robust alternatives described in the text to aggregate absolute income growth.

[20]This evidence can not be simply attributed to zero-income women in older cohorts – who therefore experienced more growth at entry – because it is also confirmed by results using only positive earnings (Figure A.2 in the Appendix).

[21]Cohorts 1940-1944 to 1945-1949 are the only two cohorts completely covered during age 35-45 by the Scala Mobile - 'elevator' - wage indexation mechanism adopted in Italy from the 1970s to the early 1990s. Since it was designed for granting the same absolute wage increase to all employees in a period of sustained inflation, the mechanism induced mechanically greater proportional wage changes at the bottom of the distribution (Manacorda, 2004).

Figure 1.3. Correlation between origin and destination earnings

a) Cohort 1940-1944

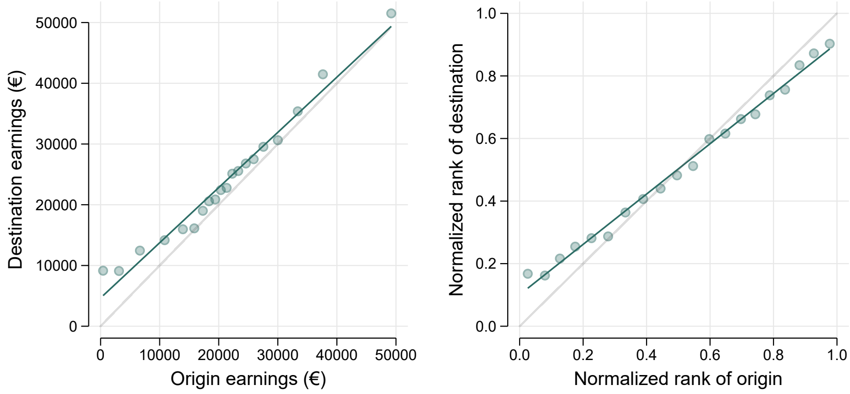

b) Cohort 1955-1959

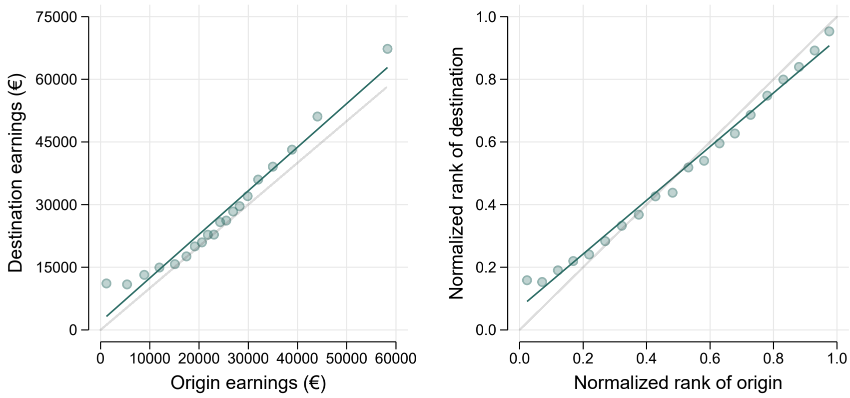

c) Cohort 1969-1973

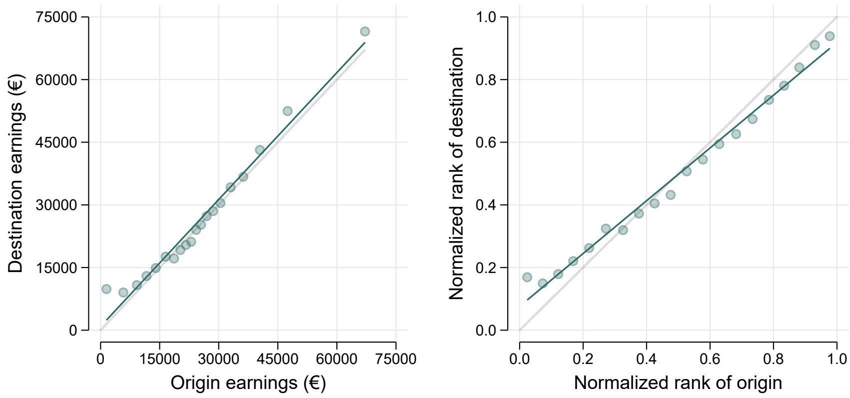

Note: The figure plots for three cohorts of birth the linear fit of destination earnings on origin earnings (left panels), and of destination income rank on origin income rank. The points are the average y-variable and x-variable inside 10 equal-sized bins. The 45-degree line is the place of perfect immobility, where destination income/rank is perfectly predicted by origin income/rank. The sample includes employees observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. Annual Earnings are real (2015 price level) and gross of personal income taxes and social contributions and include income from any source. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

bottom and the top experience the highest income growth levels. There seems to be no reversion to the mean in place – i.e. the higher the income, the lower the growth –, but rather a U-shape pattern of income growth.

A similar picture emerges from the right panels showing normalized rank mobility. Positional mobility is a zero-sum game – if someone goes up, someone else has to go down –, so the fitted line cannot be completely above or completely below it. This way, absolute income growth is ignored and the cohorts can be compared in terms of positional mobility alone. We see a process of mean reversion – those in the lower half of the origin distribution tend to improve their ranking, while those in the top half tend to worsen it –, but with the very top more sheltered from this process and closer to maintaining its position. If we exclude zero-income workers (Figure A.4 in the Appendix), we no longer see workers at the bottom improving their position on average, signalling that the 'bottom-out' phenomenon is driven mainly by the exit from non-employment, while low-income workers tend to remain low-income workers also after 10 years.

1.4.2 Income risk components

We now move to the intragenerational indices of dynamics, shown in Figure 1.4. In the left panel of the figure, we plot the within-cohort levels of inequality as measured by the GE2 index. We notice a relevant long-run trend of rising earnings inequality, increasing by 39% from the first to the last cohort (60% if excluding zero earnings) if measured by the average cross-sectional GE2. Comparing the pattern of average and overall inequality with the permanent one gives a first clue about income mobility: the more they depart from each other, the more people experience income movements, according to the decomposition in Equation (1.2). However, mobility can come from very different income trajectories, more or less growing, and more or less stable.

To inspect the details of mobility, we plot in the central panel in Figure 1.4 the three separate elements of within-worker inequality – upward linear mobility, downward linear mobility, and volatility –, and in the right panel the per cent contribution of each element to overall inequality. The intensity of upward linear mobility is always greater than that of downward mobility, but their difference becomes very narrow for recent cohorts due to a long-run trend of declining good mobility (-32% from the first to the last cohort). Moreover, the level of average individual volatility seems to be close to that of upward smooth mobility and less cyclical. In the whole period, volatility diminished by 15%. These patterns are similar if we exclude zero earnings, with the main difference being the relationship between good mobility and volatility: individual volatility remains for most cohorts lower than good mobility but for the most recent cohorts. This suggests that periods of non-employment have a strong impact on the level of volatility, but for recent cohorts the instability is greater than good mobility even without the impact of the zeros.

Figure 1.4. Income risk components

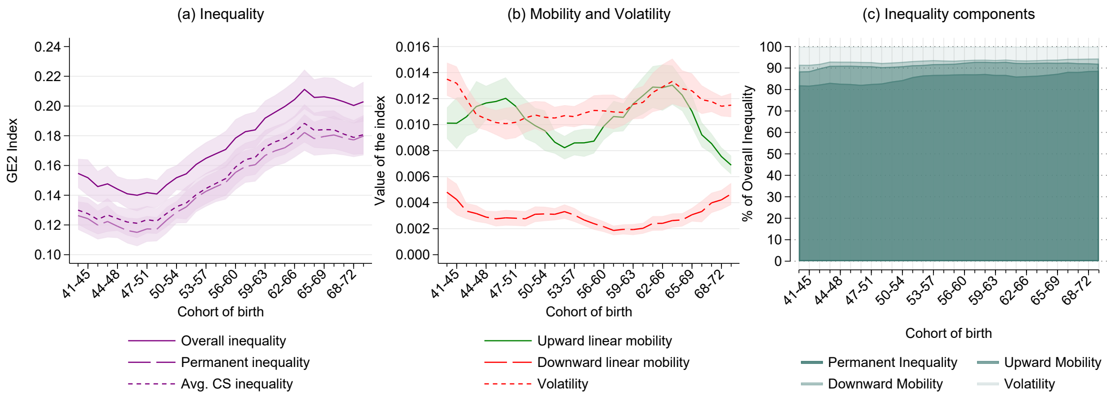

Note: The figure plots cohort-by-cohort the overall (across people and time) intragenerational inequality and its components according to the decomposition described in Section 1.2.2. The indices are computed for 30 five-year-long rolling cohorts of birth of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Inequality is measured through the general entropy index of degree 2. Workers with zero earnings for at most five years are included. IT-SILC sample weights are used to compute the indices and normal-based confidence intervals (95%) are obtained through 100 bootstrap repetitions. *Source*: AD-SILC data 1975-2018.

Finally on the income risk decomposition, the right panel of Figure 1.4 reveals that overall inequality is mainly due to persistent differences across workers (more than 80%), while the rest is for one-third 'good' and for two-thirds 'bad' individual mobility. Importantly, the share of overall inequality attributable to persistent differences increases across cohorts: from 81 to 89% including zero earnings, and from 84 to 91% for positive earnings.

1.4.3 The inequality-mobility trade off

Given the picture of inequality and mobility provided in the previous section, we now move to study the correlation between the two. Table 1.1 (and Table A.7 in the Appendix for only positive earnings) reports the correlation coefficients between each of the mobility indices explained in Section 1.2.2 and the three notions of inequality we employed in this work – namely overall inequality, permanent inequality and average cross-sectional inequality. As a first interesting result, the magnitude of the correlation does not change much for the three notions of inequality, confirming that overall and average inequality are driven by permanent differences across workers.

If we look at the two cross-country intragenerational Great Gatsby curves taken from Gangl (2005) and Guvenen et al. (2022) (Figure A.1 in Appendix), we see that Italy is a middle-high inequality country, but its relative position in terms of mobility depends on the index used. This is why we want to study the correlation between inequality and mobility by employing several different concepts of income dynamics.

Indeed, we see that the correlation is heavily dependent on the index used, but some

Table 1.1. Table of inequality-mobility correlation

			Overall Inequality	Permanent Inequality	Avg. Inequality
		1 - ρ_n	-0.001	-0.038	-0.046
		Avg. Jump up	0.380	0.353	0.380
	Relative	Avg. Jump down	0.022	0.034	0.064
	indices	Pr(upper quintile)	0.318	0.298	0.280
Biperiodical		Pr(lower quintile)	**0.560**	**0.541**	**0.556**
mobility		Pr(exit from bottom)	**0.757**	**0.762**	**0.751**
indices		Pr(falling from top)	**-0.725**	**-0.752**	**-0.753**
	Absolute	Avg. Income growth	-0.236	-0.251	-0.208
	indices	Median Income growth	-0.473	**-0.521**	-0.475
		DH Income growth	-0.080	-0.114	-0.063
		Pr(upward linear trend)	**-0.581**	**-0.623**	**-0.579**
Indices of		Avg. upward mobility	-0.087	-0.140	-0.086
dynamics		Avg. downward mobility	0.097	0.093	0.057
		Avg. Individual volatility	**0.709**	**0.669**	**0.678**

Note: The table reports the cohort-level correlation between earnings mobility and inequality indices. All the coefficients are significant at 95% confidence level unless the number is in light grey. We highlight in bold the correlations greater or equal to 0.5. The underlying basis for computing the indices are 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

regularities emerge. In terms of the sign of the association, there seems to be a trade-off between inequality and mobility as measured by the Spearman index, the average and median income growth, the probability of having an upward linear trend, and the probability of leaving the top quintile after 10 years. On the contrary, a complementarity emerges between inequality and mobility as the probability of changing quintile, the average jump, the probability of exit from the bottom quintile, the intensity of downward mobility risk, and the average individual volatility. The pattern is puzzling: while it is clear that there is a trade-off between inequality and most of the notions of ‘good’ mobility, we also find a positive association between inequality and the probability to escape from the bottom of the distribution.

The key to explaining this puzzle are the non-employed. If we compare the correlations in Table 1.1 with those for the indices excluding zeros in Table A.7, we notice that when only positive earnings are included the results are remarkably clear: there is a trade-off between inequality and every notion of good mobility, while inequality is positively linked to the three measures of bad mobility we have – namely, the average jump down, the intensity of downward mobility risk, and average volatility. This gives us two important results: first, the cohorts experiencing a higher level of inequality do not see it compensated by good mobility, but rather suffer the effects of the worst notion of mobility which is instability. Second, since we know that by including the zeros a positive correlation emerges between inequality and the probability of leaving the bottom quintile, this means that the most unequal cohorts experience more mobility at the bottom due to workers exiting the non-employment status.

In terms of magnitude, two correlations dwarf the others: the negative correlation between inequality and the probability to leave the top (-0.75 including zeros, -0.80 excluding them), and the positive one between inequality and volatility (0.69 including zeros, 0.87 excluding them). The former has no unique interpretation: from an individual-welfare perspective, leaving the top quintile after 10 years is a bad, since it is a downgrading in the income ladder. However, from a social welfare point of view having low mobility at the top – sticky ceilings – is a risk in terms of inequality of opportunity, concentration and strengthening of power, to the point of being a threat to the functioning of democracy. On the contrary, the welfare interpretation of the complementary between inequality and volatility is much simpler, being volatility undesirable for its unpredictability.

Focusing only on the mobility indices that show a high correlation ($>$0.5) with the level of inequality, we plot in Figure 1.5 and A.5 the cross-cohort intragenerational Great Gatsby curves. On the y-axis, there is always the average GE2 inequality index, while on the x-axis there are several indices of mobility. The curves are informative beyond what we already saw in Table 1.1 because we can also 'locate' the cohorts in the graph: using a darker color for the most recent cohorts, we see that there has been a gradual shift from one generation to the next toward greater inequality coupled with stickier ceilings, lower growth, and greater instability.

1.4.4 Evidence on 'unequal mobility'

Besides the aggregate dynamics, we are interested in investigating who has been most impacted by the different types of mobility we are measuring. Figure 1.6 shows for three different cohorts (1940-1944, 1955-1959, and 1969-1973) the average decile of permanent earnings when 35-45 by combinations of decile of good mobility (x-axis) and bad mobility (y-axis). Good mobility is measured as the steepness of the upward linear trend, while bad mobility is the sum of the steepness of the downward linear trend and volatility around the trend. Darker areas in the heat maps indicate the 'places' of the mobility combination where richer people are concentrated.

We see in Figure 1.6 gradually darkening colour from the upper left corner to the lower right corner, with a more distinct pattern for the two youngest cohorts: low-permanent income people tend to be concentrated in the first half of the distribution of good mobility,

Figure 1.5. intragenerational Great Gatsby curves

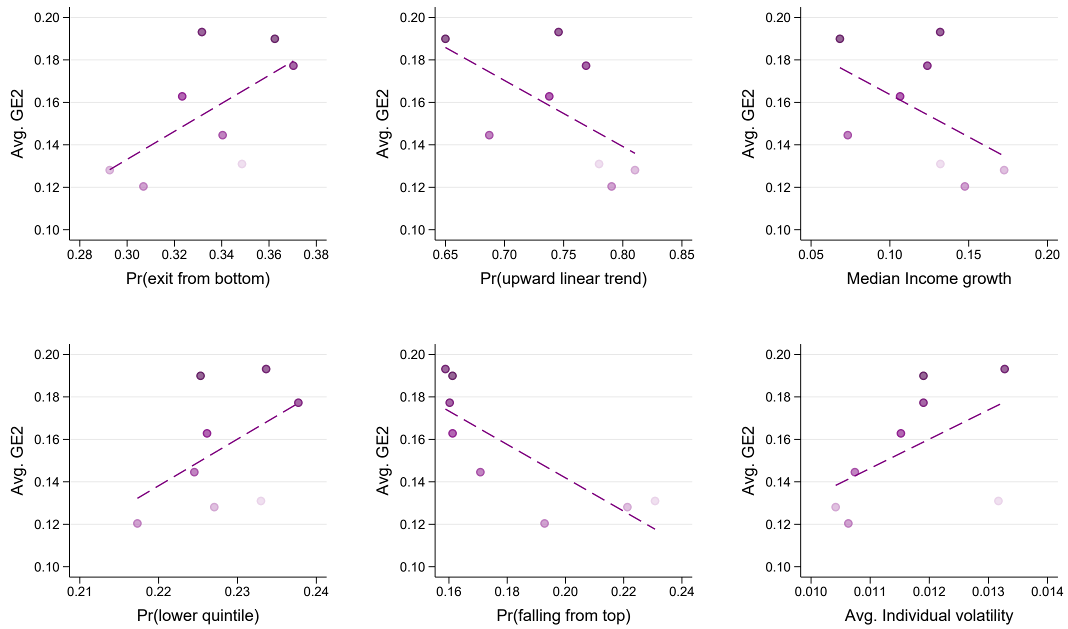

Note: The figure plots average within-cohort inequality measured through the GE2 index against several measures of intragenerational mobility. The selected measures of mobility are those with a correlation with inequality greater than 0.5 in Table 1.1. Only the cohorts of birth overlapping for one year are shown for clarity (1940-1944, 1944-1948, ..., 1968-1972), and the colour of the circle gets darker for more recent cohorts. The inequality and mobility indices are computed on a sample of employees in the private sector in Italy observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Figure 1.6. Heat map of unequal mobility

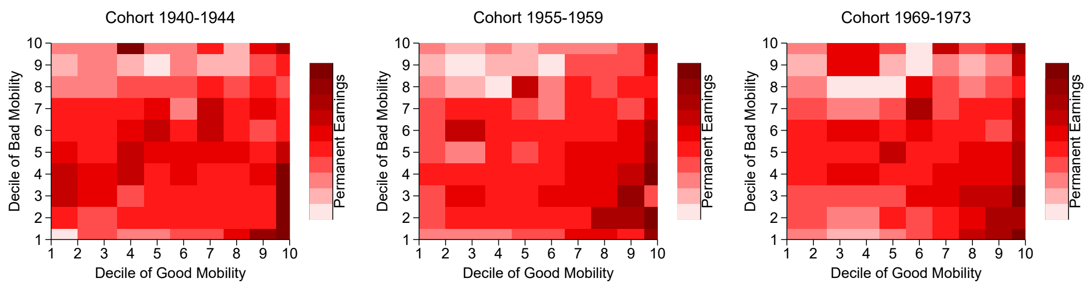

Note: The figure shows for three cohorts of birth the 'heatmap' of decile of permanent earnings – average income at age 35-45 – for the combination of deciles of 'good' (x-axis) and 'bad' (y-axis) mobility. Darker areas indicate a greater decile of permanent earnings. 'Good' and 'bad' mobility are estimated through the income risk decomposition à la Nichols described in Section 1.2.2 and measure, respectively, smooth upward income growth and individual income volatility. The sample includes employees in the private sector in Italy observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

and at the top of the distribution of instability, and the reverse is true for high-permanent income recipient. We interpret it as evidence of unequal mobility in place. Looking at figure A.6 in the Appendix to compare these results with the case of only positive earnings, we notice two interesting differences. First, for the oldest cohort (1940-1944) permanent income is distributed rather independently of mobility. Second, for the other two cohorts permanent income follows the distribution of good mobility, but not that of instability: richer people benefit on average from greater smooth growth, but the burden of volatility is shared across the distribution. Therefore, a relevant component of unequal mobility are the transitions to and from non-employment: they lead to a permanent low-income state worsened by high levels of instability.

1.5 Heterogeneity

As a further and final insight into intragenerational mobility, we look at possible heterogeneity linked to relevant socio-demographic characteristics of workers – namely the gender, the highest level of education, and the area of work. To explore the differences in mobility by these categories, we regress separately for each cohort the several individual-level measures of mobility – one at a time – on indicators for worker i being a woman (W_i), tertiary graduate (T_i), and working in the South or Islands of Italy (S_i), controlling for the rank of origin (R_i^o):

$$\text{Mobility}_i = \beta_0 + \beta_1 W_i + \beta_2 T_i + \beta_3 S_i + R_i^o + \epsilon_i \tag{1.8}$$

Figure 1.7. Gender differences in earnings mobility

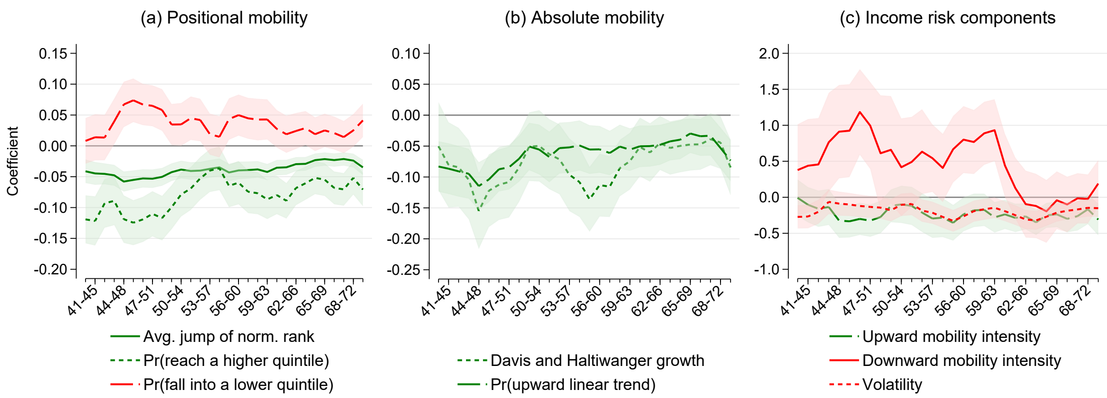

Note: The figure plots by cohort of birth the coefficient of an indicator variable for being women in several OLS linear regressions of mobility measures controlling for being a tertiary graduate, working in the South of Italy, and for the normalised rank at age 35-37. The mobility variables in panel (c) are taken in log. The regressions are fitted separately for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Figure 1.7 plots the coefficient β_1 for several measures of mobility. In terms of positional mobility (panel a), we find relevant gender asymmetries: women are always less likely to reach a higher quintile than men, while for some cohorts they are more likely to worsen their position after 10 years. Even if the gap is decreasing across cohorts, in the last cohort women are still less likely than men to step up by about 7pp, and more likely to step down by about 4pp, despite equal education, area of work, and rank of origin. If we compare these results with those excluding zero earnings (Figure A.7 in the Appendix), we discover that the higher probability to fall into a lower quintile for women is due to the transition to non-employment, while the lower probability to step up stays there: a glass ceiling makes it harder also for women attached to the labour market to improve their position as compared to a man.

Gender differences emerge also in terms of absolute mobility: women's income growth, as well as their probability of having an upward linear trend, is systematically dominated by men's one. Even in this case, there is a long-run trend of reduction of this gap, but it is still there for the recent cohorts (-8.5pp for income growth, -7.3pp for the probability of an upward trend). When excluding zeros, with less cyclicality, the gender gap remains at the same level. Finally, looking at gender differences in the income risk components, we see the impact of the increase in female participation: up to recent cohorts, women had slightly lower upward mobility risk and volatility, but a great disadvantage in terms of downward mobility risk. When excluding the zeros, most of the coefficients lose significance across cohorts, and also the gender differences in upward mobility risk and volatility seem to disappear for recent cohorts. Moving to the differences in mobility by level of education,

Figure 1.8. Education differences in earnings mobility

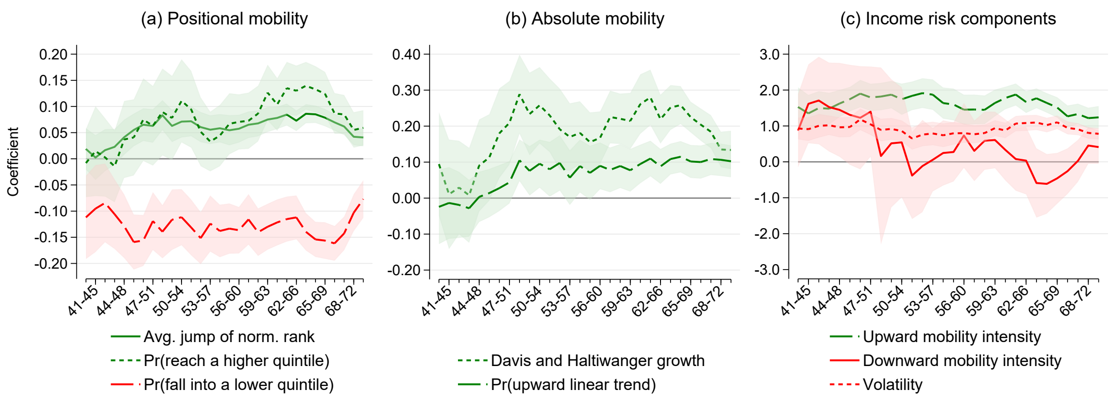

Note: The figure plots by cohort of birth the coefficient of an indicator variable for being a tertiary graduate in several OLS linear regressions of mobility measures controlling for being a woman, working in the South of Italy, and for the normalised rank at age 35-37. The mobility variables in panel (c) are taken in log. The regressions are fitted separately for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

we see in Figure 1.8 the coefficient β_2 of the indicator for tertiary education given gender, area of work and rank of origin. We find a clear and sizable advantage related to education in terms of positional mobility and earnings dynamics: tertiary graduates are more likely to improve their ranking after 10 years, but even more so they are sheltered from falling into a lower quintile (between 10 and 15pp of advantage). However, this positional mobility advantage is shrinking for recent cohorts. In terms of absolute mobility, tertiary graduates have acquired a considerable advantage of growth (between 10 and 30pp), and they also experience higher levels of upward mobility intensity while being affected by more volatile earnings. As a final dimension of heterogeneity, we look at the differences in individual-level

Figure 1.9. Geographical differences in earnings mobility

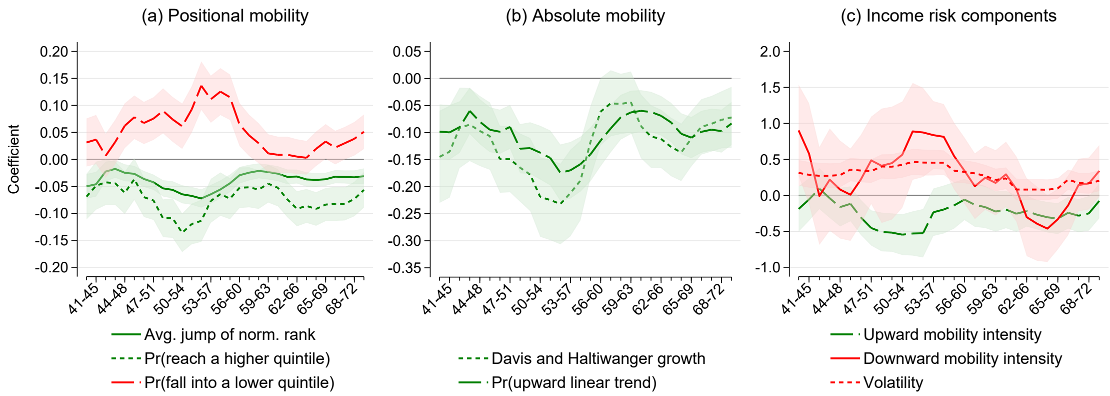

Note: The figure plots by cohort of birth the coefficient of an indicator variable for working in the South of Italy in several OLS linear regressions of mobility measures controlling for being a woman, being a tertiary graduate, and for the normalised rank at age 35-37. The mobility variables in panel (c) are taken in log. The regressions are fitted separately for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

mobility by macro area of work. Figure 1.9 shows the coefficient β_3 for workers in the South and Islands of Italy, given gender, education and rank of origin. The picture resembles that of women: there is a ‘geographical gap’ in terms of positional mobility, being workers in the South more likely to step down and less likely to step up but with a converging pattern across cohorts. The gap in absolute mobility is reducing over time but is indeed very huge (between 5 and 20pp), and also the income risk components are not randomly distributed across geographical areas: it is less likely for workers in the South to benefit from good mobility (but the gap is zero for some cohorts), while they are the ones with more volatile earnings.

1.6 Discussion and Conclusion

We started our investigation by asking whether income inequality can be more acceptable if coupled with a high degree of mobility along the income distribution that makes the inequality burden widely shared through ‘changing fortunes’ (Jenkins, 2011). Our analysis of the correlation between intragenerational inequality and mobility in the case of Italy prompts us toward a negative answer. Indeed, we find evidence of an empirical trade-off between income inequality and 'good' mobility, and complementarity with the worst notions of mobility – i.e. those related to income insecurity. Instead of being combined with more mobility, the rising inequality experienced by Italian cohorts of workers in the last decades has been increasingly set in stone: younger cohorts are burdened with greater gaps to start with that are not transitory and are reproduced even ten years later.

Being aware of the persistence of inequality is crucial for planning policies from two perspectives. First, if income differences are not transient, the initial positioning of the individual is decisive. Hence, the role of all those determinants of initial positions (education, sector and occupation, but also the quality of the employer, the socio-economic background etc...) is even more important and specific investments are needed to guarantee equal access to the best starting points. Second, the very persistence of inequality can be the subject of specific policies. In fact, we find that poor workers can improve their condition by finding a job, but when they are employed 'the floors are sticky' and it is hard to escape a low-income condition, while those at the top of the distribution remain firmly anchored there. This evidence suggests the need to address the underlying factors of low-income status to mitigate the risks of chronic poverty and social segregation related to its persistence, and to remove the factors that make top positions inaccessible to those in other parts of the distribution to avoid excessive concentration of economic power.

In a poorly mobile society like the one we have described, measuring permanent inequality or simply the cross-sectional one does not make a big difference; however, this is something that needs to be proved in the first place and not taken for granted. And it could, in any case, change from society to society and from generation to generation. We think that this work, besides shedding some new light on the link between inequality and mobility in the Italian case and the intragenerational context, has made some methodological contributions, or at least important discussion points. The first one is related to the inclusion of zeros: we showed how much sensitive the measurement of mobility is to the zeros, and we tried to interpret case-by-case the possible impact of periods of non-employment. When the focus is on income mobility, including its aspect of insecurity, we cannot leave out of the picture exactly those who are more mobile, if only on the extensive margin. Moreover, when the analysis includes both men and women, the consideration of movements in and out of employment acquires even more importance.

A second methodological contribution is related to the very notion of intragenerational mobility. In the wide range of possible definitions, concepts and methodological details, we decided not to choose so as not to create constraints. This non-choice allowed us to analyze different aspects of mobility and its association with inequality, also being able to break it down into its components that even go in opposite directions in terms of individual and social well-being. With a single concept of mobility, this would not have been possible.

Our analysis suffers from some limitations that need to be taken into account: we focused only on private employees due to data limitations, and our earnings measure suffers from misreporting due to a lack of information on informal employment. Moreover, we couldn't exploit information on family conditions, which certainly influence income mobility, and the role of taxes and transfers. Further research is needed to understand how redistribution affects the observed patterns of inequality and mobility and the role of intra-household insurance in protecting against income shocks, particularly against job loss. We also hope that further work can go up to more recent times to cover the Covid-19 crisis and the generations born later than 1973, and shed some light on the micro and macro drivers of the observed mobility trends.

Appendix A

Additional material

A.1 Additional figures and tables

Figure A.1. Intragenerational Great Gatsby curves

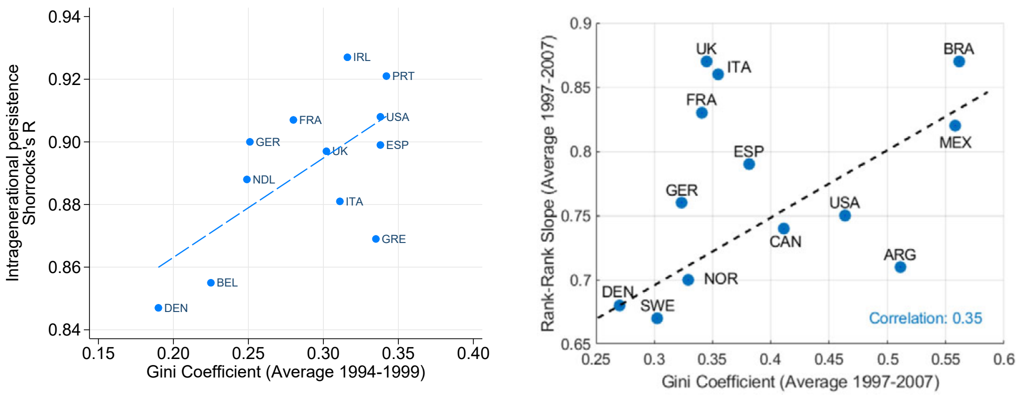

Source: Left panel: authors' elaboration from Gangl (2005) (Table 1, p. 150), Panel Study of Income Dynamics and the European Community Household Panel. Right panel: Guvenen et al. (2022) (Fig. 12, p. 1356), GRID data.
Note: Shorrocks's R index (Shorrocks, 1978) is the ratio between inequality (Gini index) computed on average income and the average cross-sectional inequality in the same period. It measures how much of the snapshot inequality is due to persistent income differences. The Rank-Rank slope is the beta coefficient of a linear regression of the income rank at the end of the period on the income rank at the beginning of the period.

Figure A.2. Mobility patterns – only positive earnings

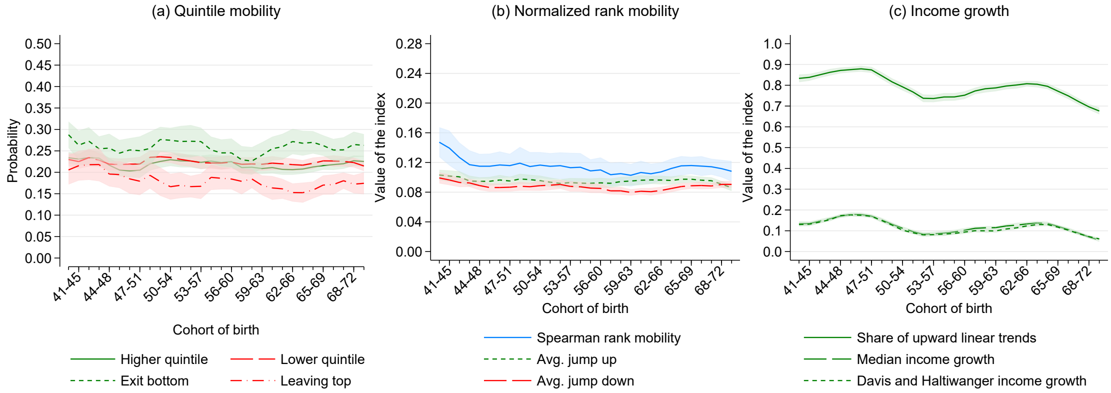

Note: The figure plots several intragenerational mobility indices for 30 five-year-long rolling cohorts of birth of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Only workers with positive earnings every year when aged 35-45 are included. IT-SILC sample weights are used to compute the indices and normal-based confidence intervals (95%) are obtained through 100 bootstrap repetitions. *Source*: AD-SILC data 1975-2018.

Figure A.3. Income risk components – only positive earnings

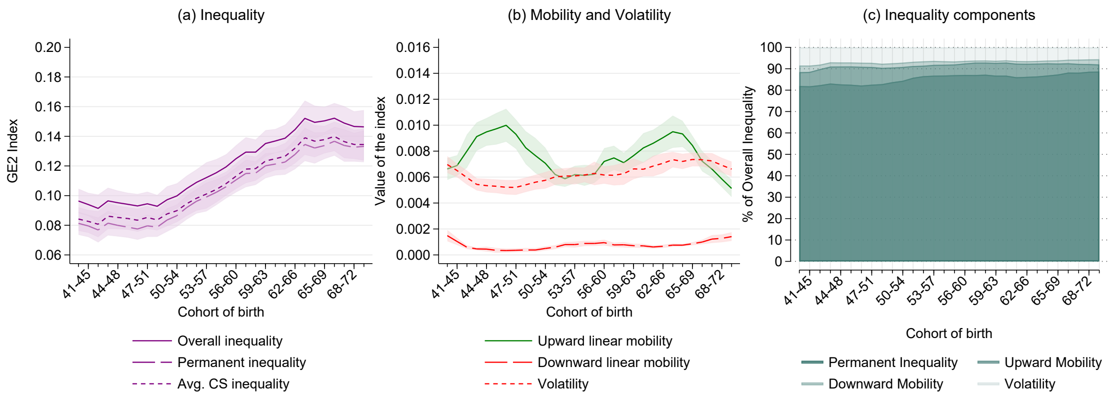

Note: The figure plots the overall (across people and time) intragenerational inequality and its components according to the decomposition described in Section 1.2.2. The indices are computed separately for 30 five-year-long rolling cohorts of birth of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Inequality is measured through the Generalized Entropy Index of degree 2. Only workers with positive earnings every year when aged 35-45 are included. IT-SILC sample weights are used to compute the indices and normal-based confidence intervals (95%) are obtained through 100 bootstrap repetitions. *Source*: AD-SILC data 1975-2018.

Table A.1. Summary statistics

		Annual earnings (€)								
Cohort	N	Mean	SD	p10	p50	p90	Zeros	Women	Tertiary	South
1940-1944	2,952	22,137	12,259	3,614	21,997	36,210	6.0	29.9	2.7	20.1
1941-1945	2,837	23,355	13,190	6,006	22,960	36,574	4.5	30.4	3.8	21.4
1942-1946	2,983	22,941	13,760	3,091	22,427	37,994	6.0	31.7	2.8	20.6
1943-1947	3,126	23,716	13,015	6,173	23,338	38,813	4.2	30.8	3.7	21.9
1944-1948	3,235	24,975	13,577	7,770	23,600	42,616	2.7	31.4	3.0	23.3
1945-1949	3,322	25,561	13,612	8,699	24,907	39,277	4.1	32.6	4.6	20.7
1946-1950	3,414	25,381	13,314	8,768	24,609	40,502	3.7	33.9	3.2	22.7
1947-1951	3,335	25,678	13,846	7,449	25,027	41,121	4.7	31.2	3.9	19.8
1948-1952	3,238	25,728	13,696	8,079	24,751	41,960	4.4	30.8	3.5	27.3
1949-1953	3,186	25,253	14,334	6,352	24,317	43,740	4.2	33.6	4.5	26.0
1950-1954	3,178	25,629	14,475	5,736	24,824	41,928	6.1	36.0	4.6	18.8
1951-1955	3,240	25,929	14,995	7,181	25,189	42,359	4.3	31.2	4.5	22.8
1952-1956	3,320	26,023	14,907	7,419	25,009	45,262	4.7	35.4	2.7	18.2
1953-1957	3,402	25,585	14,281	6,326	24,191	43,574	4.2	39.0	4.4	20.8
1954-1958	3,490	26,294	14,749	7,684	25,398	43,088	4.8	36.2	4.2	23.6
1955-1959	3,555	26,557	15,424	9,879	24,536	43,765	3.3	35.0	4.4	18.7
1956-1960	3,573	26,388	16,360	5,281	25,238	45,993	4.9	35.1	4.7	21.6
1957-1961	3,648	25,295	14,183	7,274	24,157	44,686	4.2	35.9	6.7	28.6
1958-1962	3,828	26,282	15,694	8,902	23,652	47,507	3.8	40.8	6.1	20.8
1959-1963	4,028	24,160	15,148	5,805	22,907	43,873	4.3	36.2	5.2	28.0
1960-1964	4,282	25,002	15,858	6,403	23,632	44,232	5.3	35.8	6.8	25.0
1961-1965	4,584	26,176	15,587	7,752	24,737	46,530	3.7	36.2	6.4	25.5
1962-1966	4,854	24,040	15,546	5,803	22,893	42,133	5.2	41.4	7.5	22.5
1963-1967	5,048	24,749	16,611	5,165	22,800	44,776	4.8	36.0	6.6	24.2
1964-1968	5,171	25,185	16,796	6,276	23,206	45,721	4.3	42.5	6.7	26.4
1965-1969	5,258	27,501	18,941	7,638	24,571	48,408	4.3	39.1	8.5	23.7
1966-1970	5,308	25,352	17,042	6,165	23,716	44,052	4.1	40.6	7.4	25.7
1967-1971	5,285	24,897	15,895	6,622	23,670	42,786	5.2	42.5	11.6	28.2
1968-1972	5,243	26,060	17,518	6,878	23,713	47,433	4.4	43.8	14.1	27.3
1969-1973	5,173	26,036	17,001	5,671	24,284	46,201	4.9	45.2	14.8	24.4
All	26,645	25,262	15,500	6,584	23,894	43,332	4.6	37.2	6.9	23.7

Note: The table reports the number of workers and summary statistics for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. Annual Earnings are real (2015 price level) and gross of personal income taxes and social contributions and include income from any source. The percentage of zero earnings, women, tertiary graduates, and workers in the South of Italy are reported. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Table A.2. Summary statistics – only positive earnings

		Annual earnings (€)							
Cohort	N	Mean	SD	p10	p50	p90	Women	Tertiary	South
1940-1944	2,362	25,520	11,196	14,572	23,898	38,412	24.1	3.2	18.2
1941-1945	2,278	25,567	10,997	14,364	24,113	38,300	26.0	3.4	18.8
1942-1946	2,439	26,381	11,335	15,145	24,813	39,198	26.6	3.5	19.7
1943-1947	2,590	27,202	12,127	15,595	25,321	40,571	26.0	4.1	20.8
1944-1948	2,706	27,375	12,030	15,603	25,611	41,121	26.7	3.8	21.5
1945-1949	2,800	27,700	12,118	15,766	25,976	41,458	27.1	3.8	21.7
1946-1950	2,872	28,101	12,332	16,318	26,206	42,176	27.2	3.8	21.5
1947-1951	2,793	28,333	12,399	16,259	26,466	42,707	27.4	4.3	21.4
1948-1952	2,698	28,289	12,153	16,218	26,521	42,679	28.1	4.0	21
1949-1953	2,636	28,545	12,562	16,347	26,443	43,323	29.1	4.4	19.8
1950-1954	2,616	28,977	12,908	16,602	26,738	44,495	29.4	5.0	18.8
1951-1955	2,684	29,100	13,405	16,258	26,781	44,891	31.1	5.2	18.1
1952-1956	2,733	29,251	13,729	16,156	26,797	45,593	31.3	4.9	17.2
1953-1957	2,811	29,377	13,990	15,802	26,831	46,141	32.2	5.2	17.0
1954-1958	2,874	29,423	14,171	15,226	26,932	46,413	32.7	5.1	17.1
1955-1959	2,934	29,230	14,320	14,702	26,570	46,524	33.7	5.3	17.5
1956-1960	2,934	29,423	14,604	14,541	26,519	47,530	33.2	5.9	18.5
1957-1961	2,990	29,217	14,731	13,942	26,243	47,707	34.4	6.2	19.3
1958-1962	3,115	29,310	14,750	14,062	26,252	47,845	34.1	6.3	19.2
1959-1963	3,246	28,924	14,850	13,396	25,879	47,689	34.3	6.4	19.9
1960-1964	3,429	28,801	14,968	12,916	25,857	47,790	34.7	6.4	20.6
1961-1965	3,657	28,448	14,914	12,235	25,624	47,259	34.6	6.7	21.1
1962-1966	3,896	28,317	15,200	11,875	25,507	47,175	35.3	7.1	21.1
1963-1967	4,068	28,370	15,847	11,478	25,486	47,597	36.5	8.0	21.5
1964-1968	4,225	28,725	16,018	11,706	25,779	47,963	37.0	8.9	21.8
1965-1969	4,324	28,762	16,105	11,783	25,811	47,697	37.9	9.7	22.2
1966-1970	4,405	28,861	16,367	11,782	25,935	47,886	39.5	10.3	21.7
1967-1971	4,380	28,940	16,256	11,958	26,056	48,022	40.2	11.7	22.7
1968-1972	4,356	28,944	16,140	12,086	25,977	48,328	41.0	13.2	22.9
1969-1973	4,279	29,093	16,159	12,273	26,054	48,596	41.7	15.0	22.2
All	21,849	28,520	14,449	13,547	25,972	45,675	33.5	7.0	20.5

Note: The table reports the number of workers and summary statistics for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Only workers with positive earnings in all years are included. Annual Earnings are real (2015 price level) and gross of personal income taxes and social contributions and include income from any source. The percentage of zero earnings, women, tertiary graduates, and workers in the South of Italy are reported. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Table A.3. Earnings inequality and intragenerational mobility indices

			Cohort 1940-1944	Cohort 1955-1959	Cohort 1969-1973	% Variation
Inequality		Overall GE2	0.155	0.171	0.203	*31.0*
		Permanent GE2	0.126	0.149	0.180	*42.9*
		Avg. CS GE2	0.130	0.151	0.181	*39.2*
Biperiodical mobility indices	Relative indices	1 - ρ_n	0.203	0.136	0.157	*-22.7*
		Avg. Jump up	0.110	0.100	0.094	*-14.5*
		Avg. Jump down	-0.121	-0.095	-0.117	*-3.3*
		Pr(upper quintile)	0.263	0.225	0.252	*-4.2*
		Pr(lower quintile)	0.240	0.226	0.222	*-7.5*
		Pr(exit from bottom)	0.350	0.318	0.347	*-0.9*
		Pr(falling from top)	0.236	0.177	0.170	*-28.0*
	Absolute indices	Avg. Income growth	0.613	0.469	0.423	*-31.0*
		Median Income growth	0.134	0.093	0.058	*-56.7*
		DH Income growth	0.135	0.116	0.021	*-84.4*
Indices of dynamics		Pr(upward linear trend)	0.778	0.720	0.639	*-17.9*
		Avg. upward mobility	0.0101	0.0087	0.0069	*-31.7*
		Avg. downward mobility	0.0048	0.0024	0.0047	*-2.1*
		Avg. Individual volatility	0.0135	0.0111	0.0115	*-14.8*

Note: The table reports the earnings inequality and intragenerational mobility indices for three five-year-long cohorts of birth of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Inequality is measured through the Generalized Entropy Index of degree 2. Workers with zero earnings for at most five years are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Table A.4. Earnings inequality and intragenerational mobility indices – only positive earnings

			Cohort 1940-1944	Cohort 1955-1959	Cohort 1969-1973	% Variation
		Overall GE2	0.096	0.119	0.146	*52.1*
Inequality		Permanent GE2	0.081	0.106	0.133	*64.2*
		Avg. CS GE2	0.084	0.108	0.134	*59.5*
Biperiodical mobility indices	Relative indices	1 - ρ_n	0.147	0.109	0.108	*-26.5*
		Avg. Jump up	0.103	0.092	0.086	*-16.5*
		Avg. Jump down	-0.099	-0.086	-0.090	*-9.1*
		Pr(upper quintile)	0.234	0.223	0.225	*-3.8*
		Pr(lower quintile)	0.230	0.222	0.215	*-6.5*
		Pr(exit from bottom)	0.288	0.245	0.263	*-8.7*
		Pr(falling from top)	0.206	0.186	0.174	*-15.5*
	Absolute indices	Avg. Income growth	0.191	0.136	0.095	*-50.3*
		Median Income growth	0.130	0.087	0.061	*-53.1*
		DH Income growth	0.134	0.094	0.056	*-58.2*
Indices of dynamics		Pr(upward linear trend)	0.833	0.744	0.676	*-18.8*
		Avg. upward mobility	0.0066	0.0062	0.0051	*-22.7*
		Avg. downward mobility	0.0015	0.0009	0.0014	*-6.7*
		Avg. Individual volatility	0.0070	0.0063	0.0066	*-5.7*

Note: The table reports the earnings inequality and intragenerational mobility indices for three five-year-long cohorts of birth of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Inequality is measured through the Generalized Entropy Index of degree 2. Only workers with positive earnings every year when aged 35-45 are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Table A.5. Indices of bi-periodical mobility

Cohort	$1\text{-}\rho_n$	Jump up	Jump down	Upper quintile	Lower quintile	Exit from bottom	Falling from top	Avg. growth	Median growth	Avg. DH growth
	Relative mobility							Absolute mobility		
1940-1944	0.203	0.110	-0.121	0.263	0.240	0.350	0.236	0.613	0.134	0.135
	(0.013)	(0.003)	(0.006)	(0.010)	(0.008)	(0.020)	(0.017)	(0.115)	(0.006)	(0.014)
1941-1945	0.189	0.105	-0.119	0.260	0.234	0.336	0.245	0.552	0.134	0.119
	(0.013)	(0.003)	(0.005)	(0.011)	(0.009)	(0.019)	(0.018)	(0.077)	(0.005)	(0.014)
1942-1946	0.168	0.102	-0.112	0.243	0.221	0.310	0.227	0.913	0.141	0.121
	(0.012)	(0.004)	(0.004)	(0.008)	(0.008)	(0.015)	(0.018)	(0.221)	(0.004)	(0.014)
1943-1947	0.161	0.096	-0.112	0.247	0.218	0.304	0.232	0.895	0.153	0.125
	(0.010)	(0.003)	(0.004)	(0.008)	(0.008)	(0.017)	(0.016)	(0.190)	(0.004)	(0.013)
1944-1948	0.153	0.095	-0.106	0.244	0.221	0.301	0.223	0.993	0.171	0.139
	(0.009)	(0.003)	(0.004)	(0.007)	(0.007)	(0.014)	(0.015)	(0.335)	(0.005)	(0.011)
1945-1949	0.150	0.095	-0.102	0.236	0.219	0.295	0.212	0.951	0.177	0.151
	(0.008)	(0.003)	(0.003)	(0.009)	(0.008)	(0.017)	(0.015)	(0.306)	(0.004)	(0.012)
1946-1950	0.149	0.096	-0.101	0.232	0.223	0.300	0.205	0.933	0.176	0.161
	(0.010)	(0.003)	(0.004)	(0.009)	(0.008)	(0.016)	(0.016)	(0.286)	(0.004)	(0.010)
1947-1951	0.149	0.097	-0.101	0.224	0.217	0.294	0.196	0.604	0.172	0.163
	(0.009)	(0.003)	(0.003)	(0.008)	(0.008)	(0.014)	(0.017)	(0.256)	(0.005)	(0.011)
1948-1952	0.148	0.100	-0.101	0.227	0.222	0.298	0.193	0.716	0.15	0.141
	(0.010)	(0.004)	(0.004)	(0.008)	(0.008)	(0.015)	(0.017)	(0.280)	(0.005)	(0.012)
1949-1953	0.155	0.102	-0.103	0.224	0.225	0.295	0.184	0.600	0.126	0.111
	(0.010)	(0.003)	(0.003)	(0.007)	(0.007)	(0.017)	(0.012)	(0.149)	(0.005)	(0.013)
1950-1954	0.155	0.100	-0.104	0.228	0.221	0.303	0.184	0.545	0.102	0.084
	(0.010)	(0.003)	(0.004)	(0.009)	(0.008)	(0.019)	(0.013)	(0.187)	(0.005)	(0.014)
1951-1955	0.157	0.100	-0.106	0.234	0.228	0.330	0.182	0.545	0.086	0.061
	(0.009)	(0.003)	(0.004)	(0.008)	(0.007)	(0.018)	(0.014)	(0.218)	(0.005)	(0.011)
1952-1956	0.155	0.099	-0.107	0.234	0.226	0.337	0.182	0.444	0.074	0.041
	(0.011)	(0.004)	(0.004)	(0.007)	(0.008)	(0.020)	(0.014)	(0.137)	(0.005)	(0.012)
1953-1957	0.148	0.098	-0.102	0.224	0.218	0.319	0.166	0.337	0.081	0.060
	(0.009)	(0.003)	(0.004)	(0.007)	(0.007)	(0.016)	(0.014)	(0.070)	(0.005)	(0.014)
1954-1958	0.141	0.099	-0.098	0.227	0.222	0.322	0.173	0.451	0.087	0.087
	(0.010)	(0.004)	(0.003)	(0.007)	(0.008)	(0.018)	(0.012)	(0.142)	(0.005)	(0.012)
1955-1959	0.136	0.100	-0.095	0.225	0.226	0.318	0.177	0.469	0.093	0.116
	(0.008)	(0.003)	(0.003)	(0.007)	(0.007)	(0.018)	(0.013)	(0.110)	(0.004)	(0.010)
1956-1960	0.133	0.103	-0.091	0.224	0.226	0.319	0.167	0.719	0.104	0.148
	(0.010)	(0.004)	(0.003)	(0.006)	(0.008)	(0.015)	(0.012)	(0.284)	(0.005)	(0.011)
1957-1961	0.134	0.106	-0.090	0.227	0.232	0.321	0.171	0.833	0.114	0.177
	(0.009)	(0.004)	(0.003)	(0.007)	(0.007)	(0.016)	(0.012)	(0.284)	(0.005)	(0.012)
1958-1962	0.141	0.110	-0.091	0.231	0.235	0.330	0.171	0.967	0.114	0.185
	(0.010)	(0.004)	(0.003)	(0.006)	(0.007)	(0.019)	(0.012)	(0.306)	(0.005)	(0.011)
1959-1963	0.145	0.115	-0.092	0.232	0.241	0.346	0.163	0.959	0.116	0.195
	(0.010)	(0.004)	(0.002)	(0.006)	(0.008)	(0.019)	(0.011)	(0.340)	(0.004)	(0.012)
1960-1964	0.149	0.113	-0.096	0.238	0.242	0.356	0.169	1.008	0.128	0.196
	(0.008)	(0.004)	(0.003)	(0.006)	(0.008)	(0.017)	(0.010)	(0.215)	(0.005)	(0.010)
1961-1965	0.153	0.115	-0.097	0.240	0.243	0.359	0.166	0.798	0.134	0.190
	(0.010)	(0.003)	(0.003)	(0.007)	(0.006)	(0.016)	(0.010)	(0.133)	(0.006)	(0.012)
1962-1966	0.150	0.111	-0.098	0.238	0.240	0.348	0.151	0.732	0.139	0.179
	(0.010)	(0.003)	(0.003)	(0.006)	(0.006)	(0.015)	(0.010)	(0.115)	(0.005)	(0.010)
1963-1967	0.154	0.110	-0.100	0.237	0.235	0.338	0.161	0.669	0.144	0.168
	(0.008)	(0.003)	(0.003)	(0.006)	(0.006)	(0.012)	(0.009)	(0.099)	(0.004)	(0.009)
1964-1968	0.152	0.108	-0.102	0.241	0.239	0.335	0.168	0.567	0.137	0.141
	(0.008)	(0.002)	(0.003)	(0.006)	(0.006)	(0.014)	(0.011)	(0.060)	(0.005)	(0.010)
1965-1969	0.155	0.106	-0.105	0.244	0.244	0.344	0.176	0.510	0.119	0.111
	(0.009)	(0.002)	(0.003)	(0.006)	(0.005)	(0.013)	(0.009)	(0.055)	(0.003)	(0.010)
1966-1970	0.154	0.103	-0.107	0.248	0.240	0.351	0.184	0.487	0.102	0.084
	(0.009)	(0.002)	(0.003)	(0.006)	(0.005)	(0.014)	(0.011)	(0.056)	(0.004)	(0.009)
1967-1971	0.159	0.100	-0.111	0.249	0.237	0.349	0.188	0.513	0.084	0.062
	(0.007)	(0.002)	(0.003)	(0.006)	(0.006)	(0.015)	(0.012)	(0.082)	(0.004)	(0.009)
1968-1972	0.156	0.097	-0.114	0.250	0.230	0.356	0.172	0.414	0.068	0.034
	(0.008)	(0.002)	(0.004)	(0.006)	(0.005)	(0.016)	(0.010)	(0.069)	(0.003)	(0.009)
1969-1973	0.157	0.094	-0.117	0.252	0.222	0.347	0.170	0.423	0.058	0.021
	(0.008)	(0.002)	(0.004)	(0.007)	(0.006)	(0.015)	(0.011)	(0.071)	(0.003)	(0.010)

Note: The table reports the bi-periodical indices of intragenerational mobility for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy observed from age 35 to 45. Workers with zero earnings for at most five years are included. The standard errors in parenthesis are obtained through 100 bootstrap repetitions. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Table A.6. Indices of earnings inequality and mobility

	Inequality			Mobility			
Cohort	Overall GE2	Permanent GE2	Avg. CS GE2	Upward trend	Upward mobility	Downward mobility	Volatility
1940-1944	0.155	0.126	0.130	0.778	0.0101	0.0048	0.0135
	(0.005)	(0.005)	(0.005)	(0.009)	(0.0006)	(0.0006)	(0.0007)
1941-1945	0.152	0.124	0.128	0.779	0.0101	0.0042	0.0132
	(0.006)	(0.005)	(0.006)	(0.009)	(0.0010)	(0.0006)	(0.0007)
1942-1946	0.146	0.120	0.124	0.789	0.0106	0.0034	0.0120
	(0.005)	(0.005)	(0.005)	(0.010)	(0.0008)	(0.0003)	(0.0005)
1943-1947	0.148	0.122	0.127	0.801	0.0114	0.0032	0.0108
	(0.006)	(0.005)	(0.006)	(0.009)	(0.0010)	(0.0003)	(0.0005)
1944-1948	0.144	0.119	0.124	0.811	0.0117	0.0029	0.0104
	(0.004)	(0.004)	(0.004)	(0.008)	(0.0007)	(0.0002)	(0.0005)
1945-1949	0.141	0.116	0.122	0.818	0.0118	0.0028	0.0102
	(0.005)	(0.004)	(0.005)	(0.007)	(0.0008)	(0.0003)	(0.0005)
1946-1950	0.140	0.115	0.121	0.825	0.0120	0.0028	0.0101
	(0.005)	(0.005)	(0.005)	(0.008)	(0.0008)	(0.0003)	(0.0006)
1947-1951	0.142	0.117	0.124	0.821	0.0114	0.0028	0.0102
	(0.005)	(0.004)	(0.005)	(0.008)	(0.0007)	(0.0003)	(0.0006)
1948-1952	0.141	0.117	0.123	0.793	0.0104	0.0028	0.0105
	(0.005)	(0.004)	(0.004)	(0.009)	(0.0008)	(0.0003)	(0.0006)
1949-1953	0.147	0.123	0.128	0.763	0.0099	0.0031	0.0107
	(0.005)	(0.005)	(0.005)	(0.008)	(0.0007)	(0.0003)	(0.0005)
1950-1954	0.152	0.129	0.132	0.740	0.0095	0.0031	0.0106
	(0.004)	(0.004)	(0.004)	(0.009)	(0.0008)	(0.0003)	(0.0006)
1951-1955	0.154	0.132	0.135	0.717	0.0087	0.0031	0.0105
	(0.006)	(0.006)	(0.006)	(0.009)	(0.0005)	(0.0002)	(0.0004)
1952-1956	0.161	0.138	0.140	0.690	0.0082	0.0033	0.0107
	(0.005)	(0.005)	(0.005)	(0.008)	(0.0005)	(0.0003)	(0.0004)
1953-1957	0.165	0.143	0.145	0.697	0.0086	0.0031	0.0106
	(0.005)	(0.005)	(0.005)	(0.011)	(0.0005)	(0.0003)	(0.0004)
1954-1958	0.168	0.146	0.148	0.711	0.0086	0.0027	0.0109
	(0.006)	(0.005)	(0.005)	(0.009)	(0.0005)	(0.0002)	(0.0005)
1955-1959	0.171	0.149	0.151	0.720	0.0087	0.0024	0.0111
	(0.006)	(0.006)	(0.006)	(0.008)	(0.0005)	(0.0002)	(0.0006)
1956-1960	0.178	0.155	0.159	0.736	0.0099	0.0022	0.0111
	(0.006)	(0.006)	(0.006)	(0.009)	(0.0005)	(0.0002)	(0.0006)
1957-1961	0.183	0.159	0.164	0.758	0.0106	0.0019	0.0110
	(0.006)	(0.006)	(0.006)	(0.007)	(0.0006)	(0.0002)	(0.0007)
1958-1962	0.184	0.161	0.165	0.767	0.0106	0.0019	0.0109
	(0.006)	(0.006)	(0.006)	(0.008)	(0.0005)	(0.0002)	(0.0005)
1959-1963	0.192	0.167	0.172	0.772	0.0116	0.0019	0.0116
	(0.006)	(0.006)	(0.006)	(0.008)	(0.0006)	(0.0002)	(0.0006)
1960-1964	0.196	0.170	0.176	0.777	0.0122	0.0020	0.0117
	(0.006)	(0.005)	(0.005)	(0.008)	(0.0007)	(0.0002)	(0.0006)
1961-1965	0.200	0.172	0.178	0.777	0.0129	0.0024	0.0125
	(0.006)	(0.005)	(0.006)	(0.008)	(0.0008)	(0.0003)	(0.0008)
1962-1966	0.204	0.176	0.182	0.776	0.0128	0.0024	0.0129
	(0.006)	(0.006)	(0.006)	(0.007)	(0.0009)	(0.0003)	(0.0009)
1963-1967	0.211	0.182	0.188	0.769	0.0130	0.0026	0.0133
	(0.007)	(0.006)	(0.007)	(0.006)	(0.0008)	(0.0003)	(0.0009)
1964-1968	0.206	0.178	0.184	0.754	0.0123	0.0027	0.0128
	(0.006)	(0.006)	(0.006)	(0.007)	(0.0006)	(0.0003)	(0.0007)
1965-1969	0.206	0.180	0.184	0.729	0.0110	0.0031	0.0126
	(0.007)	(0.006)	(0.006)	(0.007)	(0.0006)	(0.0003)	(0.0007)
1966-1970	0.205	0.181	0.184	0.706	0.0092	0.0033	0.0119
	(0.007)	(0.007)	(0.007)	(0.007)	(0.0005)	(0.0004)	(0.0006)
1967-1971	0.203	0.179	0.181	0.680	0.0085	0.0040	0.0118
	(0.006)	(0.006)	(0.006)	(0.008)	(0.0004)	(0.0004)	(0.0004)
1968-1972	0.200	0.177	0.179	0.656	0.0075	0.0042	0.0114
	(0.006)	(0.006)	(0.006)	(0.007)	(0.0004)	(0.0004)	(0.0004)
1969-1973	0.203	0.180	0.181	0.639	0.0069	0.0047	0.0115
	(0.007)	(0.007)	(0.007)	(0.008)	(0.0004)	(0.0004)	(0.0005)

Note: The table reports the earnings inequality indices and the indices of dynamics of intragenerational mobility for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy observed from age 35 to 45. Inequality is measured using the Generalized Entropy Index of degree 2. The standard errors in parenthesis are obtained through 100 bootstrap repetitions. Workers with zero earnings for at most five years are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Figure A.4. Correlation between origin and destination earnings – only positive earnings

a) Cohort 1940-1944

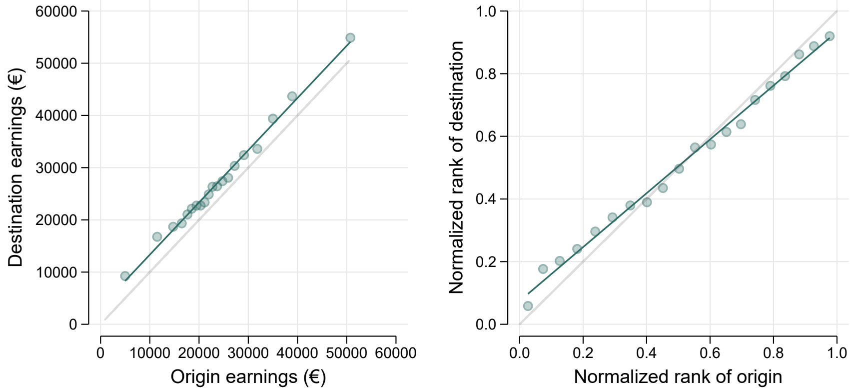

b) Cohort 1955-1959

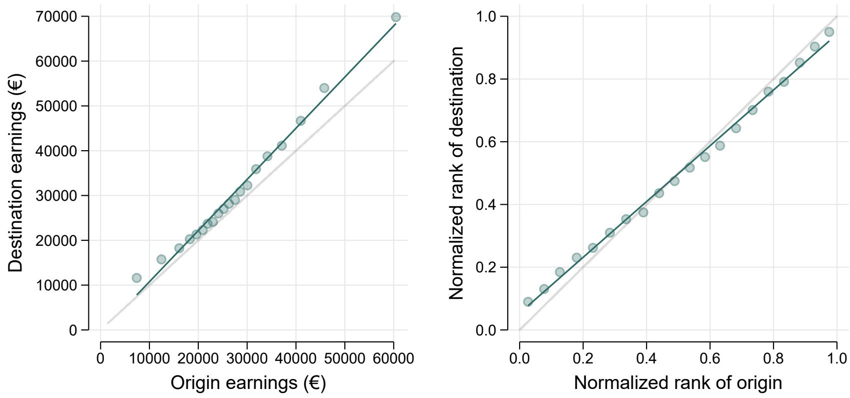

c) Cohort 1969-1973

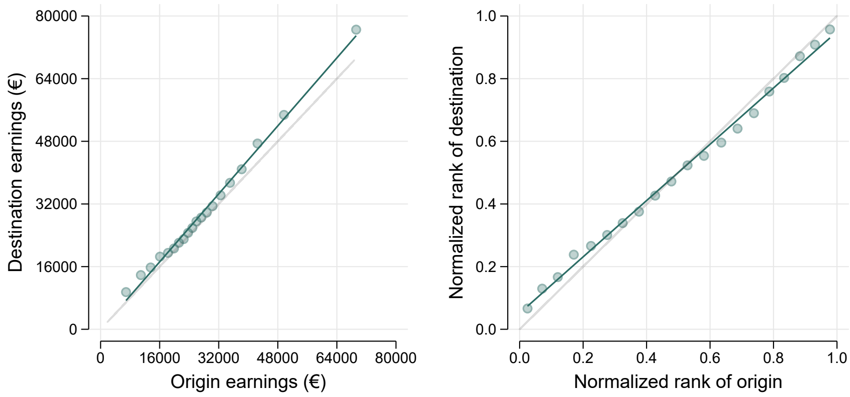

Note: The figure plots for three cohorts of birth the linear fit of destination earnings on origin earnings (left panels), and of destination income rank on origin income rank. The points are the average y-variable and x-variable inside 10 equal-sized bins. The 45-degree line is the place of perfect immobility, where destination income/rank is perfectly predicted by origin income/rank. The sample includes employees observed every year from age 35 to 45. Annual Earnings are real (2015 price level) and gross of personal income taxes and social contributions and include income from any source. Only workers with positive earnings every year when aged 35-45 are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Table A.7. Table of inequality-mobility correlation – only positive earnings

			Overall Inequality	Permanent Inequality	Avg. Inequality
Biperiodical mobility indices	Relative indices	1 - ρ_n	**-0.537**	**-0.536**	**-0.547**
		Avg. Jump up	-0.452	-0.468	-0.461
		Avg. Jump down	0.454	0.450	0.459
		Pr(upper quintile)	**-0.515**	-0.490	**-0.518**
		Pr(lower quintile)	**-0.695**	**-0.688**	**-0.694**
		Pr(exit from bottom)	-0.190	-0.176	-0.197
		Pr(falling from top)	**-0.794**	**-0.806**	**-0.795**
	Absolute indices	Avg. Income growth	-0.495	**-0.534**	-0.493
		Median Income growth	**-0.538**	**-0.577**	**-0.536**
		DH Income growth	**-0.538**	**-0.577**	**-0.537**
Indices of dynamics		Pr(upward linear trend)	**-0.682**	**-0.714**	**-0.681**
		Avg. upward mobility	-0.215	-0.259	-0.211
		Avg. downward mobility	**0.584**	**0.606**	**0.577**
		Avg. Individual volatility	**0.880**	**0.880**	**0.873**

Note: The table reports the cohort-level correlation between earnings mobility and inequality indices. All the coefficients are significant at 95% confidence level unless the number is in light grey. We highlight in bold the correlations greater or equal to 0.5. The underlying basis for computing the indices are 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Only workers with positive earnings every year when aged 35-45 are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Figure A.5. intragenerational Great Gatsby curves – only positive earnings

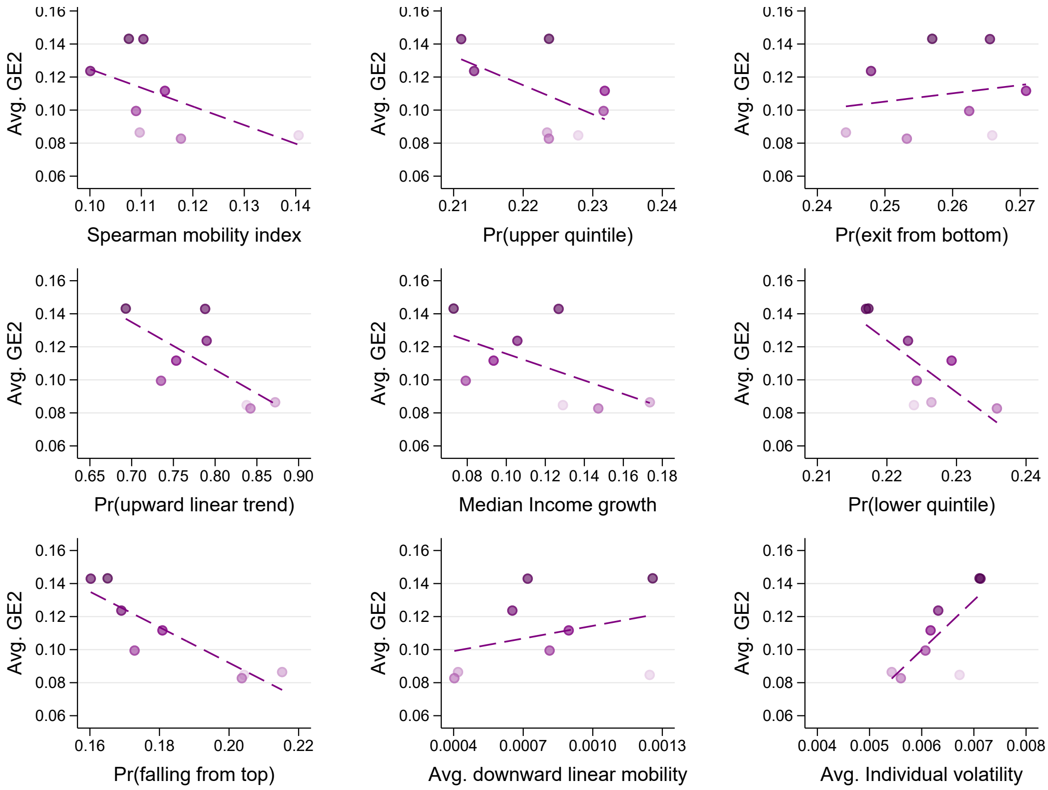

Note: The figure plots average within-cohort inequality measured through the GE2 index against several measures of intragenerational mobility. The selected measures of mobility are those with a correlation with inequality greater than 0.5 in Table 1.1. Only the cohorts of birth overlapping for one year are shown for clarity (1940-1944, 1944-1948, ..., 1968-1972), and the colour of the circle gets darker for more recent cohorts. The inequality and mobility indices are computed on a sample of employees in the private sector in Italy observed every year from age 35 to 45. Only workers with positive earnings every year when aged 35-45 are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Figure A.6. Heat map of unequal mobility – only positive earnings

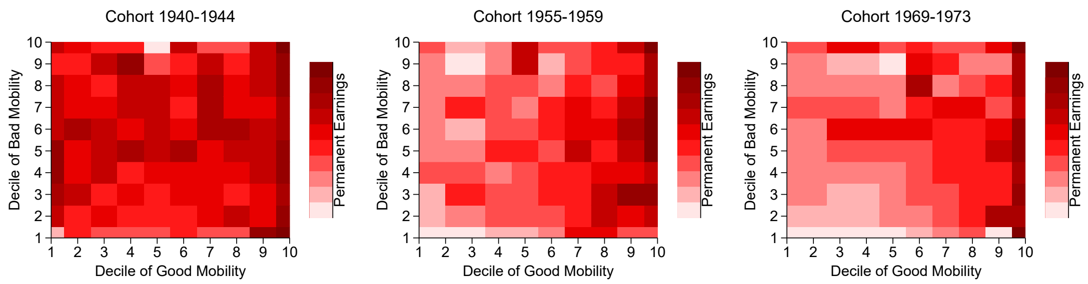

Note: The figure shows for three cohorts of birth the *heat map* of decile of permanent earnings – average income at age 35-45 – for the combination of deciles of 'good' (x-axis) and 'bad' (y-axis) mobility. Darker areas indicate a greater decile of permanent earnings. Good and bad mobility are estimated through the income risk decomposition à la Nichols described in Section 1.2.2 and measure, respectively, smooth upward income growth and individual income volatility. The sample includes employees in the private sector in Italy observed every year from age 35 to 45. Only workers with positive earnings every year when aged 35-45 are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Figure A.7. Gender differences in earnings mobility – only positive earnings

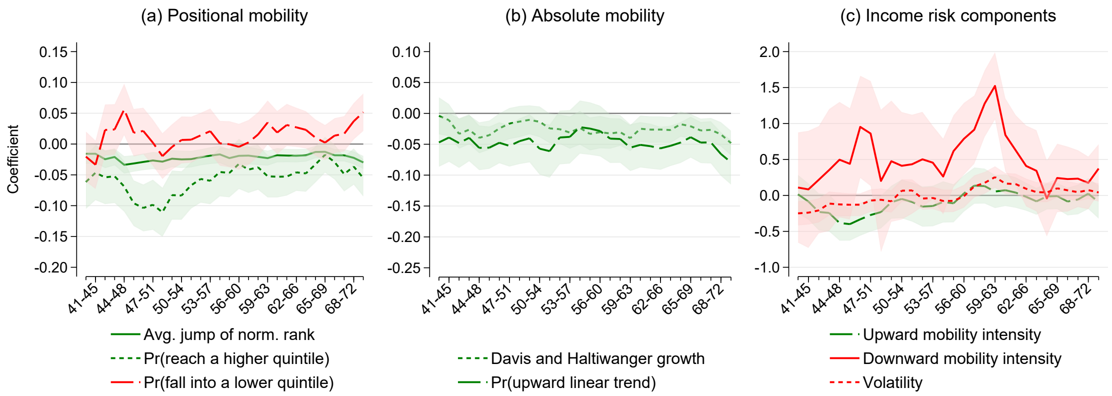

Note: The figure plots by cohort of birth the coefficient of an indicator variable for being women in several OLS linear regressions of mobility measures controlling for being a tertiary graduate, working in the South of Italy, and for the normalised rank at age 35-37. The mobility variables in panel (c) are taken in log. The regressions are fitted separately for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy observed every year from age 35 to 45. Only workers with positive earnings every year when aged 35-45 are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Figure A.8. Education differences in earnings mobility – only positive earnings

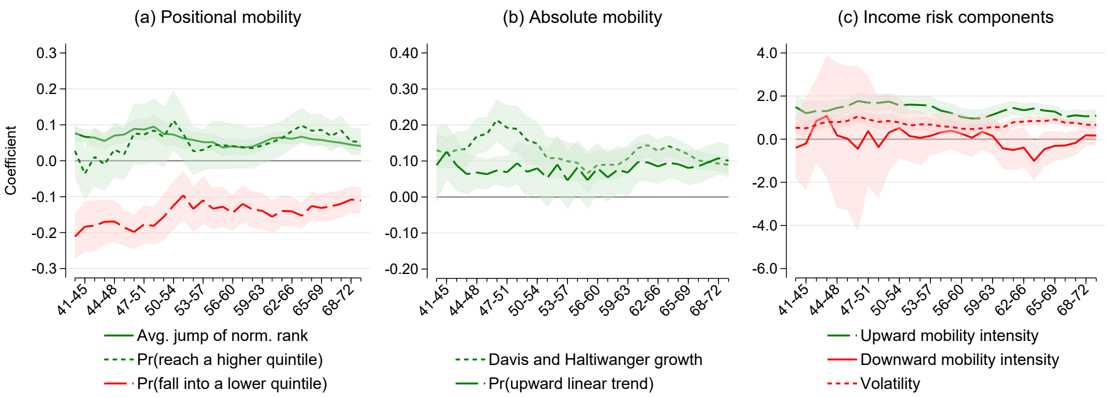

Note: The figure plots by cohort of birth the coefficient of an indicator variable for being a tertiary graduate in several OLS linear regressions of mobility measures controlling for being a woman, working in the South of Italy, and for the normalised rank at age 35-37. The mobility variables in panel (c) are taken in log. The regressions are fitted separately for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy observed every year from age 35 to 45. Only workers with positive earnings every year when aged 35-45 are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

Figure A.9. Geographical differences in earnings mobility – only positive earnings

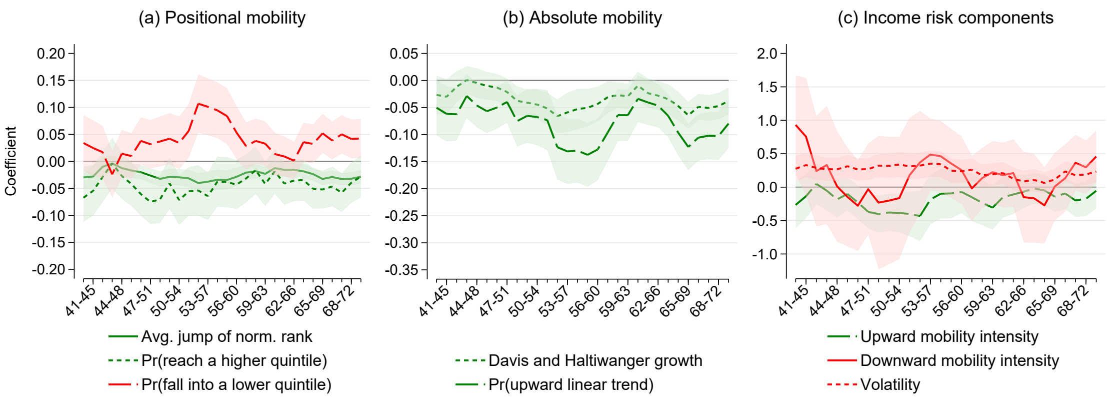

Note: The figure plots by cohort of birth the coefficient of an indicator variable for working in the South of Italy in several OLS linear regressions of mobility measures controlling for being a woman, being a tertiary graduate, and for the normalised rank at age 35-37. The mobility variables in panel (c) are taken in log. The regressions are fitted separately for 30 five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy observed every year from age 35 to 45. Only workers with positive earnings every year when aged 35-45 are included. The observations are weighted using IT-SILC sample weights. *Source*: AD-SILC data 1975-2018.

A.2 The Stata program `intramob`

In this Appendix, I show the functioning (help file and ado file) of a Stata program called `intramob` developed to allow the measurement of intragenerational mobility described in Section 1.2.2 using panel data. The program is still under review and is not yet available as a user-written command in Stata.

A.2.1 Help file

```
  ___  ____  ____  ____  ____(R)
 /__    /   ____/   /   ____/
___/   /   /___/   /   /___/
  Statistics/Data analysis
```

Intragenerational Mobility Indices

intramob - Program to measure intragenerational mobility with panel data

Syntax

intramob *varname* [*if*] [*in*] [*weight*] [, *options*]

Options

options	Description
idvar(*varname***)**	**panel** variable. If absent, data need to be **xtset**
tvar(*varname***)**	**time** variable. If absent, data need to be **xtset**
group(*varname***)**	computes the indices also by categories of *varname*
generate(*string***)**	if **yes**, asks to generate individual-level variables of mobility
restrict(*string***)**	sets that only **biperiodical** or **dynamics** indices are computed
Options for biperiodical indices	
tsmooth(*#***)**	sets how many periods are averaged for origin and destination income
nquantiles(*#***)**	sets the number of quantiles used for transition matrix; default is 5
binsvalue(*#***)**	sets the number of bins used in the binscatter of income values
binsrank(*#***)**	sets the number of bins used in the binscatter of income ranks
seed(*#***)**	to set seed for replication, needed if zeros are included
Options for indices of dynamics	
trend(*varname***)**	declares variable already containing income trend
residual(*varname***)**	declares variable already containing income residual

fweight, **aweight**, **pweight** and **iweight** are allowed; see help weights. **bootstrap** prefix is allowed; see help bootstrap.

Description

itramob computes several intragenerational income mobility indices starting from **strongly balanced** panel data in long form. The program takes the first and last periods in **tvar(***varname***)** to compute **biperiodical** indices, and all the periods to compute indices of **dynamics**. See the Remarks below for a list of the indices. It also provide a transition m d visualisation tools for absolute and relative mobility.
It works with non-negative continuous variables, so zero values are allowed.

Options

idvar(*varname***)** specifies the panel variable if data are not already **xtset**.

tvar(*varname***)** specifies the time variable if data are not already **xtset**.

group(*varname***)** in an option to compute the indices by categories of *varname*. For positional indices, the underlying ranks and quantiles are assigned using all data. It is not an index decomposition.

generate(*string***)**: *string* can be *yes* or *no* (default). If *yes*, the program generates individual-level variables of mobility while computing the aggregate indices.

restrict(*string***)**: *string* can be *biperiodical* or *dynamics*; the option restricts the program to compute only one of the two categories of mobility indices listed in the Remarks.

Options for biperiodical indices

tsmooth(*#***)** specifies the number of periods used to smooth origin and destination income; naming *s* the argument of **tsmooth(***#***)** and T the last period, for each unit in **idvar(***varname***)** origin income is defined as the average from period 1 to period 1 + *s*, while destination income is the average from period T - *s* to period T.

nquantiles(*#***)** specifies the number of quantiles for the positional indices based on the transition matrix. The number of quantiles is the same for origin and destination income.

<u>bins</u>value(#): if this option is specified, the program produces a scatterplot named *binsvalue*.gph of the average income of origin and destination in # equal-sized bins, a linear fit, and the 45 degrees line for comparison with the perfect (absolute) immobility benchmark.

<u>bins</u>rank(#): if this option is specified, the program produces a scatterplot named *binsrank*.gph of the average normalized rank of income of origin and destination in # equal-sized bins, a linear fit, and the 45 degrees line for comparison with the perfect (relative) immobility benchmark. If both **binsvalue()** and **binsrank()** options are spcificed, the program produces a combined graph named *binscatter*.gph.

seed(#) allows to set seed for replication in case there are zero values in the data; only for positional indices, the program adds random numbers from a uniform distribution to zeros to allow ranking. There is no need to set seed if the option **restrict(dynamics)** is specified because the randon numbers generation is need only for the biperiodical category of indices.

<u>Options for indices of dynamics</u>

trend(*varname***)**: if this option is not specified, the program computes the indices of dynamics by fitting a liner trend with local regressions at individual level. Since this procedure takes time for large samples, if you need to repeat the computations (e.g. for bootstrap inference) you can generate the demeaned trend once using the **gen(yes)** option, and then run the program specifying which variable contains the trend and save time.

residual(*varname***)**: as in the previous option, but for the residual from demeaned detrended income.

<u>Remarks: list of indices</u>

The program computes the indices of intragenerational mobility employed in Subioli and Raitano (2022) allowing for the inclusion of zeros.

Biperiodical indices of relative mobility: the Spearman Mobility Index, the Average Jump Up and the Average Jump Down are base on **normalized ranks** ([0;1]) of origin and destination. The Probability of Going Up, the Probability of Going Down, the Probability of Exit from the Bottom, the Probability of Falling from the Top are based on the **transition matrix** between origin and destination **quantiles**. The default is quintiles and can be changed using the **nq(#)** option.

Biperiodical indices of absolute mobility: the Average and Median Income Growth are based on a standard growth rate y2/y1-1, while the Davis and Haltiwanger Average Growth is based on the bounded ([-2;2]) growth rate (y2-y1)/((y2+y1)/2).

Indices of dynamics: the indices of dynamics are based on the 'income risk decomposition' proposed by Nichols (2008) as applied in Subioli and Raitano (2022). The program computes the Overall General Entropy Index of degree 2 and its three components: the Permanent GE2, the Upward Mobility Risk, the Downward Mobility Risk, and Volatility. It also adds the Share of Upward Trends, and the Average Cross-sectional GE2 for comparison.

<u>Stored results</u>

Scalars

r(spearman) Spearman Mobility Index (with normalized ranks)

r(jump_up) Avg. jump up of normalized ranks

r(jump_down) Avg. jump down of normalized ranks

r(prob_up) Pr(reaching a higher quantile)

r(prob_down) Pr(falling into a lower quantile)

r(exit_bottom) Pr(exit from the bottom quantile)

r(exit_top) Pr(falling from the top quantile)

r(avg_growth) Avg. income growth

r(med_growth) Median income growth

r(dh_growth) Avg. Davis and Haltiwanger income growth

r(uptrend_share) Share of upward linear trends

r(mobrisk_up) Upward mobility risk

r(mobrisk_down) Downward mobility risk

r(volatility) Volatility

r(avgcs_ge2) Average (unweighted) cross-sectional GE2

r(overall_ge2) Overall GE2

r(perm_ge2) Permanent GE2

If the **group(***varname***)** option is used, the program generates for each index a vector of dimension 1 x *k* for *k* categories of *varname* in **group**

Examples

Generate longitudinal income data

```
. clear

. set obs 100

. set seed 010101

. egen workerid = seq(), from(1) to(100)

. forvalues y = 1/3 {
2. generate income`y' = rnormal(30000, 5000)
3. }

. reshape long income, i(worker) j(year)

. xtset worker year
```

Compute all mobility indices producing and saving binscatter plots with 20 bins

```
. intramob income, binsrank(20) binsvalue(20) seed(010101)

. graph save "binscatter" mygraph.gph
```

Compute only biperiodical indices generating individual-level values of mobility and using 3-period smoothing

```
. intramob income, tsmooth(3) gen(yes) seed(010101)
```

Author

Francesca Subioli
Department of Economics and Law, Sapienza University of Rome, Italy
Contact: francesca.subioli@uniroma1.it

References

Subioli, F. and M. Raitano (2022), *Differences set in stone: evidence on the inequality-mobility trade off in Italy*.

Jäntti, M., and S. P. Jenkins (2015), *Income mobility. In Handbook of income distribution*, Elsevier, Vol. 2, pp. 807-935.

Nichols, A. et al. (2008). Trends in income inequality, volatility, and mobility risk. Technical report, IRISS at CEPS/INSTEAD.

Davis, S. J. and J. Haltiwanger (1992). Gross job creation, gross job destruction, and em- ployment reallocation. The Quarterly Journal of Economics 107(3), 819–863.

A.2.2 Ado file

intramob* - Printed on 30/11/2022 17:06:05

```
*** ADO FILE FOR INDICES OF MOBILITY ***
* November 2022
* Author Francesca Subioli

cap prog drop intramob
program define intramob, rclass sortpreserve properties(svyb svyj)
    syntax varname(numeric) [if] [in] [pw aw fw iw], ///
    [IDvar(varname numeric) ///
    Tvar(varname numeric) ///
    TSMOOTH(real 1) ///
    RESTrict(string) ///
    GENerate(string) ///
    NQuantiles(real 5) ///
    BINSvalue(real -1) ///
    BINSrank(real -1) ///
    GROUP(varname numeric) ///
    SEED(real -1) ///
    TREND(varname numeric) ///
    RESidual(varname numeric)]

    // if/in:

    marksample touse

    /////////////////////////////
    ////   Validity checks   ////
    /////////////////////////////

    // id and time variables

    if ("`idvar'" != "") {
        if ("`tvar'" == "") {
            di as error "Define time variable tvar()"
            exit 198
        }
        else {
            local id `idvar'
            local t `tvar'
        }
    }
    else {
        _xt, trequired
        local id `r(ivar)'
        local t `r(tvar)'
    }

    // Balanced panel

    tempvar N
    qui: bys `id' `touse': gen `N' = _N if `touse'
    qui: sum `N', meanonly
    local tmax = `r(max)'
    if (`r(min)' < `tmax') {
        display as error "The panel must be balanced"
        exit 198
    }

    // Periods of smoothing

    if ("`restrict'" != "" & "`restrict'" != "biperiodical" ///
        & "`restrict'" != "dynamics") {
            display as error "Argument of restrict() not allowed"
            exit 198
    }

    if ("`restrict'" == "" | "`restrict'" == "dynamics") {
        // ensure at least three periods per person
        if (`tmax' < 3) {
```

```
            display as error ///
            "At least three periods required for dynamics indices"
            exit 198
        }
    }
    else if ("`restrict'" == "biperiodical") {
        // ensure at least two periods per person
        if (`tmax' < 2) {
            display as error ///
            "At least two periods required for biperiodical indices"
            exit 198
        }
    }
    if ("`restrict'" != "dynamics") {
        if (`tsmooth' < 0) {
                display as error "Negative tsmooth() not allowed"
                exit 198
            }
        else {
            if (`tsmooth' > (`tmax'-1)) {
                display as error "tsmooth() too big"
                exit 198
            }
        }
    }

    // Missing values

    qui: count if `touse' & `varlist' == .
    if (`r(N)' != 0) {
        display as error "Missing values not allowed"
        exit 198
    }

    // Negative values

    qui: count if `touse' & `varlist' < 0
    if (`r(N)' != 0) {
        display as error "Negative values not allowed"
        exit 198
    }

    // Zeros

    qui: count if `touse' & `varlist' == 0
    if (`r(N)' != 0) {
        display as text "Warning: `r(N)' zero values used in the computations"
        if ("`restrict'" != "dynamics" & `seed' != -1) set seed `seed'
    }


    // Sample weights

    if ("`weight'" == "") {
        tempvar wvar
        qui: gen byte `wvar' = 1 if `touse'
    }
    else {
        tempvar wvar
        qui: gen `wvar' `exp' if `touse'
    }

    // Variables to generate

    if ("`generate'" != "" & "`generate'" != "yes" & "`generate'" != "no") {
            display as error "Argument of generate() not allowed"
            exit 198
    }
```

intramob* - Printed on 30/11/2022 17:06:05

```
if ("`generate'" == "") local generate no
if ("`generate'" == "yes") {
    foreach var in nrank_origin nrank_destination quant_origin ///
                    quant_destination nrank_jump stepped_down stepped_up ///
                    rel_growth upward_trend up_mobrisk down_mobrisk ///
                    dem_trend det_residual volatility {
        capture confirm var `var'
        if !_rc {
            di as error "Variable `var' already exists"
            exit 198
        }
    }
}

// Grouping variable

if ("`group'" != "") {
    if `touse' & `group' != `group'[_n-1] & `id' == `id'[_n-1] {
        display as error "Variable in group() must be time-invariant"
        exit 198
    }
}

//////////////////////
////   Commands   ////
//////////////////////
local y `varlist'

if ("`restrict'" == "" | "`restrict'" == "dynamics") {

////////////////////////
/// Dynamics indices

sum `y' [aw = `wvar'] if `touse', meanonly
local ybar = `r(mean)' // overall mean
local N = `r(N)' // number of observations

tempvar yibar
qui: bys `id': egen `yibar' = mean(`y') if `touse'

if ("`trend'" != "" & "`residual'" != "") {
    local ytrend `trend'
    local yres `residual'
}
else {
    tempvar ytrend
    qui: gen `ytrend' = .
    tempvar yres
    qui: gen `yres' = .
    tempvar demeanedy
    qui: gen `demeanedy' = `y' - `yibar' if `touse'
    tempvar c_time
    qui: by `id': center `t' if `touse', gen(`c_time')
    qui: {
    levelsof `id' if `touse', local(levels)
        foreach l of local levels {
            reg `demeanedy' `c_time' if `id' == `l'
            replace `ytrend' = e(b)[1,1]*`c_time' if `id' == `l'
            replace `yres' = `demeanedy' - `ytrend' if `id' == `l'
        }
    }
    if ("`generate'" == "yes") {
        qui: gen dem_trend = `ytrend'
        label var dem_trend "Demeaned individual income trend"
        qui: gen det_residual = `yres'
        label var det_residual "Demeaned detrended individual income residual"
    }
}
```

Page 3

intramob* - Printed on 30/11/2022 17:06:05

```

    // Overall GE2

    tempvar x1
    qui: gen `x1' = (`y'-`ybar')^2 if `touse'
    sum `x1' [aw = `wvar'] if `touse', meanonly
    local overall_ge2 = `r(mean)' / (2*(`ybar'^2))
    mat overall_ge2 = `overall_ge2'

    // Permanent GE2

    tempvar x2
    qui: gen `x2' = (`yibar'-`ybar')^2 if `touse'
    sum `x2' [aw = `wvar'] if `touse', meanonly
    local perm_ge2 = `r(mean)' / (2*(`ybar'^2))
    mat perm_ge2 = `perm_ge2'

    // Average cross-sectional GE2

    local csineq = 0
    qui: levelsof `t', local(levelst)
    local periods = `r(r)'
    foreach l of local levelst {
        sum `y' [aw = `wvar'] if `touse' & `t' == `l', meanonly
        local ybar`l' = `r(mean)'
        tempvar x2`l'
        qui: gen `x2`l'' = (`yibar'-`ybar`l'')^2 if `touse' & `t' == `l'
        sum `x2`l'' [aw = `wvar'] if `touse' & `t' == `l', meanonly
        local avgcs_ge2`l' = `r(mean)' / (2*(`ybar`l''^2))
        local avgcs_ge2 = `avgcs_ge2' + `avgcs_ge2`l''
    }
    local avgcs_ge2 = `avgcs_ge2' / `periods'
    mat avgcs_ge2 = `avgcs_ge2'

    // Mobility Risk

    sort `touse' `id' `t'
    tempvar x3
    qui: gen `x3' = (`id' == `id'[_n+1] & `ytrend'<`ytrend'[_n+1]) if `touse'
    tempvar uptrend
    qui: bys `id': egen `uptrend' = max(`x3') if `touse'
    if ("`generate'" == "yes") {
        qui: gen upward_trend = `uptrend'  if `touse'
        label var upward_trend "Indicator for upward linear trend"
    }
    sum `uptrend' [aw = `wvar'] if `touse', meanonly
    local uptrend_share = `r(mean)'
    tempvar trend2
    qui: gen `trend2' = (`ytrend'^2) if `touse'
    if ("`generate'" == "yes") {
        qui: bys `id': egen up_mobrisk = mean(`trend2') if `touse' & `uptrend'
        qui: bys `id': egen down_mobrisk = mean(`trend2') if `touse' & !`uptrend'
        qui: replace up_mobrisk = up_mobrisk / (2*(`ybar'^2))
        label var up_mobrisk "Intensity of upward linear mobility"
        qui: replace down_mobrisk = down_mobrisk / (2*(`ybar'^2))
        label var down_mobrisk "Intensity of downward linear mobility"
    }
    qui: replace `trend2' = `trend2' * `wvar' if `touse'
    tempvar sumSW
    egen `sumSW' = sum(`wvar') if `touse'
    tempvar mobup
    sum `trend2' if `uptrend' & `touse', meanonly
    local mobrisk_up = `r(sum)' / `sumSW' / (2*(`ybar'^2))
    tempvar mobdown
    sum `trend2' if !`uptrend' & `touse', meanonly
    local mobrisk_down = `r(sum)' / `sumSW' / (2*(`ybar'^2))
    matrix uptrend_share = `uptrend_share'
    matrix mobrisk_up = `mobrisk_up'
```

intramob* - Printed on 30/11/2022 17:06:05

```
    matrix mobrisk_down = `mobrisk_down'
    if ("`group'" != "") {
        foreach l of local levelsgroup {
            sum `uptrend' if `touse' & `group' == `l' [aw = `wvar'], meanonly
            matrix uptrend_share = uptrend_share, `r(mean)'
            sum `trend2' if `uptrend' & `touse' & `group' == `l' [aw = `wvar'], meanonly
            local wev = `r(sum)' / `sumSW' / (2*(`ybar'^2))
            matrix mobrisk_up = mobrisk_up, `wev'
            sum `trend2' if !`uptrend' & `touse' & `group' == `l' [aw = `wvar'], meanonly
            local wev = `r(sum)' / `sumSW' / (2*(`ybar'^2))
            matrix mobrisk_down = mobrisk_down, `wev'
        }
    }

    // Volatility

    tempvar res
    qui: gen `res' = (`yres'^2) if `touse'
    sum `res' [aw = `wvar'] if `touse', meanonly
    local volatility = `r(mean)' / (2*(`ybar'^2))
    matrix volatility = `volatility'
    if ("`group'" != "") {
        foreach l of local levelsgroup {
            sum `res' if `touse' & `group' == `l' [aw = `wvar'], meanonly
            local wev =  `r(mean)' / (2*(`ybar'^2))
            matrix volatility = volatility, `wev'
        }
    }
    if ("`generate'" == "yes") {
        qui: bys `id': egen volatility = mean(`res') if `touse'
        qui: replace volatility = volatility / (2*(`ybar'^2))
        label var volatility "Intensity of individual volatility"
    }
    }

    ///////////////////////
    /// Biperiodical indices

    if ("`restrict'" == "" | "`restrict'" == "biperiodical") {

    local wev = 1 + `tsmooth'
    local wev2 = `tmax' - `tsmooth'
    tempvar period
    qui: bys `id' `touse': gen `period' = _n if `touse'

    // Origin income
    tempvar y1
    qui: bys `id': egen `y1' = mean(`y') if `touse' & `period' >= 1 & `period' <= `wev'

    // Destination income
    tempvar y2
    qui: bys `id': egen `y2' = mean(`y') if `touse' & `period' <= `tmax' & `period' >= `wev2'

    tempfile originaldata
    qui: save "`originaldata'", replace
    qui: keep if `touse'
    collapse (min) `y1' `y2' (mean) `wvar' `group', by(`id')

    // Income growth (%), also using Davis and Haltiwanger growth rate

    tempvar growth
    qui: gen `growth' = (`y2'/`y1') -1
    qui: replace `growth' = 1 if `y1' == 0 & `y2' != 0
    qui: replace `growth' = -1 if `y2' == 0 & `y1' != 0
    if ("`generate'" == "yes") {
        qui: gen inc_growth = `growth'
        label var inc_growth "Individual income growth y2/y1-1"
    }
```

```
        qui: sum `growth' [aw = `wvar'], det
        mat avg_growth = `r(mean)'
        mat med_growth = `r(p50)'
        tempvar dhgrowth
        qui: gen `dhgrowth' = (`y2' - `y1') / ((`y2' + `y1') / 2)
        sum `dhgrowth' [aw = `wvar'], meanonly
        mat dh_growth = `r(mean)'
        if ("`generate'" == "yes") {
            qui: gen dh_growth = `dhgrowth'
            label var dh_growth "Individual income growth (y2-y1)((y2+y1)/2)"
        }
        if ("`group'" != "") {
            qui: levelsof `group', local(levelsgroup)
            foreach l of local levelsgroup {
                qui: sum `growth' if `group' == `l' [aw = `wvar'], det
                matrix avg_growth = avg_growth, `r(mean)'
                matrix med_growth = med_growth, `r(p50)'
                sum `dhgrowth' if `group' == `l' [aw = `wvar'], meanonly
                matrix dh_growth = dh_growth, `r(mean)'
            }
        }
        if (`binsvalue' != -1) {
            binscatter `y2' `y1' [aw = `wvar'], ///
                n(`binsvalue') xlabel(#8, grid format(%9.0f)) ylabel(#8, ///
                format(%9.0f)) xtitle(Origin `y') ///
                ytitle(Destination `y') ///
                title(Binscatter (`binsvalue' bins), size(medsmall)) ///
                yscale(noextend) xscale(noextend titlegap(*4)) ///
                mcol(emerald%30) lcol(emerald) xsize(4.5) ysize(4) ///
                name(binsvalue, replace)
                tempvar wev
                qui: xtile `wev' = `y1' [aw = `wvar'], nq(`binsvalue')
                qui: sum `y1' [aw = `wvar'] if `wev'==`binsvalue', meanonly
            graph addplot line `y1' `y1' if `y1' <= `r(mean)', lcol(gray%20*1.5)
        }

        // Spearman mobility index with nomalized ranks

        forvalues i = 1/2 {
            // random numbers to sort the zeros (if any)
            qui: count if `y`i'' == 0
            qui: if (`r(N)' != 0) replace `y`i'' = `y`i'' + runiform(0,1) if `y`i''==0
            sort `y`i''
            tempvar rank`i'
            qui: gen `rank`i'' = (_n-1)/(_N-1) // normalized rank
        }
        qui: reg `rank2' `rank1' [aw = `wvar']
        local spearman = 1-e(b)[1,1]
        mat spearman = `spearman'
        if ("`group'" != "") {
            foreach l of local levelsgroup {
                qui: reg `rank2' `rank1' if `group' == `l' [aw = `wvar']
                local spearman = 1-e(b)[1,1]
                mat spearman = spearman, `spearman'
            }
        }
        if (`binsrank' != -1) {
            binscatter `rank2' `rank1' [aw = `wvar'], ///
                n(`binsrank') xlabel(0(.1)1, grid format(%2.1f)) ylabel(0(.1)1, ///
                format(%2.1f)) xtitle(Normalized rank of origin) ///
                ytitle(Normalized rank of destination) ///
                title(Binscatter (`binsrank' bins), size(medsmall)) ///
                yscale(noextend) xscale(noextend titlegap(*4)) ///
                mcol(emerald%30) lcol(emerald) xsize(4.5) ysize(4) ///
                name(binsrank, replace)
            graph addplot line `rank1' `rank1', lcol(gray%20*1.5)
        }
        if (`binsvalue' != -1 & `binsrank' != -1) ///
```

intramob* - Printed on 30/11/2022 17:06:05

```
    graph combine binsvalue binsrank, rows(1) xsize(4) ysize(2) iscale(1.2) ///
    name(binscatter, replace)

    // Average normalized rank jump up and down

    tempvar jump
    qui: gen `jump' = `rank2' - `rank1'
    if ("`generate'" == "yes") {
        qui: gen nrank_origin = `rank1'
        label var nrank_origin "Normalized rank of origin"
        qui: gen nrank_destination = `rank2'
        label var nrank_destination "Normalized rank of destination"
        qui: gen nrank_jump = `jump'
        label var nrank_jump "Jump of normalized rank from origin to destination"
    }
    sum `jump' if `jump' > 0 [aw = `wvar'], meanonly
    mat jump_up = `r(mean)'
    sum `jump' if `jump' < 0 [aw = `wvar'], meanonly
    mat jump_down = `r(mean)'
    if ("`group'" != "") {
        foreach l of local levelsgroup {
            sum `jump' if `group' == `l' & `jump' > 0 [aw = `wvar'], meanonly
            matrix jump_up = jump_up, `r(mean)'
            sum `jump' if `group' == `l' & `jump' < 0 [aw = `wvar'], meanonly
            matrix jump_down = jump_down, `r(mean)'
        }
    }

    // Normalized probability of moving up or down from quantile

    forvalues i = 1/2 {
        tempvar quant`i'
        qui: xtile `quant`i'' = (`y`i'') [aw = `wvar'], nquantiles(`nquantiles')
    }
    qui: gen origin = `quant1'
    qui: gen destination = `quant2'
    display ""
    display as text "Transition matrix"
    tab origin destination [aw = `wvar'], row nofreq

    tempvar up
    qui: gen `up' = (destination > origin)
    sum `up' if origin != `nquantiles' [aw = `wvar'], meanonly
    mat prob_up = `r(mean)'
    tempvar down
    qui: gen `down' = (destination < origin)
    sum `down' if origin != 1 [aw = `wvar'], meanonly
    mat prob_down = `r(mean)'
    if ("`group'" != "") {
        foreach l of local levelsgroup {
            sum `up' if `group' == `l' & origin != `nquantiles' [aw = `wvar'], meanonly
            matrix prob_up = prob_up, `r(mean)'
            sum `down' if `group' == `l' & origin != 1 [aw = `wvar'], meanonly
            matrix prob_down = prob_down, `r(mean)'
        }
    }
    if ("`generate'" == "yes") {
        qui: gen quant_origin = `quant1'
        label var quant_origin "Quantile of origin"
        qui: gen quant_destination = `quant2'
        label var quant_destination "Quantile of destination"
        qui: gen stepped_up = `up'
        label var stepped_up "Indicator for quantile rise from origin to destination"
        qui: gen stepped_down = `down'
        label var stepped_down "Indicator for quantile drop from origin to destination"
    }

    // Normalized probability of leaving bottom and top quantiles
```

Page 7

intramob* - Printed on 30/11/2022 17:06:05

```

    tempvar exitb
    qui: gen `exitb' = (origin == 1 & destination != 1)
    sum `exitb' if origin == 1 [aw = `wvar'], meanonly
    mat exit_bottom = `r(mean)'
    tempvar exitt
    qui: gen `exitt' = (origin == 5 & destination != 5)
    sum `exitt' if origin == 5 [aw = `wvar'], meanonly
    mat exit_top = `r(mean)'
    if ("`group'" != "") {
        foreach l of local levelsgroup {
            sum `exitb' if `group' == `l' & origin == 1 [aw = `wvar'], meanonly
            matrix exit_bottom = exit_bottom, `r(mean)'
            sum `exitt' if `group' == `l' & origin == 5 [aw = `wvar'], meanonly
            matrix exit_top = exit_top, `r(mean)'
        }
    }
    if ("`generate'" == "yes") {
        keep `id' nrank_origin nrank_destination quant_origin ///
                        quant_destination nrank_jump stepped_down stepped_up ///
                        inc_growth dh_growth
        tempfile mobilitydata
        qui: save "`mobilitydata'", replace
    }
    drop _all
    use "`originaldata'"
    if ("`generate'" == "yes") qui: merge m:1 `id' using "`mobilitydata'", nogen
    }

    // Display and store results

    if ("`restrict'" == "") local listresults spearman jump_up jump_down ///
                        prob_up prob_down exit_bottom exit_top avg_growth ///
                        med_growth dh_growth overall_ge2 perm_ge2 avgcs_ge2 uptrend_share ///
                        mobrisk_up mobrisk_down volatility
    else if ("`restrict'" == "biperiodical") local listresults spearman ///
                        jump_up jump_down prob_up prob_down exit_bottom exit_top ///
                        avg_growth med_growth dh_growth
    else if ("`restrict'" == "dynamics") local listresults overall_ge2 avgcs_ge2 ///
                        perm_ge2 uptrend_share mobrisk_up mobrisk_down volatility

    display ""
    display ""
    foreach result in `listresults' {
        if ("`result'" == "spearman") local text Spearman mobility index (with normalized ranks)
        if ("`result'" == "jump_up") local text Avg jump up of normalized ranks
        if ("`result'" == "jump_down") local text Avg jump down of normalized ranks
        if ("`result'" == "prob_up") local text Pr(reaching a higher quantile)
        if ("`result'" == "prob_down") local text Pr(falling into a lower quantile)
        if ("`result'" == "exit_bottom") local text Pr(exit from the bottom quantile)
        if ("`result'" == "exit_top") local text Pr(falling from the top quantile)
        if ("`result'" == "avg_growth") local text Avg. income growth
        if ("`result'" == "med_growth") local text Median income growth
        if ("`result'" == "dh_growth") local text Avg. Davis and Haltiwanger income growth
        if ("`result'" == "uptrend_share") local text Share of upward linear trends
        if ("`result'" == "mobrisk_up") local text Upward mobility risk à la Nichols
        if ("`result'" == "mobrisk_down") local text Downward mobility risk à la Nichols
        if ("`result'" == "volatility") local text Volatility à la Nichols

        if ("`result'" != "perm_ge2" & "`result'" != "overall_ge2" & "`result'" != "avgcs_ge2") {
            display as text "`text'"
            mat list `result', noblank noheader nonames format(%10.3f)
        }
        return matrix `result' = `result'
        }
end
```

Page 8

Chapter 2

Inter-temporal income polarization

JEL Codes: D31, D63.
Keywords: Income polarization, Income mobility, Alienation, Earnings dynamics, Income inequality, Italy

Notes and Acknowledgements

This chapter is joint work with Philippe Van Kerm (University of Luxembourg and Luxembourg Institute for Socio-economic Research). I sincerely thank him for the care and welcome he gave me and for his precious mentoring, as well as for funding and making possible the three-month visit that gave rise to this work. I am also grateful to the Giacomo Brodolini Foundation for allowing us to use the AD-SILC dataset for the application, and to my supervisor Michele Raitano for the support in using it correctly. This chapter benefited from comments and suggestions on previous versions from participants at the 4th International Conference on European Studies (Milan 2022), the 7th Annual Symposium of the Society for Economic Measurement (Calgary 2022), the 37th Annual Conference of the Italian Association of Labor Economics (Salerno 2022), and members of the INEQ group at Sapienza University of Rome. Appendix B.2 contains details (ado and help files) of a Stata program we developed to make it possible for interested researchers to easily compute inter-temporal income polarization.

2.1 Introduction

Income polarization is the tendency of an income distribution to cluster around distinct income levels: individuals belonging to the same cluster share similar levels of income which are (perceived to be) clearly distinct from incomes of individuals located around other poles. The polarization of an income distribution combines elements of equality and inequality

(Esteban and Ray, 1999): intra-group homogeneity and inter-group heterogeneity are the two key features of a polarized society.

The economic literature has developed two main methods of measuring polarization based on alternative conceptualisations of polarization itself. In response to a concern about a shrinking middle class in the U.S. society at the beginning of the 1990s, one method views polarization as the hollowing out of the middle of the income distribution towards the tails. This strand of literature started with the work of Wolfson and Foster (Wolfson, 1994, 1997; Foster and Wolfson, 2010) and is referred to as *bi-polarization* – typically across both sides of the median.

Esteban and Ray developed an alternative conceptualization of income polarization as the tendency of the income distribution to concentrate around two or more poles, irrespective of where they are located (Esteban and Ray, 1994; Duclos et al., 2004; Esteban et al., 2007). They rationalize income polarization as a measure of potential conflict in a population, relating it to feelings of alienation that people feel from one another when distant in the population, and to the strength of identification that people feel toward people close by. Within this identification-alienation framework, the society is seen as made of groups which are internally homogeneous and externally heterogeneous according to a relevant characteristic: in the case of income polarization, people share similar levels of income with those belonging to the same group – with whom they feel identified – while they are far from the members of other groups – over whom they feel alienated.

Esteban and Ray (1994) refer to the Marxian theory of classes to highlight the relevance of polarization beyond inequality:

> «We begin with the obvious question: why are we interested in polarization? It is our contention that the phenomenon of polarization is closely linked to the generation of tensions, to the possibilities of articulated rebellion and revolt, and to the existence of social unrest in general.» (Esteban and Ray (1994), p. 820).

The main intuition is that polarization is related to the *effective antagonism* between individuals in a society. Such antagonism is the result of both the alienation that people feel from one another – monotonic in the absolute distance as in the classical conceptualization of inequality – and the strength of *group identity* – function of the relative size of income groups. For alienation to become effective voice, action or protest, the individual must not be alone but rather identify with others (Esteban and Ray, 1994). This makes polarization different from inequality.

Existing operationalization of the notion of polarization and empirical work have discarded any time dimension: income polarization is assessed on a snapshot of the income distribution in a cross-sectional perspective. In our view, this is unsatisfactory: we argue here that the two key ingredients of polarization—feelings of alienation and identification—are sensitive to the duration of individuals' proximity or distance. Duclos et al. (2004) point out that large-scale social unrest—strikes, demonstrations, widespread violence, revolts—are phenomena that thrive on differences but cannot exist without notions of group identity. In our view, group identity and effective antagonism are not transient phenomena: it takes time to create bonds and break them, and to develop, consolidate, and translate feelings of alienation and identification into collective action. High (income) polarization at a point in time may therefore have much fewer implications for potential conflict if individuals move around the income distribution over time than if the feelings of identification and alienation consolidate through persistent income proximity or distance.

This issue is reminiscent of the literature on poverty dynamics which distinguishes snapshot – one period – poverty from a chronic or persistent poverty status (see, among others, Calvo and Dercon, 2009; Bossert et al., 2012; Hoy et al., 2012). On similar premises, we propose a generalization of the Esteban and Ray (1994) polarization index in a temporal perspective: by introducing memory parameters of past income differences, our inter-temporal polarization index measures the concentration around poles of income *trajectories* rather than point-in-time income values. By defining proximity and distance as the closeness and the remoteness of income paths, this procedure allows the dynamics of income to mediate the identification-alienation mechanism. A parameter of memory allows to calibrate the degree of relevance of the past; setting the parameter to zero, we obtain the standard Esteban and Ray index as a limit case. This property implies that when people have no memory of their past income and of that of others, polarization can be computed on current income. On the contrary, in case people remember how much they (and others) earned in preceding periods, the resulting polarization value is based on differences between income trajectories.

We apply this longitudinal perspective to matched survey-administrative data for Italy using a cohort approach, comparing income vectors of people of the same age. We document a long-term trend of increasing alienation but decreasing identification which leads to a picture of declining polarization over time. We show that incorporating past differences affects proximity and distance differently. In the identification process, it lowers the number of people belonging to the same group, distancing from each other people who would be closer in terms of current income. Conversely, alienation between people is mitigated when computed on income trajectories, suggesting a role for income mobility in reducing long-run distances between people who would instead be far apart. The effect of allowing for memory is not the same across cohorts, as they differ in their income dynamics patterns. We also demonstrate how important it is to take into account zero earners in the measurement of polarization for the categories of the population more exposed to non-employment risk, and show that the level of effective antagonism experienced is linked to some socio-demographic characteristics.

We begin by setting notation and formally describing the original approach of Esteban and Ray (1994) and its later developments in Section 2.2. The inter-temporal polarization measure is developed in Section 2.3: we first derive and characterize a measure of interpersonal income distance between income trajectories in Section 2.3.1, and then use this model to develop our inter-temporal polarization index in Section 2.3.2. Section 2.4 discusses estimation issues. Section 2.5 provides an application to cohorts of Italian workers. Section 2.6 concludes.

2.2 The notion and measurement of polarization

Esteban and Ray (1994) (henceforth ER) model income polarization as the average effective antagonism in a society, where effective antagonism is a function of identification and alienation feelings of the members of different income groups. Initially developed for an income distribution represented by a finite set of discrete income classes and levels (Esteban and Ray, 1994), implementations of ER's model on micro-data now usually rely on continuous extensions thereof (Duclos et al., 2004; Esteban et al., 2007). We describe the discrete case for clarity of exposition but will expand on the continuous representations in the rest of the paper.

Polarization over discrete income classes With n (discrete) income groups with income levels $\{y_i\}_{i=1}^n$, ER express the antagonism T as a function of alienation A that depends on pairwise income distances $\delta(y_i, y_j)$, and identification I, that is a function of the proportion π_i of the population at the income level y_i

$$P(\pi, \mathbf{y}) = \sum_{i=1}^{n}\sum_{j=1}^{n} \pi_i \pi_j T(I(\pi_i), A(\delta(y_i, y_j))). \tag{2.1}$$

Total polarization, therefore, depends both on the distance between income groups $\delta(y_i, y_j))$ and on their relative size π_i.

Three axiomatic restrictions lead ER to model identification as the group relative numerosity itself, and alienation as the average absolute income distance, ending up, in an additive utilitarian context, with the following functional form for polarization:

$$P^{\text{ER94}}(\pi, \mathbf{y}) = K \sum_{i=1}^{n}\sum_{j=1}^{n} \pi_i^{1+\alpha} \pi_j |y_i - y_j| \tag{2.2}$$

where K is a positive constant and $\alpha \in [1, 1.6]$ is a parameter measuring the degree of sensitivity to identification of the index. The larger α, the more polarization departs from inequality by giving more weight to the group size component of the index.[1] This can be easily seen by comparing the index in Equation (2.2) and the Gini index: when $\alpha = 0$ and

[1] The boundaries for the parameter α are derived in Esteban and Ray (1994), p. 833.

$K = \frac{1}{2\mu}$ (where μ is average income), Equation (2.2) measures half the average (mean-normalized) income distance in the population, which is exactly the Gini inequality index. In summary, as explained in Duclos et al. (2004), the level of polarization depends on both the separate contributions of alienation and identification and on their co-movement. Increased alienation comes from larger income distances, while increased identification is due to population shifts from less crowded to more crowded groups. The final effect depends on the product of the two components.

From discrete groups to continuously measured incomes The original model of ER works with discrete income classes like quintile groups, assuming the income distribution to be pre-arranged in mutually exclusive groups. However, allowing identification only inside income groups causes a discontinuity problem at the boundaries: comparing two people who are very close to a threshold but on opposite sides of it leads to a violation of the assumption that the groups are internally homogeneous and externally heterogeneous. Moreover, it requires to believe that the pre-arranged grouping conforms with the psychological group identification process (Esteban et al., 2007).

To address this issue, ER propose an extension consistent with their axioms but which avoids arbitrary discretisation of incomes and discontinuities. Individuals with income y_i have a window of identification centred on y_i within which they perceive other individuals as neighbours, feeling identified with them. Each individual, therefore, has his or her own set of identification. Possibly, feelings of identification to i may be strongest with individuals j having exactly the same income ($y_i = y_j$) and weaker with individuals whose income y_j is close to the boundaries of the window. This suggests the use of weights decreasing with distance to model the identification process around individual income. This approach shifts the grouping rationale from splitting the distribution into non-overlapping groups to allowing rolling individual identification windows.

To formalize this, let $b > 0$ be an amount of money such that if an income y is within b of an income y' there is some identification between two persons earning y and y'.[2] Then, let $w(d; b)$ be a positive weighting scheme on $[0, b]$ that decreases with the distance $d \equiv |y - y'|$ and reaches zero at $d \equiv |y - y'| = b$. Outside the window bounded by $y \pm b$, the weight is always zero.

Moving from a discrete set of income groups to a rolling, individual-level identification window allows re-expressing the polarization measure over continuously distributed incomes. If F denotes the continuous income distribution function, then the extent of identification at y is given by the continuous sum of the weights over all the other income

[2] For example, if $y = 2,000$ and $b = 200$, all individuals with an income between 1,800 and 2,200 are part of y's group.

levels y':

$$I(y; F, b) = \int_{y'} w(|y - y'|; b) dF(y') \tag{2.3}$$

This framework has been adopted and further developed by Esteban et al. (2007), who measure identification as in Equation (2.3) and alienation as the average continuous distance outside the identification window

$$A(y; F) = \int_{y'} \max(|y - y'| - b, 0) dF(y') \tag{2.4}$$

Using Equation (2.3) and (2.4), total polarization for continuously measured incomes is defined as

$$P^{\text{ERG07}}(\alpha, F) = \int_y \left(\int_{y'} w(|y - y'|; b) dF(y') \right)^{\alpha} \left(\int_{y'} \max(|y - y'| - b, 0) dF(y') \right) dF(y). \tag{2.5}$$

Equation (2.5), therefore, provides a continuous version of ER's original discrete polarization index. Note that Esteban et al. (2007) model identification at income level y as the continuous sum of the weights inside the identification window defined around y, and allow alienation only *outside* the identification window, measuring income distances for alienation from the boundary b. Alternatively, Duclos et al. (2004) proposed a continuous index of the form

$$P^{\text{DER04}}(\alpha, f) = \int_y \left(\int_{y'} f(y)^{1+\alpha} f(y') |y - y'| dy' \right) dy \tag{2.6}$$

where identification experienced at income y is given by $f(y)^{\alpha}$ and alienation between individuals of incomes y and y' is given by $|y - y'|$. Averaging the product of identification and alienation over y and y' leads to (2.6). Note that Duclos et al. (2004) make no explicit reference to an identification window (and derive (2.6) from a completely different axiomatic foundation). However, an implicit identification window is indirectly introduced in the model by the estimation of f through kernel density estimation methods implemented in Duclos et al. (2004). With $f(y)$ empirically estimated by the kernel density estimator

$$\hat{f}(y) = \frac{1}{n} \sum_{i=1}^{n} \frac{1}{h} K\left(\frac{z_i - y}{h} \right)$$

the kernel function K effectively plays the role of $w(;b)$, and the kernel bandwidth h corresponds to the identification window size b.

In spite of differences in the derivation of the two indices, the key difference between $P^{\text{ERG07}}(\alpha, F)$ and $P^{\text{DER04}}(\alpha, f)$ is merely that while the former rules out alienation between any two individuals within a common identification window, the latter allows for both identification *and* alienation to be simultaneously felt between two individuals holding different incomes, albeit with identification declining in income distance and alienation increasing in income distance. The latter is therefore 'smoother' – since no discontinuity needs to be introduced at the boundaries of identification windows – and will be the starting point of our inter-temporal measure.

2.3 Inter-temporal income polarization

One limitation of ER's measurement model is the neglect of time. At the core of the model is the idea that two individuals with similar (resp. different) income feel identified (resp. alienated) to one another. Individual incomes, however, notoriously change over time. Two persons with similar income at time T—when polarization is measured—are likely to have had different incomes in the past and may have ended up around that common income 'pole' from different experiences – say, a large income drop for one and a large income rise for the other. It is natural to consider that these two individuals will not feel equally alienated or identified to one another as if they have had similar incomes for a longer time period. The degree of polarization is therefore affected by the dynamics of incomes in periods prior to T. To put it differently, given two societies with an identical distribution of income at time T, the society with higher income mobility can be expected to be less polarized.

To address this concern, we propose an extension of ER's measure of polarization incorporating a notion of inter-temporal (historical) distance in the formation of effective antagonism and its constituent notions of identification and alienation. Our conceptualization of polarization over several periods is the concentration around poles of income trajectories, rather than of income values in one period.

ER's polarization measure is the sum of effective antagonisms between all pairs of individuals in the society. Effective antagonism of a person earning y towards a person earning y' is itself a function (i) of the alienation felt vis-à-vis each other, depending on the income distance between y and y', and (ii) the strength of identification felt by y towards her own income group. Our proposed inter-temporal polarization measure starts from the same premises: polarization is the sum of effective pairwise antagonisms. However, we allow the two components of effective antagonism—alienation and identification—to depend not only on current incomes, but also on *the history of income differences* between y and y'.

2.3.1 Inter-temporal pairwise income distances

We postulate that the alienation between y and y' and the contribution of y' to y's feeling of identification depends on a combination of current and past income differences between the two individuals. More precisely, given two income trajectories over T periods (leading up to current income T) $\mathbf{y} = (y_1, y_2, ..., y_T)$ and $\mathbf{y'} = (y'_1, y'_2, ..., y'_T)$, we consider a measure of inter-temporal income distance $D(\mathbf{y}, \mathbf{y'})$ between the two as:

$$D(\mathbf{y}, \mathbf{y'}) = \sum_{t=1}^{T} \phi_t \, |y_t - y'_t| \tag{2.7}$$

with $\{\phi_t : \phi_t \geq 0\}_{t=1}^{T}$ a sequence of non-negative constants such that $\sum_{t=1}^{T} \phi_t = 1$.

The sequence of weights $\{\phi_t : \phi_t \geq 0\}_{t=1}^{T}$ reflects the relative importance of current and historical income differences in forming current (period T) feelings of alienation and identification. We make two remarks before characterising the sequence of weights $\{\phi_t\}_{t=1}^{T}$.

One, we will allow identification and alienation to be governed by different sequences of weights (and therefore different inter-temporal income distances). As we discuss below, this allows for distinct roles of past income differences in building up identification and generating alienation. For clarity, we do not introduce distinct notation at this stage.

Two, time-additivity in (absolute) income differences is consistent with a myopic cumulative absolute difference model: contemporaneous income gaps between y and y' accumulate in inter-temporal differences, independently of income differences in any other time period. At each period, agents y and y' compare their incomes and the distance accumulates over time. To be clear, this rules out inter-temporal compensation in which income differences $y_s > y'_s$ in period s could be partially 'compensated' in period t when $y_t < y'_t$. Therefore, there is no consideration of inter-temporal income smoothing in forming alienation and identification. In this respect, our approach differs from using the distance between permanent incomes as a proxy for inter-temporal distance.

Moving on, we now characterize the sequence of weights $\{\phi_t\}_{t=1}^{T}$ governing the memory process. As pointed out by Calvo and Dercon (2009) in the closely-related context of chronic poverty measurement, the choice of an inter-temporal aggregation rule requires some judgements on the relative importance of the present with respect to the past. For example, a simple unweighted sum implicitly gives the same weight to present and past periods. In our approach, the reference period for comparing incomes is the present, but past values can have relevance today through a memory process. People can have some memory of their income in previous periods, but not necessarily a full one: it is reasonable to assume that the further back in time we go, the fewer past differences matter today.

We formalize this memory process through a continuous compound discounting process through the vector of weights $\boldsymbol{\phi} = \left\{ \frac{e^{-rt}}{\sum_{t=0}^{T-1} e^{-rt}} \right\}_{t=1}^{T}$, where r is a non-negative discount rate. The inter-temporal income distance is therefore modelled as

$$D(\mathbf{y}, \mathbf{y'}; r) = \frac{\sum_{s=0}^{T-1} e^{-rs} |y_{T-s} - y'_{T-s}|}{\sum_{s=0}^{T-1} e^{-rs}} \tag{2.8}$$

such that $D : 2 \times T \mapsto R$ takes the two income vectors **y** and **y'** as input, and outputs a scalar summarizing the inter-temporal distance between the two.

The memory process For a given parameter r in Equation (2.8), the weight is decreasing going backwards in time unless $r = 0$. Since the reference period for the inter-temporal distance is the last one (T), the discount rate r is a parameter of inverse *memory*: the larger the discount rate, the faster income differences distant in the past lose their importance today.

Figure 2.1 helps visualise the discounting process plotting the sequence of discounting factors $\phi_t = (e^{-rt})/(\sum_{t=0}^{T-1} e^{-rt})$ as a function of time t for an example situation of four periods. If there is maximum memory ($r = 0$), then all the periods have the same weight

Figure 2.1. Visualising the discounting process

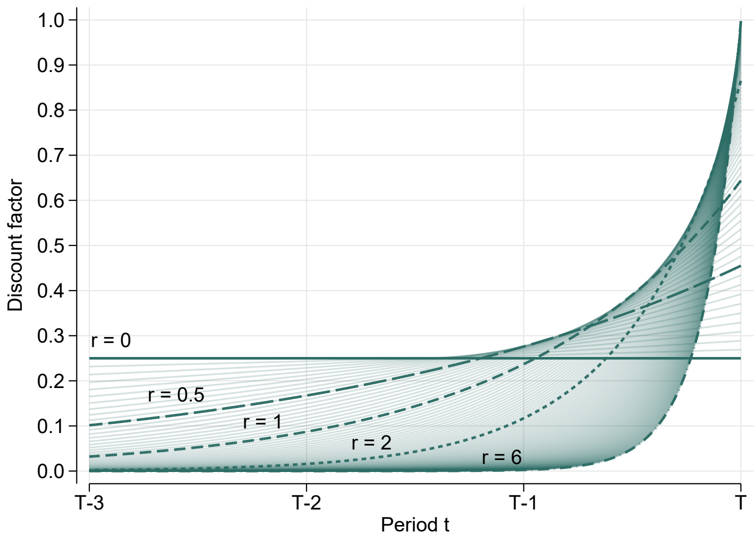

Note: The figure plots the discount factor $\phi = (e^{-rt})/(\sum_{t=0}^{T-1} e^{-rt})$ as a function of time for four periods, with T being the present. The lines are drawn using different values of the parameter of memory r, going from 0 (maximum memory) to 6 (minimum memory) by steps of 0.05.

$1/T = 0.25$, and the inter-temporal distance coincides with the simple average period distance. In all the other cases, the weight of the present period outweighs those of the past ones, and the speed of convergence to zero while going backwards in time depends on the value of the parameter r: the higher the discount rate r, the faster the convergence to zero. If there is no memory at all, a discount rate of $r = 6$ is enough to set to (practically) zero all the weights but the one for the present (0.00247 for $t = T - 1$ –one 400th of the weight at T –, 0.00000613 for $t = T - 2$ – one 163,000th of the weight at T – and so on). Intermediate values of r between zero and six allow setting different degrees of memory.

Figure 2.2 provides a further illustration of how the inter-temporal distance behaves depending on the strength of memory and on the correlation in time between the two trajectories **y** and **y**′. We see in the figure two examples of income vectors in a time frame of three periods, one in which the two income trajectories are diverging over time (upper-right panel) and one in which they are converging (lower-right panel). When r is zero (maximum memory), the inter-temporal distance coincides with the average of the period distances $|y_t - y'_t|$ since every period has the same weight $1/3$. As r increases, and the memory fades, the inter-temporal distance is closer and closer to the present one. When the two income trajectories are diverging, so that the distance in the last period is larger than the average one, the less there is memory of past closeness, the more the two trajectories are distant in a temporal perspective. On the contrary, if the two income vectors are converging, so that the distance today is lower than the distance in the past, then less memory translates into smaller inter-temporal distance.

Figure 2.2. The impact of memory on the inter-temporal distance

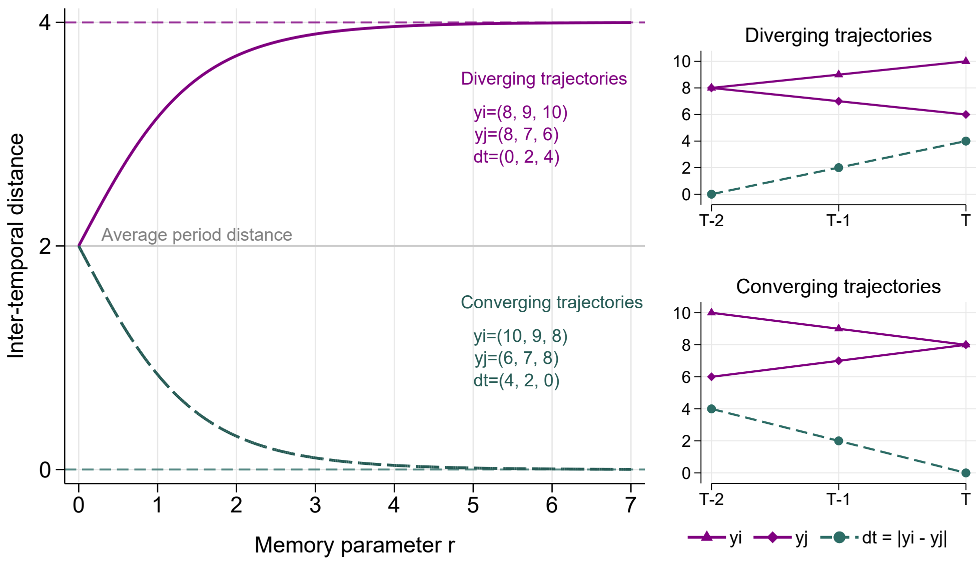

Note: On the right-hand side, the figure shows two income vectors **y** and **y’** and their period-by-period distance, separating the cases of diverging and converging trajectories. On the left-hand side, the figure plots the inter-temporal distance D (see Equation 2.8) between the two income vectors as a function of the parameter of memory r. Consider that higher r means weaker memory. The two lines represent the pattern of D when the two trajectories are diverging (solid line) and converging (dashed line).

2.3.2 Inter-temporal income polarization

Endowed with a measure for the pairwise distance between income trajectories evaluated at contemporaneous time T, we apply the identification-alienation framework as in the cross-sectional setting described in Section 2.2: we only change the notion of distance on which group membership and individual alienation are based.

Differently from the approach of Esteban et al. (2007) – and in the spirit of Duclos et al. (2004) instead – we allow alienation also inside the identification window to avoid the discontinuity at the boundaries and to make the alienation component of the index fully comparable with the Gini inequality index. Remaining in the cross-sectional context for a moment, while we take the definition of identification from Equation (2.3), alienation at y is the continuous sum of the absolute income differences with respect to the rest of the income values y':

$$A(y; F) = \int_{y'} |y - y'| dF(y') \tag{2.9}$$

Therefore, income polarization is defined as:

$$P(F; b, \alpha) = \int_{y} \left[\left(\int_{y'} w(|y - y'|; b) \right)^{\alpha} \int_{y'} |y - y'| dF(y') \right] dF(y) \tag{2.10}$$

Our strategy to incorporate the longitudinal dimension in the definition of polarization from Equation (2.10) is simply replacing the cross-sectional distance $|y - y'|$ with the inter-temporal one $D(\mathbf{y}, \mathbf{y'}; r)$ defined in Equation (2.8).

As shown in Equation (2.10), income distances are relevant for both the identification and the alienation components of antagonism: the absolute distance determines whether or not two income values belong to the same group (depending on $|y - y'| \gtreqless b$), and it is also directly the value of reciprocal alienation. We allow the discount factors' sequence $\{\phi\}_{t=1}^{T}$ to differ for the two components: we call r_I the discount rate capturing the memory for identification, and r_A that for alienation. If $r_I > r_A$, the past experience is less important—more discounted—when it comes to deciding who belongs to one's own group, rather than in determining reciprocal alienation. On the contrary, if $r_I < r_A$ the alienation component is more focused on present differences, while the proximity perceptions are more sensitive to past closeness. Of course, they may also coincide, and the choice depends on the preferred behavioural model.

Starting from identification, let $g(\mathbf{y})$ be the continuous multivariate density of income trajectories; then, the extent of identification for the income trajectory **y** is given by the continuous sum of the weights over all the other income vectors **y'**:

$$I(\mathbf{y}, G; b, r_I) = \int_{\mathbf{y'}} w(D(\mathbf{y}, \mathbf{y'}; r_I); b) dG(\mathbf{y'}) \tag{2.11}$$

Equation (2.11) gives a scalar summarizing the inter-temporal identification for the income vector **y**, given the degree of memory r_I, the identification window width b around **y**, and the weighting scheme w that assigns a positive weight to income trajectories inside b, and a weight of zero outside.

For the alienation component, we add the time dimension by computing pairwise income differences through the inter-temporal distance $D(\mathbf{x}, \mathbf{y}; r_A)$:

$$A(\mathbf{y}, G; r_A) = \int_{\mathbf{y}'} D(\mathbf{y}, \mathbf{y}'; r_A) dG(\mathbf{y}') \tag{2.12}$$

The final measure of inter-temporal polarization is the average over all income trajectories of the effective antagonisms towards the rest of the trajectory distribution, in analogy with the cross-sectional case. For each income trajectory **y**, the inter-temporal effective antagonism is defined as the product between identification to the power of α and alienation:

$$EA(\mathbf{y}, G; b, r_I, r_A, \alpha) = I(\mathbf{y}, G; b, r_I)^{\alpha} A(\mathbf{y}, G; r_A) \tag{2.13}$$

The resulting inter-temporal polarization is the average antagonism in the distribution of income trajectories, defined in Equation (2.14). To ease comparability between indices computed for different values of α, we rescale the index by raising it to the power of $1/\alpha$.

$$IP(G; b, r_I, r_A, \alpha) = \left[\int_{\mathbf{y}} \left(\int_{\mathbf{y}'} w(D(\mathbf{y}, \mathbf{y}'; r_A); b) dG(\mathbf{y}') \right)^{\alpha} \left(\int_{\mathbf{y}'} D(\mathbf{y}, \mathbf{y}'; r_A) dG(\mathbf{y}') \right) dG(\mathbf{y}) \right]^{\frac{1}{\alpha}} \tag{2.14}$$

To make the index operational, one needs to choose a domain for b and a weighting function inside it. We will discuss some of the choices for empirical analysis in the part of the paper devoted to the estimation (Section 2.4), and in the application (Section 2.5).

2.4 Estimation

2.4.1 Calculation

Assuming a sample of n individual income vectors $\{\mathbf{y}_i = y_{i1}, y_{i2}, ..., y_{iT}\}_{i=1}^{N}$, the estimation of the inter-temporal polarization index in Equation (2.14) aggregates the effective antagonism of each person in the sample. For each trajectory $\mathbf{y}_i$ in the data, we compute the pairwise inter-temporal distance with respect to all the other trajectories $\mathbf{y}_j$ in the sample as follows:

$$\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r) = \frac{\sum_{s=0}^{T-1} e^{-rs} |y_{i,T-s} - y_{j,T-s}|}{\sum_{s=0}^{T-1} e^{-rs}} \tag{2.15}$$

Since the distance $\hat{D}$ is symmetric ($\hat{D}_{ij} = \hat{D}_{ji}$), we end up with $\frac{n(n-1)}{2}$ inter-temporal distances. If the discount rate for alienation is different from that of identification, the distance in Equation (2.15) is computed twice: once using r_A and once using r_I, obtaining $n(n-1)$ values of distances. We distinguish the two distances using the notation $\hat{D}_{ij}^{I}$ for $\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r_I)$ and $\hat{D}_{ij}^{A}$ for $\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r_A)$.

The individual identification component in Equation (2.11) is estimated by using the sum of the weights of people belonging to one's group, where the group is defined as those falling inside one's identification window. Let b_i be the identification window width for person i and $w_{ij}(\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r_I); b_i)$ be the weight attached to person j depending on the inter-temporal distance between the two $\hat{D}_{ij}$ and on the identification window width of i. If $\hat{D}_{ij}$ falls inside b_i, the weight w_{ij} is positive; outside, it is zero. We discuss in Section 2.4.2 the details and implications of the weighting process.

$$\hat{I}_i(\mathbf{y}_i; b_i, r_I) = \frac{\sum_j w_{ij}(\hat{D}^I_{ij}; b_i)}{n} \tag{2.16}$$

The alienation component is computed as the average distance between the income trajectory of i and that of any other person in the population:

$$\hat{A}_i(\mathbf{y}_i; r_A) = \frac{\sum_j \hat{D}^A_{ij}}{n} \tag{2.17}$$

For each person i, the effective antagonism towards the rest of the population is given by the product of her identification to the power of α, and her average alienation from the rest of the income trajectories:

$$\hat{EA}_i(\mathbf{y}_i; \alpha, r_I, r_A, b_i) = \left[\hat{I}_i(\mathbf{y}_i; b_i, r_I)\right]^{\alpha} \hat{A}_i(\mathbf{y}_i; r_A) \tag{2.18}$$

Combining Equations (2.16), (2.17), and (2.18), the final estimate of inter-temporal polarization in Equation (2.14) is the average antagonism in the sample:

$$\hat{IP}(\alpha, r_I, r_A, b_i) = \left[\frac{\sum_i \sum_j \left(\sum_j w_{ij}(\hat{D}^I_{ij}; b_i)\right)^{\alpha} \hat{D}^A_{ij})}{n^{\alpha+2}}\right]^{\frac{1}{\alpha}} \tag{2.19}$$

The formulas above can be modified to include sample weights, and we can obtain homogeneity of degree zero (scale invariance) by mean-normalizing income in each period.

Finally, since the index in Equation (2.19) is an average, we can aggregate the antagonism at the individual level by subgroups, as long as they do not overlap and represent a complete partition of the population:

$$\hat{IP} = \left(\sum_{k=1}^{K} \hat{IP}_k\right)^{\frac{1}{\alpha}} = \left(\sum_{k=1}^{K} \frac{\sum_{i \in n_k} \hat{EA}_i}{n_k}\right)^{\frac{1}{\alpha}}, \qquad \sum_{k=1}^{K} n_k = n \tag{2.20}$$

Notice that the individual effective antagonism $\hat{EA}_i$ is computed as in Equation (2.18), thus using the whole population as the reference for identification and alienation. This implies that $\hat{IP}_k$ conveys different information from $\hat{IP}$ computed on the subset k of the population: the index computed separately for each subgroup informs on the level of inter-temporal polarization *within* each subgroup regardless the rest of the population; on the other hand, the element $I\hat{P}_k$ measures the average antagonism for group k with respect to the whole population. This decomposition allows us to compute how much of the overall antagonism can be attributed to different subgroups – of gender, level of education, occupation, geographic area etc. –, and therefore to identify 'hot spots' of potential conflict.

2.4.2 Identification weights

The width of the identification window b can be set in several ways: we group them into two main categories. One possibility is a *relative* threshold, function of one's own income, that makes the identification window width increase linearly with income. Another possibility is an *absolute* threshold fixing the identification window width at the same value for every person in the distribution. A relative threshold approach is consistent with traditional inequality and polarization analysis based on relative differences. In the inter-temporal context, the identification window should be defined around an inter-temporal notion of income: let $\hat{Y}_i$ be the average income weighted using the discount rate for identification r_I:

$$\hat{Y}_i(\mathbf{y}_i; r_I) = \frac{\sum_{s=0}^{T} e^{-r_I s} y_{i,s}}{\sum_{s=0}^{T} e^{-r_I s}} \tag{2.21}$$

Then, the identification window width for person i is defined as a fraction p of her inter-temporal income: $b = p\hat{Y}_i$.

Along with the identification window, the choice for the weighting scheme is to be made. The weight must be defined for each pair of income trajectories: w_{ij} is the weight of income trajectory $\mathbf{y}_j$ for income trajectory $\mathbf{y}_i$, and is a function of their symmetric distance D_{ij}, and of the identification window of $\mathbf{y}_i$. We set $w(0) = 1$ so that the maximum weight goes to people having exactly the same income trajectory, and the identification group is never empty (at least self-identification is true). Moreover, the function is chosen so that $w(\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r) \geq b) = 0$, where b is the maximum distance allowed for identification: outside the identification window the weight is always zero.

We propose in Equation (2.22) three possible weighting functions coherent with the conditions $w(0) = 1$, and $w(\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r) \geq b) = 0$ and plot them in Figure 2.3.

$$\begin{cases} w_{ij}^{rectangular}(\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r)) = 1 & \text{if } |y_j - y_i| \leq b_i \\ w_{ij}^{rectangular}(\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r)) = 0 & \text{otherwise} \\ \\ w_{ij}^{triangular}(\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r)) = max[0, 1 - \frac{\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r)}{b_i}] \\ \\ w_{ij}^{triweight}(\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r)) = max[0, (1 - (\frac{\hat{D}_{ij}(\mathbf{y}_i, \mathbf{y}_j; r)}{b_i})^2)^3] \end{cases} \quad (2.22)$$

Rectangular weights is reminiscent of Esteban and Ray (1994)'s original model with discrete income classes, but is the least attractive option with continuously measured incomes: it implies that the feeling of identification is not a function of distance within a window, and leads to a discontinuity at the boundary of the window. It is, however, the easiest choice in terms of interpretation: since each person counts for one unit, the identification component is simply the proportion of the population falling inside the identification window.

Figure 2.3. Possible weighting functions

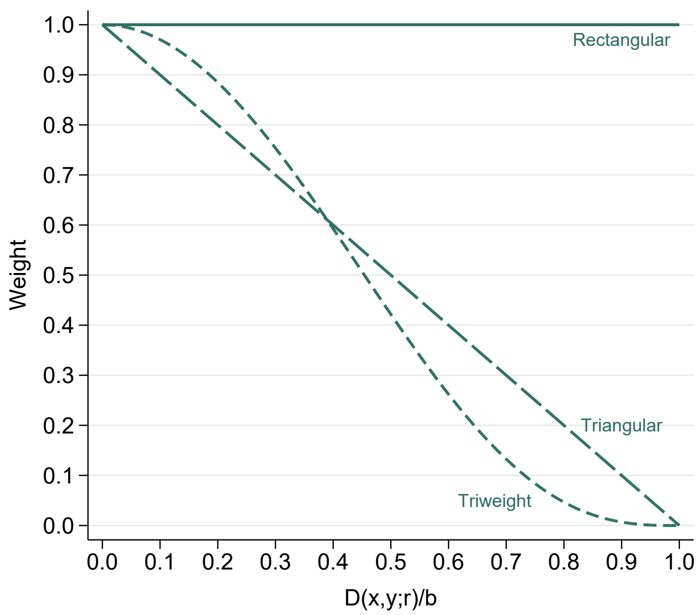

Note: The figure plots three possible weighting schemes based on Kernel functions. They share the properties of assuming the maximum value of one when $D(\mathbf{y}, \mathbf{y}'; r) = 0$, and the minimum value of zero when $D(\mathbf{y}, \mathbf{y}'; r) \geq b$.

On the contrary, when using the triangular or the triweight schemes, the inclusion into the identification group is smoother: more weight is given to income trajectories close to one's own, while for the others inside the identification window the weight falls when approaching the boundary. Moreover, while the triangular weights imply a linear decay, the triweight scheme gives more importance (as compared to the triangular) to observations below a certain value of the $\frac{\hat{D}_{ij}(\mathbf{y}_i,\mathbf{y}_j;r)}{b_i}$ ratio, and less to observations above it. Figure 2.3 illustrates this clearly. Of course, other weighting schemes are possible and those in Equations (2.22) are not exhaustive.

2.5 Inter-temporal polarization across cohorts of Italian workers

We provide in this section an application of the index developed in Section 2.3 using matched survey-administrative data for Italy. The main goal of this application is to detect whether changing the relevance of past differences for identification and alienation r_I and r_A captures different aspects of the polarization process. We compute separately for successive cohorts of Italian workers the level of inter-temporal polarization over ten years of earnings. A cohort analysis means that the reference group for identification and alienation sentiments are those having a similar age in the same years. We examine whether the measured levels and patterns of polarization are affected by varying the parameters of memory. We also show a heterogeneity analysis by gender, level of education, and geographic area of work, which we suspect to be important correlates of the polarization process.

2.5.1 Data

Data source The inter-temporal index of polarization requires income trajectories as input, so we need panel data that follow people over time. The pairwise comparisons of income trajectories in Equation (2.8) also require a balanced panel – that every person is observed in the same periods. The data used for the application are a subset of the Administrative-SILC (AD-SILC) panel dataset, developed by merging through fiscal codes the waves from 2004 to 2017 of the IT-SILC survey (the Italian component of the European Union Statistics on Income and Living Conditions, EU-SILC) with social security records collected by the Italian National Social Security Institute (INPS). The INPS archives record employment and earnings histories of all individuals working in Italy, collecting demographic characteristics, gross annual earnings, allowances, weeks worked in the year, and the type of employment contract. The EU-SILC component allows us to exploit individual-level information usually not available in administrative data, as the highest level of education.

Sample selection The sample is restricted by excluding individuals without Italian citizenship, under-represented in older cohorts. We focus on employees in the private sector, the only category covering a very long-time span in the INPS archives (from 1974 on). However, our measure of economic well-being — real annual earnings at 2015 price level, gross of personal income taxes and social contributions — includes individual income from any job, also from atypical work and self-employment, and allowances for sickness, maternity and CIG.[3] We follow workers born between 1940 and 1973 for the eleven years in which they are aged 35–45, with at least six years of positive earnings. Periods of non-employment observed in the data are counted as zero income. In terms of calendar years, we observe earnings patterns from 1975 to 2018. The bottom and top 0.1% of the earnings distribution in each year are dropped to minimize measurement errors at the tails and get rid of severe outliers. Earnings are mean-normalized within the cohort and age to ensure the scale invariance of the index and therefore its comparability over time.

The workers included are divided into thirty five-year rolling cohorts of birth, each of which overlaps with the preceding one for every year but the last one. The final sample includes 26,645 workers, of which 16,720 men and 9,925 women, divided into cohorts of birth from 1940–1944 to 1969–1973. Summary statistics on the selected sample are available in Table B.1 and Table B.2 in the Appendix.

2.5.2 Choice of parameters

The inter-temporal polarization index is computed separately for each cohort of birth using IT-SILC sample weights. We repeat the computation of the inter-temporal polarization index for $\alpha = 1$ — minimum sensitivity to identification — and for $\alpha = 1.6$ — maximum sensitivity. When setting the degree of memory, we use all the combinations of $r_A, r_I = 0, 0.5, 6$. If the discount rate is high ($r_I, r_A = 6$) only the value of income in the last period — at age 45 — has a positive weight, and the past is completely forgotten. If $r_I, r_A = 0.5$, the discount process is such that the closer the age approaches 45 the more weight the income gains, but income in the first period — at age 35 — has a weight of zero. Finally, $r_I, r_A = 0$ means no discount of the past, so that every age has the same weight equal to $1/11$.

In the baseline specification, we use all the possible combinations of $\alpha = 1$ and $\alpha = 1.6$ with $r_I, r_A = 0, 0.5, 6$, an identification window of 20% of own inter-temporal income, and a triweight Kernel to weight the neighbours inside one's identification window.[4] We provide in the Appendix also results with triangular and rectangular weights, and the results using an absolute threshold for the identification window defined as 20% of average within-cohort inter-temporal income.

Only for the specifications with $r_I = r_A$ for brevity, we report in the Appendix a plot

[3]The Cassa Integrazione Guadagni (CIG) is an Italian short-time work scheme for supporting the wages of employees of firms going through crisis events. It is limited in time and subject to specific requirements for both the employer's nature and situation and the employment contract.

[4]Since we include observations with zero income, setting a relative threshold mechanically forces the identification window for zero earners to be zero. This implies that identification for zero-income earners is only possible with other zero-income earners, which is consistent with the identification of those not working with the group of non-employed. Results with only positive earnings are reported in the Appendix.

of the indices with normal-based confidence intervals (Figure B.5) and a table with the standard errors (Table B.3) computed through 1,000 bayesian bootstrap repetitions.

2.5.3 Results

Figure 2.4 shows the patterns across cohorts of inter-temporal identification, alienation and polarization for the minimum ($\alpha = 1$) and the maximum ($\alpha = 1.6$) value of sensitivity to identification, and for all the combinations of the parameters of memory $r_I, r_A = 0, 0.5, 6$. To read the graph, consider that a darker colour indicates stronger memory for identification, while a more solid pattern indicates stronger memory for alienation. The identification component is not sensitive to the discount rate for alienation r_A; therefore, only the three lines corresponding to the three values of r_I are plotted on the left-hand panel of Figure 2.4. The same applies to the alienation component, which does not depend on the discount rate r_I but only on r_A. The degree of memory seems to be relevant for identification in terms of both level and pattern: in case of no memory ($r_I = 6$), the average proportion of people inside one's identification group at age 45 is between 12 and 17.5% depending on the cohort. If $r_I = 0.5$—the memory decays going backwards in time reaching zero at age 35—the average fraction of neighbours falls between 4.7 and 7%. Increasing memory to the maximum ($r_I = 0$), the identification is between 2.8 and 5.3%. This suggests that, in the process of including nearby people in income groups, which is the basis of the identification process, incorporating past income differences leads to a lower rate of inclusion, distancing from each other people who would be closer if only income at age 45 were taken into account.

The inter-temporal identification has a clear long-run decreasing trend, falling by 25.1% with minimum memory, 30.4% with maximum memory, with a peak around cohort 1946–1950, falling thereafter, and a recovery for younger cohorts from 1964-1968 on. The trends are comparable when changing the memory parameter r, but relevant differences emerge for the oldest cohorts: when focusing only on present income, inter-temporal identification falls by 14.1% from cohort 1940-1944 to cohort 1947-1951; on the other extreme, if the past is weighted as the present, we see an increase in identification of 19.9%. This means that, for those cohorts, people were moving from bigger to smaller sized groups in terms of earnings at age 45, reducing cross-sectional group identity in middle-career, while they were moving from smaller to bigger groups in terms of income trajectories, increasing group identity in terms of career paths.

A possible explanation for this pattern relies on the functioning of the Scala Mobile – 'elevator' – wage indexation mechanism adopted in Italy from the 1970s to the early 1990s (weakened from 1984 before final abolition in 1992). It was designed for granting the same absolute wage increase to all employees in a period of sustained inflation, inducing mechanically greater proportional wage changes at the bottom of the distribution (Manacorda,

Figure 2.4. Inter-temporal polarization and its components by cohort

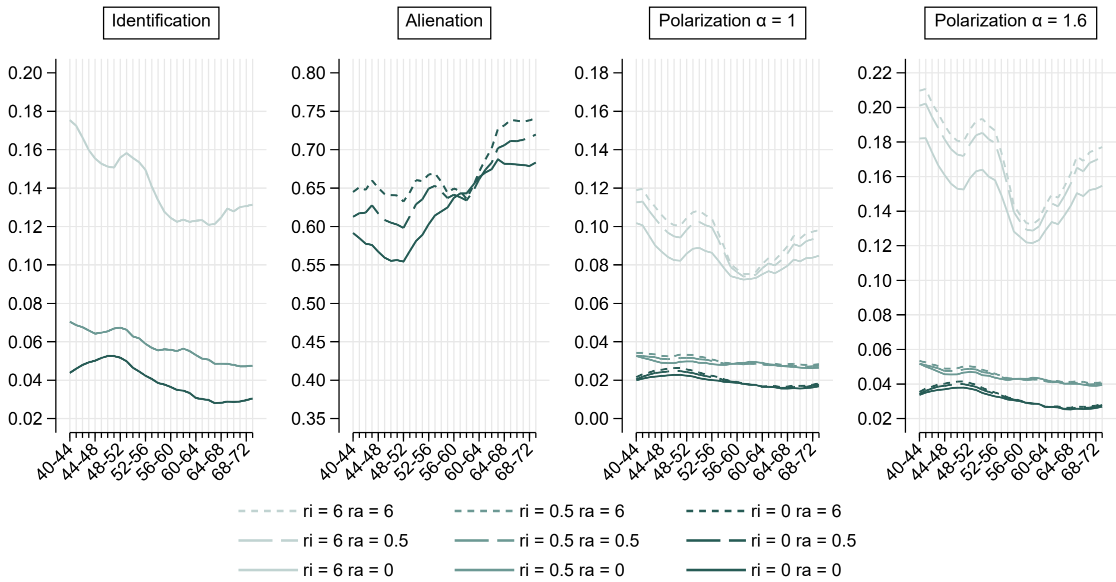

Note: The figure plots the value of average inter-temporal identification, alienation and polarization of gross real annual earnings for all the possible combinations of three values $(0, 0.5, 6)$ of the discount rates of past differences for identification r_I and for alienation r_A. The identification window width is defined as 20% of own inter-temporal earnings, and an adjusted triweight weighting scheme is applied to weight observations inside one's identification window. The sample includes employees in the private sector, and the indices are computed separately for thirty five-year rolling cohorts of birth, from 1940-1944 to 1969-1973 among workers observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. *Source*: AD-SILC data 1975–2018.

2004). The indexation may have played a role in making income trajectories more similar than before for workers affected by the Scala Mobile. Cohorts 1940–1944 to 1945–1949 in our analysis were completely covered by the mechanism during age 35–45; later on, falling inflation and the reforms aimed at reducing the wage increases made the equalizing power increasingly weak.

Moving to inter-temporal alienation, we can appreciate in the second panel from the left in Figure 2.4 that the level and the upward trend of inequality are comparable to available estimates of the Gini index for Italy in the private sector.[5] We see that increasing the memory for alienation ($\downarrow r_A$, more solid pattern in the figure) lowers the level of inequality for almost every cohort: the more people remember their past income, the less alienation they perceive, meaning that incorporating past experience reduces long distances between people. This behaviour is consistent with a world of mostly diverging trajectories: when the income distances today are larger than those in the past, including past experience makes the inter-temporal distance decrease. Interestingly, this mechanism is not in place for a bunch of cohorts around cohort 1953-1957 for which alienation is almost independent of the parameter of memory, uncovering very persistent inequalities.

[5]When comparing the alienation component of the polarization index with the Gini index, remember that the former should be divided by 2 to avoid double counting of income differences.

The two panels on the right in Figure 2.4 combine inter-temporal identification and alienation as in Equation 2.19 using two different levels of sensitivity to identification $\alpha = 1, 1.6$. We notice that polarization is influenced more by identification than by alienation; the more the inter-temporal polarization index is sensitive to group identity, the more its pattern resembles that of identification. The overall trend of inter-temporal polarization is decreasing despite rising alienation, suggesting that the effect of falling identification is predominant. To comment on the impact of the discount process on the polarization index, we focus on the case in which $\alpha = 1$. The impact on inter-temporal polarization of increasing memory follows what we saw for identification and alienation separately: when $\downarrow r_I$—stronger memory for identification—the level of polarization rapidly falls for all cohorts, while when $\downarrow r_A$—stronger memory for alienation—there is again a reduction in polarization but it does not affect heavily the cohorts characterized by very persistent inequalities.

The long-run reduction of polarization is similar in magnitude if we use minimum memory (-17.5%) or maximum memory (-15.6). However, as we have seen for identification, the pattern is different: while polarization in income at 45 for older cohorts was falling, it was actually on the rise in terms of income trajectories, and a similar decoupling emerges for the very last cohorts.

We show in the Appendix in Figure B.5 and Table B.3 that the differences commented here are statistically significant by using bayesian bootstrap inference. Moreover, we provide some sensitivity analyses changing the weighting scheme inside the identification window to triangular (Figure B.1) and rectangular (Figure B.2), changing the way of defining the identification window to an absolute threshold of income (Figure B.4), and removing zero earners (Figure B.3). As expected, a less smooth weighting scheme increases the groups' relative size, but all the main results for inter-temporal polarization are robust to different specifications. When excluding workers with even a single zero when 35-45, we report an expected large drop in the level of alienation and we notice that the impact of the zeros on identification is visible almost exclusively when $r_I = 6$; with no memory, there can be a sizable group including zero earners, while moving to income trajectories the periods spent without earnings are smoothed out and the level of identification is lower. The impact of this mechanism is particularly strong for the last cohorts from cohort 1962-1966 on, for whom there is a strong increase in identification when zero-income earners are included and a slightly decreasing trend when they are excluded.

2.5.4 Heterogeneity

In this section, we explore the heterogeneity in within-cohort inter-temporal polarization and its components by gender, level of education, and geographical area of work. We regress separately for each cohort individual identification, alienation and effective antagonism on indicators for being a woman, tertiary graduate, and working in the South or Islands of Italy, as in Equation (2.23):

$$z_i = \beta_0 + \beta_1 F_i + \beta_2 G_i + \beta_3 S_i + u_i \tag{2.23}$$

where z_i can be either identification, alienation, or effective antagonism for person i, F_i is an indicator for whether i is a woman, G_i a tertiary graduate, and S_i a worker in the South or Islands. The proportion of these categories in the sample is available in Table B.1 in the Appendix. On average, but with relevant variation across cohorts and an upward trend over time, women are 37% of workers, tertiary graduates are 6.9%, and workers in the South and Islands are 23.7%. These regressions exploit the linearity of our index of inter-temporal polarization with respect to individual alienation and identification to assess the contributions of different groups of people to aggregate inter-temporal polarization. We use three specifications changing the level of memory $r = 0, 0.5, 6$, with $r = r_I = r_A$. The identification window is set at 20% of own inter-temporal earnings and the weighting scheme inside the identification window is the triweight Kernel. An alternative specification using only positive earnings for comparison is provided in the Appendix.

Figure 2.5. Gender differences in inter-temporal polarization

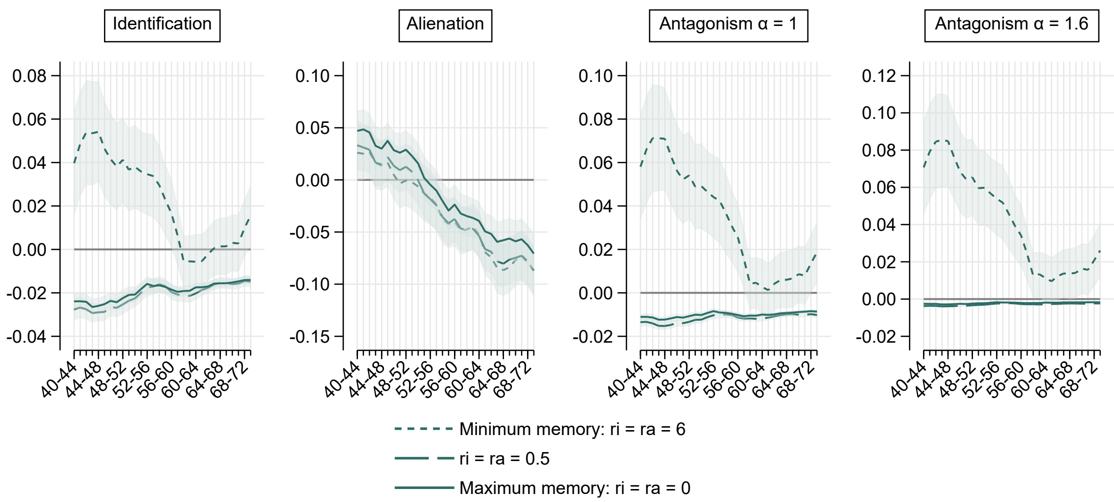

Note: The figure plots the coefficient of the indicator variable for women in a linear regression of inter-temporal identification, alienation, and effective antagonism on three indicator variables for women, tertiary educated, and workers in the South or Islands of Italy. The regressions are performed separately for each cohort of birth and level of the parameter of memory r – 0, 0.5, 6 –, and $r_I = r_A$ to make the results manageable. Consider that higher r means weaker memory. The confidence intervals plotted are at 95% confidence level computed using robust standard errors. The values of individual inter-temporal identification, alienation and polarization used on the left-hand side are computed on gross real annual earnings using an identification window of 20% of own inter-temporal earnings, an adjusted triweight weighting scheme, and including workers with zero earnings for at most five years are included. *Source*: AD-SILC data 1975–2018.

Figure 2.5 plots the coefficient for the indicator for women in the regression also including indicators for tertiary graduates and workers in the South and Islands. It is remarkable that the degree of memory is crucial to assess the role of gender on inter-temporal polarization and its components. For older cohorts, while women were more concentrated into big-sized groups than men in terms of income at age 45, they were less concentrated than men in terms of income trajectories. However, if we look at the same graph excluding zero earners (Figure B.6), the stronger cross-sectional identification of women disappears: it is entirely due to the zeros at age 45 which were more frequent among women. This explains why we see in Figure 2.5 that the identification gap is closing for subsequent cohorts, consistently with increasing participation and stability of women in the labour market.

With or without zeros, the alienation component is on average larger for men than for women, with a gap rapidly widening across cohorts. A possible explanation is linked to the nature of the rising inequality we see in Figure 2.4: if the rise is due to a right tail getting further away, and in that tail there are mostly men, then we should see an increase in alienation for men faster than that of women. If women are less identified and less alienated, they clearly also show lower levels of effective antagonism. The gap is reducing across cohorts but gender differences remain also for young cohorts. No statistically significant differences in effective antagonism emerge comparing different levels of memory, unless we include zero earners: because of the role of non-employment for the identification process, the resulting effective antagonism is stronger for women than for men when it is computed at age 45 with no memory. Using income trajectories, instead, reduces the gender gap in antagonism.

Figure 2.6. Education differences in inter-temporal polarization

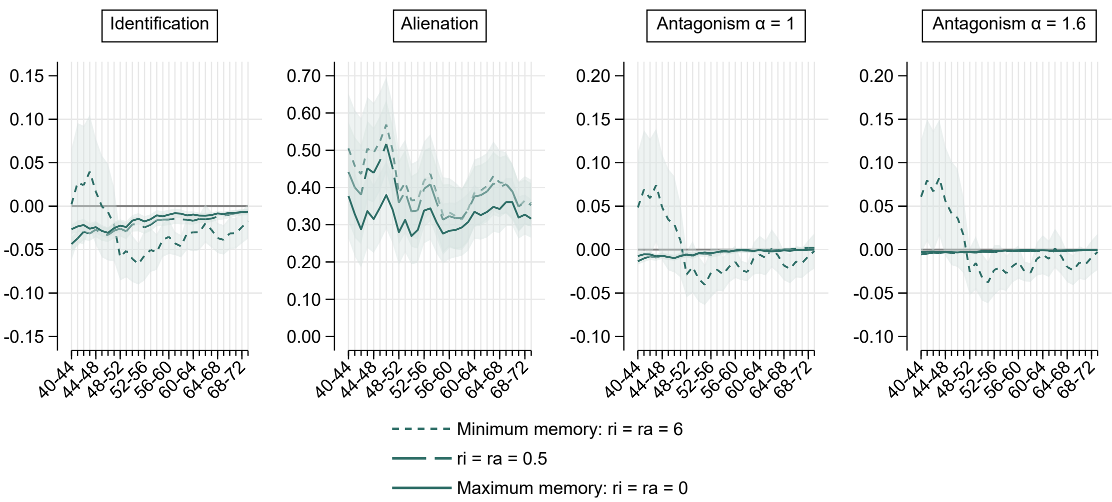

Note: The figure plots the coefficient of the indicator variable for tertiary education in a linear regression of inter-temporal identification, alienation, and effective antagonism on three indicator variables for women, tertiary educated, and workers in the South or Islands of Italy. The regressions are performed separately for each cohort of birth and level of the parameter of memory r — 0, 0.5, 6 —, and $r_I = r_A$ to make the results manageable. Consider that higher r means weaker memory. The confidence intervals plotted are at 95% confidence level computed using robust standard errors. The values of individual inter-temporal identification, alienation and polarization used on the left-hand side are computed on gross real annual earnings using an identification window of 20% of own inter-temporal earnings, an adjusted triweight weighting scheme, and including workers with zero earnings for at most five years are included. *Source*: AD-SILC data 1975–2018.

With regard to education differences, Figure 2.6 shows that tertiary graduates are less identified but far more alienated than non-graduates, consistently with a concentration of more educated workers in the upper tail of the distribution. The differences in identification are present both in terms of income at age 45 and of income trajectories, and the identification gap seems to be larger if measured only in the last year ($r_I = 6$), as it should be the case if tertiary education gives an advantage in terms of faster career progression. The large variance of the identification gap for oldest cohorts comes from the presence of graduated women with zero earnings for childcare reasons. If we look at Figure B.7 where zero earnings are excluded, the coefficients for the old cohorts are in line with the overall pattern of an existing but closing over time gap between graduates and non-graduates. The resulting effective antagonism appears to be slightly larger for non-graduates, driven by their stronger group identity counterbalancing the smaller alienation. However, the gap for the last cohorts is almost zero, suggesting that tertiary graduates are less polarized in the distribution than before.

Figure 2.7. Geographic differences in inter-temporal polarization

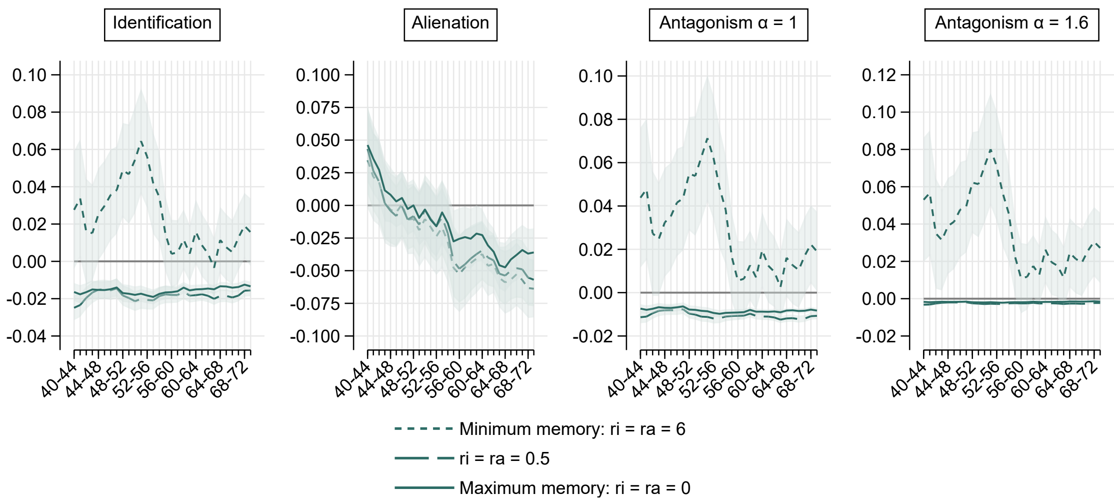

Note: The figure plots the coefficient of the indicator variable for working in the South or Islands of Italy in a linear regression of inter-temporal identification, alienation, and effective antagonism on three indicator variables for women, tertiary educated, and workers in the South or Islands of Italy. The regressions are performed separately for each cohort of birth and level of the parameter of memory r — 0, 0.5, 6 —, and $r_I = r_A$ to make the results manageable. Consider that higher r means weaker memory. The confidence intervals plotted are at 95% confidence level computed using robust standard errors. The values of individual inter-temporal identification, alienation and polarization used on the left-hand side are computed on gross real annual earnings using an identification window of 20% of own inter-temporal earnings, an adjusted triweight weighting scheme, and including workers with zero earnings for at most five years are included. *Source*: AD-SILC data 1975–2018.

Also for geographic differences, there is a large role for the zeros when there is no memory: in Figure 2.7, workers in the South and Islands seem to be part of bigger sized income groups at age 45 with respect to the other geographic areas, at least for some cohorts. However, in Figure B.8, in which zero earners are excluded, we don't see anymore this identification advantage for the South, while we see a lower or equal level of group identity depending on the cohort. The alienation component, while very similar in the past for all the areas, seems to be smaller in the South for younger cohorts. The effective antagonism is therefore weaker on average for workers in the South and Islands due to lower identification and alienation. Looking at the role of memory, when excluding the zeros geographical differences appear not to be linked to the level of memory: this suggests a relevant degree of persistence in earnings differences linked to the area of work. As we have seen for gender differences, also for the geographical area the non-employed play a crucial role in making workers in the South and Islands more polarized at age 45 than the other groups.

2.6 Discussion and Conclusions

We start from a well-established literature on income polarization modelled as the interaction between group identity and income distance (Esteban and Ray, 1994; Duclos et al., 2004; Esteban et al., 2007) to develop an index of inter-temporal polarization to explicitly incorporate the time dimension in polarization analysis. Our main claim is that cross-sectional indices fail in properly measuring identification and alienation in presence of income mobility because they do not account for the duration of individuals' proximity or distance that matters for building ties. We introduce the concept of inter-temporal distance based on discounting past income differences through a parameter of memory. By simply replacing the cross-sectional income distance with this more complex notion of reciprocal proximity or remoteness, our index measures the concentration around poles of income trajectories rather than point-in-time incomes, allowing income dynamics to mediate the identification-alienation mechanism.

The proposed framework is based on income-based processes of identification and alienation, and further research is needed to complement this approach with other dimensions of antagonism: certain socio-demographic characteristics (e.g. gender, ethnicity), but also characteristics of labour market experience (e.g. occupation, self-employed status), can certainly influence group identity in addition to (or in contrast to) income. Specific considerations can also be made about the geographical dimension: while in this paper we focus on the effect of *time* on proximity and distance, we expect also *space* to play a role.

By applying the index to a sample of about 27,000 workers in Italy covering the years 1975–2018, we provide a concrete example of how much the measure of income polarization is sensitive to the longitudinal perspective. We adopt a cohort approach and follow workers every year between ages 35 and 45, computing income polarization and its two components—identification and alienation—within each cohort separately. First of all, we document for cohorts of Italian workers a long-term trend of decreasing identification and increasing alienation, regardless of the degree of memory. The decline in identification prevails over the rise in alienation, leading to a picture of falling polarization. The role of memory is relevant for the level of identification: incorporating past income differences leads to a lower rate of inclusion of neighbours into one's identification window, uncovering that it is more rare to be close in trajectory than in mid-career earnings. On the other hand, a stronger memory mitigates alienation by reducing large income distances between people, suggesting that the high cross-sectional inequality in mid-career is due to income trajectories that diverge over time.

We also uncover in our application that the effect of memory is not constant across cohorts: for some of them, polarization in mid-career is stronger than that in income trajectory, but for others, the two are not so distant. This result can be a useful indicator of how persistent income proximity and distance are along people's careers.

Heterogeneity analysis, exploiting the linearity of the index, sheds some light on which part of the population can be defined as more antagonistic. Women seem to be less concentrated in big-sized groups than men, and also less alienated from the rest of the distribution, showing a lower level of overall antagonism with respect to men. This difference is stronger in middle-career than in income trajectory and is reversed when zero incomes are included for old cohorts: the presence of a big group of women in non-employment makes them more polarized than men. A similar situation emerges for workers in the South and Islands, for whom the level of antagonism is lower than the rest of the population as long as the zero earners are excluded.

These results lead us to discuss the possible consequences of different levels of antagonism across social groups: the literature on income polarization started by Esteban and Ray (1994) was born to measure conflict, defined as a situation in which different social groups with opposing interests suffer losses to increase the probability of obtaining their preferred outcome (Esteban and Ray, 1999). Conflict is therefore a source of resource dissipation, and we can think of effective antagonism as a correlate of bargaining power in a conflict society. If there are groups that are less polarised than others, in the sense that they are less clustered in relatively large and isolated groups, they may represent the part of the population that is deprived of a political voice and the opportunity to claim better conditions. This may make us wonder whether, for example, the decline of women's group identity over time is actually good for them.

Appendix B

Additional material

B.1 Additional figures and tables

Figure B.1. Inter-temporal polarization by cohort – triangular weighting

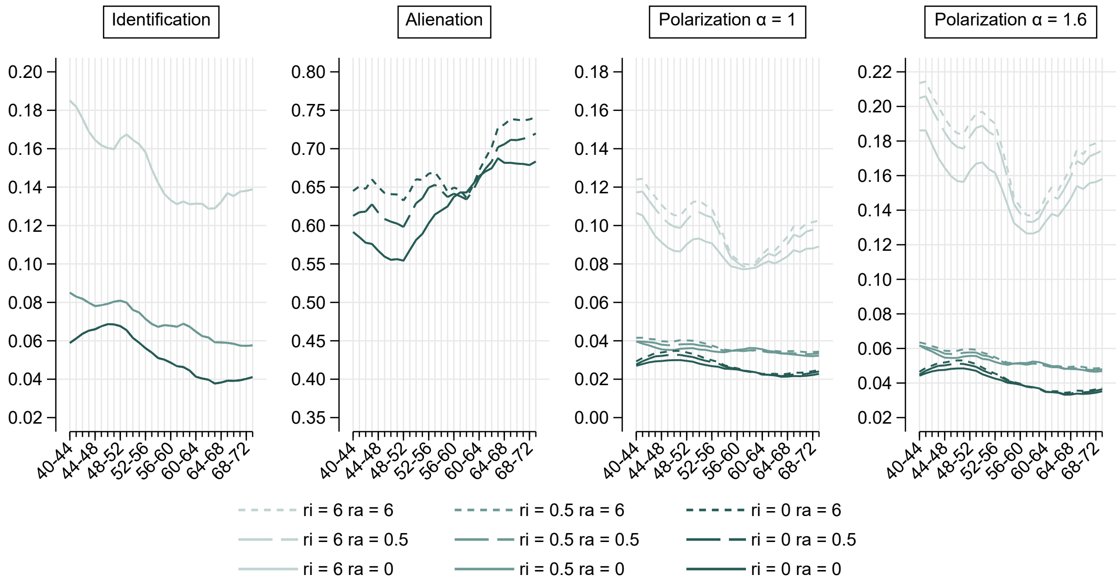

Note: The figure plots the value of average inter-temporal identification, alienation and polarization of gross real annual earnings for all the possible combinations of three values $(0, 0.5, 6)$ of the discount rates of past differences for identification r_I and for alienation r_A. The identification window width is defined as 20% of own inter-temporal earnings, and a triangular weighting scheme is applied to weight observations inside one's identification window. The sample includes employees in the private sector, and the indices are computed separately for thirty five-year rolling cohorts of birth, from 1940-1944 to 1969-1973 among workers observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. *Source*: AD-SILC data 1975–2018.

Figure B.2. Inter-temporal polarization by cohort – rectangular weighting

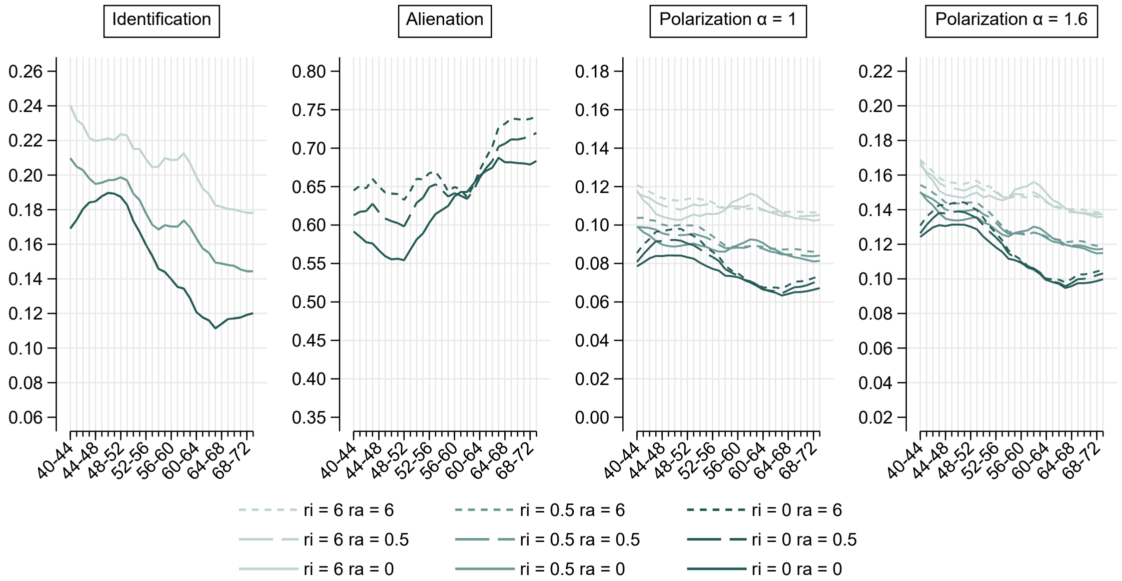

Note: The figure plots the value of average inter-temporal identification, alienation and polarization of gross real annual earnings for all the possible combinations of three values $(0, 0.5, 6)$ of the discount rates of past differences for identification r_I and for alienation r_A. The identification window width is defined as 20% of own inter-temporal earnings, and a rectangular weighting scheme is applied to weight observations inside one's identification window. The sample includes employees in the private sector, and the indices are computed separately for thirty five-year rolling cohorts of birth, from 1940-1944 to 1969-1973 among workers observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. *Source*: AD-SILC data 1975–2018.

Figure B.3. Inter-temporal polarization by cohort – only positive earnings

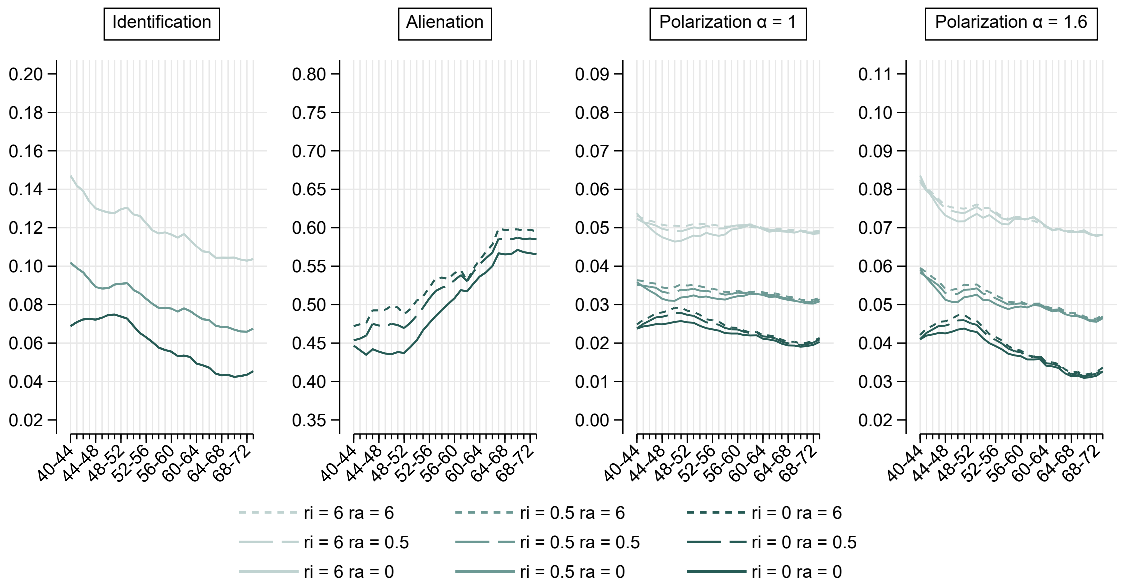

Note: The figure plots the value of average inter-temporal identification, alienation and polarization of gross real annual earnings for all the possible combinations of three values $(0, 0.5, 6)$ of the discount rates of past differences for identification r_I and for alienation r_A. The identification window width is defined as 20% of own inter-temporal earnings, and an adjusted triweight weighting scheme is applied to weight observations inside one's identification window. The sample includes employees in the private sector, and the indices are computed separately for thirty five-year rolling cohorts of birth, from 1940-1944 to 1969-1973 among workers observed every year from age 35 to 45. Workers with even one year of zero are excluded. *Source*: AD-SILC data 1975–2018.

Table B.1. Summary statistics

		Annual earnings (€)								
Cohort	N	Mean	SD	p10	p50	p90	Zeros	Women	Tertiary	South
1940-1944	2,952	22,137	12,259	3,614	21,997	36,210	6.0	29.9	2.7	20.1
1941-1945	2,837	23,355	13,190	6,006	22,960	36,574	4.5	30.4	3.8	21.4
1942-1946	2,983	22,941	13,760	3,091	22,427	37,994	6.0	31.7	2.8	20.6
1943-1947	3,126	23,716	13,015	6,173	23,338	38,813	4.2	30.8	3.7	21.9
1944-1948	3,235	24,975	13,577	7,770	23,600	42,616	2.7	31.4	3.0	23.3
1945-1949	3,322	25,561	13,612	8,699	24,907	39,277	4.1	32.6	4.6	20.7
1946-1950	3,414	25,381	13,314	8,768	24,609	40,502	3.7	33.9	3.2	22.7
1947-1951	3,335	25,678	13,846	7,449	25,027	41,121	4.7	31.2	3.9	19.8
1948-1952	3,238	25,728	13,696	8,079	24,751	41,960	4.4	30.8	3.5	27.3
1949-1953	3,186	25,253	14,334	6,352	24,317	43,740	4.2	33.6	4.5	26.0
1950-1954	3,178	25,629	14,475	5,736	24,824	41,928	6.1	36.0	4.6	18.8
1951-1955	3,240	25,929	14,995	7,181	25,189	42,359	4.3	31.2	4.5	22.8
1952-1956	3,320	26,023	14,907	7,419	25,009	45,262	4.7	35.4	2.7	18.2
1953-1957	3,402	25,585	14,281	6,326	24,191	43,574	4.2	39.0	4.4	20.8
1954-1958	3,490	26,294	14,749	7,684	25,398	43,088	4.8	36.2	4.2	23.6
1955-1959	3,555	26,557	15,424	9,879	24,536	43,765	3.3	35.0	4.4	18.7
1956-1960	3,573	26,388	16,360	5,281	25,238	45,993	4.9	35.1	4.7	21.6
1957-1961	3,648	25,295	14,183	7,274	24,157	44,686	4.2	35.9	6.7	28.6
1958-1962	3,828	26,282	15,694	8,902	23,652	47,507	3.8	40.8	6.1	20.8
1959-1963	4,028	24,160	15,148	5,805	22,907	43,873	4.3	36.2	5.2	28.0
1960-1964	4,282	25,002	15,858	6,403	23,632	44,232	5.3	35.8	6.8	25.0
1961-1965	4,584	26,176	15,587	7,752	24,737	46,530	3.7	36.2	6.4	25.5
1962-1966	4,854	24,040	15,546	5,803	22,893	42,133	5.2	41.4	7.5	22.5
1963-1967	5,048	24,749	16,611	5,165	22,800	44,776	4.8	36.0	6.6	24.2
1964-1968	5,171	25,185	16,796	6,276	23,206	45,721	4.3	42.5	6.7	26.4
1965-1969	5,258	27,501	18,941	7,638	24,571	48,408	4.3	39.1	8.5	23.7
1966-1970	5,308	25,352	17,042	6,165	23,716	44,052	4.1	40.6	7.4	25.7
1967-1971	5,285	24,897	15,895	6,622	23,670	42,786	5.2	42.5	11.6	28.2
1968-1972	5,243	26,060	17,518	6,878	23,713	47,433	4.4	43.8	14.1	27.3
1969-1973	5,173	26,036	17,001	5,671	24,284	46,201	4.9	45.2	14.8	24.4
All	26,645	25,262	15,500	6,584	23,894	43,332	4.6	37.2	6.9	23.7

Note: The table reports the number of workers and summary statistics for thirty five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. Annual Earnings are real (2015 price level) and gross of personal income taxes and social contributions. The percentage of zero earnings, women, tertiary graduates, workers in the South or Islands of Italy are reported. We use EU-SILC sample weights. *Source*: AD-SILC data 1975–2018.

Table B.2. Summary statistics - only positive earnings

Cohort	N	Annual earnings (€) Mean	SD	p10	p50	p90	Women	Tertiary	South
1940-1944	2,362	25,520	11,196	14,572	23,898	38,412	24.1	3.2	18.2
1941-1945	2,278	25,567	10,997	14,364	24,113	38,300	26.0	3.4	18.8
1942-1946	2,439	26,381	11,335	15,145	24,813	39,198	26.6	3.5	19.7
1943-1947	2,590	27,202	12,127	15,595	25,321	40,571	26.0	4.1	20.8
1944-1948	2,706	27,375	12,030	15,603	25,611	41,121	26.7	3.8	21.5
1945-1949	2,800	27,700	12,118	15,766	25,976	41,458	27.1	3.8	21.7
1946-1950	2,872	28,101	12,332	16,318	26,206	42,176	27.2	3.8	21.5
1947-1951	2,793	28,333	12,399	16,259	26,466	42,707	27.4	4.3	21.4
1948-1952	2,698	28,289	12,153	16,218	26,521	42,679	28.1	4.0	21
1949-1953	2,636	28,545	12,562	16,347	26,443	43,323	29.1	4.4	19.8
1950-1954	2,616	28,977	12,908	16,602	26,738	44,495	29.4	5.0	18.8
1951-1955	2,684	29,100	13,405	16,258	26,781	44,891	31.1	5.2	18.1
1952-1956	2,733	29,251	13,729	16,156	26,797	45,593	31.3	4.9	17.2
1953-1957	2,811	29,377	13,990	15,802	26,831	46,141	32.2	5.2	17.0
1954-1958	2,874	29,423	14,171	15,226	26,932	46,413	32.7	5.1	17.1
1955-1959	2,934	29,230	14,320	14,702	26,570	46,524	33.7	5.3	17.5
1956-1960	2,934	29,423	14,604	14,541	26,519	47,530	33.2	5.9	18.5
1957-1961	2,990	29,217	14,731	13,942	26,243	47,707	34.4	6.2	19.3
1958-1962	3,115	29,310	14,750	14,062	26,252	47,845	34.1	6.3	19.2
1959-1963	3,246	28,924	14,850	13,396	25,879	47,689	34.3	6.4	19.9
1960-1964	3,429	28,801	14,968	12,916	25,857	47,790	34.7	6.4	20.6
1961-1965	3,657	28,448	14,914	12,235	25,624	47,259	34.6	6.7	21.1
1962-1966	3,896	28,317	15,200	11,875	25,507	47,175	35.3	7.1	21.1
1963-1967	4,068	28,370	15,847	11,478	25,486	47,597	36.5	8.0	21.5
1964-1968	4,225	28,725	16,018	11,706	25,779	47,963	37.0	8.9	21.8
1965-1969	4,324	28,762	16,105	11,783	25,811	47,697	37.9	9.7	22.2
1966-1970	4,405	28,861	16,367	11,782	25,935	47,886	39.5	10.3	21.7
1967-1971	4,380	28,940	16,256	11,958	26,056	48,022	40.2	11.7	22.7
1968-1972	4,356	28,944	16,140	12,086	25,977	48,328	41.0	13.2	22.9
1969-1973	4,279	29,093	16,159	12,273	26,054	48,596	41.7	15.0	22.2
All	21,849	28,520	14,449	13,547	25,972	45,675	33.5	7.0	20.5

Note: The table reports the number of workers and summary statistics for thirty five-year rolling cohorts of birth (1940-1944 to 1969-1973) of employees in the private sector in Italy. The workers are observed every year from age 35 to 45. Only workers with positive earnings every year are included. Annual Earnings are real (2015 price level) and gross of personal income taxes and social contributions. The percentage of zero earnings, women, tertiary graduates, workers in the South or Islands of Italy are reported. We use EU-SILC sample weights. *Source*: AD-SILC data 1975–2018.

Figure B.4. Inter-temporal polarization by cohort – absolute threshold

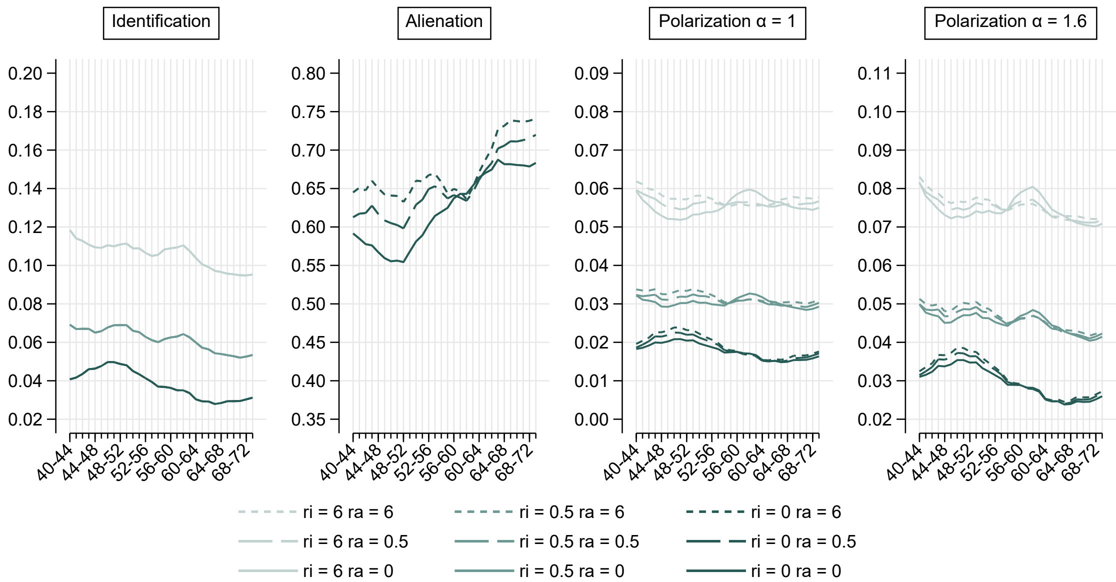

Note: The figure plots the value of average inter-temporal identification, alienation and polarization of gross real annual earnings for all the possible combinations of three values $(0, 0.5, 6)$ of the discount rates of past differences for identification r_I and for alienation r_A. The identification window width is defined as 20% of the mean inter-temporal earnings and an adjusted triweight weighting scheme is applied to weight observations inside one's identification window. The sample includes employees in the private sector, and the indices are computed separately for thirty five-year rolling cohorts of birth, from 1940-1944 to 1969-1973 among workers observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. *Source*: AD-SILC data 1975–2018.

Figure B.5. Inter-temporal polarization by cohort with bootstrap CI

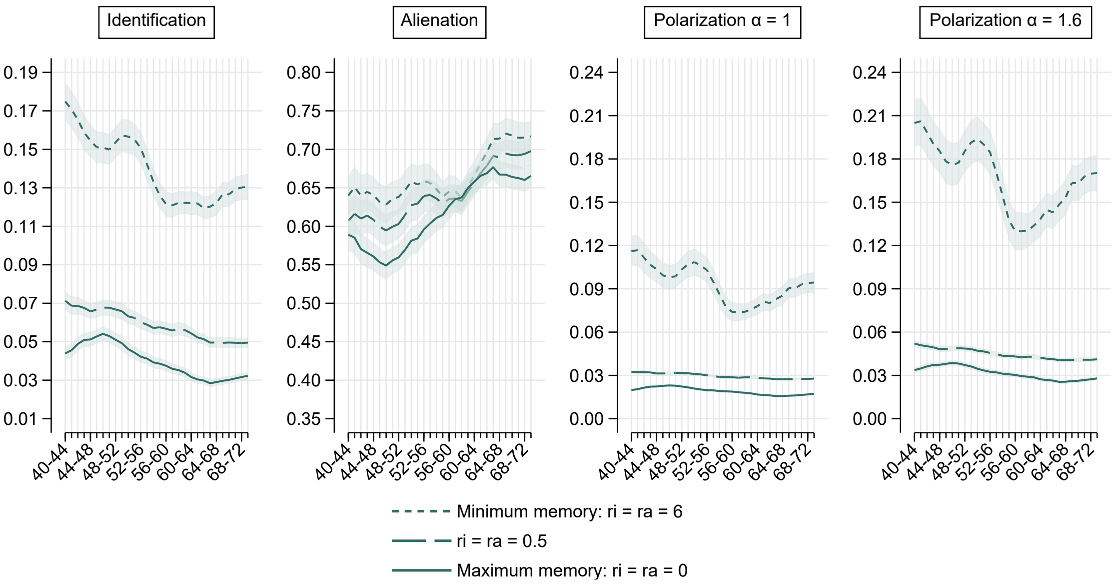

Note: The figure plots the value of average inter-temporal identification, alienation and polarization of gross real annual earnings for the three values of memory $r = r_I = r_A = 0, 0.5, 6$ for both identification and alienation. Consider that higher r means weaker memory. The identification window width is defined as 20% of own inter-temporal earnings, and a triangular weighting scheme is applied to weight observations inside one's identification window. Normal-based confidence intervals at 95% confidence level are computed using 1,000 bayesian bootstrap repetitions and reported through the shaded areas. The sample includes employees in the private sector and the indices are computed separately for thirty five-year rolling cohorts of birth, from 1940-1944 to 1969-1973 among workers observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. *Source*: AD-SILC data 1975–2018.

Figure B.6. Gender differences in inter-temporal polarization — only positive earnings

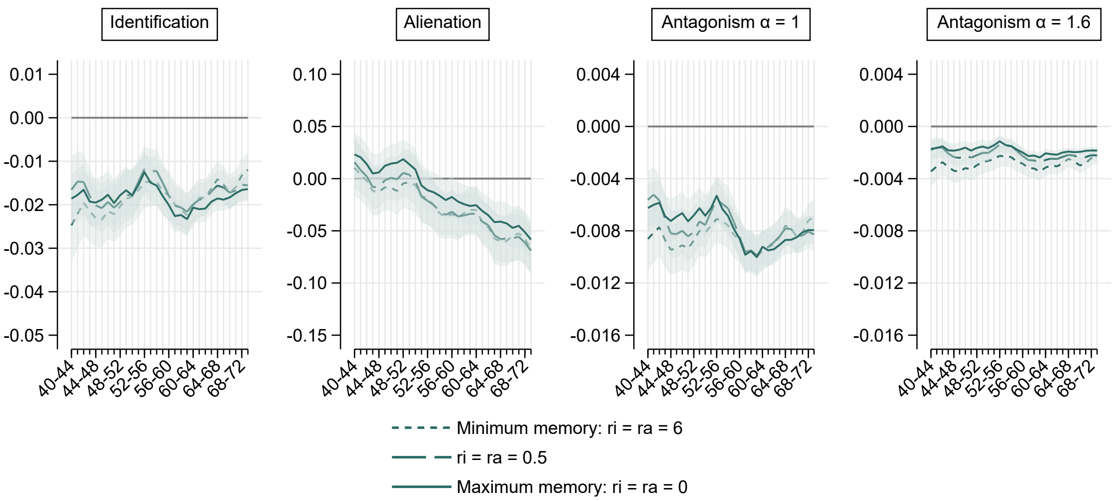

Note: The figure plots the coefficient of the indicator variable for women in a linear regression of inter-temporal identification, alienation, and effective antagonism on three indicator variables for women, tertiary educated, and workers in the South or Islands of Italy. The regressions are performed separately for each cohort of birth and level of parameter of memory r — 0, 0.5, 6 —, and $r_I = r_A$ to make the results manageable. Consider that higher r means weaker memory. The confidence intervals plotted are at 95% confidence level computed using robust standard errors. The values of individual inter-temporal identification, alienation and polarization used on the left-hand side are computed on gross real annual earnings using an identification window of 20% of own inter-temporal earnings, an adjusted triweight weighting scheme, and excluding workers with zero earnings in any year. *Source*: AD-SILC data 1975–2018.

Table B.3. Inter-temporal polarization index with bootstrap standard errors

	Identification			Alienation			Polarization $\alpha = 1$			Polarization $\alpha = 1.6$		
Cohort	$r = 0$	$r = 0.5$	$r = 6$	$r = 0$	$r = 0.5$	$r = 6$	$r = 0$	$r = 0.5$	$r = 6$	$r = 0$	$r = 0.5$	$r = 6$
1940-1944	0.044	0.071	0.175	0.589	0.607	0.640	0.020	0.033	0.116	0.034	0.052	0.205
	(0.002)	(0.002)	(0.005)	(0.010)	(0.012)	(0.013)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.008)
1941-1945	0.046	0.069	0.171	0.585	0.616	0.651	0.020	0.032	0.117	0.035	0.051	0.206
	(0.002)	(0.003)	(0.005)	(0.011)	(0.013)	(0.014)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.008)
1942-1946	0.049	0.069	0.165	0.570	0.610	0.640	0.021	0.032	0.112	0.036	0.050	0.198
	(0.002)	(0.003)	(0.005)	(0.010)	(0.012)	(0.013)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.009)
1943-1947	0.051	0.068	0.159	0.566	0.614	0.645	0.022	0.032	0.107	0.037	0.049	0.190
	(0.002)	(0.002)	(0.005)	(0.010)	(0.011)	(0.012)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.008)
1944-1948	0.051	0.066	0.155	0.561	0.608	0.641	0.022	0.031	0.104	0.037	0.048	0.185
	(0.002)	(0.002)	(0.004)	(0.009)	(0.011)	(0.012)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.008)
1945-1949	0.053	0.067	0.151	0.553	0.599	0.631	0.023	0.031	0.099	0.038	0.048	0.178
	(0.002)	(0.002)	(0.004)	(0.009)	(0.011)	(0.012)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.008)
1946-1950	0.054	0.068	0.151	0.549	0.595	0.628	0.023	0.031	0.098	0.039	0.049	0.176
	(0.002)	(0.002)	(0.004)	(0.009)	(0.011)	(0.012)	(0.001)	(0.001)	(0.004)	(0.001)	(0.001)	(0.008)
1947-1951	0.053	0.068	0.150	0.556	0.599	0.635	0.023	0.032	0.099	0.038	0.049	0.177
	(0.002)	(0.002)	(0.004)	(0.009)	(0.010)	(0.012)	(0.001)	(0.001)	(0.004)	(0.001)	(0.001)	(0.007)
1948-1952	0.051	0.067	0.154	0.560	0.603	0.638	0.022	0.032	0.103	0.037	0.049	0.185
	(0.002)	(0.002)	(0.004)	(0.009)	(0.010)	(0.011)	(0.001)	(0.001)	(0.004)	(0.001)	(0.001)	(0.007)
1949-1953	0.049	0.066	0.157	0.569	0.615	0.647	0.022	0.031	0.107	0.036	0.048	0.191
	(0.002)	(0.002)	(0.004)	(0.010)	(0.011)	(0.013)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.008)
1950-1954	0.046	0.063	0.157	0.581	0.627	0.658	0.021	0.031	0.109	0.035	0.047	0.194
	(0.002)	(0.002)	(0.005)	(0.009)	(0.011)	(0.012)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.008)
1951-1955	0.044	0.062	0.155	0.584	0.630	0.654	0.020	0.031	0.106	0.033	0.047	0.190
	(0.002)	(0.002)	(0.004)	(0.010)	(0.011)	(0.012)	(0.001)	(0.001)	(0.004)	(0.001)	(0.001)	(0.008)
1952-1956	0.042	0.060	0.151	0.596	0.639	0.659	0.020	0.030	0.103	0.033	0.046	0.185
	(0.002)	(0.002)	(0.004)	(0.009)	(0.011)	(0.012)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.008)
1953-1957	0.041	0.059	0.142	0.604	0.641	0.656	0.020	0.030	0.095	0.032	0.045	0.171
	(0.002)	(0.002)	(0.004)	(0.009)	(0.011)	(0.012)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.008)
1954-1958	0.039	0.057	0.133	0.611	0.636	0.648	0.019	0.029	0.086	0.031	0.044	0.155
	(0.002)	(0.002)	(0.004)	(0.009)	(0.010)	(0.011)	(0.001)	(0.001)	(0.004)	(0.001)	(0.001)	(0.008)
1955-1959	0.039	0.058	0.126	0.615	0.629	0.638	0.019	0.029	0.078	0.031	0.044	0.138
	(0.001)	(0.002)	(0.003)	(0.009)	(0.010)	(0.011)	(0.001)	(0.001)	(0.004)	(0.001)	(0.001)	(0.007)
1956-1960	0.038	0.057	0.122	0.627	0.636	0.645	0.019	0.029	0.074	0.03	0.043	0.130
	(0.001)	(0.002)	(0.003)	(0.009)	(0.010)	(0.011)	(0.001)	(0.001)	(0.003)	(0.001)	(0.001)	(0.007)
1957-1961	0.036	0.056	0.121	0.635	0.636	0.645	0.018	0.028	0.074	0.03	0.043	0.130
	(0.001)	(0.002)	(0.003)	(0.009)	(0.009)	(0.010)	(0.000)	(0.001)	(0.003)	(0.001)	(0.001)	(0.007)
1958-1962	0.035	0.057	0.122	0.637	0.633	0.637	0.018	0.029	0.074	0.029	0.043	0.130
	(0.001)	(0.002)	(0.003)	(0.009)	(0.009)	(0.010)	(0.000)	(0.001)	(0.003)	(0.001)	(0.001)	(0.006)
1959-1963	0.034	0.056	0.122	0.650	0.645	0.652	0.018	0.029	0.076	0.029	0.043	0.134
	(0.001)	(0.002)	(0.003)	(0.009)	(0.009)	(0.010)	(0.000)	(0.001)	(0.003)	(0.001)	(0.001)	(0.006)
1960-1964	0.032	0.054	0.122	0.658	0.656	0.665	0.017	0.028	0.078	0.027	0.042	0.138
	(0.001)	(0.002)	(0.003)	(0.008)	(0.009)	(0.010)	(0.000)	(0.000)	(0.003)	(0.001)	(0.001)	(0.006)
1961-1965	0.030	0.052	0.122	0.666	0.669	0.682	0.016	0.028	0.081	0.027	0.042	0.145
	(0.001)	(0.001)	(0.003)	(0.008)	(0.009)	(0.010)	(0.000)	(0.000)	(0.003)	(0.001)	(0.001)	(0.006)
1962-1966	0.030	0.051	0.119	0.669	0.679	0.696	0.016	0.028	0.080	0.026	0.041	0.143
	(0.001)	(0.001)	(0.003)	(0.008)	(0.009)	(0.010)	(0.000)	(0.000)	(0.003)	(0.001)	(0.001)	(0.006)
1963-1967	0.028	0.050	0.120	0.677	0.691	0.713	0.016	0.027	0.083	0.026	0.041	0.149
	(0.001)	(0.001)	(0.003)	(0.008)	(0.009)	(0.010)	(0.000)	(0.000)	(0.003)	(0.001)	(0.001)	(0.006)
1964-1968	0.029	0.050	0.122	0.667	0.690	0.714	0.016	0.027	0.085	0.026	0.041	0.153
	(0.001)	(0.001)	(0.003)	(0.008)	(0.009)	(0.010)	(0.000)	(0.000)	(0.003)	(0.001)	(0.001)	(0.006)
1965-1969	0.030	0.049	0.127	0.667	0.695	0.720	0.016	0.027	0.091	0.026	0.041	0.163
	(0.001)	(0.001)	(0.003)	(0.008)	(0.009)	(0.011)	(0.000)	(0.000)	(0.003)	(0.001)	(0.001)	(0.006)
1966-1970	0.030	0.050	0.127	0.664	0.692	0.718	0.016	0.027	0.090	0.026	0.041	0.163
	(0.001)	(0.001)	(0.003)	(0.008)	(0.010)	(0.011)	(0.000)	(0.000)	(0.003)	(0.001)	(0.001)	(0.006)
1967-1971	0.031	0.049	0.129	0.663	0.692	0.715	0.016	0.027	0.093	0.027	0.041	0.168
	(0.001)	(0.001)	(0.003)	(0.008)	(0.010)	(0.011)	(0.000)	(0.000)	(0.004)	(0.001)	(0.001)	(0.006)
1968-1972	0.032	0.049	0.130	0.660	0.694	0.715	0.017	0.028	0.094	0.027	0.041	0.170
	(0.001)	(0.001)	(0.003)	(0.008)	(0.009)	(0.010)	(0.000)	(0.000)	(0.004)	(0.001)	(0.001)	(0.006)
1969-1973	0.032	0.050	0.130	0.665	0.698	0.717	0.017	0.028	0.094	0.028	0.041	0.170
	(0.001)	(0.001)	(0.003)	(0.008)	(0.009)	(0.010)	(0.000)	(0.000)	(0.004)	(0.001)	(0.001)	(0.007)

Note: The table reports the values and standard errors of average inter-temporal identification, alienation and polarization of gross real annual earnings for the three values of memory $r = r_I = r_A = 0, 0.5, 6$ for both identification and alienation. Consider that higher r means weaker memory. The identification window width is defined as 20% of own inter-temporal earnings, and a triweight weighting scheme is applied to weight observations inside one's identification window. The standard errors are computed using 1,000 bayesian bootstrap repetitions. The sample includes employees in the private sector and the indices are computed separately for thirty five-year rolling cohorts of birth, from 1940-1944 to 1969-1973 among workers observed every year from age 35 to 45. Workers with zero earnings for at most five years are included. *Source*: AD-SILC data 1975–2018.

Figure B.7. Education differences in polarization — only positive earnings

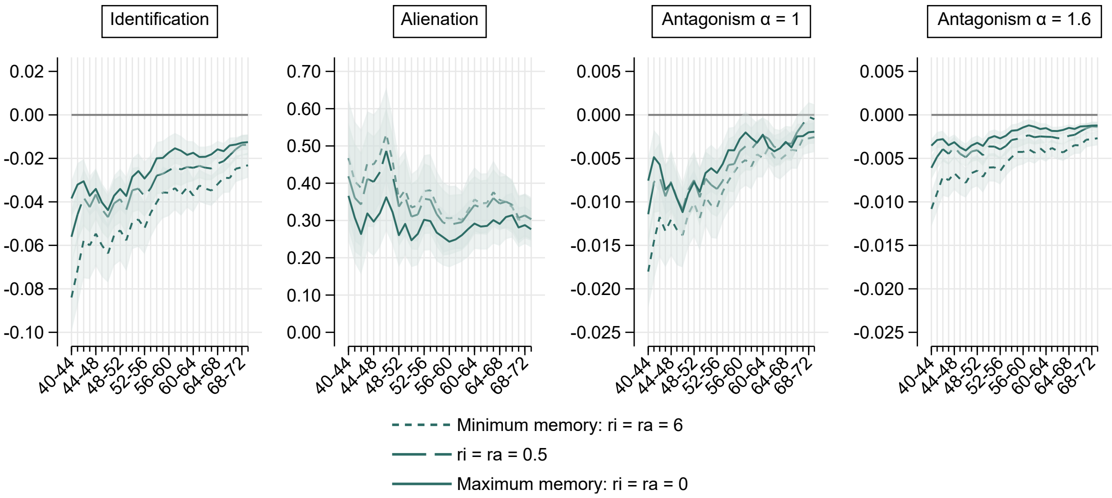

Note: The figure plots the coefficient of the indicator variable for tertiary education in a linear regression of inter-temporal identification, alienation, and effective antagonism on three indicator variables for women, tertiary educated, and workers in the South or Islands of Italy. The regressions are performed separately for each cohort of birth and level of parameter of memory r — 0, 0.5, 6 —, and $r_I = r_A$ to make the results manageable. Consider that higher r means weaker memory. The confidence intervals plotted are at 95% confidence level computed using robust standard errors. The values of individual inter-temporal identification, alienation and polarization used on the left-hand side are computed on gross real annual earnings using an identification window of 20% of own inter-temporal earnings, an adjusted triweight weighting scheme, and excluding workers with zero earnings in any year. *Source*: AD-SILC data 1975–2018.

Figure B.8. Geographic differences in polarization — only positive earnings

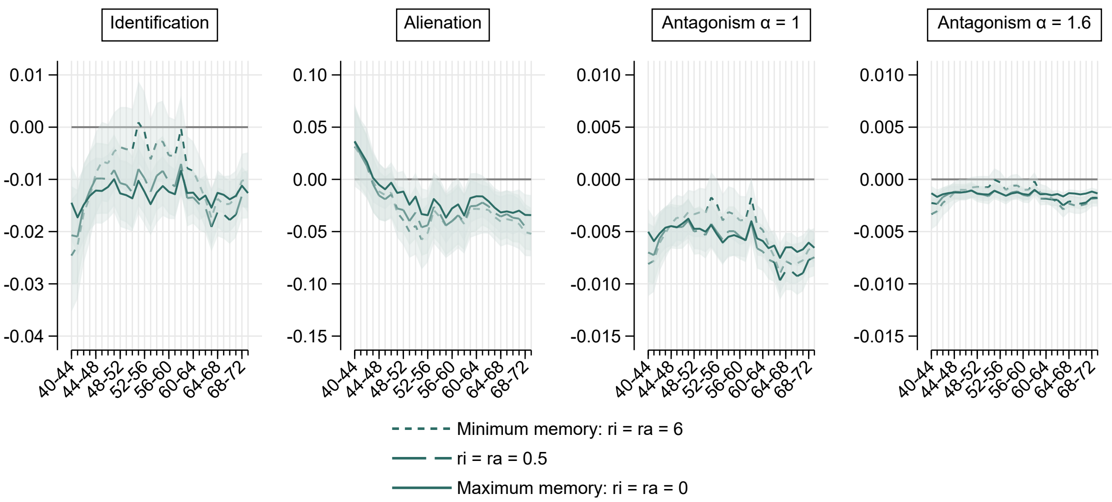

Note: The figure plots the coefficient of the indicator variable for working in the South or Islands of Italy in a linear regression of inter-temporal identification, alienation, and effective antagonism on three indicator variables for women, tertiary educated, and workers in the South or Islands of Italy. The regressions are performed separately for each cohort of birth and level of parameter of memory r — 0, 0.5, 6 —, and $r_I = r_A$ to make the results manageable. Consider that higher r means weaker memory. The confidence intervals plotted are at 95% confidence level computed using robust standard errors. The values of individual inter-temporal identification, alienation and polarization used on the left-hand side are computed on gross real annual earnings using an identification window of 20% of own inter-temporal earnings, an adjusted triweight weighting scheme, and excluding workers with zero earnings in any year. *Source*: AD-SILC data 1975–2018.

B.2 The Stata program `itempolar`

In this Appendix, I show the functioning (help file and ado file) of a Stata program called `itempolar` developed with Philippe Van Kerm to allow other researchers to compute the index proposed in this chapter using panel data. The program is still under review and is not yet available as a user-written command in Stata.

B.2.1 Help file

```
  ___  ____  ____  ____  ____(R)
 /__    /   ____/   /   ____/
___/   /   /___/   /   /___/
  Statistics/Data analysis
```

Inter-temporal Polarization Index

itempolar - Inter-temporal polarization index

Syntax

itempolar *varname* [*if*] [*in*] [*weight*] [, *options*]

Options

options	Description
idvar(*varname*)	**panel** variable. If absent, data need to be **xtset**
tvar(*varname*)	**time** variable. If absent, data need to be **xtset**
Parameters	
alphas(*numlist*)	level of **sensitivity to identification**; default is 1
ri(*#*)	discount rate for identification; default is 0.5
ra(*#*)	discount rate for alienation; default is 0.5
rt(*#*)	**relative threshold** to define the identification window
at(*#*)	**absolute threshold** to define the identification window
rtvar(*varname*)	declares which variable contains the rt() option, if different by person
atvar(*varname*)	declares which variable contains the at() option, if different by person
wtype(*string*)	**weighting scheme** inside the identification window; default is *triweight*
Variables to generate	
gen_i(*newvarname*)	generates *newvarname* for individual-level identification
gen_a(*newvarname*)	generates *newvarname* for individual-level alienation
gen_ea(*newvarname*)	generates *newvarname* for individual-level effective antagonism

fweight, **aweight**, **pweight** and **iweight** are allowed; see help weights. **by** and **bootstrap** are allowed; see help prefix. The use of weighted bootstrap is recommended (see the Remarks below).

Description

itempolar computes the inter-temporal polarization of an income distribution as defined in Subioli and Van Kerm (2022). It can be applied to other continuous variables as consumption and wealth. It works with **strongly balanced** panel data.

Options

alphas(*numlist***)** specifies a vector of **positive constants** between 1 and 1.6 capturing the importance of group identification for interpersonal effective antagonism (Esteban and Ray 1994).

ri(#) and **ra(#)** are the discount rates used to weight the period distances to get the inter-temporal distance. Setting the discount rates at zero gives **maximum memory** to the process of polarization, while setting them at 6 or more gives positive weight only to the **last** period, imposing (approximately) **zero memory** of the past. **ri(#)** and **ra(#)** may be different or not according to theoretical reasons, depending on how much past differences matter for identification and for alienation. The code allows faster computation when ri = ra, which is the default (= 0.5).

rt(#), **at(#)**, **rtvar(***varname***)**, **atvar(***varname***)**: the program allows four possible ways of defining the **bandwidth for identification** around one's income:

1. A relative threshold **rt(#)** as the fraction (0-1) of individual inter-temporal income

2. An absolute threshold **at(#)** as the fraction (0-1) of mean inter-temporal income

3. A variable **rtvar(***varname***)** containing the relative threshold for each person

4. A variable **atvar(***varname***)** containing the absolute threshold for each person

If **rtvar(***varname***)** or **atvar(***varname***)** are specified, the program takes the number defined for the last period T. A bandwidth of 0 implies that identification is possible only for exact same values of income. The default is **rt(0.2)**, implying that one feels identified with those earning from -20 to +20% of his/her mean-normalized income.

wtype(*string***)**: the program allows three possible Kernel functions **rectangular**, **triangular**, **triweight** to weight observations inside the identification window. See Jann (2007) as a reference for these Kernel functions, and consider that the triweight scheme is adjusted so that the weight is one for exact same *varname* and zero at the boundary of the identification window. The **rectangular** scheme assigns a weight of one for every observation within one's identification window. **triweight** is the default.

gen_i(*newvarname***)**, **gen_a(***newvarname***)**, **gen_ea(***newvarname***)**: it is possible to generate individual-level identification, alienation and effective antagonism. **gen_ea(***newvarname***)** requires also **gen_i(***newvarname***)** and **gen_a(***newvarname***)**, but not viceversa. Identification and alienation at the individual level do not depend on *alpha*, while effective antagonism does. If more than one **alpha** is specified, only one *newvarname* is generated for identification and alienation, while **gen_ea(***newvarname***)** will produce *newvarname1*, *newvarname2*, ..., *newvarnameA* for *A* elements of **alphas(***numlist***)**.

Remarks

Income is automatically mean-normalized period-by-period within the program to obtain **scale invariance**.

Use **bootstrap** prefix with proper clustering and removing **xt** settings if present while specifying the **idvar(***varname***)** and **tvar(***varname***)** as in the example below; see help bootstrap. The use of weighted bootstrap inference combining **exbsample** and **bs4rw** is recommended; see help bs4rw, help exbsample, and the example below.

If the **by** prefix is used, inter-temporal identification, alienation, and polarization are computed **separately** for each of the *k* subgroups.

Stored results

Scalars

r(itemident) average inter-temporal identification

r(itemalien) average inter-temporal alienation

Vectors

r(itempolar) 1 x *A* vector of values of inter-temporal polarization index for *A* elements of **alphas(***numlist***)**

If the **by** prefix is used, the program generates three matrices:

Matrices

r(itempolar) *k* x *A* matrix, inter-temporal polarization index for *k* groups and *A* elements of **alphas(***numlist***)**

r(itemident) *k* x 1 matrix, average inter-temporal identification for each of the *k* groups

r(itemalien) *k* x 1 matrix, average inter-temporal alienation for each of the *k* groups

Examples

Generate longitudinal income data

```
. clear

. set obs 100

. set seed 010101

. egen workerid = seq(), from(1) to(100)

. gen female = 0 in 1/50
```

```
. replace female = 1 in 51/100

. forvalues y = 1/3 {
2. generate income`y' = rnormal(30000, 5000)
3. }

. reshape long income, i(worker) j(year)

. xtset worker year
```

Compute inter-temporal polarization in the last year

```
. itempolar income, alphas(1) ri(6) ra(6)
```

Compute inter-temporal polarization using several values of alpha

```
. itempolar income, alphas(1 1.3 1.6)
```

Compute inter-temporal polarization separately for men and women

```
. bys female: itempolar income, alphas(1 1.6) wtype(rectangular) rt(0.1)
```

Compute the share of total inter-temporal polarization attributable to men

```
. itempolar income, alphas(1.6) gen_i(ident) gen_a(alien) gen_ea(polar)

. mat p = r(itempolar)

. preserve

. keep workerid female polar

. duplicates drop

. collapse polar, by(female)

. gen polar_all = p[1,1]

. gen share = polar/polar_all*100

. list share in 1

. restore
```

Bootstrap inference for two values of alpha

```
. xtset, clear

. tempvar tempid

. bootstrap p_a1=el(r(itempolar),1,1) p_a16=el(r(itempolar),2,1), ///
cluster(workerid) idcluster( `tempid' ): itempolar income, ///
idvar( `tempid' ) tvar(year) alphas(1 1.6)
```

Bayesian (or weighted) bootstrap inference

```
. exbsample 100, stub(rw) cluster(workerid)

. bs4rw polar = el(r(itempolar),1,1), rweight(rw1-rw100) nodots: ///
itempolar income, idvar(workerid) tvar(year) alphas(1.6)
```

Authors

Francesca Subioli
Department of Economics and Law, Sapienza University of Rome, Italy
Contact: francesca.subioli@uniroma1.it

Philippe Van Kerm
Department of Social Sciences, University of Luxembourg
Luxembourg Institute of Socio-Economic Research (LISER), Luxembourg
Contact: philippe.vankerm@uni.lu, philippe.vankerm@liser.lu

References

Subioli, F. and P. Van Kerm (2022), *Inter-temporal income polarization*.

Esteban, J. and D. Ray (1994), *On the measurement of polarization*, Econometrica, 819–851.

Jann, B. (2007), *Univariate kernel density estimation*, Boston College Department of Economics, Statistical Software Component, (S456410).

B.2.2 Ado file


```
*** ADO FILE FOR INTER-TEMPORAL POLARIZATION ***
* November 2022
* Authors: Francesca Subioli & Philippe Van Kerm

program define itempolar, rclass sortpreserve byable(recall) properties(svyb svyj)
    syntax varname(numeric) [if] [in] [pw aw fw iw], ///
    [IDvar(varname numeric) ///
    Tvar(varname numeric) ///
    Alphas(numlist >=0) ///
    RA(real 0.5) ///
    RI(real 0.5) ///
    RT(real -1) ///
    AT(real -1) ///
    RTVAR(varname numeric) ///
    ATVAR(varname numeric) ///
    WTYPE(string) ///
    gen_i(string) ///
    gen_a(string) ///
    gen_ea(string)]

    // if/in:

    marksample touse

    /////////////////////////////
    ////   Validity checks   ////
    /////////////////////////////

    // id and time variables

    if ("`idvar'" != "") {
        if ("`tvar'" == "") {
            di as error "Define time variable tvar()"
            exit 198
        }
        else {
            local id `idvar'
            local t `tvar'
        }
    }
    else {
        _xt, trequired
        local id `r(ivar)'
        local t `r(tvar)'
    }

    // Discount rates

    if (`ra' < 0 | `ri' < 0) {
        display as error "Negative discount rate not allowed"
        exit 198
    }

    // Balanced panel

    tempvar N
    qui: bys `id' `touse': gen `N' = _N if `touse'
    qui: sum `N', meanonly
    if (`r(min)' < `r(max)') {
        display as error "The panel must be balanced"
        exit 198
    }

    // Missing values

    qui: count if `touse' & `varlist' == .
    if (`r(N)' != 0) {
        display as error "Missing values not allowed"
```

itempolar - Printed on 14/11/2022 16:32:37

```
        exit 198
    }

    // Negative values

    qui: count if `touse' & `varlist' < 0
    if (`r(N)' != 0) {
        display as error "Negative values not allowed"
        exit 198
    }

    // Zeros

    qui: count if `touse' & `varlist' == 0
    if (`r(N)' != 0) {
        display as text "Warning: `r(N)' zero values used in the computations"
    }

    // Alpha

    if ("`alphas'" == "") local alphas "1" // Default

    // Identification window width

    // If more than one option specified, exit
    if ( ( (`rt'>=0) + (`at'>=0) + ("`rtvar'"!="") + ("`atvar'"!="") ) >1 ) {
        di as error "Options rt, at, rtvar and atvar are mutually exclusive"
        exit 198
    }

    qui: summ `t' if `touse', meanonly
    local T `r(max)'

    // If no window option specified, set rt=0.2 as default
    // If negative fraction of income, exit
    if ( ( (`rt' >= 0) + (`at' >= 0) + ("`rtvar'" != "") + ("`atvar'" != "") ) == 0 ) {
        di as text "Window width for group identification set to rt=0.2"
        local rt = 0.2
    }
    else if ("`atvar'" != "") {
        summ `atvar' if `touse' & `t' == `T' , meanonly
        local Nat `r(N)'
        local atmin `r(min)'
        local atmax `r(max)'
        qui: count if `touse' & `t' == `T'
        local N = `r(N)'
        if (`atmin' < 0 | `atmax' > 1) {
            di as error "`atvar' must be between 0 and 1"
            exit 198
        }
        if (`Nat' != `N') {
            di as error "The identification window must be defined for every id"
            exit 198
        }
    }
    else if ("`rtvar'" != "") {
        summ `rtvar' if `touse' & `t' == `T' , meanonly
        local Nrt `r(N)'
        local rtmin `r(min)'
        local rtmax `r(max)'
        qui: count if `touse' & `t' == `T'
        local N = `r(N)'
        if (`rtmin' < 0 | `rtmax' > 1) {
            di as error "`rtvar' must be between 0 and 1"
            exit 198
        }
        if (`Nrt' != `N') {
            di as error "The identification window must be defined for every id"
```

Page 2

itempolar - Printed on 14/11/2022 16:32:38

```
            exit 198
        }
    }
    else if (`rt' != -1) {
        if (`rt' < 0 | `rt' > 1) {
            di as error "The threshold for identification must be between 0 and 1"
            exit 198
        }
    }
    else if (`at' != -1) {
        if (`at' < 0 | `at' > 1) {
            di as error "The threshold for identification must be between 0 and 1"
            exit 198
        }
    }

    // Weighting scheme
    if ("`wtype'" != "" & "`wtype'" != "rectangular" & "`wtype'" != "triangular" & "`wtype'" != "triweight") {
            di as error "Argument of wtype() not allowed"
            exit 198
    }
    if ("`wtype'" == "") local wtype "triweight"

    // Sample weights

    if ("`weight'" == "") {
        tempvar wvar
        qui: gen byte `wvar' = 1 if `touse'
    }
    else {
        tempvar wvar
        qui: gen `wvar' `exp' if `touse'
    }

    // Variables to generate

    if ("`gen_i'" != "" | "`gen_a'" != "" | "`gen_ea'" != "") local genopt "YES"
    else local genopt "NO"
    if !_by() | (_by() & _byindex()==1) {
        foreach variable in `gen_i' `gen_a' `gen_ea' {
            if ("`variable'" != "") {
                capture confirm var `variable'
                if !_rc {
                    di as error "Variable `variable' already exists"
                    exit 198
                }
            }
        }
    }

    //////////////////////
    ////   Commands   ////
    //////////////////////

    // Generate a vector of periods (1, .., T)
    qui: levelsof `t' if `touse'
    local max `r(r)'
    tokenize `r(levels)'
    tempname tmat
    matrix `tmat' = J(`max', 1, .)
    forvalues i = 1/`max' {
        matrix `tmat'[`i',1] = ``i''
    }

    // Mean-normalize income
    tempname y
    gen `y' = `varlist' if `touse'
```

```
    forvalues i = 1/`max' {
        qui: summ `varlist' if `touse' & `t'==``i'' [aw = `wvar'], meanonly
        qui: replace `y' = `y'/`r(mean)' if `touse' & `t'==``i''
    }

    sort `touse' `id' `t'

    // Apply Mata function

    mata : itempolarfunction("`y'", "`wvar'", "`id'", "`touse'", "`tmat'", `T', ///
        "`alphas'", `ra', `ri', "`wtype'", `rt', `at', "`rtvar'", "`atvar'", "`genopt'")

    if _by() {
        local num = _byindex()
        scalar ident`num' = ident
        scalar alien`num' = alien
        matrix polar`num' = polar
    }

    // Display results

    display ""
    display "Inter-temporal Polarization Index"
    tempname polartrans
    matrix `polartrans' = polar`num''
    matrix list `polartrans', noblank noheader nonames format(%10.3f)

    // Store results

    if !_by() {
        return scalar itemident = ident
        return scalar itemalien = alien
        return matrix itempolar = polar
        scalar drop ident alien
    }
    else if (_by() & _bylastcall()) {
        matrix ident = ident1
        matrix alien = alien1
        matrix polar = polar1
        local last = `num'
        forvalues i = 2/`num' {
            matrix ident = ident, ident`i'
            matrix ident = ident, alien`i'
            matrix ident = ident, polar`i'

        }
        return matrix itemident = ident
        return matrix itemalien = alien
        return matrix itempolar = polar
        forvalues i = 1/`num' {
            mat drop polar`i'
            scalar drop ident`i' alien`i'
        }
    }

    // Generate individual-level variables

    if ("`genopt'" == "YES") {
        qui {
            tempname idfake
            egen `idfake' = group(`id') if `touse'
            tempfile originaldata
            save "`originaldata'", replace
            keep if `touse'
            keep `id'
            sort `id'
            duplicates drop
            gen `idfake' = _n
```

```
            if ("`gen_i'" != "" | "`gen_ea'" != "") {
                svmat Ii, name("`gen_i'")
                if _by() {
                    rename `gen_i'1 `gen_i'`num'
                    gen `gen_i' = .
                }
            else rename `gen_i'1 `gen_i'
            }
            if ("`gen_a'" != "" | "`gen_ea'" != "") {
                svmat Ai, name("`gen_a'")
                if _by() {
                    rename `gen_a'1 `gen_a'`num'
                    gen `gen_a' = .
                }
                else rename `gen_a'1 `gen_a'
            }
            if ("`gen_ea'" != "") {
                tokenize `alphas'
                local count: word count `alphas'
                if (`count' != 1) {
                    forvalues v = 1/`count' {
                        if _by() {
                            gen `gen_ea'_`v' =.
                            gen `gen_ea'`num'_`v' = (`gen_i'`num'^``v'')*`gen_a'`num'
                        }
                        else gen `gen_ea'_`v' = (`gen_i'^``v'')*`gen_a'
                    }
                }
                else {
                    if _by() {
                        gen `gen_ea' =.
                        gen `gen_ea'`num' = (`gen_i'`num'^`alphas')*`gen_a'`num'
                    }
                    else gen `gen_ea' = (`gen_i'^`alphas')*`gen_a'
                }
            }
            tempfile newvars
            qui: save "`newvars'", replace
            use "`originaldata'", clear
            qui: merge m:1 `id' `idfake' using "`newvars'", nogen
            if _by() & _bylastcall() {
                forvalues i = 1/`num' {
                    replace `gen_i' = `gen_i'`i' if `gen_i'`i' != .
                    replace `gen_a' = `gen_a'`i' if `gen_a'`i' != .
                    drop `gen_i'`i' `gen_a'`i'
                    if (`count' != 1) {
                        forvalues v = 1/`count' {
                            replace `gen_ea'_`v' = `gen_ea'`i'_`v' ///
                                if `gen_ea'_`v' ==. & `gen_ea'`i'_`v' != .
                            drop `gen_ea'`i'_`v'
                        }
                    }
                    else {
                        replace `gen_ea' = `gen_ea'`i' if `gen_ea' ==. & `gen_ea'`i' != .
                        drop `gen_ea'`i'
                    }
                }
            }
        }
    }
end

///////////////////////////
////   MATA FUNCTIONS   ////
///////////////////////////

mata:

```

```
// Function to get an NxN matrix of weighted distances starting from a matrix
// of inter-temporal distances and an Nx1 vector defining the identification window

real matrix distweight(real matrix D, real matrix bw, string scalar type) {

    if (type == "rectangular") {
        return(D :<= bw)
    }
    else if (type == "triangular") {
        if (sum((bw) :== 0) == 0) {
            A = 1 :- (D :/ bw)
            return((A :>= 0) :* A)
        }
        else {
            A = 1 :- editmissing((D :/ bw), 0)
            return((A :>= 0) :* A)
        }
    }
    else if (type == "triweight") {
        if (sum((bw) :== 0) == 0) {
            A = ((1 :- (D :/ bw) :^ 2) :^3)
            return((A :>= 0) :* A)
        }
        else {
            A = ((1 :- editmissing((D :/ bw), 0) :^ 2) :^3)
            return((A :>= 0) :* A)
        }
    }
}

void itempolarfunction(string scalar y, string scalar w, string scalar id, ///
        string scalar touse, string scalar tmatrix, real scalar T, string scalar alphalist, ///
        real scalar ra, real scalar ri, string scalar type, real scalar rt, ///
        real scalar at, string scalar rtvar, string scalar atvar, string scalar genopt) {

    st_view(ID, ., id, touse)
    st_view(YW, ., (y, w), touse)
    if (rtvar != "") st_view(B, ., rtvar, touse)
    if (atvar != "") st_view(B, ., atvar, touse)
    info = panelsetup(ID, 1)
    SW = YW[info[., 1], 2]
    if (rtvar != "" | atvar != "") B = B[info[., 2], .]
    alphas = (strtoreal(tokens(alphalist)))'
    tmat = st_matrix(tmatrix)
    riw = exp(-ri :* (T :- tmat)) // (T :- tmat) is a vector of lags s from T
    raw = exp(-ra :* (T :- tmat)) // riw and raw are vectors of discount factors

    Di = J(rows(info), rows(info), 0) // NxN matrix of zeros
    Y = J(rows(info), 1, 0) // Nx1 vector of zeros

    // Matrix of symmetric inter-temporal distances
    // Accumulation for subsequent periods, discounting past distances

    for(t = 0; t <= rows(tmat)-1; t++) {
        if (riw[t + 1, 1] > 0.001) {  // do calculation only if discount factor above .001
            Di = Di :+ (abs(lowertriangle((J(1, rows(info), YW[info[.,1] :+ t, 1]) ///
                        :- (J(1, rows(info), YW[info[.,1] :+ t, 1]))')')) :* riw[t + 1, 1])
            // Inter-temporal income
            Y = Y :+ (YW[info[.,1] :+ t, 1] :* riw[t :+ 1, 1])
        }
    }
    Di = Di :/ colsum(riw)
    Di = makesymmetric(Di)
    Y = Y :/ colsum(riw)

    // If ra != ri, compute also the matrix of inter-temporal distances
    // discounted through ra

```

```
        if (ra != ri) {
            Da = J(rows(info),rows(info),0)
            for(t = 0; t <= rows(tmat)-1; t++) {
                if (raw[t + 1, 1] > 0.001) {
                    Da = Da :+ (abs(lowertriangle((J(1, rows(info), YW[info[.,1] :+ t, 1]) ///
                        :- (J(1, rows(info), YW[info[.,1] :+ t, 1]))')))) :* raw[t + 1, 1])
                }
            }
            Da = Da :/ colsum(raw)
            Da = makesymmetric(Da)
        }

        // Bandwidth matrix
        if (rtvar != "") {
            B = B :* Y
            bw = B
        }
        else if (atvar != "") {
            B = B :* mean(Y, SW)
            bw = B
        }
        else if (at != -1) {
            bw = J(rows(info), 1, at*mean(Y, SW))
        }
        else if (rt != -1) {
            bw = rt :* Y
        }

        Ii = mean(distweight(Di, bw, type)', SW)'

        if (ra != ri) {
            Ai = Da*SW :/ colsum(SW)
        }
        else {
            Ai = Di*SW :/ colsum(SW)
        }
        polar = J(rows(alphas), 1 , .)

        for(i = 1; i <= rows(alphas); i++) {
            polar[i,.] = mean(Ii :^ alphas[i,] :* Ai, SW)
        }

        // Rescaling

        for(i = 1; i <= rows(alphas); i++) {
            if (alphas[i,.] > 1) polar[i,.] = polar[i,.] ^ (1/alphas[i,.])
        }

        st_numscalar("ident", mean(Ii, SW))
        st_numscalar("alien", mean(Ai, SW))
        st_matrix("polar" , polar)

        if (genopt == "YES") {
            st_matrix("Ii" , Ii)
            st_matrix("Ai" , Ai)
        }
    }
    end
```

Chapter 3

Labour market dynamics and geographical reallocation

JEL Codes: J23, J61, R23, J63.
Keywords: Labour demand, Turnover, Layoff, Geographic labour mobility.

Notes and Acknowledgements

This chapter is joint work with Gaetano Basso (Bank of Italy) and Salvatore Lo Bello (Bank of Italy). I thank them and the Households and Labour Market Division of the Directorate General for Economics, Statistics and Research of the Bank of Italy for hosting me for one year as a research assistant and for six months of collaboration as a Stringher intern. I especially thank Salvatore for the time and dedication that made my internship possible. The views expressed in this chapter are those of the authors and do not necessarily reflect those of the Bank of Italy or of the Eurosystem. The Statistical Information System of Compulsory Communications (SISCO) database was obtained thanks to an agreement between the Bank of Italy and the Italian National Agency for Active Labour Market Policies (ANPAL). I thank for their comments Federico Cingano, Francesco D'Amuri, Greg Howard, Marco Leonardi, Giovanni Peri, Eliana Viviano and participants at the Bank of Italy seminar series, the 2022 European Meetings of the Urban Economic Association, the Bank of Canada Workshop on Macroeconomic Implications of Migration, the 2022 EALE conference, and the XII fRDB Workshop.

3.1 Introduction

Understanding the responsiveness of the geographical allocation of workers to local labour market dynamics is a first-order issue in the economic literature. Indeed, it has long been argued that migration is a major mechanism to absorb labour demand variations and that people move across regions (Blanchard and Katz, 1992) or change their commuting behaviour (Monte et al., 2018) in response to employment opportunities. There is evidence of such empirical regularities in many contexts. However, labour market flows are correlated across space and time and so do migration flows, posing challenges to the identification of the causal nexus. For this reason, most of the literature focuses on specific labour demand events such as mass layoffs (Gathmann et al., 2020; Foote et al., 2019), disruptions due to international trade (Autor et al., 2013), major construction events (Carrington, 1996) or the Great Recession (Monras, 2018), often taking a reduced form approach that hardly allows generalizing the results.[1] Moreover, the literature largely overlooks whether the migration response to positive and negative demand variation is symmetric, often as a consequence of the identification strategies adopted.[2]

In this paper, we try to overcome these limitations by jointly estimating the causal impact of both positive and negative local labour demand shocks on internal migration flows. To do so, we exploit plausibly exogenous variation stemming from mass events at the establishment level. Our estimates reveal important effects of both job creation and job destruction on net migration, with the former being much larger than the latter. In particular, job creation strongly stimulates the in-migration rate, whereas job destruction has a milder effect on the out-migration rate, which appears as a less responsive margin of adjustment. Moreover, each margin has a different geographical reach, with important implications in terms of policy prescriptions. Overall, our findings highlight the importance of separately assessing the contemporaneous impact of different-sign labour demand shocks on migration patterns.

Our study uses administrative data on the universe of labour market transitions in Italy, covering the period 2010-2018. The high quality of the data allows us to precisely identify mass events at the micro level, and to track their effects on aggregate migration flows via job creation and job destruction – defined ad the sum of job flows net of establishment-level churning.[3] Our results are particularly interesting if one considers that they apply to what is typically thought to be a low-dynamism economy.[4] Surprisingly, we uncover a large amount of labour market dynamism and a substantial degree of responsiveness of internal migration to employment opportunities.

More in general, we contribute to the existing literature in two ways. First, we account

[1] The literature on the response of migration to labour demand shocks is vast and it is not fully reported here. A subset of the literature also looks at international migration responses (Beyer and Smets, 2015; Basso et al., 2019). For the Italian case, see Ciani et al. (2019). For an overview of internal migration and of the related literature, see Molloy et al. (2011).

[2] Notable exceptions are Ciani et al. (2019) and Notowidigdo (2020).

[3] As the focus of our analysis is labour demand, we adopt the definition of Davis and Haltiwanger (1992); Davis et al. (1998, 2006) that has been shown to capture well business cycle dynamics.

[4] As documented by Elsby et al. (2013), in Italy the probability of flowing into and out of employment is the smallest, among OECD countries.

for both job creation and job destruction when estimating adjustments to demand shocks, ensuring identification by instrumenting both margins of labour market dynamism. The literature, instead, typically takes a reduced-form approach, such as estimating region-level mass-layoffs, or focuses on one margin at a time. We show that all margins of job creation and destruction, properly instrumented, matter in explaining gross and net migration flows. Second, we uncover that migration adjustments are stronger for positive than for negative variations in labour demand, in line with the evidence on the US (Notowidigdo, 2020) and differently from what previously shown for Italy (Ciani et al., 2019). Migration adjustments are made up of the response of inflows and outflows. We find that the asymmetry between the effect of positive versus negative shocks is mainly due to the greater amount of inflows generated by job creation than the outflows (or reduced inflows) spurred by job destruction.

In the first part of the paper (Section 3.2), we use the detailed microdata to document novel facts on labour market flows and internal migration rates at the local level.[5] In particular, in a symmetric fashion for job and migration flows, we find that: i) the magnitude of gross flows dwarfs that of net flows, implying a large degree of excess turnover , both at the aggregate and at the local level; ii) gross and net flows feature systematic differences across space and over time; iii) both gross job creation and job destruction (in-migration and out-migration) are important determinants of aggregate fluctuations in the employment growth (net in-migration) rate.

More precisely, we find that the average excess turnover rate – the sum of gross flows over the absolute value of net flows – is 38,7 (107,6) for job (internal migration) flows at the province-year level.[6] To fix ideas, this implies for instance that when a region experiences a unit increase in employment, total gross job flows (i.e. job creation plus job destruction) amount to about 39, on average. This highlights the existence of marked differences in the magnitude of gross and net flows, suggesting that it is important to study them separately. Moreover, we document that gross flows are especially concentrated in regions where net flows are relatively small, and that their dynamics over time are very different. Again, these facts indicate that these two margins capture different aspects of worker transitions. Importantly, very similar patterns arise for both job flows and internal migration, suggesting a link between the two. Finally, we decompose the variance of net flows into each gross flow component to gauge the importance of each gross flow in determining overall employment (migration) fluctuations. We find that, on average, gross flow rates account for roughly 50 per cent of the aggregate fluctuations in net rates over time. Summing up, two important lessons can be drawn from this descriptive evidence. First, gross flows starkly differ from net flows, both in terms of levels and dynamics. Second, it is key to account for both job

[5]Our baseline geographical level is the province, but we also explore the heterogeneity at the level of region (more coarse) and municipality (more disaggregated).

[6]These rates vary greatly depending on the time and the geographical aggregation: they range between 5 and 14 (5 and 22) at the municipality level, between 17 and 39 (32 and 108) at the province level, and 16 and 40 (40 and 413) at the region level.

creation and job destruction (in-migration and out-migration), as they are both important drivers of employment (migration) growth. Consistently, in the remainder of our paper, we focus our attention separately on gross and net flows and explore potential differences between the effects on internal migration of positive and negative labour demand shocks.

Following up on the descriptive evidence, in Section 3.3 we relate internal migration to labour market dynamism. In particular, our objective is to analyze the contemporaneous effects of both job creation and job destruction. Simple OLS regressions indicate that both labour market flows are highly correlated with migration flows and that the magnitudes of such correlations are symmetric: a 1 percentage point increase in the job creation rate (job destruction rate) is associated with an increase (decrease) in the net migration rate of 0.10 percentage points. To overcome plausible endogeneity and identify causal effects (as the causality nexus between migration flows and job creation and destruction is not a priori determined), we then turn to an instrumental variable identification strategy.

We exploit changes in labour demand spurred by mass hire and layoff events at the establishment level, which we identify directly from the administrative data. We define mass events as those that involve more than 250, 500 or 1,000 workers. Even though the Italian labour market is characterized by institutions aiming at preserving employment relations (e.g., short-time work schemes) and delaying layoffs, we show that such events are quite salient (they involve between 0.1 and 6.9 per cent of the province-year employment) and spread across the country (depending on the threshold, they occur in 9 up to 69 provinces over the sample period). Most importantly, regression results indicate that they are unanticipated, a major condition for exogeneity to hold, and have strong predictive power on aggregate job flows. Moreover, there are no cross-effects (mass layoff on job creation rate and mass hire on job destruction rate), indicating that we are able to separately identify job creation and job destruction.

The 2SLS estimates provide interesting evidence. First, the magnitude of the job creation effect on net migration is about three times larger than in the OLS estimates (a 1 percentage point increase in the job creation rate leads to a .3 percentage points increase in the in-migration rate), whereas that of job destruction is about one-sixth in magnitude (.05). These differences can be traced back to the specific gross flows involved in the adjustment, as job creation is able to generate more worker flows than job destruction. Second, similarly to the literature on US and Germany (Gathmann et al., 2020), our evidence shows that movements out of a region are muted relative to inflows: however, we still find outflows to be non-negligible following negative shocks in labour demand, differently to what was found in previous papers (Monras, 2018; Notowidigdo, 2020). Third, the 2SLS estimates of the cross-effects have the expected sign – that is, job creation (job destruction) is negatively correlated with out-migration (in-migration) – contrary to the OLS. For instance, this occurs whenever positive labour demand shocks prevent people from leaving an area (though the estimates are not statistically significant).

We inspect the mechanics of the in-migration reaction by decomposing the overall inflow into two separate terms. The first one is a measure of *potential in-migrants*, i.e. the sum of workers relocating from all other locations, while the second one captures how attractive a given location is relative to all other competing destinations. By running separate regressions, we document that job creation rate shocks do *not* cause more individuals to relocate, but rather they cause relocating individuals to revise their ranking over possible destination alternatives. We see this as a very important finding for a number of reasons. First, this is a crucial piece of information to correctly model the reaction of migration choices to labour demand shocks. Second, this implies the existence of negative externalities of region-specific shocks to other competing regions, which work by altering their relative attractiveness. Third, we claim that this sheds a fundamental light on the asymmetry between the effect of positive shocks on in-migration – i.e., very large – and that of negative ones on out-migration – i.e., much smaller. Our evidence points to the fact that the effect of positive shocks on in-migration operates through a margin - the relative attractiveness of the location – that, by construction, does not exist for out-migration.

Moreover, we also study the geographical reach of the effects of labour demand shocks on internal migration. Binning the migration flows by distance, we show that the out-migration rate responds only locally to a change in job destruction rate (up to 50 km), while in-migration flows increase in response to job creation rate with a much larger reach, though with a decaying intensity over space. This is a relevant finding that speaks directly to the consequences of shocks on spatial inequality. Indeed, the welfare gains brought about by positive shocks are shared with relatively large inflows of migrants, who cover about 30 per cent of new jobs and act as a counteracting equilibrium force. In addition, our results provide useful insights into the extent to which labour markets are actually local as opposed to perfectly integrated, in the spirit of Manning and Petrongolo (2017).

Finally, we perform a heterogeneity analysis to study whether specific characteristics of the locations are associated with a higher or lower response of migration flows. Interestingly, we find that the reaction of out-migration to negative shocks is actually large and significant in the Center-North of Italy and in locations with a relatively high incidence of foreign-born and college graduates. On the contrary, a large share of young individuals or high values of the homeownership rate is associated with a sluggish reaction. This is consistent with existing evidence on heterogeneity in migration responses by demographic characteristics and homeownership status (for reviews, see Basso and Peri (2020) and Jia et al. (2022)).

We conclude the paper (Section 3.4) by laying out the policy implications of our results. First, we claim that the design of active labour market, social and housing policies should take into account the extent of gross migration flows, to minimize the potential congestion and frictions that may arise after a labour demand shock. Moreover, our results indicate that policymakers may want to act more aggressively against negative labour demand shocks – as internal migration does not help much at absorbing them – in order to mitigate the potential short-run increase in cross-regional disparities. On the contrary, the consequences of job creation are already largely shared across space due to the strong and quick reaction of in-migration even at long distances. Hence, they contribute to a lesser extent to generating spatial inequality.

3.2 Descriptive evidence

In this Section, we use highly detailed microdata to compute aggregate labour market flows at the local level. In order to do so, we need to fix a sampling interval and a geographical aggregation level. In our empirical analysis based on quasi-random variation (Section 3.3), we adopt a province-year level specification for reasons that will be clarified later. However, for most of the descriptive analysis carried out in this Section, we also explore the heterogeneity stemming from other possible aggregation levels: month and quarter for the time aggregation,[7] municipality and region for the geographical aggregation. We do this to lend further support to our findings, and to guarantee that they are not simply a by-product of our own arbitrary choices.

3.2.1 Data

SISCO data. The data we use is a selection – from January 2010 to December 2018 – from the Statistical Information System of Compulsory Communications (SISCO). The SISCO database contains all the employee and para-subordinate work relationships that have undergone an event (activation, transformation, extension, termination) since March 2008.[8] Hence, the resulting database covers the universe of labour market flows for all types of contracts and employers.[9] Overall, the dataset contains information on 119.1 million labour contracts, involving about 22 million workers throughout the sample period.[10]

A job is defined as a contractual relationship identified by the employee, the employer and the activation date, and contains all the subsequent events (extensions in the case

[7]Gomes (2015) and Bertheau and Vejlin (2022) have recently shown the importance of this source of bias when measuring labour market transitions based on individual-level data.

[8]Since that date, each hiring, separation, contract renewal and contract transformation is collected for administrative purposes by the Italian Ministry of Labour through an online communication system named 'comunicazione obbligatoria' (CO) to be filed by the employers at the time of the event. The sample at our disposal is reduced to the years 2010-2018.

[9]SISCO includes all public and private sector jobs including maids and caregivers hired by households and excludes only employment relationships in the armed forces and those involving senior figures such as presidents and CEOs of public and private companies.

[10]If we weigh the number of individuals by the time spent in employment (i.e. assigning a weight of 1 to workers who are employed continuously throughout a period), the average number of workers for which we have information is about 9.1 million per year.

of fixed-term contracts, transformations to open-ended contracts, terminations). The information collected and made available to the researchers is very detailed at the level of worker (including the municipality of residence), job (length and type of contract) and employer (5-digit sector, municipality of the production unit).[11] SISCO data are advantageous with respect to other administrative sources used to analyze the labour market – such as those collected by the Italian Social Security Administration (INPS) for social security contribution purposes – as they cover the universe of labour market flows and have additional detailed information at the worker and job level.[12] However, SISCO data do not collect any information on earnings or salary. Moreover, due to the particular structure of the data, it is not possible to observe workers who have not experienced any contractual event between January 1, 2010 and December 31, 2018. In other words, we do not have information on pre-existing stocks, i.e. these data cannot be used to assess employment stocks. We circumvent this issue by taking estimates of these stocks from external sources, namely from the Italian Labour Force Survey (ILFS). This allows us to consistently construct flow rates.

ILFS data. The Italian Labour Force Survey (ILFS) is a sample survey conducted by the Italian National Statistical Institute (ISTAT) by interviewing each year since 2004 more than 250,000 households resident in Italy (for a total of 600,000 individuals between 15 and 89 years of age), distributed in about 1,400 Italian municipalities and representative of the resident population.[13] It represents the primary source of statistical information on the Italian labour market, harmonised at the European level, and it is used for the official estimates of unemployment and the main aggregates of labour supply. For the purposes of this paper, ILFS data are used only to compute the aggregate stocks of payroll employment (as well as population and unemployment, for our robustness checks), to be used as denominators for computing rates.

3.2.2 Labour Market Dynamism

The forces behind labour market transitions can be grouped into two broad categories, commonly used in the analyses carried out by the existing literature (e.g. Davis and

[11] The worker's residence is only available for the last job recorded. Therefore, residence changes are not registered in the SISCO data. Since the SISCO data are collected continuously, the classification of the municipalities is not coherent in the whole period because of some mergers and abolition that took place between 2010 and 2018. We bring all the municipality codes to the classification in force at the end of 1 January 2019. The only geographical shifts that we cannot adequately deal with are transfers of 'districts' from one municipality to another. For 0.20% of the contracts, it was not possible to attribute the correct municipality code because of irrecoverable errors in the data.

[12] INPS data do not cover maids, caregivers and agricultural workers in the private sector or most of the public sector employment. Data on on-call and employment-agency-hired temporary workers are treated separately from other temporary workers.

[13] The survey is carried out during all weeks of the year using a uniform distribution of the sample over the weeks. The sampled households are interviewed four times over 15 months: each one is interviewed for two consecutive quarters and two more quarters after a two-quarter break.

Haltiwanger (1992); Davis et al. (1998, 2006)). On the one hand, employers create new jobs and destroy old ones every period, thus affecting the distribution of jobs across space (*demand-side*). On the other hand, for a given distribution of jobs, workers switch jobs and change employment status because of *supply-side* events (e.g., relocation, labour force entry, migration, retirement, death, change in preferences). As the focus of this paper is on the impact of labour demand shocks, we analyze exclusively job flows.[14]

Let $E_{r,t}$ be the employment level of location r at time t (where location can be a municipality, a province, a region, etc.). At any level of aggregation, the net change in employment between two points in time ($\Delta E_{r,t} = E_{r,t} - E_{r,t-1}$) satisfies the following accounting identity:

$$\Delta E_{r,t} \equiv \underbrace{JC_{r,t} - JD_{r,t}}_{\text{Job flows}} \equiv \underbrace{NJC_{r,t}}_{\text{Net Job Creation}} ,$$

where $JC_{r,t}$ denotes job creation and $JD_{r,t}$ denotes job destruction.[15] In turn, job creation is defined as the sum of net employment gains over all establishments that either expand or start up within a given time interval. In a symmetric fashion, job destruction is defined as the sum of net employment losses over all establishments that either contract or shut down in the time interval:

$$JC_{r,t} = \sum_{i \in \mathcal{G}_t} \Delta e^i_{r,t}, \qquad\qquad JD_{r,t} = -\sum_{i \in \mathcal{S}_t} \Delta e^i_{r,t},$$

where $\Delta e^i_{r,t}$ denotes the employment change at establishment i, $\mathcal{G}$ denotes the set of growing establishments ($\{i \in \mathcal{G}_t : \Delta e^i_{r,t} > 0\}$) and $\mathcal{S}$ denotes the set of shrinking establishments ($\{i \in \mathcal{S}_t : \Delta e^i_{r,t} < 0\}$). We define an establishment as the combination of a firm identification number and the municipality of the workplace, i.e., establishments of the same firm in a given municipality are pooled together.[16]

In order to obtain rates, we divide the absolute flows by the current periods' corresponding stocks of payroll employment $E_{r,t}$ as estimated from ILFS data.[17] Eventually, we define $JCR_{r,t}$ and $JDR_{r,t}$ respectively as the job creation and job destruction rate. Last, we define the excess job turnover rate $EJTR$ as the ratio between the job turnover rate JTR – that

[14]Note that job flows do *not* correspond to the aggregate gross job creation and job destruction as in the Diamond-Mortensen-Pissarides class of models (Mortensen and Pissarides, 1994). In a separate research project, we plan to investigate extensively the patterns of job and worker flows – the latter being aggregate hires and separations – both at the local and at the establishment level.

[15]Throughout the paper, we refer to the increase in the number of active (i.e., not expired) labour market contracts as the change in payroll employment. In principle, this may potentially be imprecise, if workers hold multiple jobs. However, in our data about 94% of contracts do not overlap with any other active contract at the same time. Hence, we conclude that this does not represent a major issue for our measurement.

[16]We need to do this because our data do not contain an establishment identifying number. In practice, we claim that our proxy for establishments is very precise, as we know from restricted-access ISTAT data, not available for research purposes, that only about 2% of firms have multiple establishments within the same municipality in the sample period.

[17]As already mentioned, the SISCO data cannot be used to derive measures of total stocks.

is the sum of the job creation and job destruction rates – and the absolute value of the net job creation rate:

$$EJTR_{r,t} = \frac{JCR_{r,t} + JDR_{r,t}}{|NJCR|_{r,t}}.$$

This synthetic measure captures how large the differences between gross and net flows are. Note that its minimum value is 1, and that it is large when gross flows are mostly offsetting (that is, when they roughly cancel one another, resulting in small net flow rates). At the aggregate level, excess turnover rates capture both cross-region and cross-firm reallocation, whereas at the local level these indexes capture uniquely the reallocation occurring across firms.

Aggregate flows. We first examine labour market flows rates at the national level, obtained by summing up events across locations and dividing the total by the aggregate stocks. Figure 3.1 (panel a) plots the average of yearly flows over the period 2010-2018, revealing that gross flows are generally much larger than net flows, i.e. excess turnover is high. When we split gross flows by contract type, we find that temporary jobs play a crucial role: despite representing only about 14% of the total stock of payroll employment on average during the period 2010-2018, they account for 40-45% of job flows (Figure C.1, panel a).[18] Turning to the evolution over time, we notice that the job creation and job destruction rates tend to negatively co-move (Figure 3.1, panel b). In practice, gross job flow rates seem to be affected by aggregate shocks, following the business cycle as expected, that is job creation (destruction) is high (low) in expansionary phases and low (high) during recessions.[19] This is also confirmed by the correlation between these variables and the net employment growth rate: job creation has a correlation coefficient of 0.70, while the job destruction rate of -0.84 (see Table 3.1).

Table 3.1. Correlation matrix

	JCR	JDR	NJCR (NHR)
JCR	1		
JDR	-0.197	1	
NJCR (NHR)	0.695	-0.842	1

Source: SISCO and ILFS data 2010-2018. *Note*: The table shows the correlation matrix of yearly aggregate job flow rates for Italy. Job flow rates are the sum across establishments of net activations at the establishment-level, divided by the stock of payroll employment in the current period taken from the ILFS data. All correlations are statistically significant at the 99% level.

[18]Temporary contracts are even more prevalent for worker flows, i.e. total hires and separations, representing the lion's share of such flows (82-84%); see Figure C.1, panel b. This suggests that a large chunk of hires and separations involving temporary contracts simply represents intra-establishment churn.

[19]Our sample period covers only the 2011-2013 crisis.

Figure 3.1. Job flows in Italy

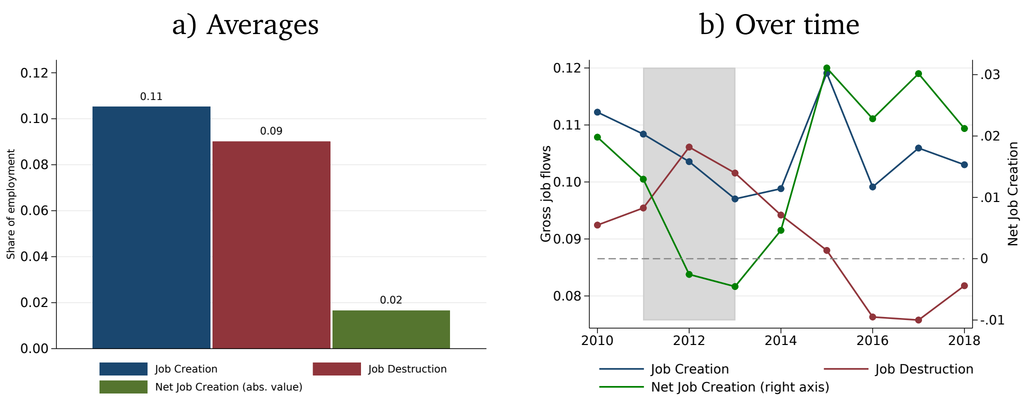

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows averages (panel a) and the time series (panel b) of yearly job flow rates for Italy. Job flow rates are the sum across establishments of net job activations at the establishment-level, divided by the stock of payroll employment in the current period taken from the ILFS data. The shaded area represents the 2011-2013 recession.

Location-specific flows. We now turn to the analysis of location-specific flows. Figure 3.2 plots the geographical distribution of average job flows at the province level. It is immediately apparent that these flows are highly correlated across space. In particular, we find that Southern provinces are characterized by a remarkably high level of gross flows, apparent for both job creation and job destruction. However, many of these flows are almost exactly offsetting each other, so that net job creation is concentrated in the Northern provinces, with the exception of some other specific locations. When we investigate these patterns distinguishing jobs by contract type (temporary vs. open-ended), we find that the larger degree of labour market dynamism (job creation and destruction) in Southern regions was mainly due to the dynamics of open-ended contracts, perhaps surprisingly (see Figure C.2). This was likely due to a number of policy interventions that targeted those regions with subsidies that incentivized the creation of open-ended positions during our sample period (Camussi et al., 2022). Instead, net job creation was by and large driven by temporary contracts, that expanded especially in the North.

Table 3.2 shows summary statistics of job flows for different geographical (municipality, province and region) and time (month, quarter, year) aggregation levels. Overall, it confirms the patterns already uncovered at the aggregate level, namely that gross flows are much larger than net flows, implying high levels of excess turnover. In particular, depending on the geographical and time aggregation level, average excess turnover rates range between 5.2 and 40.2. For instance, this implies that, on average, if the employment stock of a given province expands or shrinks by 1%, the cumulative flow of jobs being created and destroyed within the year in that province will be equal to 39%. This highlights the large differences between gross and net flows. It is important to notice that excess turnover at

Figure 3.2. Average job flow rates across provinces

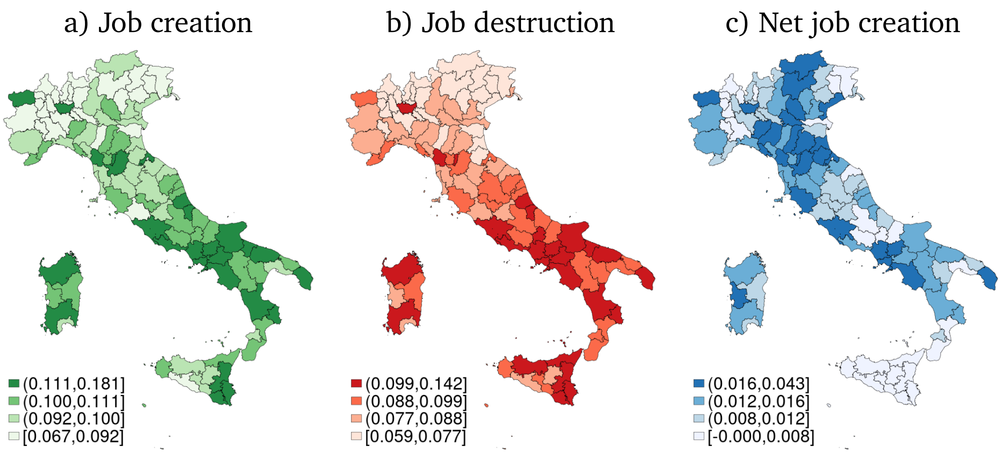

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the geographical distribution of average job flow rates. Job flow rates are the sum across establishments of net activations at the establishment-level, divided by the stock of payroll employment in the current period taken from the ILFS data.

the location level can be only brought about by job reallocation across firms. These features of our data echo results by Davis et al. (2006), who study flows at the establishment level in the US, finding high excess turnover. Moreover, from Table 3.2 we also notice that the distributions of both gross and net labour flows are systematically more dispersed the finer the geographical level (one can see this by comparing vertically the standard deviations of flows, for a given time-frequency). The interpretation of this fact is that relatively larger shocks, i.e. involving a relatively higher share of the local employment stock, are more likely to happen in areas identified at a more disaggregated level.

To get further insights on the relationship between gross and net job flows, in Figure 3.3 we plot average gross rates against the employment growth rate at the province level.[20] The shape of the gross-to-net flows relationship is very relevant for our exercise because it reveals potential asymmetries across different flows. The graph shows that job creation and destruction change in different directions vis-à-vis employment growth changes. Moreover, the average gross rates lie far above the 45-degree line, reflecting the high excess turnover.[21] In other words, provinces where employment is expanding (shrinking) still experience, on average, a substantial amount of job destruction (creation). Overall, these pieces of evidence confirm the previous results on aggregate flows. From Figure 3.3, we also notice the presence of pronounced non-linearities in gross flows, namely that the job creation (destruction) rate is roughly constant for provinces where employment is shrinking (expanding). These non-linearities are even more apparent when studying flows at the monthly frequency

[20]Qualitative results are essentially unchanged at the region or municipality level.

[21]The 45-degree line represents the minimum necessary level of job creation (job destruction) for provinces where employment is expanding (shrinking).

Table 3.2. Summary statistics, job flow rates

Location	Frequency	Job flows			
		JCR	JDR	|NJCR|	$\frac{\text{JTR}}{\lvert\text{NJCR}\rvert}$
Municipality	Monthly	0.029	0.028	0.031	5.2
		(0.063)	(0.073)	(0.089)	(15.0)
	Quarterly	0.062	0.059	0.061	7.1
		(0.107)	(0.108)	(0.136)	(28.9)
	Yearly	0.079	0.069	0.033	13.9
		(0.065)	(0.056)	(0.047)	(47.0)
Province	Monthly	0.032	0.031	0.019	17.0
		(0.020)	(0.029)	(0.027)	(151.2)
	Quarterly	0.069	0.065	0.041	16.1
		(0.039)	(0.045)	(0.044)	(83.7)
	Yearly	0.102	0.089	0.018	38.7
		(0.022)	(0.020)	(0.013)	(138.8)
Region	Monthly	0.033	0.031	0.018	15.8
		(0.017)	(0.025)	(0.024)	(103.0)
	Quarterly	0.070	0.066	0.035	21.7
		(0.029)	(0.036)	(0.032)	(162.3)
	Yearly	0.105	0.091	0.018	40.2
		(0.015)	(0.015)	(0.012)	(157.7)

Source: SISCO and ILFS data 2010-2018. *Note*: The table shows summary statistics of job flow rates for different combinations of geographical (municipality, province, region) and time (monthly, quarterly, yearly) aggregation levels. Job flow rates are the sum across establishments of net activations at the establishment-level, divided by the stock of payroll employment in the current period taken from the ILFS data.

(panel (b) of Figure 3.3), thanks to the larger support of the distribution of net flow rates (see Figure C.3 for the whole scatter). For the purpose of our main exercise, evidence of pronounced non-linearities calls for separate analyses of the effects of specific job flows. This is because a given change in net rates cannot be unambiguously traced back to a given change in gross flows.[22]

The above exercises deal with both time and space variation. A different, though related,

[22]More in general, the relationship between gross and net flows also offers useful insights on the nature of aggregate fluctuations. Indeed, aggregate rates are the result of the combination between these average gross rates and the underlying employment growth rate distribution. In particular, strong non-linearities in the gross flow rates imply that even small changes in the underlying employment growth rate distribution may bring about large movements in aggregate rates. Moreover, when we investigate the extent to which the employment growth rate distribution is subject to swings over time in Figure 3.4, we find that these movements are actually very large. Therefore, this implies that shifts in the employment growth rate distribution over time represent a primary source of fluctuations in labour market flow rates, in line with Davis et al. (2006).

Figure 3.3. Average gross vs. net job flow rates

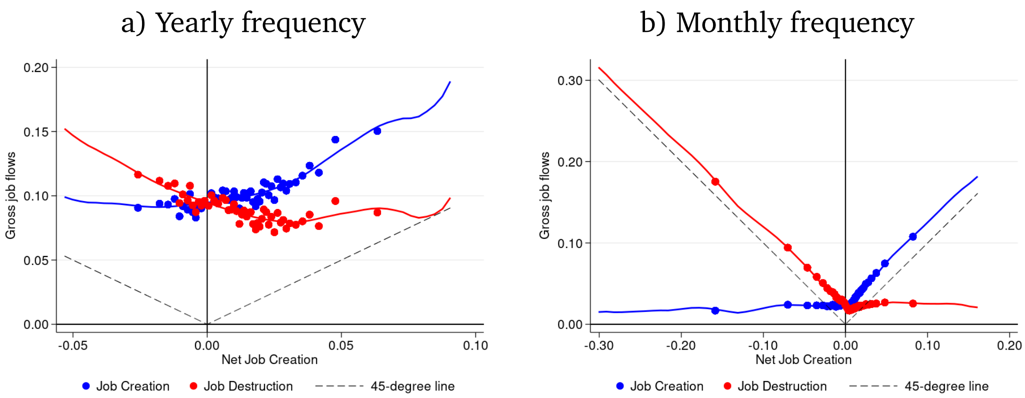

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows average province-level gross job flow rates against the corresponding net flows, at yearly (panel a) and monthly (panel b) frequencies. Solid lines are the prediction of second-degree local polynomial regressions. Scatter points represent averages of two percentiles of the underlying distribution. Dashed lines represent the 45-degree lines. Gross and net flows are divided by the stock of payroll employment in the current period taken from the ILFS data.

Figure 3.4. The employment growth rate distribution over time

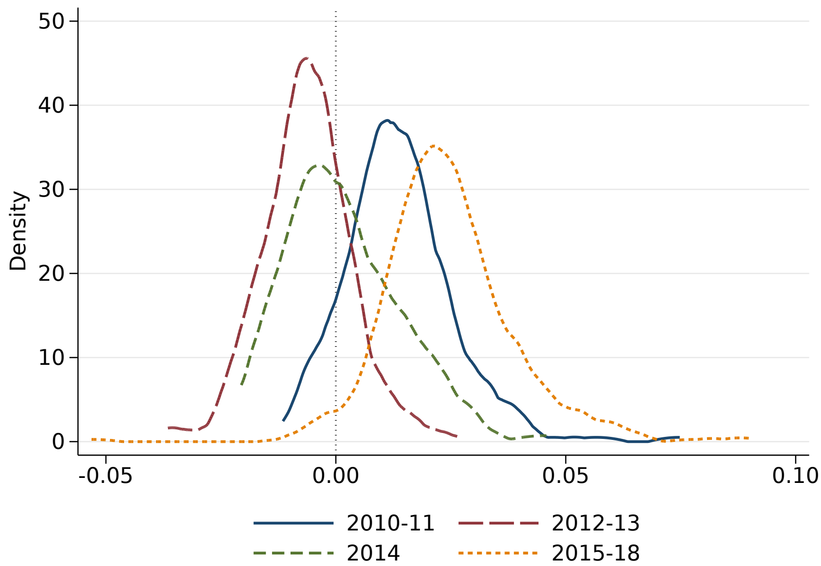

Source: SISCO data 2010-2018. *Note*: The figure shows the employment growth rate distribution at the province-year aggregation level, for different years.

question is to ask which gross flows drive the variation over time in net flows at the local level. In order to shed light on this, we employ a simple statistical decomposition, as proposed by Monras (2018). By regressing gross flow rates on the corresponding net flow rates, one can measure the extent to which each flow contributes to aggregate fluctuations in the net rates. For instance, to decompose employment growth dynamics, we run the following regressions:

$$JCR_{r,t} = \beta_1 NJCR_{r,t} + \alpha_r + \gamma_t + \epsilon_{r,t} \tag{3.1}$$

$$JDR_{r,t} = \beta_2 NJCR_{r,t} + \alpha_r + \gamma_t + \epsilon_{r,t} \tag{3.2}$$

Given the definition of the variables ($NJCR = JCR - JDR$), in this setup the condition $\beta_1 + \beta_2 = 1$ must hold.[23] It is important to notice that this exercise decomposes within-location variation in employment growth rates. Hence, we can directly interpret the estimated coefficients as the share of variance of net rates accounted for by the specific gross flow. Table C.1 shows the results of this decomposition for all combinations of time and geographical aggregations. At the yearly frequency, the estimated value of β_1 ranges between 0.5 and 0.7, implying that job creation accounts for 50 to 70% of overall fluctuations. Overall, the share of variance accounted for job creation is between 35 and 71%, if one extends the analysis to the other possible combinations of sampling interval and geographical aggregation level.[24]

Summing up, two important lessons can be drawn from this descriptive evidence. First, gross flows starkly differ from net flows, both in terms of levels and dynamics. Second, it is crucial to account for both job creation and job destruction, as they are both important drivers of employment growth. Consistently with this, our subsequent analysis in Section 3.3 will focus separately on specific gross job flows, allowing for potentially asymmetric effects.

3.2.3 Internal Migration

To measure internal migration flows, we leverage the SISCO individual-level data exploiting the information of the workplace location, which we use as a proxy for residence.[25] Compared to using the more standard residence measures, this has a number of impli-

[23] To see this, notice that $\hat{\beta}_1 = \frac{cov(JCR,NJCR)}{var(NJCR)}$, and $\hat{\beta}_2 = \frac{cov(JDR,NJCR)}{var(NJCR)}$.

[24] Differences across time aggregation levels can be understood referring to Figure 3.3. To a first approximation, gross flows are important determinants of net flows if and only if they systematically vary with the latter, i.e. for instance job creation is increasing in employment growth rates (not necessarily true by construction). Moreover, for a given shape of the gross-to-net-flows relationship, the distribution of employment growth rates also matters for the decomposition results. This is because, depending on the actual realizations of shocks, more or less weight is given to parts of the support where job creation (destruction) is more (less) correlated with the employment growth rate. In our data, such distribution at the yearly frequency has much more mass in the positive region than the one at the monthly frequency (see Figure C.4), explaining the different decomposition results.

[25] This choice is mainly due to the fact that in our data the information on the individuals' residence is not updated over time. Identifying the residence through the workplace is a strategy that was adopted also by Bartolucci et al. (2018).

cations, which we discuss at length in Appendix C.3. Overall, we claim that SISCO data are actually better suited to measure internal migration with respect to traditional data sources based on administrative data (e.g., changes of residence) or surveys (e.g., ILFS). Indeed, SISCO data do not suffer from (i) misreporting, known to be potentially quite large in other sources (Rubolino, 2020), (ii) under-counting (residence-based data are likely not to record short-distance transfers), which we find to be a very severe problem in ILFS data, or (iii) attrition, that has been documented for labour force surveys (Martí and Ródenas, 2007). The main drawback of SISCO data is that they do not record movements of non-employed people within their unemployment or inactivity spells (i.e. until they find a new job in a new location). In practice, we find that such a limitation is likely to be very small, possibly because most of the internal migration also entails a job change. Indeed, in Appendix C.3 we show that our measures of internal migration are very highly correlated with administrative-based residence changes from the ISTAT, which is reassuring of our proxy being valid.

To construct the migration measures we first need to transform the contract-level dataset into a worker-level panel. This involves assigning to each worker-period combination the prevalent job, for all those cases with multiple contracts within a given time interval. Details on how we pick the prevalent job for each period can be found in Appendix C.2.1. In the baseline version of our dataset, we focus only on *direct* transitions, that is we do not consider transitions that involve non-employment spells. We do this to avoid having to impute the exact timing of the transitions, as well as to avoid the possibility of spurious transitions, given that we do not observe residence changes of non-employed individuals.[26]

With the worker-level panel dataset, it is straightforward to compute aggregate migration flows through individual transitions. We define a dummy $m^j_{s\to r,t}$ that takes the value of 1 if worker j has made a transition between location s (i.e., any region different from r) to location r at time t.

Therefore, location-specific inflows $IM_{r,t}$, outflows $OM_{r,t}$ and net inflows $NIM_{r,t}$ are simply defined as:

$$IM_{r,t} = \sum_s \sum_j m^j_{s\to r,t}, \qquad OM_{r,t} = \sum_s \sum_j m^j_{r\to s,t}, \qquad NIM_{r,t} = IM_{r,t} - OM_{r,t}.$$

[26]Notice that, at the yearly frequency, this is only excluding workers who completely leave employment for at least one full calendar year. However, we also compile another version of our dataset in which we keep this type of transitions, assuming that the worker's location corresponds to her last workplace location until the new job is found. That is, we assign the location switch at the end of the non-employment spell. In this alternative dataset, we include all cases in which the non-employment spell covers exactly one yearly observation (i.e. we retrieve all histories of the type E-N-E), which represent about 51% of all cases with non-employment yearly observations. Notice that this implies that the non-employment spell can practically last for up to almost two full calendar years. For the purposes of our main empirical exercise, in a robustness check we show that including also indirect transition does not affect our results (Table C.2).

To better interpret magnitudes, we divide these flows by the previous period's payroll employment stocks (derived from ILFS data), thus obtaining inflow, outflow and net inflow rates: $IMR_{r,t}$, $OMR_{r,t}$ and $NIMR_{r,t}$. Last, we define the excess migration turnover rate $EMRT$ as the ratio between the migration turnover rate MTR – that is the sum of the inflow and outflow rates – and the absolute value of the net inflow rate:

$$EMTR_{r,t} = \frac{IMR_{r,t} + OMR_{r,t}}{|NIMR|_{r,t}}.$$

In the context of migration, excess turnover is a particularly relevant statistic, since it embeds information on the nature of shocks that possibly cause migration. For instance, if migration were driven almost only by aggregate (i.e. equal for all workers) shocks, excess turnover should be minimal. Instead, if idiosyncratic shocks (i.e. heterogeneous across workers) were prevalent for migration choices, then excess turnover would be large. In the first case, most of the migration choices would be nearly identical across individuals, whereas in the second case migration decisions would be heterogeneous, resulting in largely offsetting flows.

Aggregate flows. Before delving into location-specific flows, we first study aggregate migration flows, which we obtain by simply summing up all events and dividing them by the corresponding aggregate stock. Figure 3.5 shows trends of internal migration in Italy for the whole combination of geographical and time aggregation levels. In our sample period, gross migration rates at the yearly frequency were about 7.5% at the municipality level, 3.4% at the province level and 2.0% at the region level. Nonetheless, the overall dynamics are quite robust across aggregation levels. We observe a substantial drop in internal mobility during the recession (2011-2013) and subsequent recovery, especially apparent in the last two years. These dynamics are more clearly detected in high-frequency data (panel (b) of Figure 3.5), whereas yearly data tend to smooth out these changes over time.[27] Importantly, virtually identical patterns are found also using other traditional data sources on internal mobility (Figure C.16, panel a), which is reassuring that the cyclicality is not simply a by-product of our definition of residence linked to the workplace.

Location-specific flows. Figure 3.6 shows the distribution of relocation flows across space, plotting the average migration rates at the province level. We see that gross migration flows are larger in the Southern provinces, while net flows are higher in the North, which is a net receiver of internal migration flows in our sample period. Once again, very similar patterns are detected using administrative data on residence changes (Figure C.5), namely a disconnection between gross and net flows. These results are suggestive of an important link between internal migration and labour market flows, given that similar geographical patterns were uncovered for job flows.

[27]In parallel ongoing work, we are investigating more in-depth the implications of the data frequency and of the time horizon for the measurement of labour market flows based on individual transitions.

Figure 3.5. Time series of internal migration

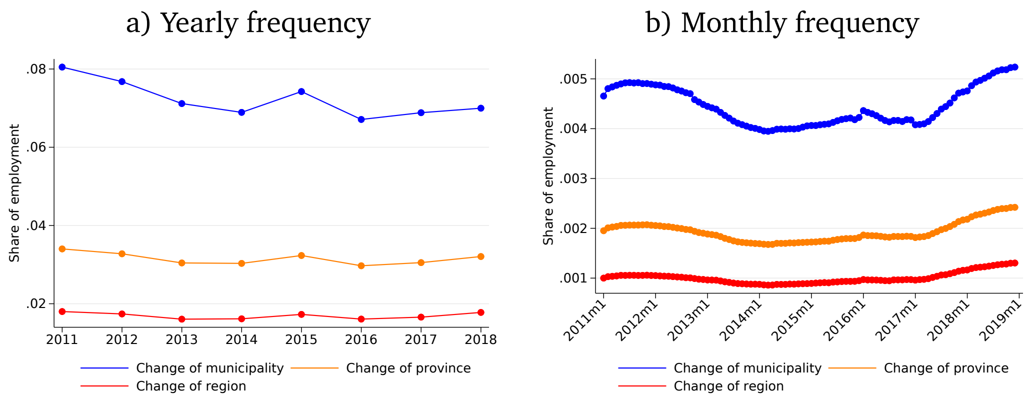

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the time trends of internal migration rates for different geographical (municipality, province, region) at yearly (panel a) and monthly (panel b) frequencies. Migration rates are computed by dividing the number of observed transitions in the SISCO microdata by the corresponding stocks of payroll employment in the previous period taken from the ILFS data.

Figure 3.6. Average internal migration rates across provinces

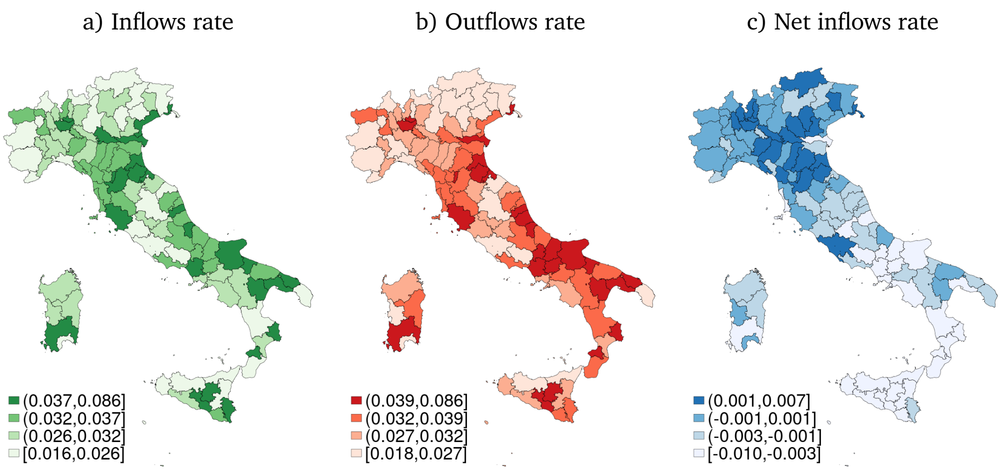

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the geographical distribution of average internal migration rates. Migration rates are computed by dividing the number of observed transitions in the SISCO microdata by the corresponding stocks of payroll employment in the previous period taken from the ILFS data.

When comparing gross and net flows, we again find a very large degree of excess turnover (see Table 3.3 and Figure 3.7). For instance, at the province level, the average excess turnover rate ranges between 31.6 and 107.6, depending on the time interval. These large numbers imply that internal migration flows in Italy systematically go in opposite directions, reflecting a large degree of heterogeneity in workers' choices. Importantly for our analysis, this calls for a separate assessment of gross vs. net flows.

Table 3.3. Summary statistics, Internal migration

Location	Frequency	Internal migration flows			
		IMR	OMR	|NIMR|	$\frac{MTR}{\|NIMR\|}$
Municipality	Monthly	0.004	0.004	0.004	5.2
		(0.008)	(0.010)	(0.010)	(11.9)
	Quarterly	0.016	0.017	0.010	9.9
		(0.026)	(0.026)	(0.026)	(29.2)
	Yearly	0.074	0.076	0.022	22.1
		(0.076)	(0.079)	(0.042)	(68.1)
Province	Monthly	0.002	0.002	0.001	31.6
		(0.001)	(0.002)	(0.001)	(73.9)
	Quarterly	0.008	0.008	0.002	50.5
		(0.004)	(0.004)	(0.003)	(163.1)
	Yearly	0.033	0.034	0.003	107.6
		(0.011)	(0.011)	(0.003)	(481.4)
Region	Monthly	0.001	0.001	0.000	40.3
		(0.001)	(0.001)	(0.001)	(115.1)
	Quarterly	0.005	0.005	0.002	54.9
		(0.004)	(0.004)	(0.003)	(256.6)
	Yearly	0.020	0.021	0.002	412.8
		(0.008)	(0.008)	(0.002)	(4324.8)

Source: SISCO and ILFS data 2010-2018. *Note*: The table shows summary statistics of internal migration rates for different combinations of geographical (municipality, province, region) and time (monthly, quarterly, yearly) aggregation levels. Migration rates are computed by dividing the number of observed transitions in the SISCO microdata by the corresponding stocks of payroll employment in the previous period taken from the ILFS data.

In order to shed light on the relationship between net and gross migration flows over time, we now perform the same decomposition exercise carried out in the previous section for labour market flows. In particular, we run the following regressions:

$$IMR_{r,t} = \beta_1 NIMR_{r,t} + \alpha_r + \gamma_t + \epsilon_{r,t}, \tag{3.3}$$

$$OMR_{r,t} = \beta_2 NIMR_{r,t} + \alpha_r + \gamma_t + \epsilon_{r,t}, \tag{3.4}$$

where the estimated β's represent the share of variance of the net inflow rate accounted for by fluctuations in the specific gross flow rate. Table C.3 reports the results of this decomposition. We can see that the split between inflow and outflow rates is overall very balanced: both at the quarterly and yearly frequency, inflow rates account for between 43

Figure 3.7. Gross vs. net internal migration flows

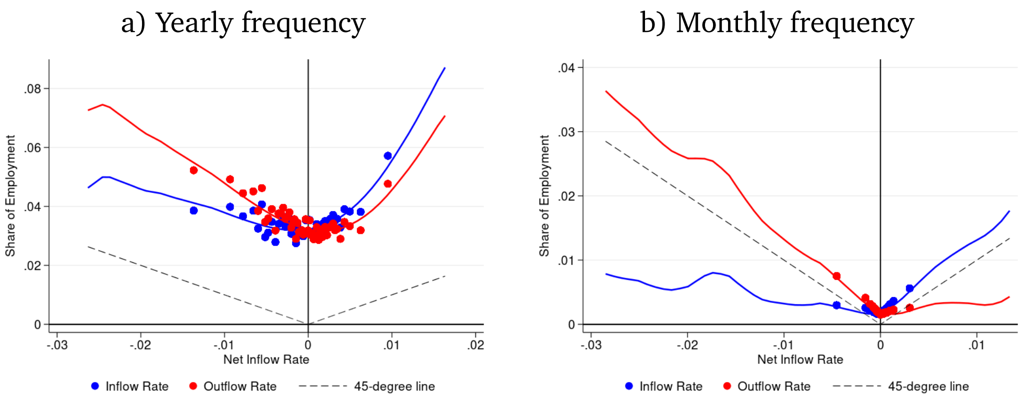

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows average province-level gross migration flow rates against the corresponding net flows, at yearly (panel a) and monthly (panel b) frequencies. Solid lines are the prediction of second-degree local polynomial regressions. Scatter points represent averages of two percentiles of the underlying distribution. Dashed lines represent the 45-degree lines. Gross and net flows are divided by the stock of payroll employment in the previous period taken from the ILFS data.

and 52% of the overall variance.

Last, we investigate the geographical reach of the observed internal migration transitions. Figure C.6 shows the density of distance (measured in kilometres) involved in our transitions.[28] It is immediately apparent that the lower the geographical aggregation level, the lower the distance involved. This effect is mainly mechanical, as progressively broader definitions of the location tend to exclude shorter moves. At any rate, differences are very large: for instance, the median move at the municipality level involves a distance of just 24.9 km, whereas it corresponds to 119.2 km and 300.7 km at the province and region level, respectively (Table C.4). Given that the focus of our study is on geographical relocation, as opposed to commuting, we decide to adopt the province level as the baseline for our analysis. Otherwise, there would be important concerns that our measures of internal migration actually capture changes in commuting patterns.[29]

Overall, in a symmetric fashion as for labour market flows, this Section has shown that it is key to distinguish gross and net flows, and that both in-migration and out-migration rates are important determinants of net variations. Consistently with this, our empirical

[28] See Appendix C.2.3 for details on the distance statistics.

[29] With SISCO data, we can study commuting patterns for a cross-section of workers (as already mentioned, this is because the residence variable is only available for the last job). Conditional on commuting, we find that the median distance between residence and workplace municipality is 19.4 km. This implies that a large part of the distribution of workplace municipality changes may indeed capture commuting. Instead, in our data, the residence and workplace province do coincide for 84% of the workers (as opposed to only 54% for the municipality). This lends support to our choice of using the workplace province as a reliable proxy for the residence province.

analysis in the next Section will separately study the behaviour of in- and out-migration rates.

3.3 Causal evidence

Following up on the descriptive evidence of Section 3.2, we now analyze whether labour market flows correlate with internal migration flows and whether the former drives the latter. We are particularly interested in analyzing the contemporaneous effects of both job creation and job destruction, as we saw that the two gross flows are spatially correlated. We first show simple associations by means of OLS regressions, and then provide causal evidence by instrumenting job creation and job destruction with sudden and plausibly exogenous mass hire and mass layoff events.

3.3.1 OLS Regressions

Using the same province-level yearly data presented in Section 3.2, we relate gross and net internal migration flows with gross job flows.[30] The empirical model we base our analysis on is the following:

$$OMR_{r,t} = \beta_1 JCR_{r,t} + \beta_2 JDR_{r,t} + \alpha_r + \gamma_t + \varepsilon_{r,t} \quad (3.5)$$

$$IMR_{r,t} = \beta_1 JCR_{r,t} + \beta_2 JDR_{r,t} + \alpha_r + \gamma_t + \varepsilon_{r,t} \quad (3.6)$$

$$NIMR_{r,t} = \beta_1 JCR_{r,t} + \beta_2 JDR_{r,t} + \alpha_r + \gamma_t + \varepsilon_{r,t} \quad (3.7)$$

where $OMR_{r,t}$, $IMR_{r,t}$ and $NIMR_{r,t}$ are, respectively, the gross out-, gross in- and net in-migration flows as a percentage of local employment in the previous year, and $JCR_{r,t}$ and $JDR_{r,t}$ are the gross job flows as a percentage of local employment in the current year.[31] We further account for time and location fixed effects; finally, the regressions are weighted using the stock of payroll employment in the current period taken from the ILFS data.[32]

The results, presented in Table 3.4, indicate that labour market dynamism is highly

[30]We base our analysis on yearly data, as our instrumental variables, mass hire and layoff events, can be credibly defined only at the yearly level (see 3.3.2 for more details). In terms of the geographical unit of interest, we replicate our results at the municipality level: the results, reported in Table C.5 are similar to those of our main specification, though the magnitude of the JCR effect is about two-thirds of that of the baseline.

[31]Unlike migration flows, which are measured with respect to the previous period's stock of employment, job flows are the cumulative sum within a period. Therefore, it is conceptually more correct to use the current period's stocks for the latter flows (consistent with Davis et al. (1998, 2006)). However, we have also verified that this choice does not affect our results. More generally, our estimates are also robust to using total population as the relevant stock, instead of employment (not reported, available upon request).

[32]In other specifications (not reported, but available upon request), we further control for labour market characteristics and most notably for the unemployment rate. The results are qualitatively similar to those reported here.

correlated with migration flows.[33] In terms of net migration flows (column (3)), the signs are as expected and the correlations are symmetric: a 1 percentage point increase in the JCR is associated with an increase in net migration of 0.10 percentage points, while a similar increase in the JDR is associated to an analogous drop in net migration. Two main considerations arise. First, the estimated associations are rather small in magnitude suggesting that jobs are mainly filled in by local workers. Second, the JCR is slightly positively correlated with outflows (though not statistically significant), a rather puzzling result. This association is likely due to the fact that JCR and JDR are spatially correlated, as shown in Figures 3.2 and 3.3. This may cause issues in the interpretation of these coefficients. More in general, the OLS results cannot ensure that migration follows, in a causal sense, changes in labour demand. It could well be that the effects are partly due to reverse causality as changes in labour supply due to movements of workers across areas determine the growth or reduction of jobs in a given location. Moreover, as highlighted above, the concurrence of job creation and job destruction flows does not allow to separately identify the drivers of gross migration flows. For all these reasons, we turn to an instrumental variable identification strategy.

Table 3.4. Internal migration flows and labour market dynamism, OLS

	(1)	(2)	(3)
	OMR	IMR	NIMR
JCR	0.021	0.125**	0.104**
	(0.016)	(0.021)	(0.013)
JDR	0.102**	-0.009	-0.111**
	(0.019)	(0.019)	(0.017)
N	856	856	856
R^2	0.954	0.961	0.629

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the OLS estimates of the effect of job creation and job destruction rates on internal migration flow rates (out-migration, in-migration, net in-migration) from SISCO data as described in Section 3.3.1. The specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level. $^+p < .10, ^*p < .05, ^{**}p < .01$.

[33] As explained in Section 3.2, we do not account for the entirety of the migration movements as we do not observe people who do not participate in the labour market.

3.3.2 IV Regressions

Borrowing the idea of leveraging variation in mass layoffs in large firms from Gathmann et al. (2020), we further augment the specification by also including mass hire events. Such a strategy allows us to contemporaneously identify the effects of job creation and job destruction on migration flows.

We isolate large mass layoffs and mass hiring events in the SISCO data by focusing on establishment-level net terminations and hires of more than 250, 500 o 1,000 workers in a given year.[34] The idea is that large firm mass layoff and entry (or enlargement) events are able to affect the labour demand of the entire area due to their direct effect, but most importantly because of spillovers at the local level (e.g., through local value chains or shifts in local goods and services demand). To exclude confounding factors generated by mergers and acquisitions, sales of business units or temporary contracts with employment agencies, we exclude layoff (hire) events for which we observe most of the same workers being hired in (laid off from) another firm in the same municipality during the next (previous) year.[35] To prove that these events are both relevant – i.e., they matter for the local labour demand – and plausibly exogenous – to satisfy the exclusion restriction the events must not follow pre-existing trends – we use an event study approach.

First, let us note that the Italian labour market is characterized by various policies and institutions that foster employment protection and increase the costs of mass layoffs. Various schemes, such as short-time work (*Cassa integrazione guadagni*), allow to preserve employment relations during downturns and tend to be used to protect labour even when a firm crisis is permanent rather than transitory. Such policies reduce the ex-ante likelihood of observing mass layoffs. Nonetheless, we find that both mass layoff and hire events are rather common and spread across the country. Table 3.5 reports the main characteristics of the events we analyze. We observe mass hires and layoffs involving more than 250, 500 and 1,000 workers occurred in 69 provinces (for 250 events) and about 30 provinces (for 500 and 1,000 events) over the sample period and mainly concentrated in the private services sector: the average size of the events is about 473, 969 and 1,960 workers, depending on the selected threshold. The mass hire events involve between 0.07 per cent and 4 per cent of the province-year employment, while the mass layoff ones can affect up to almost 7 per cent of the local workforce. Figure C.7, Figure C.8, and Figure C.9 further show the geographical annual distribution of mass events in Italy, for each selected threshold. In Table C.6 we report the summary statistics of mass events by geographical macro area

[34] As already mentioned, we define an establishment as the combination of a firm and a municipality of work, i.e., establishments of the same firm in a given municipality are pooled together.

[35] More specifically, we exclude events for which we observe 70 (for 250-unit events), 50 (for 500) and 30% (for 1,000) of the workforce being employed in (or laid off from) another firm in the same municipality the next (previous) year. The results are robust to the inclusion of these events (not reported, but available upon request).

(North-East, North-West, Center, South and Islands): we detect a concentration of the events (40-45% on average) in the north-west of the country, in particular of large (more than 1,000 employees involved) mass hires. However, all the areas are hit by shocks comparable in absolute and relative size.[36]

We estimate establishment-level event studies according to the following specification:

$$\begin{aligned}\Delta Empl_{i,t} &= \sum_{k=1}^{3} \beta_k 1[Massevent_i = t-k] + \sum_{k=1}^{4} \beta_k 1[Massevent_i = t+k] + \\ &+ \alpha_i + \gamma_t + \varepsilon_{i,t}\end{aligned} \tag{3.8}$$

where α_i and γ_t are establishment and year fixed effects and $\Delta Empl_{i,t}$ are the net activations in each establishment i and year t. The standard errors are clustered at the establishment level. We define a $Massevent_i$ as either an annual increase in employment of more than 250, 500 or 1,000 units, or an annual decrease in employment of more than 250, 500 or 1,000 units. The treated units are all firms in the data where a mass event occurs (i.e., we do not impose any minimum or maximum firm size): in the main specification, we do not include untreated units, allowing the not-yet treated and already-treated firms to act as control units.[37]

The regression results of the event studies for the three thresholds (250, 500, and 1,000) are reported graphically in Figure 3.8, where the left-hand panels show the results for mass hire events, and the right-hand panels for mass layoff events. Both types of events show a similar pattern. First, and most notably, the pre-trend is flat indicating that both mass hires and mass layoffs are not consequences of pre-existing trends and that firms where these events do not occur are observationally equivalent to the affected firms before the events. This test reassures us about the plausible exogeneity of the instruments used for the identification strategy and confirms that such events occur despite the existence of labour institutions aimed at reducing, or spreading over time, the extent of workers' layoffs. Second, mass hires imply, on average, an increase in employment of just above the relative thresholds (250, 500 or 1,000 units) in the event year, which only partially offsets future hires. In the case of mass layoffs, the average decline in the establishment workforce is larger than the relative threshold in each specification and it keeps declining though at lower levels in the subsequent years.

Both types of events are indeed salient and involve a non-negligible share of local employment (Table 3.5). To formally test the relevance of the events for local job creation

[36]The south of the country experiences, on average, bigger shocks than the other areas in terms of relative employment, due to an outlier mass layoff of more than 8,000 workers dismissed following the takeover of a large Italian steel producer in the Taranto province in 2018.

[37]Following a recent growing literature that highlights the potential pitfalls of having only treated units in difference-in-difference regressions (see de Chaisemartin and D'Haultfoeuille (2022) for a recent survey), we also included untreated units in the control group: the results, available upon request, are robust to such specification of the event studies.

Table 3.5. Event studies summary statistics

	Mass hire events	Mass layoff events
Panel (a). Events with at least 250 workers involved		
Number of events	548	520
Number of provinces hit by events	69	69
Min (Avg) Max size of events	251 (482) 4,861	251 (463) 8,414
% of province employment	0.07 (0.39) 4.15	0.07 (0.37) 6.87
Number of events by industry		
Manufacturing	22	45
Construction	6	6
Private services	439	327
Public services	66	95
Other sector/not specified	15	47
Panel (b). Events with at least 500 workers involved		
Number of events	130	110
Number of provinces hit by events	30	31
Min (Avg) Max size of events	501 (965) 4,861	501 (973) 8,414
% of province employment	0.04 (0.52) 3.78	0.04 (0.43) 6.87
Number of events by industry		
Manufacturing	5	8
Construction	0	0
Private services	115	74
Public services	7	18
Other sector/not specified	3	10
Panel (c). Events with at least 1,000 workers involved		
Number of events	31	26
Number of provinces hit by events	11	9
Min (Avg) Max size of events	1,023 (1,911) 4,861	1,013 (2,009) 8,414
% of province employment	0.10 (0.72) 3.78	0.09 (0.59) 6.87
Number of events by industry		
Manufacturing	2	3
Construction	0	0
Private services	28	17
Public services	1	2
Other sector/not specified	0	4

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the summary statistics for mass layoff and hiring events with thresholds set at 250 (panel a), 500 (panel b) and 1,000 (panel c) units calculated from SISCO data as described in Section 3.3.2. Public services include privately-provided education and health services.

Figure 3.8. Establishment-level mass hires and layoffs in Italy, event studies

Events involving more than 250 workers

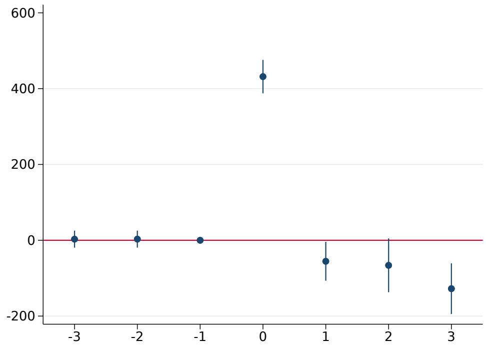

(a) Mass hire events

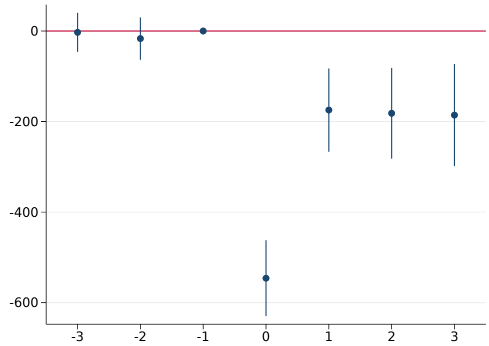

(b) Mass layoff events

Events involving more than 500 workers

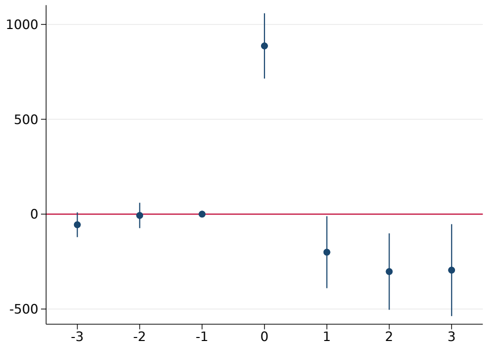

(c) Mass hire events

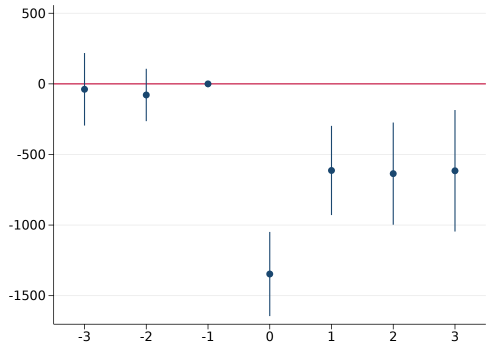

(d) Mass layoff events

Events involving more than $1,000$ workers

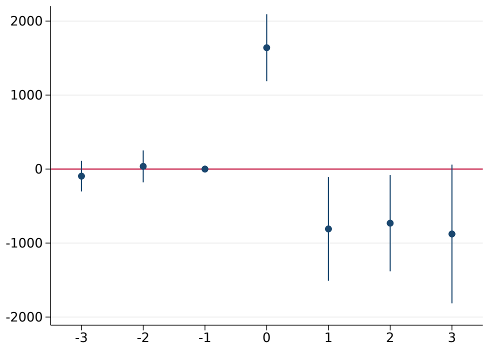

(e) Mass hire events

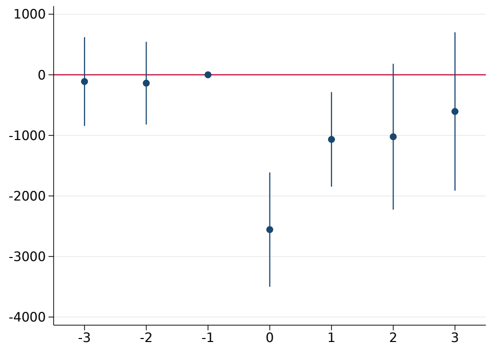

(f) Mass layoff events

Source: SISCO data 2010-2018. *Note*: The figure shows the results of the event studies for establishment-level mass hires and layoffs (equation (3.8)), at different thresholds. We isolate large mass layoff and mass hire events in the SISCO data by focusing on establishment-level terminations and hires above the specified threshold (250, 500, 1,000) in a given year. The 95 per cent confidence intervals (bars) are clustered at the establishment level.

and job destruction, we run first-stage regressions. Based on the event-study analysis, we define our exogenous shifter as the number of employees involved in the mass hire and layoff events over the province employment, i.e. we consider the intensive margin of the events.[38] Therefore, the IV variables are defined as:

$$MH^{IV}_{r,t} = \sum_j \frac{empl_j(r) \cdot 1[Masshire_j(r) = t]}{empl_{r,t}} \tag{3.9}$$

$$ML^{IV}_{r,t} = \sum_j \frac{empl_j(r) \cdot 1[Masslayoff_j(r) = t]}{empl_{r,t}} \tag{3.10}$$

where establishments located in province r are indexed by j.

The estimation results in Table 3.6 are extremely reassuring: the mass hire events are positively correlated with local job creation, while the mass layoff events are positively and strongly correlated with job destruction (for all three mass events' sizes). Importantly, there are basically no cross-effects indicating that mass hire and layoff events are able to separately identify JC and JD, respectively. To make sure that the instrument is exogenous also at the provincial level, we further test the absence of pre-trends in the correlation between the mass hires/layoffs as a percentage of local employment and the JC and JD rates. The local projection estimates (Jordà, 2005), reported graphically for different thresholds in Figure C.10, Figure C.11, and Figure C.12, confirm that both instruments have a one-time impact on the endogenous variables and are not affected by anticipation effects or confounding trends.[39]

While mass layoffs have been widely studied in the literature, we are the first to our knowledge to use mass hires as an instrument for job creation. There are mainly two categories of mass hire events in our data: i) the opening or expansion of establishments, with the consequent creation of mainly permanent jobs, concentrated in the manufacturing and transport sectors; ii) the hiring en masse of temporary staff for specific projects and short-term needs of expansion of the production capacity. This second type of event mainly concerns business services (marketing, business consulting, but also cleaning and software development). Representative examples in our data are increases in the workforce in 2015 at a plant in the province of Potenza (+1837) of a well-known car company and

[38]The results are qualitatively similar if we include a dummy variable for the province-year of the events instead of the intensive margin. The 2SLS are reported in Table C.7. In an additional robustness check, we test whether our results hold if using the change of residence administrative data. Table C.8 indicates that while some of the main effect hold (notably that of JDR on out-migration), most of the results are wiped out. As we claim at length in Appendix C.3, this is due to problems in the residence data in tracking mobility over time, but not in terms of flow magnitudes or bilateral geographical correlations between origin and destination provinces in SISCO vis-à-vis administrative residence data.

[39]Similar to distributed-lag regressions, the local projections allow to estimate the impact of a treatment on an outcome over time (for more details see Jordà (2005) and Jordà et al. (2020)). The specification used here is the same as that of the first stage reported in Table 3.6, though the outcome variable is lagged or forwarded accordingly (i.e., $y_{r,t+h} = \beta \Delta Empl^{MassEv}_{i(r),t} + \alpha_r + \gamma_t + \varepsilon_{r,t+h}$ for $h = -3, .., 0, 1, .., 4$, where $y_{r,t+h}$ represents either JCR or JDR). The results are robust to the inclusion of lags in the explanatory variable.

Table 3.6. Internal migration flows and labour dynamism, 'mass events' IV first-stage

Threshold IV	250 workers		500 workers		1,000 workers	
	(1)	(2)	(3)	(4)	(5)	(6)
	JCR	JDR	JCR	JDR	JCR	JDR
Employment in mass hire	0.707**	0.106	0.800**	0.084	0.736**	-0.011
	(0.211)	(0.064)	(0.199)	(0.069)	(0.271)	(0.100)
Employment in mass layoff	0.155*	1.182**	0.063	1.174**	0.071	1.183**
	(0.077)	(0.088)	(0.063)	(0.090)	(0.052)	(0.119)
N	856	856	856	856	856	856

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the first stage estimates of mass layoff and hiring events on job creation and job destruction rates from SISCO data as described in Section 3.3.2. We selected events that involve more than 250 (specifications (1)-(2)), more than 500 (specifications (3)-(4)), or more than 1,000 (specifications (5)-(6)) net activations or terminations. Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level. $^{+}p < .10, ^{*}p < .05, ^{**}p < .01$.

at a regional public transport company in the province of Cagliari (+520). On the other hand, examples of massive fixed-term hires are the recruitment by marketing and business support companies linked to the Expo Milano 2015 event, and the temporary hiring of more than 1,000 people in Turin in 2017 for a project launched by a well-known sporting club. Interestingly, these events show a certain symmetry with respect to mass layoffs: where the closure of a plant or its downsizing generates a mass layoff, the opening or expansion generates a mass hire. Similarly, mass hires for specific business projects translate (one or more years later, depending on the duration of the need for more capacity) into mass layoffs.

Having shown the plausibility of mass hire and layoff events as independent determinants of local job creation and destruction, we now turn to the 2SLS estimation. Table 3.7 partly confirms the OLS results, but provides further interesting evidence. First, the magnitude of the job creation effect on net migration is about three times larger than the OLS estimates, possibly reflecting a local effect, while that of job destruction is about half as large than before, indicative of an upward bias in the OLS. More specifically, a 1 percentage point increase in the job creation rate leads to a .3 percentage points increase in the in-migration rate; the same change in the job destruction rate causes an increase in the out-migration rate of just .05 percentage points. This shows that JC has greater power in generating migration flows than JD does. Importantly, we separately identify the effects of each labour market flow on gross migration flows. The cross-effects (i.e., a change in JCR on outflows and of JDR on inflows) have always the expected signs, though they are not statistically different from zero. Moreover, the analysis of gross flows indicates that JD per se does not cause large out-migration flows, though the effect is statistically significant in two specifications out of three, something rarely observed in the existing literature on the US and Germany (Notowidigdo, 2020; Monras, 2018). Such results highlight the importance of accounting for both margins of labour market dynamics when estimating adjustments to demand shocks.

Table 3.7. Internal migration and labour dynamism, 'mass events' IV 2SLS

	250 workers			500 workers			1,000 workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	OMR	IMR	NIMR	OMR	IMR	NIMR	OMR	IMR	NIMR
JCR	-0.032	0.355*	0.387**	-0.031	0.319**	0.350**	0.004	0.367^{+}	0.363**
	(0.048)	(0.145)	(0.126)	(0.041)	(0.118)	(0.106)	(0.076)	(0.196)	(0.137)
JDR	0.053^{+}	-0.032	-0.085**	0.034	-0.008	-0.042^{+}	0.049^{+}	-0.002	-0.050*
	(0.031)	(0.035)	(0.033)	(0.022)	(0.027)	(0.023)	(0.026)	(0.037)	(0.021)
N	856	856	856	856	856	856	856	856	856

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the 2SLS estimates of job creation and job destruction rates, instrumented by mass hiring and layoff events, on migration flow rates (out-migration, in-migration, net in-migration) from SISCO data as described in Section 3.3.2. We selected events that involve more than 250 (specifications (1)-(3)), more than 500 (specifications (4)-(6)), or more than 1,000 (specifications (7)-(9)) net activations or terminations. Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level. $^{+}p < .10,^{*} p < .05,^{**} p < .01$.

3.3.3 Inspecting the asymmetry of the effect

One of the key results presented above was that the effect of JC (i.e. positive shocks) on in-migration is much larger than the one of JD (i.e. negative shocks) on out-migration. We believe that this is a very important finding, with clear implications on how different shocks are absorbed through migration flows and on how we should think of the reaction of migration choices to labour market shocks. In this Section, we try to shed some light on the difference in the magnitude of the responses. First of all, it is useful to note that, by construction, in-migration flows reflect decisions of households of *all other* locations, implying a much larger originating stock than out-migration, which is instead bounded by the local population. From this observation, we write down a simple decomposition equation of the in-migration rate that allows us to distinguish the extent of the population at risk of in-migrating from the actual decision of migrating towards a given region. That is:

$$IMR_{r,t} = \frac{\sum\limits_{s\neq r}\sum\limits_{j} m^{j}_{s\to r,t}}{E_{r,t-1}} = \underbrace{\frac{\sum\limits_{s\neq r}\sum\limits_{s'}\sum\limits_{j} m^{j}_{s\to s',t}}{E_{r,t-1}}}_{Migrants\,from\,s\neq r} \underbrace{\frac{\sum\limits_{s\neq r}\sum\limits_{j} m^{j}_{s\to r,t}}{\sum\limits_{s\neq r}\sum\limits_{s'}\sum\limits_{j} m^{j}_{s\to s',t}}}_{Share\,picking\,r} \tag{3.11}$$

$$= \underbrace{\frac{\sum\limits_{s\neq r} E_{s,t-1}}{E_{r,t-1}}}_{\equiv ER} \underbrace{\frac{\sum\limits_{s\neq r}\sum\limits_{s'}\sum\limits_{j} m^{j}_{s\to s',t}}{\sum\limits_{s\neq r} E_{s,t-1}}}_{=OMR_{s\neq r}} \underbrace{\frac{\sum\limits_{s\neq r}\sum\limits_{j} m^{j}_{s\to r,t}}{\sum\limits_{s\neq r}\sum\limits_{s'}\sum\limits_{j} m^{j}_{s\to s',t}}}_{\equiv RAR}. \tag{3.12}$$

Equation (3.11) makes clear that the in-migration rate of region r is equal to the product of two terms. The first one represents *potential* in-migrants (i.e., workers relocating from all other locations), while the second one is the decision of actually migrating into r, which we can interpret as the probability that region r represents the best choice for relocating households. Loosely speaking, the first term has to do with the binary decision of whether to relocate or not, whereas the second term reflects the ranking of different alternatives (we, therefore, name it *relative attractiveness ratio*, RAR). In fact, the first term can be decomposed further into a component that represents the relative size effect (that we term ER – *employment ratio* – i.e. how large the stock of employment in all other locations is relative to the one of the region r) and a component that corresponds to the overall out-migration rate of a macro-region made of all locations minus r (Equation (3.12)). Note that, by construction, it is not possible to derive any decomposition of this kind for the out-migration rate.

From (3.12), we derive three counterfactual series of the IMR for each province generated by letting only one component vary over time and fixing the others at their average. Regressing the actual series on these counterfactual ones allows us to gauge the quantitative role of these components in generating movements of the IMR over time. Table 3.8 reveals that all components are responsible for part of the overall dynamics, but with marked differences. In particular, the RAR stands out as the most important component, with a coefficient of the counterfactual series of about 0.92, very close to unity. Even more compellingly, the share of variance accounted for only by movements in RAR is about 74%, substantially more than the one explained by changes in the ER (47%) or in the OMR of the other locations (31%). Having established that the relative attractiveness ratio is a key source of variation in the in-migration rate, we now study whether this component can also explain the stronger reaction of the IMR to job creation shocks than that of the OMR to job destruction.

To do so, we repeat the IV regressions of Section 3.3.2 by using as a LHS variable the

Table 3.8. Relevance of the in-migration rate components

	(1)	(2)	(3)
	IMR	IMR	IMR
IMR keeping $OMR_{s\neq r}$ and RAR constant	0.846**		
	(0.090)		
IMR keeping ER and RAR constant		0.533**	
		(0.171)	
IMR keeping ER and $OMR_{s\neq r}$ constant			0.918**
			(0.038)
N	856	856	856
Within-R^2	0.469	0.309	0.744

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the correlation between in-migration rates at the province-year level and three counterfactual in-migration rate variables constructed by moving one of its three components at a time and keeping the others constant at the average of the period 2011-2018. The three components, namely the employment ratio, the out-migration rate of other provinces, and the relative attractiveness ratio, are computed using SISCO data as described in Section 3.3.3. Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level. $^{+}p < .10, ^{*}p < .05, ^{**}p < .01$.

three components of the in-migration rate (ER, $OMR_{s\neq r}, RAR$) after taking logs.[40] Table 3.9 reveals that virtually the whole positive reaction of IMR to JCR shocks can in fact be traced back to changes in the relative attractiveness ratio. This means that positive shocks do *not* cause more workers to decide to relocate, but rather that they cause relocating households to change their desired destination. We believe that this is a crucial finding that enhances our understanding of the migration reaction to labour demand shocks. First, we document that the extensive margin (the decision to in-migrate or not) is almost unresponsive to shocks, whereas the ranking of destinations in a revealed preference sense is highly sensitive to local economic conditions. Moreover, we uncover negative externalities, as positive shocks in a region relatively worsen the RAR for other competing regions. Last, we claim that the strong reaction of the RAR largely explains the asymmetry between the effect of positive shocks on in-migration and the one of negative shocks on out-migration. In fact, through the lens of our separate regressions, we show that positive shocks also do not cause more relocation, but rather they act on a substantially different margin, one that is not at play for out-migration.

[40]Taking logs is needed to make the decomposition into a sum. Furthermore, this also helps at making the scale of the variables comparable, being otherwise very different.

Table 3.9. In-migration rate components and labour dynamism, IV 2SLS

Threshold IV	250 workers			500 workers			1000 workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	ER	$OMR_{s \neq r}$	RAR	ER	$OMR_{s \neq r}$	RAR	ER	$OMR_{s \neq r}$	RAR
JCR	-0.043	-0.058	10.154*	0.102	-0.061	9.598*	0.048	-0.130	10.178^{+}
	(0.843)	(0.077)	(4.397)	(0.750)	(0.065)	(3.807)	(0.998)	(0.160)	(6.040)
JDR	1.029^{+}	-0.036	-2.833*	0.772	-0.022	-1.678**	0.860	-0.029	-1.500*
	(0.616)	(0.034)	(1.294)	(0.533)	(0.025)	(0.635)	(0.539)	(0.038)	(0.622)
N	856	856	856	856	856	856	856	856	856

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the 2SLS estimates of job creation and job destruction rates, instrumented by mass layoff and hiring events, on the three components of yearly in-migration flow rates from SISCO data as described in Section 3.3.3, namely the employment ratio, the out-migration rate of other provinces, and the relative attractiveness ratio. The three ratios are taken in logarithms to allow comparability as they are scaled differently. Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level. $^{+}p < .10, ^{*} p < .05, ^{**} p < .01$.

3.3.4 Heterogeneity analysis

Effect by Distance of Migration

The richness of the SISCO data allows us to further characterize the flows into and out of an area in response to job creation and job destruction. In particular, we focus on the distance to see how local the effects of labour demand are: this margin is relevant from a policy perspective, as it informs about the extent of geographical mobility in response to different labour demand shocks.

First, we show two case studies to exemplify our empirical strategy. Figure 3.9, panel a, shows the percentage point change in migration rate from each province to Milan after a mass hire event in Milan in 2012. The strongest effect is concentrated in the nearby provinces of Como and Monza-Brianza, north of Milan, and Lodi and Cremona, south-east of Milan, while also other nearby provinces experience smaller impacts and more distant provinces, such as Genova, south of Milan, experience only a minor rise in migration flows towards Milan. Interestingly, also the provinces of Livorno and Grosseto, in the Central region of Tuscany, experience a rise in migration directed to Milan. In panel b, we show the same estimates, but for an opposite-sign event, i.e., a mass layoff occurred in Palermo, Sicily, in 2013. Interestingly, the outflows of workers from Palermo to other provinces are much smaller and less dispersed across space: people moved mainly to the province of Caltanissetta, south of Palermo.

Second, we show more systematically the geographical extent of migration flows in response to labour demand shocks to confirm that the examples are not driven by idiosyn-

Figure 3.9. The geographical impact of local mass events

(a) In-migration rate in Milan after a mass hire event in 2012 (b) Out-migration rate out of Palermo after a mass layoff event in 2013

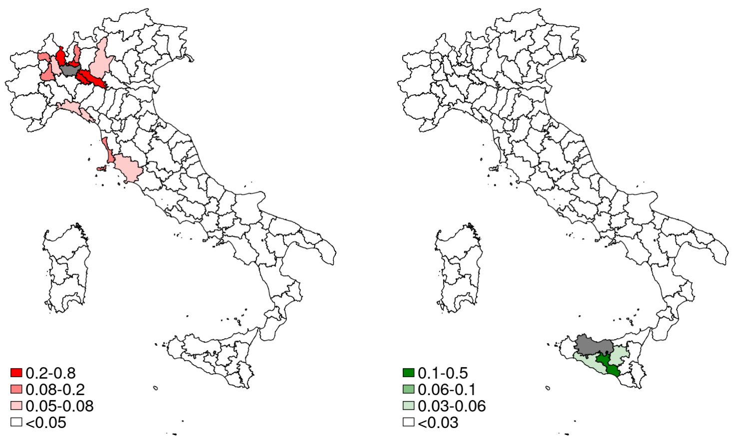

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the geographical distribution of the effects of two large mass events in the SISCO data. Panel (a) shows the percentage point change in the in-migration rate in Milan from each Italian province following a mass hire event that occurred in 2012 in Milan. Panel (b) shows the percentage point change in the out-migration rate from Palermo to each Italian province following a mass layoff event that happened in 2013 in Palermo. All underlying regressions include time and origin-destination fixed effects.

cratic factors (e.g., regional differences). We replicate the analysis of equations (3.5)-(3.7), estimated by 2SLS, binning the migration flows into five groups: less than 50 km of distance from the province main city, between 50 and 100 km, between 100 and 200 km, between 200 and 400 km and more than 400 km away. The results, reported in Table C.9 for different thresholds, confirm the descriptive patterns presented in the previous exercise. The out-migration rate responds only to a change in JDR and is a rather local phenomenon (panel a, columns (1) and (11), less than 50 km). On the contrary, in-migration flows (panel b) decline in response to a rise in JDR, but the effects are never statistically significant. The increase in IMR in response to JCR has instead a much larger geographical reach, though with a decaying intensity.[41]

Other Heterogeneity

We finally study whether particular characteristics of the locations, mainly related to the composition of their population, are associated with different responsiveness of migration flows. To do so, for each of the variables considered we run separate regressions splitting our sample into locations that are above or below the median of the distribution of each

[41] Note that the first bin is rather peculiar, as not all provinces have neighbouring provinces within a 50 km radius. The number of observations in the first bin (columns (1),(6) and (11)) drops to less than two-thirds of those of the other bins and the estimates are rather imprecise.

characteristic of interest. In particular, we study whether the reaction to shocks differs by macro-regions, using the split Center-North vs. South; we further investigate whether the local share of international migrants, highly educated (with or without a college degree), young (below the age of 35) individuals, and the average homeownership rate are associated to different coefficients (the overall trends for these groups are plotted in Figure C.13).

As shown in Figure C.14, we find some evidence that the reaction of out-migration to negative shocks depends on these variables (panel d). More precisely, out-migration reacts more strongly to JDR in the Center-North and in locations where the incidence of foreign-born and highly educated is larger. This is consistent with the evidence shown in the literature by Cadena and Kovak (2016); Basso et al. (2019); Basso and Peri (2020), among others, who show greater relocation responsiveness of foreign-born to local shock. Moreover, we also find that the homeownership rate is negatively associated with the reaction of out-migration, suggesting that households may be locked in, possibly due to housing market frictions. This finding echoes the results of an expanding literature on migration frictions associated with homeownership recently reviewed by Jia et al. (2022). Finally, and perhaps surprisingly, we find that the reaction of out-migration is actually more muted in locations populated by a larger share of young individuals. Indeed, in our data, prime-aged individuals (with an age between 35 and 54) account for the bulk of the reaction in out-migration flows. Instead, we do not find any significant differences in all other coefficients (see panels a,b and c of Figure C.14).

Finally, given that the reaction of in-migration may depend on the average characteristics of the origin locations, we repeat the same separate regressions splitting locations according to the characteristics of neighbouring regions, weighting them with the inverse of the distance in kilometres. However, we do not find any significant association between the in-migration reaction and these characteristics (see in particular panels a and c of Figure C.15).

3.4 Concluding remarks

Exploiting high-quality administrative data on the universe of labour market flows, this paper studies the effect of local labour demand shocks on internal migration in Italy. We start by documenting several novel facts on job flows as well as on internal migration. In particular, we uncover systematic differences between gross and net flows, suggesting that they capture different aspects of worker transitions. Moreover, we show that gross flows going in opposite directions (e.g. job creation and job destruction) are both important to quantitatively account for fluctuations of net rates at the local level. Consistently with this descriptive evidence, we then turn to estimate the causal impact of both positive and negative labour demand shocks on gross internal migration flows.

Our empirical strategy is based on plausibly exogenous shifts of labour demand stemming from mass hire or layoff events at the establishment level. We validate our strategy by means of event studies both at the firm and at the province level. Our 2SLS estimates reveal that both job creation and job destruction have important (opposite) effects on net in-migration, with the former being four times larger than the latter. These differences are due to the greater effect of job creation on in-migration (along with a muted response of out-migration), as opposed to the smaller – though not negligible – impact of job destruction on the out-migration rate. The literature typically takes a reduced-form approach focusing on one margin at a time, such as estimating region-level mass-layoffs. Papers such as Gathmann et al. (2020) and Monras (2018) find that reduced inflows are by and large the most relevant margin of adjustment after a local employment crisis, while we show that, more generally, all margins of labour creation and destruction matter in explaining gross and net migration flows.

We dig deeper into the large reaction of in-migration to positive shocks, documenting that it is not brought about by an increase in the number of relocating workers, but rather by a reshuffling among their preferred alternatives. We see this as a crucial finding, both for correctly modelling migration choices and for better understanding the different reactiveness to labour market shocks of in- vs. out-migration. Finally, we investigate whether the estimated effects vary by distance. We find that job creation induces larger in-migration responses even from relatively distant locations, whereas job destruction causes out-migration which remains locally concentrated. These findings are highly relevant for the spatial distribution of welfare gains following local labour demand shocks.

Overall, our results remark on the importance of accounting for the gross job and internal migration flows when designing labour market and social policies. The extent of gross migration flows – much larger than that of net flows (in the order of 2 per cent across regions) – suggests that considerable frictions could occur in accommodating incoming workers into a region. The design of active labour market, social and housing policies should take into account the extent of overall people's movements and minimize the potential congestion and frictions that may arise. This issue could not be assessed if policymakers were to look only at low-frequency net flows. Moreover, and most relevant, our regression analysis highlights the asymmetry generated by job creation and job destruction in generating movements of people across areas. Indeed, we document that the consequences of positive shifts in labour demand are more largely shared across space through internal migration; the incidence of negative labour demand is instead mainly, though not completely, local. Such asymmetry might be a source of inequality of job opportunities across areas in the short run. Therefore, policymakers may want to pose greater attention to the realization of negative events and act accordingly to avoid the insurgence of cross-regional disparities. Future research will uncover relevant underlying determinants of geographical mobility that will better inform policy decisions. The richness of the SISCO data, in particular, will allow investigation of the dimensions of heterogeneity such as workers' education level, demographic characteristics and the geographical distribution of occupations and job opportunities.

Appendix C

Additional material

C.1 Additional figures and tables

Figure C.1. Average job and worker flows by contract type

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the average aggregate yearly job and worker flow rates, by type of contract. Job flows are the sum across establishments of net flows at the establishment-level, while worker flows are the sum of establishment-level gross flows. The flows are divided by the stock of payroll employment in the current period taken from the ILFS data.

Figure C.2. Average job flow rates across provinces, by contract type

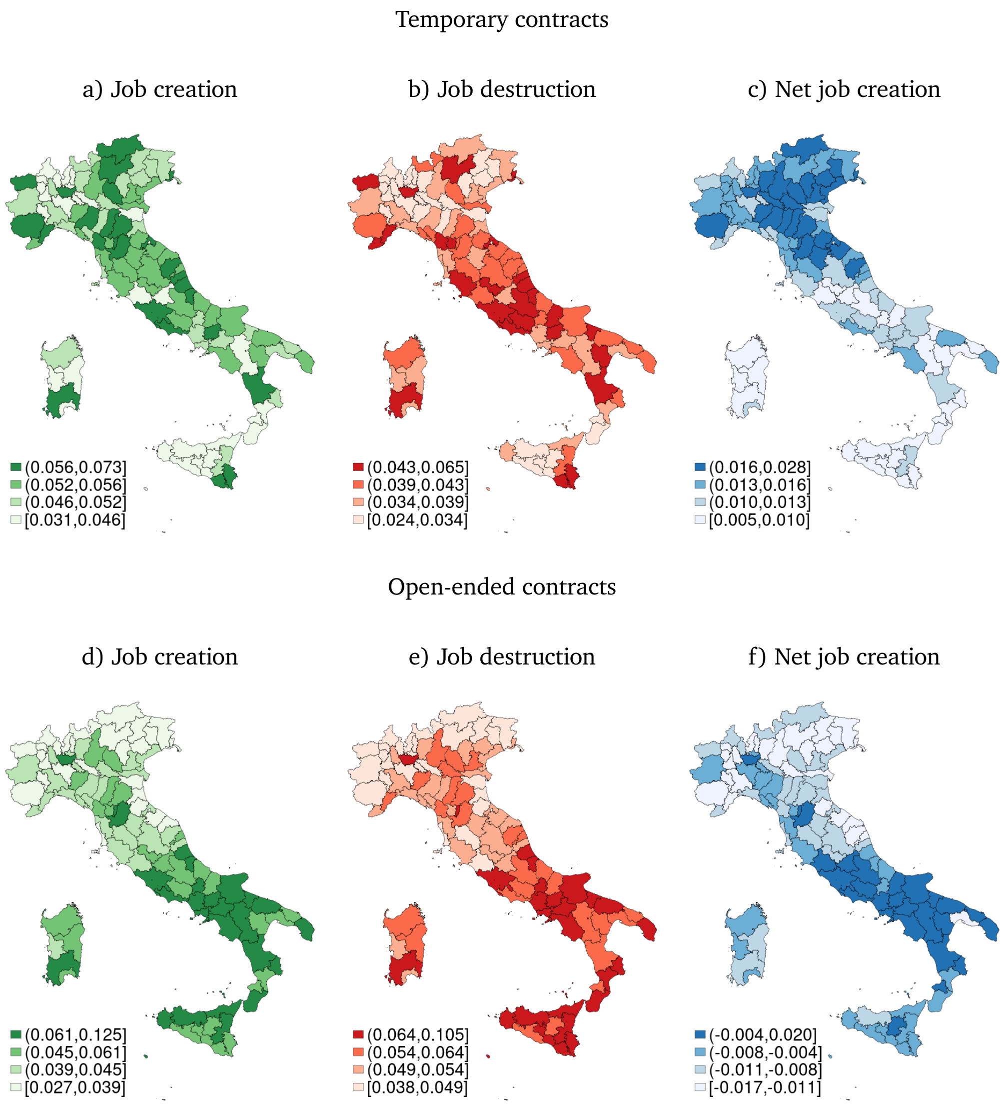

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the geographical distribution of average job flow rates. Job flow rates are the sum across establishments of net flows at the establishment-level, divided by the stock of payroll employment in the current period taken from the ILFS data.

Figure C.3. Gross vs. net job flow rates

a) Yearly frequency b) Monthly frequency

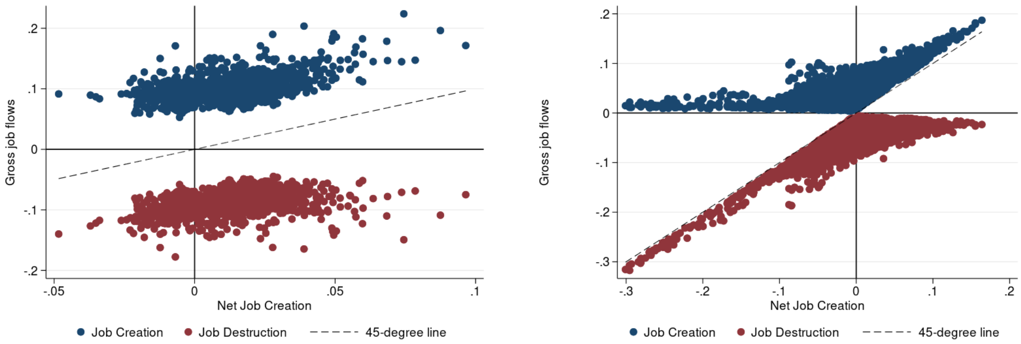

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows province-level gross job flow rates against the corresponding net flows, at yearly (panel (a)) and monthly (panel (b)) frequency. The flows are divided by the stock of payroll employment in the current period taken from the ILFS data. Dashed lines represent the 45-degree lines.

Figure C.4. Average gross vs. net job flow rates, with density

a) Yearly frequency b) Monthly frequency

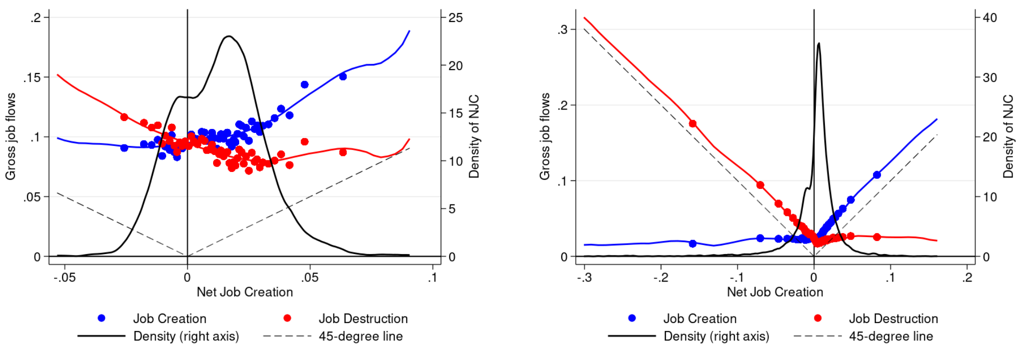

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows province-level gross job flow rates against the corresponding net flows. The flows are divided by the stock of payroll employment in the current period taken from the ILFS data. Solid red and blue lines are the predictions of second-degree local polynomial regressions. Scatter points represent averages of two percentiles of the underlying distribution. The black solid line represents the kernel density of the employment growth rate distribution. Dashed lines represent the 45-degree lines.

Table C.1. Decomposition, Labour market dynamism

	Monthly data		Quarterly data		Yearly data	
Panel (a). Municipality level						
	(1)	(2)	(3)	(4)	(5)	(6)
	JCR	JDR	JCR	JDR	JCR	JDR
JNCR	0.433**	-0.567**	0.502**	-0.498**	0.529**	-0.471**
	(0.014)	(0.014)	(0.009)	(0.009)	(0.013)	(0.013)
N	760,800	760,800	253,600	253,600	63,400	63,400
R^2	0.572	0.713	0.752	0.769	0.535	0.485
Panel (b). Province level						
	(1)	(2)	(3)	(4)	(5)	(6)
	JCR	JDR	JCR	JDR	JCR	JDR
JNCR	0.347**	-0.653**	0.464**	-0.536**	0.632**	-0.368**
	(0.016)	(0.016)	(0.018)	(0.018)	(0.039)	(0.039)
N	10,272	10,272	3,424	3,424	856	856
R^2	0.717	0.906	0.845	0.903	0.760	0.790
Panel (c). Region level						
	(1)	(2)	(3)	(4)	(5)	(6)
	JCR	JDR	JCR	JDR	JCR	JDR
JNCR	0.373**	-0.627**	0.475**	-0.525**	0.713**	-0.287**
	(0.060)	(0.060)	(0.056)	(0.056)	(0.044)	(0.044)
N	1,920	1,920	640	640	160	160
R^2	0.771	0.921	0.868	0.929	0.911	0.917

Source: SISCO and ILFS data 2010-2018. *Note*: The table shows the results of linear regressions of gross against net labour market flows, for different geographical and time aggregation levels. Estimated coefficients can be interpreted as the share of the total variance of net flows accounted for by the variation in the specific gross flow. All regressions include time and location fixed effects. Standard errors reported in parentheses are clustered at the location level. $^+p < .10, ^* p < .05, ^{**} p < .01$.

Table C.2. Internal migration and labour dynamism, 'full' dataset 2SLS

Threshold IV	250 workers			500 workers			1,000 workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	OMR	IMR	NIMR	OMR	IMR	NIMR	OMR	IMR	NIMR
JCR	-0.013	0.429*	0.441**	0.001	0.394*	0.393*	0.026	0.422^{+}	0.396*
	(0.052)	(0.182)	(0.157)	(0.047)	(0.156)	(0.133)	(0.080)	(0.226)	(0.163)
JDR	0.049^{+}	-0.049	-0.097**	0.038^{+}	-0.017	-0.055^{+}	0.051*	-0.011	-0.062*
	(0.026)	(0.036)	(0.037)	(0.021)	(0.026)	(0.028)	(0.023)	(0.037)	(0.027)
N	856	856	856	856	856	856	856	856	856

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the 2SLS estimates of job creation and job destruction rates, instrumented by mass layoff and hiring events, on migration flow rates from SISCO data as described in Section 3.3.2. The 'full' dataset is built assuming that the worker's location corresponds to her last workplace location until a new job is found. We allow for a maximum of one year out of employment before relocation. We selected events that involve more than 250 (specifications (1)-(3)), more than 500 (specifications (4)-(6)), or more than 1,000 (specifications (7)-(9)) net activations or terminations. Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level. The standard errors in parentheses are clustered at the province level. $^{+}p < .10, ^{*}p < .05, ^{**}p < .01$.

Figure C.5. Average internal migration rates across provinces using residence-based data

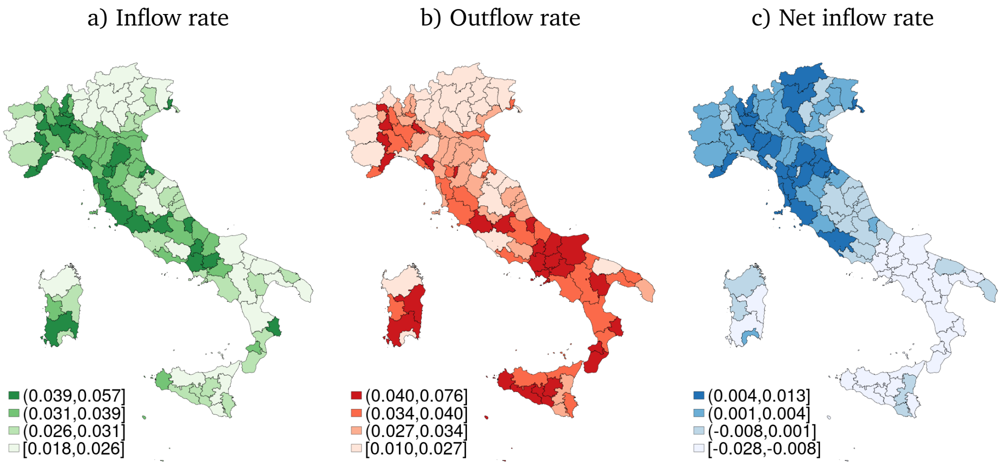

Source: Istat data on residence changes and ILFS data 2010-2018. *Note*: The figure shows the geographical distribution of average internal migration rates across provinces, computed using administrative data on residence changes (Istat). The migration flows are divided by the stock of payroll employment in the previous period taken from the ILFS data.

Table C.3. Decomposition, Internal migration

Panel (a). Municipality level

	Monthly data		Quarterly data		Yearly data	
	(1)	(2)	(3)	(4)	(5)	(6)
	IMR	OMR	IMR	OMR	IMR	OMR
NIMR	0.384**	-0.616**	0.503**	-0.497**	0.487**	-0.513**
	(0.082)	(0.082)	(0.016)	(0.016)	(0.017)	(0.017)
N	760,800	760,800	253,600	253,600	63,400	63,400
R^2	0.351	0.576	0.581	0.575	0.484	0.507

Panel (b). Province level

	Monthly data		Quarterly data		Yearly data	
	(1)	(2)	(3)	(4)	(5)	(6)
	IMR	OMR	IMR	OMR	IMR	OMR
NIMR	0.288**	-0.712**	0.483**	-0.517**	0.434**	-0.566**
	(0.074)	(0.074)	(0.011)	(0.011)	(0.083)	(0.083)
N	10,272	10,272	3,424	3,424	856	856
R^2	0.505	0.743	0.710	0.741	0.428	0.460

Panel (c). Region level

	Monthly data		Quarterly data		Yearly data	
	(1)	(2)	(3)	(4)	(5)	(6)
	IMR	OMR	IMR	OMR	IMR	OMR
NIMR	0.216**	-0.784**	0.524**	-0.476**	0.429**	-0.571**
	(0.036)	(0.036)	(0.027)	(0.027)	(0.072)	(0.072)
N	1,920	1,920	640	640	160	160
R^2	0.532	0.862	0.823	0.791	0.627	0.651

Source: SISCO and ILFS data 2010-2018. *Note*: The table shows the results of linear regressions of gross against net internal migration flows, for different geographical and time aggregation levels. Estimated coefficients can be interpreted as the share of the total variance of net flows accounted for by variation in the specific gross flow. All regressions include time and location fixed effects. Standard errors reported in parentheses are clustered at the location level. $^{+}p < .10, ^{*} p < .05, ^{**} p < .01$.

Figure C.6. Distribution of distance of internal migration

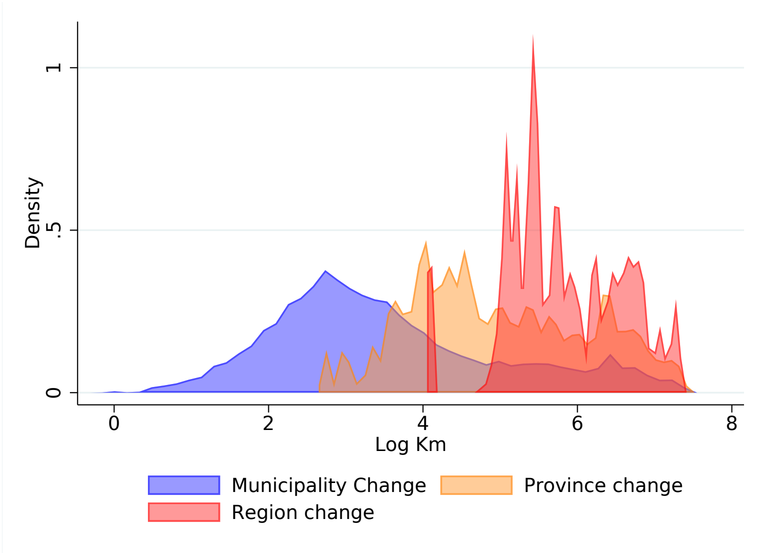

Source: SISCO data 2010-2018. *Note*: The figure shows the distribution of the distance between origin and destination location (log km) of internal mobility individual transitions for different geographical aggregation levels (municipality, province, region).

Table C.4. Summary statistics, Distance of internal migration moves (km)

Statistics	Municipality	Province	Region
Mean	119.0	271.6	439.3
Min	0.8	14.8	59.7
Max	1,809.6	1,756.6	1,598.7
P1	2.3	15.8	59.7
P5	4.4	27.3	59.7
P10	6.4	36.6	160.1
P25	11.9	57.2	193.5
P50	24.9	119.2	300.7
P75	75.8	364.3	619.4
P90	372.1	755.8	929.0
P95	666.9	989.2	1,164.5
P99	1,248.5	1,447.7	1,439.7
N	10,679,725	4,291,503	2,303,722

Source: SISCO data 2010-2018. *Note*: The table shows summary statistics of the distribution of distance (km) of internal migration transitions at the yearly frequency identified from the SISCO microdata, for different geographical levels (municipality, province, region).

Table C.5. Internal migration and labour dynamism, IV 2SLS at municipality-level

Threshold IV	250 workers			500 workers			1,000 workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	OMR	IMR	NIMR	OMR	IMR	NIMR	OMR	IMR	NIMR
JCR	-0.017^{+}	0.286**	0.303**	-0.006	0.261**	0.267**	-0.008	0.217**	0.225**
	(0.010)	(0.050)	(0.050)	(0.010)	(0.054)	(0.049)	(0.012)	(0.032)	(0.029)
JDR	0.065**	-0.004	-0.069^{+}	0.040	0.008	-0.032	0.042*	0.018	-0.025*
	(0.022)	(0.029)	(0.040)	(0.030)	(0.048)	(0.059)	(0.017)	(0.014)	(0.010)
N	63,288	63,288	63,288	63,288	63,288	63,288	63,288	63,288	63,288

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the 2SLS estimates of job creation and job destruction rates, instrumented by mass layoff and hiring events, on yearly migration flow rates at the municipality level. Each specification includes municipality and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the municipality level. $^{+}p < .10, ^{*}p < .05, ^{**}p < .01$.

Figure C.7. Geographical distribution of mass hires and layoffs in Italy. Events > 250

a) Mass hire events

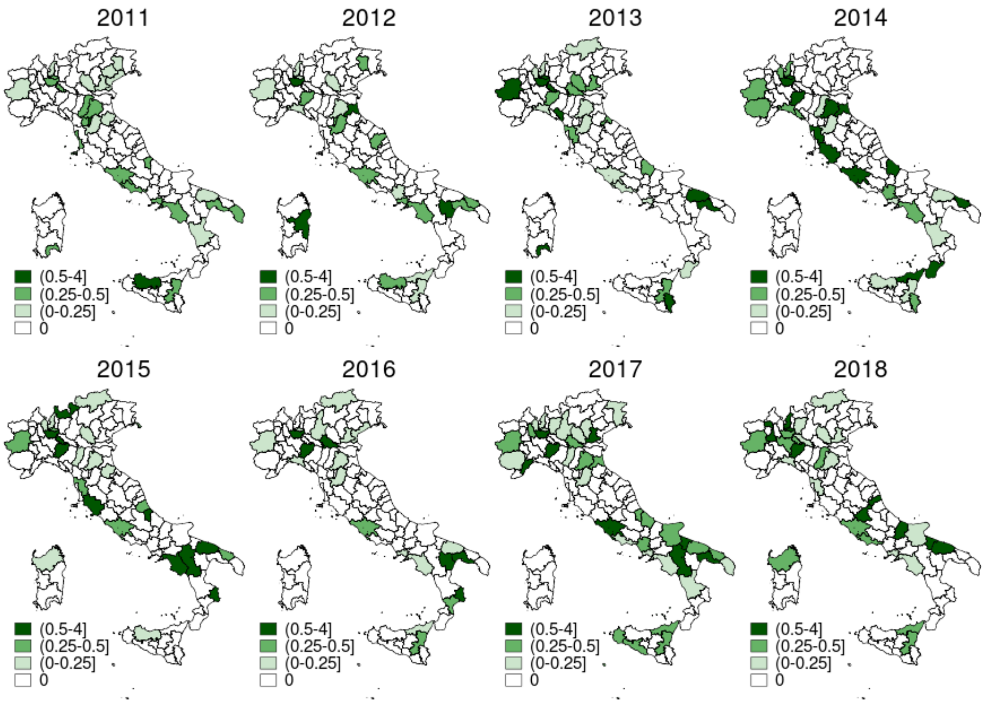

b) Mass layoff events

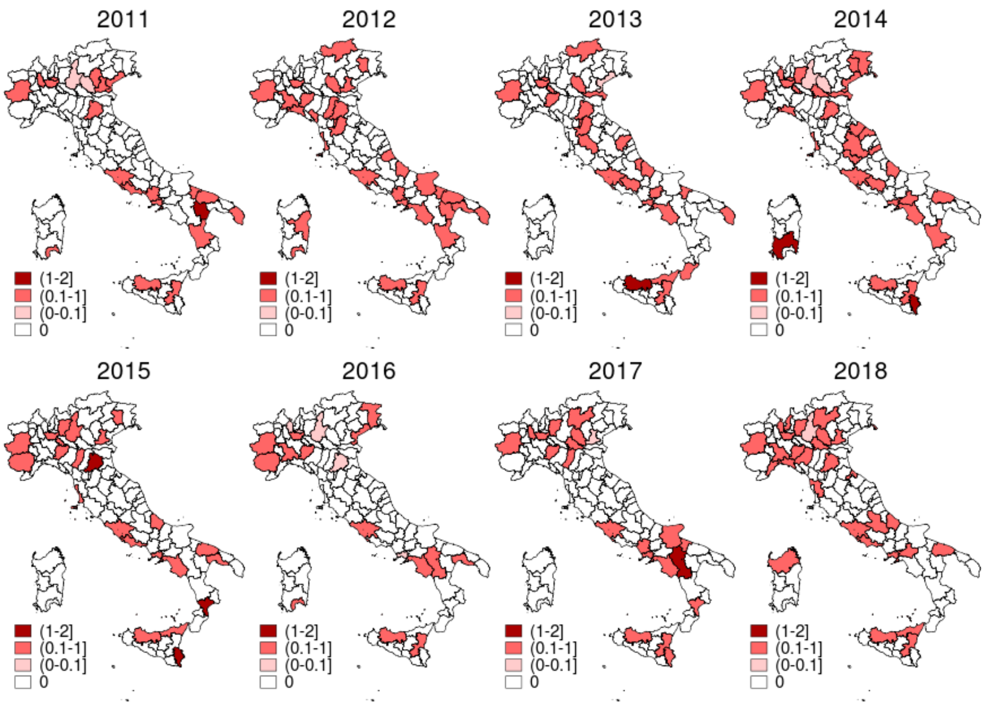

Source: SISCO and ILFS data 2011-2018. *Note*: The figure shows the maps of mass hires and layoffs as a percentage of local employment, only for events that involve more than 250 net activations or terminations. Local employment is taken from the ILFS data. The unit of analysis is the establishment.

Figure C.8. Geographical distribution of mass hires and layoffs in Italy. Events > 500

a) Mass hire events

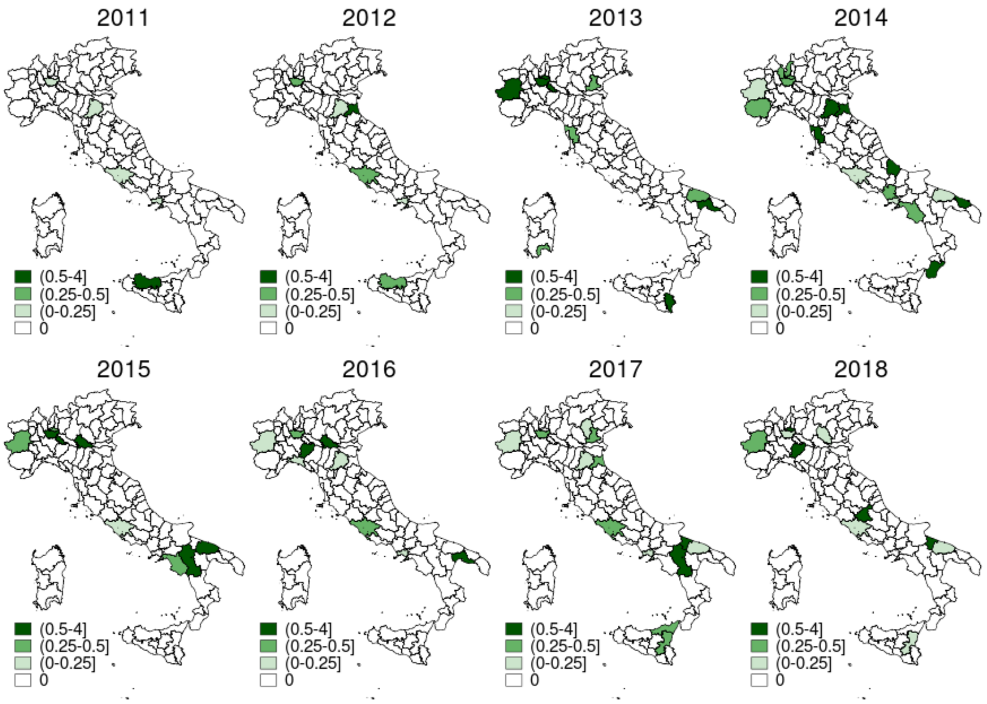

b) Mass layoff events

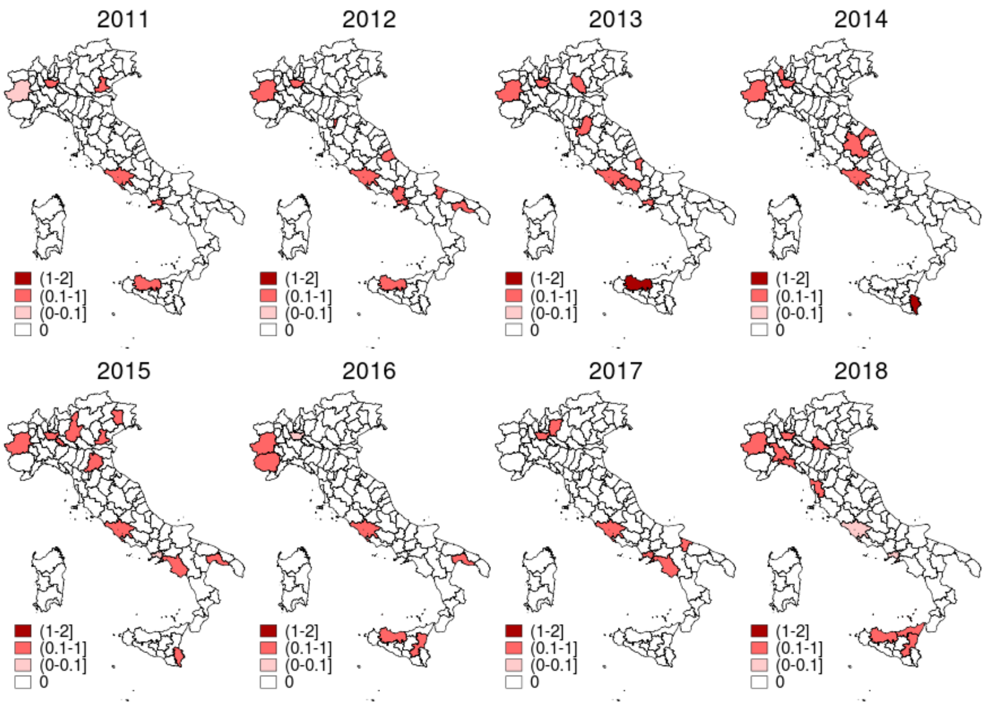

Source: SISCO and ILFS data 2011-2018. *Note*: The figure shows the maps of mass hires and layoffs as a percentage of local employment, only for events that involve more than 500 net activations or terminations. Local employment is taken from the ILFS data. The unit of analysis is the establishment.

Figure C.9. Geographical distribution of mass hires and layoffs in Italy. Events $> 1,000$

a) Mass hire events

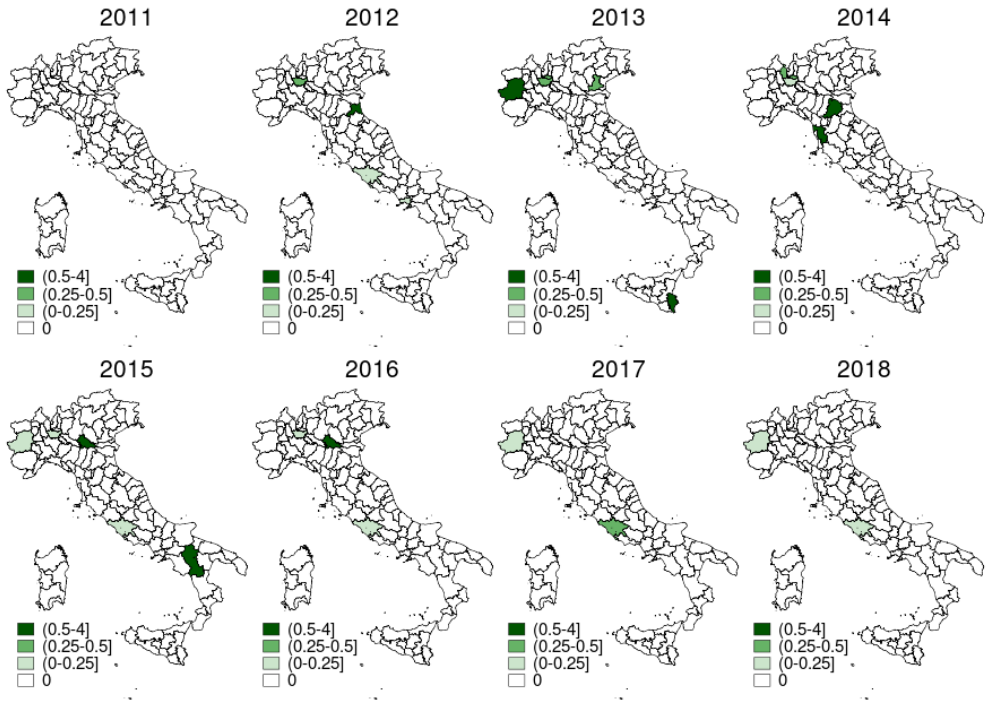

b) Mass layoff events

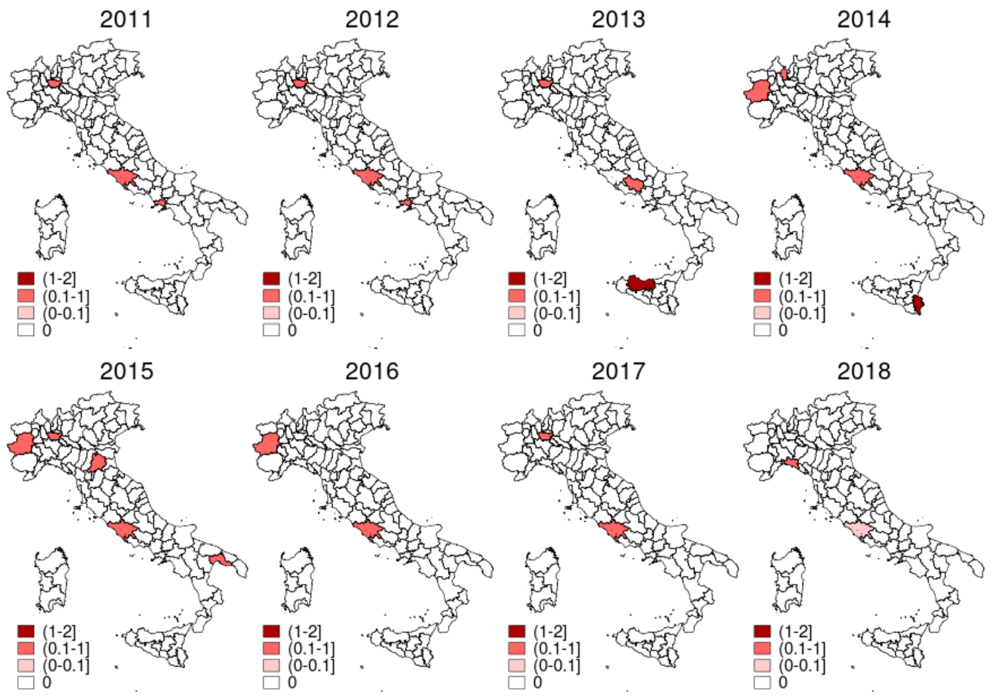

Source: SISCO and ILFS data 2011-2018. *Note*: The figure shows the maps of mass hires and layoffs as a percentage of local employment, only for events that involve more than 1,000 net activations or terminations. Local employment is taken from the ILFS data. The unit of analysis is the establishment.

Table C.6. Event studies summary statistics by macro area

	Mass hire events					Mass layoff events				
	North-East	North-West	Center	South	Islands	North-East	North-West	Center	South	Islands
Panel (a). Events with at least 250 workers involved										
Events	71	228	123	86	40	56	186	132	98	48
Provinces hit	14	18	12	17	8	16	14	15	15	9
Min (Avg) Max size	251(411)2055	251(521)4861	252(488)2713	251(427)1837	251(485)3169	251(370)2057	251(488)3460	252(470)2487	251(464)8414	252(463)3111
% of prov. empl.	0.09(0.31)1	0.07(0.43)4.15	0.08(0.40)1.98	0.08(0.38)2.02	0.10(0.44)2.05	0.07(0.22)0.80	0.07(0.31)0.92	0.10(0.41)0.94	0.08(0.49)6.87	0.11(0.43)1.75
By industry										
Manifacturing	6	9	5	2	0	4	16	10	11	4
Construction	0	2	0	1	3	2	1	1	0	2
Private services	55	200	108	55	21	40	135	97	42	13
Public services	10	16	7	22	11	8	20	16	34	17
.	0	1	3	6	5	2	14	8	11	12
Panel (b). Events with at least 500 workers involved										
Events	14	60	28	19	9	4	45	30	18	13
Provinces hit	6	8	4	8	4	4	10	6	7	4
Min (Avg) Max size	513(806)2055	501(1030)4861	509(1020)2713	527(769)1837	520(1025)3169	555(1073)2057	508(983)3460	501(946)2487	516(1101)8414	522(796)3111
% of prov. empl.	0.16(0.43)1	0.06(0.53)3.78	0.04(0.49)1.98	0.09(0.49)2.02	0.24(0.86)2.05	0.25(0.43)0.59	0.05(0.29)0.71	0.04(0.36)0.94	0.09(0.72)6.87	0.20(0.41)1.48
By industry										
Manifacturing	1	1	2	1	0	0	2	3	2	1
Construction	0	0	0	0	0	0	0	0	0	0
Private services	12	59	24	14	6	3	36	23	7	5
Public services	1	0	1	4	1	0	3	3	7	5
.	0	0	1	0	2	1	4	1	2	2
Panel (c). Events with at least 1,000 workers involved										
Events	2	1	16	12	2	10	1	2	2	1
Provinces hit	2	1	4	3	10	2	1	2	2	1
Min (Avg) Max size	1262(1659)2055	2057	1023(2020)4861	1013(1796)3460	1028(1688)2713	1099(1583)2487	1837	1263(4839)8414	1728(2449)3169	3111
% of prov. empl.	0.62(0.81)1	0.59	0.14(0.64)3.78	0.13(0.24)0.40	0.10(0.37)1.08	0.09(0.24)0.87	2.02	0.22(3.54)6.87	1.21(1.63)2.05	1.25
By industry										
Manifacturing	0	0	1	1	0	1	1	1	0	0
Construction	0	0	0	0	0	0	0	0	0	0
Private services	2	1	15	8	10	8	0	0	1	0
Public services	0	0	0	0	0	1	0	1	1	0
.	0	0	0	3	0	0	0	0	0	1

Source: SISCO data 2010-2018. *Note*: The table reports the summary statistics for mass layoff and hiring events with thresholds set at 250 (panel a), 500 (panel b) and 1,000 (panel c) units calculated from SISCO data as described in Section 3.3.2 by geographic macro area. Public services include privately-provided education and health services.

Table C.7. Internal migration and labour dynamism, 'mass events dummy' IV 2SLS

Threshold IV	250 workers			500 workers			1,000 workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	OMR	IMR	NIMR	OMR	IMR	NIMR	OMR	IMR	NIMR
JCR	-0.022	0.206^{*}	0.228^{+}	0.045	0.334^{**}	0.289^{**}	0.340	0.953	0.613^{+}
	(0.104)	(0.094)	(0.121)	(0.098)	(0.112)	(0.091)	(0.656)	(0.961)	(0.332)
JDR	-0.026	-0.093	-0.067	-0.056	-0.036	0.020	0.185	0.265	0.080
	(0.099)	(0.087)	(0.103)	(0.060)	(0.071)	(0.084)	(0.123)	(0.244)	(0.129)
N	856	856	856	856	856	856	856	856	856

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the 2SLS estimates of job creation and job destruction rates, instrumented by mass layoff and hiring events expressed as dummies rather than as a fraction of local employment, on migration flow rates from SISCO data as described in Section 3.3.2. We selected events that involve more than 250 (specifications (1)-(3)), more than 500 (specifications (4)-(6)), or more than 1,000 (specifications (7)-(9)) net activations or terminations. Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level. $^{+}p < .10, ^{*}p < .05, ^{**}p < .01$.

Table C.8. Internal migration and labour dynamism, IV 2SLS using data on residence changes (ISTAT)

Threshold IV	250 workers			500 workers			1,000 workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	OMR	IMR	NIMR	OMR	IMR	NIMR	OMR	IMR	NIMR
JCR	0.038	-0.014	-0.053	-0.001	-0.001	0.000	0.049^{*}	-0.076	-0.125
	(0.033)	(0.058)	(0.070)	(0.047)	(0.061)	(0.077)	(0.025)	(0.084)	(0.086)
JDR	0.064^{+}	0.032	-0.033	0.056^{+}	0.033	-0.022	0.042^{+}	0.050	0.008
	(0.033)	(0.033)	(0.044)	(0.031)	(0.029)	(0.040)	(0.024)	(0.030)	(0.035)
N	856	856	856	856	856	856	856	856	856

Source: Residence changes (ISTAT) and ILFS data 2010-2018. *Note*: The table reports the OLS and 2SLS estimates of job creation and job destruction rates, instrumented by mass layoff and hiring events, on yearly migration flow rates from residence changes administrative data (ISTAT). Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level. $^{+}p < .10, ^{*}p < .05, ^{**}p < .01$.

Figure C.10. Mass events intensive margin on JCR and JDR, local projection estimates. Events > 250

(a) Effect of mass hire on JCR

(b) Effect of mass hire on JDR

(c) Effect of mass layoff on JCR

(d) Effect of mass layoff on JDR

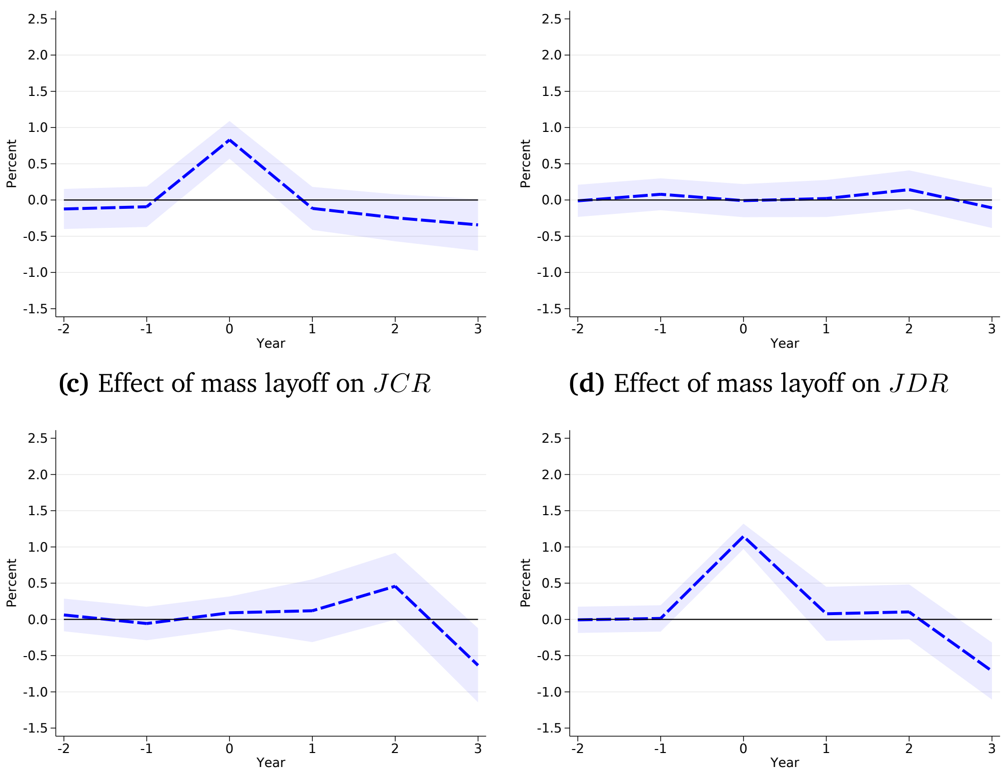

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the local projection estimates of the variation in employment following an establishment-level mass event of more than 250 net activations or terminations on the job creation and the job destruction rates at the province-year level. The intensive margin is obtained by dividing the employment change by the stock of current employment taken from ILFS data. The 95 percent confidence intervals (shaded areas) are estimated by clustering pointwise standard errors at the provincial level.

Figure C.11. Mass events intensive margin on JCR and JDR, local projection estimates. Events > 500

(a) Effect of mass hire on JCR

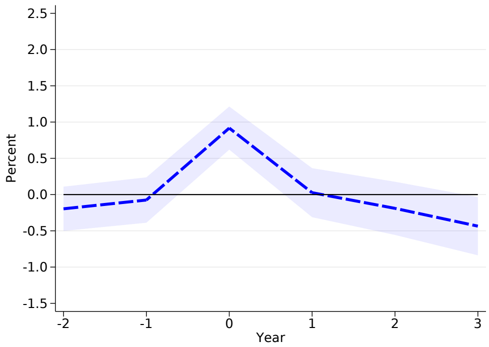

(b) Effect of mass hire on JDR

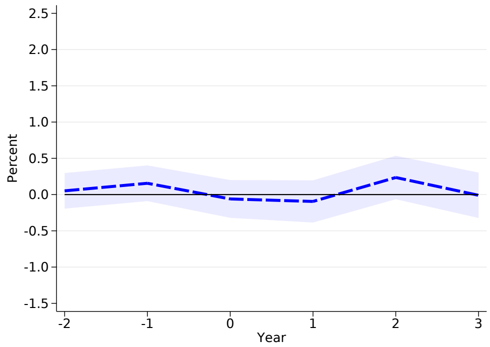

(c) Effect of mass layoff on JCR

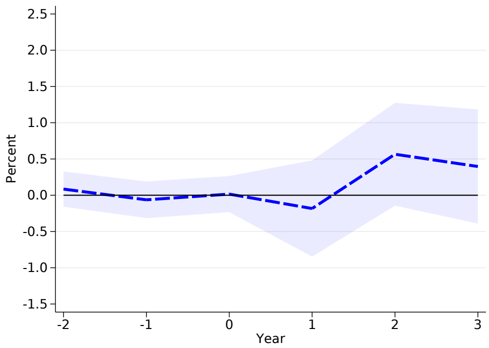

(d) Effect of mass layoff on JDR

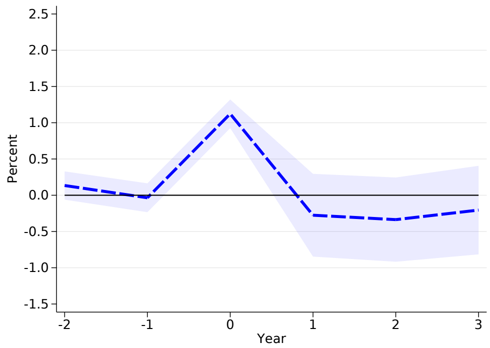

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the local projection estimates of the variation in employment following an establishment-level mass event of more than 500 net activations or terminations on the job creation and the job destruction rates at the province-year level. The intensive margin is obtained by dividing the employment change by the stock of current employment taken from ILFS data. The 95 percent confidence intervals (shaded areas) are estimated by clustering pointwise standard errors at the provincial level.

Figure C.12. Mass events intensive margin on JCR and JDR, local projection estimates. Events $> 1,000$

(a) Effect of mass hire on JCR

(b) Effect of mass hire on JDR

(c) Effect of mass layoff on JCR

(d) Effect of mass layoff on JDR

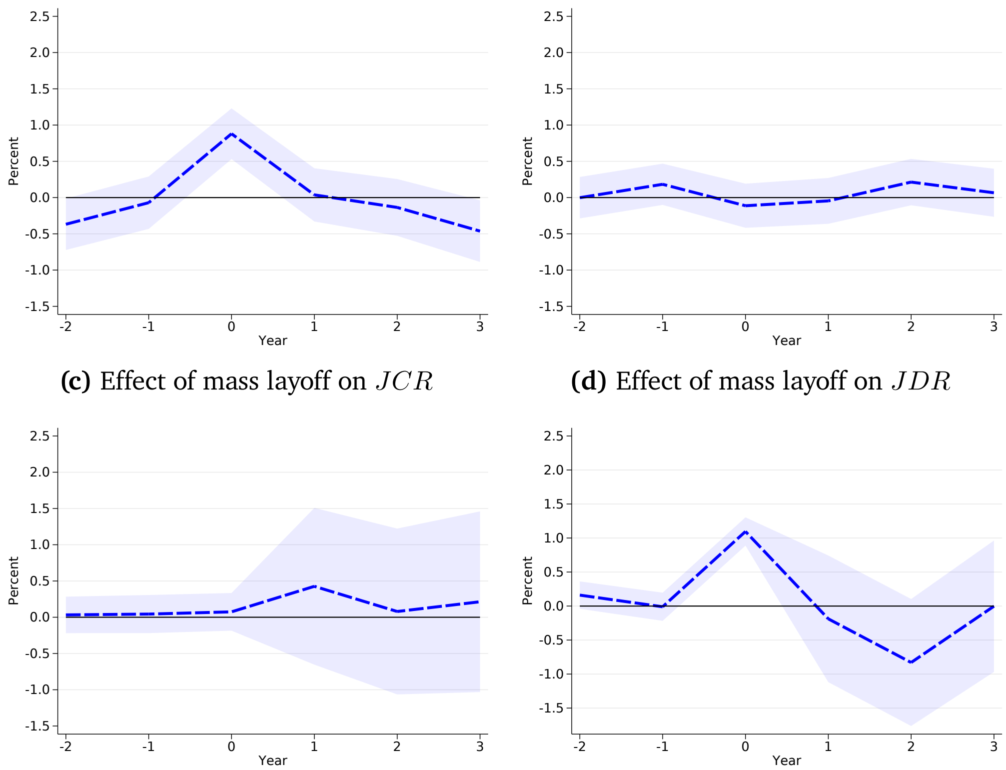

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the local projection estimates of the variation in employment following an establishment-level mass event of more than 1,000 net activations or terminations on the job creation and the job destruction rates at the province-year level. The intensive margin is obtained by dividing the employment change by the stock of current employment taken from ILFS data. The 95 percent confidence intervals (shaded areas) are estimated by clustering pointwise standard errors at the provincial level.

Table C.9. Internal migration and labour dynamism by distance bin, 2SLS

Threshold IV	250 workers					500 workers					1,000 workers				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
km $\in$	[0, 50)	[50, 100)	[100, 200)	[200, 400)	$\geq$ 400	[0, 50)	[50, 100)	[100, 200)	[200, 400)	$\geq$ 400	[0, 50)	[50, 100)	[100, 200)	[200, 400)	$\geq$ 400
Panel (a). OMR															
JCR	-0.006	-0.041*	0.006	0.010	-0.006	-0.013	-0.023	0.012	0.008	-0.009	-0.008	-0.032*	0.008	0.019	0.004
	(0.019)	(0.018)	(0.010)	(0.013)	(0.027)	(0.019)	(0.016)	(0.010)	(0.011)	(0.024)	(0.012)	(0.014)	(0.012)	(0.022)	(0.043)
JDR	0.052^{+}	0.014	0.009	0.008	0.007	0.022	0.010	0.007	0.005	0.005	0.070^{+}	0.019**	0.007	0.008	0.004
	(0.029)	(0.010)	(0.005)	(0.009)	(0.015)	(0.045)	(0.011)	(0.005)	(0.007)	(0.012)	(0.042)	(0.006)	(0.005)	(0.011)	(0.012)
Panel (b). IMR															
JCR	0.039	0.071	0.081**	0.082**	0.085*	0.023	0.047	0.094**	0.077**	0.088**	0.002	0.046	0.120**	0.088*	0.099^{+}
	(0.031)	(0.056)	(0.021)	(0.030)	(0.039)	(0.027)	(0.046)	(0.017)	(0.029)	(0.038)	(0.018)	(0.064)	(0.035)	(0.044)	(0.059)
JDR	0.011	-0.014	-0.014	-0.000	-0.010	0.008	-0.007	-0.006	0.005	-0.005	0.052	-0.001	-0.005	0.003	-0.009
	(0.035)	(0.011)	(0.009)	(0.009)	(0.011)	(0.043)	(0.008)	(0.005)	(0.006)	(0.008)	(0.036)	(0.012)	(0.006)	(0.006)	(0.008)
Panel (c). NIMR															
JCR	0.044^{+}	0.112^{+}	0.075**	0.072**	0.091**	0.036	0.070	0.082**	0.069**	0.097**	0.011	0.077	0.112**	0.068*	0.095**
	(0.027)	(0.062)	(0.018)	(0.027)	(0.033)	(0.024)	(0.047)	(0.014)	(0.026)	(0.033)	(0.013)	(0.063)	(0.028)	(0.031)	(0.035)
JDR	-0.041^{+}	-0.028^{+}	-0.022^{+}	-0.008	-0.017	-0.014	-0.017	-0.013*	-0.000	-0.011	-0.018^{+}	-0.020	-0.012*	-0.005	-0.013
	(0.022)	(0.015)	(0.011)	(0.010)	(0.019)	(0.022)	(0.015)	(0.007)	(0.006)	(0.015)	(0.009)	(0.017)	(0.005)	(0.006)	(0.014)
N	520	824	856	848	856	520	824	856	848	856	520	824	856	848	856

Source: SISCO and ILFS data 2010-2018. *Note*: The table reports the 2SLS estimates of job creation and job destruction rates, instrumented by mass layoff and hiring events, on migration flow rates from SISCO data as described in Section 3.3.2. The specification differentiates migration flows into five distance bins (km $\in [0, 50)$, km $\in [50, 100)$, km $\in [100, 200)$, km $\in [200, 400)$, km ≥ 400). Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level. $^{+}p < .10,^{*} p < .05,^{**} p < .01$.

Figure C.13. Migration rates by socio-demographic groups.

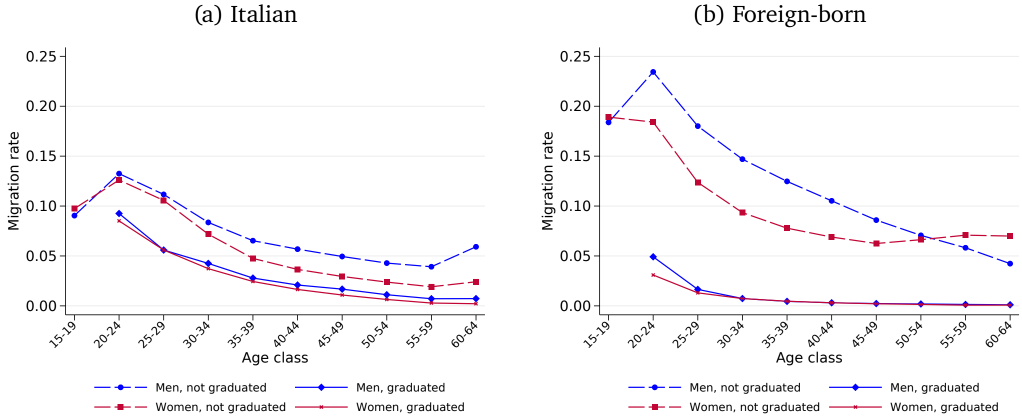

Source: SISCO and ILFS data 2010-2018. *Note*: The figure shows the average age profile of the yearly migration rate between provinces computed by dividing the number of observed transitions in the SISCO microdata by the corresponding stocks of payroll employment in the previous period taken from the ILFS data for each socio-demographic group.

C.2 Geographical mobility: data construction

C.2.1 Prevalent job definition in the SISCO Data

The standard procedure in the economic literature for selecting the prevalent job in a period is to keep the contract with the highest wage and/or duration. Since the SISCO database does not record wages, we rely on the duration as measured in days. However, selecting as prevalent in each period (year, quarter or month) the contract with the longest spell *within* the period may introduce spurious mobility in our sample. Indeed, when more than one contract covers more than one period, the overlapping contracts may have the same duration and be located in different places; if we randomly choose one contract per period, the duration within the period being equal, we may select different workplaces in different times, without the worker having actually changed the main place of work.

To solve this problem and avoid this bias, we select the prevalent contract in each period looking at the *overall* duration of each contract: we delete the contracts whose entire spell is strictly contained in another contract (that started earlier and finished later); in case of partial overlapping, we keep for the overlapping periods only the contract with the longest overall duration; in case of perfect overlapping and exact same duration, we give priority to full-time jobs, open-ended contracts, and jobs that started earlier, following this ordering of criteria. Finally, in the residual cases in which it is still not possible to choose a prevalent contract (jobs started the same day, with the same characteristics and duration), we proceed with a random selection; again, the selection covers the entire overlapping period to avoid the bias in mobility mentioned above.

Figure C.14. Heterogeneity analysis, by characteristics of the province.

(a) Effect of JCR on IMR

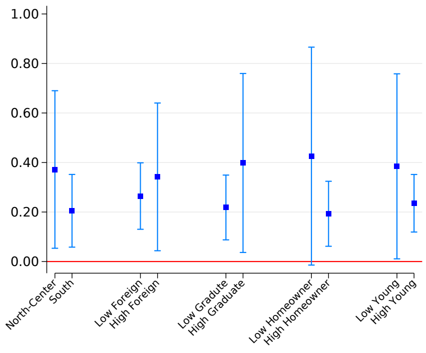

(b) Effect of JCR on OMR

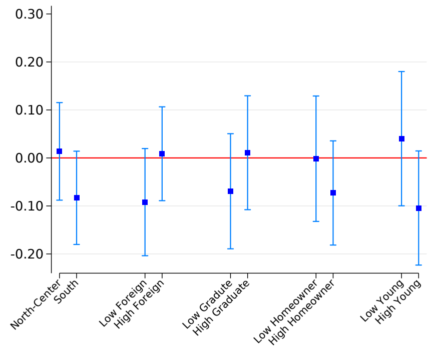

(c) Effect of JDR on IMR

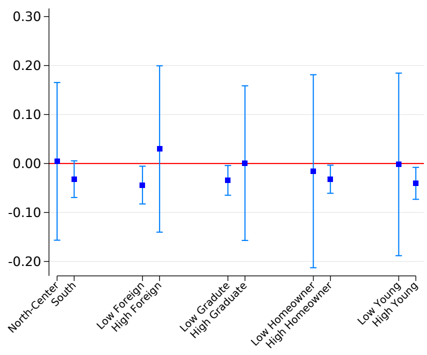

(d) Effect of JDR on OMR

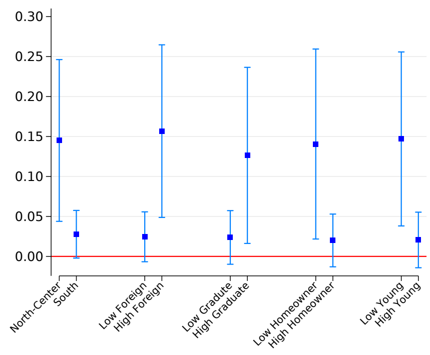

Source: SISCO and ILFS data 2010-2018. *Note*: The figure reports the 2SLS estimates of job creation and job destruction rates at province and year level of aggregation, instrumented by mass layoff and hiring events involving more than 250 workers, on migration flow rates from SISCO data as described in Section 3.3.2. The models are applied separately for each group so as to appreciate any differences in estimates. The "high" categories include those provinces whose average share between 2011 and 2018 of the characteristic under consideration in the population is higher than the median. Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level, and plotted confidence intervals are at 90% confidence level.

C.2.2 Mobility definition in the SISCO Data

We record an internal mobility flow from location A to location B in period t whenever a worker has a prevalent contract in location A in period $t-1$ and a prevalent contract in location B in period t. An outflow for location A and an inflow for location B are registered in t. Using this procedure, in every period the sum of the inflows equals the sum of the outflows. In our robustness checks, we use an extended version of the data set in which the missing observations (non-employment spells) in the career are filled assuming that the worker is looking for a job in the location of the last employment relationship: therefore, in the case of a job flow from unemployment in period t, the outflow in t is attributed to the last location of work even if it is distant in time. Since we assume that the worker remains

Figure C.15. Heterogeneity analysis, by characteristics of the neighbouring provinces.

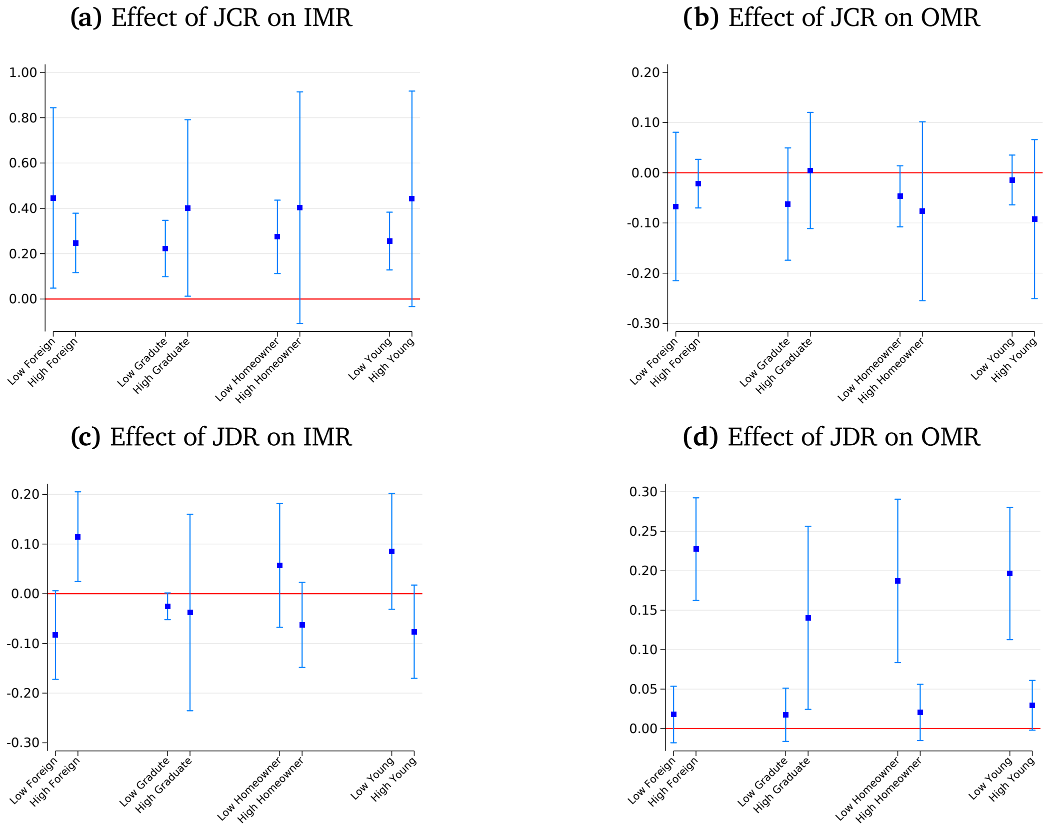

Source: SISCO and ILFS data 2010-2018. *Note*: The figure reports the 2SLS estimates of job creation and job destruction rates at province and year level of aggregation, instrumented by mass layoff and hiring events involving more than 250 workers, on migration flow rates from SISCO data as described in Section 3.3.2. The models are applied separately for each group so as to appreciate any differences in estimates. The "high" categories include those provinces whose neighbouring provinces have an average share between 2011 and 2018 of the characteristic under consideration in the population higher than the median. The share is obtained by averaging the shares of the other provinces weighted by the inverse of the distance in km. Each specification includes province and year fixed effects, and the observations are weighted using the stock of payroll employment in the current period taken from the ILFS data. The standard errors in parentheses are clustered at the province level, and plotted confidence intervals are at 90% confidence level.

at the place of his last job until a new contract is activated, we restrict the analysis to cases where the unemployment period is only one year in order to avoid introducing spurious mobility.

C.2.3 Details on distance statistics

ISTAT releases origin-destination matrices of distances in metres and travel times in minutes between all Italian municipalities (using the centroids in 2013), computed using a commercial road graph. For the islands, it provides the internal distances and those between the main ports of connection with the peninsula. We therefore complement information about internal distances with those between internal municipalities and the nearest port. From the end of 2013 and 1 January 2019, 687 transformations of municipalities took place, of which: 156 changes of province due to the creation and then suppression of new provinces in Sardinia; 1 change of region of the municipality of Sappada from Veneto to Friuli-Venezia-Giulia; 260 new institutions from mergers of pre-existing municipalities; 270 consequent terminations of merged municipalities. We adjust the distances taking into account these transformations, and use symmetric distances for simplicity. For provinces and regions, we use the distance between the capital municipalities. For the municipalities of Monte Isola and Campione d'Italia, for which ISTAT does not provide the distances from the rest of the Italian municipalities, we use the data of the near municipalities of Sulzano and Alta Valle Intelvi, respectively.

C.3 Measuring internal migration through workplace-based data: a comparison with other traditional data sources

Measuring internal migration through SISCO microdata entails several benefits with respect to existing datasets (ILFS microdata or aggregate administrative data from ISTAT). First, the information on geographical location is very detailed (i.e., the municipality of work both at origin and at destination) and so is the frequency of movements (potentially at the daily level). Moreover, the data allow to construct aggregate gross (and net) migration flows accounting for many individual and job characteristics. More traditional data sources such as the Italian LFS data and administrative data on residence transfers fall short in several dimensions. The former only asks retrospective questions on the residence the year before at the provincial level. Such a rather coarse measure might fail to capture intra-annual movements across provinces as well as all within-province movements. Moreover, the data are survey-based and there is evidence that attrition due to internal mobility might induce compositional bias (Martí and Ródenas, 2007). On the other hand, changes of residence data register all changes of official residence within a year summing up individual movements (the unit of observation is the change, and not the individual). Furthermore, the data are available only at the annual cell level (defined by year, municipality and demographic characteristics — some demographic characteristics are incomplete, especially for foreigners). Most importantly, these data might be biased due to misreporting for tax purposes (Rubolino, 2020). With respect to the existing dataset, SISCO data contain far richer information, in terms of both timing and geography, and record the universe of movements that also entail a job change. However, they do not record the movements of non-employed persons (within their periods of unemployment or inactivity) or at the time of a transition from employment to non-employment, nor do they record information on the change of residence of the worker. In practice, the drawbacks of the SISCO data are likely to be very small, if the bulk of internal migration is indeed job-related.

All of these differences notwithstanding, we check whether major differences arise between migration patterns detected in the SISCO data and the ones from the ILFS and administrative data on changes of residence. Panel (a) of Figure C.16 shows that the overall extent of migration is remarkably similar between SISCO and residence changes data in terms of levels. The dynamics in the two series are, however, quite different as transfer of residence data report an abnormal drop in 2011, a peak in 2012 and a drop again in 2015. With regard to the 2011-12 drop and peak, this is due to a well-known change in the method of collecting residence data to correct misreporting prior to 2011, which itself led to record transfers in 2012 even though they occurred in earlier years (ISTAT, 2016). The drop in 2015 appears instead an anomaly as the labour market was particularly healthy in that year.

Figure C.16. Comparison between workplace-based, residence-based and ILFS migration

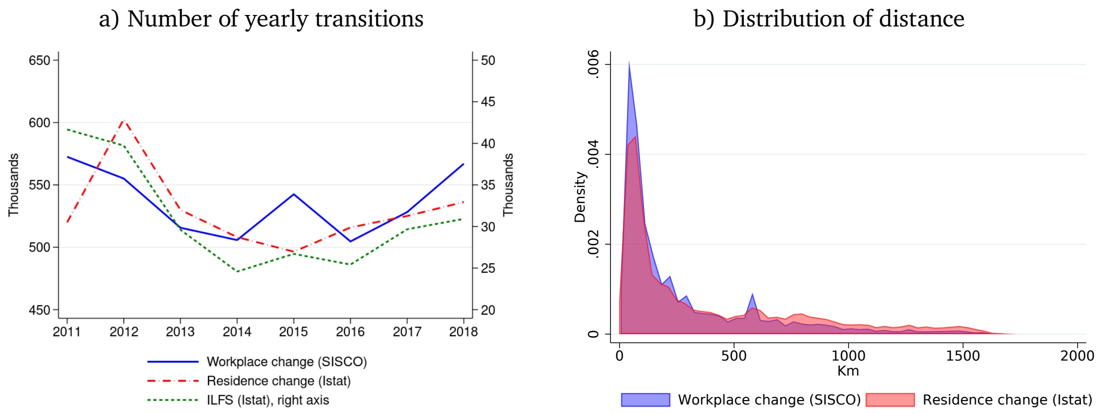

Source: SISCO, residence changes (ISTAT) and ILFS data, 2010-2018. *Note*: The figure shows a comparison between workplace-based (SISCO), residence-based (ISTAT) and ILFS (ISTAT) internal migration. Panel (a) plots the time series of total yearly location switches. Panel (b) plots the distribution of the distance between origin and destination province, pooling data from all the years.

The ILFS data, instead, systematically record only a minor share of all annual moves across provinces, consistently with Martí and Ródenas (2007). In this sense, the ILFS seems to underestimate geographical mobility in a significant way, but tracks the dynamics of SISCO data remarkably well.

Moreover, Figure C.16, panel (b), shows that the distribution of the distance of the migration moves is almost identical between SISCO and residence-based change data, with some minor differences in the lower part of the distribution (residence changes are underrepresented within shorter distance bins). Finally, we also test whether SISCO and residence-based bilateral flows across provinces are correlated, finding a large elasticity of about 0.6-0.7 (Table C.10). We conclude that the two administrative measures of internal mobility line up well across space, though they might have limitations in capturing the same dynamics over time because of the problem with the residence data highlighted above. On the contrary, survey-based mobility underestimates actual movements, but is in line with SISCO in terms of the dynamics over time. Hence, this evidence is very reassuring regarding the use of a workplace-based measure of migration from SISCO data.

We further test whether the main 2SLS results based on SISCO data hold if using changes of residence data (Table C.8). While some of the effects do (e.g., that of JDR on OMR), most do not. We believe that the difference between our baseline result and the robustness check lies in the inability of residence data to correctly capture the dynamics of internal migration across years as pointed out above. Such a source of mis-measurement causes the estimates to be biased towards zero. On the contrary, we already observed that the geographical correlation between the two sources is remarkably high.

Table C.10. Workplace-based vs. residence-based migration flows

	(1)	(2)	(3)
	Log flows	Log flows	Log flows
Log flows (workplace-based)	0.775**	0.642**	0.642**
	(0.002)	(0.012)	(0.012)
Log population of origin			-0.138
			(0.103)
Log population of destination			-0.081
			(0.096)
N	84,088	84,088	84,088
Origin FE	NO	YES	YES
Destination FE	NO	YES	YES
R^2	0.680	0.783	0.768

Source: SISCO and residence changes (ISTAT) data, 2010-2018. *Note*: The table shows the results of log-log regressions between residence-based (ISTAT source) and workplace-based (SISCO source) yearly migration bilateral flows across provinces. Standard errors reported in parentheses are clustered at the origin and destination province level. $^{+}p < .10, ^{*}p < .05, ^{**}p < .01$.

Bibliography

Aaberge, R. and M. Mogstad (2014). Income mobility as an equalizer of permanent income. Technical report, Discussion Papers.

Alves, N. and C. Martins (2012). Mobility and income inequality in the european union and in portugal. *Economic Bulletin and Financial Stability Report Articles*.

Atkinson, A. B. (2016). Inequality: What can be done. *Practice 40*(2), 289–292.

Autor, D. H., D. Dorn, and G. H. Hanson (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review 103*(6), 2121–68.

Barbieri, P. and S. Scherer (2009). Labour market flexibilization and its consequences in italy. *European Sociological Review 25*(6), 677–692.

Bartholomew, D. J. (1968). Stochastic models for social processes. *The Australian and New Zealand Journal of Sociology 4*(2), 171–172.

Bartolucci, C., C. Villosio, and M. Wagner (2018). Who migrates and why? Evidence from Italian administrative data. *Journal of Labor Economics 36*(2), 551–588.

Basso, G., F. D'Amuri, and G. Peri (2019). Immigrants, Labor Market Dynamics and Adjustment to Shocks in the Euro Area. *IMF Economic Review 67*(3), 528–572.

Basso, G. and G. Peri (2020). Internal mobility: The greater responsiveness of foreign-born to economic conditions. *Journal of Economic Perspectives 34*(3), 77–98.

Bertheau, A. and R. Vejlin (2022). Employer-to-employer transitions and time aggregation bias. *Labour Economics 75*, 102–130.

Beyer, R. C. M. and F. Smets (2015). Labour market adjustments and migration in Europe and the United States: how different? *Economic Policy 30*(84), 643–682.

Blanchard, O. and L. Katz (1992). Regional evolutions. *Brookings Papers on Economic Activity 23*(1), 1–76.

Boeri, T. and P. Garibaldi (2007). Two tier reforms of employment protection: a honeymoon effect? *The Economic Journal 117*(521), F357–F385.

Böhlmark, A. and M. J. Lindquist (2006). Life-cycle variations in the association between current and lifetime income: Replication and extension for sweden. *Journal of Labor Economics 24*(4), 879–896.

Bossert, W., S. R. Chakravarty, and C. D'Ambrosio (2012). Poverty and time. *Journal of Economic Inequality 10*(2), 145–162.

Burkhauser, R. V. and K. A. Couch (2009). Intragenerational inequality and intertemporal mobility. In W. Salverda, B. Nolan, and T. M. Smeeding (Eds.), *The Oxford Handbook of Economic Inequality*. Oxford University Press.

Cadena, B. C. and B. K. Kovak (2016). Immigrants equilibrate local labor markets: Evidence from the great recession. *American Economic Journal: Applied Economics 8*(1), 257–90.

Calvo, C. (2018). Vulnerability to poverty: theoretical approaches. In *Handbook of research on economic and social well-being*, pp. 272–283. Edward Elgar Publishing.

Calvo, C. and S. Dercon (2009). Chronic poverty and all that. *Poverty dynamics: interdisciplinary perspectives*, 29–58.

Calvo, C. and S. Dercon (2013). Vulnerability to individual and aggregate poverty. *Social Choice and Welfare 41*(4), 721–740.

Camussi, S., F. Colonna, and F. Modena (2022). Temporary contracts: An analysis of the North-South gap in Italy. Bank of Italy, unpublished manuscript.

Carrington, W. J. (1996). The Alaskan labor market during the pipeline era. *Journal of Political Economy 104*(1), 186–218.

Ceriani, L. (2018). Vulnerability to poverty: empirical findings. In *Handbook of research on economic and social well-being*, pp. 284–299. Edward Elgar Publishing.

Ciani, E., F. David, and G. de Blasio (2019). Local responses to labor demand shocks: A Re-assessment of the case of Italy. *Regional Science and Urban Economics 75*(C), 1–21.

Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives 27*(3), 79–102.

Davis, S. J., R. J. Faberman, and J. Haltiwanger (2006). The flow approach to labor markets: New data sources and micro-macro links. *Journal of Economic perspectives 20*(3), 3–26.

Davis, S. J. and J. Haltiwanger (1992). Gross job creation, gross job destruction, and employment reallocation. *The Quarterly Journal of Economics 107*(3), 819–863.

Davis, S. J., J. C. Haltiwanger, S. Schuh, et al. (1998). Job creation and destruction. *MIT Press Books 1*.

de Chaisemartin, C. and X. D'Haultfoeuille (2022). Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey. unpublished manuscript available on ssrn.

Duclos, J.-Y., J. Esteban, and D. Ray (2004). Polarization: concepts, measurement, estimation. *Econometrica 72*(6), 1737–1772.

Elsby, M. W., B. Hobijn, and A. Şahin (2013). Unemployment dynamics in the oecd. *Review of Economics and Statistics 95*(2), 530–548.

Esteban, J., C. Gradín, and D. Ray (2007). An extension of a measure of polarization, with an application to the income distribution of five OECD countries. *Journal of Economic Inequality 5*(1), 1–19.

Esteban, J. and D. Ray (1994). On the measurement of polarization. *Econometric*, 819–851.

Esteban, J. and D. Ray (1999). Conflict and distribution. *Journal of Economic Theory 87*(2), 379–415.

Fields, G. S. and E. A. Ok (1999). The measurement of income mobility: an introduction to the literature. *Handbook of Income Inequality Measurement*, 557–598.

Foote, A., M. Grosz, and A. Stevens (2019). Locate your nearest exit: Mass layoffs and local labor market response. *ILR Review 72*(1), 101–126.

Foster, J. E. and M. C. Wolfson (2010). Polarization and the decline of the middle class: Canada and the US. *Journal of Economic Inequality 8*(2), 247–273.

Gangl, M. (2005). Income inequality, permanent incomes, and income dynamics: Comparing europe to the united states. *Work and Occupations 32*(2), 140–162.

Gathmann, C., I. Helm, and U. Schönberg (2020). Spillover effects of mass layoffs. *Journal of the European Economic Association 18*(1), 427–468.

Gomes, P. (2015). The importance of frequency in estimating labour market transition rates. *IZA Journal of Labor Economics 4*(1), 1–10.

Gottschalk, P. and E. Spolaore (2002). On the evaluation of economic mobility. *The Review of Economic Studies 69*(1), 191–208.

Guvenen, F., L. Pistaferri, and G. L. Violante (2022). Global trends in income inequality and income dynamics: New insights from grid. *Quantitative Economics 13*(4), 1321–1360.

Haider, S. and G. Solon (2006). Life-cycle variation in the association between current and lifetime earnings. *American Economic Review 96*(4), 1308–1320.

Hart, P. E. (1976). The dynamics of earnings, 1963-1973. *The Economic Journal 86*(343), 551–565.

Hoffmann, E. B., D. Malacrino, and L. Pistaferri (2022). Earnings dynamics and labor market reforms: The italian case. *Quantitative Economics 13*(4), 1637–1667.

Hoffmann, F., D. S. Lee, and T. Lemieux (2020). Growing income inequality in the united states and other advanced economies. *Journal of Economic Perspectives 34*(4), 52–78.

Hoy, M., B. S. Thompson, and B. Zheng (2012). Empirical issues in lifetime poverty measurement. *The Journal of Economic Inequality 10*(2), 163–189.

ISTAT (2016). La revisione post censuaria delle anagrafi: 2012-2014. Technical report.

Jäntti, M. and S. P. Jenkins (2015). Income mobility. In *Handbook of Income Distribution*, Volume 2, pp. 807–935.

Jäntti, M., R. Kanbur, M. Nyyssölä, and J. Pirttilä (2014). Poverty and welfare measurement on the basis of prospect theory. *Review of Income and Wealth 60*(1), 182–205.

Jenkins, S. P. (2011). *Changing fortunes: Income mobility and poverty dynamics in Britain*. OUP Oxford.

Jia, N., R. Molloy, C. Smith, and A. Wozniak (2022). The economics of internal migration: Advances and policy questions. *Journal of Economic Literature , forthcoming*.

Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review 95*(1), 161–182.

Jordà, O., M. Schularick, and A. M. Taylor (2020). The effects of quasi-random monetary experiments. *Journal of Monetary Economics 112*(C), 22–40.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic perspectives 5*(1), 193–206.

Kahneman, D. and A. Tversky (1984). Choices, values, and frames. *American psychologist 39*(4), 341.

Latner, J. P. (2018). Income volatility and mobility: A conceptual exploration of two frameworks. *Research in Social Stratification and Mobility 53*, 50–63.

Manacorda, M. (2004). Can the Scala Mobile explain the fall and rise of earnings inequality in Italy? A semiparametric analysis, 1977–1993. *Journal of Labor Economics 22*(3), 585–613.

Manning, A. and B. Petrongolo (2017). How local are labor markets? evidence from a spatial job search model. *American Economic Review 107*(10), 2877–2907.

Martí, M. and C. Ródenas (2007). Migration estimation based on the labour force survey: An eu-15 perspective. *International Migration Review 41*(1), 101–126.

Molloy, R., C. L. Smith, and A. Wozniak (2011). Internal migration in the united states. *Journal of Economic Perspectives 25*(3), 173–96.

Monras, J. (2018). Economic shocks and internal migration. CEPR Discussion Paper 12977 (updated), CEPR.

Monte, F., S. J. Redding, and E. Rossi-Hansberg (2018). Commuting, migration, and local employment elasticities. *American Economic Review 108*(12), 3855–90.

Mortensen, D. T. and C. A. Pissarides (1994). Job creation and job destruction in the theory of unemployment. *The review of economic studies 61*(3), 397–415.

Naticchioni, P., M. Raitano, and C. Vittori (2016). La meglio gioventù: Earnings gaps across generations and skills in italy. *Economia Politica 33*(2), 233–264.

Nichols, A. (2008). Trends in income inequality, volatility, and mobility risk. Technical report, IRISS at CEPS/INSTEAD.

Nichols, A. (2010). Income inequality, volatility, and mobility risk in china and the us. *China Economic Review 21*, S3–S11.

Nichols, A. and P. Rehm (2014). Income risk in 30 countries. *Review of Income and Wealth 60*, S98–S116.

Notowidigdo, M. (2020). The incidence of local labor demand shocks. *Journal of Labor Economics 38*(3), 68–725.

Nybom, M. and J. Stuhler (2016). Heterogeneous income profiles and lifecycle bias in intergenerational mobility estimation. *Journal of Human Resources 51*(1), 239–268.

OECD (2008). *Growing Unequal? Income Distribution and Poverty in OECD Countries*. OECD Publishing.

OECD (2011). *Divided we Stand. Why Inequality Keeps Rising*. OECD Publishing.

OECD (2018). *A Broken Social Elevator? How to Promote Social Mobility*. OECD Publishing.

Raitano, M. and M. Fana (2019). Labour market deregulation and workers' outcomes at the beginning of the career: Evidence from italy. *Structural Change and Economic Dynamics 51*, 301–310.

Raitano, M., E. Karagiannaki, T. Premrov, L. Geyer, M. Fuchs, F. Bloise, J. Costa-Font, F. Iudicone, and B. De Micheli (2021). Study on intergenerational fairness. final report. *Brussels: European commission*.

Rosolia, A. and R. Torrini (2007). The generation gap: relative earnings of young and old workers in italy. *Bank of Italy Temi di Discussione (working paper) n. 639*.

Rosolia, A. and R. Torrini (2016). The generation gap: a cohort analysis of earnings levels, dispersion and initial labor market conditions in italy, 1974-2014. *Bank of Italy Occasional Paper* (366).

Rubolino, E. (2020). Tax-induced transfer of residence: Evidence from tax decentralization in italy. *Available at SSRN 3710932*.

Salverda, W., B. Nolan, D. Checchi, I. Marx, A. McKnight, I. G. Tóth, and H. van de Werfhorst (2014). *Changing inequalities in rich countries: Analytical and comparative perspectives*. OUP Oxford.

Shorrocks, A. F. (1978). The measurement of mobility. *Econometrica: Journal of the Econometric Society*, 1013–1024.

Shorrocks, A. F. (1984). Inequality decomposition by population subgroups. *Econometrica: Journal of the Econometric Society*, 1369–1385.

Wolfson, M. C. (1994). When inequalities diverge. *American Economic Review 84*(2), 353–358.

Wolfson, M. C. (1997). Divergent inequalities: theory and empirical results. *Review of Income and Wealth 43*(4), 401–421.