BETAM WORKING PAPER SERIES

The hidden panel of the Turkish Household Labour Force Survey: Recovering household links from public microdata, 2004–2025

Luis Pinedo Caro

Working paper

#202604

Bahçeşehir University Centre for Economic and Social Research (BETAM) Osmanpaşa Mektebi Sokak No:4-6 34353 Beşiktaş, İSTANBUL, TÜRKİYE Phone: +902123815962

E-mail:

betam@eas.bau.edu.tr

25 November 2026

The hidden panel of the Turkish Household Labour Force Survey: Recovering household links from public microdata, 2004–2025

Luis Pinedo Caro[1]

Bahçeşehir University

Abstract. The Turkish Household Labour Force Survey (HLFS) is a rotating panel but its public microdata are released as independent cross-sections without the longitudinal identifiers. This article proposes a strategy to recover the rotation from the order of the records in the public files and, with it, to reconstruct household and person identifiers for the 2004–2025 period. Families are linked through a bounded search where the published geography and the structure of the files reduce the possible matches of each household to a handful of candidates. The members of these candidate households are then compared using pre-specified agreement weights and matched one to one. Dates of birth, which are never used by the algorithm, confirm more than 99 per cent of the links, while the survey's own recall module bounds the share of false links at 0.3 per cent. As a by-product, the reconstruction documents a number of changes in survey operations over two decades, including the restart of the panel in 2014 and a new assignment of interviews to months from 2023 onwards. At last, as an application, I use the panel to measure the fall and recovery of working hours after the February 2023 earthquakes in the most affected region, finding results close to those of a dedicated ILO–IOM survey.

Keywords: Labour Force Survey, Data Reconstruction, Short-Term Panel Data, Turkish Statistical Institute (TurkStat).

J.E.L. Classification: C81, C83, J21, Q54

[1] Luis Pinedo Caro is a researcher at the Bahçeşehir University Center for Socio-Economic Research (BETAM). *Email: Luis.pinedo@bau.edu.tr.*

1. Introduction

Short panels, where households are interviewed on a monthly or quarterly basis for a year or two, are a valuable resource for labour economics research. These datasets allow researchers to follow the same individuals through crises, disasters or life-changing events like having a baby and to measure flows and changes in their status. Even though they are less popular than other well known long panel datasets such as the PSID in the United States, the SOEP in Germany or the KLIPS in South Korea, they have proved to be critical to our understanding of labour market dynamics, including worker flows between employment and unemployment. In the United States, for instance, these flows have shaped the debate on whether unemployment rises because more workers lose their jobs or because fewer of the unemployed find one; estimates based on unemployment duration gave conflicting answers (Shimer, 2012; Elsby, Michaels and Solon, 2009), whereas linked data assign a substantial role to job losses (Fujita and Ramey, 2009). Likewise, in the United Kingdom, whose Labour Force Survey follows dwellings over five consecutive quarters, the same flows have been measured using a rotating design close to the Turkish one (Gomes, 2012).

Türkiye has two such short panels, both produced by the Turkish Statistical Institute (TurkStat) and harmonised with their Eurostat counterparts, the Survey of Income and Living Conditions (SILC) and the Household Labour Force Survey (HLFS). The SILC is released both as cross-sections and as a panel, and it has supported research on labour mobility between formal and informal work (Tansel and Kan, 2012) as well as on the dynamics of poverty, ranging from the persistence of early childhood poverty (Dayıoğlu and Demir Şeker, 2016) to multidimensional poverty (Acar, 2014). Still, the SILC has two limitations when it comes to labour-market research. On the one hand, its sample is too small for many groups of interest. On the other hand, being released annually, it is not suitable for observing high-frequency transitions.

In contrast, the panel dimension of the HLFS suffers from none of the mentioned limitations. It has between 80,000 and 130,000 unique households interviewed each year, making it the largest recurrent household survey in the country. Moreover, the HLFS provides quarterly transitions thanks to its 2-(2)-2 design, households are interviewed for 2 consecutive quarters, left outside for 2 more and then interviewed again for an extra 2 quarters. On top of that, the fifteen months that separate the first and the fourth interview are long enough to observe, for example, the effect of a natural disaster, the transition from school to work, the return of mothers to employment after childbirth, movements in and out of informality, the duration of unemployment spells or the response of households to a pandemic.

The downside of the HLFS panel is that its unique IDs are not released in the public microdata. In fact, the HLFS microdata is exclusively provided to the public as an independent annual cross-section. The few studies that had access to TurkStat's internal files (Taşçı and Tansel, 2005; Tunalı, 2008; Tunalı, Ekinci and Yavuzoğlu, 2012) required a special permission and, for everyone else, the panel has remained, in Tunalı's words, "beyond the reach of micro-data users because of TurkStat's data dissemination policy" (Tunalı, 2023). Indeed, it is as if the panel did not exist for users of the

public files, to the point that recent work comparing data sources describes the HLFS as a survey that "does not have a panel structure" (Ayhan, Lehmann and Pelek, 2023).

This article aims precisely at filling this gap by recovering the HLFS panel from the public files alone. The method I propose has two steps. The first one is blocking and is based on reducing the number of available candidates for a match to a point where the link becomes almost trivial in certain cases. In the second step, a comparison between a given family and the remaining candidates for a link is carried out using certain observable, relatively stable characteristics of households and individuals. Ex-ante fixed agreement weights following Fellegi and Sunter (1969) are applied to these variables, and families receive a score indicating how similar they are to each other before performing a one-to-one match, as proposed by Jaro (1989). The result is a succession of short panels made of 3.4 million household interviews from 2004 to 2025 with identifiers for both households and individuals.

There is a precedent in the US that illustrates the importance of following families and not addresses (as this article does by construction). The Current Population Survey (CPS) also produces a short panel with a slightly different rotation pattern, four months in, eight out, and then four in again, and, unlike in TurkStat's panel, key identifiers are provided to the public. However, users need to drop links between different families that were interviewed in the same house in different months (Madrian and Lefgren, 2000), a process that is now standardised thanks to IPUMS (Drew, Flood and Warren, 2014). A similar setting where panel IDs were missing is found in India's Periodic Labour Force Survey, where Enevoldsen (2023) recovers an urban panel after the sampling-unit codes of the first wave were scrambled. When linkage is not possible, as was the case in Türkiye until now, researchers resort to building synthetic panels from repeated cross-sections (Dang et al., 2014), which recover transition probabilities only under certain modelling assumptions.

The work of this article, in particular the way it uses the structure of the microdata to gain information, also connects with a literature focused on showing that answering enough aggregate queries allows a database to be reconstructed almost entirely, see Dinur and Nissim (2003). Garfinkel, Abowd and Martindale (2019) show that this feature may also apply to tables extracted from the census, arguing that combinations of counts and average ages by sex, race and marital status may leave only one possible set of individual records. In fact, the U.S. Census Bureau found that its own 2010 Census tables allowed 97 million records to be reconstructed exactly (Abowd et al., 2023), and decided to adopt differential privacy for the 2020 Census. The rationale behind the above-mentioned articles is that a few attributes may be enough to single out a person (Sweeney, 2000; Narayanan and Shmatikov, 2008; Rocher, Hendrickx and de Montjoye, 2019).

This article also shows the potential limits of such reconstruction efforts. Across the whole HLFS microdata file, households with the same size, sex and ages are common, and a search based on brute force returns many false matches. This applies with particular strength to LFS panels, where the linkage is intertemporal and, thus, researchers should refrain from using most of the information that is subject to change so as not to compromise the results. Once sufficiently small bounds exist, though, a few stable characteristics are enough to create reliable transitions.

Finally, this paper builds on earlier work that also used the HLFS public files to generate information that is not available. Pieces of information such as Syrian refugees (Pinedo Caro, 2020), the month of interview (Pinedo Caro, 2021, 2026a), the provinces in 2004–2013 (Pinedo Caro, 2026b) and the degree of urbanisation in 2025 (Pinedo Caro, forthcoming) all enrich potential analysis. These studies share a common premise, using the structure of the data and the available information in a way it can recover valuable information for a better understanding of Türkiye. The panel presented in this article is probably the most demanding example; recovering a province requires placing a record correctly once, but recovering a family requires finding the same record in two files, months apart, without any common key and potential alternative candidates.

This article provides a twofold contribution. The first one is the identification of the four visits of each family from 2004 to 2025, allowing the HLFS to be used as a short panel from the public files. It should be noted that the links follow families instead of dwellings. In this regard, a household that replaces another one in a sampled dwelling starts a new history. The second contribution is to show that the structure of a microdata file can hold information about the sampling methodology and other variables that are not released in the file. It is, thus, recommendable to pay attention to the data structure in similar exercises. For the HLFS the order in which the data are loaded not only reveals the rotation panel, but also where families should reappear in the file during their next visit. This characteristic bounds the search of each link from some 8,200 households in its region to about 90. Due to this reduction, dates of birth (not used in the algorithm) confirm more than 99 per cent of the links.

The article is organized as follows. Section 2 describes the rotation design of the HLFS, the content of its public files and the linkage algorithm. Section 3 shows how the structure of the public files narrows down the search for each household. Section 4 presents the resulting panel and assesses its reliability. Section 5 uses the panel to measure the impact of the 2023 earthquakes on hours of work in the south eastern region. Finally, Section 6 concludes.

2. The HLFS panel and the linkage problem

2.1 The rotating panel

The HLFS is a rotating panel. Since 2000, each dwelling taking part in the sample is interviewed four times following a 2-(2)-2 pattern, i.e., in two consecutive quarters and then, after two quarters out of the sample, in two more (Taşçı and Tansel, 2005; Tunalı, 2008). Even though the rule is stated in quarters, the schedule is actually set more finely, by month until 2013 and, as the re-interviews suggest, by week since 2014. Until 2013, a new group of dwellings entered the sample every month and was re-interviewed three, twelve and fifteen months later. From 2014 onwards, the re-interviews observed in the files are consistent with fieldwork organised by week, with each dwelling re-interviewed in the same week of the quarter one, four and five quarters after its first visit. Throughout the paper, I group dwellings by the month of their first visit and call each of these groups a rotation group. Every monthly sample thus consists of four rotation groups, at their first, second, third and fourth visit, respectively.

The aim behind the 2-(2)-2 rotation is an increase in the precision of published quarterly and yearly changes. Based on this rule, half of the sample is interviewed again in consecutive quarters and in consecutive years (TurkStat, 2017). Moreover, it should be noted that the survey follows dwellings, not families. If, for example, a family were to leave their house before the third visit and that house was rented to another family, the third visit would be unrelated to the second visit and yet linked in TurkStat's panel. With regards to the sample size the survey has been growing over time, from 120,000–146,000 household interviews per year during the 2004-2013 period, to around 150,000 in the 2014-2019 period and circa 200,000 households after the latest changes attested since 2020.

2.2 What the public files contain

TurkStat provides the HLFS microdata to the public as annual cross-section files. This article uses these microdata sets for the years 2004-2025, including 3.4 million households and 11.7 million individuals. Each household interview has a unique identifier that carries no relationship with other interviews given by the same household in the past, present or future year. The datasets include demographic information regarding all the members in a family except for the 2014-2020 period, where children are excluded and the only available information is the number of children aged 14 or less in the household. It also includes a labour market module that is asked to everyone aged 15 years-old or older and information on the subregion (26 NUTS2) where the interview was carried out.

What is missing from the files matters more than what is in them, though. There is no leftover of the panel structure such as the visit number and the unique family ID. Moreover there is no time variable, i.e. the month or the quarter, that could provide information on when a household is expected to be reinterviewed. A rural/urban indicator is only available from 2004 to 2013, after which it is no longer reported, further complicating the location of the households that are to be linked. Also of relevance for the algorithm, there are exact dates of birth in 2014 for adults and 2021-2025 for all members of the household.

2.3 A linkage algorithm

Linking a household to its next visit requires two interviews to be matched based on characteristics that they share (or not) with each other. The linkage used in this article borrows three elements from Fellegi and Sunter (1969): blocking, a characteristic-by-characteristic comparison of the interviews, and an additive score to which each agreement adds and from which each disagreement subtracts.

At the same time, the method departs from Fellegi and Sunter in that I use fixed weights instead of estimating them. There are two reasons for this. First, estimating weights requires pairs of households whose match status is known. In the HLFS, such pairs can only be identified with dates of birth, which the files record only in 2014 for adults and from 2021 onwards.[2] Second, estimating

[2] Appendix C estimates Fellegi-Sunter weights using dates of birth to create matches that are "certain".

weights without known matches assumes that the variables used for the comparison are independent. This assumption is often false (Winkler, 1995), and it is especially so within households. For instance, a person appearing as the spouse of the head of the household is likely to be married as well. As a result, the same fact (being married to someone) counts more than it should. As a note of caution, fixed weights do not solve this issue by construction; they may count the same fact twice using correlated variables, but in this article the weights given to variables tied to each other are kept small to avoid this double counting.

A second departure concerns the decision rule. Since a household reappears at most once at its next visit, I impose a one-to-one assignment, as in Jaro (1989). This uses knowledge of the design to discard candidates that score well but are not the best, which would otherwise become false links. Jaro solves the assignment globally, by maximising the total score of the accepted pairs. For simplicity, and since with a handful of candidates per household the two approaches rarely differ, I make each household keep its best candidate, and each candidate keep the best of the households that chose it. In addition, I require the best candidate to beat the second best by a margin, leaving near ties unlinked so as not to create a false sense of security.

The comparison between households relies, as much as possible, on characteristics that are not expected to change between visits. Variables include sex, the position within the household, age, the place where the person was born, the year of arrival in the province (and, in 2004–2013, whether the person has always lived there), marital status, relationship with the head of household, and education at the individual level, and the number of children (2014–2020) and the position within the file (expected vs actual) at the household level. Moreover, for households with only one person, and with the intention of dissolving ties, detailed education, literacy, citizenship, time since leaving the last job and the date the current job started are added to the mix. Many more variables exist in the file, such as those related to employment status, occupations, hours or earnings, which would make the matches easier to find, but the purpose of a panel is to observe change and adding these variables would leave many households unlinked, precisely those with changing lives. The only exception are the two job dates of one-person households, which only break ties.

For a household h interviewed in month t, and given a gap $g \in \{3,9,12,15\}$, which refers to months between visits to a given household, I define the box $\mathcal{B}_g(h)$ as the set of households interviewed g months later with potential to be the same household at its next visit. Once the box is defined, h is compared with each candidate h' by adding up the scores of the pairs of persons that the two households have in common.

Person-pair score. A person *p* in household *h* and month *t* is paired with the person who occupies the same roster line in household *h'* at month *t* + *g* provided that both have the same sex and that their age difference is within one year of the ageing implied by the gap (given by ⌊g/12⌋, the number of full years in the gap). If a pair is formed, it receives the following score

$$s_p = a_p - 0.5c_p + 0.1q_p,$$

where a_p = 1 if the age difference equals the implied ageing or exceeds it by one year, and $a_p = 0.5$ if it falls one year short. In turn, c_p counts the conflicts on characteristics that cannot change, namely whether the person was born abroad, the year since which the person lives in the current province and (only in the 2004–2013 period) whether the person has lived in the current province since birth. Then, q_p = 1[same year of arrival to the current province] + 0.5·1[same marital status] + 0.5·1[same relationship to the head of the household] + 0.25·1[educational level not lower at the later visit]. The residence variables are only recorded from 2009 onwards, so that in 2004–2008 there are no conflicts and agreements rest on marital status, relationship to the head and education.

Household-pair score. Two households that are candidates to be linked receive a score defined by

$$S(h, h') = \sum_p s_p - d(h, h'),$$

where $d(h, h')$ penalises differences between the positions of the two households in the file. From 2014 onwards, $d(h, h') = |r - r'|$, where *r* and *r′* are the relative positions of the two households within the records of their subregion in the cycle. In 2004–2013, the penalty becomes $d(h, h') = 5e(1 - F)$, where *e* is the distance to the expected position,[3] and F = 1 if both households have the same size, are of at least three persons, and all their members are paired (0 otherwise).

Acceptance rule. Let $k = \sum_p a_p$ be the sum of the points allocated to the age similarity among paired persons (which equals their number when all ages agree as expected) and let n and n′ be the household sizes. A link between two multi-person households is accepted when three conditions are met: first, $k \geq 2$ and $k \geq \frac{1}{2} \cdot \max(n, n')$, meaning that at least two people are matched and that k is at least half of the larger household size. Second, *h′* is the candidate with the highest score for *h*, and among the households whose best candidate is *h′*, *h* has the highest score. Lastly, whenever there is a second-best candidate, its score must be exceeded by at least 0.3 at each of these two steps, so as to avoid false links.

Chains. Accepted links are then assembled into household histories. Primary links connect consecutive visits (first to second at +3 months, second to third at +9 and third to fourth at +3). When an intermediate visit is missing, a bridge link at +12 or +15 months connects the visits on either side. A link joins two histories only if they share no month and the combined set of months fits the schedule of a single rotation group, ensuring a history contains at most four visits.

Special cases. Two exceptions are made to the linking rule stated above. First of all, in families that entered the sample between 2014 and 2020, and in all households interviewed in July 2021, whose rotation group is unknown, households are compared on adults' characteristics only. This is due to children not appearing in the public files except for their number. As a consequence, the number of children enters the household-pair score as a separate term, which adds or subtracts points

[3] The expected position is defined in Section 3.2 together with the box of plausible matches.

depending on how much it differs between the two candidate households. Denoting by Δ the difference in the number of children between two given visits, the term becomes

$$b(\Delta) = \begin{cases} +0.3 & \Delta = 0 \\ 0 & |\Delta| = 1 \\ -0.5(|\Delta| - 1) & |\Delta| \geq 2 \end{cases}$$

meaning that an equal number of children adds to the score, a difference of one child is neutral, and larger differences subtract from it. In these groups, households with a single adult are treated under the rule for one-person households described below.

Another exception is made for one-person households, which carry less information. A pair of one-person households must show no conflict ($\sum_p c_p$ = 0) and is scored using the following expression,

$$S_1(h, h') = S(h, h') + 0.3v(h, h'),$$

where *v* may be regarded as a tie-breaker with additional agreements on eight characteristics that hardly change over a short period, namely detailed education, literacy, marital status, time since leaving the last job, the date the current job started, living in the province since birth, year of arrival in the province and citizenship. Where it applies, S includes the children term b(Δ). As with multi-person households, each household keeps its best candidate and each candidate keeps the best of the households that chose it, but the margin required at each step is 0.2 instead of 0.3.

3. Box, admissible candidates and linking

The accuracy of the algorithm described in Section 2.3 depends heavily on the size of the box, i.e., the number of plausible candidates for a link. The files publicly released by TurkStat are not particularly suitable to create small boxes. Indeed, only the NUTS-2 subregion[4] could potentially restrict the search due to the fact that the HLFS follows addresses and not families (i.e. links never cross subregions). Other variables that could help narrow it further, such as the province, or the month in which the interview was taken are not released. As a result, the average box of the data "as is" contains around 8,200 households in 2025. In a box of this size, the number of candidates for a link with the same size, sex and age composition is common enough to raise doubts with regards how much links can be trusted.

3.1 Blocks, cycles and months

In order to shrink the size of the box I exploit the order in which the data is loaded onto the public files. I take each annual file as a sequence of blocks, hereinafter defined as runs of consecutive records belonging to the same NUTS-2 subregion. In the period 2004-2013 there are several hundred blocks a year while the figure goes into the thousands since 2014. These blocks follow a recurring pattern which I call a cycle. A cycle is defined as a stretch of the file in which every

[4] And the rural/urban variable only until 2013.

subregion appears, usually more than once. The way to identify cycles uses the omnipresence of a large İstanbul block, which serves as the anchor of a given cycle.

The large İstanbul block appears 12 times a year in 2004–2013 and 96 times in the 2014-2025 period. It is as if it were appearing once a month and eight times a month, respectively. The idea that the cycles are related to the month can further be seen in several variables; monthly earnings rise as the file advances, with jumps in the months in which the minimum wage was raised (for instance, in July 2022), and agricultural employment peaks in the cycles identified as the summer months. Based on these findings, I assume throughout the article that counting cycles from the start of the file gives the month of the interview.

Table 1. Pairs of households whose members all match exactly, by gap between interviews, 2022–2025

Gap (months)	Pairs	Gap (months)	Pairs
1	501	**9**	**106,771**
2	4,628	10	369
3	**275,598**	11	8,447
4	1,919	**12**	**176,584**
5	453	13	342
6	475	14	5,023
7	410	**15**	**73,014**
8	3,608	16	322

Source: TurkStat, HLFS public microdata, 2022–2025, and author's own calculations. ***Notes:*** The table counts the pairs of households of two or more members whose members all match exactly on roster line, sex, and year and month of birth, and which are unique on both sides. The search is done within NUTS-2 subregions. The pairs at gaps of 2, 4, 8, 11 and 14 months correspond to re-interviews that fall in the neighbouring month (Section 3.3). Results for 2004–2021, by age, are shown in Table B1 (Appendix B).

Once the month is known, the calendar followed by the households forming the rotating panel becomes clear and a strategy to shrink the box of plausible matches can be devised. As can be seen in Table 1, at gaps 3, 9, 12 and 15 months (which coincide with the times of panel rotation) it is much more common to find pairs than at any other gap. For instance, in 2022–2025, 275,598 pairs are found at a gap of 3 months (potentially the second or the fourth visit), but only 453 at a gap of 5 months. The table is also informative about the danger of creating wrong links. At every other gap, there are still a few hundred pairs whose members all coincide. Since no re-interview takes place at those gaps, these are different households that happen to look alike, and they give an idea of how many false links a search over the whole subregion would produce.

3.2 Period 2004–2013

From 2004 to 2013, the box of potential matches is formed by the urban or rural records of the same province (or pair of provinces) in the target month.[5] The month is given by the cycle, which is identified by the large İstanbul block that lies between a block of TR62 and one of TR31. Each cycle starts at the first of the two TR62 blocks that precede the İstanbul block and ends just before the next TR62 block, which opens the following cycle. Exceptions to this rule, in 2006 and 2008, are documented in Appendix A.

Provinces are identified as in Pinedo Caro (2026b), who shows that, since the blocks follow the plate order, most of them can be traced to a province. Indeed, this order identifies 69 out of the 81 provinces deterministically. The remaining twelve form six pairs with consecutive plate numbers within the same subregion. Contrary to Pinedo Caro (2026b), I do not try to split them and rather treat them as single units, since the order cannot separate them. In the 2004-2013 period the public datasets include a variable that specifies whether the household lives in a rural or an urban area, a split that is also used to further shrink the size of the box since households are not reclassified.

Rotation groups. In this period the four rotation groups cannot be separated using the blocking strategy, however they appear in order within the box. In most months, the fourth visit comes first and the other visits follow it in descending numerical order. New rotation groups enter the box at the end, and move one segment forward at each visit.[6] This order has exceptions found in 2004 and early 2005, in 2009 and in the first quarter of 2010 and the algorithm adapts to them (see Appendix B for more information). The rotation is intended to be complete at the subregion level, so it is not surprising to see missing rotation groups at small provinces in certain months. This does not affect the algorithm nor the place where the remaining households are expected to reappear, it just means provinces are not suitable for the production of monthly statistics.

Interestingly, within each rotation group, households keep their relative order from one visit to the next. This can be seen in Figure 1, which shows, for the box of urban İstanbul, scatter plots with the relative position within the box at the January 2011 visit and, respectively, the April 2011 and January 2012 visits.[7] True links lie on straight lines, as it would be expected of groups taking the space of previous groups as their visit count increases. For instance, those interviewed as their first visit in January 2011 take the final deciles of the box, whereas at their second visit in April their position has moved as a whole towards the upper-mid deciles. The importance of this relationship can also be seen in the figures, with isolated points representing coincidental matches between

[5] In 2012, each province appears in four or five blocks per month, each of them holding a single rotation group, as after 2014 (Appendix B). Ordering the households within the box reproduces the standard sequence of the four rotation groups, so the 2012 files are linked with the same method as the other years.

[6] See Appendix B, Figure B.2 for the relative position of each rotation group observation in Istanbul.

[7] It should be noted that households interviewed in January and April 2011 may correspond to first and second or third and fourth visits.

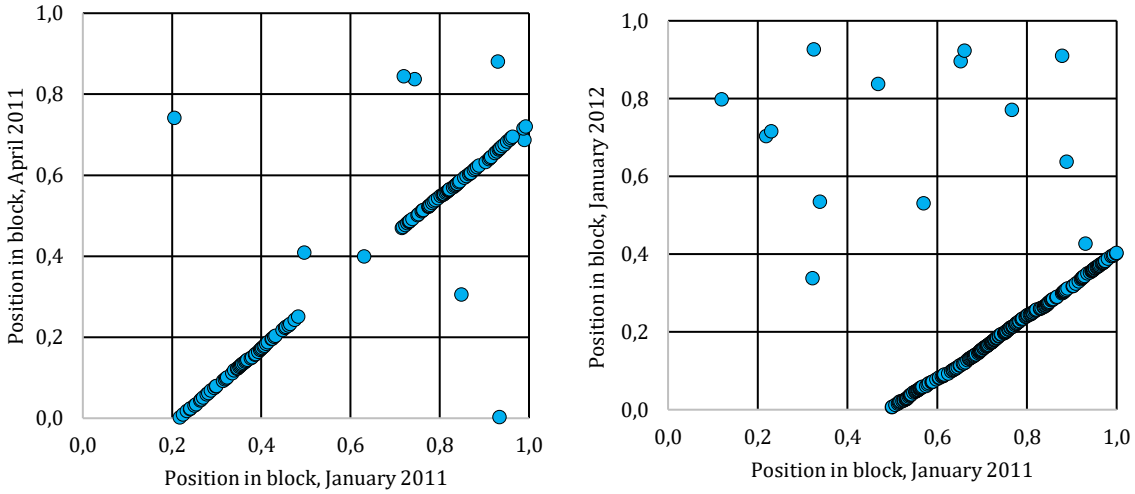
unrelated households with the same size, sex and age composition, which a search ignoring positions would accept.

Untangling the rotation groups. In order to minimize the number of plausible candidates I design a mechanism to draw the bounds of the four rotation groups and the corresponding visit in each box. The rationale is the following. Once I know the place in the box, the number of comparisons decreases, and knowing the visit allows the algorithm to skip certain searches where no candidate should exist (for instance, a second-visit household should never have a match three months later). The process starts with a trial linkage without restrictions, which searches each household at every gap and whose only purpose is to reveal which part of the box holds each visit. It is complemented with a second pass that repeats the linkage knowing the visit and the position of the rotation groups in the box.

The first pass starts with so-called anchors, pairs of households linked with high certainty thanks to their characteristics, which include having three or more persons, the same size in both months and all members paired uniquely on both sides. Their role is to trace how positions in one box map into positions in the other, as shown in Figure 1. In order to discard the few anchors that could be coincidental matches, I keep, for each box and pair of months, the longest subset of anchors whose positions increase together in both months.

Figure 1. Position in the box of households whose members all match exactly at two interviews, urban İstanbul, 2011–2012

(a) January–April 2011 (3 months) *(b) January 2011 – January 2012 (12 months)*



Source: TurkStat, HLFS public microdata, 2011–2012, and author's own calculations. ***Notes:*** The figure shows households of three or more members whose members all coincide in roster line, sex and age (age is equal at 3 months, and equal or one year higher at 12 months), and which are unique on both sides. Position in the box is defined as the rank of the household in file order within the İstanbul urban records of the month, divided by their number. Panel (a) shows first and third visits in January 2011 re-interviewed at the second and fourth visits in April 2011, while panel (b) shows first and second visits in January 2011 re-interviewed at the third and fourth visits in January 2012.

The anchors give the position where each household is expected to reappear, and the distance to that position serves two purposes, to discard the candidates that are very unlikely to be the same household, typically households of another rotation group, and to rank the remaining ones. Formally, let *A* be the set of anchors of a box for a given pair of months, with positions (r_a, r'_a), and let $\hat{f}_A$ be the piecewise-linear interpolation through them.[8] For a candidate pair (*h*, *h'*), the residual is defined as

$$e(h, h') = \left| r'_{h'} - \hat{f}_A(r_h) \right|.$$

Candidates are discarded when this distance is too large. Let *F*(*h*, *h'*) be an indicator that takes value one when both households have the same size, of at least three persons, and all their members are paired. A candidate pair is admissible only if

$$e \leq 0.05 \text{ or } (F = 1 \text{ and } e \leq 0.15),$$

a tolerance that is wider for whole-household matches, since these are unlikely to be coincidental. The same residual is then used to rank the admissible candidates, as it enters the household-pair score of Section 2.3 for 2004–2013 as the position penalty $d(h, h') = 5e(1 - F)$, so that the candidate closest to its expected position is preferred.

The links of the first pass may include some coincidences, but they are only used collectively, to give each household its visit. The visit comes first from the histories, since many households have a visit that follows from the months in which they are observed. For instance, a household observed at t, t + 3 and t + 12 can only be at its first, second and third visits, and a household linked at nine months can only be at its second and third. Doubtful cases are households observed only once, or twice three months apart, since the latter may be at their first and second visits or at their third and fourth. The clear cases show in which order the four rotation groups appear in the box that month, since their average positions rank them. The order is usually the standard one, with the fourth visit first and the first visit last. It changes, however, in 2004–2005 and in 2009–2010, and the size of each segment varies from box to box, so I detect it every month. Each box is then cut into four segments that agree with the clear cases, and the doubtful cases take the visit of their segment.

The detected order is only used when it is consistent across boxes, that is, when the median Kendall rank correlation is at least 0.6 over the boxes with eight or more clear cases. When it is not, or when there are too few clear cases, as at the very start and end of the period, the visit is taken from the household's own links only when they leave a single possible value, and otherwise it remains unknown.

Once the visit of each segment is known, I repeat the whole search for all households. Each household is now searched only at the gaps that its visit implies and in the segment of the target

[8] If *h* is itself an anchor and $|A| \geq 3$, $\hat{f}_A$ is computed from *A* after excluding *h*. Beyond the first and last anchors, $\hat{f}_A$ is constant, and boxes with fewer than two anchors for a pair of months produce no candidates.

box where its rotation group should reappear, so candidates of other rotation groups no longer compete with the true match. The links of this second pass replace those of the first one, and the histories and visits are built again from them. Candidates whose visit is unknown are kept in the search. Households whose visit is still unknown are searched in the whole box at 3 and 12 months, and at 9 and 15 months only when $F(h, h') = 1$.

3.3 Period 2014–2025

From 2014 onwards the box where candidates are searched for is formed by the households of the same subregion and the same half rotation group, interviewed in the month or months where the next interview may fall. The subregion comes with the public files, and it is readily available. The half groups are found with the structure of the data and the way its blocks are organized. It turns out that from 2014 a large İstanbul block appears 96 times a year, eight times as often as in 2004–2013. Divided by 12, this hints that each month has eight cycles, and counting cycles from the start of the file gives the month.[9] These eight cycles turn out to be half of a rotation group a piece. The specific visit that each subgroup holds is learnt empirically from the data, following exactly matched pairs from one interview to the next. In most months, cycles 6 and 8 hold the rotation groups at their first visit, cycles 5 and 7 those at their second, cycles 2 and 4 those at their third, and cycles 1 and 3 those at their fourth. Table 2 shows the resulting rule together with the share of these pairs that follow it, which is between 99.6 and 99.9 per cent at every gap.[10]

Table 2. Cycle within the month at the later interview given the cycle at the earlier one, 2022–2025

Gap	Rule	Pairs (Table 1)	Following the rule (%)
+3	2 → 1, 4 → 3, 6 → 5, 8 → 7	275,598	99.67
+9	5 → 2, 7 → 4	106,771	99.76
+12	5 → 1, 6 → 2, 7 → 3, 8 → 4	176,584	99.93
+15	6 → 1, 8 → 3	73,014	99.68

The position of the household within the box also helps determine the veracity of the link, although its role is smaller than in 2004–2013. Since each cycle holds a single half rotation group, households keep the same relative order at every visit. The expected position at the next visit is therefore the same relative position, and anchors are not needed. The household-pair score of Section 2.3 simply penalises the distance to it, $|r - r'|$.

In certain months the box may also need to cover the neighbouring month. This has to do with an inconsistency between how the data is ordered (monthly) and the precision of the re-interview

[9] Usually true with exceptions detected with information like wages that is expected to monotonically increase over time.

[10] The rule does not always hold. In 2020 and the first half of 2021, part of the households that enter in cycle 8 reappear in cycle 8 again at their second visit, instead of cycle 7. Since the logic behind this deviation is not clear from the files, the box does not follow it, and these households are left unlinked. Other exceptions are documented in Appendix A.

calendar, which has weekly precision. For instance, a household hypothetically interviewed for the first time in the week going from May 29 to June 4 of 2023 is assigned to June, since its Thursday falls on June 1. However, its re-interview in the next quarter (13 weeks later) taking place at the August 28-September 3 week, is assigned to August even though September would have been expected if the jump were to be measured in months. This inconsistency does not affect TurkStat's work, but affects the formation of links due to the lack of information on the exact date when the re-interview happens (I operate with monthly precision but the calendar requires weekly).

A question that arises next is about which weekly interviews are moved to a neighbouring month. Quarters have 13 weeks, so one month must be left with 5 weeks and two with 4 weeks. The effect of this distribution can be seen in the data, as the months with an extra week have about 25 per cent more households. The exact months with 5 weeks have changed over time, though, and three forms are observed (see Table 3). Until 2022 the long month that holds the 5th week keeps the same place in every quarter, with an exception at the end of 2020, a year with 53 weeks. From 2023, a week belongs to the month that contains its Thursday, so the long months change from year to year. As a result, a re-interview is moved to the neighbouring month only when the long month is not in the same place in the quarter of origin and in the quarter of destination.

Table 3. Assignment of interview weeks to months implied by the re-interviews

Period	Weeks per month within each quarter	Months with five weeks
2014–2019	4-5-4	Feb, May, Aug, Nov
2020	4-5-4 in Q1–Q3; 5-4-5 in Q4 (53 weeks)	Feb, May, Aug, Oct, Dec
2021–2022	4-4-5	Mar, Jun, Sep, Dec
2023–2025	month containing the week's Thursday	2023: Mar, Jun, Aug, Nov; 2024: Feb, May, Aug, Oct; 2025: Jan, May, Jul, Oct

Since the calendar predicts when a candidate may fall in the neighbouring month, the algorithm dynamically adapts to this, allowing the box to widen when needed.[11] In practice, when the calendar is expected to move, the algorithm allows households interviewed at month t to be searched for in t + g and t + g ± 1. The best candidate is chosen over both target months, but that means the box grows and with it the number of coincidences. In an attempt to help the record linkage algorithm I use the fact that shifted interviews come in clusters, assuming that each interview week covers a group of neighbouring dwellings. Given a candidate link, let $N(h)$ be the nearest households in the box, up to three on each side in file order, that also have a link at the same nominal gap. A link at the neighbouring month is kept if more than 20 per cent of $N(h)$ are also linked at the neighbouring month. The rationale for imposing this rule is simple. Isolated links found in the neighbouring month

[11] How well the calendar reflects the rotating structure is shown in Appendix B.4.

are almost always coincidences, with only 9 per cent of them being confirmed by dates of birth, against 94 to 98 per cent of those within clusters. This rule excludes 4,300 of the 53,000 candidates for a link, and, thus, about 390 true links (9 per cent).

A word regarding an algorithm limitation. The cycles in the data are ordered in two ways, a so-called plate order and a subregion order (see Table A1 for more details) and this affects the efficacy of the linking algorithm. When an order changes abruptly, the relative position of a household in the months of origin and destination changes, increasing the penalty provided by the algorithm and slightly reducing its chances of being matched. Since usually the number of candidates is small (Table 4), its impact is limited.

3.4 How much the box narrows the search

The structure of the files narrows the search a great deal. In 2025, the naive box holds about 8,200 households, whereas $\mathcal{B}_g(h)$ holds about 90. Table 4 shows what this means when household members are compared. It takes the households interviewed in 2024 and due back twelve months later, when every member is one year older, and counts the households of 2025 with exactly the same composition by sex and age, first in the naive box and then in $\mathcal{B}_g(h)$.

Table 4. Households of 2025 with exactly the same composition by sex and age as a household of 2024, links at 12 months

		Naive box		Box $\mathcal{B}_g(h)$		
Household size	# of origin households	False twins per HH	With a false twin (%)	False twins per HH	With a false twin (%)	Only outside the box (%)
1	15,134	14.15	99.4	0.17	14.7	19.9
2	24,602	3.43	82.6	0.08	6.6	14.0
3	15,944	0.43	39.5	0.00	0.3	7.3
4 or more	24,487	0.23	22.8	0.00	0.1	4.3
All	80,167	3.88	58.9	0.06	4.9	10.8
Mean number of households in each group		*8,198*			*87*	

Notes: The table refers to households interviewed in 2024 whose next visit, twelve months later, falls in 2025 and that have at least one household with the same composition in their NUTS-2 region in 2025. Composition is defined by the sex and age of every member, with ages advanced by one year. A false twin is a household of 2025 with the same composition that is not the household itself, the latter being identified by the panel at 12 months (or in the neighbouring month). Mean number of households in each group refers to all households interviewed in the NUTS-2 subregion in 2025 (naive box), or in $\mathcal{B}_g(h)$, per origin household. Only outside the box refers to the origin households with no household of the same composition in their box; the neighbouring months of the implied calendar are not included in this table.

In the naive box, more than half of the households find another household with the same sex and age composition. These so-called twins are particularly common among single-person households, with 99.4 per cent of them encountering a doppelganger. The risk of using the naive box becomes apparent; an algorithm would have to either accept many coincidental matches or keep adding characteristics until ties are dissolved (and, thus, kill the very changes a panel is expected to

measure). In the box used in this article only 4.9 per cent of households (as opposed to 58.9 per cent in the naive one) still have a twin, mostly people living on their own. Incidentally, 11 per cent of households have no household of their composition in their box (last column of Table 4), mostly because they did not come back or their composition changed, and the naive box would link many of them to the wrong household.

4. The panel and its reliability

4.1 The panel

The 3.4 million household interviews of the 2004-2025 HLFS public files make up around 1.27 million household histories[12] (unique families), each having 2.7 interviews on average. It must be noted that the linkage mechanism does not return the original panel of TurkStat (which is address-based) but a panel that follows the same families over time. In practice, this means that transitions are never computed between members of different families. The resulting panel is summarized in Table 5, where results are split into the two main periods, 2004-2013 and 2014-2025.

Table 5. The reconstructed panel

	2004–2013	2014–2025	Total
Household interviews	1,351,243	2,084,740	3,435,983
Household histories (unique families)	531,453	736,039	1,267,492
1 interview (% of histories)	31.3	22.0	25.9
2 interviews (% of histories)	20.8	20.3	20.5
3 interviews (% of histories)	10.3	10.3	10.3
4 interviews (% of histories)	37.6	47.5	43.3
Followed 3 months apart (% of families)	67.4	76.6	72.8
Followed 12 months apart (% of families)	49.8	58.7	55.0
Followed 15 months apart (% of families)	39.4	49.0	45.0

Source: TurkStat HLFS microdata 2004-2025 and author's own calculations. Notes: Followed 3, 12 and 15 months apart refers to families observed at two visits with that gap (first and second or third and fourth visits; first and third or second and fourth; first and fourth), among the families whose visit is known. The visit of the rotation group is known for 98.0 per cent of the interviews of 2004–2013 and 99.9 per cent of those of 2014–2025, where it follows from the cycle except in July 2021, when it is imputed from the household's links.

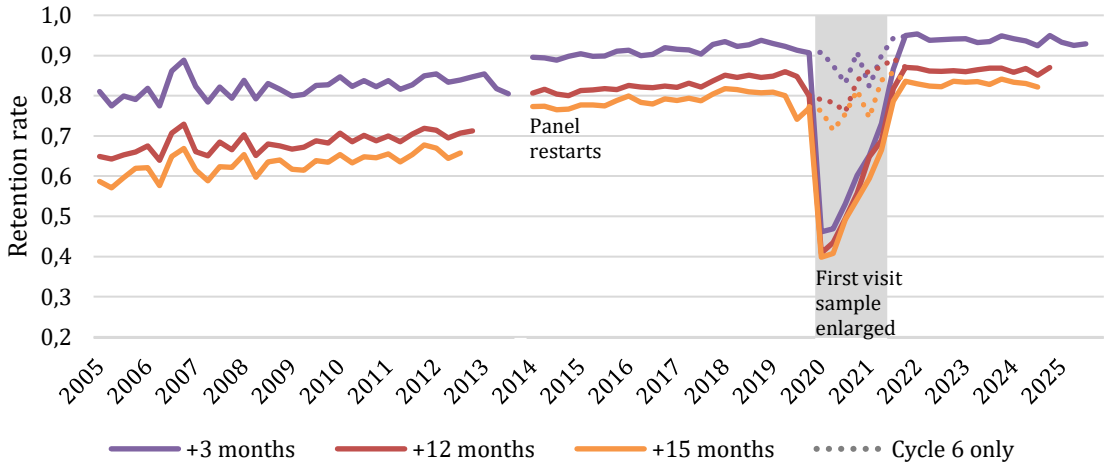
Based on the reconstructed linkages nearly half of the families in the 2014-2025 period and 39.4 per cent in 2004-2013 can be followed for 15 months. For comparisons requiring interviews to be a year apart, the shares go up to 58.7 and 49.8 per cent, respectively. It also strikes that about one family in four is observed only once, however, this should not be read as attrition as part of it comes from

[12] A household history is the set of interviews joined by links. It should be noted that the number of histories is close to, but not identical with, the number of distinct households in the sample. A household whose link to its next visit is not found starts a second history and, as a result, the number of distinct households is somewhat lower. In turn, a new household moving into a dwelling already in the sample starts its own history, as it should.

the way the panel is built; households towards the start or the end of each period cannot possibly complete all four visits, [13] families moving to another address create a break in the history and the creation of a new one with the moving in family, and, of course, some links may not have been detected. The actual loss of families from one visit to the next is shown in Figure 2. Interestingly, one-person households make up 17 per cent of the families observed once, but only 9 per cent of those observed four times. A plausible explanation lies in the fact that they move more often but also the algorithm may link them using less information, thus skipping some.

The actual attrition rate of the panel is shown in Figure 2. It should be noticed that the graph follows the households entering the sample from 2005 onwards, once the panel has run for a full year, and that each series stops when the later visit would fall outside its period, so retention at 12 and 15 months is not available for the households entering towards the end of 2013 and 2025. It is also apparent that there is a break in the panel, TurkStat restarted it altogether alongside certain major changes that were applied to the HLFS in 2014. Retention is lower in the first period, with 79-85 per cent retention at 3 months, 65 to 71 per cent after a year and 59 to 66 after 15 months. Rates are noticeably higher in the second period by about 10 percentage points on average, a feature associated with the method in which the rotation groups are found in each period. Attrition is also seen to decrease over time in both periods, reaching in the latest available period retention rates close to 95 per cent at 3 months.

Figure 2. Household retention by quarterly entry cohort, 2005–2025



Source: *TurkStat, HLFS public microdata, 2004–2025, and author's own calculations.* ***Notes:*** *The figure shows the share of households of each quarterly entry cohort observed again 3, 12 and 15 months after their first visit. The series break in 2014, when the panel restarts. Dotted lines refer to households of cycle 6, the first-visit cycle that keeps its usual size while the sample is enlarged.*

[13] Histories are truncated at the edges of each period. Households interviewed in the first fifteen months of 2004 or 2014 may have had earlier visits, and those first interviewed after September 2012 or September 2024 miss their last ones. The two periods do not overlap, since no household interviewed in 2013 is found in the 2014 files.

There is a feature in the panel that affects attrition and is worth noting; retention falls abruptly among households entering the sample in 2020, with only 46 to 61 per cent of them observed three months later. This fall seems to be related to an enlargement of the first visit sample and a limit in the way the algorithm is constructed. Indeed, TurkStat increased the size of the sample at that time, and newcomers were added to one of the half rotation groups assigned for first visits (cycle 8). These newcomers do not move downstream to cycle 7 after 3 months, but some seem to remain in cycle 8. Since the rationale behind this change is unclear the algorithm was not modified to expand the box to search at a different place. In Figure 2 retention rates for cycle 6 (the other half for first visits) are shown as a benchmark for the period, since its dynamics were not altered.

4.2 Reliability

Dates of birth. Since testing the accuracy of the family histories against the original panel is not possible, I kept dates of birth away from the algorithm to use them as an independent check. These variables are only available for individuals aged 15 or over in 2014, and for everyone from 2021 to 2025. I define a link to be correct when at least half of the persons paired across two visits report the same year and month of birth, or the same year when the month is unknown at one of the visits.

Table 6. Share of links confirmed by dates of birth

Link type	2014: links	2014: % correct	2021–2025: links	2021–2025: % correct
Multi-person households	50,224	99.1	532,258	99.7
One-person households	5,836	97.6	89,715	99.4
Bridges across a missing visit	—	—	9,867	91.6
Links to the neighbouring month (calendar of Table 3)	—	—	28,861	97.8

Notes: 2014 links are those between two visits within 2014, and dates of birth in 2014 are recorded for adults only. Bridges cannot occur within a single year. Links to the neighbouring month overlap with the other rows, and there are none within 2014. They are pairs of interviews of the same household at any of the scheduled gaps, so one shifted interview can enter more than one pair.

Table 6 reports the share of links (at family level) that appear to be truthful, with figures close to certainty regarding links built with no break in series (those in the first two rows of the table) to 91.6 per cent in the few cases where the family was not linked in a given visit and rescued at a later one. The precision of these so-called bridges is highest for the 12-month ones, which make up most of them (92.9 per cent), and lowest for the rare 15-month bridges that skip two middle visits (68 per cent, 242 links). In this sense, bridges can be thought of as being mostly genuine and a feature worth keeping in the algorithm to save atypical transitions and avoid bias towards stable transitions.

The HLFS recall module. Dates of birth are restricted to 2014 (adults only) and 2021–2025, leaving most survey years unassessed. In order to make up for it, I use what usually would have been an alternative to the panel itself, a recall module asking individuals about their situation a year earlier. This recall module took place in the HLFS from 2004 to 2020, asking every respondent aged 15 or over about their labour-market status and, in certain years, about the characteristics of their employment. For the consistency check I use households linked at 12 months from the panel and

compare their recalled status with the status they themselves provided at the earlier visit. If the link is correct, both refer to the same person and month and should only differ due to recall error. The comparison is done on an individual basis, distinguishing between employed, unemployed and inactive, for 2.07 million person-year comparisons. Agreement is benchmarked against a ceiling and a floor. Said ceiling is the agreement among households of three or more members whose members all match exactly, reweighted to the age, sex and status composition of the persons compared. The floor is the agreement with a lookalike of the same sex and age from another household of the same month and geographic unit. The implied share of false links is then given by (ceiling − panel) / (ceiling − floor).

Table 7. Agreement between observed and recalled labour-market status (%), links at 12 months

Period	Links	Households	Ceiling	Panel	Floor	Implied false links (95% CI)
2004–2013	All	453,879	84.4	84.5	70.1	−0.26 (−0.67; 0.25)
	Full chains	423,210	84.3	84.3	69.8	−0.11 (−0.68; 0.36)
	One-person households	27,584	91.3	92.8	82.9	—
	Bridges	3,085	81.0†	82.0	70.4	—
2014–2020	All	364,308	85.1	85.2	64.6	−0.50 (−0.84; −0.05)
	Full chains	323,546	84.9	84.9	64.0	−0.27 (−0.67; 0.08)
	Households with children	159,771	84.6	84.7	64.6	−0.33 (−0.97; 0.23)
	One-person households	36,820	91.1	93.1	78.3	—
	Bridges	3,942	80.5†	81.6	63.9	—
	Neighbouring month (2019 → 2020)	1,306	84.7	83.2	63.5	6.9 (1.0; 12.7)

Notes: Implied false links are expressed in per cent and computed as (ceiling − panel) / (ceiling − floor), without truncating at zero; a negative value means that the panel agrees with recall slightly more than the ceiling. Confidence intervals are obtained from 200 bootstrap replications clustered by household. Households with children in 2014–2020 are linked on adults only. † Ceiling computed from bridges whose members all match exactly. For one-person households, the ceiling, which is drawn from households of three or more, understates the stability of persons living alone and, hence, no false-link rate is computed.

Agreement in linked households equals the ceiling to within 0.1 points in both periods (Table 7), and the implied share of false links is slightly negative, −0.3 per cent in 2004–2013 and −0.5 per cent in 2014–2020, with 95 per cent upper bounds of 0.3 and −0.05 per cent. A negative value means that linked households agree with their recall slightly more often than the exact matches of the ceiling, possibly because the reweighting of the ceiling does not fully reproduce the composition of the panel. In either case, the estimates leave no room for false links beyond a few tenths of a per cent. Bridges agree with the recall less often than other links, but so do bridges whose members all match exactly (81.0 and 80.5 per cent). A plausible explanation could go along the lines of households that miss an interview changing status more often. Links to the neighbouring month can only be tested for households followed from 2019 into 2020. Their implied share of false links is 6.9 per cent, with a wide interval (1.0 to 12.7), which is consistent with the 2 to 3 per cent measured with dates of

birth. Recall itself is, though, imperfect. Even in the ceiling, about 15 per cent of persons recall a status different from the one they reported a year earlier, and only 41 to 43 per cent of those unemployed at the earlier visit recall being unemployed.

5. An application: the 2023 earthquakes

In the two decades covered by the panel, Türkiye went through the Van earthquake of 2011, the Elazığ and İzmir earthquakes of 2020, the pandemic, the wildfires and floods of 2021 and the earthquakes of February 2023. The damages caused by these events were quickly assessed, but information on job losses usually came incomplete and late. Indeed, in the context of natural disasters measuring the impact on the labour market gets tricky. The Social Security Institution does not register employment losses of informal workers, dismissals of formal workers may be banned for a while, and, thus, not reflected in official statistics, and to make things worse, data from labour force surveys reports on the current stock of workers, eliminating the possibility of getting to know whether recoveries imply people getting back their jobs or new people entering the employment.

These difficulties were encountered in the aftermath of the 2023 earthquakes, alongside an HLFS recall module that could have been useful (it used to ask respondents about their situation a year earlier) but was discontinued in 2020. In contrast, the existence of a rotating panel where the same families provide information about their labour force status evolution solves many of the above-mentioned concerns; it covers informal employment, hours of work, and the possibility of assessing whether the affected families are also part of the recovery or whether there exist scarring effects after the disaster.

5.1 Impact measurement

The earthquakes of February 2023 struck eleven provinces in southern and eastern Türkiye but most of the damage was attested in four of them, Adıyaman, Malatya, Kahramanmaraş and Hatay. Unfortunately the HLFS only provides information at the NUTS-2 level, which packs these provinces with other, not-so affected ones. In order to make the impact as neat as possible the analysis focuses on TR63, the subregion that includes Hatay (arguably the province with the hardest hit districts), Kahramanmaraş (the province where the earthquake epicentres were located) and Osmaniye, with less damage but making up for only 15 per cent of the subregion population. Other affected provinces such as Adıyaman or Malatya are part of, respectively, TRC1 (with Gaziantep and Kilis) and TRB1 (with Elazığ, Bingöl and Tunceli), thus sharing the subregion with large populations where the earthquakes had a less severe impact.

In the rotating panel of the HLFS I choose households whose first visit took place in December 2022 and January 2023 so that their second visit takes place after the earthquake in March/April 2023 and their fourth visit 15 months later, in March/April 2024. These two scenarios allow me to estimate the short-term or immediate impact of the earthquake as well as the midterm impact after the emergency has settled down. Employment losses are measured in terms of hours of work lost taking first visit figures as the benchmark for both formal and informal workers.

5.2 Results

Table 8. Weekly hours of work of persons employed before the earthquakes (before = 100)

	Panel, TR63	Panel, rest of Türkiye	ILO–IOM, Hatay and Kahramanmaraş
All			
Right after (+3 months)	71.7	93.3	69.5
One year after (+15 months)	77.6	85.4	78.4
Formal before			
Right after (+3 months)	76.8	95.1	74.1
One year after (+15 months)	86.0	87.8	80.9
Informal before			
Right after (+3 months)	62.9	87.2	54.9
One year after (+15 months)	66.2	76.9	70.7

Notes: The table refers to persons aged 15 or over employed in the week before the earthquakes (ILO–IOM, by recall) or at their first visit in December 2022 or January 2023 (panel), who are followed to their second and fourth visits. Formal refers to jobs registered with social security. Figures are weighted using panel weights of the first visit. The sample consists of 318 persons in the panel (193 formal, 125 informal) and 1,634 in the ILO–IOM survey (1,183 and 451). The 95 per cent confidence intervals of the panel, obtained from bootstrap replications clustered by household, are about ±8 points for all persons and for formal workers, and ±11 to ±14 points for informal workers. Rest of Türkiye excludes TR62, TR63, TRB1, TRC1 and TRC2.

The column labelled Rest of Türkiye in Table 8 is used as a naive counterfactual. Since the group assessed only contains people employed before the earthquakes, their employment share may only fall. Ordinary turnover in the rest of Türkiye brings the index down to 93.3 after three months and to 85.4 after fifteen months. The immediate impact of the earthquakes in TR63 is estimated at 23.2 per cent of hours of work lost (result obtained as one minus 71.7 divided by 93.3 from Table 8). The loss is concentrated among workers without social security, with hours of work lost estimated at 27.9 per cent versus the 19.2 per cent found among formal workers. One year later, the impact after one year is estimated at 9.1 per cent overall, with formal workers having almost recovered the pre-earthquake level, with a minimal 2.1 per cent loss and informal workers still bearing a 13.9 per cent loss in hours worked. Total hours of all persons aged 15 or over, including those who started working after the earthquakes and largely replaced the jobs lost, stand at 94.7 per cent of their level before the earthquakes, against the 99.6 held by the rest of Türkiye.

Panel samples are small and intervals are wide. In addition, the panel does not follow those who left the region; no household interviewed in the three subregions before the earthquakes reappears in another subregion in the 2023 and 2024 files. Attrition between the first and second visits is 13 per cent in TR63, TRB1 and TRC1, against 5 per cent in the rest of Türkiye, so the panel is more likely to understate the impact than to overstate it.

5.3 A comparison with a dedicated survey

A year after the earthquakes, in February and March 2024, the ILO and IOM interviewed 3,749 households (13,554 persons) in Adıyaman, Hatay, Kahramanmaraş and Malatya (ILO & IOM, 2024). The ILO–IOM survey measures each person at three moments: before the earthquakes, right after them (both by recall) and at the time of the survey (Pinedo Caro, 2026c). So as to keep the comparison with the HLFS as close as possible I combine the data of Hatay and Kahramanmaraş to put it against that of TR63. Malatya and Adıyaman, on the other hand, are a small part of their subregions and have no reliable counterpart in the HLFS, so they are left out.[14] A few adjustments are made to the ILO–IOM survey figures so that they are comparable with those of the HLFS. Employment before the earthquakes is recalled, but hours are not: I use the current hours of those still employed, or I impute them for those who lost their job. Moreover, hours worked right after the earthquakes are derived from the job-loss question, which distinguishes job loss, fewer hours (assumed to account for a 50 per cent loss) and no change.

Both sources describe the same population of persons who lived in the region before the earthquakes and still lived there a year later, so it should not come as a surprise that, as can be seen in Table 8, they give very similar results. Right after the earthquakes, the hours of persons previously employed fall to 69.5 in the ILO–IOM survey, against 71.7 in the panel, and one year later the index stands at 78.4 and 77.6, respectively. In both sources, workers without social security are the ones who lose most. A year later, total hours stand at 99.4 per cent of their earlier level in the survey, against 94.7 in the panel. In each of the six comparisons, the estimate of the ILO–IOM survey lies within the confidence interval of the panel.

Admittedly, a dedicated survey goes further than the panel can. It was designed to understand how the earthquakes changed people's lives and, thus, it records what happened to each household's home and where it lives a year later, which allows employment losses to be related to building damage (Pinedo Caro, 2026c). None of this is available in the HLFS. Still, such surveys are not always available, or are carried out once the emergency has settled, relying on recall. Moreover, even when a dedicated survey exists, the panel can provide researchers with a baseline.

6. Conclusions

Even though the Turkish Household Labour Force Survey is a rotating panel, the public microdata files are released as independent cross-sections with no panel identifier. This article proposes the use of a bounded search to recover the panel identifiers from the files as they are released by TurkStat (the public files). I use geographic indicators as well as the order of the records in the files to reduce the number of candidates with which each household could potentially be matched. The choice of the winning candidate is carried out with a comparison of rather stable individual and

[14] Using the same measure as in Table 8, the hours of persons employed before the earthquakes fall, according to the ILO–IOM survey, to 79.5 in Malatya and 63.4 in Adıyaman right after the earthquakes, and stand at 83.5 and 80.2 one year after. In the panel, the same figures fall to 87.0 in TRB1 and 86.5 in TRC1, and stand at 85.3 and 78.2 one year after.

household-level variables, avoiding as much as possible the inclusion of labour market variables. As a result, it is found that the 3.4 million household interviews that were carried out between 2004 and 2025 represent, roughly speaking, 1.3 million families (household histories), i.e., each history comes with 2.7 interviews on average. A reliability assessment yields 99 per cent of correct links when dates of birth (which are not used in the algorithm) are used to check the links.

The identification of the panel identifiers has also shed light on how the HLFS is conducted, revealing a total renewal of the sample in 2014, and a sizeable increase in the number of first visits in 2020 – a transition period with an aim at monthly estimates. The public files show in 2012 an organization similar to the one of 2014, with rotation groups appearing separately. Last but not least, the transition from the rural/urban classification to the degree of urbanisation (DEGURBA) seems to have been carried out between January 2024 and October 2025, one half rotation group at a time. Of course these are interpretations that can only be confirmed by TurkStat itself; in any case, the above mentioned regularities were used to shrink the box of plausible matches.

As a proof of concept, I use the resulting panel to estimate the impact of the 2023 earthquakes that struck south-eastern Türkiye on working hours. Results are computed for the TR63 subregion that comprises Hatay, Kahramanmaraş and Osmaniye, and it is found that hours worked fell by 28 per cent, a real drop of 23.2 per cent once the rest of Türkiye is taken as reference. A comparison with an ad-hoc ILO–IOM survey carried out in the area in early 2024 yields similar results with a 31 per cent drop just in Hatay and Kahramanmaraş. Both sources also agree in informal workers having borne a larger drop in hours worked, and, in addition, also point at a near recovery in hours of work a year after the disaster.

The identification of the panel IDs is part of an effort to find variables and dimensions that the public files of the HLFS microdata do not include. Past attempts include the month of interview (Pinedo Caro, 2021, 2026a), the provinces in 2004–2013 (Pinedo Caro, 2026b), the degree of urbanisation (Pinedo Caro, forthcoming) and Syrian refugees (Pinedo Caro, 2020). These dimensions allow researchers to work on a number of relevant topics that would not have been possible otherwise and the panel IDs now join this literature.

Finally, I would like to express that none of the work carried out in this article would have been possible without the work of TurkStat. The fact that surveys have been consistently carried out for two decades, passing through several natural disasters and a pandemic is what makes links recoverable and what makes the panel useful. I also hope this work makes the case for releasing the panel officially, in line with what other agencies do.

References

Abowd, J. M., Adams, T., Ashmead, R., Darais, D., Dey, S., Garfinkel, S. L., Goldschlag, N., Hawes, M. B., Kifer, D., Leclerc, P., Lew, E., Moore, S., Rodríguez, R. A., Tadros, R. N., & Vilhuber, L. (2023). *A simulated reconstruction and reidentification attack on the 2010 U.S. Census: Full technical report*. CES Working Paper 23-63, U.S. Census Bureau.

Acar, A. (2014). *The dynamics of multidimensional poverty in Turkey*. Betam Working Paper No. 014, Bahçeşehir University.

Ayhan, S. H., Lehmann, H., & Pelek, S. (2023). *Job creation and job destruction in Turkey: 2006–2021*. IZA Discussion Paper No. 16491.

Dang, H.-A., Lanjouw, P., Luoto, J., & McKenzie, D. (2014). Using repeated cross-sections to explore movements into and out of poverty. *Journal of Development Economics*, 107, 112–128.

Dayıoğlu, M., & Demir Şeker, S. (2016). Social policy and the dynamics of early childhood poverty in Turkey. *Journal of Human Development and Capabilities*, 17(4), 540–557.

Dinur, I., & Nissim, K. (2003). Revealing information while preserving privacy. In *Proceedings of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (pp. 202–210).

Drew, J. A. R., Flood, S., & Warren, J. R. (2014). Making full use of the longitudinal design of the Current Population Survey: Methods for linking records across 16 months. *Journal of Economic and Social Measurement*, 39(3), 121–144.

Elsby, M. W. L., Michaels, R., & Solon, G. (2009). The ins and outs of cyclical unemployment. *American Economic Journal: Macroeconomics*, 1(1), 84–110.

Enevoldsen, N. (2023). *Filling gaps in India's official labor force survey: Constructing an urban panel*. EGC Perspectives, Yale University.

Fellegi, I. P., & Sunter, A. B. (1969). A theory for record linkage. *Journal of the American Statistical Association*, 64(328), 1183–1210.

Fujita, S., & Ramey, G. (2009). The cyclicality of separation and job finding rates. *International Economic Review*, 50(2), 415–430.

Garfinkel, S., Abowd, J. M., & Martindale, C. (2019). Understanding database reconstruction attacks on public data. *Communications of the ACM*, 62(3), 46–53.

Gomes, P. (2012). Labour market flows: Facts from the United Kingdom. *Labour Economics*, 19(2), 165–175.

ILO & IOM (2024). *One year after: Living and working conditions in Adıyaman, Hatay, Kahramanmaraş and Malatya after the February 2023 earthquakes in Türkiye*. International Labour Organization and International Organization for Migration, Ankara.

Jaro, M. A. (1989). Advances in record-linkage methodology as applied to matching the 1985 Census of Tampa, Florida. *Journal of the American Statistical Association*, 84(406), 414–420.

Madrian, B. C., & Lefgren, L. J. (2000). An approach to longitudinally matching Current Population Survey (CPS) respondents. *Journal of Economic and Social Measurement*, 26(1), 31–62.

Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *IEEE Symposium on Security and Privacy* (pp. 111–125).

Pinedo Caro, L. (2020). Identifying Syrian refugees in Turkish microdata. *IZA Journal of Development and Migration*, 11(1), 1–21.

Pinedo Caro, L. (2021). *The impact of COVID-19 on the Turkish labour market during 2021: Results from a nowcasting model using Google mobility data*. ILO Research Brief, ILO Office for Turkey.

Pinedo Caro, L. (2026a). *A tale of two methodologies: The impact of inflation on wages in Türkiye*. Betam Research Brief.

Pinedo Caro, L. (2026b). *Recovering province identifiers from Turkish Household Labour Force Survey microdata*. Working paper, Bahçeşehir University.

Pinedo Caro, L. (2026c). *The link between building damage and employment losses: Evidence from the 2023 earthquakes in Türkiye*. Working paper, Bahçeşehir University.

Pinedo Caro, L. (forthcoming). *Twelve years in the dark: Are we recovering the urban-rural divide in Türkiye's Household Labour Force Survey?* Betam Research Brief.

Rocher, L., Hendrickx, J. M., & de Montjoye, Y.-A. (2019). Estimating the success of re-identifications in incomplete datasets using generative models. *Nature Communications*, 10, 3069.

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15(2), 127–148.

Sweeney, L. (2000). *Simple demographics often identify people uniquely*. Data Privacy Working Paper 3, Carnegie Mellon University.

Tansel, A., & Kan, E. Ö. (2012). *Labor mobility across the formal/informal divide in Turkey: Evidence from individual level data*. MPRA Paper No. 35672.

Taşçı, H. M., & Tansel, A. (2005). *Unemployment and transitions in the Turkish labor market: Evidence from individual level data*. IZA Discussion Paper No. 1663.

Tunalı, İ. (2008). *Analysis of attrition patterns in the Turkish Household Labor Force Survey, 2000–2002*. ERF Working Paper 393. Published in *METU Studies in Development*, 36(1), 217–252, 2009.

Tunalı, İ. (2023). Unofficial user's guide to the Household Labor Force Survey, Turkey. *Ekonomi-tek*, 12(2), 35–76.

Tunalı, İ., Ekinci, E., & Yavuzoğlu, B. (2012). *Rescaled additively non-ignorable (RAN) model of attrition and substitution*. Koç University–TÜSİAD ERF Working Paper 1220.

TurkStat (2017). *Household Labour Force Survey 2017: Survey metadata*. IHSN Survey Catalog.

Winkler, W. E. (1995). Matching and record linkage. In B. G. Cox, D. A. Binder, B. N. Chinnappa, A. Christianson, M. J. Colledge, & P. S. Kott (Eds.), *Business survey methods* (pp. 355–384). Wiley.

Appendix A. Reconstructing cycles and months: rules and exceptions

As described in Section 3.1, cycles and months are recovered from the order of the records, where the large block of İstanbul locates each cycle and cycles are counted from the start of the file. Reproducing them from the public files requires the exact rules and exceptions listed below, which are followed by the checks that the reconstruction passes.

Locating the cycles. Each cycle contains one large block of İstanbul, which serves as its anchor. A block of İstanbul is considered large when it holds at least a quarter of the largest İstanbul block of the year; the smaller İstanbul blocks that appear in some cycles are treated as ordinary blocks. In cycles in plate order, the large İstanbul block is followed by İzmir (TR31), and the cycle starts a few blocks earlier. The file is read backwards from the anchor until the first of the two blocks of TR62 (until 2013) or of TR22 (from 2014) that precede it. In cycles in subregion order, which start with İstanbul (TR10), the cycle starts at the anchor itself. From 2014, when the TR22 block is immediately preceded by a block of Adana (TR62), the cycle starts one block earlier, at the Adana block.

Table A1 gives the order of the cycles in each month. The change from one order to the other is sometimes gradual, one cycle per quarter starting with cycle 8, and sometimes abrupt, from one month to the next.

Table A1. Order of the cycles by month, 2014–2025

Period	Plate order	Subregion order
Jan 2014	1–8	–
Feb–Mar 2014	–	1–8
Apr–Jun 2014	8	1–7
Jul–Sep 2014	7–8	1–6
Oct–Dec 2014	6–8	1–5
Jan–Mar 2015	5–8	1–4
Apr–Jun 2015	4–8	1–3
Jul–Sep 2015	3–8	1–2
Oct–Dec 2015	2–8	1
Jan 2016 – Sep 2019	1–8	–
Oct–Dec 2019	–	1–8
Jan 2020 – Jun 2021	1–8	–
Jul 2021	–	–
Aug–Sep 2021	–	1–8
Oct–Dec 2021	1–8	–
Jan–Mar 2022	1–7	8
Apr–Jun 2022	1–6	7–8
Jul 2022 – Dec 2025	–	1–8

Source: TurkStat, HLFS public microdata, 2014–2025, and author's own calculations. Notes: The columns give the cycles of the month that follow each order. In July 2021 the records are not ordered by cycle.

Files that do not start in January. Two annual files do not begin with January. In the 2019 file, the records of July to December come first, followed by those of January to June, whereas in the 2023

file the records of April to December come first, followed by those of January to March. In both cases, cycles are counted from the first record of the file, whose month is July (2019) or April (2023).

A month without cycles: July 2021. In July 2021 the records are not ordered by cycle, and consecutive records belong to different subregions. The month is identified as the stretch of the file between the last cycle of June and the first cycle of August. Since its households have no cycle, their visit is imputed (Section 4.1) and their box is formed without the restriction to a cycle (Section 3.3).

Irregular months in 2004–2013. Three months need a special rule. In March 2006, İstanbul appears in several smaller blocks instead of one large block. For that year, blocks of at least a tenth of the largest İstanbul block count as anchors, and a small one only starts a new month if the previous cycle was large. In March 2008, İstanbul does not lie between Mersin and İzmir as usual, so the cycle starts at the last block of Adana before it. Finally, in January 2008, a few records precede the first cycle and are assigned to it.

Subregions missing from a cycle. A subregion with no households in one cycle does not break the count, because cycles are located by the position of İstanbul. In 2015–2019 and 2023–2025, every subregion appears in all eight cycles of every month. In 2014 and 2020–2022, between 0.3 and 7 per cent of subregion-months appear in only seven of the eight cycles, most of them in Ankara (TR51), and no subregion is missing from more than one cycle of a month.

Checks. Every month from 2014 to 2025, except July 2021, contains exactly eight cycles, while every month from 2004 to 2013 contains exactly one. In addition, the longer months of Table 3 contain about 25 per cent more households than the others, in accordance with what the implied calendar predicts (Section 3.3), and the cycle of each household predicts its visit with the precision shown in Table 2.

Appendix B. The structure of the files

The box of Section 3 rests on four regularities of the public files. These are the imprint of the rotation in every period (B.1), the order of the rotation groups within the boxes of 2004–2013 (B.2), the cycles of 2014–2025 (B.3) and the calendar implied by the re-interviews (B.4). All of them can be checked using the public files alone, together with the month of interview obtained in Appendix A.

B.1 Exact matches in every period

Table B1 repeats the count of Table 1 for 2004–2021, in pairs of consecutive years. Before 2021 the files record no date of birth (2014 records it only for adults), so exact coincidence is defined on roster line, sex and age among households of three or more members. Age must be equal at gaps of up to five months, and equal or one year higher from six months on. Until 2013, the search is carried out within the urban or rural part of each province and, from 2014, within NUTS-2 subregions.

Table B1. Pairs of households whose members all match exactly, by gap between interviews, 2004–2021

Gap	2004–2005	2006–2007	2008–2009	2010–2011	2012–2013	2014–2015	2016–2017	2018–2019	2020–2021
1	288	256	282	327	302	6,609	6,482	6,038	4,158
2	250	258	281	290	308	6,244	6,134	5,725	4,440
3	**20,833**	**20,585**	**21,683**	**23,881**	**23,787**	**31,287**	**31,355**	**31,744**	**30,868**
4	236	227	236	286	283	5,783	5,608	5,223	3,188
5	242	213	231	259	240	5,382	5,290	4,891	2,672
6	509	518	532	548	502	5,266	5,155	4,884	2,644
7	369	364	401	447	457	4,963	5,038	4,541	2,171
8	402	387	372	446	433	4,632	4,491	4,191	1,880
9	**5,758**	**6,271**	**6,330**	**7,624**	**7,576**	**7,371**	**7,559**	**7,255**	**5,629**
10	326	294	325	336	343	4,068	4,036	3,753	1,255
11	283	310	305	350	333	3,719	3,680	3,438	716
12	**12,929**	**20,116**	**21,626**	**24,876**	**24,946**	**11,284**	**11,910**	**12,635**	**10,968**
13	227	252	254	281	280	3,195	3,196	2,939	1,139
14	239	217	214	254	242	2,926	2,795	2,714	260
15	**2,046**	**2,529**	**2,820**	**3,154**	**3,189**	**3,727**	**3,750**	**3,729**	**1,745**
16	178	182	188	199	240	2,290	2,362	2,109	239

Source: TurkStat, HLFS public microdata, 2004–2021, and author's own calculations. ***Notes:*** The table counts households of three or more members which are unique on both sides, matched on roster line, sex and age (equal at gaps of up to five months, and equal or one year higher from six months onwards). The search is done within the urban or rural part of each province (2004–2013) and within NUTS-2 subregions (2014–2021). Each column counts the pairs whose two interviews both fall in the two years of the column. For 2022–2025, see Table 1.

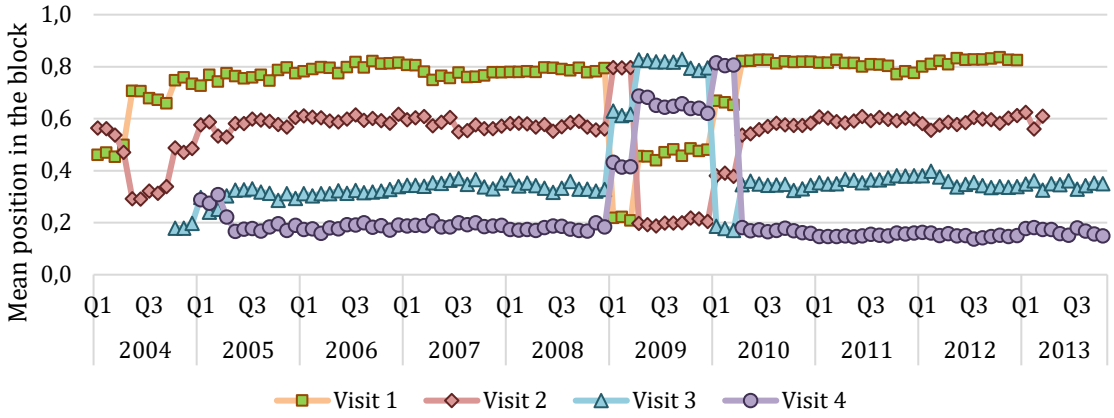
In every period, the gaps of 3, 9, 12 and 15 months stand out from a background of coincidental matches, although less sharply from 2014 onwards. From that year, the background is an order of magnitude higher for three reasons. First, the search runs over the whole subregion rather than the province. Second, age is the only characteristic of the persons that is compared. Third, until 2020 the files record children only as a number, so households differ less from each other. At the gap of

15 months, which fits in only nine of the 24 origin months of a pair of years, the excess over the neighbouring gaps is small in 2014–2019. The cycles of the month, however, remove almost all of this background (Section 3.3), and Table 2 shows that, within the designated cycle, more than 99.6 per cent of the exact matches of 2022–2025 follow the rule.

B.2 The order of the rotation groups, 2004–2013

From 2004 to 2013, each box holds the four rotation groups of the month, one after the other (Section 3.2), and Figure B1 shows what their order looks like, providing an independent check of the order assumed by the box. For every month, the figure plots the mean relative position in the box of the households whose visit is determined by the shape of their history alone, that is, by the months in which the household is observed, without using its position in the file.

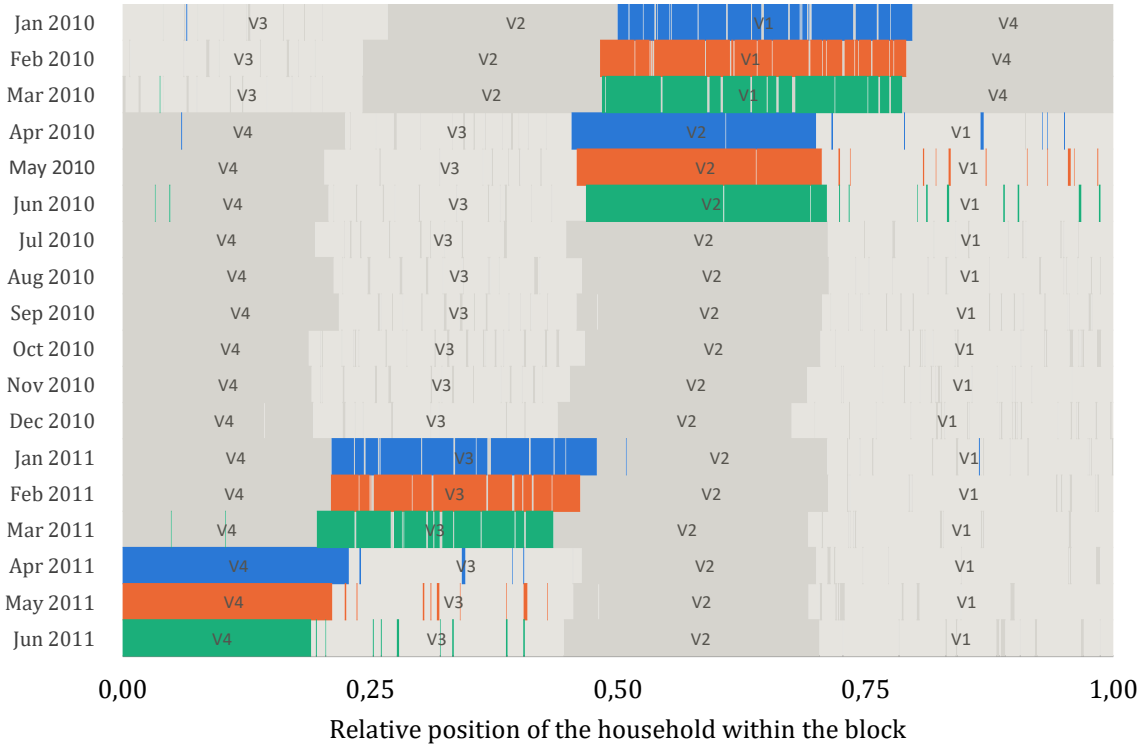
Figure B1. Mean relative position in the box of the households at each visit of their rotation group, 2004–2013



Source: *TurkStat, HLFS public microdata, 2004–2013, and author's own calculations.* ***Notes:*** *Each line shows, for the households at a given visit of their rotation group, their mean position in the box (i.e., the rank of the household in file order within the province × urban/rural records of the month, divided by their number, where 0 = start of the box and 1 = end). The visit is inferred only from the months in which the household is observed, and not from its position in the file.*

In most months, the four groups appear in the standard order, with the fourth visit first and the first visit last; in other words, groups are ordered by entry, the oldest first. The exceptions are those listed in Section 3.2. Figure B2 illustrates this order in one box of İstanbul followed over eighteen months.

Figure B2. The four rotation groups within the İstanbul urban box, January 2010 – June 2011



Source: TurkStat, HLFS public microdata, 2010–2011, and author's own calculations. ***Notes:*** Each row represents the box of İstanbul urban households in the month, in file order. Coloured segments show the rotation groups that entered in January (blue), February (orange) and March 2010 (green), as linked by the algorithm, while grey segments show the other rotation groups. V1–V4 denote the visit of each group. The groups follow the standard order (visit 4, 3, 2, 1) from April 2010, whereas in January–March 2010 the groups that entered in 2009 still lead the box (Appendix B.2). About 2 per cent of the households of the coloured groups lie outside their segment, since the boundaries between rotation groups, especially between the fourth and third visits, are not perfectly sharp.

In 2004 and early 2005, only the groups entered from January 2004 onwards can be followed, so the figure shows the first and second visits from January 2004, the third from October 2004 and the fourth from 2005. Until April 2004, the first and second visits overlap in the box, while in early 2005 the third and fourth visits are not yet separated.

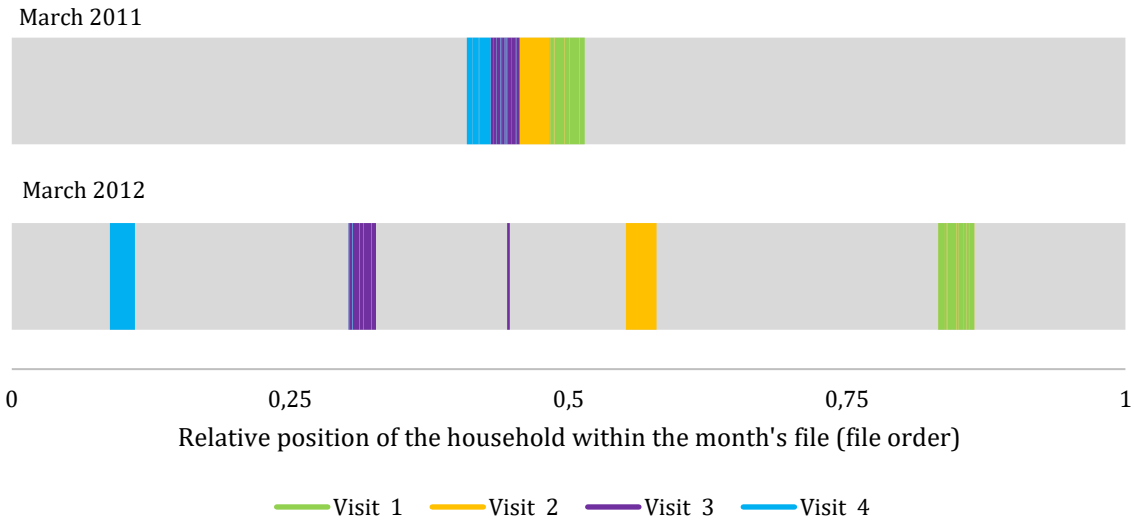
From January 2009, the groups entering the sample were placed at the start of the box, ahead of the groups that had entered before 2009, and kept that place at their later visits. As a result, the order of the box changed as these groups advanced through their visits and the older ones left the sample. In particular, the order was first visit, fourth, third and second in January–March 2009; second, first, fourth and third from April 2009; and third, second, first and fourth in January–March 2010. In April 2010, the last group that had entered before 2009 left the sample, and the standard order came back.

At the other end of the period, the visits of the last months of 2013 cannot be determined from histories that end with the 2013 file, which explains why the lines stop before the end of the period.

In 2012, the layout of the file changes, but the order does not (Section 3.2). Instead of one block per province and month, each province appears in four or five blocks, each of them holding a single

rotation group. Figure B3 shows the records of İstanbul within the file of March 2011 and of March 2012.

Figure B3. İstanbul records within the whole month's file, March 2011 and March 2012



Source: *TurkStat, HLFS public microdata, 2004–2013, and author's own calculations.* ***Notes:*** *Each bar represents the file of the month (all households in Türkiye, in file order), and coloured segments show the İstanbul households by visit of their rotation group, as determined by the panel. In March 2011 İstanbul forms one block, while in March 2012 it forms one block per rotation group.*

B.3 The cycles of 2014–2025

From 2014, each month consists of eight cycles, and each cycle holds half of one rotation group at a given visit (Section 3.3). Figure B4 follows the eight blocks of İstanbul, one in each cycle, over eighteen months. Each rotation group enters in cycles 6 and 8, moves to cycles 5 and 7 three months later, to cycles 2 and 4 nine months after that, and to cycles 1 and 3 for its last visit, as shown in Table 2, with the exception of part of cycle 8 in 2020 and the first half of 2021 (Section 3.3). Between 79 and 98 per cent of the households of each block at the second to fourth visit, and about 90 per cent in most blocks, are linked to their previous visit. From 2023, each cycle covers the subregions twice and, for the half rotation groups entering from January 2024, three times, presumably once for each degree of urbanisation (Pinedo Caro, forthcoming). The box of a subregion in a cycle may therefore span several blocks.

Figure B4. İstanbul (TR10): the eight blocks of each month, January 2016 – June 2017

Month	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6	Cycle 7	Cycle 8
Jan 2016	V4 · 93%	V3 · 89%	V4 · 93%	V3 · 90%	V2 · 89%	V1	V2 · 87%	V1
Feb 2016	V4 · 93%	V3 · 88%	V4 · 96%	V3 · 88%	V2 · 89%	V1	V2 · 94%	V1
Mar 2016	V4 · 95%	V3 · 91%	V4 · 95%	V3 · 92%	V2 · 95%	V1	V2 · 92%	V1
Apr 2016	V4 · 97%	V3 · 95%	V4 · 95%	V3 · 88%	V2 · 92%	V1	V2 · 92%	V1
May 2016	V4 · 92%	V3 · 91%	V4 · 93%	V3 · 86%	V2 · 94%	V1	V2 · 89%	V1
Jun 2016	V4 · 93%	V3 · 89%	V4 · 92%	V3 · 94%	V2 · 88%	V1	V2 · 87%	V1
Jul 2016	V4 · 97%	V3 · 93%	V4 · 92%	V3 · 87%	V2 · 92%	V1	V2 · 94%	V1
Aug 2016	V4 · 95%	V3 · 92%	V4 · 93%	V3 · 93%	V2 · 92%	V1	V2 · 90%	V1
Sep 2016	V4 · 95%	V3 · 89%	V4 · 96%	V3 · 90%	V2 · 89%	V1	V2 · 89%	V1
Oct 2016	V4 · 96%	V3 · 90%	V4 · 94%	V3 · 88%	V2 · 79%	V1	V2 · 87%	V1
Nov 2016	V4 · 96%	V3 · 88%	V4 · 90%	V3 · 90%	V2 · 84%	V1	V2 · 86%	V1
Dec 2016	V4 · 96%	V3 · 92%	V4 · 90%	V3 · 90%	V2 · 91%	V1	V2 · 90%	V1
Jan 2017	V4 · 93%	V3 · 92%	V4 · 95%	V3 · 87%	V2 · 92%	V1	V2 · 92%	V1
Feb 2017	V4 · 91%	V3 · 89%	V4 · 98%	V3 · 91%	V2 · 91%	V1	V2 · 95%	V1
Mar 2017	V4 · 95%	V3 · 89%	V4 · 94%	V3 · 88%	V2 · 92%	V1	V2 · 95%	V1
Apr 2017	V4 · 96%	V3 · 92%	V4 · 94%	V3 · 89%	V2 · 95%	V1	V2 · 89%	V1
May 2017	V4 · 96%	V3 · 87%	V4 · 95%	V3 · 90%	V2 · 95%	V1	V2 · 94%	V1
Jun 2017	V4 · 97%	V3 · 90%	V4 · 94%	V3 · 91%	V2 · 95%	V1	V2 · 94%	V1

Source: TurkStat, HLFS public microdata, 2016–2017, and author's own calculations. ***Notes:*** Each cell represents the block of İstanbul (TR10) households in one cycle of the month. V refers to the visit of the rotation group (with visit 1 shaded in green, 2 in orange, 3 in purple and 4 in blue), and % to the share of the block's households linked to their previous visit, which is not shown at the first visit.

B.4 The calendar implied by the re-interviews

Suppose that, from 2014, a re-interview takes place 13, 39, 52 or 65 weeks after the previous visit, in the same week of the quarter, and that the calendar of Table 3 assigns each week to a month. Then, when a week that belongs to month *t* belongs, 13 or 52 weeks later, to a month other than *t* + *g*, its households are re-interviewed in the neighbouring month. This hypothesis predicts where such re-interviews occur, and the exact matches of the public files confirm it. In 2021–2025, using matches on date of birth among households of two or more members, there are 193 combinations of origin month and gap with a complete target window. In 168 of them, the calendar predicts no shift and the share of exact matches found in the neighbouring month is at most 0.2 per cent. In the remaining 25, the calendar predicts that one re-interview in five moves, and the observed share lies between 18 and 21 per cent (Table B2).

Table B2. Re-interviews in the neighbouring month predicted by the calendar of Table 3, 2021–2025

Origin month	Gap	Re-interview month	Pairs	Predicted (%)	Observed (%)
Jun 2022	+15	Aug 2023	2,713	20	19.2
Sep 2022	+12	Aug 2023	6,005	20	19.3
Sep 2022	+15	Nov 2023	2,647	20	19.5
Dec 2022	+9	Aug 2023	3,247	20	19.8
Dec 2022	+12	Nov 2023	5,841	20	20.1
Dec 2022	+15	Feb 2024	2,771	20	20.4
Mar 2023	+9	Nov 2023	3,150	20	20.3
Mar 2023	+12	Feb 2024	5,822	20	19.4
Mar 2023	+15	May 2024	2,692	20	19.0
Jun 2023	+3	Aug 2023	6,918	20	20.5
Jun 2023	+9	Feb 2024	3,213	20	19.4
Jun 2023	+12	May 2024	5,823	20	20.3
Jun 2023	+15	Aug 2024	2,604	20	21.1
Aug 2023	+15	Oct 2024	2,776	20	19.4
Nov 2023	+12	Oct 2024	6,040	20	19.5
Nov 2023	+15	Jan 2025	2,725	20	19.5
Feb 2024	+9	Oct 2024	3,156	20	19.0
Feb 2024	+12	Jan 2025	5,908	20	20.1
May 2024	+9	Jan 2025	3,211	20	21.2
May 2024	+15	Jul 2025	2,548	20	18.0
Aug 2024	+3	Oct 2024	7,021	20	20.1
Aug 2024	+12	Jul 2025	5,569	20	19.0
Aug 2024	+15	Oct 2025	2,558	20	20.0
Jan 2025	+3	May 2025	7,053	20	20.7
May 2025	+3	Jul 2025	6,775	20	19.2

Notes: The table uses exact matches on roster line, sex and year and month of birth among households of two or more members, within the cycle designated by the rule of Table 2. Pairs refers to all exact matches at months $t + g - 1$, $t + g$ and $t + g + 1$. Predicted refers to the share of the interview weeks of the origin month that the calendar of Table 3 places in the neighbouring month (one out of five). In the remaining 168 combinations of origin month and gap, the calendar predicts no shift.

The cluster rule of Section 3.3 exploits the fact that the re-interviews that fall in the neighbouring month are not spread at random over the households of a box. Instead, they come in runs of contiguous records, as if each interview week covered a group of neighbouring dwellings. In the 25 combinations of Table B2, the households re-interviewed in the neighbouring month form runs of eight households on average in file order, and 99 per cent of them belong to runs of three or more. Among the nearest neighbours of a shifted household (up to three on each side, in the same box), 69 per cent are shifted as well, compared with 8 per cent among the neighbours of a household that is not shifted.

Appendix C. Fixed vs. Fellegi–Sunter weights

As explained in Section 2.3, the weights are fixed in advance because Fellegi–Sunter weights cannot be estimated for the whole period; they require pairs of known match status, which dates of birth only provide in 2014 for adults and from 2021 for everyone. Nonetheless, where dates of birth exist, estimated weights can be compared with the fixed ones. To do so, I take every candidate pair formed by the algorithm for origin months in 2023–2024 (5.1 million person pairs in 3.6 million household pairs) and classify each household pair as a match or a non-match based on dates of birth. For each characteristic, I estimate the probability of agreement among matches (m) and among non-matches (u). Its agreement weight is then $\log_2(m/u)$ and its disagreement weight $\log_2\big((1-m)/(1-u)\big)$. Relative position within the box is treated as one more comparison, with weights by distance class, and the weight W of a household pair is the sum of the weights of its comparisons.

Table C1. Agreement probabilities and weights: Fellegi–Sunter estimates and the fixed weights of Section 2.3

Characteristic	m (agree)	u (agree)	FS weight, agree	FS weight, disagree	Fixed weight, agree	Fixed weight, disagree
Age as expected (vs one year short)	0.993	0.655	0.60	−5.63	1	0.5 (partial)
Born abroad	0.996	0.954	0.06	−3.44	0	−0.5
Year of arrival in the province	0.281	0.007	5.37	−2.99	0.1	−0.5
Marital status	0.969	0.854	0.18	−2.24	0.05	0
Relationship to the head	0.921	0.869	0.08	−0.72	0.05	0
Education not lower	0.958	0.703	0.45	−2.84	0.025	0
Relative position, |Δr| ≤ 0.02 / > 0.2			3.95	−2.32	0	−0.2 to −1

Notes: FS weights in bits ($\log_2$). Year of arrival in the province: agreement requires both visits to report it; otherwise the comparison is missing and carries no weight. The utilized position term is $-|r - r'|$. Age: the "disagree" columns refer to an age one year short of the expected ageing, which the algorithm's score treats as partial agreement; persons outside the age window are not paired.

Both sets of weights order the evidence in the same way (Table C1) and, since so few candidates remain, the differences between them hardly matter. The strongest evidence in favour of a match is agreement on rare characteristics, such as the year of arrival in the province, whereas the strongest evidence against it is disagreement on characteristics that cannot change, such as age or country of birth. Compared with the estimates, the weights used in the algorithm penalise disagreement on marital status, relationship to the head and education less, and give more weight to the matched age itself.

Table C2. Links obtained with fixed weights and with the estimated Fellegi–Sunter weights, origin months 2023–2024

Weights and rule	Links	Confirmed by dates of birth (%)	Same links as fixed weights (%)
Section 2.3 weights	495,672	99.73	100
FS weights, two-step best, $W > 0$	501,371	99.77	98.8
FS weights, two-step best with a margin of 1 bit, $W > 0$	500,272	99.81	98.7
FS weights, two-step best with a margin of 1 bit, $W > 2$	490,346	99.87	96.8

Notes: The table shows links at the scheduled gap, and the matched-member conditions of Section 2.3 apply in all rows. Since FS weights are estimated on the same months, their precision should be taken as an in-sample upper bound.

Estimated weights cannot be obtained for 2004–2013 or 2015–2020, where they would be needed, but where they can be estimated they change little. Replacing the weights of the algorithm with the estimated Fellegi–Sunter weights changes about one link in a hundred (Table C2), and precision moves by less than a tenth of a point. This result illustrates what a bounded search means in practice since, once the structure has reduced each household's box to a handful of candidates, any reasonable weighting of the evidence ends up choosing the same candidate. As expected, stricter thresholds trade a few links for slightly higher precision.