A study of the labour market trajectories in the Grand Duchy of Luxembourg

Nicholas T. Longford
SNTL and UPF, Barcelona, Spain,

Ioana C. Salagean
CEPS/INSTEAD, Esch-sur-Alzette, Luxembourg

Abstract

We present an analysis of the register of all unemployment episodes in the Grand Duchy of Luxembourg over a recent period of 55 months. We apply propensity score matching to account for the systematic differences among the groups of subjects (registrants) and unemployment spells. We devise graphical and tabular summaries for describing the sequences of employment states of the members of the labour force who register at *Agence pour le développement de l'emploi*, the Luxembourg Public Unemployment Agency. Some employment-related information about them is collected by linking their records to the national register of social security contributions, maintained by *Inspection générale de la sécurité sociale.* A class of univariate indices for characterising the sequences of labour force states is defined.

Keywords: *Administrative register; index; labour force status; propensity matching; sequences of states; unemployment spells.*

N. T. Longford, SNTL and Departament d'Economia i Empresa, Universitat Pompeu Fabra, c/ Ramon Trias Fargas 25 – 27, 08005 Barcelona, Spain. Email: `sntlnick@sntl.co.uk` .

I. C. Salagean, CEPS/INSTEAD, Avenue de la Fonte 3, L–4364, Esch-sur-Alzette, Luxembourg.
Email: `ioana.salagean@ceps.lu` .

1 Introduction

Unemployment is a concern in most developed countries, especially in periods of slower economic growth and recession. The unemployed represent a waste of skills and expertise, and of human capital in general, that is not engaged in the generation of wealth and in self-realisation. A fraction of the income lost may be replaced by unemployment insurance and social security payments. These payments are a considerable drain on public resources and a threat to the government's balance of payments when the rate of unemployment increases. This induces the public unemployment agencies, where the newly unemployed submit their applications for income support and seek assistance with job search, to monitor the labour force careers of the unemployed. The ideal career is that a worker would find a job soon after becoming unemployed and would stay employed, not necessarily in the same job, for a long time.

The unemployment rate summarises the unemployment cross-sectionally, treating a single time point (a month) as the domain. More detail, related to the persistence of unemployment and turnover among the unemployed, can be obtained by studying the unemployment spells and their outcomes, treating each individual as a member of a panel for the period of his or her unemployment. This report presents such a study and explores how it can contribute to a better assessment of unemployment as an economic and social problem and to devising more effective measures for dealing with unemployment. The principal motive for the study are the budgetary pressures on the system of unemployment benefits, although the social consequences and the financial impact of (long-term) unemployment on the household also merit a consideration.

In the Grand Duchy of Luxembourg, the expenditure on employment policies reached 514 million Euros in 2010, 1.2% of the gross domestic product, and about 1000 Euro per head of population. Of this amount, 63% is spent on unemployment benefits and related payments. It has increased by 46% between 2007 and 2011 (Ries, 2012). The rate of unemployment in Luxembourg has risen in the recent years and it reached 6.1% in May 2012, but it remains very low in relation to other countries of the European Union. (The Union-wide rate was 10.3% in May 2012.) However, the unemployment benefits are rather generous, and the Government of the Grand Duchy may not be able to maintain them if the rate of unemployment increases further or remains elevated for a long time.

The unemployment agency of Luxembourg, *Agence pour le développement de l'emploi* (ADEM), maintains a database of all its appointments with applicants for assistance with job seeking and for unemployment benefit. A case file in ADEM corresponds to a spell of unemployment (job seeking) of an individual. A list of unemployed is compiled at the end of every month. An individual can have several spells of unemployment in the studied period. A spell may conclude with a switch to (new) employment or to economic inactivity, but this outcome is not recorded by ADEM. Economic inactivity arises, for example, when the person is discouraged from further search by failures or reaches the age of retirement. A detailed description of the database is given in Section 2.1.

The General Social Security Registry of Luxembourg, *Inspection générale de la sécurité sociale* (IGSS), maintains records of all social security contributions, based on the employment contracts and contributions from the self-employed, and indicates the outcome of the job search of a person who was unemployed earlier. The two databases, ADEM and IGSS, can be linked and the employment status of individuals with a spell of unemployment established for both the future (following the spell) and the past (preceding the spell). From these databases, we constructed the sequences of labour force states (employed, unemployed and not in the labour force) recorded at the end of each of the 55 months from January 2007 to July 2011. The records are incomplete for subjects who entered the labour force during the period of the study (mostly new members of the labour force, such as school leavers and graduating students). Many commuters from the neighbouring countries (Belgium, France and Germany) and residents of other countries (mainly other members of the European Union) are employed in Luxembourg. When their tenure of employment expires and they return to their home countries, their employment status can no longer be traced. To maintain integrity of the database and to address the main concerns of the Luxembourg social security and unemployment support system, we restrict our study to the residents of Luxembourg with a record in ADEM; it comprises 62 720 records (sequences of length 55) which involve nearly 100 000 unemployment spells, nearly 44 000 for persons aged 30 years or younger at the time of registration in ADEM, at the beginning of the unemployment spell.

The study described in this paper is concerned with labour market histories and with unemployment spells, their precursors and spells of subsequent employment, if any. Studying spells of unemployment on their own is not satisfactory, because the spells that follow them are also of interest. We regard these spells as the resolution of unemployment, and classify them as successes (long-term employment) or failures (return to unemployment). Further, spells that precede an unemployment spell have a role to play akin to covariates (conditioning variables) in a regression model. The intent of the study is to propose methods for a package of analyses that could be conducted regularly (e.g., annually) to assess the success of integrating unemployed and school leavers in the labour force, estimate the distribution of unemployment spells (of their duration, starting dates and other attributes) within groups defined by the basic socio-demographic categories, and inform the process of allocating resources for dealing with the financial and other committments made to the unemployed as well as for measures to combat unemployment. These analyses are meant to supplement the information gained from cross-sectional analyses which are limited to estimates of unemployment rates (and their changes over time) for the entire labour force and its subpopulations defined by socio-demographic variables.

Although the rate of unemployment in Luxembourg is much lower than in the neighbouring countries, and the adjoining regions in particular, the unemployment benefits are rather liberal and generous by European standards. The main features of the unemployment benefit scheme are summarised in the Appendix. Preparation for contingencies if the current economic crisis persists and affects the country's labour market more seriously than it has thus far (May 2012) are an important aspect of the social security policies.

We develop a set of methods for the analysis of sequences of (discrete) labour force states (employed, unemployed, economically inactive and in transition between states) that would assess the phenomenon of unemployment in a much finer detail than by the monthly rates of unemployment. Of particular concern are the lengths of unemployment spells, the success of their resolution — stability of the succeeding employment and the distribution of these phenomena across socio-demographic groups defined by age, sex, educational level, and the like.

Our first goal is a compact graphical summary of the sequences of labour force states of length up to 55 months, together with tabular summaries of the lengths of stay in the particular states and changes of states from one month to the next. Then we study indices that summarise such sequences, and relate them to background variables. The next section describes the data sources and their purpose. The following section discusses discrete longitudinal data (sequences of states). Section 4 presents a set of analyses of the data compiled from the ADEM and IGSS databases. The paper is concluded by a discussion.

2 The roles of ADEM and IGSS

The mission of ADEM is to promote the optimal use of the potential labour supply in the country by providing assistance to job seekers and forwarding and mediating job offers to them. ADEM also offers job counselling and training to selected subpopulations, including those with handicap and limited capacity to work. It determines whether job seekers are entitled to receive unemployment (insurance) benefits and sets the amounts of benefits they should receive and over what time period (the payment regime).

Registration of job seekers at ADEM is voluntary. There are two main incentives for the unemployed to register and keep regular appointments with ADEM throughout the process of job search: first, unemployment-related benefits are received only after a (successful) application and compliance with ADEM's requirements; second, all employers are, at least in theory, required to inform ADEM of their job openings before making them public, without any restrictions on how that information can be used, so ADEM is, in effect, also an employment agency. Employers are compelled by law to treat a job application mediated by ADEM without any prejudice.

ADEM maintains a record (a case file) of each application submitted by a resident of Luxembourg until its resolution. A case file is resolved when the applicant is no longer classified as unemployed: when he or she gains employment, gives up the job search (for instance, reaches the statutory age of retirement), or materially breaches the regulations of ADEM. We combine the information in this register with the information in the IGSS database to form personal sequences of labour force states over a period of 55 months (January 2007− July 2011).

A part of the Ministry of Social Security in Luxembourg, IGSS collects and analyses the statistical information needed to steer and assess progress towards the government's objectives in matters concerning social security, both in the short and long term. To this aim, IGSS records all social security contributions associated with paid (employment) contracts in Luxembourg. We derive all employment-related information from this database, which we describe, together with the procedures applied, in the next three sections.

2.1 ADEM database

All job seekers in Luxembourg may register in ADEM, whether previously employed or not. Upon registration, ADEM opens an unemployment file in order to monitor the progress of the applicant's job search. This case file is updated at each appointment, typically at monthly intervals, until it is closed, either successfully when the applicant finds a job, or unsuccessfully when he or she fails to attend three consecutive appointments, in which case ADEM classifies the job search as abandoned. Discouragement by (repeated) failure to secure employment, reaching the age of retirement and falling ill are some of the reasons for closing a case file.

An ADEM case file is updated at each monthly meeting with the applicant, without any follow up (e.g., by telephone) if the appointment is not kept. Therefore, some of the case files are not complete; in particular, the dates when some unemployment spells (uninterrupted periods of time spent entirely in unemployment) are concluded are not established with certainty. Below we describe how we resolve such ambiguity and delimit an unemployment spell by its start and end dates.

We have a series of 55 datasets, one for each month between January 2007 and July 2011. Each dataset informs about all ADEM case files that are (still) open on the last working day of the month. The status, unemployed or not unemployed, is assigned to a person for the whole month, irrespective of the exact date when a change in the labour force status may have taken place. The assignment is made on the last working day of the month. A person still unemployed that day is regarded as unemployed for the whole month; otherwise he or she is regarded as not unemployed. In general, the ADEM records do not contain information about the outcome of the case (employment, economic inactivity, emigration or death).

The uniquely identified ADEM case files contain information concerning an applicant's date of registration in ADEM, the socio-economic characteristics of the applicant, the characteristics of previous employment and of the employment sought, as well as the participation in job counselling, training and subsidized employment schemes, including those that are part of active labour market policies (ALMP).

A spell of unemployment is considered to start in the month when the application in ADEM is submitted. When the applicant attends every scheduled appointment with ADEM until the closure of the case file, the end of the unemployment spell is defined as the month in which the last entry in the case file is made. In case of one or several absences (non-attendance), the end of the unemployment spell is established by the following rules:

- If an individual has fewer than three consecutive absences preceded and followed by ADEM appointments that were kept, the unemployment spell is considered to have lasted throughout the period of absence — the 'gaps' in the case file are filled accordingly.

- If a case file records three or more consecutive absences which are preceeded and followed by ADEM appointments that were kept, the case file is broken into two separate unemployment spells: the first one ends in the month of the first absence and the second starts after the first (renewed) contact with ADEM following the string of absences.

- If two case files related to the same person are adjacent, that is, a file is opened in the month immediately following the closure of the first file, the two files are merged into a single case file, corresponding to a single unemployment spell.

An individual may experience several unemployment spells, separated by at least one month. Such case files have distinct identification numbers but refer to the same person, identified by his or her social security number.

2.2 IGSS database

We establish whether an individual with an ADEM case file closed at some point is employed at another (given) time point by searching the database obtained from IGSS. The database comprises 55 monthly datasets with records of all social security contributions related to employment contracts or paid by the self-employed in the period January 2007 – July 2011. There is a separate record for each contract for which an individual is paying (self-)employment related social security contributions in Luxembourg. The contracts include voluntary contributions, parental leave, and the like.

An individual may have several records in the IGSS database at any given month, or none at all. We reduce each monthly dataset to the person level by defining a variable that indicates whether the person has made any (self-)employment contributions or not. We do this by the following steps:

1. we discard all records related to voluntary contributions to the social security by students, retirees, and the like;
2. if there are several contracts at the end of month, we take the main one, i.e., the first one after sorting by the following criteria:
   - indicator of the level of occupation (main vs. secondary);
   - number of hours of work;
   - tenure of the occupation; and
   - social security number of employer;
3. if no job is held at end of the month, we apply the same rules as in par. 2 to all the jobs held at any time during the month.

A spell of employment starts in the month in which social security contributions are paid for the first time (after an interruption). The spell is considered to end in the month when the individual pays no employment-related social security contributions for the first time after a sequence of payments. Change of jobs, or even of employers, entails no end to an employment spell if one job succeeds the other without an intervening period.

2.3 Combining ADEM and IGSS data

The (monthly) sequences of labour force states are compiled by linking the records in the ADEM and IGSS databases using the person's identifier (social security number). We establish the labor force status for each person-month by inspecting ADEM and IGSS datasets for the month. We apply the following rules:

- an entry in IGSS, but not in ADEM — employed (EM);
- an entry in ADEM, but not in IGSS — unemployed (UN);
- entries in both ADEM and IGSS — transition (T);
- entry in neither ADEM nor IGSS, but an entry in either of them in an earlier month — economically inactive (IA);
- entry in neither ADEM nor IGSS, and no entry in either of them at any time in the past — absence (AB). This is relevant only to persons who appear in a dataset for a later month.

The transition status (T) covers employment and training subsidised by the Government and administered by ADEM, and registration in ADEM shortly before the (anticipated) loss of employment. It also includes temporary employment arranged by ALMP, typically short-term work appointments (e.g., for a month or less). Searching for a job without the assistance of ADEM and waiting for an accepted job or contract to commence are classified as economic inactivity. A person has status AB (absence) until his or her first registration in either ADEM (as unemployed) or IGSS (as employed).

A large fraction of the labour force in Luxembourg (44% in 2011) are not residents of the country. Most of them commute to work from the regions of Belgium, France and Germany that adjoin Luxembourg. They make social security contributions and have entries in IGSS, but their absence from IGSS and ADEM in any particular month cannot be interpreted as economic inactivity, because they may be employed or registered as unemployed in another country, where they may be eligible for unemployment benefits. (We do not have access to such information.) We exclude them from our study because they are subject to a different social security legislation in Luxembourg, which has been altered recently.

3 Sequences of labour force states

Sequences of labour force states (employment, unemployment, economic inactivity and possibly some other transitory states) have some distinct properties. First, many members of the labour force are continually employed, and others are employed for most of the time, except for occasional short spells of unemployment. Next, younger members of the labour force tend to have more changes of status, and shorter stays in any state (including short-term employment spells), than older members many of whom hold permanent jobs. For older workers, unemployment is often followed by economic inactivity (retirement). Thus, in the terminology of Elzinga (2010), complexity of the sequences (frequent changes of states) is concentrated in a small subsample. Among the states, employment is regarded as positive and unemployment as negative, but economic inactivity and other states are difficult to place on an (ordered) scale.

For changes (development) over time, it is essential to study sequences, that is, (discrete) longitudinal data. The basic unit of analysis is a sequence of T states, $\mathbf{x} = (x_1, x_2, \ldots, x_T)$, where each x_t is one of a (finite) set of K (labour force) states $S = (r_1, r_2, \ldots, r_K)$. In our case, $T = 55$ and $K = 5$. Analysis of such sequences is of interest in a variety of areas of social sciences, because many phenomena are recorded by categorical variables, not necessarily with ordered categories. Any cross-sectional analysis, of the states at a given time, is of limited utility, because many important questions relate to the patterns of the sequences, and spells (sequences of consecutive time points in the same status) in particular. They include persistence (tendency to stay in a status for a long time) and patterns of changes (switches from one status to another).

A prominent example of discrete longitudinal analysis arises in the study of transition between the states of poverty and prosperity (non-poverty). A borderline (or transition) status can be added to these two states, but the three (or more) states are ordered. The states considered in studies of cohabitation and child-bearing (living alone, cohabiting, married, widowed, etc.) are without an obvious ordering, and transition from some states to some others is either very rare or impossible.

Methods for longitudinal analysis are better established and are closer to the mainstream statistics for continuous data (with multivariate normal distribution, possibly after a transformation); compare, for instance, Verbeke and Molenberghs (2000) and Molenberghs and Verbeke (2005). One way of taking advantage of these methods is by regarding the observed (discrete) vectors as coarse versions of latent (continuous) vectors. Apart from methods for latent variables, methods for missing data can be applied, generating several plausible latent vectors, analysing each set, and summarising the results; see Heitjan (1989) and Rubin (2002). This approach is attractive when an underlying continuous vector is easy to interpret, or even refers to a well defined quantity, such as income.

The key operations in our analysis are comparisons of the sets (collections) of sequences for a small number of (sub-)populations. These populations are defined in the same period of time (e.g., members of the labour force in the period 2007 − 2011, classified according to their level of education), but also in different time points, for example, to compare the labour force of a country in two distinct periods of time.

3.1 Describing a sequence

We use the term *spell* for an uninterrupted sequence of months spent in the same labour force status, such that the immediately preceding and following states are different. A contiguous subset of a spell is called a *censored* or incomplete spell. A sequence can be summarised by its spells. For example, [EM (1–5), UN (6–27), IA (28–55)] describes (completely) the sequence that starts with five months of employment, follows by 22 months of unemployment, and concludes with economic inactivity (months 28–55). As an alternative, the switches (changes from one state to another) can be listed together with time points (months) that delimit the spells. For the example above, [5 (EM–UN), 27 (UN–IA)] is the alternative description. We do not want to ignore the order or the states, so this summary differs essentially from [UN (1–22), EM (23–27), IA (28–55)]. Neither do we want to add up the lengths of separate spells in a status, so the sequence defined by the summary [UN (1–12), EM (13–17), UN (18–27), IA (28–55)] is also different. The variable length of these summaries presents a practical difficulty in describing large sets of sequences.

A sequence can be represented graphically by a segment with sub-segments of distinct colours to indicate the spells according to status. Figure 1 presents an example of three individuals. The first was employed until month 27 (March 2009), then spent one month in transition, followed by five months of unemployment. After four months of employment (October 2009 − January 2010), he or she was unemployed till the end of our records in July 2011. This person had four changes of labour force status, that is, five spells, two each in employment and unemployment, and one in transition.

The vector of times spent in each of the five states, (31, 23, 0, 1, 0), for employment, unemployment, inactivity, transition and absence, is incomplete because it does not inform about the spells — the order in which the states occur and how frequently they change (the *sequencing*) is lost. It is well recognised in the labour market literature that, say, a half-year spell of unemployment, preceded and followed by periods of employment is qualitatively different from two spells of unemployment lasting three months each and separated by a spell of employment, because the disadvantage of unemployment accumulates and its effects are easier to undo for several small quanta than for one large dose.

Spells of transition are rare and usually last only a month or two. They could be redefined to the immediately preceding status. To simplify a summary, we remove in a sequence any spell lasting a

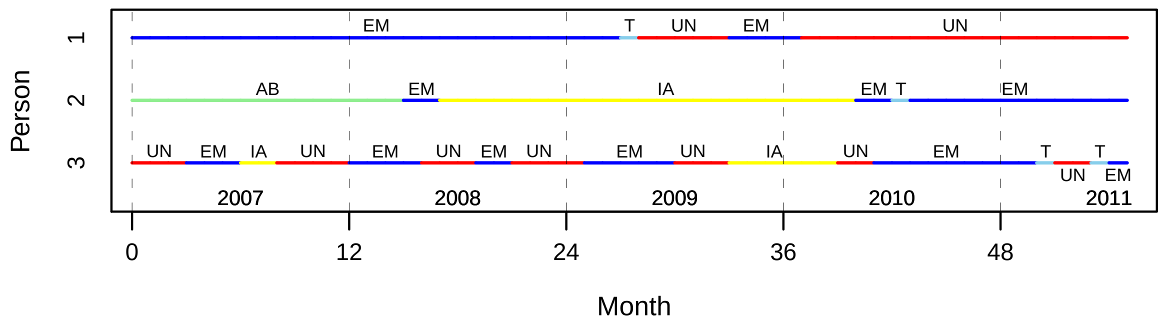

Figure 1: Examples of sequences of labour force states. The states are: EM (blue) — employment; UN (red) — unemployment, IA (yellow) — inactivity; T (skyblue) — in transition; AB (green) — absent.

single month if it is preceded and followed by at least two months in the same status. For example, the second person in Figure 1 was in transition in month 43 (July 2010), after two months of employment (EM) and before several (at least 12) months of another EM spell. Thus, we change the status for month 43 to EM, so that the person has a single (censored) spell of EM from month 41 (May 2010) till the end of our records (July 2011). We refer to such changes as *status smoothing*. No smoothing is applicable for the third individual because the two cases of a single-month spell (months 52, and 54) are preceded and succeeded by spells in different states.

The *majority status* of a sequence is defined as the status in which the individual spent more time than in any other status during the 55 months in our records. In case of a tie, the status is selected at random. We use the term *preamble* for the status in and the length of the first spell in the records. Thus, the first individual in Figure 1 has an EM preamble of length 27. A preamble is said to be *majority* if it is in the majority status. Note that such a preamble does not have to be long, because there may be other spells in the majority status. The *postamble* is defined similarly, relating to the last spell. Preambles and postambles are in general censored or incomplete spells, because a preamble may be preceded by months in the same status, in late 2006, and postamble may be followed by more time in the status, from August 2011 on.

A spell that starts after the first month and ends before the last month is called a *mid-spell*, and is said to be majority if the person has that status for more months than any other status. The second person in Figure 1 has a majority mid-spell of IA of length 23 (months 18 – 40). The main roles of preambles, postambles and longest mid-spells are for ordering the individuals so that their graphical representations, as illustrated in Figure 1, would be easy to interpret.

In summary, a sequence comprises either a single censored spell, a preamble and a postamble, or a preamble, one or several mid-spells (uncensored spells) and a postamble.

3.2 Indices

A set of sequences can be represented graphically by segments coloured according to the states. Such diagrams represent the sequences completely, without discarding any information, but their interpretation requires fine judgment. Tabular summaries of such sequences are usually impractical because there are too many unique sequences (patterns). Tabulation of some of the within-sequence summaries might be useful but they are manageable only for summaries that are too simple. For example, there are 23 700 distinct quintets of numbers of months spent in the five labour force states and 16 000 of them occur only once each. If we count the numbers of completed half-years, we have 1017 distinct patterns, 210 of which occur only once. Counting the completed years yields 115 distinct patterns, only four of which occur once each, but such rounding is too coarse for most purposes.

We use the term *index* for a univariate function $g(\mathbf{x})$ defined in the space of sequences. Simple examples include the count of elements with a particular value (e.g., the number of months in UN) and the longest contiguous subsequence with the value (e.g., the longest spell of UN). An index is closely connected with a particular perspective in which states, their order, switches and concentration (lengths of spells) have a particular meaning. We regard neither any perspective as superior to others, nor any single index as superior to others for a wide range of perspectives, and assess merely the appropriateness of an index to a particular perspective, the elucidation of which is an important input into the analysis. In our perspective, the harm caused by UN to the well-being of the individual, to the labour market, the economy, and the society at large, accumulates out of proportion with the length of the UN spell; two consecutive months of UN amount to more harm than two isolated months of UN separated by one or several spells in other states, with at least one of them in EM. The 'harm' of UN persists past its spell, so two UN spells separated by only a short period in another status cause more harm that two spells separated by a longer period. Further, only EM in this period reduces the harm, and does so only gradually; other states (IA in particular) do not change it. See Shorrocks (2009) for a related discussion.

This perspective motivates the following scheme for summarising a sequence. First we assign a score s_t to each element $t = 1, \dots, T$ of the sequence. Starting with $s_0 = 1$, we form s_{t+1} by adjusting s_t for the previous state x_t; increasing it if $x_t =$ UN, decreasing it if $x_t =$ EM, and leaving it intact otherwise. The simplest way of adjusting s_t is by multiplying it by a factor that depends on the status.

In Section 4.2, we apply the factors $\mathbf{f} = (1/1.1,\ 1.1,\ 1.0,\ 1.0,\ 1.0)$ for the respective states EM, UN, IA, T and AB. Instead of a multiplicative factor, an additive term can be used, and more complex schemes can be devised. A (complex) formula for s_t could be replaced by a look-up table.

In principle, the vector $\mathbf{f}$ may be specified as a function of the entire history (the sub-sequence from month 1 to t), and of the states in the previous two months in particular. Also, the factors may depend on variables defined for the individuals, the time point (season) and time-varying covariates. There may be a rationale for different scoring for young and older members of the labour force, as well as for men and women. A score is truncated from below at 1.0 — that is the minimum harm done by an isolated month of UN. We also assume a maximum score, which is attained after about ten months of UN. That is, from ten months on, each subsequent month of UN entails equal harm.

The index for a sequence is formed by adding up the scores for the months spent in UN. In a more general scheme, a weighted total of the scores is formed, with the weights determined by the current status. No generality is lost by setting the weight for EM to zero, but IA can be associated with a positive weight. The weight for UN should be higher than for any other status. Further generalisation is achieved by defining the weights in a more complex way. For example, the weights may be adjusted for the recency, to reflect the greater importance of the later time points, or for working in the winter months. In summary, we define the class of indices by a scoring scheme and a summary of the scores. First, a sequence of states $\mathbf{x}$ is converted to a sequence of scores $\mathbf{s}$, using a vector of status-specific factors $\mathbf{f}$ or, more generally, functions $F(s, x)$ that define the scores recursively as $s_{t+1} = F(s_t\, , x_t)$. The index is formed from these scores by the formula

$$g(\mathbf{x}) \;=\; \sum_{t=1}^{T} w_{x_t} s_t \tag{1}$$

for a suitable vector of weights $\mathbf{w}$ defined for the states. Some trivial summaries, such as the number of months in UN, or not in EM, are special cases of this scheme. The length of the longest spell of UN cannot be expressed in the form (1). The index can be standardised, to be in the range (0, 1) or (0, 100), by dividing it (or the weights) by the maximum possible value attained by a person who is in UN throughout the studied period. Comparisons of the values of the index and its summaries (e.g., averages) are appropriate only for periods of equal length and with equal frequency of records (e.g., months).

The index defined by (1) satisfies the axioms of Mendola, Busetta and Milito (2011). Their scoring scheme caters for perspectives similar to ours and has considerable flexibility, but we regard it not suitable for our context of long sequences because the scoring involves operations on pairs of time points, of which there are $T(T-1)/2$, in our case 1485. The axioms formulated by Mendola *et al.* (2011) can be adapted to our setting. With appropriate choice of the coefficients f_s and w_s they are satisfied.

Although an index is defined for individuals, its use is principally for summarising sets of individuals (subpopulations), and averaging is a key device in forming such summaries. An important property of an index is *additivity*. An index is said to be additive if a pair of its values v_1 and v_2 amounts to the same quantity of the measured attribute as any other pair v_3 and v_4 for which $v_1 + v_2 = v_3 + v_4$. Without additivity, averaging is not meaningful. An assessment as to whether an index is (close to being) additive requires a judgement by an expert, relating the scoring scheme to the perspective. The data contains no information that would contribute to making such a judgement.

A score should have the property of continuity or *inertia*, that is, it cannot change abruptly from one month to the next. The index should have a similar property; two sequences that differ only in one month should have similar values of the index. The material and economic state of affairs of an individual does not change suddenly from one month to the next as a result of a change of the labour force status, even though the prediction (anticipation) of changes in the future further ahead may be altered substantially. In the score we consider, we want to discount any psychological and related factors (shocks) that have a transient nature, even if they may have a strong effect in the short term.

Any index is a simplification (reduction) of the data, and the data is often a simplified representation of the reality. Several concerns arise in our context. First, some short-term and part-time employment, such as arranged by ALMP, cannot be regarded on par with employment of full value, as intended to be measured by the index. Next, economic inactivity covers a wide range of circumstances which should be reflected in an ideal scoring scheme. Education and training may be regarded as constructive for the person's labour market position, and would reduce or diminish the disadvantage accumulated by earlier unemployment. Electing to look after the family instead of pursuing a professional career is in general difficult to judge and representing the choice equally for everybody is too simplistic. Similarly, (old-age) retirement, when its election is prompted by (anticipated) failure to find a new job, could well be regarded as a form of unemployment, because the skills and expertise of the new retiree are lost to the economy and a new demand is presented to the social security system. We note that the benefits received during a spell of UN may have an impact on the labour market behaviour (job search) of the individual. We study this issue in Section 4.5.

We have no details about the spells of IA. Some information can be gleaned from the persons' background variables, age and sex in particular, but in the narrow context of an unemployment agency, success is measured by how rapidly unemployment case files are closed (cases resolved) and how much time passes before the workers involved return to open new case files.

3.3 Comparing two populations

For any comparison of two populations (or samples), we consider two modes. In the *raw* mode, we compare a set of summaries for the two samples in their entirety. In the *matched* mode, we use background variables to select matched pairs of units (sequences), with one unit from each sample, and compare these two matched subsamples. Matching is applied extensively in causal analysis with observational data (Rubin, 1973, 2006). In some of the analyses we do not have a 'causal' agenda, but regard matching as a way of coming closer to the ideal of comparing two populations *like with like*, as if they did not differ in the distribution of their backgrounds. Multinomial regression is an alternative to this approach to adjusting for background. We prefer matching because the assumptions of a regression are often problematic, model-selection issues are rarely attended to with integrity and the results of multinomial regression are often difficult to interpret. See also McCullagh (2008) for a discussion of some related issues.

A practical advantage of matched-pairs analysis is that it can conclude with one or several summaries of the pairwise comparisons of the units. With the state space comprising K states, we have $K(K-1)$ distinct (ordered) pairs of states. Their frequency within matched pairs can be tabulated and these tables summarised across the pairs. This is ineffective when there are many possible pairs (e.g., for $K > 4$). However, often only some pairs of states are of interest, such as employment–unemployment (EM–UN) and UN–EM. Their monthly balance, defined as the monthly differences of the number of pairs EM–UN and the number of pairs UN–EM, is of obvious interest. We note that measures of distance between two sequences (McVicar and Anyadike and Danes, 2002; Studer *et al.*, 2011) are not useful in this context because we want to treat a pair of states as ordered, and distinguish a pair from its reverse.

4 Analysis

We study the set of 62 720 sequences for the individuals represented in ADEM records in the period January 2007 – July 2011 (55 months). Table 1 displays their monthly summaries. It contains no information about the spells because the link between any two consecutive months of an individual is not maintained. The table may be collapsed to quarters or years; then it is easier to comprehend, but

Table 1: Number of individuals in the labour force states by month; ADEM registrants.

Month	EM	UN	IA	T	AB	Month	EM	UN	IA	T	AB
2007											
1	23997	8851	0	1835	28037	7	26895	7404	4418	1466	22537
2	24316	8725	1137	1774	26768	8	26727	7638	5093	1379	21883
3	25224	8378	1877	1608	25633	9	27492	7690	5454	1476	20608
4	25833	7997	2603	1534	24753	10	27830	7905	5755	1607	19623
5	26395	7626	3234	1439	24026	11	27867	8100	6092	1677	18984
6	26803	7432	3722	1461	23302	12	27613	8170	6743	1593	18601
2008											
13	27536	8503	7091	1581	18009	19	30094	7531	8534	1654	14907
14	27904	8419	7310	1638	17449	20	29835	7872	9116	1487	14410
15	28714	7881	7604	1578	16943	21	30422	7981	9203	1780	13334
16	29173	7813	7707	1597	16430	22	30437	8478	9410	1805	12590
17	29630	7582	7957	1574	15977	23	30011	8798	9860	1936	12115
18	30083	7427	8155	1576	15479	24	29280	9410	10268	2040	11722
2009											
25	28386	10670	10506	1955	11203	31	29135	10458	11984	2240	8903
26	28223	10920	10774	1977	10826	32	28606	10860	12696	2015	8543
27	28504	10573	11059	2181	10403	33	29198	10980	12438	2344	7760
28	28806	10534	11163	2152	10065	34	29191	11439	12498	2497	7095
29	29063	10448	11430	2049	9730	35	28981	11724	12742	2585	6688
30	29221	10328	11667	2130	9374	36	28513	12197	13122	2550	6338
2010											
37	28188	12874	13463	2206	5989	43	30675	11018	14551	2559	3917
38	28352	12868	13622	2256	5622	44	30210	11386	15206	2378	3540
39	29113	12029	13788	2537	5253	45	30710	11373	15041	2661	2935
40	29963	11658	13806	2334	4959	46	30766	11774	15019	2677	2484
41	30427	11160	14135	2324	4674	47	30626	12391	14929	2719	2055
42	30705	11006	14285	2386	4338	48	30160	12812	15514	2451	1783
2011											
49	30410	13009	15775	2047	1479	55	31828	11381	16780	2731	0
50	30746	12735	15990	2006	1243						
51	31477	12087	16155	2069	932						
52	32093	11736	16196	1982	713						
53	32430	11340	16454	2021	475						
54	32523	11202	16721	2002	272						

Note: EM — employed; UN — unemployed; IA — economically inactive; T — in transition; AB — not in the records for the month nor earlier.

some information is lost by the aggregation. Similar summaries can be compiled for subpopulations, such as age groups.

In total, there are $62\,720 \times 55 = 3.450$ million person-months; 1.597 million are in EM (46.3%), 0.551 million in UN (16.0%), 0.578 million in IA (16.8%), 0.110 million in T (3.2%) and 0.614 million (17.8%) AB, that is, with an incomplete (or no) record in IGSS prior to their first registration in ADEM at some point in January 2007 – July 2011. Of course, these figures do not reflect the UN and IA rates in the population because of the absences counted and the continually employed who have no ADEM records. Persons who were not unemployed in any month in the period of the study have no record in ADEM for that period, so they are not included in the dataset. However, persons who were in transition in the first or last month of the study are included because their case files were open at the time. Our dataset has the form of a cross-section. For some analyses, we select from it cohorts defined by the start of a UN spell or a similar attribute.

The solid lines in Figure 2 reproduce the information in the table graphically. The diagram shows that the number of employed in the records (blue) has been growing, except for year 2009 and a few months following it. The number of unemployed (red) tends to rise in winter and decline in summer. The number of economically inactive persons in the ADEM records has been increasing steadily, with small peaks in September of every year, which are followed by slower increases or even small declines. The number of individuals in transition has been rising steadily, except for a drop in the first half of 2011. Absence as a status can be entered only in month 1 and the status is left in the month of the first registration. That is why the number of persons absent from the database, who become present later, is steadily decreasing. The (downward) slope of the curve for status AB can be interpreted as the number of new entries in the records. Short-term increase in UN tends to be accompanied by lower count of employed. The curves are slightly smoothed to remove some sharp edges. Smoothing is performed by averaging each observation with its neighbours. We use the scheme

$$\dot{y}_t \;=\; b y_{t-1} + (1-2b) y_t + b y_{t+1} \tag{2}$$

with $b = 0.25$; $\dot{y}_t$ is the smoothed version of y_t. The first and last value, y_1 and y_T, are not altered. The official UN rate is derived from the monthly count of unemployed prior to smoothing, but is it adjusted for the season of the year.

The changes from one month to the next can be summarised similarly, although there are many more categories, 22 altogether. Five of them are for no change, such as EM – EM. Their counts for the five states are represented by dashed lines in Figure 2. Except for status T, changes from one state to another are much less common than staying in a labour force status — there is a lot of inertia in the

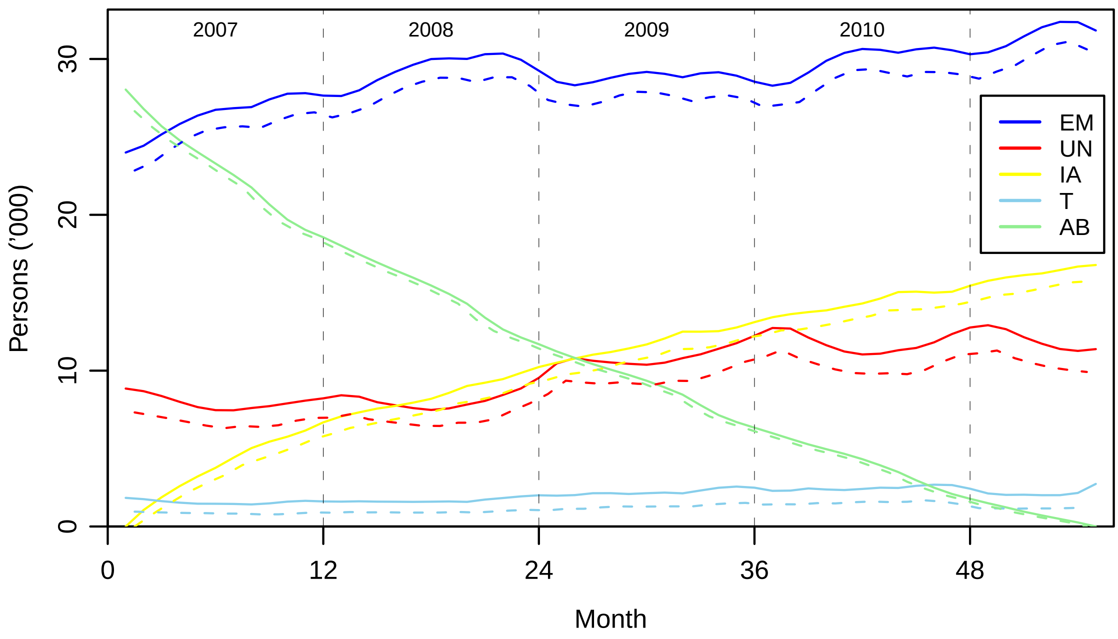

Figure 2: Numbers of individuals in the labour force states by month (solid lines) and the numbers of persons staying in the same labour force status (dashes).

states. That is why the dashes in the diagram closely follow from below the trend of the solid lines of the same colour.

A graphical summary of some of the most frequent changes of labour force status is presented in Figure 3. The counts of switches from EM to UN are drawn in Panel A by black, and the counts of the 'reverse' switches, from UN to EM, by gray colour. Thin lines are used for the exact numbers and thick lines for their smoothed versions which better reflect the longer-term trend. The number of individuals who switch to EM exceeds the number of those becoming unemployed, but the two curves have similar features. However, UN is not decreasing as much as the difference of the two curves suggests, because the pool of unemployed is replenished from other states, including AB. Job losses reach their peak at the end of every year, but they are preceded by peaks in the switches to EM.

The counts for changes between EM and IA have a regular annual pattern. Retirements (EM–IA) have a deep trough every late winter and spring and increase sharply in summer. The numbers of switches IA–EM have peaks in February and March of every year, after declines at the end of the previous year. There are much fewer switches from IA to UN than in reverse; the latter have peaks in January and troughs in summer. The switches from AB are highly irregular. This may be affected

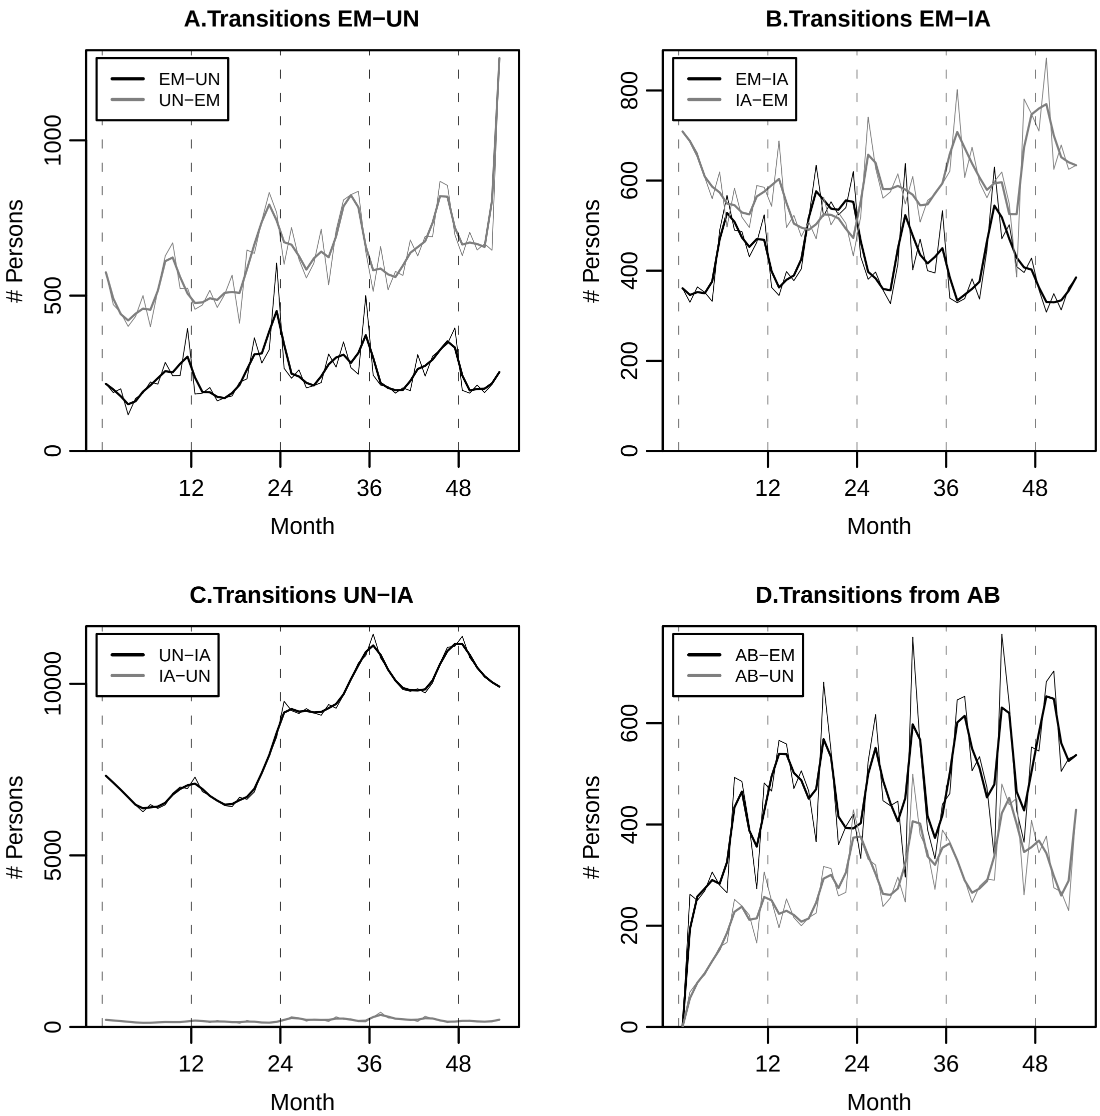

Figure 3: The monthly counts of changes of employment status. Thin lines represent the exact counts and thick lines their smoothed versions.

by our rate of success in recovering information about individuals prior to their first registration in ADEM. Note in particular that many switches AB – EM are made without being recorded by ADEM.

All the summaries presented in this section require a careful deduction for making any meaningful conclusions about the cases handled by ADEM and their outcomes. A complete summary of the data can be provided only by the sequences of states for every individual on record or for a random sample of sufficiently large size.

4.1 The labour force carpet

This section presents a complete graphical summary of the sequences of persons with ADEM records. Each person is associated with a horizontal segment of length $T = 55$ comprising sections coloured according to the labour force status at the time. A small example was given in Figure 1. We refer to the diagram as the *labour force carpet*. When drawn for many sequences, it is essential to order the sequences so that many features of their distinct patterns would be transparent and would enable us to formulate substantive hypotheses and conclusions. For example, it is desirable to have large patches of the same status, representing long spells in the status in the same period of time. There is no such unique ordering. After trial and error, we selected the following. At the bottom of the carpet, we place all the sequences that have majority postamble in EM, followed by those with majority preamble in EM. They are followed by the same majority patterns for UN, IA, T and AB. These patterns account for about 90% of the sequences. The remainder is ordered by the status of the longest mid-spell with majority in EM, UN and IA. Within the patterns of majority pre- and postamble, the ordering is by the length of the pre- or postamble, as applicable. More detailed ordering is implemented within these composite patterns, but it has only slight impact on the overall appearance of the carpet. The segments are drawn in Figure 4 only for a simple random sample of 10 000 sequences. The uncertainty associated with sampling can be ignored. The individual segments are drawn by very thin lines; there are no vertical gaps between the segments.

A spell of EM in the carpet ends with a switch to status T in about half the cases (47.9%), to IA in about a third (32.9%) and to UN in about a fifth of the case (19.2%). New arrivals (switches from A) enter EM and UN in about equal measure, but more recent switches are almost exclusively into UN. This is linked to economic crisis and diminished opportunities for employment for the new entrants into the labour force, even though Luxembourg is not affected by the current crisis as strongly as most other countries of the European Union.

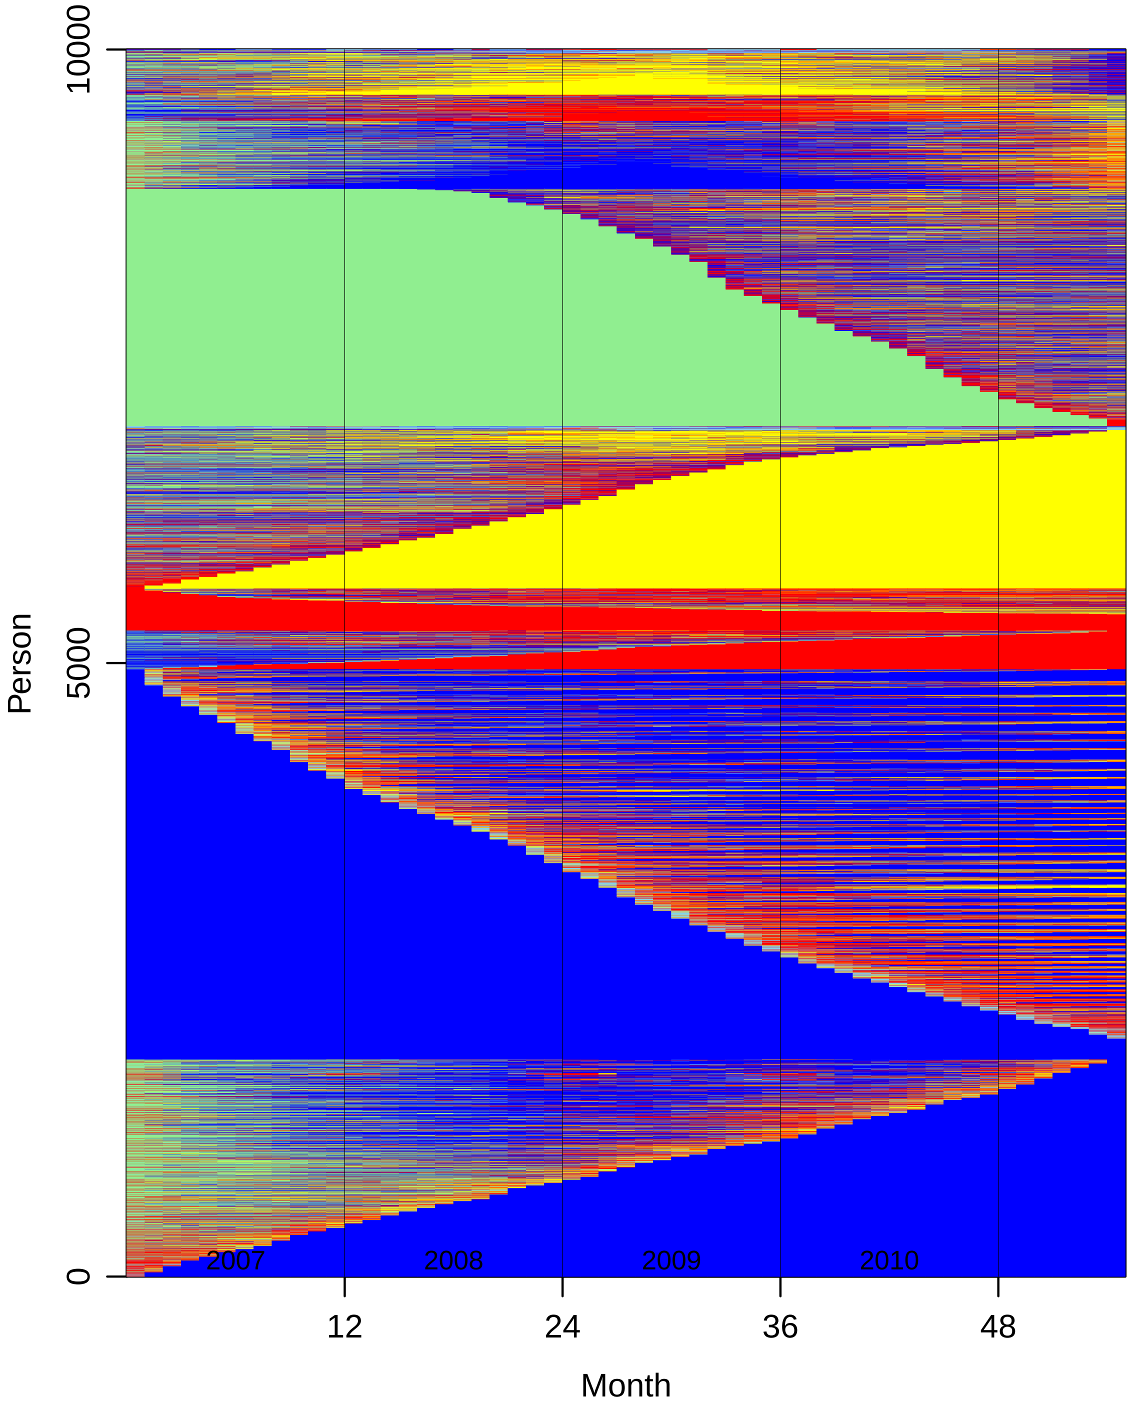

Figure 4: The labour force carpet for the ADEM records. Based on a simple random sample of size 10 000 sequences (16%). The states are: employment (EM) — blue; unemployment (UN) — red; economic inactivity (IA) — yellow; transition (T) — skyblue; absence (AB) — light-green.

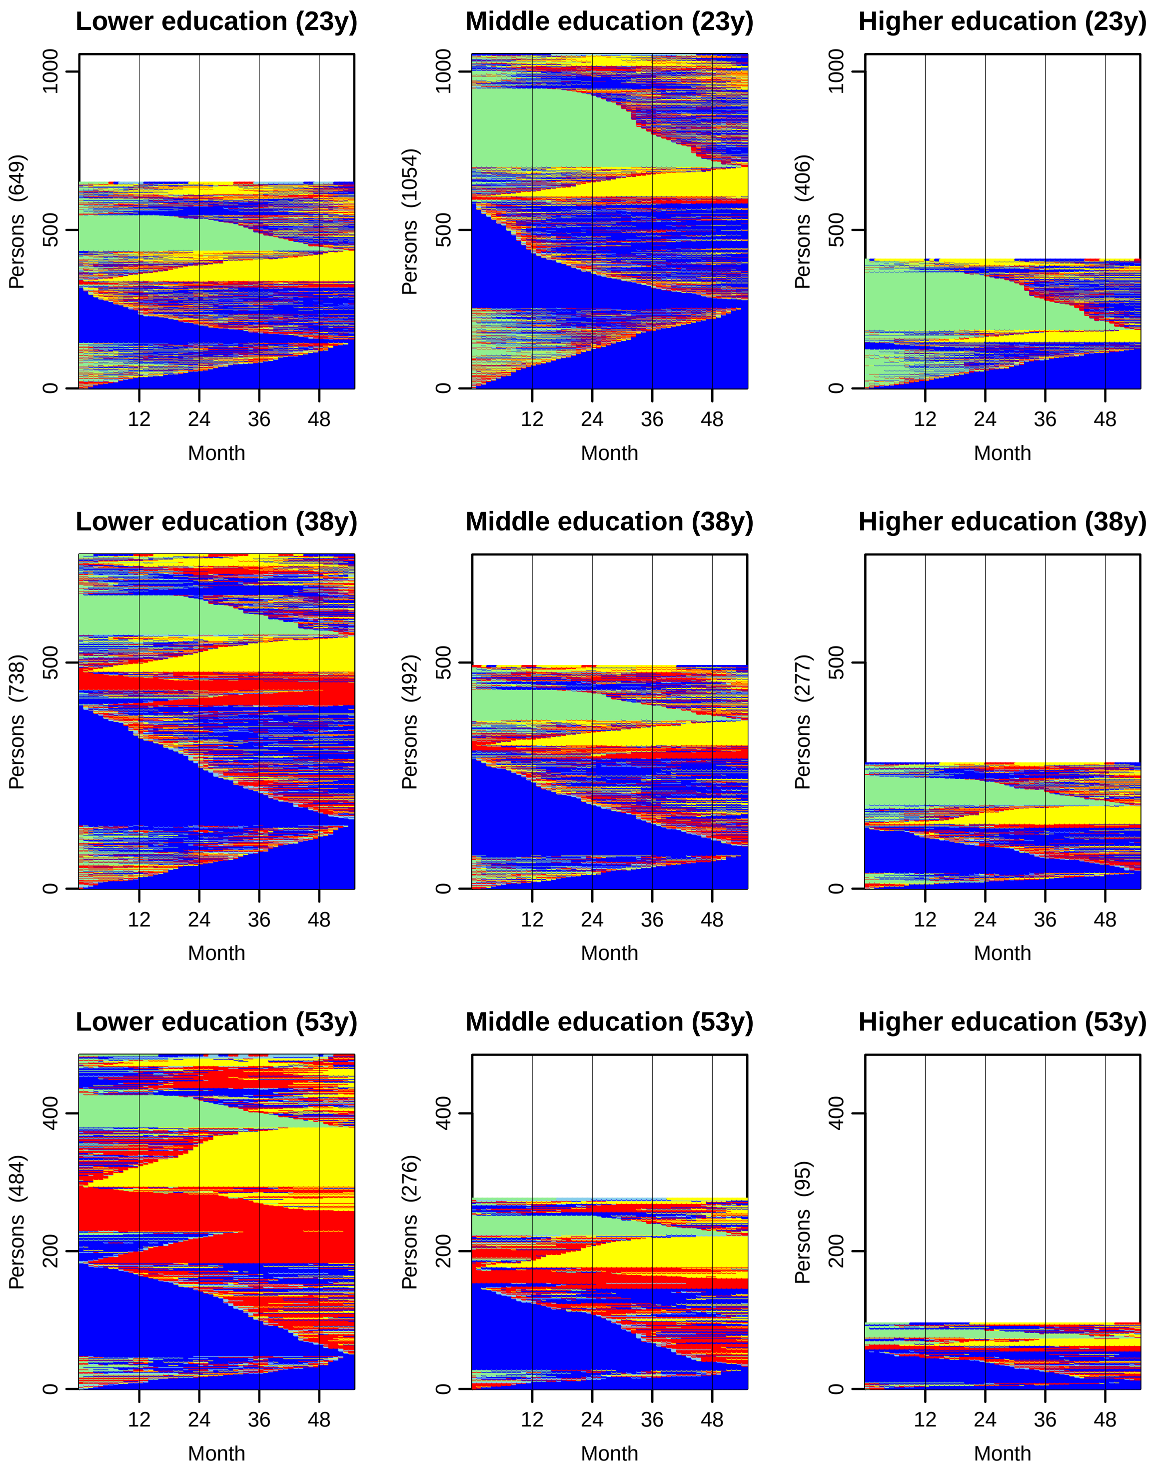

Figure 5: The labour force carpet for three age groups and three levels of education.

The labour force carpet can be drawn for various subpopulations, defined by age, sex, education and employment history. Comparison of a set of carpets is facilitated by using the same ordering in each (sub-)carpet and by using identical vertical scales. Figure 5 displays a set of carpets for ages of 23, 38 and 53 years (in 2007; rows) and lower, middle and higher levels of education (columns). The carpets in a row have identical vertical scales. Among the young, the top row, those with mid-level of education are in a majority, while among the older persons, and the 53-year-olds in particular, those with lower level are in a majority. Among the young there are very few sequences with long spells of UN (large red patches are absent). There are many such sequences among the 53-year-olds with lower level of education. Among the young, many return to EM after a short spell of UN or IA (blue colour dominates in mixed patches), whilst among the 53-year-olds long-term UN after a spell of EM (red patches) is quite common.

We emphasise that the persons with a record in ADEM are not a good representation of the country's labour force. Young are over-represented in the records, but their records tend to be short (case files closed after switching to EM). Older workers are under-represented, but their case files tend to be open for longer periods of time, owing to long spells of UN.

4.2 Vulnerability index

For vulnerability in the labour market (the accumulated harm done by UN), we adopt the index defined by (1) with factors $\mathbf{f} = (1/1.1, 1.1, 1, 1, 1)$, weights $\mathbf{w} = (0, 1, 0, 0, 0)$ for the respective states EM, UN, IA, T and AB, and maximum score $M = 2.5$, as described in Section 3.2. Each month of a delineated period (not necessarily the entire period of 55 months in our case) is associated with a score. For every month in the first spell of UN the score is multiplied by a 1.1. So, a spell of four months of UN results in the sequence of scores (1.00, 1.10, 1.21, 1.33) and after ten consecutive months of UN in a single spell the score reaches $1.1^9 = 2.358$. If the person is still unemployed in the next and following months, the score remains at 2.5 for each month, because $1.1^{10} > 2.5$. The choice of this maximum is subjective. If the status in a month is IA or T, the score is copied from the previous month without any change. For every month of EM, the score is reduced by the factor 1.1. Thus, after ten months of EM, or earlier, the score reaches the minimum of 1.0. The index is defined as the total of the scores over the months of UN in the designated period. Thus, an upper limit on the score is the 2.5-multiple of the number of months. This limit is not achieved by anybody when the score in the first month, when no employment history is available, is set to 1.0.

Figure 6 displays the histograms of the vulnerability indices for persons in ADEM aged 23, 38 and

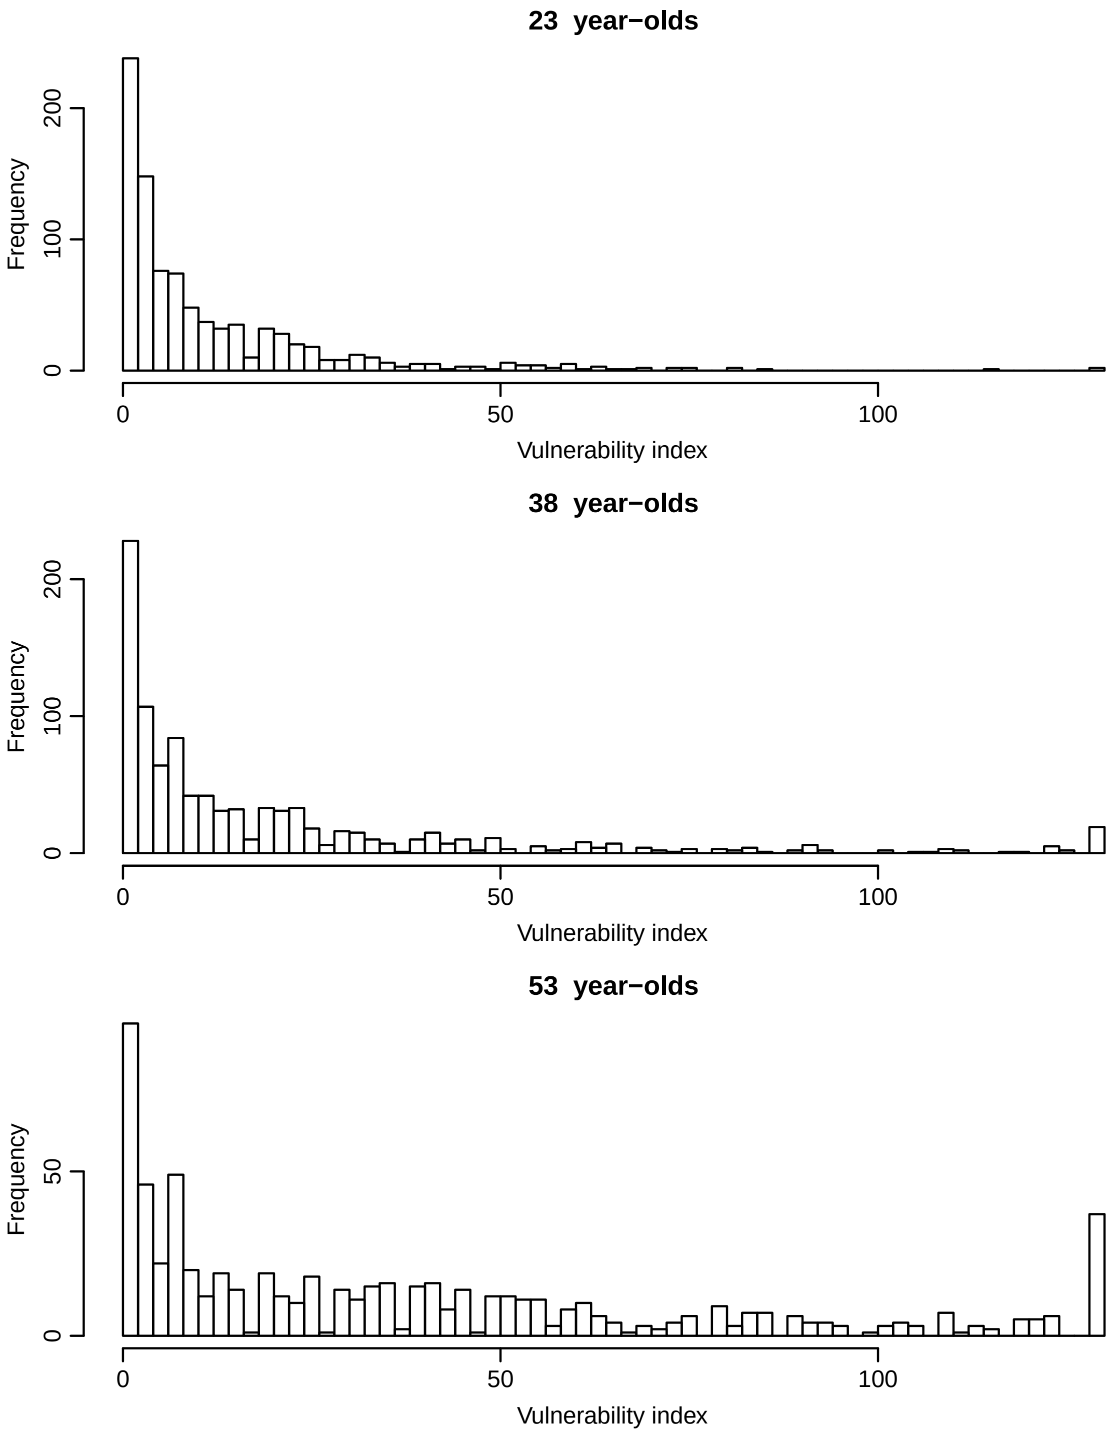

Figure 6: Histograms of the vulnerability index for 23-, 38- and 53-year-olds registrants in ADEM.

53 at their first registration in ADEM. The maximum value of the index, 128.44, is attained by those who were unemployed throughout. It is quite frequent among the 53-year-olds (5.58%), and very rare among the 23-year-olds (0.22%). No (statistical) distribution adequately describes the values of this index, although a mixture of a few distributions, one of them degenerate (for the maximum) may fit it adequately.

An index can be applied to a part of the period for which the records are available. In our case, the records are for the period from January 2007 to July 2011. If we define the index for the period July 2007 – June 2011 (4 years), the scores are calculated by starting in January 2007. Then the first score that is counted, in July 2007 for a person unemployed at that time, depends on the past employment history and may be greater than 1.0. A similar argument can be presented for adjusting the score depending on the conclusion of the record (in June or July 2011). If the record is concluded with UN, it is more likely to continue with UN, but this would not be reflected in the index defined by (1).

This issue can be addressed by imputation for the (future) scores. In imputation we distinguish recipients, whose records and scores are supplemented, and donors, whose sub-sequences of scores are inserted into the recipients' records. Matching is the general term for finding a donor for every recipient. For a (recipient) record with a particular sub-sequence of states in 2011 (January – July), we find a match among the sub-sequences in January – July 2010, and impute for the recipient's record in August – December 2011 the donor's subsequence from August – December 2010. Donors are usually assigned to recipients without replacement; that is, a sequence is a donor for at most one recipient. In our setting, there are often more recipients than donors, so we have to break this rule. With a sequence extended by imputation, we evaluate the index for the entire period July 2007 – December 2011. In principle, we could impute also for the past employment history, e.g., for January – June 2007, using donor subsequences from January – June 2008.

We emphasise that the parameters of the index used in our example, and in Figure 6, have not been set by any profound deliberation or elicitation from experts. In an application, sensitivity analysis, comparing the conclusions based on parameters with similar values, should be conducted. For example, the conclusions are not changed materially with the setting $\mathbf{f} = (1/1.2, 1.2, 1, 1, 1)$, $\mathbf{w} = (0, 1, 0.4, 0.2, 0)$ and $M = 3$, in which the score increases faster during a UN spell and decreases faster during an EM spell, but spells of IA and T also contribute to the index, with respective weights 0.4 and 0.2. Although M is increased, the maximum score is reached after only six months of UN. ($1.2^6 \doteq 3.0$). Details of the results are omitted.

4.3 Comparing two groups

A generic inferential task with ADEM records is to compare two subpopulations. We identified in Section 3.3 two versions of this problem, raw and matched. In the first, we compare the values of a summary for one and the other group. If no summary is adequate, the carpets for the two groups can be compared informally, by placing them side-by-side. Such a comparison can be made more formal by the following procedure. Suppose first that the two groups have the same number of individuals. Then we apply the ordering used in the construction of the carpet, and compare the individuals with the same rank order pairwise. We tabulate pairs of states by the two codes, and summarise these pairs (note the affinity with tabulation of changes from one month to the next), or define a measure for the distance between the two sequences. Such a distance should be directional, distinguishing between dominance of pairs UN–EM over EM–UN and the like.

When the two groups have unequal sizes, we select a simple random sample from the larger group to match the size of the smaller group. This is nearly equivalent to pro-rating the rank order of the smaller group to match it with the larger group. For example, if the ratio of the sizes is r, then the first sequence from the smaller group is paired with the kth sequence from the second group for $k < r$ selected arbitrarily; sequence m is paired with one of the sequences $[(m-1)r] + 1, \ldots, [rm]$, where $[\,]$ denotes the floor (the largest integer smaller than its argument). The choice between the integers can be made at random, with the probabilities proportional to the length of the representation in the range $\{(m-1)r, rm\}$.

If an individual-level (sequence-level) summary is defined, then a comparison of two groups can be based on this summary. An index, such as the vulnerability index, is an obvious candidate. Individuals in the two groups can be matched on their past employment records and the summaries thereof, and other attributes, age, sex, level of education, and the like, also can be used. It is preferable to match sequences using a lot of matching information, to satisfy a higher standard of matching 'like with like'. In that case, propensity score matching (Rosenbaum and Rubin, 1983, and Rubin, 2006), provides a practical solution. Although propensity matching is usually applied in causal analysis of observational data, our inferences do not relate to any causes because the membership of the groups we compare cannot be manipulated, nor do we want to make references to any hypothetical manipulation of it.

The propensity with respect to a division of the sample to two (treatment) groups is defined as the conditional probability of being in one group given a set of covariates. Random assignment corresponds to (known) propensities, equal to a constant different from zero and unity in each group. In a typical observational study, the fitted propensity scores are in a wide range in both groups, and

Table 2: Covariates used in matching.

Variable	Type	Values	Notes
Age	Cont.	Years	Range 15 – 71; median 32; mean 33.9
Level of education	Categ.	1 – 3	Lower (41.0%), middle (38.5%) and higher (17.8%); missing (2.7%)
Civil status	Categ.	1 – 3	Single (49.6%), married (incl. remarried; 40.8%), and other (divorced, separated or widowed; 9.2%), missing (0.4%)
Sex	Binary	0, 1	Women (46.3%), men (53.7%)
Nationality	Categ.	1, 3, 7	Luxembourgois (33.8%), Portuguese (29.8%) and other (36.4%); the largest contribution to 'other' is French (8.4%)
Date of completing education	Cont.	Years since 1960	Range $(7.6, 50.9)$, median 21.5, mean 21.4, st. dev. 0.6
First record in IGSS	Cont.	Years since 1960	Range $(-0.3, 52.0)$, median 40.8, mean 35.0, st. dev. 15.4

these ranges (score distributions) differ, but have a substantial overlap.

Example

We compare the labour force states in months 25 – 55 (over the period of 31 months) of the subjects in ADEM younger than 40 years and those aged 40 or over when they registered in ADEM. As covariates, we use the labour force states in the preceding twelve months (months 13 – 24), and the covariates listed in Table 2, except for Age. To relate the age group, a dichotomy, to the covariates, we use logistic regression. Model selection issues (quality of the model fit) are not relevant because the sample size, 62 720, is substantial, and the degrees of freedom lost by including a few covariates with small (or zero) coefficients are inconsequential. The purpose of the propensity scores, obtained by the regression, is to achieve a *balance* of the distributions of the covariates in the matched pairs, see below. That is why we are liberal with including covariates in the propensity model.

The fitted propensity scores are summarised by histograms for the two groups in Figure 7. Among those below 40 years of age, a substantial fraction of individuals has very small propensity, indicating that their configurations of the background (including labour force states in 2008, the preceding year) are very rare among those aged 40 or over. Nearly half of the younger group (20 530, that is, 48.65%) have fitted propensities smaller than 0.02. Their configurations of the values of background variables (background profiles) are very rare in the older group.

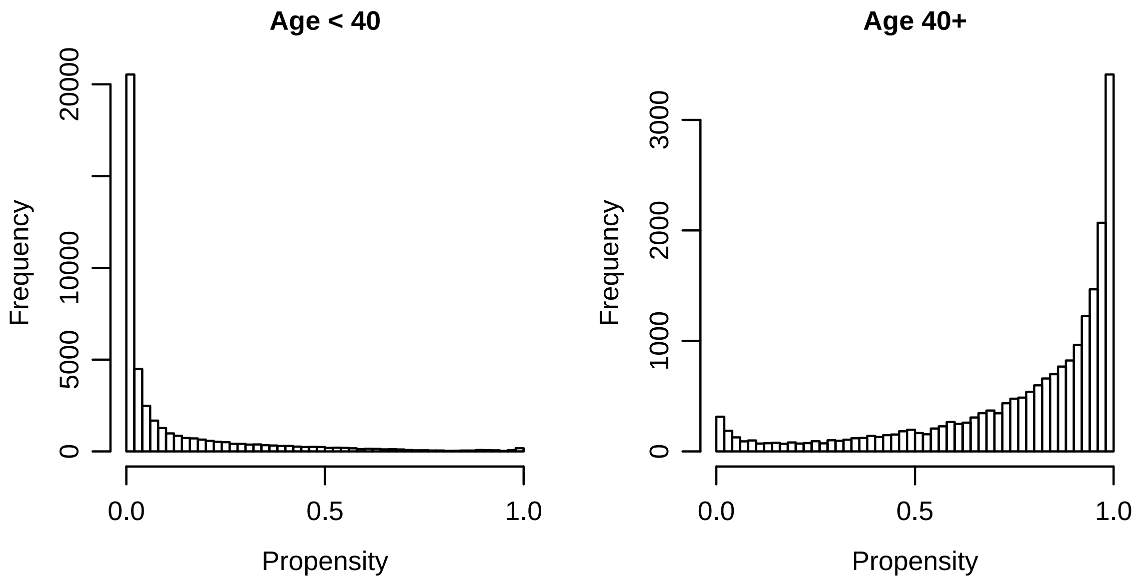

Figure 7: Fitted propensity scores for individuals in ADEM aged below and above 40 years of age.

In the histogram for the older group, many individuals have high propensities; their backgrounds and configurations of labour force states in 2008 are rare in the younger group. The distribution of the propensities appears to be not as extreme as for the younger group; 3456 (16.8%) of them have propensities greater than 0.98. Only about 6% of individuals in the younger group have propensity greater than 0.50, whereas nearly 15% of the older group have propensity smaller than 0.50. Such a comparison is misleading, however, because the mean propensity is equal to 32.7%, the percentage of those in the older group; 11.5% of the younger group have propensities greater than the mean, whereas only 8.5% of the older group have propensities smaller than the mean. These individuals can be regarded as having unusual configurations of the covariates, relatively more frequent in the other group.

The propensities are used to form matched pairs of individuals, one from each group in a pair. We classify the propensities to 20 intervals, $(0.0, 0.05), (0.05, 0.10), \ldots, (0.95, 1.0)$. The number of intervals is set so that there would be sufficiently many individuals in each interval. Usually fewer intervals are used, but we have relatively large samples. Suppose $n_{k\mathrm{Y}}$ and $n_{k\mathrm{O}}$ individuals from the younger (Y) and older group (O) fall into propensity interval k. If $n_{k\mathrm{Y}} \leq n_{k\mathrm{O}}$, then we pair every younger individual in propensity group k with a randomly selected older individual from the same propensity group; thus, we form $n_{k\mathrm{Y}}$ pairs. If $n_{k\mathrm{Y}} > n_{k\mathrm{O}}$, then we reverse the roles of the two age groups and form $n_{k\mathrm{O}}$ matched pairs. The units to be paired are selected at random and the pairs are formed also at random, with no regard for their fitted propensities other than belonging to the same interval. This is substantially different from the nearest-neighbour and related methods. We obtained 5386 pairs, so $100 \times (1 - 2 \times 5386/62720) = 82.8\%$ of the sequences were discarded as unpaired.

The phenomenon of failing to find many matches for two groups is related to the common support problem (Lechner, 2002), in which certain combinations of background variables are not possible in some treatment groups. In our case, the number of years since completing education has an upper bound in the younger group, and most subjects in the older group have completed their education much earlier. Therefore, most of the matches are formed by the oldest subjects in the younger group and the youngest in the older group. The concern arises that these matches represent subpopulations that differ substantially from the populations we intended to compare originally. However, the matches still represent the best pick from the two groups in terms of the distribution of the auxiliary variables. The problem is not in the decision to match, but in innate lack of comparability of the groups, and the reliance on the auxiliary variables is more onerous.

Next, we draw the labour force carpets for the two groups and the months $25-55$, side-by-side, to assess their differences among the matched pairs. For completeness, we add the carpets for the entire groups, for which we apply simple random sampling, so that the sequences are represented by a sample of size 10 000, about 6700 from the younger and 3300 from the older group. The resulting diagram is drawn in Figure 8. It confirms that there are relatively more new arrivals to the database (first-time unemployed) among the young (more dominant yellow patches), long-term unemployment is more frequent among the older workers (a relatively wider horizontal red strip) and loss of employment is followed by long-term unemployment more often among the older workers (more red colour to the right of the blue patch).

Even though it is obvious that matching on the background reduces the differences between the two age groups substantially, detailed comparison of the carpets within rows requires finely trained judgement. The carpets can be summarised by comparing the monthly composition of the states for months $25-55$. This is done in Figure 9. In the top left-hand panel, the number of pairs in which the 'young' member is employed but the 'older' one is unemployed is plotted, as a function of the month, by black colour. The gray curve is the corresponding function for the pairs in which the young member is unemployed and the older one is employed. At the beginning, in early 2009 (month 25), there are more pairs with the young member of the pair unemployed, but after month 36 (early 2010), many more young than older members are employed. The two curves have similar shapes, although

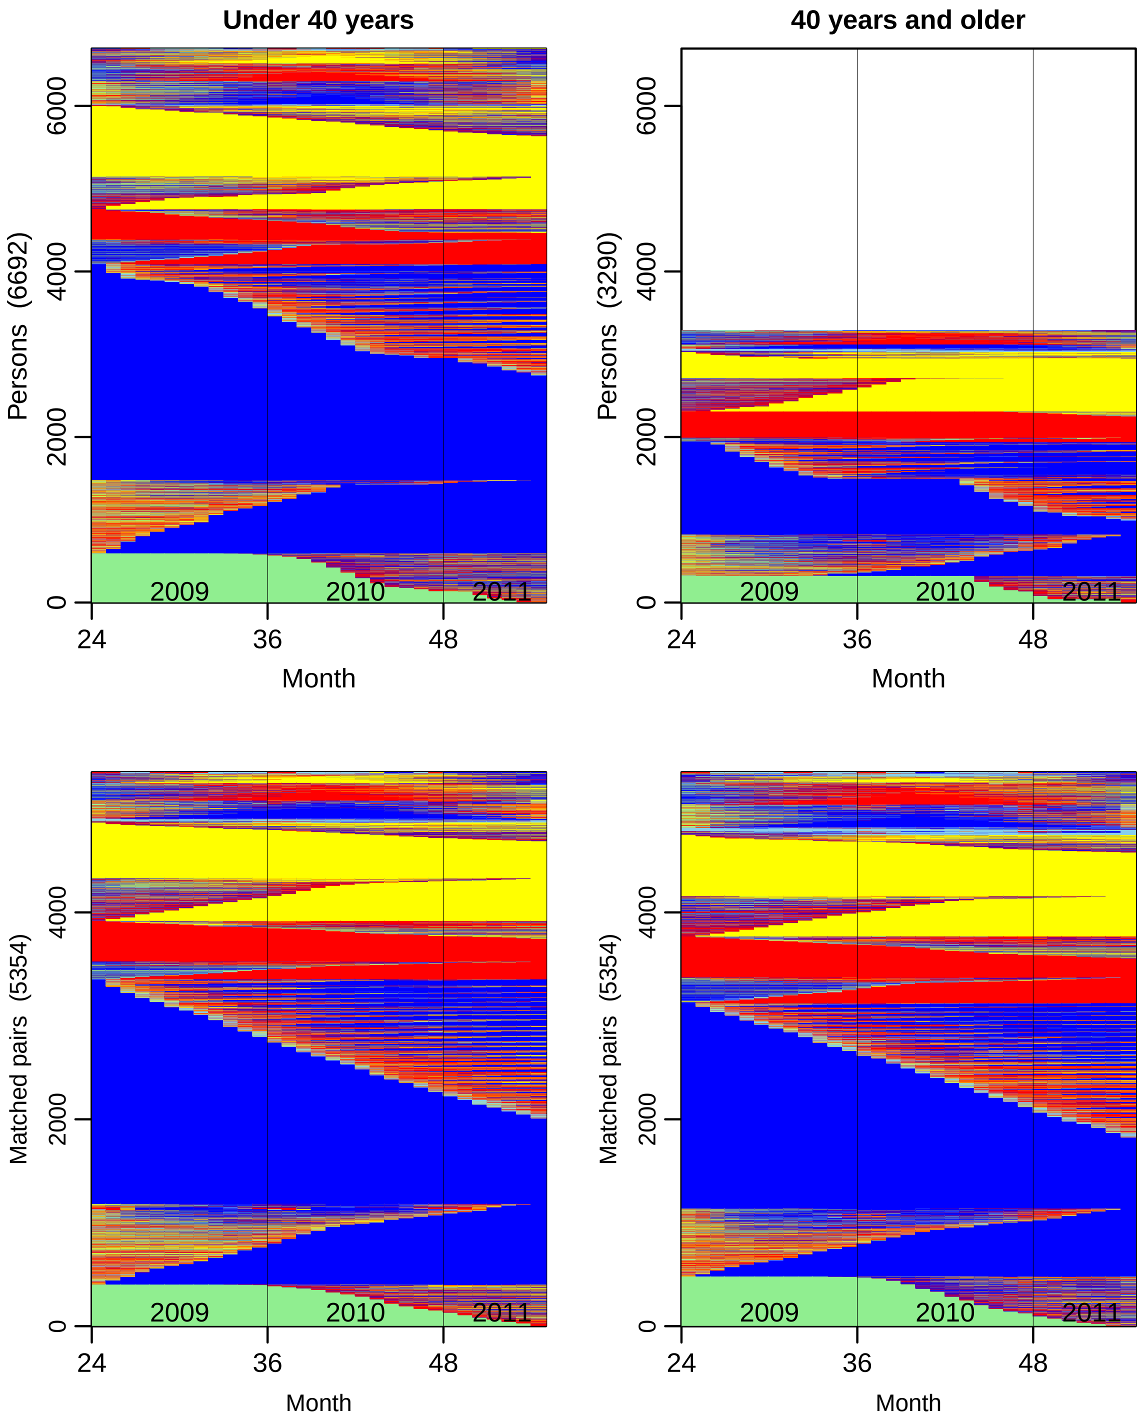

Figure 8: Labour force carpets for the entire groups of younger and older individuals in ADEM (top panels), and the subsets of matched pairs (bottom panels).

their difference declines in 2011. From these comparisons we conclude that in 2009 and 2010 older unemployed are less likely to find employment than younger unemployed, when matched on their background. This does not directly imply that long-term unemployment is a greater problem among the older workers, because the comparisons in Figure 9 are made separately for each month.

The number of pairs with pattern EM−IA, with one member employed and the other inactive, increases throughout the period, but the differences between the two counts are small, except perhaps in late 2010 and early 2011 when more young matched workers are employed. The number of pairs EM−T (young member employed and the older in transition) is much greater than the number of pairs T−EM throughout the period, except for the first two months.

The vulnerability index can be compared similarly. Figure 10 displays the histograms of the index for all the sequences, separately for the younger and older individuals (top panels), and their counterparts for the matched pairs. The index has much more similar distributions in the groups after matching. In particular, the discrepancy in the percentage of the maximum score is eradicated; in fact, the frequency among the older individuals in the matched pairs is slightly lower.

The means and standard deviations of the respective young and old groups are 7.73 (13.07) and 16.12 (22.86). The corresponding figures for the matched pairs are 11.26 (18.49) and 11.94 (18.58), so by matching on the covariates and labour force history (year 2008) we have also obtained a close agreement of the distributions of the vulnerability index for the period January 2009 − July 2011. The differences within the matched pairs have zero median, mean -0.69 and standard deviation 26.27. This conclusion might appear to contradict the conclusion based on the switches summarised in Figure 8.

As an alternative to the histogram and summaries of the values of the index within the groups, which rely on the assumption of additivity, we may compare the matched pairs by tabulating the sign of the within-pair differences. Among the matched pairs, there are 2346 negative differences (younger vs. older), 2305 positive differences and 735 ties. The difference of two binomial variables (for ± 1) with denominators 2350 and 2300 and both with probabilities $p = 0.5$ has standard deviation $0.5\sqrt{4650} = 34.1$. The realised difference, 41, is only slightly greater, so it could have arisen purely by chance. In the analysis of the states of the matched pairs in Figures 8 and 9, we found substantial differences between the two groups. Thus, some information is lost, or distorted, by summarising the sequences of states by the vulnerability index and treating the index coarsely.

The main purpose of propensity matching is to obtain two groups that appear to be equivalent with regard to their background — that have as similar distributions of the covariates as one would expect in a hypothetical experiment in which the treatment is assigned to the individuals completely

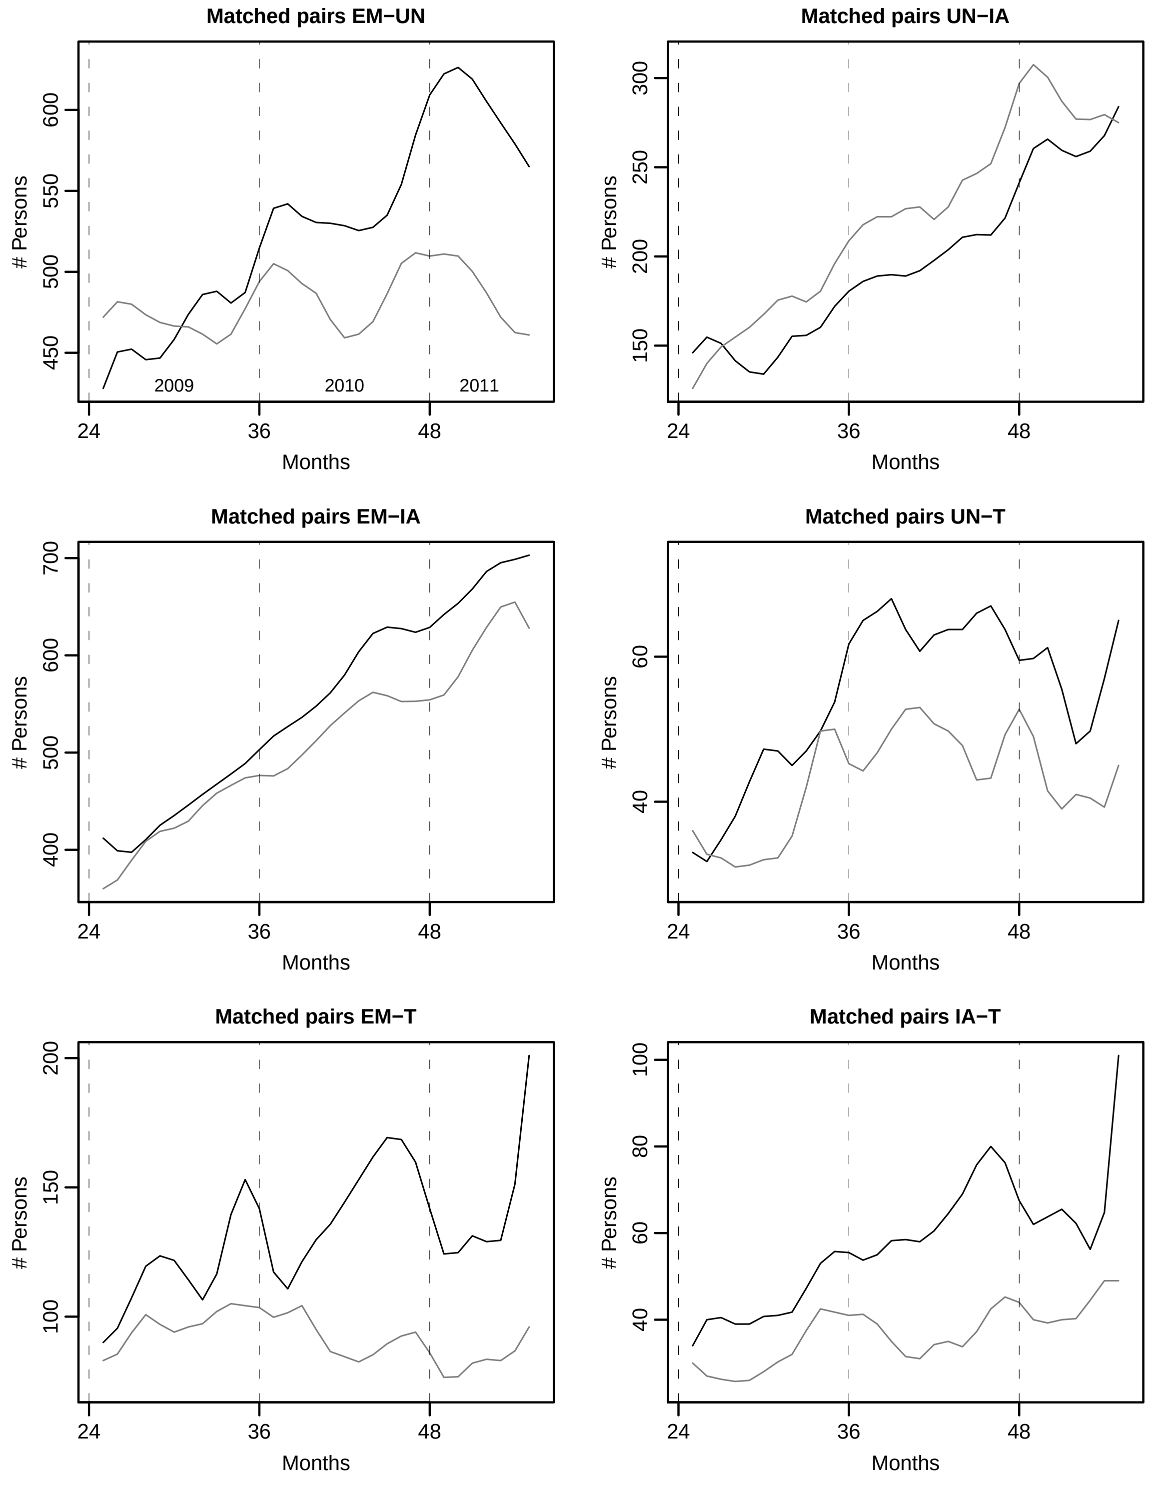

Figure 9: Summary of the matched labour force carpets in Figure 8. Black colour is for the pairs Younger–Older in respective groups A–B, and gray for the reverse.

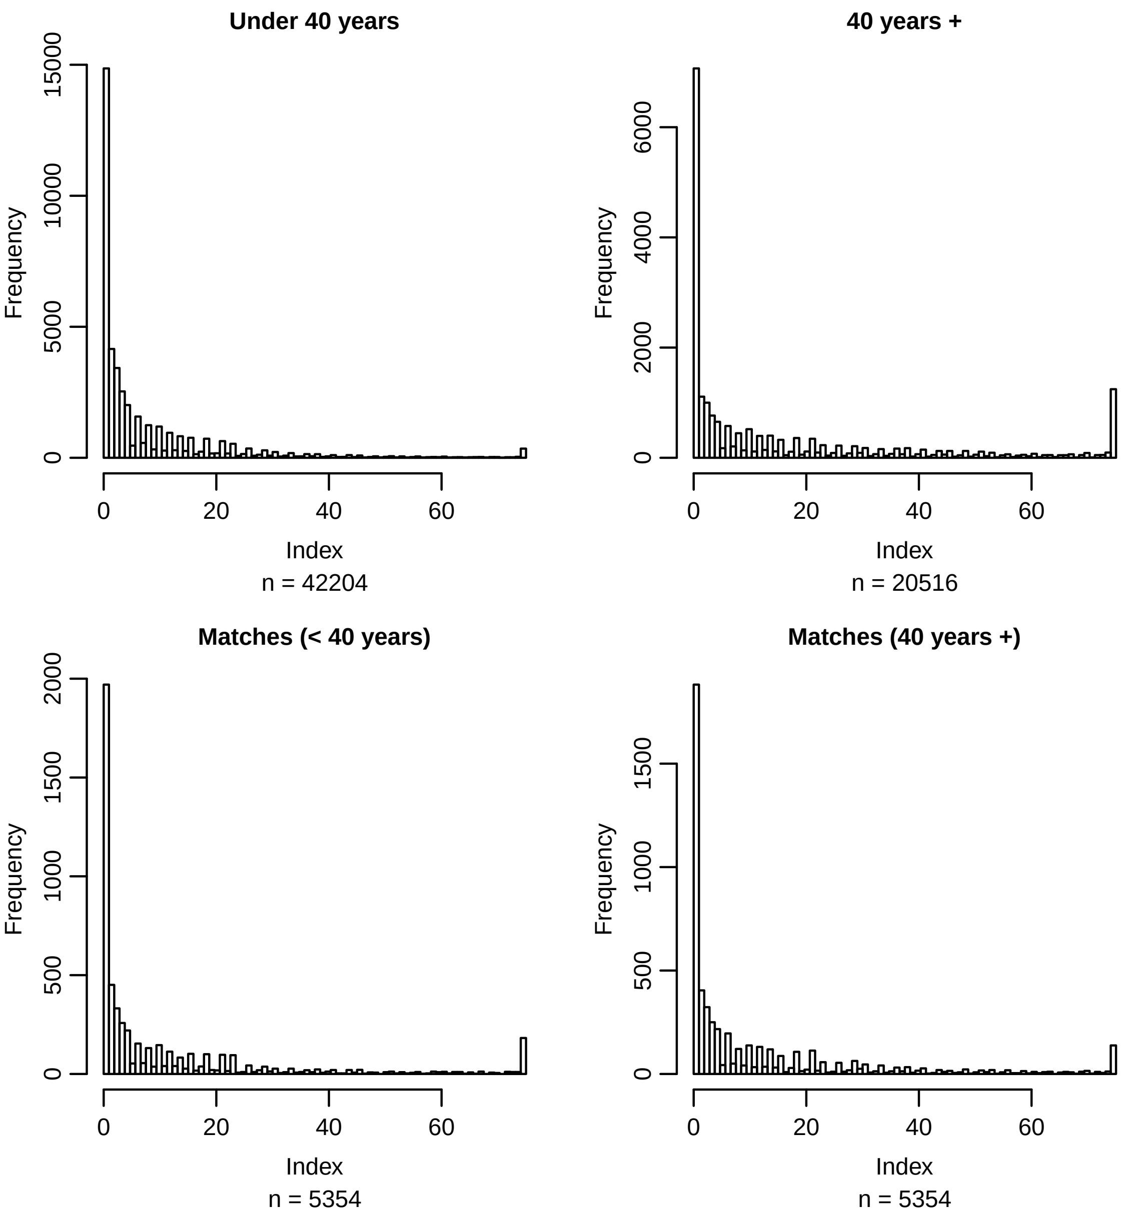

Figure 10: Distribution of the vulnerability index for younger and older individuals before and after matching.

at random. We assess whether this goal has been achieved by comparing the distributions of the two groups on all the variables used for matching. These comparisons are listed in Table 3 for the covariates that are categorical. The left-hand part of the table lists the percentages of individuals younger and older than 40 years in each group of a categorical covariate, including a category for missing values, and their within-category differences. The right-hand part of the table lists the corresponding differences between the matched subsets of the groups. The largest absolute difference in the latter is 3.3% (Luxembourgois nationality), which is even higher than in the entire database. However, some substantial differences, such as for singles and married (Marital status 1 and 2) and lower level of education, are reduced substantially by matching.

Figure 11 summarises the differences before and after matching on the 12-month labour force history of the matched and unmatched groups. Each month used for matching is represented by a set of five segments, one for every labour force status. The segment is drawn between the raw difference of the two groups, e.g., −6.6 for employment (EM) in month 13 (January 2008), and its negative, 6.6, which corresponds to the discrepancy of the same magnitude. The discrepancy for the matched pairs is marked by the black disc (0.18 for M13–EM). The diagram shows that the differences have been reduced substantially by matching, but some small differences remain.

For the two continuous covariates, date of completing education and the first record in IGSS, we have to assess the agreement of the entire distributions. We reduce this to comparing the means and standard deviations. For the date of completing education, the respective means (standard deviations) for the younger and older are 1999.75 (9.30) and 1978.31 (12.00), in calendar years, whereas after matching they are 1987.44 (12.40) and 1987.29 (13.44). The corresponding figures for the first record in IGSS are 1998.49 (13.74) and 1986.88 (15.44) for unmatched (entire) groups and 1992.15 (15.59) and 1993.14 (14.58) for the matched pairs. We conclude that most of the disparity has been removed by matching on the propensities, but the differences in the percentages of the categories between one group and the other have not been eradicated.

4.4 Comparing time periods

In this section, we address the issue of comparing the sequences of states in two periods. The periods may overlap. As an example, we choose the two-year periods August 2009 – July 2011 and August 2008 – July 2010. As in the earlier example, the periods may be compared without or with matching. For matching we use the same covariates (see Table 2) and twelve months of labour force history.

Table 3: The balance of the categories of the background variables in the unmatched and matched groups; all the entries are percentages: the composition of the younger and older groups and their within-category differences.

Variable	All $(42\,204 + 20\,516)$			Matched pairs (2×5386)		
Category	Younger	Older	Difference	Younger	Older	Difference
Education						
Missing	2.4	3.3	−0.9	6.0	6.7	−0.7
1	36.4	50.6	−14.2	45.4	43.1	2.3
2	41.4	32.3	9.1	31.6	30.7	0.9
3	19.8	13.8	6.0	17.0	19.5	−2.5
Marital status						
Missing	0.5	0.3	0.3	0.6	0.5	0.1
1	66.2	15.4	50.8	27.4	30.4	−3.0
2	28.6	58.5	−29.9	55.6	53.2	2.5
3	0.0	0.1	−0.1	0.1	0.1	0.1
4	0.1	1.6	−1.4	0.6	0.8	−0.3
5	3.5	18.3	−14.8	12.9	10.4	2.5
6	0.9	5.7	−4.8	2.6	4.4	−1.9
7	0.0	0.1	−0.1	0.1	0.1	0.0
Sex						
Women	47.3	44.4	2.9	45.7	46.8	−1.1
Men	52.7	55.6	−2.9	54.3	53.2	1.1
Nationality						
1	34.7	32.0	2.7	26.9	23.6	3.3
2	8.3	8.5	−0.1	9.2	10.4	−1.2
3	29.9	29.6	0.3	31.7	30.9	0.7
4	1.9	4.4	−2.6	3.7	3.6	0.1
5	2.9	5.0	−2.1	4.5	4.9	−0.4
6	3.8	5.0	−1.3	4.4	4.6	−0.2
7	18.6	15.6	3.0	19.6	22.0	−2.4

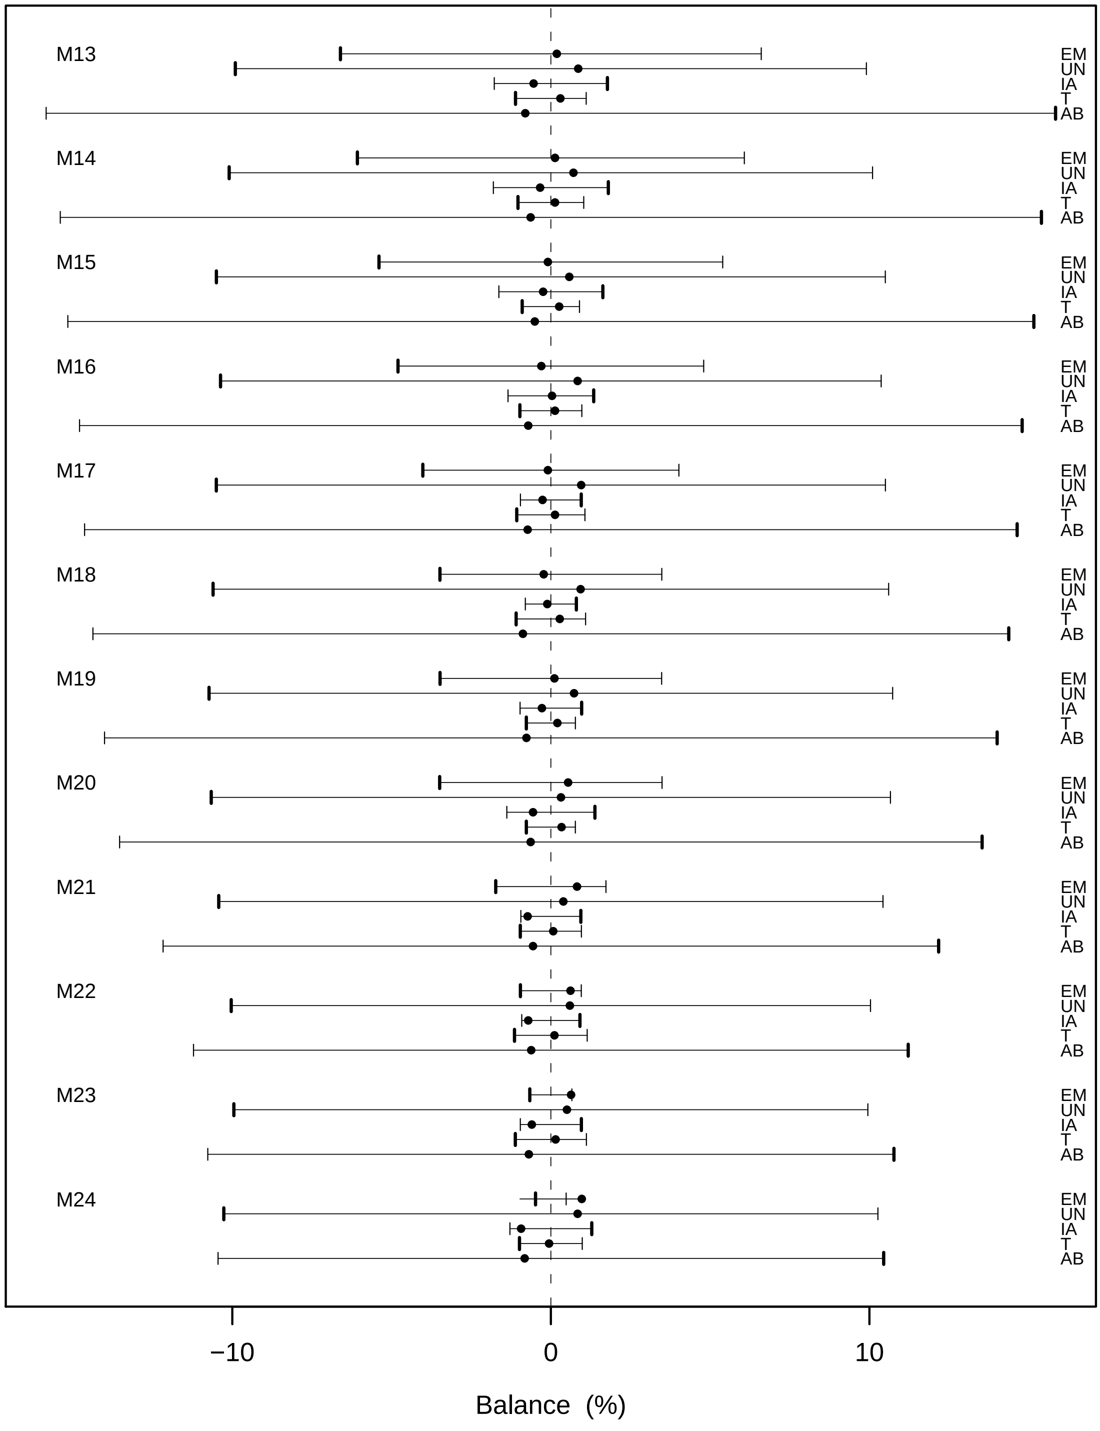

Figure 11: Differences of the percentages of the labour force states in months 13 − 24 for unmatched and matched groups.

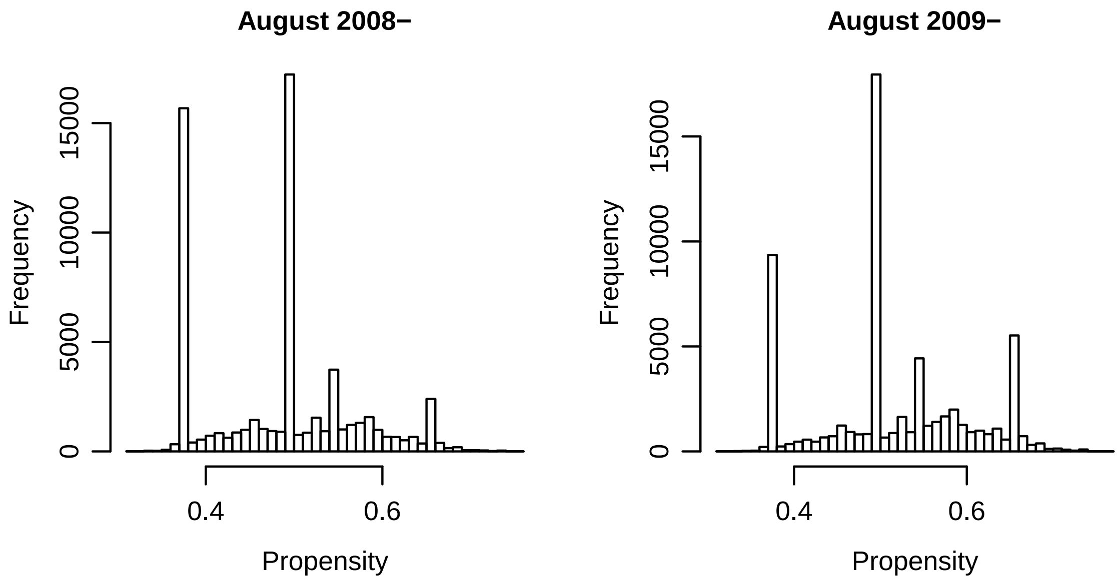

Figure 12: Propensity scores in two two-year periods starting in August 2008 and August 2009.

In the matching procedure, we enforce the additional rule that the donor and recipient have to be distinct individuals.

Figure 12 displays the fitted propensities for the two-year periods. The propensities are spread much less than in the comparison of the two age groups in Figure 7, indicating that the distribution of the labour force profiles has not changed substantially by the passage of one year. The propensities are in the range (0.31, 0.75) and (0.32, 0.76) in the respective groups August 2008– and August 2009–. They are in close agreement on other summaries, such as the median (0.499 for both groups), mean (0.484 vs. 0.516) and standard deviation (0.086 and 0.089). Two values of the propensity, 0.499 and 0.373 occur for many spells in both periods. The spells can be paired by the person from whom they originate; however, the corresponding pairs of propensities are not related in any straightforward way, even though they share a segment of their spells — the second yearly spell in August 2008– coincides with the first yearly spell in August 2009–. The pairs of propensities differ by less than 10^{-4} for 21 200 persons (34%); they include 8800 and 9230 persons whose both propensities are equal to 0.4991 and 0.3732 (after rounding to four decimal places), respectively.

We classified the propensities into 100 intervals of equal width and obtained 53 969 matched pairs. The labour force carpet for the unmatched and matched sets of spells is displayed in Figure 13. The carpets are perceptibly more similar for the matched than unmatched pairs, and this can be confirmed

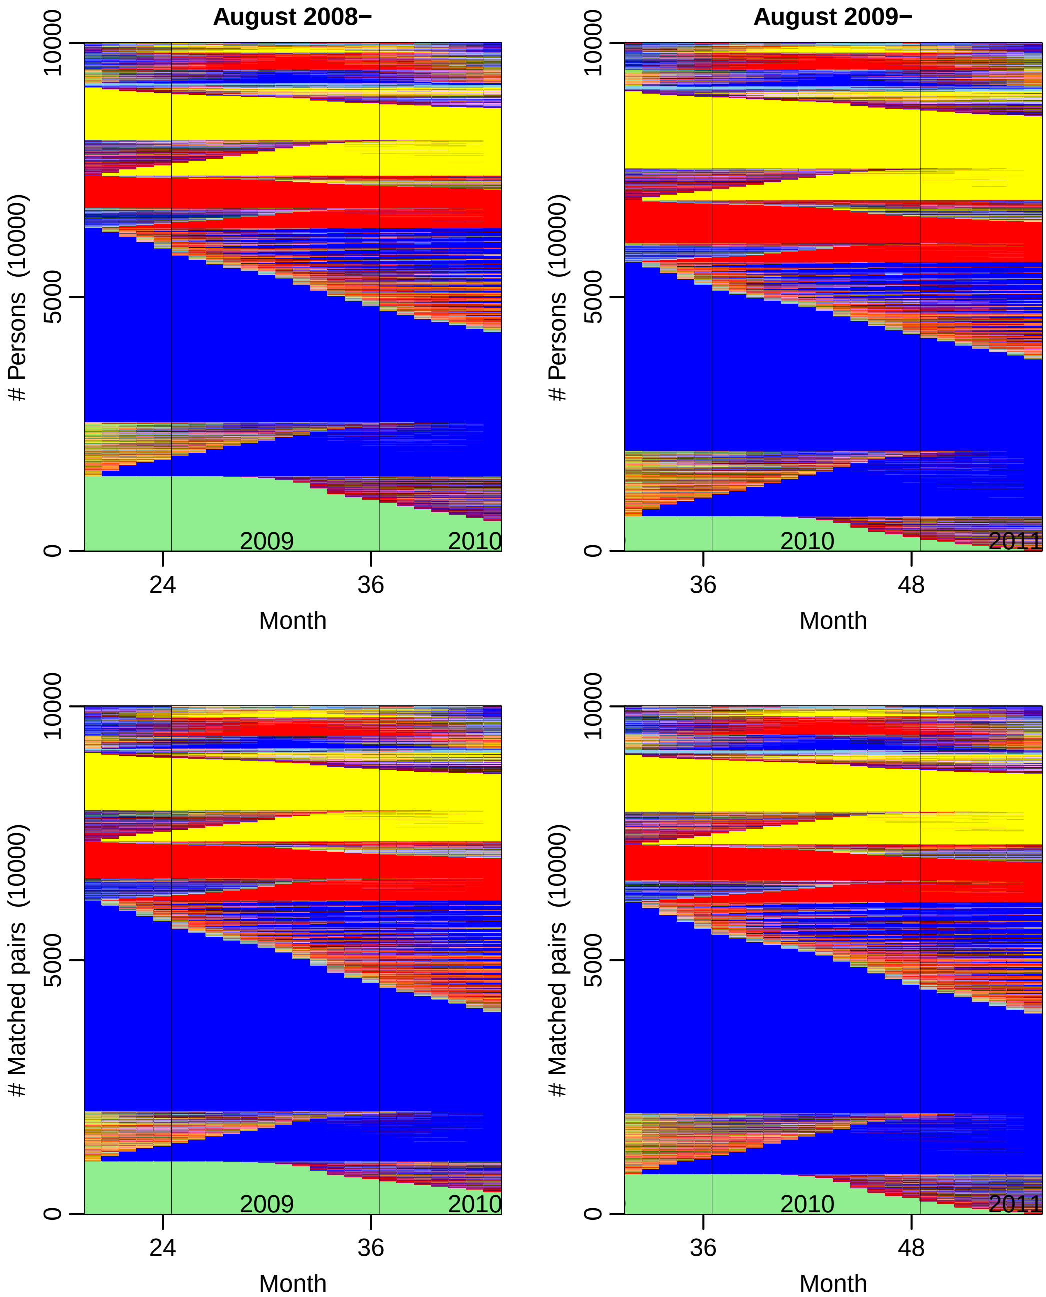

Figure 13: Unmatched (top) and matched labour force carpets (bottom) for the two-year periods starting in August 2008 and August 2009.

by summarising the within-pair differences in the states by month.

Figure 14 presents the monthly summaries of the discrepancies between the matched pairs, using the same layout as in Figure 9. The counts of pairs EM–UN and UN–EM differ only slightly throughout the period of months 20 – 43 for the spells in group 1 and their matches in months 32 – 55. The largest relative differences are in months 22 and 23 (EM–UN in excess of UN–EM by over 10%) and in months 37 and 38 (UN–EM in excess by over 10%). The panels EM–IA and UN–IA indicate a clear excess of switches from UN to IA in the earlier period, but also a consistent excess of switches from EM to IA (retirements) in the later period. Comparisons involving switches to T display no clear trend except for excess departures in the later period from all three states in months 36 – 42 over months 48 – 54.

4.5 Employment following an unemployment spell

In this section, we study the lengths of the EM spells that follow a spell of UN recorded by ADEM. We reduce our attention to persons who were below 30 years of age at the beginning of the UN spell, because more senior members of the labour force tend to have difficulties in the search for new employment but, when successful, they often secure long-term employment; see Figure 5. The unit of this analysis is a UN spell (an ADEM case file), and we study the sequence of labour force states following the end of such a *reference spell.* We define the outcome as the length of the succeeding EM spell (*the outcome spell*). The outcome is set to zero if the reference spell is followed by IA. Some persons have several UN spells in the period January 2007 – July 2011. There are 43 825 reference spells. They involve 26 835 unique individuals, 16 698 (62.2%) of whom have a single record each, 5957 (22.2%) two records and 2481 (9.2%) three records each. One person has eleven reference spells, the maximum, and four persons have ten spells each.

We classify the reference spells of UN according to whether the person received any unemployment benefit or not. The eligibility and the amount are determined by a set of rules (see Appendix), and the payment regime, although quite complex, is entirely predictable once a decision is made by ADEM. Economic theories argue that persons who receive benefit are under less pressure to resume employment, and are therefore more selective in their job search. As a result, their reference spells tend to be longer, but are followed by longer outcome spells. We test this hypothesis by comparing the two groups on the lengths of their outcome spells. That is, receiving benefits is regarded as the treatment, the length of the EM spell is the outcome variable, and the covariates are the variables listed in Table 2, which have their values fixed at the beginning of the reference spell. The length of the outcome spell is truncated at 24 months, if applicable. For the outcome spells that ended in

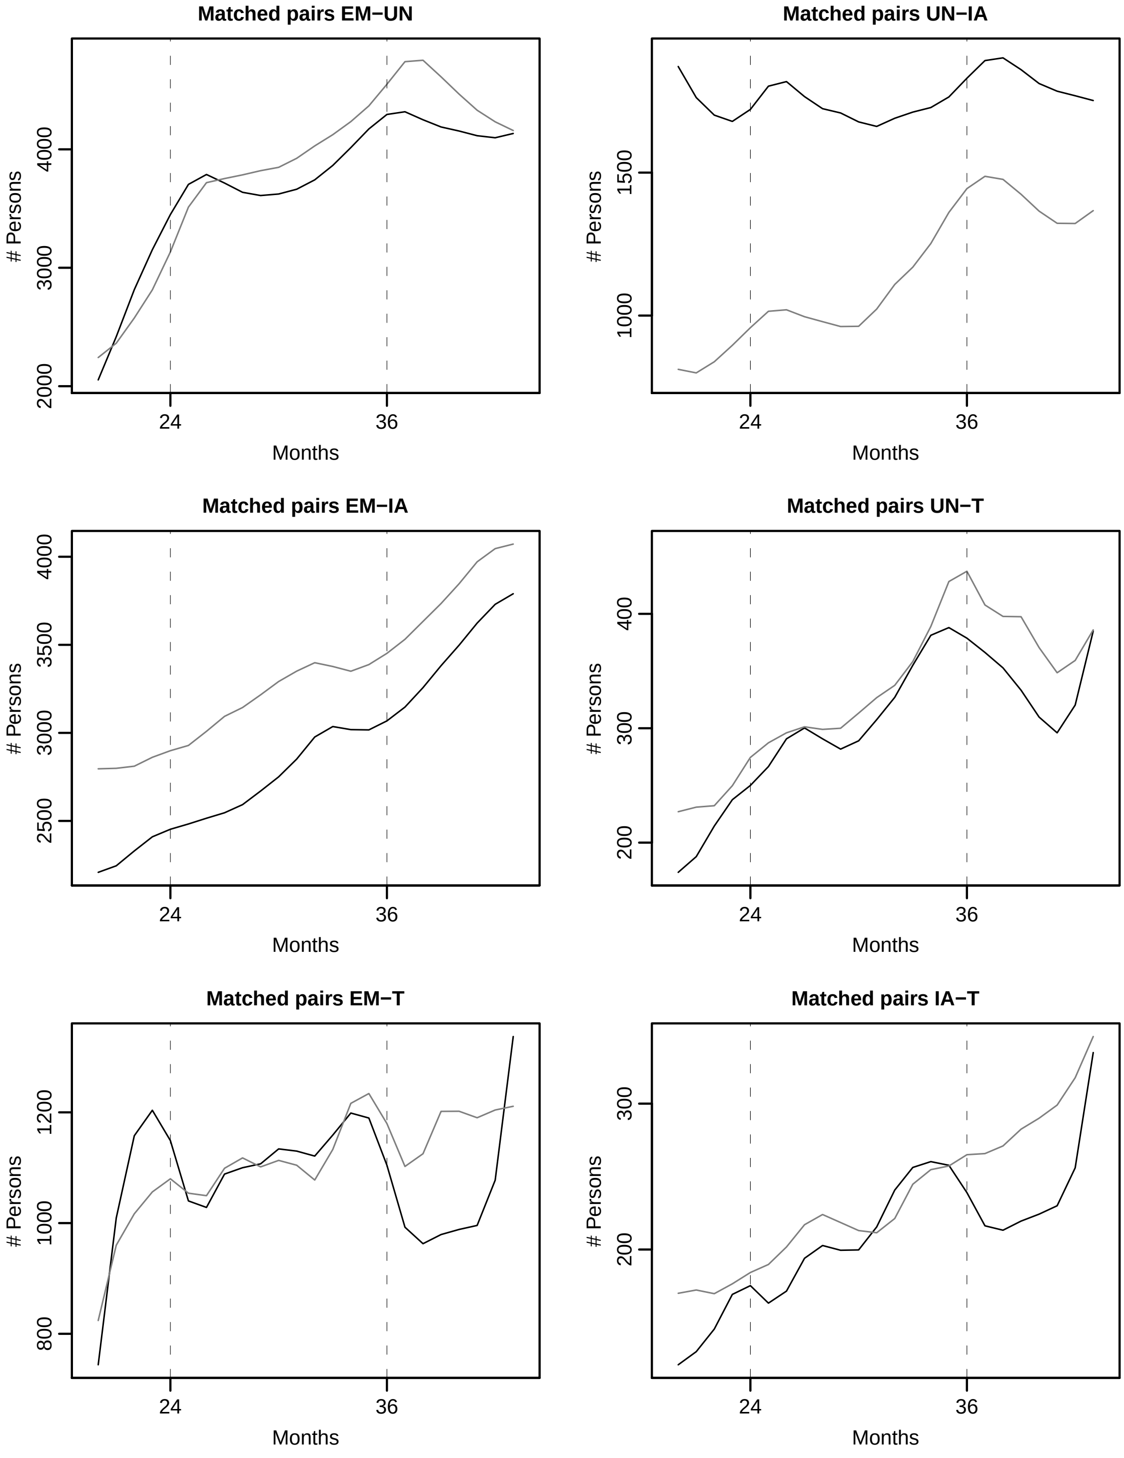

Figure 14: Monthly summaries of discrepancies between the matched pairs of spells starting in August 2008 and August 2009.

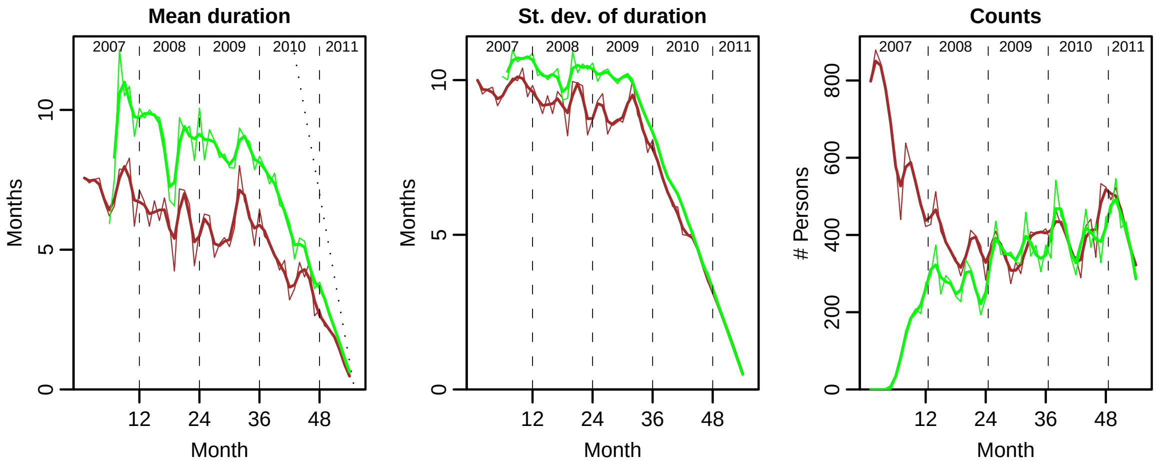

Figure 15: Monthly summaries of the outcome (employment) spells for reference (unemployment) spells with benefits (green) and without (brown). The dots mark the truncation due to the horizon in July 2011.

August 2009 or later, further truncation takes place because our records end in July 2011, less than 24 months later.

Without any adjustment for covariates and without any matching on their values, the outcomes are summarised by their within-group means. We evaluate these summaries separately for each month in which the reference spell ended. Figure 15 presents these summaries as a pair of functions in the left-hand panel. It shows that reference spells associated with benefits (green) tended to have longer outcome spells than reference spells with no benefits (brown), and the differences persisted until about mid-2010. From then on, many outcomes are truncated by the data horizon of July 2011. In the diagram, thin lines connect the exact within-group and within-month means, and the thick lines are obtained by smoothing, applying (2). The annual mean lengths of the outcome spells starting in the respective years 2007, 2008 and 2009 are 9.23, 8.32 and 8.46 for reference spells with benefits and 6.68, 6.63 and 6.40 for spells with no benefits. The means and their difference in 2010, $5.55 - 4.50 = 1.05$, were much smaller as a result of substantial truncation.

The middle panel plots the standard deviations of the employment spells. It shows that the standard deviations have a pattern similar to the means, although the between-group differences in standard deviations are much smaller than for the means. The right-hand panel plots the numbers of spells involved in the comparisons. In total, there are 15 329 spells (38.1%) without a benefit and 24 888 spells with benefit. There is a clear seasonal pattern, with sharp increases in the numbers of reference spells ending in February and October of each year.

We cannot attribute the observed differences solely to the group membership, because the rules for awarding unemployment benefits are intentionally selective. Neither can we make any references to causal analysis, because we regard as the treatment the award of a benefit as defined by the (ADEM) rules in force. However, we can speculate what would happen if the group membership were altered in complete isolation from the past labour-force and related history and background of the persons registered in ADEM. In such a process, we have to discount the possibility of so-called *interference*, that the change of rules would affect others, including those with no ADEM records, for instance, through altered policies of employers.

We set aside these concerns and compare the two groups after matching them on the set of available covariates. Even without a reference to causal analysis, we regard such a comparison as more appropriate because by matching we come closer to the ideal of comparing like with like. We match on the variables listed in Table 2, supplemented by the labour force states in the 12 months preceding this spell, the month when the reference spell starts and by the length of the reference (UN) spell. We truncate the length of the reference spell at 13 months, to avoid values with disproportional influence on the estimates and to better reflect the seriousness of the spell in relation to its length.

Missing values, which arise by reference to months prior to January 2007, are treated as a separate category. This would make the matching more stringent, but we match not only on the propensity score, but also on the month in which the outcome spell starts (fine matching; Rosenbaum, Ross and Silber, 2007). We fit a single propensity model, but pair the outcome spells within intervals of fitted propensity scores within the 55 months. The propensities are divided into 20 equidistant intervals that cover their range.

The matched pairs can be compared by the average difference of the lengths of their outcome spells. Using a linear scale for the length of the outcome spell is problematic, because it would be weighed too heavily by the long spells. In our case, every spell is truncated, after 24 months by our choice, and also by the horizon at July 2011. A non-parametric version of our analysis is based on within-pair comparisons. We count the number of pairs in which the reference spell with a benefit and in which its match with no benefit has a longer outcome spell and compare the corresponding percentages. We refer to this comparison (or analysis) as being of dominance. Ties are quite common, because many UN spells end with IA, when the value of the outcome is zero. The two percentages therefore do not add up to 100.

The within-month comparisons of the lengths of EM spells on the matched pairs are displayed in Figure 16. The results differ from their counterparts with no matching (Figure 15) substantially. Only

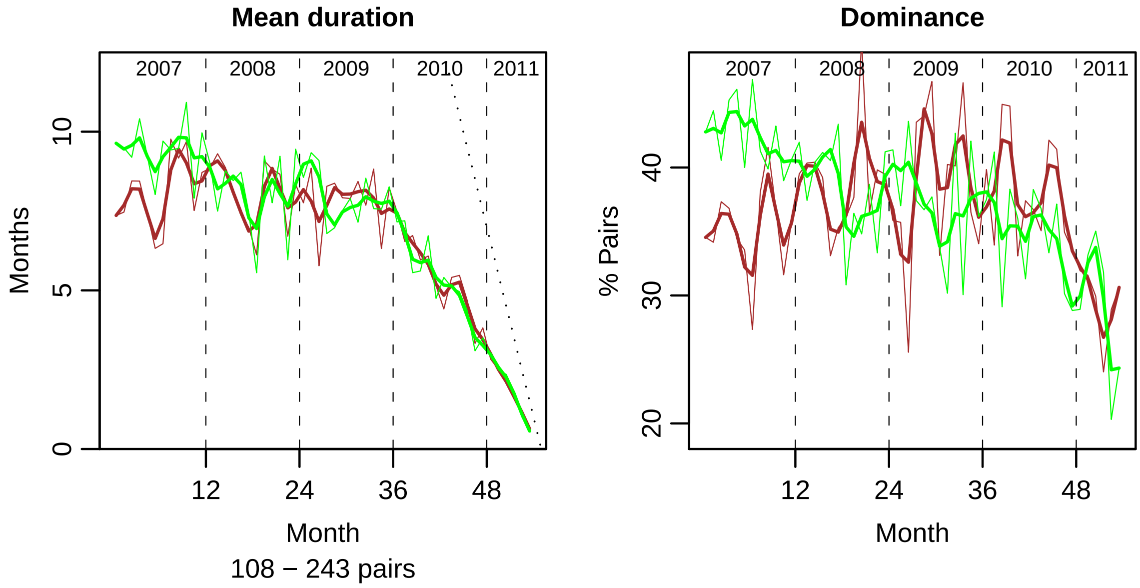

Figure 16: The mean lengths of outcome spells and dominance for pairs matched by propensity categories and month of leaving the reference spell of unemployment; reference spells with benefit — green; without benefit — brown.

7110 pairs, involving about 35% of all the spells that ended in months $1-54$, are formed. Employment tended to last longer after spells that ended in 2007, but toward the end of the year the difference shrunk to a few months, and from then on the differences have been very small, with spells without benefit having higher mean on average in a majority of the months. The comparison of the dominance in the right-hand panel appears to be somewhat clearer, partly owing to the choice of the vertical scale, but the two panels are largely in agreement. Some of the volatility of the means and dominances can be attributed to the relatively small effective sample sizes — there are about 130 matched pairs in a month on average, and these counts vary a great deal. There are only 86 pairs in November 2008 and 235 pairs in February 2007. Smoothing injects some stability in the comparisons in the diagram.

As an alternative to more extensive smoothing, we summarize the differences and dominance by periods longer than one month, such as a quarter or half a year. We illustrate this on more refined comparisons, done separately for each of the three educational groups. We use the same set of fitted propensities as in the earlier analyses of the 'effect' of the unemployment benefit, but match on the level of education in addition to the propensity scores, and evaluate summaries within these levels. We obtained 2673, 3408 and 812 matched pairs for the respective lower, middle and higher educational levels, that is, 6893 pairs in total, 217 (3.0%) fewer than without the constraint on the level of education. The three pairs of panels in Figure 17 present the results using the same layout as in Figure 15, but with (smoothed) quarterly summaries.

The averages of the outcome spells are quite smooth for the lower and middle levels of education, for which we have many more matched pairs than for the higher level. The average is perceptibly lower for the lower level of education, for which benefit is associated with longer outcome spells, except for year 2008 and the quarter preceding it. In contrast, benefit is associated with shorter outcome spells for those with higher level of education. Although dominance is a coarser way of comparing EM spells, the comparison for the higher level of education is clearer, albeit only after smoothing. In summary, the effects of the unemployment benefit differ substantially among the levels of education.

Another alternative is motivated by survival analysis. For each length of the outcome spell, we consider as the basis the number of spells 'at risk', which have not yet been concluded, and evaluate the percentage of them that conclude in the coming month. These percentages are then drawn for each reference month (end of the reference spell) in separate panels for spells with and without benefit. From such a diagram we have to exclude reference months $31-55$, because of earlier censoring. The versions of this diagram for quarters or half-years can be drawn. We have found them to have much poorer resolution than the diagrams for averages and dominance. For brevity, they are omitted.

In our analysis, the persons registered in ADEM and their records are fixed, and therefore so are the fitted propensities, even though the model underlying them entails uncertainty. However, in a replication of the study (analysis) we would use the same register database and the same set of propensities would be obtained. Uncertainty arises only in the next step when matched pairs are formed. This source of uncertainty can be assessed simply by replications. Figure 18 displays the estimated averages and dominances for ten replications, using the same scales and conventions, including smoothing, as in Figure 16. The variation associated with matching does not in any way challenge the conclusions we drew about the comparison of the mean outcome spells from a single replication in Figure 16 (left-hand panels of the two diagrams). The variation associated with dominance is much greater and some details observed in a single replication are not reproduced in other replications. However, the overall conclusions are not contradicted by the replications. For orientation, the averages over the ten replications, drawn by thick dashes, are added in the panel for dominance.

The paperwork with an application takes a few weeks, and in some cases longer, depending on the level of cooperation of the applicant, complexity of the case and the like. If the applicant secures a (new) employment in the meantime, the file associated with the application is closed, and the applicant may receive a benefit payment, but this would happen during his/her EM spell, and would

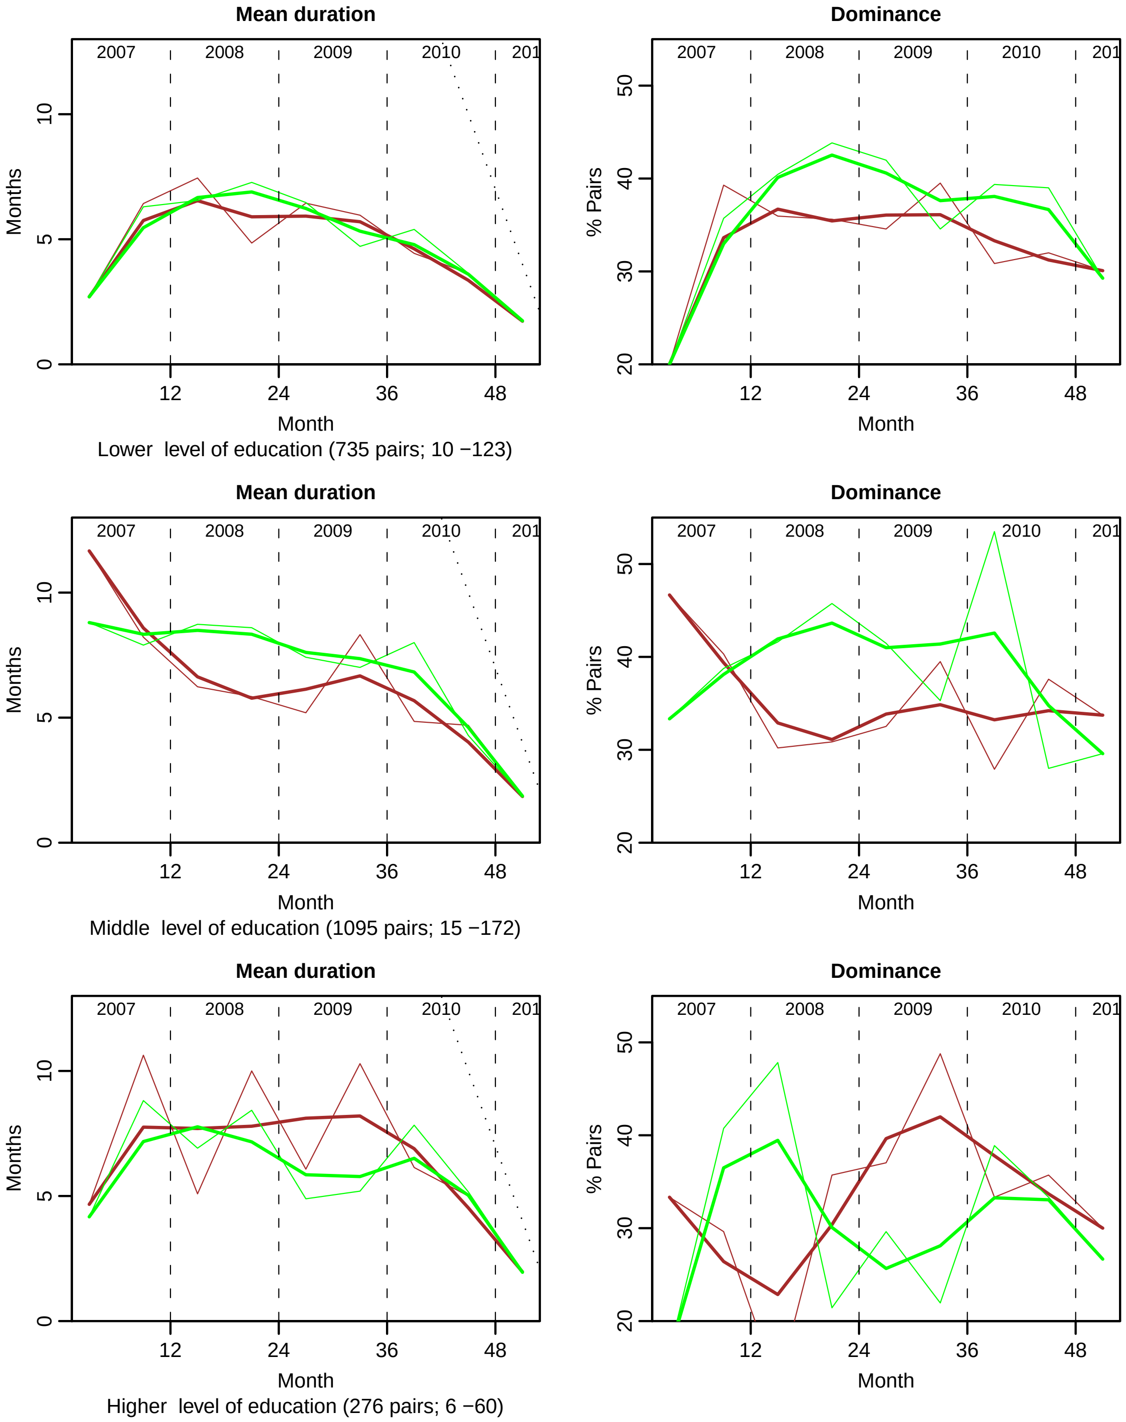

Figure 17: Comparisons of the mean lengths of EM spells and dominance for pairs matched on propensity categories, level of education and the month of leaving the reference (UN) spell.

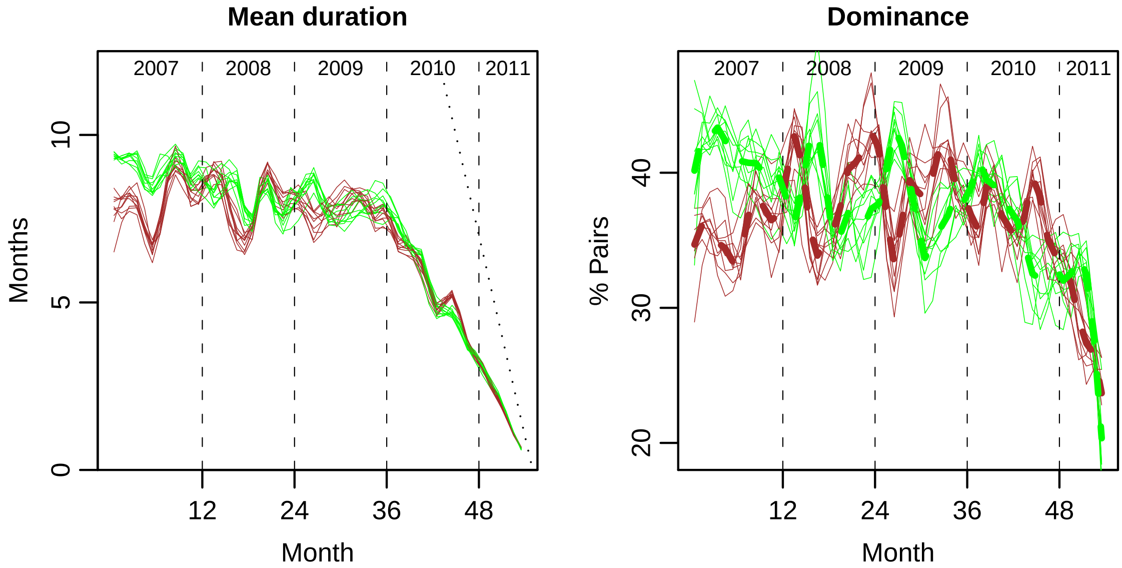

Figure 18: Replications of the propensity matching analysis for the average length of the outcome spell and dominance. Thick dashes in the right-hand panel mark the average over the ten replications.

not be connected in the dataset with the reference UN spell. Also, some short spells that are entirely contained in a single month are not recorded because the labour force status is established only at the end of each month. The treatment status (benefit paid or not) is therefore not recorded perfectly, especially for short spells (up to three months). Unfortunately, such spells are in a majority among the young (nearly 30 000 out of 43 825, that is, 68%). Among the spells of the older unemployed, this percentage is 49.4%.

4.6 Results for the older workers

This section summarises the results for the outcome spells of persons over the age of 30 years. There are 56 102 such (reference) spells involving 37 181 distinct persons; 25 213 of them (67.8%) have one spell each and 7580, 2764 and 992 (20.4, 7.4 and 2.7%, respectively) have two, three and four spells each. The largest number of spells of a person is eleven, in one instance.

Figure 19 plots the monthly summaries of the outcome spells that immediately follow a reference spell. Among the older workers, the spells without benefit were in excess of 50% only in 2007; since 2008 their numbers have stayed in a narrow band, 275 – 440, whereas the number of spells with benefits has increased throughout, with distinct peaks in February and troughs in July. The mean length of the outcome spells for reference spells with benefits is greater than for spells without benefits only

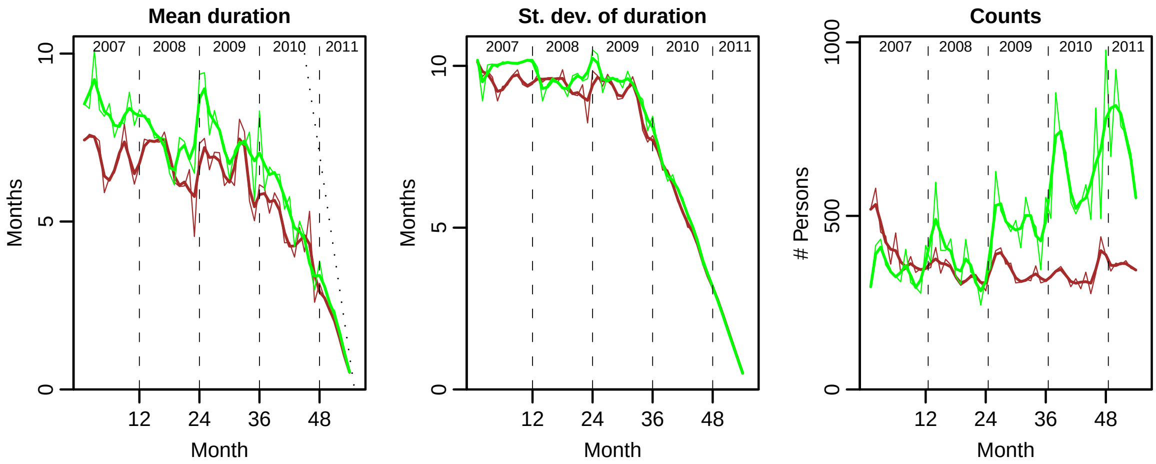

Figure 19: Monthly summaries of the outcome spells for reference spells with benefits (green) and without (brown); persons over 30 years of age.

in some periods, such as the entire year 2007, late 2008 and the first half of 2009, but is never much smaller. The standard deviations of the lengths are greater throughout, but only by a narrow margin in most months.

Using the same rules as for the younger persons, propensity matching yields 10 380 pairs, involving 37.0% of the reference spells. The average differences and dominances for the matched pairs, the counterpart of Figure 16 for the older workers, are displayed in Figure 20. Most of the differences in the averages observed for the unmatched groups are substantially reduced after matching, except for those in the first half of 2007. For dominance, there are some large differences in relatively short periods of time, but they change abruptly from one month to the next.

4.7 School leavers

Among the young members of the labour force, the school leavers, who have minimal (or no) record of employment, are a particularly vulnerable subpopulation. A school leaver is defined as a person who completed his or her education no more than 12 months earlier than the reference time point. In our context, this time point is the beginning of the UN spell (and of a record in the ADEM register). Our database contains 6754 UN spells of school leavers in period January 2007 – July 2010. They involve 6079 unique persons. We summarise their labour force states in the 12 months following the end of the UN spell by the vulnerability index (Section 4.2). There are very few 15- and 16-year olds in our records, so we truncate the age at 17 years. Similarly, the age is truncated from above at 27

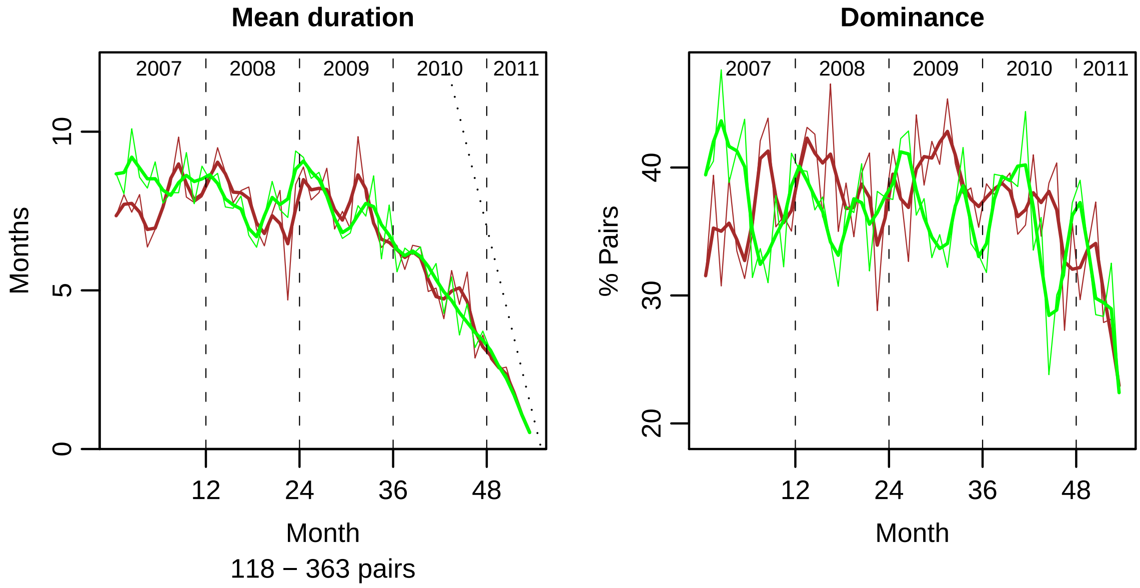

Figure 20: Comparisons of the mean lengths of EM spells and dominance for pairs matched by propensity and month of leaving the reference (UN) spell; persons over 30 years of age.

years of age. Figure 21 displays the means of the vulnerability indices by age. It shows that the younger entrants tend to have higher scores, although the difference of the means between 20- and 25-year-olds is not substantial. The diagram confirms the 'common knowledge' that completing high school education (leaving school after the age of 18) confers a distinct advantage in the labour market.

This is reinforced by the means of the vulnerability index within the educational groups; the means (standard deviations) in the respective educational categories 1 (lower), 2 (middle) and 3 (higher) are 1.42 (3.33), 0.94 (2.51) and 0.45 (1.79). The corresponding numbers of records are 1391, 3786 and 1577; around 10% of the spells are second or later spells of the same person. We note that the comparisons made are without any matching. The age and educational-level groups cannot be meaningfully matched on some background variables, because a change in these variables, if any control were feasible, would be associated also with a change in other background variables. However, it is meaningful to consider the effect of education on the labour force profile in the future.

The observed trends are reproduced in gross features for the outcomes within the calendar years 2007 − 2010, although the sample sizes, especially for 2010, are rather small (35–85 within months). For the educational level, the annual mean vulnerability index has risen over the years 2007 − 2009 for school leavers with lower and middle level of education, but has dropped slightly for those with higher level of education. This suggests that coping with unemployment has become more difficult for those

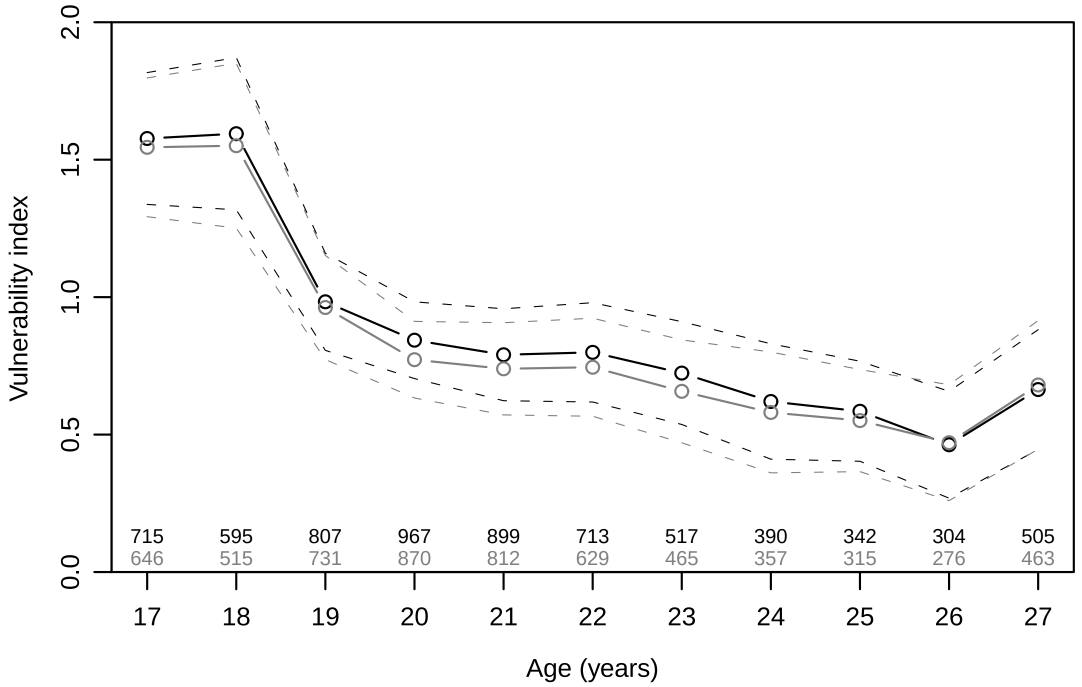

Figure 21: Mean vulnerability index for one year after a UN spell, as a function of age; school leavers. All spells in January 2007 – July 2010 (black lines); first spells (gray lines). The dashes mark the pointwise 95% confidence limits. The counts of records are given at the bottom margin.

Table 4: Annual mean vulnerability index for educational groups; spells of unemployed of school-leavers.

Educational level	Calendar year 2007		2008		2009		2010	
Lower	1.28	(0.14)	1.53	(0.18)	1.72	(0.21)	1.00	(0.16)
Middle	0.78	(0.06)	0.91	(0.07)	1.20	(0.10)	0.94	(0.13)
Higher	0.55	(0.09)	0.47	(0.09)	0.44	(0.07)	0.18	(0.07)

with lower and middle level of education, but those with higher level have not been affected, or their position in the labour market has even strengthened. The corresponding values for 2010 are much lower for all three educational groups, but the data are not complete and the summaries are distorted by the absence of data from the second half of the year. The details are given in Table 4.

5 Conclusion

We have presented a set of analyses of the unemployment and social security registers in Luxembourg. Their purpose is to monitor the return of (recently) unemployed to employment and inform the employment policy of the country's government. The outputs of the analysis are mainly graphical, in the form of labour force carpets and their summaries, although (univariate) analyses in which one or several indices are compared across subpopulations are also important.

For comparing two subpopulations, we prefer comparisons of pairs matched on propensity scores, drawing on advice and experience collected in Rubin (2006). Multinomial regression would have to be adapted for the essentially multivariate comparisons for the sequences. Model choice, involved in the propensity analysis, is independent of the outcomes, so it can be done even before the outcomes are realised.

The list of analyses presented here is not complete, and is likely to evolve over time as the overall inferential agenda is continually updated. The analyses suggest that some differences between subpopulations defined by age and educational level have alternative explanations in terms of other background variables. One purpose of the analyses is to inform the process of devising and assessing new interventions in the labour market. We note that such analyses have to take into account the interventions already taking place; that is, a new treatment is devised for a population that is already subjected to other treatments. These treatments and rules related to them are transparent, so they are subject to mutual interference.

All computing described in this paper was done by custom-written functions in `R`. The functions are available from the first author on request, but the datasets are proprietary and cannot be distributed. A key rule adopted in these functions is to avoid operations on sequences, of which there are many, in favour of operations on time points. Some of the algorithms used are implemented in the `R` package `TraMineR` (Gabadinho *et al*, 2009), but may require some adaptation for large numbers of sequences.

Acknowledgements

We express our gratitude to ADEM and IGSS for granting us access to their databases. The first author acknowledges hospitality of CEPS/INSTEAD during his tenure as a Visiting Professor in May and June 2012. Insightful and constructive comments and suggestions by Jacques Blosius and Philippe van Kerm are acknowledged.

References

Elzinga, C. (2010). Complexity of categorical time series. *Sociological Methods and Research* **38**, 463–481.

Gabadinho, A., Ritschard, G., Studer, M., and Müller, N.S. (2009). Mining sequence data in `R` with the `TraMineR` package. A user's guide. University of Geneva.

Hoy, M., Thompson, B.S., Zheng, B. (2011). Empirical issues in lifetime poverty measurement. *Journal of Economic Inequality.* DOI 10.1007/s10888-011-9192-1.

La Sécurité Sociale (2011). Code de la Sécurité Sociale. Lois et Réglaments. Ministère de la Sécurité Sociale. The Grand Duchy of Luxembourg.

Lechner, M. (2002). Some practical problems in the evaluation of heterogeneous labour market programmes by matching methods. *Journal of the Royal Statistical Society* Series A **165**, 59–82.

McCullagh, P. (2008). Sampling bias and logistic regression. *Journal of the Royal Statistical Society* Series B **70**, 643–677.

McVicar, D., Anyadike-Danes, M. (2002). Predicting successful and unsuccesful transitions from school to work by using sequence methods. *Journal of the Royal Statistical Society* Series A **165**, 317–334.

Mendola, D., Busetta, A., and Milito, A.M. (2011). Combining the intensity and sequencing of the poverty experience: a class of longitudinal poverty indices. *Journal of the Royal Statistical Society* Series A **174**, 953–973.

Molenberghs, G., and Verbeke, G. (2005). *Models for Discrete Longitudinal Data.* Springer-Verlag, New York.

Ries, J. (2012). Regards sur le politiques de l'emploi. *Regards* **7**, 1–4. STATEC, Institut national de la statistique et des études économiques, Luxembourg.

R Development Core Team (2009). R: A language and environment for computing and graphics. R Foundation for Statistical Computing, Vienna, Austria.

Rosenbaum, P.R., and Rubin, D.B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**, 41–55.

Rosenbaum, P.R., Ross, R.N., and Silber, J.H. (2007). Minimum distance matched sampling with fine balance in an observational study of treatment for ovarian cancer. *Journal of the American Statistical Association* **102**, 75–83.

Rubin, D.B. (1973). Matching to remove bias in observational studies. *Biometrics* **29**, 159–183.

Rubin, D.B. (2006). *Matched Sampling for Causal Effects.* Cambridge University Press. New York.

Shorrocks, A. (2009). Spell incidence, spell duration and the measurement of unemployment. *Journal of Economic Inequality* **7**, 295–310.

Studer, M., Ritschard, G., Gabadinho, A., and Müller, N.S. (2011). Discrepancy analysis of state sequences. *Sociological Methods and Research* **40**, 471–510.

Verbeke, G., and Molenberghs, G. (2000). *Linear Mixed Models for Longitudinal Data.* Springer-Verlag, New York.

Appendix. Rules for compensating unemployed

This Appendix describes the main features of the unemployment benefit scheme of the Grand Duchy of Luxembourg.

All residents aged $16-64$, who involuntarily lose employment after having worked for at least six months in the twelve months before registering in ADEM, and who are available for work and are able to take up employment, are eligibile for insurance benefits. Benefits are paid without any delay for a maximum of 24 months, but the duration of compensation cannot exceed the number of months spent in employment prior to entry into unemployment. Compensation is extended by up to twelve months for unemployed aged 50 years or over who have contribution records of at least 20 years. The monthly payments are equal to 80% of the wage earned prior to the job loss, and to 85% if the unemployed person has dependents living in the same household. The amount of monthly benefit payments is subject to limits. The lower limit is implied by the statutory minimum wage (SMW), and the upper limit is set to 250% of SMW for the first twelve months of payments and to 200% thereafter.

The unemployment benefit scheme in Luxembourg caters also for those leaving education who have never been employed before. School-leavers and graduates up to 28 years of age who fail to find work within a year of completing their education are eligible for unemployment benefits. After a waiting period of ten months, graduates are awarded benefits that amount to 70% of SMW for unqualified workers, while school leavers receive only 40% of SMW. The waiting period is reduced to 6.5 months if the young unemployed enrols in an internship or training. The period is skipped altogether if he or she enrols in community service.

ADEM imposes sanctions on applicants who fail to attend the scheduled monthly meetings. A typical sanction is the loss of the benefits pro-rated for a set number of days, and three missed appointments in succession result in closure of the case file and loss of all benefits. These and related rules are published in the social security code for Luxembourg (*Code de la Sécurité Sociale; Lois et Réglaments*, 2011). They relate mostly to the information recorded in the ADEM register. Exceptions are the information about rejecting a job offer mediated by ADEM and excepting a job offer before the application for unemployment benefits has been processed. Our attempts to recover this classification were not successful. Our classification agrees with the classification by ADEM in only about 70% of the spells, and the discrepancy cannot be accounted for by the information that is not available to us.