

A Bayesian Analysis of Female Wage Dynamics Using Markov Chain Clustering

Christoph Pamminger¹ and Regina Tüchler²

¹ Vienna University of Economics and Business, Austria

² Wirtschaftskammer Österreich, Vienna, Austria

Abstract: In this work, we analyze wage careers of women in Austria. We identify groups of female employees with similar patterns in their earnings development. Covariates such as e.g. the age of entry, the number of children or maternity leave help to detect these groups. We find three different types of female employees: (1) “high-wage mums”, women with high income and one or two children, (2) “low-wage mums”, women with low income and ‘many’ children and (3) “childless careers”, women who climb up the career ladder and do not have children.

We use a Markov chain clustering approach to find groups in the discrete-valued time series of income states. Additional covariates are included when modeling group membership via a multinomial logit model.

Zusammenfassung: In dieser Arbeit analysieren wir Einkommensverläufe von Frauen in Österreich. Unser Ziel ist es, mit Hilfe von zusätzlichen erklärenden Variablen wie z.B. dem Alter bei Berufseintritt, der Kinderanzahl oder einer Karenzvariable, Gruppen von erwerbstätigen Frauen mit ähnlichen Einkommensmustern zu finden. Es ergeben sich drei verschiedene Gruppen: (1) Frauen mit hohem Einkommen und ein bis zwei Kindern, (2) Frauen mit geringem Einkommen und ‘vielen’ Kindern und (3) Frauen, die im Laufe des Erwerbslebens die “Karriereleiter” hinaufklettern und keine Kinder haben.

Wir verwenden eine *Markov Chain Clustering* Methode, um auf Basis der diskreten Einkommensdaten Gruppen von Frauen zu identifizieren. Die Gruppenzugehörigkeit wird mit Hilfe eines multinomialen logistischen Modells, in dem zusätzliche erklärende Variablen berücksichtigt werden, geschätzt.

Keywords: Income Career, Transition Data, Multinomial Logit, Auxiliary Mixture Sampler, Markov Chain Monte Carlo.

1 Introduction

In this paper, we analyze individual wage careers of Austrian women who started to work between 1980 and 1985. Our administrative data set covers a long time period of about twenty years and ends in 2001.

With this work we contribute to the discussion of the amount and development of wages in Austria. One focus of the ongoing discussion of that topic is how individual attributes and decisions influence the wages of *female* employees. It is well known that female employment is characterized by a huge amount of part-time employment – in 2010

43.8 percent of female employees – as well as by a concentration on a few professions in the tertiary sector (Statistik Austria, 2011). Both of these characteristics have a large influence on the amount of earnings of women (Statistik Austria, 2009a). In the labor force survey (Statistik Austria, 2011) respondents working part-time are also asked to state the reason for their part-time instead of full-time work. The results show that for females care of children and of the elderly are the most important reasons for part-time employment. From the time-use-survey (Statistik Austria, 2009b) we know that care of children is mainly carried out by females. As a consequence of these results from official statistics we include variables concerning maternity leave and children in our study.

Many studies analyze the dependence of the amount of earnings on different attributes at one specific point of time. On the one hand, individual attributes are assumed to be correlated with individual productivity and therefore to be influential on the amount of individual earnings. On the other hand, it was found out that general economic conditions and firm-specific attributes are also important. Individual attributes, which are often included in such studies, are for example age, gender, education, professional status of occupation, whereas economic activity is an influential firm-specific attribute and the general situation of the labor market is for example reflected in the unemployment rate. For Austrian employees the interrelationship between earnings and different attributes was for example investigated on the basis of structure of earnings data by Geisberger (2007), on the basis of data on income and living conditions by Grünberger and Zulehner (2009) and for a combined data set from Austrian administrative data sources by Böheim et al. (2011).

In the present paper, we take a different approach. We are interested in the development of earnings over time and how this development depends on individual attributes. We aim at finding typical female career patterns and at classifying women into groups with similar patterns. We will see that career patterns of women highly depend on individual attributes and that variables like maternity leave or motherhood are especially important. Therefore, we concentrate on individual attributes in this work and do not take into account the economic situation. A similar analysis including individual as well as economic variables was carried out for male Austrian employees by Frühwirth-Schnatter et al. (2011). They focus their investigation on the effects of the individual's observable characteristics and labor market conditions at the time of entry into the labor market on the probability of belonging to specific career patterns.

For building and estimating our model we follow the ideas of Frühwirth-Schnatter et al. (2011). Clustering of Markov chain models is used to find homogeneous groups of the discrete-valued income data. It is assumed that the income states of all individuals of one specific group are modeled by the same group-specific transition matrix. This model was presented by Pamminer and Frühwirth-Schnatter (2010). As in Frühwirth-Schnatter et al. (2011) we additionally include covariates via a multinomial logit model in the prior of the cluster weights.

The paper is structured as follows. In Section 2, the administrative data set is introduced. Section 3 gives an overview over the model used to analyze female wage dynamics and the results are described in Section 4. The paper is summarized in Section 5.

2 Data Description

For our investigation, we consider data from the Austrian Social Security Data Base (ASSD), which contains detailed longitudinal information on employment and earnings of all private sector employees in Austria since 1972 (Zweimüller et al., 2009).

Our cohort study is based on a data set consisting of $N = 183\,805$ female Austrian employees, who entered the labor market for the first time in the years 1980 to 1985 and were 14 to at most 25 years old at entry. Two out of three women were between 17 and 19 years old at entry.

Due to a change of the qualifying conditions for maternity leave in the beginning of 2002 we had to cut the observation period after the year 2001 to omit inconsistencies concerning maternity leave as since then more ‘newly mums’ became eligible for maternity leave and hence childcare benefit (Kinderbetreuungsgeld).

As in Frühwirth-Schnatter et al. (2011) we take yearly earnings observations measured by gross monthly wages representing May. The wages are observed for a period between 2 (to have at least one transition observed) to 22 successive years for each individual. The median time length in our panel is equal to 14 years.

Following Weber (2001), we divide the gross monthly wage into six categories labeled with 0 up to 5. Category zero corresponds to zero-income, i.e. unemployment or out of labor force. The categories one to five correspond to the quintiles of the income distribution which are calculated for each year from all non-zero wages observed in that year for the total population of female employees in Austria. Note that the data set excludes marginal employees and, as hours of work are not available in our data set, we are not able to distinguish between part- and full-time jobs.

There are two important advantages by using wage categories: we do not need any inflation adjustment; and we do not have to be worried about the fact that recorded wages are right-censored because of a social security payroll tax cap.

As in Frühwirth-Schnatter et al. (2011) we cut the time series of individuals after observing more than five consecutive years with zero income, because these individuals have most likely stopped being employed. For individuals first observed in the data as apprentices, we consider their first wage after finishing education as the point of job entry.

As we are interested in characterizing the wage path of women since their first job with respect to covariates, we include the following predetermined variables:

We use a dummy for motherhood, which indicates if the (female) job entrant has ever been at least once on maternity leave (that are 72.6 % of the women in our data). This means that there was either a notice of the (live) birth of a child followed by maternity leave within nine months or an adoption leave.

The number of children indicates how often a person has had a live birth announcement. The average number of children is equal to 1.417 and the median equals 1.

We incorporate the age at entry into the model and introduce a linear and a quadratic age effect.

We include the color of the collar at job entry. 59.2 % of the women in our data started as white collar worker and 40.8 % as blue collar worker.

To include interaction effects between maternity leave and the color of the collar we build four categories out of these two variables: we define the combination blue collar &

no maternity leave as baseline category (that are 9.49 % of the women in our data), further blue collar & maternity leave (31.35 %), white collar & no maternity leave (17.90 %) and white collar & maternity leave (41.26 %).

As in Frühwirth-Schnatter et al. (2011), we also introduce dummies for the initial state (wage category) to take the initial conditions problem into account (see Appendix A).

3 The Model

A common problem in many areas of applied statistics is to identify groups of similar time series in a panel of time series. Clustering techniques are a widely used statistical tool to determine subsets in a given data set. Many clustering criteria have been developed which try to group the data as to maximize or minimize a certain numerical criterion defined solely from the data. However, distance-based clustering methods cannot be easily extended to time series data, where an appropriate distance-measure is rather difficult to define, particularly for discrete-valued time series.

Model-based clustering represents an alternative approach. It is based on formulating a clustering kernel for an individual time series \mathbf{y}_i in terms of a sampling density $p(\mathbf{y}_i|\boldsymbol{\xi})$, where $\boldsymbol{\xi}$ is an unknown model parameter. It is assumed that the N time series arise from H hidden groups, where within each group, say h , the corresponding clustering kernel $p(\mathbf{y}_i|\boldsymbol{\xi}_h)$ can be used to sufficiently describe all time series within this group.

Recently, Frühwirth-Schnatter and Kaufmann (2008) demonstrated that model-based clustering based on finite mixture models extends to time series data in a natural way. In this section we describe the approach of Frühwirth-Schnatter et al. (2011) for model-based clustering methods specifically designed for categorical (discrete-valued) time series data.

3.1 Markov Chain Clustering

Let $\mathbf{y}_i = \{y_{i1}, \dots, y_{iT_i}\}$ denote for each individual $i = 1, \dots, N$ the time series of income states, excluding the initial state y_{i0} . The categorical income variables y_{it} , $i = 1, \dots, N$, $t = 1, \dots, T_i$ take one of K states labeled $\{1, \dots, K\}$. We classify the individuals into H groups and assume that the clustering kernel $p(\mathbf{y}_i|\boldsymbol{\xi}_h)$ depends on the cluster-specific unknown parameter vector $\boldsymbol{\xi}_h$ for each of the $h = 1, \dots, H$ hidden clusters. As we introduce a group indicator $S_i \in \{1, \dots, H\}$ for each individual i we obtain $p(\mathbf{y}_i|S_i, \boldsymbol{\xi}_h) = p(\mathbf{y}_i|\boldsymbol{\xi}_{S_i})$.

For Markov chain clustering we introduce separate transition processes for each (hidden) group through a first-order time-homogeneous Markov chain model with a cluster-specific transition matrix $\boldsymbol{\xi}_h$. The elements of $\boldsymbol{\xi}_h$ are equal to

$$\xi_{h,jk} = \Pr(y_{it} = k | y_{i,t-1} = j, S_i = h), \quad j, k = 1, \dots, K.$$

Hence, each row of $\boldsymbol{\xi}_h$ represents a probability distribution over the discrete set $\{1, \dots, K\}$, i.e. $\sum_{k=1}^K \xi_{h,jk} = 1$. The clustering kernel $p(\mathbf{y}_i|\boldsymbol{\xi}_h)$ reads:

$$p(\mathbf{y}_i|\boldsymbol{\xi}_h) = \prod_{t=1}^{T_i} p(y_{it}|y_{i,t-1}, \boldsymbol{\xi}_h) = \prod_{j=1}^K \prod_{k=1}^K \xi_{h,jk}^{N_{i,jk}}, \quad (1)$$

where $N_{i,jk} = \#\{y_{it} = k, y_{i,t-1} = j\}$ is the number of transitions from state j to state k observed in time series i . Note that we condition in (1) on the first observation y_{i0} and the actual number of observations is equal to T_i for each time series.

3.2 Mixture-of-Experts Model for Modeling Prior Group Membership

We assume that individual attributes help to identify groups of female employees with similar career patterns. Therefore, we use a mixture-of-experts model (see e.g. Frühwirth-Schnatter, 2006, Chapter 8.6.3; Peng et al., 1996) to specify the prior for the individual group indicators. Prior group membership $\Pr(S_i = h)$ is modeled as a multinomial logit model (MNL):

$$\Pr(S_i = h | \beta_2, \dots, \beta_H) = \frac{\exp(\mathbf{x}_i \beta_h)}{1 + \sum_{l=2}^H \exp(\mathbf{x}_i \beta_l)}. \quad (2)$$

\mathbf{x}_i is the row vector of predetermined individual covariates, including 1 for the intercept. To determine group $h = 1$ as baseline group in the MNL we set $\beta_1 = \mathbf{0}$ and the unknown group-specific regression coefficients β_2, \dots, β_H are the effects on the log-odds ratio relative to the baseline group.

3.3 Model Estimation

For model estimation we pursue the Bayesian approach of Frühwirth-Schnatter et al. (2011) for fixed H . \mathbf{S} is estimated along with the group-specific transition matrices ξ_1, \dots, ξ_H and regression coefficients β_2, \dots, β_H from the data. Details on the initial conditions problem and on the choice of the prior may be found in Appendix A.

After choosing initial values for the group indicators \mathbf{S} we repeat the following steps:

(i) **Sample the cluster-specific transition matrices ξ_1, \dots, ξ_H given \mathbf{S} and \mathbf{y} :**

The conditionally independent rows are sampled from a total of $K \cdot H$ Dirichlet distributions:

$$\xi_{h,j \cdot} | \mathbf{S}, \mathbf{y} \sim \mathcal{D}(e_{0,j1} + N_{j1}^h(\mathbf{S}), \dots, e_{0,jK} + N_{jK}^h(\mathbf{S})),$$

$$j = 1, \dots, K, \quad h = 1, \dots, H,$$

where $N_{jk}^h(\mathbf{S}) = \sum_{i:S_i=h} N_{i,jk}$ is the total number of transitions from j to k observed in group h .

(ii) **Sample the regression coefficients β_2, \dots, β_H conditional on \mathbf{S} :**

The likelihood $p(\mathbf{S} | \beta_2, \dots, \beta_H)$ is obtained from the multinomial logit model (2). To sample β_2, \dots, β_H we apply auxiliary mixture sampling in the differenced random utility model representation as introduced by Frühwirth-Schnatter and Frühwirth (2010). Details are given in Appendix B and C.

(iii) **Sample S conditional on $\beta_2, \dots, \beta_H, \xi_1, \dots, \xi_H$ and y :**

For each individual $i = 1, \dots, N$ we draw S_i from:

$$\Pr(S_i = h | y_i, \mathbf{x}_i, \beta_2, \dots, \beta_H, \xi_1, \dots, \xi_H) \propto p(y_i | \xi_h) \cdot \frac{\exp(\mathbf{x}_i \beta_h)}{1 + \sum_{l=2}^H \exp(\mathbf{x}_i \beta_l)},$$

$$h = 1, \dots, H.$$

The MCMC algorithm is carried out for a fixed number H of clusters. This fixed number H is selected by using model selection criteria summarized and described in Frühwirth-Schnatter et al. (2011). We use the integrated classification likelihood (ICL) and the approximate weight of evidence (AWE) to determine the number of clusters H .

4 Results

We analyze patterns in the earnings development of young female labor market entrants in Austria over their life cycle. Further, we investigate the effects of the individual's observable characteristics such as number of children or maternity leave on the probability of belonging to each transition type.

Therefore, we applied Markov chain clustering for two up to four groups and simulated 10000 MCMC draws including a burn-in of 5000 draws. Finally, we applied a thinning parameter equal to 5 and used the remaining 1000 draws for posterior inference. Results were confirmed by starting the algorithm from different starting values. We used the priors specified in Appendix A. Additionally, we ran the algorithm with an uninformative prior for the transition matrices and the results proved to be extremely robust.

The analysis was carried out in R (R Development Core Team, 2011). The add-on software package is available from the authors.

4.1 Choosing the Number of Groups

To choose the number of clusters we apply two criteria, namely the integrated classification likelihood criterion (ICL) and the approximate weight of evidence criterion (AWE). Both criteria agree in choosing three groups as the optimal number of groups, see Figure 1. As this choice additionally proves to yield meaningful results from an economic point of view we carry out the rest of the analysis for this three-cluster solution.

4.2 Results of Markov Chain Clustering

Figure 2 visualizes the posterior transition probabilities ξ_1 , ξ_2 and ξ_3 . In these balloon plots, the size of each circle corresponds to the respective entry in the transition matrix. Numerical results for the posterior transition probabilities as well as their standard deviations are given in Table 3 in Appendix D. On the other hand, in Figure 3 the balloon plots represent the contingency tables. This figure indicates where the transitions mainly occur.

The largest group is the “low-wage” group with 54.62 %. Individuals of that group have a higher risk to move into (next) lower or no income categories. Transitions occur

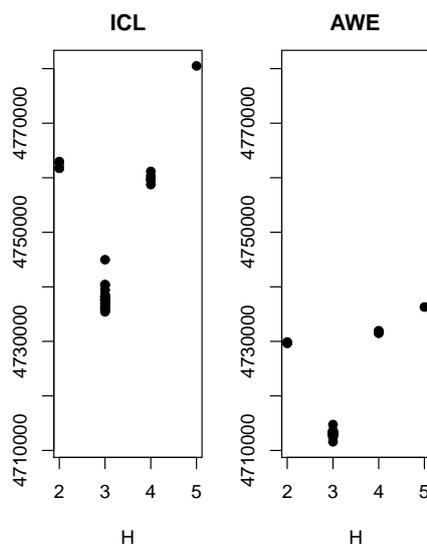


Figure 1: Model selection criteria ICL (integrated classification likelihood) and AWE (approximate weight of evidence) for various numbers H of clusters and several independent MCMC runs.

mainly in the lower wage categories. On the contrary, the “high-wage” and the “career ladder” clusters are of group size 28.02 % and 17.36 %, respectively. Women in these two clusters have a higher chance to move upwards into (the next) higher wage categories. The “career ladder” group more likely stays within the (same) wage category. We observe transitions mainly in the higher wage categories, whereas, for the “high-wage” cluster we find transitions in either end of the income distribution. This fact separates the “high-wage” from the “career ladder” group. Both end up in the higher wage categories, but contrary to the “career-ladder” group “high-wage” employees also move into the no income category.

Figure 4 shows the most typical group members and confirms the interpretation of the transition probabilities. “High-wage” women earn high wages but also spend time in the no income category whereas the “career ladder” employees climb up the ladder, what can be perfectly seen from these pictures. The “low-wage” group remains mainly in the lower quintiles of the income distribution with ‘visits’ to the no income category.

Figure 5 shows the development of the distribution of the wage categories over the years starting in the initial distribution (determined from the initial states in the data) up to the steady state. The majority of the “career ladder” group ends up in the highest wage category whereas the “low-wage” women stay within the lower income categories as well as in the no income state. The “high-wage” employees show some kind of mixed tendencies: approximately 50 % end up in the two highest wage categories and approximately 30 % in the lowest or no income category. The “high-wage” and “low-wage” employees reach the steady state within a decade, whereas the “career ladder” individuals maybe will never reach the steady state within their working life.

An important criterion indicating the quality of the clustering procedure is the seg-

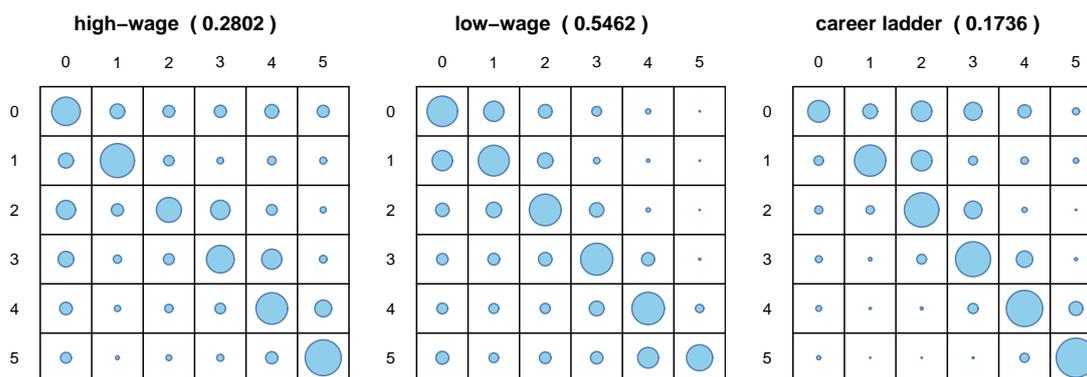


Figure 2: Balloon plots of posterior expectation of the transition matrices ξ_1 , ξ_2 and ξ_3 obtained by Markov chain clustering. The circular areas are proportional to the size of the corresponding entry in the transition matrix. The corresponding group sizes are calculated based on the posterior classification probabilities and are indicated in parentheses.

mentation power (which equals one minus the misclassification risk) shown in Table 1: three out of four individuals are assigned with at least 73.8 % probability to their respective cluster. The segmentation power for the “low-wage” cluster is much higher than for the other groups, whereas the segmentation power for the “high-wage” group is slightly lower than for the “career ladder” cluster. Three in four “low-wage” women are assigned with at least 79.6 % to their group. Three-quarters of the “high-wage” individuals have an assignment probability of at least 65.7 % and three-quarters of the “career ladder” women have an assignment probability of at least 67.9 %.

Table 1: Segmentation power of Markov chain clustering; reported are the lower quartile, the median and the upper quartile of the individual posterior classification probabilities for all individuals within a certain cluster as well as for all individuals.

	Markov chain clustering		
	1st Qu.	Median	3rd Qu.
“high-wage”	0.6574	0.8613	0.9770
“low-wage”	0.7956	0.9323	0.9843
“career ladder”	0.6790	0.8703	0.9674
overall	0.7378	0.9110	0.9807

4.3 Results of the Multinomial Logit Model

Table 2 shows the posterior regression coefficients for the MNL (2). We choose the “low-wage” group as the baseline group. The 95 % credibility intervals do not cover zero for any coefficient.

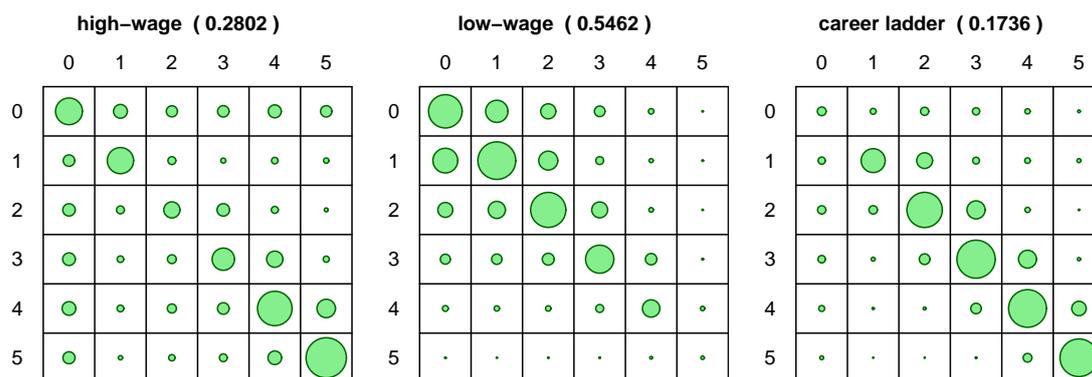


Figure 3: Balloon plots of cluster-specific contingency tables. For each cluster we find in cell (j, k) the probability $\Pr(y_{i,t-1} = j, y_{it} = k | S_i = h)$ of observing wage categories (j, k) in consecutive years for an individual in this cluster. The entries of this table sum to one.

The interaction model shows that blue collar workers who were at least once on maternity leave (compared to blue collar workers who have never been on maternity leave) more likely belong to the “low-wage” group than to the other groups. White collar workers who have never been on maternity leave (compared to blue collar workers who have never been on maternity leave) belong more likely to the “high-wage” or the “career ladder” cluster than to the “low-wage” cluster, whereas white collar workers having been on maternity leave have higher probability to belong to the “high-wage” group than to the “low-wage” group and much less likely to the “career ladder” than to the “low-wage” group. To sum up, having been on maternity leave reduces the chance to belong either to the “high-wage” or to the “career ladder” group (instead of belonging to the “low-wage” group) except for white collar workers.

Not surprisingly, with growing number of children it becomes less likely to belong to either “high-wage” or “career ladder” group. The ex post analyses show that the median number of children within group “career ladder” is equal to 0, in the “high-wage” group 1 and in the “low-wage” group 2, which confirms these results.

We include the age and the quadratic age effect in our model and find that with growing age at entry the allocation to either “high-wage” or “career ladder” group is more likely.

Finally, the higher a woman starts in the income distribution the more likely she belongs to the “high-wage” or to the “career ladder” group instead of belonging to the “low-wage” group.

To summarize the interpretation of the estimation results we finally found the following ‘labeling’ for the three types of earnings careers:

“**high-wage mums:**” women with high income and one or two children

“**low-wage mums:**” women with low income and ‘many’ children

“**childless careers:**” women who climb up the career ladder and do not have children

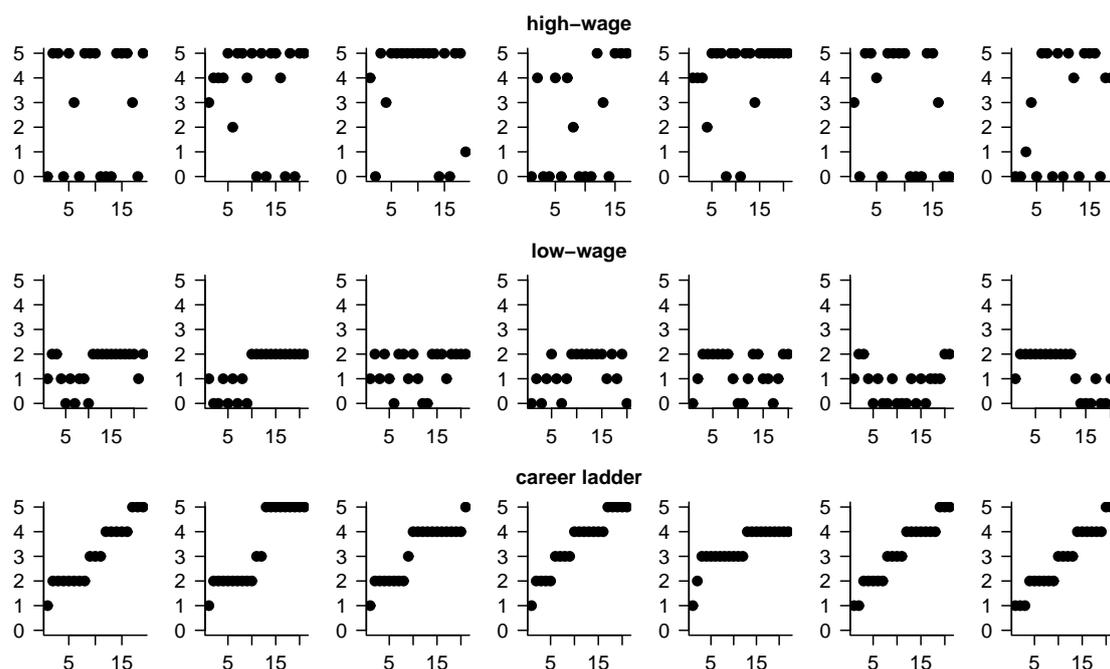


Figure 4: Wage careers of typical group members showing the highest classification probabilities.

4.4 A Comparison With the Results for Male Labor Market Entrants

Frühwirth-Schnatter et al. (2011) analyze patterns in the earnings development of young *male* labor market entrants in Austria over their life cycle. They identify four distinctly different types of transition patterns. Further, they investigate the effects of the individual's observable characteristics and labor market conditions at the time of entry on the probability of belonging to each transition type.

The four types of earnings careers they identified are: “upward”, which are ‘social climbers’ with consistent upward mobility having a higher chance to move forward and/or out of the no-income category; “static”, a cluster consisting of employees having stable employment relationships and high persistence, respectively; “downward”, a group of ‘losers’ with high downward mobility as well as high risk to move into the next lower wage category and high risk to drop into and/or remain in the no-income category; and finally “mobile”, which move easily between wage categories and have a higher chance to move forward.

It turned out that the individual attributes education, age at entry and type of the first job as well as the economic conditions at the year of entry, which are reflected in the unemployment rate, are influential on the earnings careers of males. Including these variables leads to a clustering of the male time series with high quality. Three out of four individuals were assigned with at least 63.8 % probability to their respective cluster. In the present paper we apply the clustering approach to female wage career data. Interestingly the most important variables that help to find the three clusters are variables that are related to children. As soon as we include at least one of the two variables maternity leave

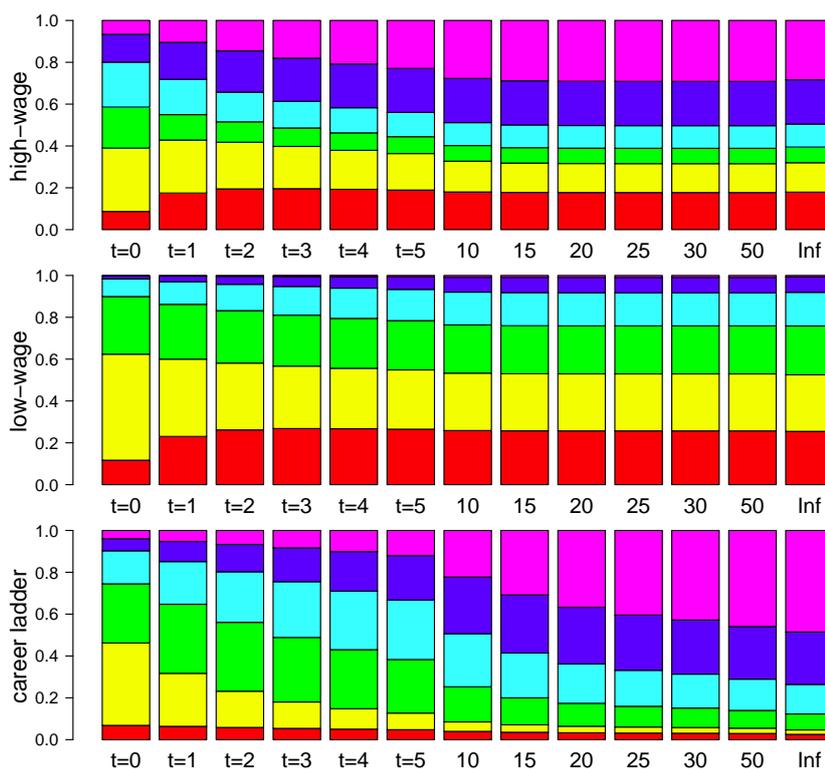


Figure 5: Posterior expectation of the wage distribution for wage category 0 (bottom) to category 5 (top) after a period of t years in the various clusters.

or number of children the three clusters can be identified. The quality of the female model is even higher than the one for the male data. Three out of four females are assigned with at least 73.8 % probability to their respective cluster.

5 Conclusions

In this paper, we investigated wage careers of women in Austria. More precisely, we analyzed patterns in the earnings development of female employees in order to identify groups of women with similar transition patterns between the discrete earnings states. We were interested whether covariates such as age at entry, maternity leave, collar of work, etc., which are commonly known to be influential on the amount of the wage, help to find those groups of employees with similar wage careers.

The results, which are based on a large administrative data set, suggest that motherhood in general, the number of children, the color of the collar and the entry age have a strong impact on the mobility pattern of women throughout their working life. We are able to identify three distinctly different types of female employees: (1) “high-wage mums”, women with high income and one or two children; (2) “low-wage mums”, women with low income and ‘many’ children and (3) “childless careers”, women climbing up the career ladder and do not have children.

Table 2: Multinomial logit model to explain group membership in a particular cluster (baseline: “low-wage” cluster); the numbers are the posterior expectation and, in parenthesis, the posterior standard deviation of the various regression coefficients.

	“high-wage”	“career ladder”
Intercept	7.988 (0.604)	5.051 (0.902)
blue collar × maternity leave	−1.111 (0.049)	−6.609 (0.184)
white collar × no maternity leave	2.383 (0.076)	2.661 (0.067)
white collar × maternity leave	0.430 (0.042)	−1.800 (0.094)
Number of children	−0.227 (0.014)	−0.638 (0.030)
Age at start	−1.101 (0.064)	−0.536 (0.096)
Age at start (squared)	0.034 (0.002)	0.013 (0.003)
Start in wage category 1	−0.084 (0.035)	0.246 (0.056)
Start in wage category 2	−0.209 (0.035)	0.594 (0.059)
Start in wage category 3	0.847 (0.041)	1.310 (0.065)
Start in wage category 4	1.804 (0.054)	1.951 (0.088)
Start in wage category 5	2.051 (0.104)	2.587 (0.133)

We found out that female wage careers and the identification of our three groups are strongly influenced by the fact whether a woman is mother and/or was on maternity leave. This finding is supported by administrative data about earnings in Austria. Thereby our study also contributes to the ongoing discussion about high quality child care places that allow females to fully take part in the labor market.

To estimate our statistical model we use Markov chain clustering, proposed by Paminger and Frühwirth-Schnatter (2010), which is a model-based clustering approach for grouping discrete-valued time series obtained by observing a categorical variable with several states. This method is based on finite mixtures of first-order time-homogeneous Markov chain models. In order to analyze group membership in dependence on additional covariates, we follow Frühwirth-Schnatter et al. (2011) and include a probabilistic model for the latent group indicators within the Bayesian classification rule using a multinomial logit model.

The choice of the number of groups was based on the AWE (approximate weight of evidence) and ICL (integrated classification likelihood) to take also into account the quality of the partitioning. Remarkably, the segmentation power is quite high: three out of four individuals are assigned with at least 73.8 % probability to their respective cluster.

Acknowledgments

We thank Sylvia Frühwirth-Schnatter for her helpful comments and discussions as well as Bettina Grün and her colleagues at the Department of Applied Statistics (IFAS) at the Johannes Kepler University, Linz, Austria, for their remarks.

The research was funded by the Austrian Science Fund (FWF): S10309-G16 (NRN “The Austrian Center for Labor Economics and the Analysis of the Welfare State”, Sub-project “Bayesian Econometrics”).

References

- Böheim, R., Himpele, K., Mahringer, H., and Zulehner, C. (2011). The gender pay gap in Austria: Tamensi movetur! *WIFO Working Papers*, 394/2011, 194–209.
- Frühwirth-Schnatter, S. (2006). *Finite Mixture and Markov Switching Models*. New York: Springer.
- Frühwirth-Schnatter, S., and Frühwirth, R. (2010). Data augmentation and MCMC for binary and multinomial logit models. In T. Kneib and G. Tutz (Eds.), *Statistical Modelling and Regression Structures – Festschrift in Honour of Ludwig Fahrmeir* (pp. 111–132). Heidelberg: Physica-Verlag. (Also available at http://www.ifas.jku.at/e2550/e2756/index_ger.html, IFAS Research Paper Series 2010-48)
- Frühwirth-Schnatter, S., and Kaufmann, S. (2008). Model-based clustering of multiple time series. *Journal of Business & Economic Statistics*, 26, 78–89.
- Frühwirth-Schnatter, S., Pamminger, C., Weber, A., and Winter-Ebmer, R. (2011). Labor market entry and earnings dynamics: Bayesian inference using mixtures-of-experts Markov chain clustering. *Journal of Applied Econometrics*, forthcoming. Available from <http://dx.doi.org/10.1002/jae.1249>
- Geisberger, T. (2007). Geschlechtsspezifische Lohn- und Gehaltsunterschiede. *Statistische Nachrichten*, 7, 633–642.
- Grünberger, K., and Zulehner, C. (2009). Geschlechtsspezifische Lohnunterschiede in Österreich. *Wifo Monatsberichte*, 2, 139–150.
- McFadden, D. (1974). Conditional logit analysis of qualitative choice behaviour. In P. Zarembka (Ed.), *Frontiers of Econometrics* (pp. 105–142). New York: Academic.
- Pamminger, C., and Frühwirth-Schnatter, S. (2010). Model-based clustering of categorical time series. *Bayesian Analysis*, 5(2), 345–368. Available from <http://ba.stat.cmu.edu/journal/2010/vol05/issue02/pamminger.pdf>
- Peng, F., Jacobs, R. A., and Tanner, M. A. (1996). Bayesian inference in mixtures-of-experts and hierarchical mixtures-of-experts models with an application to speech recognition. *Journal of the American Statistical Association*, 91, 953–960.
- R Development Core Team. (2011). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. Available from <http://www.R-project.org/> (ISBN 3-900051-07-0)
- Statistik Austria. (2009a). *Verdienststrukturerhebung 2006 – Struktur und Verteilung der Verdienste in Österreich*. Wien.
- Statistik Austria. (2009b). *Zeitverwendung 2008/09 – Ein Überblick über geschlechtsspezifische Unterschiede*. Wien.
- Statistik Austria. (2011). *Arbeitskräfteerhebung 2010, Ergebnisse des Mikrozensus*. Wien.
- Weber, A. (2001). *State dependence and wage dynamics: A heterogeneous Markov chain model for wage mobility in Austria* (Research report). Vienna: Institute for Advanced Studies.
- Zweimüller, J., Winter-Ebmer, R., Lalive, R., Kuhn, A., Wuellrich, J.-P., Ruf, O., et al. (2009). *The Austrian Social Security Database (ASSD)* (Working Paper No. 0903).

Linz, Austria: NRN: The Austrian Center for Labor Economics and the Analysis of the Welfare State.

Appendix

A The Initial Conditions Problem and the Prior

Frühwirth-Schnatter et al. (2011) found a simple solution to the initial conditions problem for the discrete case present in this paper. They allow for dependence between the initial states y_{i0} and the group indicators S_i which capture unobserved heterogeneity. The initial states are included in the design matrix \mathbf{x}_i of equation (2). Since the two distributions $p(S_i|y_{i0}, \cdot)$ and $p(y_{i0}|\cdot)$ appearing in the factorization of the joint distribution $p(y_{i0}, S_i|\cdot) = p(S_i|y_{i0}, \cdot) \cdot p(y_{i0}|\cdot)$ do not have parameters in common the marginal distribution $p(y_{i0}|\cdot)$ cancels out and we do not need to specify it.

We assume prior independence between ξ_1, \dots, ξ_H and β_2, \dots, β_H . The a priori independent regression coefficients β_{hj} follow a standard normal prior distribution and the a priori independent K rows $\xi_{h,1}, \dots, \xi_{h,K}$ of ξ_h following a Dirichlet distribution, i.e. $\xi_{h,j} \sim \mathcal{D}(e_{0,j1}, \dots, e_{0,jK})$, with prior parameters $\mathbf{e}_{0,j} = (e_{0,j1}, \dots, e_{0,jK}) = N_0 \cdot \xi_j^*$, where $N_0 = 10$ and

$$\xi^* = \begin{pmatrix} 0.7 & 0.2 & 0.025 & 0.025 & 0.025 & 0.025 \\ 0.15 & 0.6 & 0.15 & 0.03 & 0.03 & 0.03 \\ 0.03 & 0.15 & 0.6 & 0.15 & 0.03 & 0.03 \\ 0.03 & 0.03 & 0.15 & 0.6 & 0.15 & 0.03 \\ 0.03 & 0.03 & 0.03 & 0.15 & 0.6 & 0.15 \\ 0.025 & 0.025 & 0.025 & 0.025 & 0.2 & 0.7 \end{pmatrix}$$

as in Frühwirth-Schnatter et al. (2011).

B MNL as Random Utility Model (RUM)

We rewrite model (2) as random utility model (RUM) as it was introduced by McFadden (1974). Let y_{hi}^u be the utility of choosing group h . Then

$$y_{hi}^u = \mathbf{x}_i \beta_h + \delta_{hi}, \quad h = 1, \dots, H$$

$$S_i = h \Leftrightarrow y_{hi}^u = \max_{l \in \{1, \dots, H\}} y_{li}^u, \quad i = 1, \dots, N.$$

If $\delta_{1i}, \dots, \delta_{Hi}$ are i.i.d. following a type I extreme value distribution, the MNL (2) results as the marginal distribution of S_i .

An alternative way to write the MNL as an augmented model involving random utilities is the differenced RUM (dRUM), which is obtained by choosing a baseline category (here $h_0 = 1$) and considering the model involving the differences of the utilities: $z_{hi} = \mathbf{x}_i \beta_h + \varepsilon_{hi}$, where $z_{hi} = y_{hi}^u - y_{1i}^u$. Marginally, the errors $\varepsilon_{hi} = \delta_{hi} - \delta_{1i}$ follow a logistic distribution but are no longer independent across categories.

It can be shown (see Frühwirth-Schnatter and Frühwirth, 2010) that for each h , the MNL has the following representation as partial (binary) dRUM:

$$z_{hi} = \mathbf{x}_i \boldsymbol{\beta}_h - \log \left(\sum_{l \neq h} \lambda_{li} \right) + \varepsilon_{hi}, \quad (3)$$

where ε_{hi} , $h \neq 1$ are now i.i.d. following a logistic distribution.

C Auxiliary Mixture Sampling

The logistic distribution is approximated for each ε_{hi} by a finite scale mixture of six normal distributions with zero means, variances s_r^2 and weights w_r , $r = 1, \dots, 6$ (for the values of the parameters see Frühwirth-Schnatter and Frühwirth, 2010). Conditional on the latent utilities $\mathbf{z} = \{z_{2i}, \dots, z_{Hi}, i = 1, \dots, N\}$ and the indicators $\mathbf{R} = \{r_{2i}, \dots, r_{Hi}, i = 1, \dots, N\}$ the dRUM (3) reduces to a Gaussian regression model:

$$z_{hi} = \mathbf{x}_i \boldsymbol{\beta}_h - \log \left(\sum_{l \neq h} \lambda_{li} \right) + \varepsilon_i, \quad \varepsilon_i | r_{hi} \sim \mathcal{N}(0, s_{r_{hi}}^2). \quad (4)$$

Thus, we select starting values for \mathbf{z} and \mathbf{R} and repeat the following steps:

- (ii-a) Sample the regression coefficients $\boldsymbol{\beta}_2, \dots, \boldsymbol{\beta}_H$ conditional on \mathbf{z} and \mathbf{R} based on the normal regression model (4) from a multivariate normal density.
- (ii-b) Sample the latent variables z_{hi} and r_{hi} conditional on $\boldsymbol{\beta}_2, \dots, \boldsymbol{\beta}_H$ and \mathbf{S} for $i = 1, \dots, N$ and $h = 2, \dots, H$ with $\lambda_{hi} = \exp(\mathbf{x}_i \boldsymbol{\beta}_h)$:
 - (ii-b-1) Sample all utilities z_{2i}, \dots, z_{Hi} simultaneously for each i from:

$$z_{hi} = \log(\lambda_{hi}^* U_{hi} + I\{S_i = h\}) - \log(1 - U_{hi} + \lambda_{hi}^* I\{S_i \neq h\}),$$

where $U_{ih} \sim \mathcal{U}[0, 1]$ and $\lambda_{hi}^* = \lambda_{hi} / (\sum_{l \neq h} \lambda_{li})$.

- (ii-b-2) Sample the component indicators r_{hi} conditional on z_{hi} from:

$$\Pr(r_{hi} = j | z_{hi}, \boldsymbol{\beta}_h) \propto \frac{w_j}{s_j} \exp \left\{ -\frac{1}{2} \left(\frac{z_{hi} - \mathbf{x}_i \boldsymbol{\beta}_h + \log \left(\sum_{l \neq h} \lambda_{li} \right)}{s_j} \right)^2 \right\}.$$

D Numerical Results for the Transition Probabilities

Table 3: Posterior expectation $E(\xi_h|y)$ and, in parenthesis, posterior standard deviations $SD(\xi_h|y)$ (multiplied by 100) of the transition matrices ξ_h in the various clusters.

"high-wage"						
	0	1	2	3	4	5
0	0.489(0.691)	0.131(0.473)	0.086(0.218)	0.090(0.198)	0.115(0.242)	0.089(0.200)
1	0.135(0.667)	0.702(1.213)	0.064(0.435)	0.026(0.151)	0.044(0.227)	0.029(0.147)
2	0.219(0.679)	0.090(0.904)	0.375(1.235)	0.226(0.726)	0.070(0.227)	0.022(0.091)
3	0.148(0.283)	0.040(0.230)	0.072(0.206)	0.460(0.715)	0.245(0.377)	0.035(0.091)
4	0.095(0.123)	0.023(0.063)	0.042(0.088)	0.068(0.118)	0.601(0.294)	0.172(0.180)
5	0.070(0.102)	0.009(0.030)	0.019(0.049)	0.030(0.063)	0.090(0.132)	0.782(0.263)
"low-wage"						
	0	1	2	3	4	5
0	0.564(0.272)	0.248(0.214)	0.115(0.116)	0.057(0.074)	0.015(0.046)	0.001(0.012)
1	0.248(0.335)	0.574(0.546)	0.145(0.229)	0.025(0.052)	0.007(0.030)	0.001(0.011)
2	0.110(0.111)	0.145(0.197)	0.608(0.214)	0.125(0.130)	0.011(0.028)	0.001(0.008)
3	0.078(0.102)	0.084(0.130)	0.111(0.123)	0.621(0.265)	0.103(0.123)	0.003(0.019)
4	0.071(0.145)	0.059(0.126)	0.062(0.136)	0.133(0.202)	0.635(0.335)	0.041(0.155)
5	0.101(0.569)	0.056(0.386)	0.081(0.532)	0.092(0.536)	0.261(1.069)	0.409(1.563)
"career ladder"						
	0	1	2	3	4	5
0	0.288(1.391)	0.131(1.023)	0.252(0.892)	0.196(0.642)	0.104(0.482)	0.029(0.335)
1	0.054(0.464)	0.591(2.031)	0.258(1.322)	0.048(0.319)	0.032(0.282)	0.017(0.203)
2	0.040(0.138)	0.041(0.256)	0.714(0.351)	0.189(0.316)	0.016(0.064)	0.001(0.021)
3	0.028(0.083)	0.008(0.052)	0.058(0.121)	0.738(0.261)	0.163(0.198)	0.005(0.035)
4	0.020(0.081)	0.003(0.028)	0.005(0.040)	0.062(0.113)	0.794(0.221)	0.116(0.176)
5	0.010(0.080)	0.001(0.018)	0.001(0.017)	0.002(0.028)	0.050(0.148)	0.936(0.182)

Authors' addresses:

Christoph Pamminger

Institute for Statistics and Mathematics

Vienna University of Economics and Business

Augasse 2-6

A-1090 Wien

E-Mail: Christoph.Pamminger@gmail.com

Regina Tüchler

Wirtschaftskammer Österreich

Stabsabteilung Statistik

Wiedner Hauptstraße 63

A-1045 Wien

E-Mail: Regina.Tuechler@wko.at