Multilevel Discrete-Time Event History Analysis

Fiona Steele

Centre for Multilevel Modelling
University of Bristol, UK

July 2010

Overview of Course

- Multilevel discrete-time models for recurrent events
  - Time-to-event (duration) models
- Models for transitions between states
  - Event history models for transitions between two states
  - Multiple states and competing risks models
  - State dependence (autoregressive) models
- Methods to handle endogenous predictors
  - Simultaneous equation models

Modelling Recurrent Events

Multilevel Event History Data

Multilevel event history data arise when events are repeatable (e.g. births, partnership dissolution) or individuals are organised in groups.

Suppose events are repeatable, and define an **episode** as a continuous period for which an individual is at risk of experiencing an event, e.g.

Event	Episode duration
Birth	Duration between birth $k-1$ and birth k
Marital dissolution	Duration of marriage

Denote by y_{ij} the duration of episode i of individual j, which is fully observed if an event occurs ($\delta_{ij} = 1$) and right-censored if not ($\delta_{ij} = 0$).

Discrete-Time Data

In this course, we focus on discrete-time methods.

In social research, event history data are usually collected:

- retrospectively in a cross-sectional survey, where dates are recorded to the nearest month or year, *OR*
- prospectively in irregularly-spaced waves of a panel study (e.g. annually)

Both give rise to discretely-measured durations.

We can convert the observed data (y_{ij}, δ_{ij}) to a sequence of binary responses $\{y_{tij}\}$ where y_{tij} indicates whether an event has occurred in time interval $[t, t+1)$.

Data Structure: The Person-Period-Episode File

individual j	episode i	y_{ij}	δ_{ij}
1	1	2	1
1	2	3	0

$$\downarrow$$

individual j	episode i	t	y_{tij}
1	1	1	0
1	1	2	1
1	2	1	0
1	2	2	0
1	2	3	0

Discrete-Time Hazard Function

Denote by p_{tij} the probability that individual j has an event during interval t of episode i, given that no event has occurred before the start of t.

$$p_{tij} = Pr(y_{tij} = 1 | y_{t-1,ij} = 0)$$

p_{tij} is a discrete-time approximation to the continuous-time hazard function.

Call p_{tij} the discrete-time hazard function.

Problem with Analysing Recurrent Events

We cannot assume that the durations of episodes from the same individual are independent.

There may be unobserved individual-specific factors (i.e. constant across episodes) which affect the hazard of an event for **all** episodes, e.g. 'taste for stability' may influence risk of leaving a job.

The presence of such unobservables, and failure to account for them in the model, will lead to correlation between durations of episodes from the same individual.

Multilevel Discrete-time Model for Recurrent Events

Multilevel (random effects) discrete-time logit model:

$$\log\left(\frac{p_{tij}}{1-p_{tij}}\right) = \mathbf{D}_{tij}\alpha + \mathbf{x}_{tij}\beta + u_j$$

p_{tij} is the probability of an event during interval t

$\mathbf{D}_{tij}$ is a vector of functions of the cumulative duration by interval t with coefficients α

$\mathbf{x}_{tij}$ a vector of covariates (time-varying or defined at the episode or individual level) with coefficients β

$u_j \sim N(0, \sigma_u^2)$ allows for unobserved heterogeneity ('shared frailty') between individuals due to time-invariant omitted variables

Multilevel Model for Recurrent Events: Notes

- The model for recurrent events is essentially the same as the (single-level) model for unobserved heterogeneity
  - Both can be estimated using multilevel modelling software
- Recurrent events allow better identification of the random effect variance σ_u^2
- Allow for non-proportional effects of covariate x by including interaction between x and functions of t in $\mathbf{D}$
- Can allow duration and covariate effects to vary across episodes
  - Include a dummy for order of event and interact with t and x

Modelling the Time-Dependency of the Hazard

Changes in p_{tij} with t are captured in the model by $\mathbf{D}_{tij}\alpha$, the baseline hazard function.

$\mathbf{D}_{tij}$ has to be specified by the user. Options include:

Polynomial of order p

$$\mathbf{D}_{tij}\alpha = \alpha_0 + \alpha_1 t + \ldots + \alpha_p t^p$$

Step function

$$\mathbf{D}_{tij}\alpha = \alpha_1 D_1 + \alpha_2 D_2 + \ldots + \alpha_q D_q$$

where $D_1, \ldots, D_q$ are dummies for time intervals $t = 1, \ldots, q$ and q is the maximum observed event time. If q large, categories may be grouped to give a piecewise constant hazard model.

Testing for Unobserved Heterogeneity ($\mathbf{H}_0 : \sigma_u^2 = 0$)

Likelihood ratio test

- Preferred method if maximum likelihood (usually via numerical quadrature) used
- Compare multilevel model ($\sigma_u^2 \neq 0$) with single-level model ($\sigma_u^2 = 0$), and compare difference in model deviances to χ^2 on 1 d.f. Take p-value/2 for one-sided test as σ_u^2 must be > 0
- Software: Stata (`xtlogit`, `xtmelogit`, `gllamm`), aML, SAS

Bayesian credible intervals for σ_u^2

- Analogous to classical confidence intervals
- Available if model estimated using Markov Chain Monte Carlo (MCMC) methods
- Software: MLwiN and WinBUGS

Example: Women's Employment Transitions

- Analyse duration of non-employment (unemployed or out of labour market) episodes
  - Event is entry (1st episode) or re-entry (2nd + episodes) into employment
- Data are subsample from British Household Panel Study (BHPS): 1399 women and 2284 episodes
- Durations grouped into years $\Rightarrow$ 15,297 person-year records
- Baseline hazard is step function with yearly dummies for durations up to 9 years, then single dummy for 9+ years
- Covariates include time-varying indicators of number and age of children, age, marital status and characteristics of previous job (if any)

Multilevel Logit Results for Transition to Employment: Baseline Hazard and Unobserved Heterogeneity

Variable	Est.	(se)
Duration non-employed (ref is $<$ 1 year)		
[1,2) years	−0.646*	(0.104)
[2,3)	−0.934*	(0.135)
[3,4)	−1.233*	(0.168)
[4,5)	−1.099*	(0.184)
[5,6)	−0.944*	(0.195)
[6,7)	−1.011*	(0.215)
[7,8)	−1.238*	(0.249)
[8,9)	−1.339*	(0.274)
$\geq$ 9 years	−1.785*	(0.175)
σ_u (SD of woman random effect)	0.662*	(0.090)

* $p < 0.5$

Multilevel Logit Results for Transition to Employment: Presence and Age of Children

Variable	Est.	(se)
Imminent birth (within 1 year)	−0.842*	(0.125)
No. children age $\leq$ 5 yrs (ref=0)		
1 child	−0.212*	(0.097)
$\geq$ 2	−0.346*	(0.143)
No. children age $>$ 5 yrs (ref=0)		
1 child	0.251	(0.118)
$\geq$ 2	0.446*	(0.117)

* $p < 0.5$

Multilevel Logit Analysis of Employment: Main Conclusions

- Unobserved heterogeneity. Significant variation between women. Deviance = 23.5 on 1 df; $p<0.01$
- Duration effects. Probability of getting a job decreases with duration out of employment
- Presence/age of children. Probability of entering employment lower for women who will give birth in next year or with young children, but higher for those with older children
- Other covariates. Little effect of age, but increased chance of entering employment for women who are cohabiting, have previously worked, whose last job was full-time, and whose occupation is 'professional, managerial or technical'

Impact of Adding Individual Random Effects

Coefficients in a random effects (RE) model may be different from those in the corresponding single-level (SL) model for two reasons:

- RE model allows for selection effect due to unobserved heterogeneity
  - Most impact on duration effects α
  - Estimates of α can increase or decrease depending on direction of α
- Scaling effect of introducing random effect
  - Impact on both α and β
  - Estimates will usually increase in magnitude

Consequences of Unobserved Heterogeneity

If there are individual-specific unobserved factors that affect the hazard, the observed form of the hazard function at the aggregate population level will tend to be different from the individual-level hazards.

For example, even if the hazards of individuals in a population are constant over time, the population hazard (averaged across individuals) will be time-dependent, typically decreasing. This may be explained by a selection effect operating on individuals.

Selection Effect of Unobserved Heterogeneity

If a population is heterogeneous in its susceptibility to experiencing an event, high risk individuals will tend to have the event first, leaving behind lower risk individuals.

Therefore as t increases the population is increasingly depleted of those individuals most likely to experience the event, leading to a decrease in the population hazard.

Because of this selection, we may see a decrease in the population hazard even if individual hazards are constant (or even increasing).

Illustration of Selection for Constant Individual Hazards

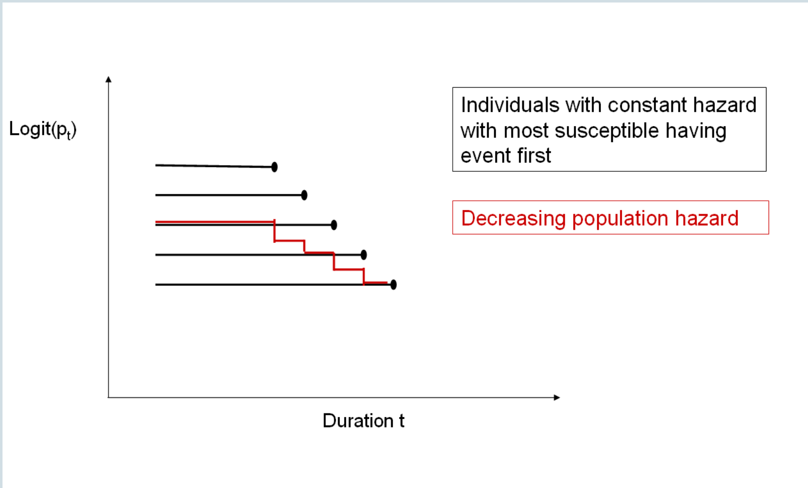

Impact of Unobserved Heterogeneity on Duration Effects

If unobserved heterogeneity is incorrectly ignored:

- A positive duration dependence will be understated (so positive α become more strongly positive when u_j added)
- A negative duration dependence will be overstated

Changes in β are more likely to be due to scaling.

Note also that coefficients from random effects and single-level models have a different interpretation (see later).

Scaling Effect of Introducing u_j (1)

To see the scaling effect, consider the latent variable (threshold) representation of the discrete-time logit model.

Consider a latent continuous variable y^* that underlies observed binary y such that:

$$y_{tij} = \begin{cases} 1 & \text{if } y^*_{tij} \geq 0 \\ 0 & \text{if } y^*_{tij} < 0 \end{cases}$$

Threshold model

$$y^*_{tij} = \mathbf{D}_{tij}\alpha + \mathbf{x}_{tij}\beta + u_j + e^*_{tij}$$

- $e^*_{tij} \sim$ standard logistic (with variance $\simeq 3.29$) $\rightarrow$ logit model
- $e^*_{tij} \sim N(0, 1) \rightarrow$ probit model

So the level 1 residual variance, $\text{var}(e^*_{tij})$, is fixed.

Scaling Effect of Introducing u_j (2)

Single-level logit model expressed as a threshold model:

$$y^*_{ti} = \mathbf{D}_{ti}\alpha + \mathbf{x}_{ti}\beta + e^*_{ti}$$

$$\text{var}(y^*_{ti}|x_{ti}) = \text{var}(e^*_{ti}) = 3.29$$

Now add random effects:

$$y^*_{tij} = \mathbf{D}_{tij}\alpha + \mathbf{x}_{tij}\beta + u_j + e^*_{tij}$$

$$\text{var}(y^*_{tij}|x_{tij}) = \text{var}(u_j) + \text{var}(e^*_{tij}) = \sigma^2_u + 3.29$$

Adding random effects has increased the residual variance
$\rightarrow$ scale of y^* stretched out
$\rightarrow$ α and β increase in absolute value.

Scaling Effect of Introducing u_j (3)

Denote by β^{RE} the coefficient from a random effects model, and β^{SL} the coefficient from the corresponding single-level model.

The approximate relationship between these coefficients (for a logit model) is:

$$\beta^{RE} = \beta^{SL}\sqrt{\frac{\sigma^2_u + 3.29}{3.29}}$$

Replace 3.29 by 1 to get expression for relationship between probit coefficients.

Note that the same relationship would hold for duration effects α if there was no selection effect. In general, both selection and scaling effects will operate on α.

Interpretation of Coefficients (1)

Suppose we have a continuous x with coefficient β.

In single-level model, β compares the log-odds of an event for two randomly selected individuals with x-values 1 unit apart (and with the same values for all other covariates). $\exp(\beta)$ is the odds ratio. We call β from this model the population-averaged or marginal effect of x.

In a random effects model, $\exp(\beta)$ compares the odds of an event for two hypothetical individuals with the same value of u_j. If x varies within individuals, β is the effect on the log-odds of a 1-unit increase in x for a given individual. We call β from this model the cluster-specific or individual-specific effect of x.

For an individual-level x (i.e. with no within-individual variation), the population-averaged effect may be of more interest.

Population-Averaged Predicted Probabilities (1)

Fortunately we can calculate predicted probabilities from a random effects model that have a population-averaged interpretation.

The probability of an event in interval t of episode i for individual j is:

$$p_{tij} = \frac{\exp(\mathbf{D}_{tij}\alpha + \mathbf{x}_{tij}\beta + u_j)}{1 + \exp(\mathbf{D}_{tij}\alpha + \mathbf{x}_{tij}\beta + u_j)}$$

where we substitute estimates of α, β, and u_j to get predicted probabilities.

Rather than calculating probabilities for each record tij, however, we often want predictions for specific values of $\mathbf{x}$. But what do we substitute for u_j?

Population-Averaged Predicted Probabilities (2)

Suppose we want predictions for a particular combination of values of $\mathbf{x}$ denoted by $\mathbf{x}^*$. What do we do about the individual random effects u?

1. Substitute the mean $u = 0$. But predictions are not the mean response probabilities for $\mathbf{x} = \mathbf{x}^*$. Because p_t is a nonlinear function of u, the value of p_t at mean of $u \neq$ mean of p_t. Predictions at $u = 0$ are medians.
2. Integrate out u to obtain an expression for mean p_t that does not involve u. Leads to mean predicted p_t that have a PA interpretation, but requires some approximation.
3. Average over simulated values of u. Also gives PA probabilities, but easier to implement.

Population-Averaged Predictions via Simulation

Suppose we have 2 covariates, x_1 and x_2, and we want mean p_t for values of x_1 holding x_2 constant.

To get predictions for $t = 1, \ldots, q$ and $x_1 = 0, 1$:

1. Set $t = 1$ and x_{1tij}=0 for each record tij, retaining observed x_{2tij}
2. Generate u_j for each individual j from $N(0, \hat{\sigma}_u^2)$
3. Compute predicted p_t for each record tij based on $x_{1tij} = 0$, observed x_{2tij}, generated u_j, and $(\hat{\alpha}, \hat{\beta})$
4. Take mean of predictions to get mean p_t for $t = 1$ and $x_1 = 0$
5. Repeat 1-4 for $t = 2, \ldots, q$
6. Repeat 1-5 for $x_{1tij} = 1$

Transitions to Employment: Duration Effects

Variable	Single-level Est.	(se)	Random effects Est.	(se)
Duration non-employed				
(ref is < 1 year)				
[1,2) years	−0.788*	(0.095)	−0.646*	(0.104)
[2,3)	−1.144*	(0.122)	−0.934*	(0.135)
[3,4)	−1.499*	(0.154)	−1.233*	(0.168)
[4,5)	−1.400*	(0.167)	−1.099*	(0.184)
[5,6)	−1.276*	(0.177)	−0.944*	(0.195)
[6,7)	−1.346*	(0.197)	−1.011*	(0.215)
[7,8)	−1.573*	(0.233)	−1.238*	(0.249)
[8,9)	−1.686*	(0.258)	−1.339*	(0.274)
≥ 9 years	−2.156*	(0.143)	−1.785*	(0.175)

* $p < 0.5$

Transitions to Employment: Selected Covariate Effects

Variable	Single-level Est.	(se)	Random effects Est.	(se)
Imminent birth (within 1 year)	−0.798*	(0.116)	−0.842*	(0.125)
No. children ≤ 5 yrs (ref=0)				
1 child	−0.151	(0.088)	−0.212*	(0.097)
≥ 2	−0.240	(0.128)	−0.346*	(0.143)
No. children > 5 yrs (ref=0)				
1 child	0.234*	(0.107)	0.251*	(0.118)
≥ 2	0.413*	(0.101)	0.446*	(0.117)
Ever worked	2.675*	(0.121)	2.936*	(0.151)
Previous job part-time	−0.339*	(0.085)	−0.441*	(0.100)

* $p < 0.5$

Transitions to Employment: Comparison of SL and RE Coefficients

- Negative duration dependence is overstated in SL model
- All covariate effects are larger in magnitude in RE model
- Coefficients from the two models have a different interpretation, e.g. for imminent birth
  - Population-averaged estimate of −0.798 compares two randomly selected women, one who is due to give birth and one who is not. Alternatively, it is the average birth effect.
  - Cluster-specific estimate of −0.842 is the effect of changing 'birth status' for a given woman.

Predicted Probabilities of Entering Employment by Duration Non-Employed and Imminent Birth

Duration (yrs)	Median, $u = 0$ No birth	Birth	Mean, $u \sim N(0, \sigma_u^2)$ No birth	Birth
< 1	0.165	0.092	0.177	0.099
[1,2)	0.106	0.055	0.114	0.061
[2,3)	0.086	0.043	0.095	0.051
[3,4)	0.068	0.033	0.077	0.039
[4,5)	0.076	0.037	0.083	0.043
[5,6)	0.085	0.043	0.095	0.050
[6,7)	0.081	0.040	0.089	0.046
[7,8)	0.068	0.033	0.073	0.038
[8,9)	0.062	0.030	0.070	0.035
≥ 9	0.043	0.020	0.049	0.025

Predicted Survival Probabilities

We can calculate survival probabilities S_t from p_t.

$$\begin{aligned} S_t &= \text{probability event occurs during or after interval } t \\ &= \text{probability no event before start of } t \\ &= (1-p_1)(1-p_2)\dots(1-p_{t-1}) \\ &= S_{t-1}(1-p_{t-1}) \end{aligned}$$

Mean p_t and S_t for No Birth

S_t = probability of entering employment during or after interval t

Duration (yrs)	p_t	S_t
< 1	0.177	1
[1,2)	0.114	0.823
[2,3)	0.095	0.730
[3,4)	0.077	0.660
[4,5)	0.083	0.609
[5,6)	0.095	0.559
[6,7)	0.089	0.506
[7,8)	0.073	0.461
[8,9)	0.070	0.427
≥ 9	0.049	0.397

The Proportional Odds Assumption

The discrete-time logit models considered so far assume that the effects of covariates $\mathbf{x}$ are constant over time. This is known as the proportional odds assumption. (Analogous to the proportional hazards assumption in models where the log-hazard is modeled, and $\exp(\beta)$ is a ratio of hazards rather than odds.)

We can relax this assumption by introducing interactions between the duration variables $\mathbf{D}$ and $\mathbf{x}$.

Testing for an Interaction between Duration Non-Employed and Ever Worked

In our transitions to employment example, we will test whether the effect of duration non-employed on the probability of entering employment depends on whether a woman has worked before. (Equivalently, we are testing whether the effect of having worked before depends on the duration non-employed.)

We need to include the products of the 'ever worked' dummy with each of the 9 duration dummies.

Using a likelihood ratio test, the change in deviance is 68.1 which is compared to a χ^2_9 distribution. The p-value is 4×10^{-11}, so strong evidence of non-proportional effects of 'ever worked'.

Log-odds of Employment by Duration and Ever Worked

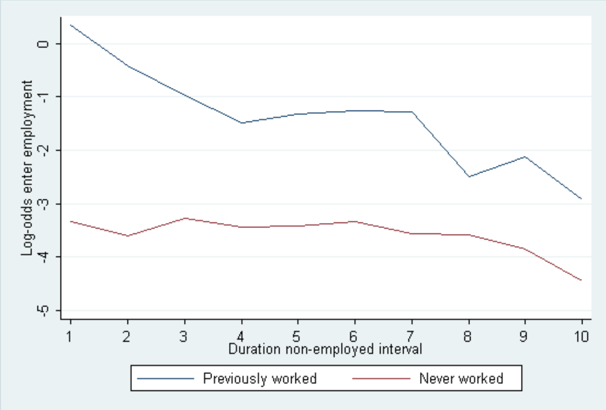

Stronger negative effects of duration non-employed if previously worked. Gap between previously and never worked decreases with duration.

Note: Log-odds calculated at observed values of other x and u=0.

Grouping Time Intervals

When we move to more complex models, a potential problem with the discrete-time approach is that the person-episode-period file can be very large. The size of the file will depend on sample size and the length of the observation period relative to the width of discrete-time intervals.

It may be possible to group time intervals, e.g. using 6-month rather than monthly intervals. In doing so, we must assume the hazard and values of covariates are constant within grouped intervals.

Analysing Grouped Intervals

If we have grouped time intervals, we need to allow for different lengths of exposure time within these intervals. For example, for any 6-month interval, some individuals will have the event or be censored after the 1st month while others will be exposed for the full 6 months.

Denote by n_{tij} the exposure time for individual j in grouped interval t of episode i. (**Note:** Intervals do not need to be the same width.)

Fit binomial logit model for grouped binary data, with response y_{tij} and denominator n_{tij}. Currently random effects binomial logit models can be fitted in Stata (`xtmelogit`), MLwiN and SAS (`PROC NLMIXED`).

Example of Grouped Time Intervals

Suppose an individual is observed to have an event during the 17th month of exposure, and we group durations into six-month intervals (t). Instead of 17 monthly records we would have three six-monthly records:

j	i	t	n_{tij}	y_{tij}
1	1	1	6	0
1	1	2	6	0
1	1	3	5	1

Modelling Transitions Between States

States in Event Histories

In the models considered so far, there is a single event (or transition) of interest. We model the duration to this event from the point at which an individual becomes "at risk". We can think of this as the duration spent in the same state.

E.g.

- In the analysis of transitions into employment we model the duration in the non-employment state
- In a study of marital dissolution we model the duration in the marriage state

More generally, we may wish to model transitions in the other direction (e.g. into non-employment or marriage formation) and possibly other transitions.

Examples of Multiple States

Usually individuals will move in and out of different states over time, and we wish to model these transitions.

Examples:

- Employment states: employed full-time, employed part-time, unemployed, out of the labour market
- Partnership states: marriage, cohabitation, single (not in co-residential union)

We will begin with models for transitions between two states, e.g. non-employment (NE) $\leftrightarrow$ employment (E)

Transition Probabilities for Two States

Suppose there are two states indexed by s $(s = 1, 2)$, and S_{tij} indicates the state occupied by individual j during interval t of episode i.

Denote by y_{tij} a binary variable indicating whether any transition has occurred during interval t, i.e. from state 1 to 2 or from state 2 to 1.

The probability of a transition from state s during interval t, given that no transition has occurred before the start of t is:

$$p_{stij} = Pr(y_{tij} = 1 | y_{t-1,ij} = 0, S_{tij} = s), \qquad s = 1, 2$$

Call p_{stij} a transition probability or discrete-time hazard for state s.

Event History Model for Transitions between 2 States

Multilevel two-state logit model:

$$\log\left(\frac{p_{stij}}{1-p_{stij}}\right) = \mathbf{D}_{stij}\alpha_s + \mathbf{x}_{stij}\beta_s + u_{sj},$$

p_{stij} is the probability of a transition from state s during interval t

$\mathbf{D}_{stij}$ is a vector of functions of cumulative duration in state s by interval t with coefficients α_s

$\mathbf{x}_{stij}$ a vector of covariates affecting the transition from state s with coefficients β_s

u_{sj} allows for unobserved heterogeneity between individuals in their probability of moving from state s. Assume $\mathbf{u}_j = (u_{1j}, u_{2j}) \sim$ bivariate normal.

Random Effect Covariance in a Two-State Model

We assume the state-specific random effects u_{sj} follow a bivariate normal distribution to allow for correlation between the unmeasured time-invariant influences on each transition.

For example, a highly employable person may have a low chance of leaving employment and a high chance of entering employment, leading to $\text{cov}(u_{1j}, u_{2j}) < 0$.

Allowing for $\text{cov}(u_{1j}, u_{2j}) \neq 0$ means that the equations for states $s = 1, 2$ must be estimated jointly. Estimating equations separately assumes that $\text{cov}(u_{1j}, u_{2j}) = 0$.

Data Structure for Two-State Model (1)

Start with an episode-based file.

E.g. employment (E) $\leftrightarrow$ non-employment (NE) transitions

j	i	State$_{ij}$	t_{ij}	y_{ij}	Age$_{ij}$
1	1	E	3	1	16
1	2	NE	2	0	19

Note: (i) t in years; (ii) $y_{ij} = 1$ if a transition (event) occurs, 0 if censored; (iii) Age in years at start of episode

Data Structure for Two-State Model (2)

Convert episode-based file to discrete-time format with one record per interval t:

t	y_{tij}	E$_{ij}$	NE$_{ij}$	E$_{ij}$Age$_{ij}$	NE$_{ij}$Age$_{ij}$
1	0	1	0	16	0
2	0	1	0	16	0
3	1	1	0	16	0
1	0	0	1	0	19
2	0	0	1	0	19

Note: E$_{ij}$ a dummy for employment, NE$_{ij}$ a dummy for non-employment.

Example: Non-Employment $\leftrightarrow$ Employment

- corr(u_{1j}, u_{2j})=0.59, se=0.13, so large positive residual correlation between E $\longrightarrow$ NE and NE $\longrightarrow$ E
- Women with high (low) chance of *entering* E tend to have a high (low) chance of *leaving* E
- Positive correlation arises from two sub-groups: short spells of E and NE, and longer spells of both types

Comparison of Selected Coefficients for NE $\longrightarrow$ E

Only coefficients of covariates relating to employment history change:

	Single-state	Multistate
Ever worked	2.936	2.677
Previous job part-time	−0.441	−0.460

So positive effect of 'ever worked' has weakened, and negative effect of 'part-time' has strengthened.

Why Decrease in Effect of 'Ever Worked' on NE $\longrightarrow$ E?

Direction of change from single-state to multistate (2.936 to 2.677) is in line with positive corr(u_{1j}, u_{2j}) in multistate model.

- Women in 'ever worked' must have made E $\longrightarrow$ NE transition
- Positive correlation between E $\longrightarrow$ NE and NE $\longrightarrow$ E leads to disproportionate presence of women with high NE $\longrightarrow$ E rate among 'ever worked'
- These women push up odds of NE $\longrightarrow$ E among 'ever worked' (inflating estimate) if residual correlation uncontrolled

Why Increase in Effect of Previous Part-Time Job on NE $\longrightarrow$ E?

Strengthening of negative effect when moving to the multistate model (−0.441 to −0.460) is also in line with positive corr(u_{1j}, u_{2j}).

- Women with tendency towards less stable employment (with high rate of E $\longrightarrow$ NE) selected into part-time work
- Positive correlation between E $\longrightarrow$ NE and NE $\longrightarrow$ E leads to disproportionate presence of women with high NE $\longrightarrow$ E rate in 'previous PT' category
- These women push up odds of NE $\longrightarrow$ E in 'previous PT' (reducing 'true' negative effect of PT) if residual correlation uncontrolled

Autoregressive Models for Two States

An alternative way of modelling transitions between states is to include the lagged response as a predictor rather than the duration in the current state.

The response y_{tij} now indicates the state occupied at the start of interval t rather than whether a transition has occurred, i.e.

$$y_{tij} = \begin{cases} 1 & \text{if in state 1} \\ 0 & \text{if in state 2} \end{cases}$$

1st Order Autoregressive Model

An AR(1) model for the probability that individual j is in state 1 at t, p_{tj} is:

$$\log\left(\frac{p_{tj}}{1-p_{tj}}\right) = \alpha + \mathbf{x}_{tj}\beta + \gamma y_{t-1,j} + u_j$$

α is an intercept term

γ is the effect of the state occupied at $t-1$ on the log-odds of being in state 1 at t

$u_j \sim N(0, \sigma_u^2)$ is an individual-specific random effect

Interpretation of AR(1) Model

Suppose states are employment and unemployment. Common to find those who have been unemployed in the past are more likely to be unemployed in the future. Three potential explanations:

- A causal effect or state dependence (γ)
- Unobserved heterogeneity, i.e. unmeasured individual characteristics affecting unemployment probability at all t (stable traits u_j)
- Non-stationarity, e.g. seasonality (not in current model)

The AR(1) model is commonly referred to as a state dependence model.

Transition Probabilities from the AR(1) Model

We model $p_{tj} = Pr(\text{state 1 at start of interval } t) = Pr(y_{tj} = 1)$

Suppose we fix $\mathbf{x}_{tj} = \mathbf{0}$ and $u_j = 0$.

Probability of moving from state 1 to 2

$$\begin{aligned} Pr(y_{tj} = 0 | y_{t-1,j} = 1) &= 1 - Pr(y_{tj} = 1 | y_{t-1,j} = 1) \\ &= 1 - \frac{\exp(\alpha+\gamma)}{1+\exp(\alpha+\gamma)} \end{aligned}$$

Probability of moving from state 2 to 1

$$Pr(y_{tj} = 1 | y_{t-1,j} = 0) = \frac{\exp(\alpha)}{1+\exp(\alpha)}$$

Initial Conditions (1)

y may not be measured at the start of the process, e.g. we may not have entire employment histories.

Can view as a missing data problem. Suppose we observe y at the start of T intervals:

Observed	$(y_1, \ldots, y_T)$
Actual	$(y_{-k}, \ldots, y_0, y_1, \ldots, y_T)$

where first $k+1$ measures are missing.

We need to specify a model for y_1 (not just condition on y_1).

Initial Conditions (2)

In a random effects framework, we can specify a model for y_{1j} and estimate jointly with the model for $(y_{2j}, \ldots, y_{Tj})$, e.g.

$$\begin{aligned} \text{logit}(p_{1j}) &= \alpha_1 + \mathbf{x}_{t1j}\beta_1 + \lambda u_j \\ \text{logit}(p_{tj}) &= \alpha + \mathbf{x}_{tj}\beta + \gamma y_{t-1,j} + u_j, \qquad t > 1 \end{aligned}$$

Variants on the above are to set $\lambda = 1$ or to include different random effect in equation for $t = 1$, e.g. u_{1j}, and allow for correlation with u_j in equation for $t > 1$.

Initial Conditions (3)

Initial conditions may also be a problem in the duration model.

If we do not observe an individual from the start of the process of interest, the state occupied at $t = 1$ may be informative.

As in the AR(1) model, we can model the initial state jointly with subsequent transitions.

Key Features of the AR(1) Model

- All relevant information about the process up to t is captured by y_{t-1} (1st order Markov assumption). This is why duration effects are not included.
- Because of the 1st order Markov assumption, there is no concept of an 'episode' (which is why we drop the i subscript)
- Effects of x (and time-invariant characteristics u_j) are the same for transitions from state 1 to 2, and from state 2 to 1

Which Model?

Consider AR(1) model when:

- Interested in separation of causal effect of y_{t-1} on y_t from unobserved heterogeneity
- Frequent movement between states (high transition probabilities)
- Duration in state at $t = 1$ is unknown, e.g. in panel data

Consider duration model when:

- Expect duration in state to have an effect on chance of transition
- More stable processes with long periods in the same state (low transition probabilities)

More than Two states

In general there may be multiple states, possibly with different destinations from each state. E.g. consider transitions between marriage (M), cohabitation (C) and single (S).

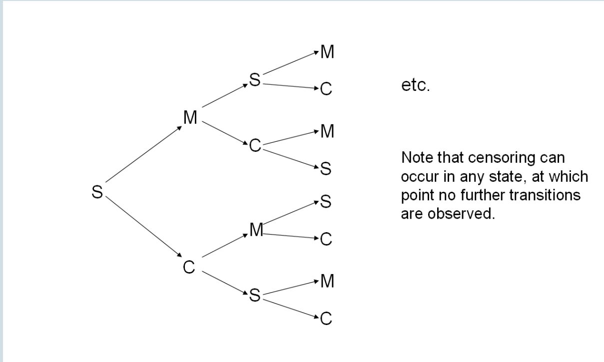

Competing risks

We will begin with the special case where we are interested in transitions from one origin state, but there is more than one destination or type of transition.

Assume these destinations are mutually exclusive. We call these 'competing risks'.

Origin state	Competing risks
Alive	Different causes of death
Employed	Sacked, redundancy, switch job, leave labour market
Single	Marriage, cohabitation

Approaches to Modelling Competing Risks

Suppose there are R types of transition/event. For each interval t (of episode i of individual j) we can define a categorical response y_{tij}:

$$y_{tij} = \begin{cases} 0 & \text{if no event in } t \\ r & \text{if event of type } r \text{ in } t \ (r = 1, \ldots, R) \end{cases}$$

Analysis approaches

1. Multinomial model for y_{tij}
2. Define binary response $y_{tij}^{(r)}$ for event type r, treating all other types of event as censored. Analyse using multivariate response model

Multinomial Logit Model

Define $p_{tij}^{(r)} = \Pr(y_{tij} = r | y_{t-1,ij} = 0)$ for $r = 1, \ldots, R$.

Estimate R equations contrasting event type r with 'no event':

$$\log\left(\frac{p_{tij}^{(r)}}{p_{tij}^{(0)}}\right) = \mathbf{D}_{tij}^{(r)}\alpha^{(r)} + \mathbf{x}_{tij}^{(r)}\beta^{(r)} + u_j^{(r)}, \qquad r = 1, \ldots, R$$

where $(u_j^{(1)}, u_j^{(2)}, \ldots, u_j^{(R)}) \sim$ multivariate normal.

Correlated random effects allows for shared unobserved risk factors.

Multivariate Binary Response Model

In the second approach to modelling competing risks we define, for each interval t, R binary responses coded as:

$$y_{tij}^{(r)} = \begin{cases} 1 & \text{if event of type } r \text{ in } t \\ 0 & \text{if event of any type other than } r \text{ or no event in } t \end{cases}$$

and estimate equations for each event type:

$$\log\left(\frac{p_{tij}^{(r)}}{1 - p_{tij}^{(r)}}\right) = \mathbf{D}_{tij}^{(r)}\alpha^{(r)} + \mathbf{x}_{tij}^{(r)}\beta^{(r)} + u_j^{(r)}, \qquad r = 1, \ldots, R$$

where $(u_j^{(1)}, u_j^{(2)}, \ldots, u_j^{(R)}) \sim$ multivariate normal.

Logic Behind Treating Other Events as Censored

Suppose we are interested in modelling partnership formation, where an episode in the 'single' state can end in marriage or cohabitation.

For each single episode we can think of durations to marriage and cohabitation, $t^{(M)}$ and $t^{(C)}$.

We cannot observe both of these. If a single episode ends in marriage, we observe only $t^{(M)}$ and the duration to cohabitation is censored at $t^{(M)}$. A person who marries is removed from the risk of cohabiting (until they become single again).

For uncensored episodes we observe $\min(t^{(M)}, t^{(C)})$.

Comparing Methods

Coefficients and random effect variances and covariances will be different for the two models because the reference category is different:

- 'No event' in the multinomial model
  - Coefficients are effects on the log-odds of an event of type r relative to 'no event'
- 'No event + any event other than r' in the multivariate binary model
  - Coefficients are effects on the log-odds of an event of type r relative to 'no event of type r'

However, predicted transition probabilities will in general be similar for the two models.

Transition Probabilities

Multinomial model

$$p_{tij}^{(r)} = \frac{\exp(\mathbf{D}_{tij}^{(r)}\alpha^{(r)} + \mathbf{x}_{tij}^{(r)}\beta^{(r)} + u_j^{(r)})}{1 + \sum_{k=1}^{R}\exp(\mathbf{D}_{tij}^{(k)}\alpha^{(k)} + \mathbf{x}_{tij}^{(k)}\beta^{(k)} + u_j^{(k)})}$$

Multivariate binary model

$$p_{tij}^{(r)} = \frac{\exp(\mathbf{D}_{tij}^{(r)}\alpha^{(r)} + \mathbf{x}_{tij}^{(r)}\beta^{(r)} + u_j^{(r)})}{1 + exp(\mathbf{D}_{tij}^{(r)}\alpha^{(r)} + \mathbf{x}_{tij}^{(r)}\beta^{(r)} + u_j^{(r)})}$$

In each case, the 'no event' probability is $p^{(0)} = 1 - \sum_{k=1}^{R} p^{(k)}$.

To calculate probabilities for specific values of $\mathbf{x}$, substitute $u^{(r)} = 0$ or generate $u^{(r)}$ from multivariate normal distribution.

Example: Transitions to Full-time and Part-time Work

Selected results from bivariate model for binary responses, $y^{(FT)}$ and $y^{(PT)}$

	NE $\rightarrow$ FT		NE $\rightarrow$ PT	
Variable	Est	(se)	Est	(se)
Imminent birth	-1.19*	(0.18)	-0.26	(0.15)
1 kid ≤ 5	-1.27*	(0.15)	0.60*	(0.12)
2+ kids ≤ 5	-1.94*	(0.27)	0.81*	(0.17)
1 kid > 5	-0.42*	(0.19)	0.80*	(0.14)
2+ kids > 5	-0.26	(0.18)	1.24*	(0.15)

So having kids (especially young ones) reduces chance of returning to FT work, but increases chance of returning to PT work.

Example: Random Effect Covariance Matrix

	NE $\rightarrow$ FT	NE $\rightarrow$ PT
NE $\rightarrow$ FT	1.49 (0.13)	
NE $\rightarrow$ PT	-0.05 (0.11)	0.98 (0.11)

Note: Parameters on diagonal are standard deviations, and off-diagonal parameter is the correlation. Standard errors in parentheses.

Correlation is not significant (deviance test statistic is < 1 on 1 d.f.).

Dependency between Competing Risks

- A well-known problem with the multinomial logit model is the 'independence of irrelevant alternatives' (IIA) assumption
- In the context of competing risks, IIA implies that the probability of one event relative to 'no event' is independent of the probabilities of each of the other events relative to 'no event'
- This may be unreasonable if some types of event can be regarded as similar
- Note that the multivariate binary model makes the same assumption

Dependency between Competing Risks: Example

Suppose we wish to study partnership formation: transitions from single (S) to marriage (M) or to cohabitation (C).

- Under IIA, assume probability of C vs. S is uncorrelated with probability of M vs. S
- E.g. if there is something unobserved (not in $\mathbf{x}$) that made M infeasible, we assume those who would have married distribute themselves between C and S in the same proportions as those who originally chose not to marry
- But as M and C are similar, we might expect those who are precluded from marriage to be more likely to cohabit rather than remain single (Hill, Axinn and Thornton, 1993, *Sociological Methodology*

Relaxing the Independence Assumption

- Including individual-specific random effects allows for dependence due to time-invariant individual characteristics (e.g. attitudes towards marriage/cohabitation)
- But it does not allow for unmeasured factors that vary across episodes (e.g. marriage is not an option if respondent or their partner is already married)

Modelling Transitions between More than 2 States

So far we have considered (i) transitions between two states, and (ii) transitions from a single state with multiple destinations.

We can bring these together in a general model, allowing for different destinations from each state.

Example: partnership transitions

- Formation: S $\longrightarrow$ M, S $\longrightarrow$ C
- Conversion of C to M (same partner)
- Dissolution: M $\longrightarrow$ S, C $\longrightarrow$ S (or straight to new partnership)

Estimate 5 equations simultaneously (with correlated random effects)

Example of Multiple States with Competing Risks

- Contraceptive use dynamics in Indonesia. Define episode of use as continuous period of using *same* method of contraception
  - 2 states: use and nonuse
  - Episode of use can end in 2 ways: discontinuation (transition to nonuse), or method switch (transition within 'use' state)
- Estimate 3 equations jointly: binary logit for nonuse $\longrightarrow$ use, and multinomial logit for transitions from use
- Details in Steele et al. (2004) *Statistical Modelling*

Selected Results: Coefficients and SEs

	Use → nonuse (Discontinuation)		Use → new method (Method switch)		Nonuse → use	
Urban	0.13	(0.14)	0.06	(0.05)	0.26	(0.04)
SES						
Medium	−0.12	(0.05)	0.35	(0.07)	0.24	(0.05)
High	−0.20	(0.05)	0.29	(0.08)	0.45	(0.05)

Random Effect Correlations from Alternative Models

	Discontinuation	Method switch	Nonuse → use
Discontinuation	1		
Method switch	0.020 0.011	1	
Nonuse → use	−0.783* −0.052	0.165* 0.095	1

Model 1: Duration effects only
Model 2: Duration + covariate effects

*Correlation significantly different from zero at 5% level

Random Effect Correlations: Interpretation

- In 'duration effects only' model, there is a large negative correlation between random effects for nonuse → use and use → nonuse
  - Long durations of use associated with short durations of nonuse
- This is due to short episodes of postnatal nonuse followed by long episodes of use (to space or limit future births)
  - Correlation is effectively zero when we control for whether episode of nonuse follows a live birth (one of the covariates)

Software for Recurrent Events and Multiple States

- Recurrent events
  - Essentially multilevel models for binary responses
  - Stata (xtlogit), SAS (proc nlmixed) and all specialist multilevel modelling software (e.g. MLwiN, SABRE, aML)
- Multiple states
  - Two-state and bivariate competing risks fitted as random coefficient models; options as above but xtmelogit in Stata
  - Multinomial competing risks in SAS, MLwiN and aML

Modelling Interdependent Events

Endogeneity in a 2-Level Continuous Response Model

Consider a 2-level random effects model for a continuous response:

$$y_{ij} = \mathbf{x}_{ij}\beta + u_j + e_{ij}$$

where $\mathbf{x}_{ij}$ is a set of covariates with coefficients β, u_j is the level 2 random effect (residual) $\sim N(0, \sigma_u^2)$ and e_{ij} is the level 1 residual $\sim N(0, \sigma_e^2)$.

One assumption of the model is that $\mathbf{x}_{ij}$ is uncorrelated with both u_j and e_{ij}, i.e. we assume that $\mathbf{x}_{ij}$ is exogenous.

This may be too strong an assumption. If unmeasured variables affecting y_{ij} also affect one or more covariates, then those covariates will be endogenous.

2-Level Endogeneity: Example

Suppose y_{ij} is birth weight of child i of woman j, and z_{ij} is the number of antenatal visits during pregnancy (an element of $\mathbf{x}_{ij}$).

Some of the factors that influence birth weight may also influence uptake of antenatal care; these may be characteristics of the particular pregnancy (e.g. woman's health during pregnancy) or of the woman (health-related attitudes/behaviour). Some of these may be unobserved.

i.e. y and z are to some extent jointly determined, and z is endogenous.

This will lead to correlation between z and u and/or e and, if ignored, a biased estimate of the coefficient of z and possibly covariates correlated with z.

Illustration of Impact of Endogeneity at Level 1

Suppose the 'true' effect of z_{ij} on y_{ij} is positive, i.e. more antenatal visits is associated with a higher birth weight.

Suppose that w_{ij} is 'difficulty of pregnancy'. We would expect $\text{corr}(w, y) < 0$, and $\text{corr}(w, z) > 0$.

If w is unmeasured it is absorbed into e, leading to $\text{corr}(z, e) < 0$.

If we ignore $\text{corr}(z, e) < 0$, the estimated effect of z on y will be biased downwards.

The disproportionate presence of high w women among those getting more antenatal care (high z) suppresses the positive effect of z on y.

Illustration of Impact of Endogeneity at Level 2

As before, suppose the 'true' effect of z on y is positive, i.e. more antenatal visits is associated with a higher birth weight.

Suppose that w_j is 'healthcare knowledge' which is constant across the observation period. We would expect $\text{corr}(w, y) > 0$, and $\text{corr}(w, z) > 0$.

If w is unmeasured it is absorbed into u, leading to $\text{corr}(z, u) > 0$.

Question: What effect would ignoring $\text{corr}(z, u) > 0$ have on the estimated effect of z on y?

Handling Endogeneity in a Single-Level Model

To fix ideas, we will start with the simplest case: outcome y and endogenous predictor z both continuous.

E.g. y_i birth weight of last born child of woman i, z_i number of antenatal visits.

We specify a simultaneous equations model for z and y:

$$\begin{aligned} z_i &= \mathbf{x}_i^z \beta^z + e_i^z \\ y_i &= \mathbf{x}_i^y \beta^y + z_i \gamma + e_i^y \end{aligned}$$

where $\mathbf{x}_i^z$ and $\mathbf{x}_i^y$ are exogenous covariates (assumed to be uncorrelated with e_i^z and e_i^y).

Estimation

If $\text{corr}(e_i^z, e_i^y) = 0$, OLS of the equation for y_i is optimal.

Endogeneity of z_i will lead to $\text{corr}(e_i^z, e_i^y) \neq 0$ and an alternative estimation procedure is required. The most widely used approaches are:

- 2-stage least squares (2SLS)
- Joint estimation of equations for z and y (Full Information Maximum Likelihood, FIML)

Estimation: 2-Stage Least Squares

1. OLS estimation of equation for z_i and compute $\hat{z}_i = \mathbf{x}_i^z \hat{\beta}^z$
2. OLS estimation of equation for y_i replacing z_i by prediction $\hat{z}_i$
3. Adjust standard errors in (2) to allow for uncertainty in estimation of $\hat{z}_i$

Idea: $\hat{z}_i$ is 'purged' of the correlated unobservables e_i^z, so $\hat{z}_i$ uncorrelated with e_i^y.

Estimation: FIML

Treat z_i and y_i as a bivariate response and estimate equations jointly.

Usually assume e_i^z and e_i^y follow a bivariate normal distribution with correlation ρ_e^{zy}.

- Can be estimated in a number of software packages (e.g. `mvreg` in Stata or Sabre)
- Sign of $\hat{\rho}_e^{zy}$ signals direction of bias
- Generalises to mixed response types (e.g. binary z and duration y)
- Generalises to clustered data (multilevel multivariate model)

Testing for Exogeneity of z

To test the null hypothesis that z is exogenous:

2SLS

Estimate $y_i = \mathbf{x}_i^y \beta^y + z_i\gamma + \hat{e}_i^z\delta + e_i^y$ via OLS

where $\hat{e}_i^z$ is the estimated residual from fitting the 1st stage equation for z_i

Test $\mathrm{H}_0 : \delta = 0$ using t (or Z) test.

FIML

Jointly estimate equations for z_i and y_i to get estimate of residual correlation ρ_e^{zy}.

Test $\mathrm{H}_0 : \rho_e^{zy} = 0$ using likelihood ratio test.

Identification

Whatever estimation approach is used, identification of the simultaneous equations model for z and y requires covariate exclusion restrictions.

$\mathbf{x}_i^z$ should contain at least one variable that is not in $\mathbf{x}_i^y$.

In our birth weight example, need to find variable(s) that predict antenatal visits (z) but not birth weight (y).

Call such variables instruments.

Note: The term ‘IV estimation’ is commonly used interchangeably with 2SLS, but both methods require instruments.

Requirements of an Instrument (1)

Need to be able to justify, on theoretical grounds, that the instrument affects z but not y (after controlling for z and other covariates).

E.g. indicator of access to antenatal care may be suitable instrument for no. visits, but only if services are allocated randomly (rare). Instruments can be very difficult to find.

If there is > 1 instrument, the model is said to be over-identified.

Requirements of an Instrument (1)

Testing over-identifying restrictions

Instruments should not affect y after controlling for z.

Fit the SEM with all but one instrument in the equation for y and carry out a joint significance test of the included instruments. If the restrictions are valid, they should not have significant effects on y.

Instruments should be correlated with z

Carry out joint significance test of effects of instruments on z.

Also check how well instruments (together with other covariates) predict z. Bollen et al. (1995) suggest a simple probit for y is preferred if $R^2 < 0.1$.

Effect of Fertility Desires on Contraceptive Use (1)

Reference: Bollen, Guilkey and Mroz (1995), *Demography*.

Interested in the impact of number of additional children desired (z, continuous) on use of contraception (y, binary).

Unmeasured variable affecting both z and y could be 'perceived fecundity'.

Women who believe they have low chance of having a(nother) child may lower fertility desires and not use contraception $\rightarrow$ $\text{corr}(e_i^z, e_i^y) > 0$.

Effect of Fertility Desires on Contraceptive Use (2)

Expect 'true' effect of fertility desires (z) on contraceptive use (y) to be negative.

- If residual correlation ignored, negative effect of z on y will be understated (may even estimate a positive effect)
- Estimated effects of covariates correlated with z also biased (e.g. whether heard family planning message)

Effect of Fertility Desires on Contraceptive Use (3)

Cross-sectional data: z and y refer to time of survey.

Use 2SLS: OLS for z equation, probit for y.

Instruments: Indicators of health care facilities in community when woman was age 20 (supplementary data).

Results:

- Residual correlation estimated as 0.07
- Stronger negative effect of z after allowing for endogeneity (changes from -0.17 to -0.28), but large increase in SE
- But fail to reject null that z is exogenous, so simple probit for contraceptive use is preferred

Handling Endogeneity in a Multilevel Model

Let's return to the multilevel case with y_{ij} the birth weight of child i of woman j, z_{ij} number of antenatal visits.

We specify a multilevel simultaneous equations (multiprocess) model for z and y:

$$\begin{aligned} z_{ij} &= \mathbf{x}_{ij}^{z}\beta^{z} + u_j^z + e_{ij}^z \\ y_{ij} &= \mathbf{x}_{ij}^{y}\beta^{y} + z_{ij}\gamma + u_j^y + e_{ij}^y \end{aligned}$$

where u_j^z and u_j^y are normally distributed woman-level random effects, and $\mathbf{x}_{ij}^z$ and $\mathbf{x}_{ij}^y$ are exogenous covariates (assumed to be uncorrelated with u_j^z, u_j^y, e_{ij}^z and e_{ij}^y).

Estimation

If $\text{corr}(u_j^z, u_j^y) = 0$ and $\text{corr}(e_{ij}^z, e_{ij}^y) = 0$, the equation for y_{ij} can be estimated as a standard multilevel model.

However, endogeneity of z_{ij} will lead to $\text{corr}(u_j^z, u_j^y) \neq 0$ or $\text{corr}(e_{ij}^z, e_{ij}^y) \neq 0$ (or both).

If z_{ij} is endogenous we need to estimate equations for z and y jointly.

In the most general model, we assume $(u_j^z, u_j^y) \sim$ bivariate normal and $(e_{ij}^z, e_{ij}^y) \sim$ bivariate normal. The SEM is a multilevel bivariate response model.

Identification (1)

Identification of the full multilevel SEM for z and y, with $\text{corr}(u_j^z, u_j^y) \neq 0$ and $\text{corr}(e_{ij}^z, e_{ij}^y) \neq 0$, requires covariate exclusion restrictions:

$\mathbf{x}_{ij}^z$ should contain at least one variable (an instrument) that is not in $\mathbf{x}_{ij}^y$.

e.g. need to find variable(s) that predict antenatal visits (z) but not birth weight (y).

Call such variables instruments.

BUT if one of the residual covariances is assumed equal to zero, covariate exclusions are not strictly necessary for identification.

Identification (2)

Suppose we are prepared to assume that endogeneity of z is due to a residual correlation at the woman level but not at the pregnancy level, i.e.

$\text{corr}(u_j^z, u_j^y) \neq 0$ but $\text{corr}(e_{ij}^z, e_{ij}^y) = 0$.

We are then assuming that bias in the estimated effect of number of antenatal visits on birth weight is due to selection on unmeasured maternal characteristics that are fixed across pregnancies.

Identification (3)

Given the difficulty in finding instruments, allowing only for selection on time-invariant unobservables (in a longitudinal design) is a common identification strategy BUT:

- It does not allow for selection on time-varying unobservables so some bias may be remain
- Some within-individual variation in z and y is required because we are estimating the effect of a change in z on y for a given woman (i.e. conditioning on woman-specific unobservables)

Allowing for Endogeneity in an Event History Model

Suppose that y_{ij} is the duration of episode i of individual j and z_{ij} is an endogenous variable. We first consider case where z is continuous and measured at the episode level.

We can extend our earlier recurrent events model to a SEM:

$$\begin{aligned} z_{ij} &= \mathbf{x}^z_{ij}\beta^z + u^z_j + e^z_{ij} \\ \log\left(\frac{p_{tij}}{1-p_{tij}}\right) &= \mathbf{D}^y_{tij}\alpha^y + \mathbf{x}^y_{tij}\beta^y + z_{ij}\gamma + u^y_j \end{aligned}$$

where p_{tij} is the probability of an event during interval t, $\mathbf{D}^y_{tij}\alpha^y$ is the baseline hazard, and $\mathbf{x}^y_{tij}$ a vector of exogenous covariates.

We assume $(u^z_j, u^y_j) \sim$ bivariate normal, i.e. we allow for selection on time-invariant individual characteristics.

Examples of Multilevel SEM for Event History Data

More generally z can be categorical and can be defined at any level (e.g. time-varying or a time-invariant individual characteristic).

We will consider two published examples before returning to our analysis of women's employment transitions:

- The effect of premarital cohabitation (z) on subsequent marital dissolution (y)
  - Lillard, Brien & Waite (1995), *Demography*
- The effect of access to family planning (z) on fertility (z)
  - Angeles, Guilkey & Mroz (1998), *Journal of the American Statistical Association*

Example 1: Premarital Cohabitation and Divorce

Couples who live together before marriage appear to have an increased risk of divorce.

Is this a 'causal' effect of premarital cohabitation or due to self-selection of more divorce-prone individuals into premarital cohabitation?

The analysis uses longitudinal data so observe women in multiple marriages (episodes). For each marriage define 2 equations:

- A probit model for premarital cohabitation (z)
- A (continuous-time) event history model for marital dissolution (y)

Each equation has a woman-specific random effect, u^z_j and u^y_j, which are allowed to be correlated

Premarital Cohabitation and Divorce: Identification

Lillard et al. argue that exclusion restrictions are unnecessary because of 'within-person replication'.

Nevertheless they include some variables in the cohabitation equation that are not in the dissolution equation:

- Education level of woman's parents
- Rental prices and median home value in state
- Sex ratio (indicator of 'marriageable men/women')

They examine the robustness of their conclusions to omitting these variables from the model.

Premarital Cohabitation and Divorce: Results (1)

Correlation between woman-specific random effects for cohabitation and dissolution estimated as 0.36.

Test statistic from a likelihood ratio test of the null hypothesis that $\text{corr}(u_j^z, u_j^y) = 0$ is 4.6 on 1 d.f. which is significant at the 5% level.

"There are unobserved differences across individuals which make those who are most likely to cohabit before any marriage also most likely to end any marriage they enter."

Premarital Cohabitation and Divorce: Results (2)

What is the impact of ignoring this residual correlation, and assuming premarital cohabitation is exogenous?

Estimated effect of cohabitation on log-hazard of dissolution

0.37 and strongly significant if $\text{corr}(u_j^z, u_j^y) = 0$ assumed

-0.01 and non-significant if $\text{corr}(u_j^z, u_j^y)$ allowed to be non-zero

Conclude that, after allowing for selection, there is no association between premarital cohabitation and marital dissolution.

Example 2: Access to Family Planning and Fertility

Does availability of family planning (FP) services lead to a reduction in fertility in Tanzania?

Problem: FP clinics are unlikely to be placed at random. They are likely to be targeted towards areas of greatest need, the type of area with high fertility.

Question: If true impact of access to FP is to increase birth spacing, how will ignoring targeted placement affect estimates of the impact?

Data and Measures

Birth histories collected retrospectively in 1992. Women nested within communities, so have a 3-level structure: births (level 1), women (level 2), communities (level 3).

Constructed woman-year file for period 1970-1991 with y_{tij}=1 if woman i in community j gave birth in year t. (Could have extra subscript for birth interval as we model duration since last birth.)

Community survey on services conducted in 1994. Construct indicators of distance to hospital, health centre etc. in year t, z_{tj}. Time-varying indicators derived from information on timing of facility placement.

Multiprocess Model for Programme Placement and Fertility

The model consists of 4 equations:

- Discrete-time event history model for probability of a birth in year t with woman and community random effects
- Logit models for placement of 3 types of FP facility in community j in year t with community random effects

Allow for correlation between community random effects for fertility and FP clinic placement.

Rather than assume normality, the random effects distribution is approximated by a step function using a 'discrete factor' method.

Programme Placement and Fertility: Identification

The following time-varying variables are included in the FP placement equations but not the fertility equation:

- National government expenditure on health
- Regional government expenditure on health
- District population as fraction of national population

These are based on time series data at the national, regional and community levels.

Programme Placement and Fertility: Findings

- From simple analysis (ignoring endogeneity of programme placement) find hospitals have more impact on reducing fertility than health centres
- But this analysis overstates impact of hospitals and understates effects of health centres
- Controlling for endogenous programme placement reveals that health centres have more impact than hospitals
- After controlling for endogenity, impact of FP facilities was 45% larger than in simpler analysis

Modelling Correlated Event Processes

Now suppose that z_{tij} is a time-varying endogenous predictor.

z_{tij} is often the outcome of a related event process.

Example: Marital dissolution and fertility

y_{ij} is duration of marriage i of woman j

z_{tij} is number of children from marriage i of woman j at time t, the outcome of a birth history

See Lillard (1993) and Lillard & Waite (1993).

Multiprocess Model SEM for 2 Interdependent Events

Simultaneous discrete-time event history equations:

$$\begin{aligned} \text{logit}(p^z_{tij}) &= \mathbf{D}^z_{tij}\alpha^z + \mathbf{x}^z_{tij}\beta^z + u^z_j \\ \text{logit}(p^y_{tij}) &= \mathbf{D}^y_{tij}\alpha^y + \mathbf{x}^y_{tij}\beta^y + z_{ij}\gamma + u^y_j \end{aligned}$$

We assume $(u^z_j, u^y_j) \sim$ bivariate normal, i.e. we allow for selection on time-invariant individual characteristics.

The model can be extended to include outcomes of the y process in the model for z.

Example: Marital Dissolution and Fertility

Lillard's model has 2 (continuous-time) event history equations for:

- hazard of conception (leading to a live birth) at time t of marriage i of woman j
- hazard that marriage i of woman j ends at time t

Consider dummies for z_{tij}, the number of children from marriage i, in dissolution equation.

Marital Dissolution and Fertility: Results (1)

Lillard (1993) finds that the residual correlation between hazard of dissolution and hazard of a conception is estimated as -0.75 (se=0.20).

$\Rightarrow$ women with a below-average risk of dissolution ($u^y_j < 0$) tend to have an above-average chance of a marital conception ($u^z_j > 0$).

$\Rightarrow$ selection of women with a low dissolution risk into having children.

Question: If the 'true' effect of having children is to reduce the risk of dissolution, what impact would this type of selection have on estimates of this effect?

Marital Dissolution and Fertility: Results (2)

Estimated effects (se) of number of children from current marriage on log-hazard of dissolution before and after accounting for residual correlation:

# children	Before		After	
0 (ref)	0		0	
1	−0.56	(0.10)	−0.33	(0.11)
2+	−0.01	(0.05)	0.27	(0.07)

Selection of low dissolution risk women in categories 1 and 2+

Other Examples of Correlated Event Histories

- Employment transitions and fertility (next example)
- Partnership formation and employment
- Residential mobility and fertility
- Residential mobility and employment
- Residential mobility and partnership formation/dissolution

Multiprocess Model for Entry into Employment and Fertility (1)

At the start of the course (and in computer exercises) we fitted multilevel models for the transition from non-employment (NE) to employment (E) among British women.

Among the covariates was a set of time-varying fertility indicators:

- Due to give birth within next year
- Number of children aged ≤ 5 years
- Number of children aged > 5 years

These are outcomes of the fertility process which might be jointly determined with employment transitions.

Multiprocess Model for Entry into Employment and Fertility (2)

Denote by y_{tij}^{NE} and y_{tij}^{B} binary indicators for leaving non-employment and giving birth during year t.

Estimate 2 simultaneous equations (both with woman-specific random effects):

- Discrete-time logit for probability of a birth
- Discrete-time logit for probability of NE $\rightarrow$ E (with fertilty outcomes as predictors)

Note: While we could model births that occur during non-employment, it would be more natural to model the whole birth process (in both NE and E). In the following analysis, we consider *all* births.

Estimation of Multiprocess Model

We can view the discrete-time multiprocess model as a multilevel bivariate response model for the binary responses y_{tij}^{NE} and y_{tij}^{B}.

- Stack the employment and birth responses into a single response column and define an index r which indicates the response type (e.g. $r = 1$ for NE and $r = 2$ for B)
- Define dummies for r which we call r_1 and r_2 say
- Multiply r_1 and r_2 by the covariates to be included in the NE and B equations respectively
- Fit woman-level random effects to r_1 and r_2 and allow to be correlated

Entry into Employment and Fertility: Residual Correlation

Likelihood ratio test statistic for test of null hypothesis that $\text{corr}(u_j^{NE}, u_j^{B}) = 0$ is 8 on 1 d.f.

$\Rightarrow$ reject the null and choose the multiprocess model.

Correlation between woman-level random effects, u_j^{NE} and u_j^{B}, estimated as 0.34 (se=0.11).

The positive correlation implies that women whose unobserved characteristics are associated with a high probability of a birth (e.g. latent preference for childbearing) tend also to enter employment quickly after a spell of non-employment.

Effects of Fertility Outcomes on Entry into Employment

	Single process		Multiprocess	
Variable	Est.	(se)	Est.	(se)
Imminent birth (within 1 year)	-0.84^*	(0.13)	-1.01^*	(0.14)
No. children $\leq$ 5 yrs (ref=0)				
1 child	-0.21^*	(0.10)	-0.35^*	(0.11)
≥ 2	-0.35^*	(0.14)	-0.60^*	(0.17)
No. children $>$ 5 yrs (ref=0)				
1 child	0.25^*	(0.12)	0.18	(0.12)
≥ 2	0.45^*	(0.12)	0.27^*	(0.13)

* $p < 0.5$

Selection of women with high NE $\rightarrow$ E probability into categories 1 and ≥ 2

Multiple States and Correlated Processes

We can extend the multiprocess model to include transitions between multiple states and further correlated processes.

E.g. we could model two-way transitions between NE and E jointly with births, leading to 3 simultaneous equations and 3 correlated random effects.

Stack employment transition and birth responses into a single column with a 3-category response indicator r (e.g. r=1 for employment episodes, r=2 for non-employment episodes, r=3 for birth intervals).

Create dummies for r and interact with covariates as for two-state and multiprocess models.

Employment Transitions and Fertility: Random Effects Correlation Matrix

	NE → E	E → NE	Birth
NE → E	1		
E → NE	0.62 (0.12)	1	
Birth	0.45 (0.11)	0.23 (0.08)	1

Standard errors in brackets

Effects of Fertility Outcomes on Exit from Employment

	Single process		Multiprocess	
Variable	Est.	(se)	Est.	(se)
Imminent birth (within 1 year)	2.31*	(0.14)	2.23*	(0.15)
No. children $\leq$ 5 yrs (ref=0)				
1 child	0.41*	(0.10)	0.31*	(0.11)
≥ 2	0.33	(0.17)	0.15	(0.18)
No. children > 5 yrs (ref=0)				
1 child	−0.35*	(0.12)	−0.34	(0.12)
≥ 2	−0.28*	(0.12)	−0.37*	(0.13)

* $p < 0.5$

Selection of women with high NE → E probability into categories 1 and ≥ 2

Example: Family Disruption and Children's Education

Research questions:

- What is the association between disruption (due to divorce or paternal death) and children's education?
- Are the effects of disruption the same across different educational transitions?
- To what extent can the effect of divorce be explained by selection?
  - There may be unobserved factors affecting both parents' dissolution risk and their children's educational outcomes

Reference: Steele, Sigle-Rushton and Kravdal (2009), *Demography*.

Strategies for Handling Selection

- Exploit longitudinal study designs
  - e.g. measures of child wellbeing before divorce, measures of parental conflict and family environment
- Exploit differences across space
  - compare children living in places with differences in availability of divorce (e.g. US states)
- Compare siblings
  - Siblings share parents (or parent) but may have different exposure to disruption
- Multiprocess (simultaneous equations) models

SEM for Parental Divorce and Children's Education

- Selection equation: event history model for duration of mother's marriage(s)
- Sequential probit model for children's educational transitions (nested within mother)
- Equations linked by allowing correlation between mother-specific random effects (unmeasured maternal characteristics)
- Estimated using aML software

Simple Conceptual Model

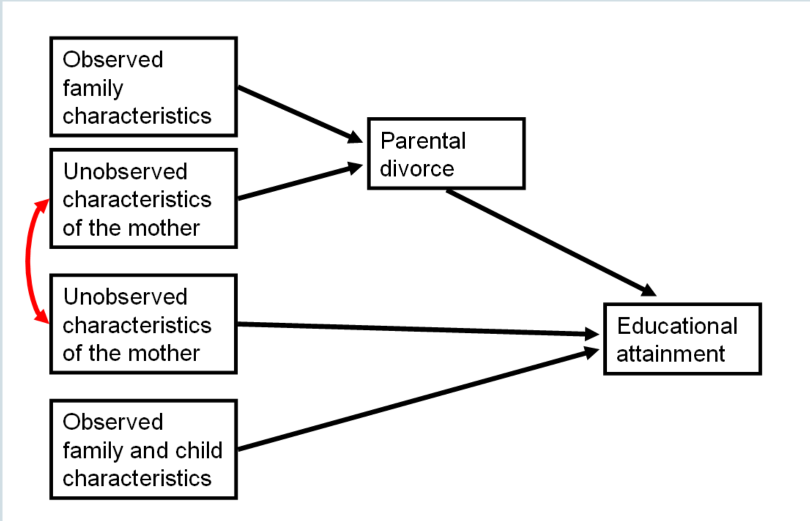

Sequential Probit Model for Educational Transitions (1)

- View educational qualifications as the result of 4 sequential transitions:
  - Compulsory to lower secondary
  - Lower to higher secondary (given reached lower sec.)
  - Higher secondary to Bachelor's (given higher sec.)
  - Bachelor's to postgraduate (given Bachelor's)
- Rather like a discrete-time event history model
- Advantages:
  - Allow effects of disruption to vary across transitions
  - Can include children who are too young to have made *all* transitions

Sequential Probit Model for Educational Transitions (2)

Transition from education level r for child i of woman j indicated $y_{ij}^{(r)}$ =1 if child attains level $r+1$ and 0 if stops at r.

$$y_{ij}^{(r)*} = \mathbf{x}_{ij}\beta^{(r)} + \mathbf{z}_{ij}\gamma^{(r)} + \lambda^{(r)}u_j + e_{ij}^{(r)}, \qquad r = 1, \ldots, 4$$

$y_{ij}^{(r)*}$ latent propensity underlying $y_{ij}^{(r)}$

$\mathbf{z}_{ij}$ potentially endogenous indicators of family disruption

$\mathbf{x}_{ij}$ child and mother background characteristics

u_j mother-specific random effect

$e_{ij}^{(r)}$ child and transition-specific residual

Unobserved Heterogeneity: Educational Transitions

Parameter	Transition	Estimate	(SE)
Random effect loading			
$\lambda^{(1)}$	To low secondary	$1^{\dagger}$	-
$\lambda^{(2)}$	To high secondary	1.078*	(0.041)
$\lambda^{(3)}$	To Bachelor's	0.864*	(0.041)
$\lambda^{(4)}$	To Master's +	0.584*	(0.077)
Standard deviation σ_u		0.498*	(0.014)

*Significant at 1% level; †Constrained to equal 1

Effect of mother-level unobservables less important for later transitions.

Evidence for Selection

- Residual correlation between dissolution risk and probability of continuing in education estimated as -0.43 (se=0.02)
- Suggests mothers with above-average risk of divorce tend to have children with below-average chance of remaining education
- Note that we are controlling only for selection on unobservables at the mother level (i.e. fixed across time)

Effects of Disruption on Transitions in Secondary School

	Compulsory to lower secondary		Lower to higher secondary	
	Model 1 ($\rho_{uv} = 0$)	Model 2 ($\rho_{uv} \neq 0$)	Model 1 ($\rho_{uv} = 0$)	Model 2 ($\rho_{uv} \neq 0$)
Parents separated	-0.580*	-0.349*	-0.631*	-0.386*
Age at separation	0.019*	0.013*	0.018*	0.013*
Father died	-0.201*	-0.178*	-0.318*	-0.295*
Female	0.217*	0.217*	0.318*	0.320*
Female×separation	0.034	0.034	0.005	0.003
Female×father died	-0.126	-0.125	0.107	0.110

*Significant at 1% level

Predicted Probabilities of Continuing Beyond Lower Secondary (Before and After Allowing for Selection)

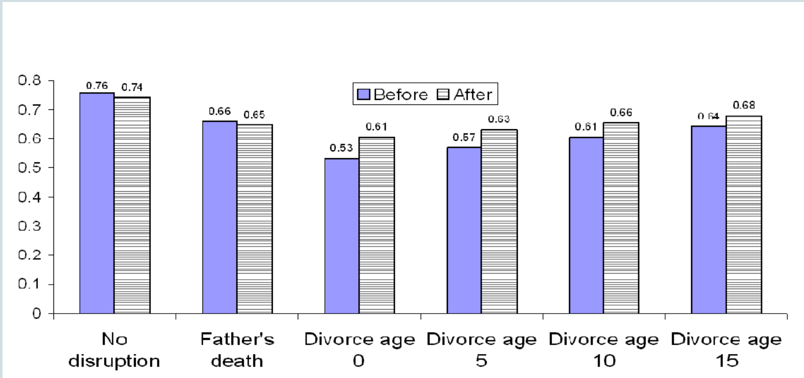

Software for multiprocess modelling

- Sabre
  - developed for analysis of recurrent events
  - handles mixtures of response types; up to 3 processes
  - 2 levels; no random coefficients
- aML
  - designed specifically for multilevel multiprocess modelling
  - mixtures of response types; multiple processes and levels
  - DOS-based; user needs to specify starting values
- MLwiN
  - designed for multilevel modelling; multiple levels and random coefficients
  - handles mixtures of continuous and binary responses
  - Markov chain Monte Carlo estimation

Bibliography: Recurrent Events

Davis, R.B., Elias, P. and Penn, R. (1992) The relationship between a husbands unemployment and a wifes participation in the labour-force. *Oxford Bulletin of Economics and Statistics*, 54:145-71.

Kravdal, . (2001) The high fertility of college educated women in Norway: An artefact of the separate modelling of each parity transition. *Demographic Research*, 5, Article 6.

Bibliography: Multiple States & Competing Risks

Enberg, J., Gottschalk, P. and Wolf, D. (1990) A random-effects logit model of work-welfare transitions. *Journal of Econometrics*, 43: 63-75.

Goldstein, H., Pan, H. and Bynner, J. (2004) A flexible procedure for analysing longitudinal event histories using a multilevel model. *Understanding Statistics*, 3: 85-99.

Hill, D.H., Axinn, W.G. and Thornton, A. (1993) Competing hazards with shared unmeasured risk factors. *Sociological Methodology*, 23: 245-277.

Steele, F., Goldstein, H. and Browne, W. (2004) A general multistate competing risks model for event history data, with an application to a study of contraceptive use dynamics. *Statistical Modelling*, 4: 145-159.

Bibliography: Multiprocess Modelling (1)

Angeles, G., Guilkey, D.K. and Mroz, T.A. (1998) "Purposive program placement and the estimation of family planning program effects in Tanzania." *Journal of the American Statistical Association*, 93: 884-899.

Bollen, K.A., Guilkey, D.K. and Mroz, T.A. (1995) "Binary outcomes and endogenous explanatory variables: tests and solutions with an application to the demand for contraceptive use in Tunisia." *Demography*, 32: 111-131.

Brien, M.J., Lillard, L.A. and Waite, L.J. (1999) Interrelated family-building behaviors: Cohabitation, marriage, and nonmarital conception. *Demography*, 36: 535-551.

Lillard, L.A., Brien, M.J. and Waite, L.J. (1995) "Premarital cohabitation and subsequent marital dissolution: a matter of self-selection?" *Demography*, 32: 437-457.

Bibliography: Multiprocess Modelling (2)

Steele, F., Kallis, C., Goldstein, H. and Joshi, H. (2005) The relationship between childbearing and transitions from marriage and cohabitation in Britain. *Demography*, 42: 647-673.

Steele, F., Sigle-Rushton, W. and Kravdal, . (2009) “Consequences of Family Disruption on Children’s Educational Outcomes in Norway”, *Demography*, 46: 553-574.

Upchurch, D.M., Lillard, L.A. and Panis, C.W.A. (2002) Nonmarital childbearing: Influences of education, marriage, and fertility. *Demography*, 39: 311-329.

Waite, L.J. and Lillard, L.A. 1991. Children and marital disruption. *American Journal of Sociology*, 96: 930-53.