Case study: Covariates and staged estimation in course engagement profiles with the latents R package

Mohammed Saqr, University of Eastern Finland
Sonsoles López-Pernas, University of Eastern Finland

26 November 2026

Contents

Data and engagement profiles 2
Accounting for classification error 4
Previous grade as a membership predictor 4
Comparing previous grades across profiles 6
Estimating profiles and grade associations together 8
Holding profile definitions fixed across student classes 9
Interpretation and limitations 11
References 12

Three-step covariate analysis is a method for examining associations between latent profile membership and variables that are external to the measurement model. An important feature of this approach is that it adjusts for uncertainty in profile classification. The adjustment reduces the bias that arises when individuals are assigned to latent classes with imperfect certainty.

In the course-engagement example, the primary research question is whether students' previous grades are associated with their engagement patterns in a subsequent course. This relationship can be investigated in several ways. One approach is to regress latent profile membership on previous grades in order to assess whether prior academic performance predicts the likelihood of belonging to a particular engagement profile. A second approach is to compare mean previous grades across the identified profiles. A third approach is to estimate the model jointly by including previous grades directly in the latent profile model.

These approaches differ in the extent to which the external variable is allowed to influence the latent profile solution. In the three-step approaches, the measurement model is estimated first, and the resulting profile definitions are held fixed while associations with previous grades are examined. This preserves the original interpretation of the engagement profiles. In one-step estimation, previous grade enters the latent profile model and can change the profiles. A later section turns to staged estimation, which holds the profile definitions fixed while the student classes are estimated.

```
library(latents)
```

Data and engagement profiles

The bundled synthetic `engagement_momentum` dataset contains enrolment records nested within students, following the design of the engagement studies of Saqr and López-Pernas (2021) and López-Pernas and Saqr (2024). Engagement is represented by five activity indicators: page views (`browse`), lecture videos viewed (`lectures`), forum posts read (`forum_read`), forum posts written (`forum_post`), and active days (`active_days`). Each indicator is simulated as a continuous value representing a log-activity scale and standardized within course to make activity levels comparable across courses.

The variable `previous_grade` represents the standardized grade obtained in the preceding course. For a student's first observed course, an entry grade is used instead. The variable `grade` is the standardized grade of the course itself, which is analysed as a distal outcome below. Because the previous grade is obtained before engagement in the current course, it can be treated as a predictor of later engagement patterns. The introductory vignette provides additional explanation of the longitudinal structure of the data and the distinction between the student and enrolment levels.

The indicator vector `vars` specifies the activity measures used to estimate the latent engagement profiles. In keeping `previous_grade` separate from these indicators, the engagement profiles can first be established using activity data alone. The relationship between previous academic performance and profile membership can then be examined in a subsequent stage without allowing grades to influence the initial definition of the profiles.

```
vars <- c("browse", "lectures", "forum_read", "forum_post", "active_days")
```

The `multilpa()` function estimates the means and variances that define each latent profile, together with the probability of observing each profile within each student class. Let this conditional probability be written as

$$\pi_{k|m} = P(C_{it} = k \mid G_i = m),$$

where C_{it} denotes the latent engagement profile of student i in course t, and G_i denotes the student's latent class. For each student class, the profile probabilities sum to one across all profiles. These probabilities therefore describe how the distribution of course-level engagement profiles differs across student-level classes.

The model considered here includes three engagement profiles and three student classes. Within each profile, the five engagement indicators are modeled as independent Gaussian variables. The profile definitions are held common across student classes, meaning that the same profile-specific means and variances apply to all students. Student classes differ instead in the probabilities with which the engagement profiles occur.

Model estimation relies on several assumptions. Students are assumed to be independent of one another, while enrolments are assumed to be conditionally independent given student class. The model also assumes conditional independence among the engagement indicators within each profile. Violations of this latter assumption can affect the profiles that are recovered from the data. In addition, because latent profile models use likelihood-based estimation, the optimization procedure may converge to different local maxima depending on its starting values.

The evaluation vignette discusses indicator dependence, likelihood maxima, and related model-evaluation issues in greater detail. In the present tutorial, the same candidate model is retained throughout so that the different covariate methods can be compared on a common basis. Ten starts (the first from Ward's clustering, the rest random) are used to reduce sensitivity to local solutions, and a fixed random seed ensures that the numerical results are reproducible.

```
fit <- multilpa(engagement_momentum, vars, id = "student", time = "sequence",
                n_profiles = 3, n_group_classes = 3, n_starts = 10,
                max_iter = 2000, tol = 1e-10, seed = 1)
fit
#> Two-level latent profile analysis: 3 profiles, 3 group classes
#> 1310 individuals in 106 groups; varying diagonal residual covariance (VVI)
#> Log likelihood: -7487.954642 | AIC: 15051.909 | BIC (groups): 15153.120
#> Converged: TRUE | iterations: 24 | best log likelihood reached by 10 of 10 starts (10 converged)
#> Smallest standardized Gaussian eigenvalue: 0.3005 (profiles, profile_1); Gaussian mass: 442.5; weight effec
#>
#> Class sizes   profile_1 33.8% (442)   profile_2 24.4% (320)   profile_3 41.8% (548)
#> Group classes group_class_1 19.7% (21)   group_class_2 44.6% (47)   group_class_3 35.7% (38)
#>
#> Means
#>    indicator profile_1 profile_2 profile_3
#>       browse    -0.928    -0.012     0.756
#>     lectures    -0.794    -0.142     0.724
#>   forum_read    -0.980    -0.079     0.837
#>   forum_post    -0.873    -0.055     0.737
#>  active_days    -0.961    -0.029     0.793
#>
#> Variances and standard errors: get_results(x, "profiles").
#> Every other table: get_results(x, what = ), or get_results(x, "all").
```

The three profiles represent the **disengaged**, **average** and **active** engagement states, respectively. This ordering is evident in the estimated means for `active_days`, which are -0.961 for Profile 1, -0.029 for Profile 2, and 0.793 for Profile 3. Profile 3, the active profile, is the last profile and therefore the reference category of the profile-membership regressions presented below.

These labels describe the current fitted solution; the profile numbers themselves carry no fixed meaning. Because latent profile labels are determined by the estimated parameters in each model, the profile-specific means should be examined again whenever a new model is fitted to confirm that the same substantive interpretation still applies.

Accounting for classification error

The three-step approach accounts for uncertainty in profile assignment by using a classification-error matrix. This matrix describes the probability that an individual is assigned to a particular profile given their underlying latent profile. Let

$$D_{ks} = P(\widehat{C} = s \mid C = k),$$

where C denotes the underlying latent profile and $\widehat{C}$ denotes the assigned profile, defined as the profile with the largest posterior probability. Each row of D sums to one because it represents the distribution of possible assignments for a given underlying profile. The diagonal elements therefore give the estimated probabilities of being assigned to the correct profile, whereas the off-diagonal elements represent misclassification between profiles. In `multilpa`, this matrix can be extracted with `get_results()` using `"classification_errors"`.

The correction relies on the assumption that, conditional on the underlying latent profile, classification errors are independent of the auxiliary variable being studied. In other words, once true profile membership is taken into account, the probability of being misclassified is assumed not to depend systematically on the covariate or distal outcome.

The classification-error matrix is estimated from the initial measurement model and is subsequently treated as fixed during the covariate analysis. As a result, confidence intervals and standard errors from the later stages do not incorporate uncertainty associated with estimating this matrix. This distinction is important when interpreting inferential results, particularly when profile separation is limited. Poorly separated profiles produce greater classification uncertainty and can make the error correction less stable.

```
get_results(fit, "classification_errors", level = "individuals")
#>         level true_class assigned_class probability
#> 1 individuals          1              1   9.495e-01
#> 2 individuals          1              2   5.048e-02
#> 3 individuals          1              3   1.786e-06
#> 4 individuals          2              1   6.525e-02
#> 5 individuals          2              2   8.274e-01
#> 6 individuals          2              3   1.073e-01
#> 7 individuals          3              1   1.314e-05
#> 8 individuals          3              2   6.073e-02
#> 9 individuals          3              3   9.393e-01
```

The estimated probabilities of correct profile assignment are 0.950 for the disengaged Profile 1, 0.827 for the average Profile 2, and 0.939 for the active Profile 3. Classification is therefore least precise for the average profile, which lies between the other two.

The main source of ambiguity is the overlap of the average profile with its two neighbours. An enrolment from the average profile is assigned to the active profile with probability 0.107 and to the disengaged profile with probability 0.065, whereas the disengaged and active profiles are almost never confused with each other. These off-diagonal probabilities indicate that a non-negligible proportion of observations may be classified into a neighbouring profile rather than their underlying profile.

Treating the assigned profiles as if they were observed without error would ignore this uncertainty. The three-step correction instead carries the estimated misclassification structure forward into the subsequent covariate analysis, so associations with external variables are estimated with classification error taken into account.

Previous grade as a membership predictor

The `r3step()` function estimates the association between an external predictor and latent profile membership while keeping the profile definitions and classification-error matrix fixed. In this example, `previous_grade`

is used to predict membership in the engagement profiles identified in the measurement model. `r3step()` uses the last profile as the reference category; here that is Profile 3, the active profile.

Let x denote standardized previous grade. For Profiles 1 and 2, the multinomial regression can be written as

$$\log \frac{P(C = k \mid x)}{P(C = 3 \mid x)} = \alpha_k + \beta_k x, \qquad k = 1, 2.$$

The intercept α_k represents the log odds of belonging to Profile k, rather than the active reference profile, when $x = 0$. Because `previous_grade` is standardized, $x = 0$ corresponds to the average previous grade. The slope β_k represents the change in these log odds associated with a one-standard-deviation increase in previous grade. A negative slope therefore indicates that higher previous grades are associated with lower relative odds of belonging to Profile k rather than the active profile.

The model assumes that previous grade has a linear effect on the log-odds scale. For analyses conducted at the enrolment level, `vcov_type = "robust"` can be used to account for the clustering of multiple enrolments within the same student. Under this option, score contributions are clustered by student when estimating the covariance matrix. Reliable estimation therefore requires a sufficient number of independent students.

The `r3step()` regression describes the marginal association between previous grade and profile membership. This differs from the one-step model considered later, in which profile membership is modeled conditional on student class. The two approaches therefore answer related but distinct questions about how previous academic performance is associated with subsequent engagement.

The reported 95% confidence intervals are based on the normal approximation. When several slope coefficients are tested simultaneously, the reported adjusted p-values use the Benjamini–Hochberg procedure to control the false discovery rate across those tests. The confidence intervals remain pointwise intervals and are not adjusted for multiple comparisons.

```
r3step(fit, data = engagement_momentum, covariates = "previous_grade", vcov_type = "robust")
#>        level   outcome          term estimate standard_error statistic   p_value p_adjusted
#> 1 individuals profile_1   (Intercept)  -0.2686        0.1478    -1.817 6.922e-02        NA
#> 2 individuals profile_1 previous_grade  -1.0345        0.1076    -9.619 6.664e-22  1.333e-21
#> 3 individuals profile_2   (Intercept)  -0.4519        0.1202    -3.760 1.696e-04        NA
#> 4 individuals profile_2 previous_grade  -0.4328        0.1099    -3.939 8.193e-05  8.193e-05
#>   conf_low conf_high
#> 1  -0.5583   0.02113
#> 2  -1.2453  -0.82372
#> 3  -0.6874  -0.21635
#> 4  -0.6482  -0.21745
```

In the fitted model, a one-standard-deviation increase in previous grade changes the log odds of belonging to the disengaged profile (Profile 1) rather than the active profile by -1.035 (95% confidence interval -1.245 to -0.824). For the average profile (Profile 2), the estimated decrease in log odds is 0.433, with a 95% confidence interval from -0.648 to -0.217.

Both coefficients are negative: higher previous grades are associated with lower relative odds of membership in the average or disengaged profiles compared with the active profile. The association is strongest for the disengaged profile. The results therefore show a consistent association between stronger previous academic performance and more active engagement in the subsequent course.

These findings are observational associations. Although `previous_grade` is measured before engagement in the subsequent course, temporal ordering alone does not establish that higher grades cause students to adopt a more active engagement pattern. In the generating model of these data, the previous grade has no effect on the engagement state: the association arises because the previous grade reflects the engagement state of the previous course, and engagement states persist.

Comparing previous grades across profiles

The `three_step()` function estimates profile-specific means of a numeric auxiliary variable and tests differences between those means while accounting for uncertainty in profile classification. In this example, `previous_grade` is treated as the auxiliary variable, so prior academic performance can be compared across the engagement profiles identified in the measurement model.

By default, `three_step()` uses the Bolck–Croon–Hagenaars (BCH) method. This approach constructs signed weights from the inverse of the classification-error matrix and uses those weights to correct for imperfect profile assignment. Let x_u denote the previous grade for enrolment u, and let w_{uk} denote its BCH weight for Profile k. The estimated profile-specific mean is

$$\hat{\mu}_k = \frac{\sum_u w_{uk} x_u}{\sum_u w_{uk}}.$$

Differences between profiles are then expressed as pairwise contrasts. For Profiles k and l, the contrast is

$$\hat{\mu}_k - \hat{\mu}_l.$$

The BCH correction requires the classification-error matrix to be invertible and the resulting weighted profile sizes to be sufficiently well behaved for estimation. The default standard errors account for the clustering of multiple enrolments within students, while the previously estimated classification-error matrix is treated as fixed. With the default `step_one = "fixed"`, uncertainty from estimation of that matrix is not propagated into the reported standard errors; `step_one = "estimated"` adds its first-order contribution.

Although `three_step()` can also be used with distal outcomes measured after profile membership, `previous_grade` occurs before the engagement measures in this example. The estimated contrasts should therefore be interpreted as differences in prior academic performance across subsequently observed engagement profiles rather than as effects of profile membership on grades.

The reported 95% confidence intervals are pointwise normal intervals. Adjusted p-values use the Benjamini–Hochberg procedure across the set of pairwise profile contrasts.

```
three_step(
  fit,
  data = engagement_momentum,
  outcome = "previous_grade",
  contrast = "pairs"
)
#> Distal outcome: class differences (bch, 95% CI)
#>
#> Class  vs  Difference        95% CI      p  p (BH)
#> -----  --  ----------  ------------  -----  ------
#>     2   1        0.53  [0.34, 0.71]  <.001   <.001
#>     3   1        0.89  [0.71, 1.07]  <.001   <.001
#>     3   2        0.36  [0.18, 0.55]  <.001   <.001
```

The BCH-adjusted estimates show that previous grades in the disengaged profile (Profile 1) are 0.89 standard deviations below those in the active profile (Profile 3) and 0.53 standard deviations below those in the average profile (Profile 2). These contrasts indicate that enrolments in which students are disengaged tend to follow courses with lower grades.

The difference between the average and active profiles is smaller. The estimated Profile 3 versus Profile 2 contrast is 0.36, with a 95% confidence interval from 0.18 to 0.55: the active profile follows higher grades than the average one, but less markedly than either differs from the disengaged profile.

Comparison with modal assignment

A useful comparison is provided by modal analysis, in which each enrolment is assigned with certainty to the profile having the largest posterior probability. The same profile means and pairwise contrasts are then calculated using these hard assignments rather than classification-error-adjusted weights.

```
three_step(
  fit,
  data = engagement_momentum,
  outcome = "previous_grade",
  contrast = "pairs",
  method = "modal"
)
#> Distal outcome: class differences (modal, 95% CI)
#>
#> Class  vs  Difference        95% CI      p  p (BH)
#> -----  --  ----------  ------------  -----  ------
#>     2   1        0.50  [0.34, 0.66]  <.001   <.001
#>     3   1        0.84  [0.67, 1.01]  <.001   <.001
#>     3   2        0.34  [0.19, 0.49]  <.001   <.001
```

Modal analysis is simple to interpret, but it treats estimated profile assignments as if they were observed without error. Its standard errors therefore do not incorporate uncertainty associated with imperfect classification.

The effect of this simplification is particularly visible for the average profile, which showed the greatest classification overlap. Under modal assignment, the estimated Profile 3 versus Profile 2 difference is 0.340, compared with 0.364 under the BCH correction, and the corresponding standard error is 0.076 under modal assignment and 0.095 under BCH. Modal assignment therefore both attenuates the contrast and understates its uncertainty.

Thus, accounting for classification error affects both the estimated contrast and its estimated uncertainty. The understatement of uncertainty is largest for the two profiles that are hardest to distinguish. Classification adjustment matters most when external variables are compared across partially overlapping latent profiles.

The course grade as a distal outcome

The grade of the course itself is measured after the engagement it follows, and can therefore be treated as a distal outcome of the engagement profile. The same BCH comparison applies.

```
three_step(fit, data = engagement_momentum, outcome = "grade",
           contrast = "pairs")
#> Distal outcome: class differences (bch, 95% CI)
#>
#> Class  vs  Difference        95% CI      p  p (BH)
#> -----  --  ----------  ------------  -----  ------
#>     2   1        0.81  [0.68, 0.95]  <.001   <.001
#>     3   1        1.52  [1.39, 1.65]  <.001   <.001
#>     3   2        0.71  [0.55, 0.87]  <.001   <.001
```

Course grades are 1.52 standard deviations higher in active than in disengaged enrolments, 0.81 higher in average than in disengaged enrolments, and 0.71 higher in active than in average enrolments. The engagement state of a course is therefore strongly associated with the grade obtained in it, in line with the link between engagement and achievement described by López-Pernas and Saqr (2024). As with previous grade, the comparison is observational.

Estimating profiles and grade associations together

One-step estimation differs from the staged approaches considered above because the latent profile measurement model and the regression of profile membership on the covariate are estimated simultaneously. In this example, specifying `profile_covariates = "previous_grade"` in `multilpa()` allows information about previous grades to contribute directly to estimation of the engagement profiles as well as to the association between grades and profile membership.

Let x denote standardized previous grade and $G = m$ denote membership in student class m. With the active Profile 3 as the reference category, the profile-membership model is

$$\log \frac{P(C = k \mid G = m, x)}{P(C = 3 \mid G = m, x)} = \alpha_{km} + \beta_k x, \qquad k = 1, 2.$$

The intercepts α_{km} are allowed to vary across student classes. They therefore describe differences in the baseline profile probabilities between student classes when previous grade is at its standardized mean, $x = 0$. In contrast, the grade slopes β_k are shared across student classes. Each slope represents the change in the log odds of belonging to Profile k, rather than Profile 3, associated with a one-standard-deviation increase in previous grade, conditional on student class.

The one-step model retains the same Gaussian measurement assumptions used in the original latent profile model, including conditional independence of the engagement indicators within profiles. The important difference is that previous grade is now part of the model during estimation rather than being introduced only after the profiles have been established.

As a consequence, the profiles obtained from the one-step model need not be identical to those from the original measurement model. Allowing previous grade to contribute to estimation can alter profile-specific means, variances, or membership probabilities. The estimated regression coefficients also have a different interpretation from the three-step coefficients because they describe the association between previous grade and profile membership conditional on student class, whereas the earlier `r3step()` analysis describes marginal profile membership.

For these reasons, the magnitude of the one-step coefficients should not be expected to match the corresponding three-step estimates exactly. Before interpreting the coefficients, the fitted profile means should be inspected to determine whether Profiles 1, 2, and 3 still correspond to the disengaged, average and active engagement patterns. Profile numbers are labels generated by the fitted model and do not carry fixed substantive meanings across separate model fits.

```
covariate_fit <- multilpa(
  engagement_momentum,
  vars,
  id = "student",
  n_profiles = 3,
  n_group_classes = 3,
  profile_covariates = "previous_grade",
  n_starts = 6,
  max_iter = 2000,
  tol = 1e-10,
  seed = 1
)

covariate_fit
#> Multilevel LPA with covariates: 3 profiles, 3 group classes
#> Log likelihood -7479.952583; AIC 15039.905; BIC (groups) 15146.443; converged TRUE
#> Smallest standardized Gaussian eigenvalue: 0.3001 (profiles, profile_1); Gaussian mass: 442.6; weight effec
#>
```

```
#> Class sizes   profile_1 33.8% (443)   profile_2 24.6% (322)   profile_3 41.6% (545)
#> Group classes group_class_1 18.8% (20)  group_class_2 45.5% (48)  group_class_3 35.7% (38)
#>
#> Means
#>    indicator profile_1 profile_2 profile_3
#>       browse    -0.928    -0.012     0.760
#>     lectures    -0.794    -0.138     0.726
#>   forum_read    -0.980    -0.073     0.839
#>   forum_post    -0.873    -0.050     0.738
#>  active_days    -0.959    -0.026     0.794
#>
#> Variances and standard errors: get_results(x, "profiles").
#> Every other table: get_results(x, what = ), or get_results(x, "all").
```

The fit converges, and the profiles keep the ordering of the measurement-only model: Profile 1 is disengaged, Profile 2 average and Profile 3 active, with `active_days` means of -0.959, -0.026 and 0.794. The means have moved slightly because previous grade now contributes to their estimation, and Profile 3, the active profile, remains the reference category.

The `"coefficients"` table holds one intercept for each profile and student class, one `previous_grade` slope for each profile relative to the reference profile, and the intercepts of the student-class model.

```
get_results(covariate_fit, "coefficients")
#>      level       outcome           term parameter estimate
#> 1  profile     profile_1  group_class_1     logit  3.74685
#> 2  profile     profile_1  group_class_2     logit -0.01407
#> 3  profile     profile_1  group_class_3     logit -2.23602
#> 4  profile     profile_1 previous_grade     logit -0.42185
#> 5  profile     profile_2  group_class_1     logit  1.86423
#> 6  profile     profile_2  group_class_2     logit  0.08488
#> 7  profile     profile_2  group_class_3     logit -1.28612
#> 8  profile     profile_2 previous_grade     logit -0.21134
#> 9    group group_class_1    (Intercept)     logit -0.63975
#> 10   group group_class_2    (Intercept)     logit  0.24165
```

The `previous_grade` slopes are -0.422 for the disengaged profile and -0.211 for the average profile, each relative to the active profile and conditional on student class. Higher previous grades therefore go with lower odds of the less active profiles, the same direction as the `r3step()` slopes of -1.035 and -0.433. The magnitudes differ because these slopes are conditional on student class and the profiles are estimated together with the covariate.

Holding profile definitions fixed across student classes

Staged estimation separates the estimation of the course-level engagement profiles from the estimation of the higher-level student-class structure. In the first stage, a common set of engagement profiles is estimated without allowing student classes to alter the profile measurement parameters. In the second stage, the composition of student classes is estimated while the profile means and variances from the first stage are held fixed.

The `fit_staged()` function implements this procedure by first fitting a model with a single student class and then estimating the requested three-class student-level structure. During the second stage, the model estimates the profile probabilities within each student class,

$$\pi_{k|m} = P(C_{it} = k \mid G_i = m),$$

together with the marginal student-class probabilities,

$$\omega_m = P(G_i = m).$$

The profile-specific means and variances remain fixed at their first-stage estimates. Staged estimation therefore isolates the estimation of the higher-level class structure from the estimation of the profile measurement model.

This separation has an important inferential consequence. Uncertainty in the first-stage measurement parameters is not propagated into the second-stage estimates. The resulting model is therefore conditional on the fixed first-stage profile solution and should not be interpreted as the joint maximum-likelihood solution that would be obtained by estimating all parameters simultaneously.

Profile numbering also requires careful attention when comparing staged and joint models. Latent profile and student-class labels are arbitrary numerical identifiers, and separate estimation procedures may assign different numbers to substantively similar groups. Comparisons should therefore be based on the estimated activity patterns and class compositions rather than on the numerical labels alone.

```
staged <- fit_staged(
  engagement_momentum,
  vars,
  id = "student",
  n_profiles = 3,
  n_group_classes = 3,
  n_starts = 6,
  max_iter = 2000,
  tol = 1e-10,
  seed = 1
)

staged
#> Two-level latent profile analysis: 3 profiles, 3 group classes
#> 1310 individuals in 106 groups; varying diagonal residual covariance (VVI)
#> Log likelihood: -7489.094665 | AIC: 14994.189 | BIC (groups): 15015.497
#> Converged: TRUE | iterations: 21 | best log likelihood reached by 5 of 6 starts (6 converged)
#> Held fixed, not estimated here: means, variances (first stage)
#> Parameters estimated here: 8; with the held measurement: 38
#> Smallest standardized Gaussian eigenvalue: 0.3053 (profiles, profile_1); Gaussian mass: 434.1; weight effec
#>
#> Class sizes   profile_1 33.1% (434)   profile_2 24.5% (321)   profile_3 42.4% (555)
#> Group classes group_class_1 44.6% (47)   group_class_2 19.5% (21)   group_class_3 35.8% (38)
#>
#> Means
#>    indicator profile_1 profile_2 profile_3
#>       browse    -0.944    -0.031     0.752
#>     lectures    -0.805    -0.167     0.725
#>   forum_read    -0.984    -0.128     0.841
#>   forum_post    -0.882    -0.086     0.737
#>  active_days    -0.989    -0.025     0.783
#>
#> Variances and standard errors: get_results(x, "profiles").
#> Every other table: get_results(x, what = ), or get_results(x, "all").
```

In the staged solution, Profile 1 represents the disengaged state, Profile 2 the average state, and Profile 3 the active state. Their estimated `active_days` means are -0.989, -0.025, and 0.783, respectively. The profile numbers happen to match those of the earlier joint model, which separate fits do not guarantee; the student-class numbers do not, since staged class 1 (47 students) corresponds to joint class 2 and staged class 2 (21 students) to joint class 1. In both models the same broad engagement patterns remain recognizable. The estimated means differ slightly because the profiles were obtained under a different estimation procedure.

The parameters estimated at each stage can be examined through the `"stages"` results table:

```
get_results(staged, "stages")
#>       stage group_classes        fixed log_likelihood parameters parameters_with_measurement
#> 1 measurement            1           <NA>         -7682        32                         32
#> 2  membership            3 means, variances         -7489         8                         38
#>   converged
#> 1      TRUE
#> 2      TRUE
```

The measurement stage estimates 32 parameters. The subsequent membership stage estimates 8 additional parameters governing the student-class composition. When the profile parameters retained from the measurement stage are also counted, the membership-stage model contains 38 `parameters_with_measurement`. This latter quantity is useful when making descriptive comparisons with a jointly estimated model because it reflects both the newly estimated higher-level parameters and the measurement parameters carried forward from the first stage.

The second-stage log likelihood is -7489.095, compared with -7487.955 for the joint model reported earlier. The lower log likelihood is consistent with the additional restriction imposed by staged estimation: the second stage cannot adjust the profile means and variances to improve the overall fit because those parameters have already been fixed. The two likelihoods should therefore be interpreted in light of the different estimation constraints.

Overall, staged estimation provides a way to preserve a previously established set of engagement profiles while estimating how those profiles are distributed across student classes. Its main advantage in this context is interpretive stability: the higher-level class structure is estimated without allowing the course-level profile definitions to change during the second stage.

Interpretation and limitations

The membership regressions and grade contrasts indicate a consistent relationship between academic performance and engagement. Disengaged enrolments follow lower grades than active and average ones, and the grade of a course is strongly associated with the engagement state in it. The difference between the average and active profiles is smaller, and it is estimated less precisely because these two profiles overlap most in classification.

The different estimation strategies should not be interpreted as interchangeable versions of the same analysis. Three-step methods retain the profile definitions established in the measurement model and then examine their associations with external variables. In contrast, one-step estimation allows the covariate to contribute directly to estimation of the profiles and estimates the membership regression conditional on student class. The choice between these approaches therefore depends on the analytic objective: whether the priority is to preserve an existing profile solution or to estimate the measurement and covariate components jointly.

All of these results remain conditional on the adequacy of the underlying measurement model. As discussed in the evaluation vignette, assumptions such as diagonal within-profile covariance structures, the selected numbers of profiles and student classes, and the stability of the likelihood solution require careful evaluation. Classification-error corrections can account for uncertainty in assigning observations to latent profiles, but they cannot compensate for a poorly specified profile model.

Several inferential limitations also remain. Robust standard errors that cluster enrolments within students account for dependence in the observed data structure, but they do not convert an observational association into a causal effect. In addition, both staged and three-step procedures treat quantities estimated in earlier stages as fixed. Their reported standard errors and confidence intervals therefore do not propagate all of the uncertainty associated with estimating the initial measurement model or classification-error matrix.

References

López-Pernas, S., & Saqr, M. (2024). How the dynamics of engagement explain the momentum of achievement and the inertia of disengagement: A complex systems theory approach. *Computers in Human Behavior, 153*, 108126. doi:10.1016/j.chb.2023.108126.

Saqr, M., & López-Pernas, S. (2021). The longitudinal trajectories of online engagement over a full program. *Computers & Education, 175*, 104325. doi:10.1016/j.compedu.2021.104325.