Week 12 Exercises

Research Paper. Continue to make progress on your research paper. The next draft is due Tuesday of Thanksgiving week. This draft should be close to finalized, so that I can make suggestions to help polish the paper.

Workshop. Finalize your materials. Some comments from me soon.

Rehearsals:

- *Cox*. Monday, November 10 at 9am (to about 10am).
- *Unordered* and *ordered*. Friday, November 14 at 12pm (to about 2pm).

Workshops: I'll solicit RIBC students.

- *Unordered*. Tuesday, November 18 from 1pm to 2pm.
- *Ordered*. Wednesday, November 19 from 1pm to 2pm.
- *Cox*. Friday, November 21 from 11am to 12noon.

Exercise 1 Simulate and recover; zero-inflated negative binomial

One important way to test and develop your understandings of statistical models and quantities of interest is to *simulate* a fake data set with known true quantities of interest (e.g., coefficients or first difference) and then *recover* those quantities of interest though your estimation procedure.

Do this for a ***zero-inflated* negative binomial** regression.

Quantities of Interest

As your quantity of interest, use the percent increase in the expected count. That's $\frac{E(Y|X_{hi})}{E(Y|X_{lo})} - 1$.

1. Compute this percent increase setting the non-focal variables at their means (or modes, if qualitative).

2. Compute the "average percent increase" setting the non-focal variables at all observed values and then averaging the estimates.

Simulate

Simulate a fake data set with a known quantity of interest.

1. Choose a number of observations.
2. Create several predictors.
3. Choose values for all the parameters. Let the zero-inflation depend on covariates.
4. Simulate an outcome variable. For the ZINB, you can do this in three steps:
    a. First, simulate a Bernoulli variable to capture whether this observation is a structural zero or not.
    b. Second, simulate a NB outcome.
    c. If the case is a structural zero, then set y equal to zero. If not, then set y equal to the NB outcome.

Recover

1. Fit the zero-inflated negative binomial model and check that you have recovered the model parameters (including the coefficients for the zero-inflation portion).
2. Use {marginaleffects} to recover the two quantities of interest described above.
    a. The percent increase for a typical case.
    b. The average percent increase across the observed values.

Exercise 2 IRT

1. Read Gelman and Hill (2007), pp. 314-317 (section 14.3 up to the subsection titled "An ideal point model for Supreme Court voting.")
2. The R code here fits the Rasch 1-PL model shown in Eqn 14.10 (or 14.11) to the `answers.csv` dataset. Run this script to replicate my results (e.g., from the slides).
3. Modify the code to also fit the 2-PL model shown in Eqn. 14.13.
4. Compare the ability parameters from the two models, especially the rankings of the abilities.
5. ~~Which model has a lower LOOIC?~~

Exercise 3 Ideal Point Estimation

1. Read Gelman and Hill (2007), pp. 317-319 (the subsection titled "An ideal point model for Supreme Court voting.")

2. Make sure to understand how the "ability" and "difficulty" parameters translate into ideology.
3. Use the IRT models in Eqns. 14.11 and/or 14.13 to estimate the ideology of the justices in the dataset below. This is a recoded version of `?MCMCpack::SupremeCourt`.

- `case_id`: is an arbitrary case number.
- `justice`: is the name of the justice.
- `appointed_by`: is the name of the president that initially appointed the justice.
- `cons_vote`: is an indicator for a *conservative* vote in the case.

```
sc <- read_csv("https://pos5747.github.io/data/sc2000.csv") |>
  glimpse()
```

```
Rows: 387
Columns: 4
$ case_id      <chr> "Case 1", "Case 1", "Case 1", "Case 1", "Case 1", "Case 1~
$ justice      <chr> "Rehnquist", "Stevens", "O'Connor", "Scalia", "Kennedy", ~
$ appointed_by <chr> "Nixon", "Ford", "Reagan", "Reagan", "Reagan", "Bush (41)~
$ cons_vote    <dbl> 1, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 1, 1, 1, 0, 1, 0, 0, 1, ~
```

Exercise 4 MRP

1. Read Gelman and Hill (2007), Section 14.1, pp. 301-310. Make sure to understand the *goals* of MRP and the role of hierarchical model in the MRP approach.
2. Read chapter 12 of *Model to Meaning*, which has an example of MRP.
3. You can find the core code from chapter 12 referenced above here. Reproduce the results.
4. Now tinker with the model.
    a. Try a model with more or different random slopes or different predictors. (Hint: See Gelman and Hill (2007), p. 303 for a 'fuller' model that includes, for example, age as `+ (1 | age)` rather than `+ age`.)
    b. Examine how these changes affect the estimates of state-level opinion. Create a plot that allows you to easily compare the two sets of estimates and see how they change.
    c. Use the LOOIC to help you identity the "better" set of estimates.

None planned. See example code here.
