

Econ 21410 - Problem Set V

Paper Replication*

May 6, 2014

This homework should be done in LaTeX. The homework will be graded on correctness, but will also heavily weight clarity and professionalism. It's better to not do some of the harder parts than to turn in an incomprehensible document. Your R script as well as a log-file should be submitted. Alternatively, use knitr to print your code inline in your latex document.

SUBMISSION: The homework must be emailed to Oliver and myself by 2p.m. Monday, May 12th. The email must include a pdf with the filename `lastname_pset6.pdf` and R code called `lastname_pset6_code.R` where "lastname" should be replaced with your last name. The subject of your email should be [ECON 21410: pset6 submission]

If you are struggling, please use the github page to ask for help!* Remember that asking and answering questions on our github page, coming to office hours to ask questions, and contributing to the class wiki are all worth participation credit, which is 10% of your grade in this class.

In this homework we will be replicating a classic paper from Industrial Organization: Bresnahan and Reiss (1991) "Entry and Competition in Concentrated Markets". It can be downloaded here: <http://pages.stern.nyu.edu/~acollard/bresnahan-reiss.pdf>

The data from the paper can be downloaded from my site or github.

Replicating Bresnahan and Reiss 1991

Using the provided data, replicate Table 4, Table 5-A, and Figure 4 from the paper but only for tire dealers. Describe how you performed the replication, what your calculations were, how they were implemented, the optimizer you used, etc. Your write-up should be written as if you were performing this replication as part of your research, or part of your Job. Include clear details of what you did, but do not be overly verbose.

To calculate your standard errors, use the "hessian()" command in the "numDeriv" library. This command will give a reasonably accurate numerical approximation to the matrix of second derivatives. Recall that your maximum likelihood standard errors will be the square-root of the diagonal terms of the inverted Hessian evaluated at your estimated parameter values.

Some hints and suggestions:

*Please email johneric@uchicago.edu and obrowne@uchicago.edu if you have questions.

1. Solve the problem using maximum likelihood. Make a function that takes data and parameter values as inputs and returns the negative log-likelihood.
2. To test your log-likelihood function, try using the parameter values reported in their paper. You should get roughly the same log-likelihood (differences larger than 1 may indicate a problem in your code).
3. In the paper, the α and γ parameters are constrained to be greater than or equal to 0 (except γ_L). These constraints are “Bound Constraints” and should not be too difficult to impose (but will force you to use a particular optimizer in `optim()`).
4. This problem is solvable with `optim()`, but you will need to give it smart starting values. You can make sure it works by starting with the parameters from the paper, but then you should find alternative starting values that may be what you would have used if writing this paper yourself and did not know the true solutions.
5. Your function will most likely produce some values of “-Inf”. These can cause problems and produce NaNs in later steps. You may wish to replace -Inf values with large negative numbers like -1000000000000000000 to avoid this problem.
6. It should be possible to write a relatively short log-likelihood function. Without comments or trying to make my code particularly compact, mine in 19 lines long and references two other functions which are both 9 lines long (without comments).

One Additional step

The optimizer called by “`optim()`” is approximating the gradient with finite-differences. The library `numDeriv` offers an improved way to numerically approximate derivatives that tend to be more accurate. Using the “`grad()`” function in the `numDeriv` library, write a new function that returns the numeric gradient.

Use this function in the “`gr`” option of “`optim()`” to have `optim()` use this function rather than the built in finite-differences method. Does this change your solutions? Are they closer to the ones reported in the paper? Use “`system.time`” to time both methods, report and discuss (1-2 sentences).

Side Projects

- Write a function above that returns the analytic gradient. Use it in the step above and report the timing (up to 2.5 points).
- Program some or all of this assignment in C++ using `Rcpp` or in Julia (up to 6 points).
- Program some or all of assignment 2 or 3 in C++ using `Rcpp` or in Julia (up to 6 points).
- Write down 3 research ideas. For each explain the idea in 4-10 sentences. Write 1-4 sentences on why it is interesting, write 1-4 sentences on what data you would use, write 1-4 sentences on what methods it would use. These ideas can be very rough. If you would like feedback on one or all of them, let us know (up to 1 point each, whole assignment can be done twice)

- Find a paper that looks at market entry and write a 1-3 page report on it (up to 3.5 points).
- Find up to 4 articles in the news and generate a related research ideas. Explain the idea in 1-4 sentences and how you might execute the research idea (1-3 sentences) (up to 1 point per idea).
- Read any published academic article by Gary Becker and write a 2-3 page summary and response (up to 3 points).
- Commonly for grants or applications you will be required to write research proposals. Write a research proposal that follows the guidelines set out by the NSF Graduate Research Fellowship. Under “Graduate Research Statement” here:
http://www.nsfgrfp.org/how_to_apply/application_materials#proposed. (up to 5 points).