RESERVE BANK OF NEW ZEALAND
TE PŪTEA MATUA

VICTORIA UNIVERSITY OF WELLINGTON
TE HERENGA WAKA

Chair in Public Finance

TaxSim NZ - new horizons in tax-and-transfer microsimulation modelling

Christopher Ball
chrispball1@gmail.com

2 June, 2021

Outline

1. Introduction
2. Data
3. Labour supply responses
4. Choice set
5. Extra TaxSim features
6. TaxSim demonstration & TaxSim development

Introduction

TaxSim - Microsimulation Modelling

- TaxSim is a microsimulation model of the New Zealand tax-and-transfer system.
- TaxSim uses individual and family characteristics to calculate disposable income, or income after taxes and transfers. This is similar to Treasury's TAWA model, and it's predecessors Taxwell and Taxmod.
- TaxSim also uses a discrete choice model to calculate labour supply responses (or changes in hours worked) in response to tax-and-transfer policy changes. This is similar to Treasury's defunct Taxwell-B and Taxmod-B models.
- TaxSim is an easy-to-use graphical interface, which simplifies running the policy changes and viewing the output.

TaxSim Timeline

March 2018 Treasury release source code for their microsimulation model TAWA to Github:

https://github.com/Treasury-Analytics-and-Insights/TAWA_microsimulation_model.

July 2018 Treasury and CPF agree to a code sharing arrangement for the Treasury's preliminary work adding a behavioural extension to TAWA. Treasury advise they will no longer maintain a labour supply microsimulation model.

October 2019 First "working" version of TaxSim created (version 0.3). This version was not released.

February 2020 First TaxSim workshop with Government agencies. Version 0.4.1 released and deployed to shinyapps.io.

July 2020 Latest version of TaxSim released (version 0.4.2). At this time the next version was expected in August 2020. Source code for TaxSim 0.4.2 released to agencies shortly after.

TaxSim - Main Developments

- TaxSim provides an easy-to-use graphical interface which allows for running policy changes with and without labour supply changes.
- Instead of using a Stats-NZ database, TaxSim uses a synthetic database (based on HLFS). This means that anyone can use the model and share results without reference to Stats-NZ.
- TaxSim uses a novel choice set that more accurately represents the labour market opportunities available to individuals and families.
- Significant improvements to the preference estimates and labour supply simulation methodology have been made.
- There is consistency between the assessment with and without labour supply changes.
- TaxSim contains a flexible templating tool which allows for easy creation of tables and data visualisation.

One slide on labour supply responses

- Assuming no labour supply responses (or change in hours worked) is heroic - particularly for policies targeted to low-income households.

Household Incomes Report (2004) - page 98

The estimates are that by 2008 the **poverty reduction impact** *of the whole WFF package at a 50% BHC threshold will be* **around 70%**...

- TaxSim estimates labour supply responses to tax-and-transfer policy that change disposable income.
- Some labour demand impacts are captured through the labour supply choices (or opportunities) available to families.

Data

TaxSim - Data

- TaxSim is based on Economic Family Units consisting of an adult principal, possibly a spouse and possibly children. Each year contains around 25,000 adults in 15,000 families.
- TaxSim models labour supply responses based on synthetic Household Labour Force Survey data, a survey which captures family composition, hours and wages over a one week period in the June quarter.
- Benefit receipt is extracted from the HLFS-IDI link.

TaxSim data

- TaxSim uses synthetically-augmented data derived from the Household Labour Force Survey (HLFS) available in the Integrated Data Infrastructure (IDI).

- These data meet all of the conditions in the Microdata Output Guide and have been approved by Stats NZ for release.

Labour supply responses

Why care about labour supply responses?

- Labour supply responses occur when people change their hours worked.
- Labour supply responses occur when a labour market option is relatively more appealing, which may happen even if their disposable income would have stayed the same at their pre-reform hours/wage choice.
- Labour supply responses can support or work against common policy objectives like increasing income support, reducing poverty, or increasing work incentives.

Utility - Desirable properties

Let utility $U = f(h, y)$, where h is hours worked and y is net income. Then we would expect to observe:

- $\frac{\partial U}{\partial y} > 0$ (why?)
- $\frac{\partial U}{\partial h} < 0$ (why?)

What is the policy (or practical) relevance of these criteria?
What about the second derivatives?

Choice set

Overview of labour supply and microsimulation modelling

- Start with an arithmetic microsimulation model to calculate disposable income at a given hours/gross wage combination.
- Calculate net wage at each discrete choice for the base and reform scenario.
- Calculate utility at each of the choices.
- Maximise utility across the choices to estimate labour supply response.

Taxwell-B

As a starting point it is useful to recap what was done in Taxwell-B

- Single adult families could choose to hours worked of 0, 5, 10, …, 50 at a fixed wage rate. Couples could choose any combination of hours worked from 0, 10, 20, …, 50 for the male and 0, 5, 10, …, 50 for the female.
- The trade-off between disposable income and hours worked was captured through a quadratic utility function.
- The labour supply responses were “calibrated" so, for each family, the observed hours worked had the maximum utility of all hours points.
- 100 labour supply simulations were run, taking about 90 minutes for 3,000 families.

Taxsim

- **R**andom **U**tility **R**andom **O**pportunity (RURO) model of job choice - both wages and hours can vary in response to policy change.
- Choices for each family are sampled from other families similar on a range of demographic and socioeconomic characteristics.
- Wages and hours in the pre-reform are the same for the fixed and variable labour supply models. They are not rounded to the nearest 5 hours.
- 1,000 labour supply simulations are run in less than 10 seconds.

Extra TaxSim features

Bootstrap measures

- TaxSim has the ability to calculate bootstrap weights on-the-fly using user-supplied calibration totals.

- Confidence intervals are calculated using Efron's bias-correction technique. Results are suppressed if the bias is too strong relative to the number of bootstrap iterations.

Custom user code execution

- Taxsim can run custom user code in a "safe" environment and return the results (with one tab for each scenario).
- Functions are white-listed for use. Generally can't use non-standard evaluation or anything that interacts with shiny directly, the file system or environments.
- Most functions in the following R packages are available: `base`, `data.table`, `ggplot2` and `scales`.

TaxSim demonstration & TaxSim development

TaxSim Development (as at June 2020)

- **Delayed due to COVID-19.** Combination of lack of access to IDI, delayed December IDI refresh and demands from my day job.
- Next version (intended to be 0.5 available at the end of March 2020) was going to:
  - Add bootstrap weights to the labour supply simulations.
  - Improve the preference equation estimation (by including new variables and removing influential observations).
  - Automatically adjust for labour supply changes between the base year and the model year.
  - Improve the output tables and charts for the labour supply simulations.

Version 0.5 -> 0.4.2

Version 0.5 (released as 0.4.2) is a major update to the labour supply response modelling, including:

- Removing the full take-up assumption for benefits from the labour supply model by allowing benefit receipt to vary with hours and wages.
- Linking the labour supply simulations with the bootstrap weights to capture both sampling error and simulation volatility.
- Automatic 'chaining' of simulations to adjust for changes between the base year and the modelled year.
- Adding extra outputs for the variable labour supply runs.

The real Version 0.5

- Better preference estimates that vary by family. See Mixed Logit Model section for more information.
- Less hidden sheets in the parameter spreadsheet.
- Tables export to Excel workbooks instead of PDF. This also removes the webshot dependency and the associated installation issues.
- Redesigned front page of TaxSim interface and removed right side-menu.
- Bootstrap output is consistent with how the rest of the outputs are displayed.
- Fixes an issue where the labour supply simulations would randomly crash.

Annual maintenance - minimum

- Develop a new synthetic database using HLFS June Quarter in the IDI (available September/October for a given year).
- Estimate preference parameters for the new synthetic database
- Update the base policy parameters, inflators and demographic projections.
- Update TaxSim to address any changes that have been made to the imported R packages (e.g. shinydashboardPlus in version 0.5).
- (Occasionally) update the calculation procedures for new policies (e.g. Best Start, Winter Energy Payment).

Development options

- Move the preference estimation onto AWS (or similar) to access more compute resources. Current laptop is 8C/16T/32GB RAM, AWS offers 96C/192T/200GB RAM.
- Could also move TaxSim from shinyapps.io to AWS.
- TaxSim could be developed to run on the real database in the IDI.
- Improve and share data preparation code. The first stage involves preparing the synthetic HLFS data in the IDI. The second stage uses the IDI output to prepare databases for TaxSim.
- Investigate applying TaxSim to policy changes that are difficult with existing microsimulation capability, for example, social insurance, UBI, universal credit.
- Analysis of hypothetical families to determine the precise factors leading to simulated labour supply changes.

Where to from here?

- Likely extend source code agreements for Government agencies from 30 June 2021 to 30 June 2022.
- The public release will also be delayed. Still working on what this looks like for each of the model components (code, databases, interface).
- Thesis delayed as well. The mixed logit model approach to preference estimation took a lot longer than originally expected.
- Next major version is expected in October 2021. See next slide for details.
- Focus for me over the next 6 months is finishing the thesis.

TaxSim Version 0.6?

- Add the 2021 databases. Maybe add 2020 as well.
- Improve the HLFS data in three ways:
  - Better family construction from HLFS relationship information. This has already been done for another project.
  - Fix two recently identified issues with the HLFS to IDI linking. Essentially, (1) people who skip a quarter are not linked properly, and (2) some people change their date of birth through the survey.
  - The match rate can be improved slightly by using the address data to look for people who have ever lived at the residence with the same year and month of birth (and same gender).
- Work through Microdata Output Guide V5 changes for synthetic data.

RESERVE BANK OF NEW ZEALAND
TE PŪTEA MATUA

VICTORIA UNIVERSITY OF WELLINGTON
TE HERENGA WAKA

Chair in Public Finance

TaxSim NZ - new horizons in tax-and-transfer microsimulation modelling

Christopher Ball
chrispball1@gmail.com

2 June, 2021

Outline

7. Choice Set

8. Estimation

9. Chaining

10. Mixed Logit Model Extensions

Choice Set

Key characteristics

- Each choice/opportunity consists of a wage rate, hours worked and indicators of benefit receipt.
- Choices are sampled from families similar on a range of demographic and socioeconomic variables.
- Hours, wages and benefit receipt can take any value observed in these data, and the empirical correlation between these variables is preserved.
- Couples make a joint choice, although each individual is independently sampled.

What do the choice sets look like?

TaxSim

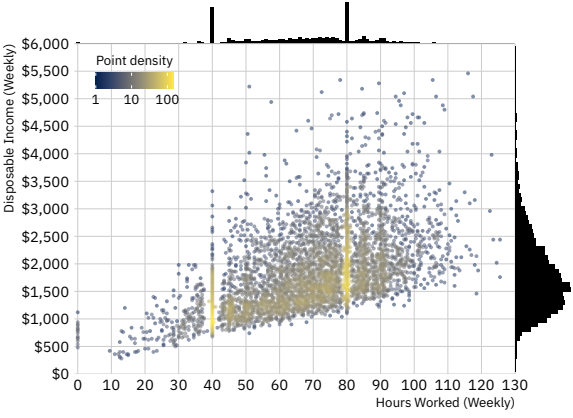

Taxwell-B/Taxmod-B

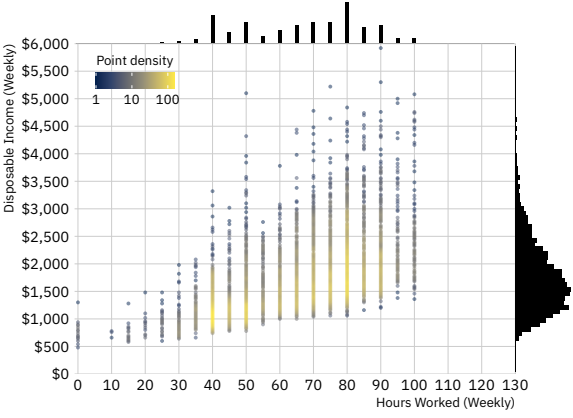

Estimation

Maximum likelihood and discrete choice

- The number of choices and the choice set composition are unobservable.
- Maximum likelihood does not include the number of choices n. This can cause problems if n varies across families.
- Adding an additional choice that is close to the optimal choice lowers the maximum likelihood. This essentially penalises choice sets with “good" alternatives.
- Maximum likelihood can (and often does) lead to parameter estimates with $\frac{\partial U}{\partial y} < 0$ and $\frac{\partial U}{\partial h} > 0$.

An alternative

$$\max \left\{ \sum_{i=1}^{N} \log\left(p_i\right) \right\} \rightarrow \max \left\{ \sum_{i=1}^{N} \sqrt{n_i p_i \left(\mathbb{I}\left[\frac{\partial U}{\partial y} > 0 \right] + \mathbb{I}\left[\frac{\partial U}{\partial h} < 0 \right] \right)} \right\}$$

where

$$n_i = \left(\sum_{i=1}^{n} p_i^2 \right)^{-1}$$

Explanation

- $n_i p_i > 1$ implies better than uniformly picking between the choices.
- $\mathbb{I}(\cdot)$ - only count if the first derivatives make sense.
- $n_i = f(p_i)$ - give more weight to good choices and less weight to bad choices.

Chaining

Chaining?

- Databases containing the observed labour supply outcomes are available at a lag, and the policy changes of interest are usually implemented in the future
- This can lead to significant differences between the observed starting point in the database and the starting point of most interest for future policy, namely the future distribution of labour supply outcomes in the absence of the policy change of interest.
- Preference estimates are assumed to be fixed and exogenous as they are generally slow to change and reasonably exogenous to tax-and-transfer policy changes in New Zealand.

Labour supply comparisons for analysing future policy changes

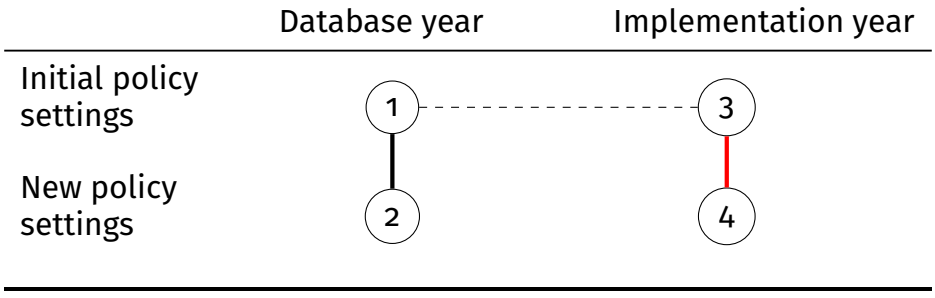

Chaining illustrative example

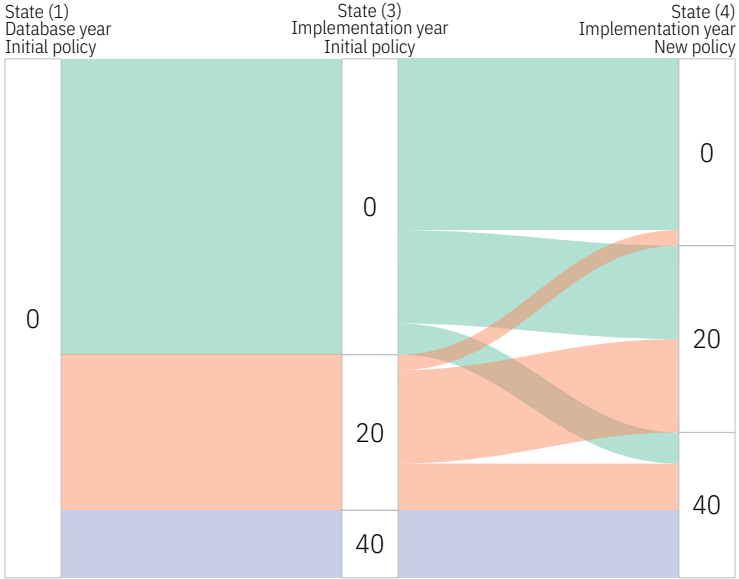

Chaining in TaxSim

- TaxSim implements a two-stage labour supply simulation (i.e. chaining): the first stage going from the database year to the implementation year (absent the policy change of interest) and the second stage comparing with and without the policy change of interest in the implementation year.
- This chaining isolates labour supply changes associated with the policy of interest by removing labour supply changes expected to occur between the database year and the implementation year.
- This adjustment will be more important if there have been large changes in the tax-and-transfer system between the database year and the implementation year, or if there is a large gap between the database year and the implementation year.

Mixed Logit Model Extensions

Preference estimation revisited

Mixed logit models generalise the logit model, typically used for discrete choice models, by allowing the coefficient estimates to vary across decision-makers. The choice probability for the mixed logit model is described in Train (2009) as:

$$p_i = \int \frac{\exp(\beta' x_i)}{\sum_j \exp(\beta' x_i)} f(\beta) d\beta \tag{1}$$

where p_i is the probability of choice i, x_i are individual characteristics and $f(\beta)$ is the density of β. This is essentially a weighted average over the distribution of β of the logit probabilities conditional on β.

Why would we do this?

Start by looking at a specification for the utility function where each of the quadratic terms depends on individual characteristics:

$$U = (\beta_1 * X_1)h + (\beta_2 * X_2)h^2 + (\beta_3 * X_3)y + (\beta_4 * X_4)y^2 + (\beta_5 * X_5)hy$$

Let's see what happens with a comprehensive estimation exercise using every covariate available with these data interacted with all of the quadratic terms. The covariates used are 5 year age bands, region, youngest child age bands, number of children capped at 3 or more, 6 binary ethnic indicators and highest qualification. For couples, the individual level variables such as age and ethnicity are included separately for both males and females.

Couples with children

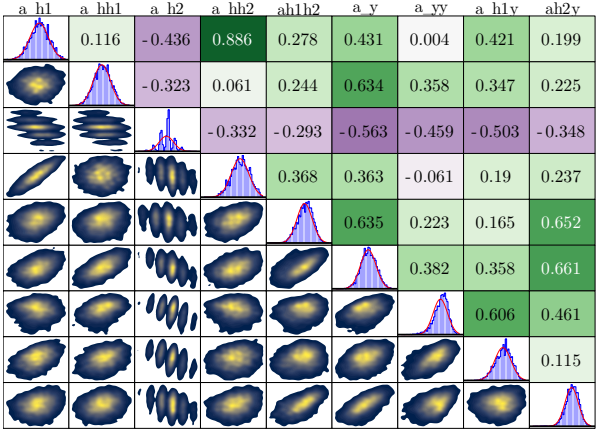

The diagonal shows a histogram of the quadratic term in the header, with the red line showing the probability density function of the closest normal distribution (MLE). The lower triangular of the matrix shows a 2D kernel estimate of the relationship between the two variables, for example the second row and first column shows the relationship between α_h and α_{hh}. The upper triangular of the matrix shows the correlation between the two variables shown in the respective lower triangular cell.

Everything is Normal

- Adding more variables to the estimation tends to make the distributions for each of the preference parameters closer to a Normal distribution.
- The mixed logit model would “average" the probabilities across the parameter distribution.
- The current approach in TaxSim gives a normal distribution to each parameter, and then assigns realisations to individuals/families using a quantile matching technique.
- The preference estimation optimises the normal distribution parameters (mean and variance) for each quadratic preference parameter (10 for singles and 18 for couples). The normal distribution parameters are then passed down to the quantile matching technique to assign realisations to each individual/family for each parameter.

Couples with children - mixed logit

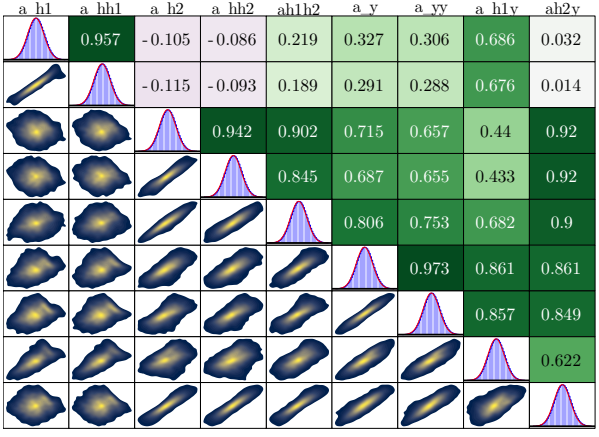

The diagonal shows a histogram of the quadratic term in the header, with the red line showing the probability density function of the closest normal distribution (MLE). The lower triangular of the matrix shows a 2D kernel estimate of the relationship between the two variables, for example the second row and first column shows the relationship between α_h and α_{hh}. The upper triangular of the matrix shows the correlation between the two variables shown in the respective lower triangular cell.

What does this achieve?

- Estimation is much more computationally expensive. Instead of optimising a 5 to 350 variable optimisation, the quantile matching technique has a top-level optimisation of the distribution parameters (10 or 18 variables) with each function calculation requiring a 5 or 9 variable optimisation for each individual or family.
- Labour supply responses tend to be more muted with the new approach.
- The fit (determined by the optimisation function in the estimation) is generally a lot better with this approach. The quantile matching makes determining the degrees of freedom challenging.
- And I haven't even gotten into the work I have been doing on the functional form of the preference function…(don't ask).

Thank You!

Appendix slides

IDI Disclaimer

These results are not official statistics. They have been created for research purposes from the Integrated Data Infrastructure (IDI) and Longitudinal Business Database (LBD) which are carefully managed by Stats NZ. For more information about the IDI and LBD please visit https://www.stats.govt.nz/integrated-data/.

The results are based in part on tax data supplied by Inland Revenue to Stats NZ under the Tax Administration Act 1994 for statistical purposes. Any discussion of data limitations or weaknesses is in the context of using the IDI for statistical purposes, and is not related to the data's ability to support Inland Revenue's core operational requirements.