Discussion of

Nowcasting Global Poverty

by Castaneda Aguilar, Mahler and Newhouse

Natalie Naïri Quinn

Lady Margaret Hall/OPHI
University of Oxford

7 November 2019

Outline

Outline of the Paper

- Documentation of current methods for nowcasting, highlighting three important and contestable assumptions (pass-through, restricted information, flat GIC).
- Data: PovcalNet for poverty estimates, World Economic Outlook for macro outcomes, World Development Indicators for broad development statistics.
- Introduction and discussion of methods: choice of target variable (mean welfare or poverty rate), machine-learning algorithms, approach to missing data, division of observations (training/ testing/ nowcasting), choice of loss function (mean absolute deviation).
- Presentation of results: assessment of status quo, comparison of nowcast errors and current nowcasts under alternative ML algorithms, exploration of greatest divergences, identification of most-informative variables.
- Discussion of potential reasons for insubstantial improvements; exploration of direct prediction of poverty rate, multiple imputation, better prediction of growth in mean and – finally (!) – relaxing assumption of flat GIC.

General Comments

I really enjoyed reading this paper and found it fascinating for several reasons:

- It was good to learn more about the current methods for nowcasting (status quo).
- Interesting to see the application of machine learning algorithms in this context.
- Parallels with some of the challenges faced by my colleagues at OPHI who lead the Global MPI computations: I'm sure they will be very interested in this study.

General Comments

I really enjoyed reading this paper and found it fascinating for several reasons:

- It was good to learn more about the current methods for nowcasting (status quo).
- Interesting to see the application of machine learning algorithms in this context.
- Parallels with some of the challenges faced by my colleagues at OPHI who lead the Global MPI computations: I'm sure they will be very interested in this study.

My biggest comment:

- You saved the best bit for last!
- My immediate concern on learning about the status quo was the flat-GIC assumption.
- I'm not at all surprised that relaxing this (even in a very structured way) improves prediction accuracy.
- It would be great to see this element placed more centrally in the paper.
- (Although, perhaps a conflict with the pragmatic focus of the paper; could this be implementable?)

Comparison of perfomance

- It was surprising that *all* of ML-algorithms nowcasted more poverty than the status quo:
  - Is this because of the apparent systematic bias for long extrapolations with the status quo method?
- It might have been helpful to report a similar figure to the RHS of Figure 3, for the test sample; it's hard to make a status quo/ ML-algorithm comparison without knowing the 'true' poverty rates.

Country-Specific Nowcasts

- I was confused about the comparisons between ML-algorithm and status quo nowcasts for specific countries.
- There seems to be an untested assumption that the ML-algorithm nowcasts are more accurate, in these specific cases.
- Surely that is not necessarily the case?
- It could equally well be that, taking the example of recent-conflict countries, because the ML-algorithms are trained mainly on non-conflict data they systematically underestimate poverty in countries in conflict?
- Perhaps more informative to report and discuss actual/ nowcast comparisons for the *test* sample (both status quo and ML-algorithm nowcasts) rather than status quo/ ML-algorithm comparisons in the nowcast sample.

Identification of Best Predictors

- Interesting (and perhaps not surprising) that growth in employment is such an effective predictor.
- It would be great to explore the mechanism further:
  - Is it a measurement story: that HH surveys better capture the welfare of formal-sector workers?
  - Is it a determinants of poverty story: that formalisation is one of the most effective mechanisms for moving people out of extreme poverty?
  - Maybe a bit of both, or neither...

Predicting Headcount Directly

- The more direct approach seems attractive: surely the best results *should* be obtained through the most direct approach?
- But predicting levels of the poverty headcount discards critical information that the status quo and alternative methods utilise: the (recent) distribution.
- So unsurprising that this yields worse predictions.
- Why not predict the change in headcount? Have you tried this?

Predicting Mean and Gini

- Great to have explored relaxing the flat-GIC assumption
- I was sceptical about the linear-GIC assumption, and still am having taken a look at Lakner et al (2019)...
- Still seems indirect and assumption-heavy: effectively trying to predict the whole distribution in order to recover one centile.
- Why not do that directly? (Suggestion on prediction of *change* in headcount, above)