B. Kreider, J. V. Pepper, C. Gundersen and D. Jolliffe: Identifying the Effects of SNAP (Food Stamps) on Child Health Outcomes When Participation Is Endogenous and Misreported

Klara Röhrl[1], Ximeng Fang[1], eisenhauer[2], Rocío Baeza[2], and Sebastian Becker[1]

[1]University of Bonn
[2]Affiliation not available

February 1, 2018

Short Summary

The authors evaluate the Supplemental Nutrition Assistance Program (SNAP), formerly known as food stamps. They do so by bounding the Average Treatment Effect (ATE) of the program on the population of eligible children. They evaluate SNAP this way with respect to food insecurity, obesity, anemia and subjective health.

The paper has a theoretical and an empirical contribution. The theoretical contribution is that it extends the literature on partial identification bounding methods to account for misreporting (in addition to selection) and derive bounds on the ATE under different sets of assumptions in this setting. The empirical contribution is that the authors apply these bounds to SNAP using the National Health and Nutrition Examination Survey (NHANES) and show that the derived bounds can be informative.

The authors find evidence that SNAP improves the health outcomes of eligible children. For anemia and subjective health the bounds point at substantial health improvements due to SNAP and rule out even small deteriorations in health. For obesity and food insecurity this is only the case after excluding the possibility of false positive reports of SNAP participation.

Contents

1 Integration into the Generalized Roy Model Framework **2**

2 Theoretical Analysis **2**
2.1 Abstracting from Measurement Error . . . 3
2.2 Accounting for Measurement Error . . . 4

3 The Data and Institutional Setting **5**

4 Results **5**
4.1 Food Insecurity . . . 6
4.2 Poor Health . . . 7
4.3 Anemia . . . 7
4.4 Obesity . . . 8

5 Critical Assessment **10**

1 Integration into the Generalized Roy Model Framework

This section classifies which type of policy evaluations this paper performs, how its model relates to the generalized Roy model, what the challenges are and how they are addressed.

The task that the authors set for themselves is the first of the three evaluation tasks set out by Heckman and Vytlacil 2007: They evaluate SNAP in its existing form on the eligible population.

The population of interest are eligible children. The treatment is whether the family participates in SNAP or not which is decided endogenously by the family. The authors start without any assumptions on the process by which treatment is selected or how health is affected by observable characteristics, the treatment or unobservables. Relating the notation of the paper to the notation from the Generalized Roy Model yields:

Potential Outcomes

$Y_1 = H(1) = f_1(X, Z, U)$

$Y_0 = H(0) = f_0(X, Z, U)$

Cost

$C = c(X, Z, U)$

Choice

$S = s(X, Z, U)$

$D = 1\{S > 0\}$

Observed Outcomes

$Y = D \cdot H(1) + (1 - D) \cdot H(0)$

Note, that the authors make less assumptions than we did in the Generalized Roy Model:

Firstly, they do not assume additive separability of the unobserved components U. Secondly, they do not make any assumptions on how households decide whether to participate. In particular, they make no assumption of optimizing behavior regarding that decision.

To avoid having to put more structure on the model, the authors use three tools: Firstly, they choose an aggregate object of interest, namely the ATE. Secondly, they only look at binary outcome variables which simplifies their analysis. Thirdly, they do not aim for point identification but only derive bounds on the ATE.

The authors face the two classical econometric problems of observing every individual in only one of the two treatments (the evaluation problem) and the selection problem that individuals endogenously decide their treatment status. In addition, the authors also face non-classical measurement error because SNAP participation has been shown to be underreported in surveys and the decision whether or not to report true participation is correlated with observable characteristics.(Bollinger and David, 1997)

2 Theoretical Analysis

This section gives an introduction to partial identification bounding methods. Firstly, it explains the general procedure to derive possibly informative bounds. It then applies this procedure to the problem of identifying the ATE of SNAP without non-classical measurement error and introduces several assumptions one might be willing to make. The section then shows how the procedure can be adjusted when allowing for non-random misreporting. It concludes with proposing assumptions on the misreporting that can be used to tighten the bounds.

Partial identification bounding methods use two steps:

Step 1: Starting from the expression of interest, unobserved quantities are first decomposed using the Law of Total Probability to introduce quantities observed in the data and reduce the role of the remaining unobserved quantities.

Step 2: One uses well-founded assumptions that allow to further bound the unobserved quantities that have shown up during the decomposition.

In the following, these steps are applied to the ATE of SNAP on health outcomes.

2.1 Abstracting from Measurement Error

Because of binary outcome variables the formula for the ATE is simplified:

$$\text{ATE} = P(H(1) = 1) - P(H(0) = 1) \tag{1}$$

where H(t)=1 denotes the bad health outcome (later the analysis will look at food insecurity, obesity, anemia and reported subjective health) in treatment t.

Without any data, one can already see that the ATE will lie between -1 and 1. Due to the selection problem neither of the two probabilities can be directly inferred from the data.

However, applying the law of total probability with respect to the treatment status FS^*, one can decompose each of the two probabilities:

$$\begin{aligned} P[H(1) = 1] &= P[H(1) = 1|\text{FS}^* = 1] \cdot P[\text{FS}^* = 1] \\ &\quad + P[H(1) = 1|\text{FS}^* = 0] \cdot P[\text{FS}^* = 0] \\ P[H(0) = 1] &= P[H(0) = 1|\text{FS}^* = 0] \cdot P[\text{FS}^* = 0] \\ &\quad + P[H(0) = 1|\text{FS}^* = 1] \cdot P[\text{FS}^* = 1] \end{aligned} \tag{2}$$

For both decompositions there is one quantity that cannot be estimated by a simple sample analog from the data: The probability of bad health under the treatment that was not chosen $P[H(1) = 1|\text{FS}^* = 0]$ and $P[H(0) = 1|\text{FS}^* = 1]$. As they are probabilities, each must lie in the [0, 1] interval. In addition, each gets multiplied by the probability of the conditioning event. This yields intervals for the two unconditional probabilities of length $P[\text{FS}^* = 0]$ and $P[\text{FS}^* = 1]$, respectively, yielding a total length of the ATE interval of 1. Thus, using the data without any additional assumption we get an interval of length 1 for the ATE. The formulae of its bounds are:

$$\begin{aligned} &P[H(1) = 1|\text{FS}^* = 1] \cdot P[\text{FS}^* = 1] \\ &- (P[H(0) = 1|\text{FS}^* = 0] \cdot P[\text{FS}^* = 0] + P[H(0) = 1|\text{FS}^* = 1] \cdot P[\text{FS}^* = 1]) \\ &\leq \text{ATE} \leq \\ &P[H(1) = 1|\text{FS}^* = 1] \cdot P[\text{FS}^* = 1] + P[H(1) = 1|\text{FS}^* = 0] \cdot P[\text{FS}^* = 0] \\ &- (P[H(0) = 1|\text{FS}^* = 0] \cdot P[\text{FS}^* = 0]) \end{aligned} \tag{3}$$

Next, there are assumptions that the authors consider to further bound the counterfactuals and the unconditional probabilities:

Assumption 1: Monotone Treatment Selection (MTS)

The first assumption puts bounds on the counterfactual conditional probabilities. Applied to our example the MTS assumption states that children that participate in SNAP have weakly worse latent health outcomes than non-participating children:

$$P[H(1)=1|\mathrm{FS}^*=0] \leq P[H(1)=1|\mathrm{FS}^*=1]$$
$$P[H(0)=1|\mathrm{FS}^*=0] \leq P[H(0)=1|\mathrm{FS}^*=1]$$

Assumption 2: Income is a Monoton Instrumental Variable (MIV)

The second assumption is that the probability of bad health decreases weakly with income. Assuming this allows the derivation of additional bounds on P[H(0)=1] and P[H(1)=1].

To do so, decompose P[H(0)=1] and P[H(1)=1] by applying the Law of Total Probability with respect to the instrumental variable:

$$P[H(0)=1] = \int P[H(0)=1|Y=y]f(y)dy$$
$$P[H(1)=1] = \int P[H(1)=1|Y=y]f(y)dy$$

where f(y) is the density of income Y at y.

The conditional probabilities suffer from the same selection problem as the unconditional probabilities but using the MIV assumption one can bound them and in turn also bound P[H(0)=1], P[H(1)=1].

Assumption 3: Monotone Treatment Response (MTR)

The strictest assumption that the authors consider is that SNAP weakly improves health status: $H(1) \leq H(0)$

Using this assumption, the ATE cannot be negative. Thus, to some degree this presumes the result. However, this assumption can interact with other assumptions to yield strictly negative upper bounds. It is also well founded as there is a broad consensus among policy makers and researchers that SNAP does not worsen food insecurity. (Bitler et al., 2003)

2.2 Accounting for Measurement Error

To account for measurement error, we introduce reported treatment status FS in addition to the actual treatment status FS*. Let Z*=1 denote a correct classification.

Allowing for this, none of the quantities in 2 can be directly inferred from the data.

However, note that the event of true participation $\mathrm{FS}^*=1$ can be decomposed into a true positive $\mathrm{FS}^*=1 \wedge \mathrm{Z}^*=1$ and a false negative $\mathrm{FS}^*=1 \wedge \mathrm{Z}^*=0$. Using this decomposition of $\mathrm{FS}^*=t$ each of the eight quantities in 2 can be decomposed by the Law of Total Probability.

The resulting formula for the ATE can be rearranged to get analogous bounds to before:

$$\begin{aligned} &-P[H(0)=1|FS=0]\cdot P[FS=0] - P[H(1)=0|FS=1]\cdot P[FS=1] + \Theta \\ \leq \mathrm{ATE} \leq &P[H(1)=1|FS=1]\cdot P[FS=1] + P[H(0)=0|FS=0]\cdot P[FS=0] + \Theta \end{aligned}$$

with:

$$\Theta \equiv (\theta_1^- - \theta_0^+) - (\theta_0^- - \theta_1^+)$$
$$\theta_j^+ = P(H=j, FS=1, Z^*=0)$$
$$\theta_j^- = P(H=j, FS=0, Z^*=0)$$

Note that these bounds depend only on observable probabilities, unobserved probabilities for which we have developed assumptions earlier and additional unobservable probabilities in Θ.

Table 1: Characteristics of SNAP Eligible Children by Participation

Variable	Recipients	Non-Recipients
Age (in years)	8.6***	9.5
Ratio of income to the poverty line	0.64**	0.86
Food-insecure	0.45**	0.35
Poor or fair health	0.09	0.07
Obese	0.19	0.18
Anemia	0.013	0.010

To get bounds on the θs, the authors use administrative data to estimate the true participation rate among eligible children which they calculate to be 50%. This true participation rate together with the reported participation rate of 46.5% yields three bounds on the θs.

Assumption 4: No False Positives

Additionally, the θs can be bounded further, if one is willing to assume a maximal amount of data corruption. Since SNAP data suffers from underreporting this is equivalent to making an assumption on which rates of false positives one allows. Assuming that there are no false positives is motivated by validity studies (Bollinger and David, 1997) that link surveys to administrative records and find that rates of false positive reports are negligibly small for SNAP.

3 The Data and Institutional Setting

The authors rely mainly on the National Health and Nutrition Examination Survey (NHANES) to evaluate SNAP because of the rich information on children's health. They focus on children from ages 2-17 whose household income is below 130% of its poverty line. These are children that would satisfy the income requirement to be eligible for food stamps. However, there is also a strict wealth requirement, compliance with which the authors cannot verify in their data. As in the theoretical derivation above, SNAP participation is binary, not accounting for differences in the awarded benefits between participating households.

While they do not observe the amount of benefits nor the assets of the household, the NHANES provides them with accurately measured and rich information on the health of the child. In addition to the parents' subjective perception of the child's health, children are measured and weighed by professional nurses and they are tested for anemia.

This yields a sample of 4418 children who were interviewed between 2001 and 2006.

45.6% of these children reported receiving food stamps. There were significant and important differences between participating children and non-participating children as can be seen in the table 1. The comparison shows that recipients are younger, poorer and more food insecure. Nearly one half of all participating households were still food insecure despite the benefits provided by SNAP. Among the non participating eligible households only one in three was classified as food-insecure. Participating children also seem to be in worse health than their non-participating peers but these differences are not statistically significant.

4 Results

This section presents the estimated ATE of SNAP on food insecurity, poor heath, anemia and obesity. For each outcome variable it first reports the estimates ignoring misreporting with varying assumptions and then shows the results allowing for different degrees of misreporting and different assumptions.

4.1 Food Insecurity

The estimated bounds on the ATE for food insecurity assuming there were no measurement error are displayed in figure 1 for the NHANES data. The bound estimates are shown as bars, with whiskers indicating the 95% confidence intervals. One can clearly see the identifying power of the MTS and MIV assumptions in reducing the width of the interval. Maintaining these two assumptions the data show that the effect of SNAP on food insecurity is substantial, reducing food insecurity by at least 12% points and possibly by over 35% points.

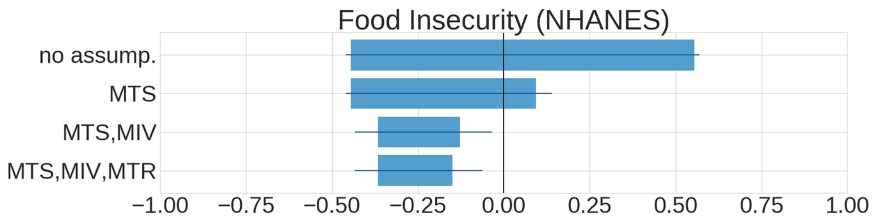

Figure 1: Bounds for the ATE of SNAP on food insecurity in the NHANES data

Since food insecurity is also elicited in the Current Population Survey, the authors estimate the same bounds using this alternative data set (figure 2). The results are qualitatively similar but the intervals are larger such that the MRT assumption is necessary for ruling out that SNAP increases food insecurity on average. These larger intervals could be due to the smaller sample size in the CPS. Another reason could be the large amount of underreporting present in the CPS (nearly 50% according to Meyer et al. (2009)).

Since the CPS allows for the construction of instrumental variables that - if exogenous and valid - could also identify a causal effect of SNAP on food insecurity, the authors construct two instruments that are common in the literature that exploit interstate variability in the implementation of SNAP:

1. About one half of states has simplified reporting requirements
2. About a third of states exempts cars from the asset-test

The authors estimate the ATE from these instrumental variables in two different ways: Firstly, they estimate it using a linear response model. Secondly, they follow Shaikh and Vytlacil 2011 and use the instrumental variables to non-parametrically bound the ATE from above. These four estimates are shown as lines in figure 2, where the thinner lines show the linear response IV estimates and the thicker lines show the non-parametric upper bound estimates. One can clearly see that at least one of the assumptions underlying the estimate of the reporting instrumental variable in the linear response model must be violated since the estimate lies outside the estimate that relies on no assumptions.

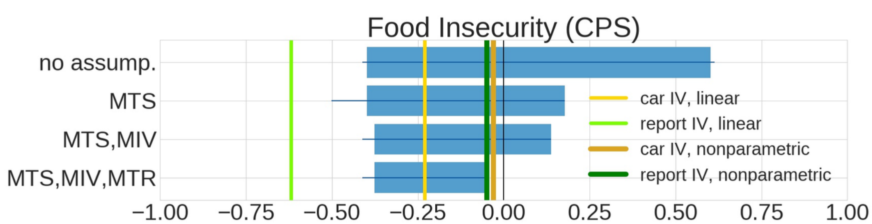

Figure 2: Bounds for the ATE of SNAP on food insecurity in the CPS data

Allowing measurement errors widens the bounds on the ATE substantially as can be seen in figures 3 and 4 for a true participation rate of 50% and 70% respectively. The MIV and MTR assumptions still hold substantive identifying power. Even more identifying power lies now in assuming that there are no false positives - especially for a true participation rate of 50%. Assuming this leaves the bounds only slightly larger than the ones without measurement error. Assuming no false positives is less powerful when the true and reported participation rate are further apart.

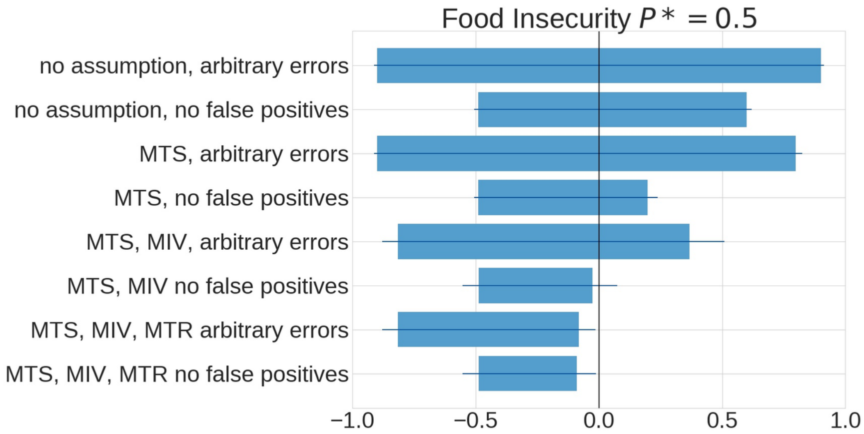

Figure 3: Bounds for the ATE of SNAP on food insecurity in the NHANES data, allowing for misreporting with a true participation rate of 50%

Figure 4: Bounds for the ATE of SNAP on food insecurity in the NHANES data, allowing for misreporting with a true participation rate of 70%

4.2 Poor Health

The estimates for the ATE on subjective health are very similar to those on food insecurity ignoring possible measurement error (Fig. 5). However, the results for subjective health are much more robust to allowing for misclassification (Fig. 6 and 7)[1] and the assumption of no false positives has mostly no noticeable identifying power for this outcome.

4.3 Anemia

The effect of SNAP on anemia is nearly indistinguishable from that on food insecurity when we abstract from measurement error. After accounting for misclassification the bounds are only marginally wider and are basically unaffected by the amount of data corruption one allows. With a higher true participation rate the interval includes much more negative values, including that food stamps reduce the incidence of anemia by 50% points on average.

[1] for subjective health, anemia and obesity only the results under the two strongest assumptions were published

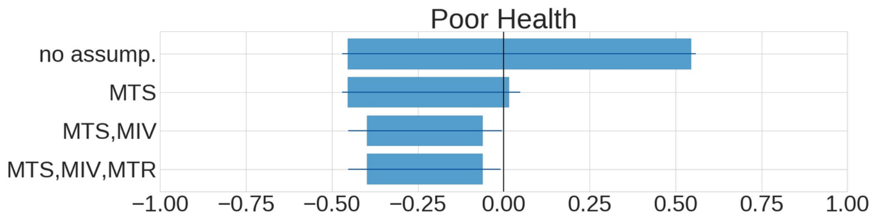

Figure 5: Bounds for the ATE of SNAP on subjective child health in the NHANES data

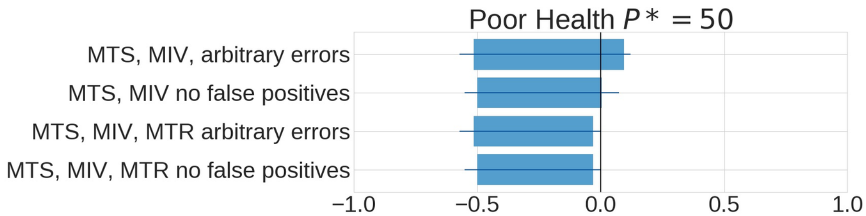

Figure 6: Bounds for the ATE of SNAP on subjective child health in the NHANES data, allowing for misreporting with a true participation rate of 50%

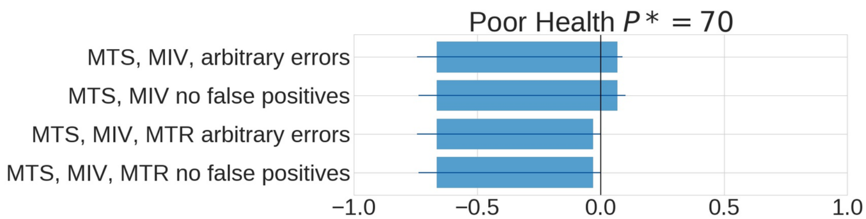

Figure 7: Bounds for the ATE of SNAP on subjective child health in the NHANES data, allowing for misreporting with a true participation rate of 70%

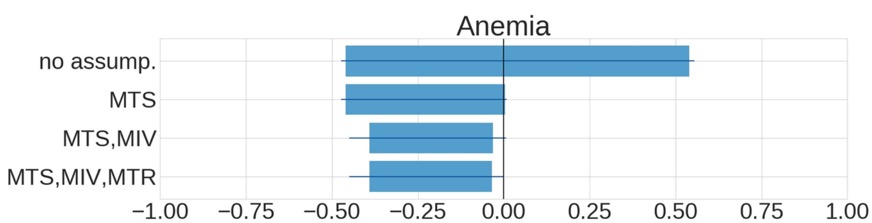

Figure 8: Bounds for the ATE of SNAP on anemia in the NHANES data

4.4 Obesity

Assuming measurement errors away, the estimates suggest that SNAP is as good at reducing the incidence of obesity as it is in reducing the incidence of anemia. Even without assuming a monotone treatment response - which one might not wish to uphold for obesity - the effects of SNAP on obesity appear to be strictly negative (although very small positive effects cannot be rejected at 95%).

However, allowing for measurement error changes these estimates. The intervals grow much larger[2] such that effects

[2]Keeping in mind that only the results when maintaining the MTS and MIV assumptions are reported when accounting for

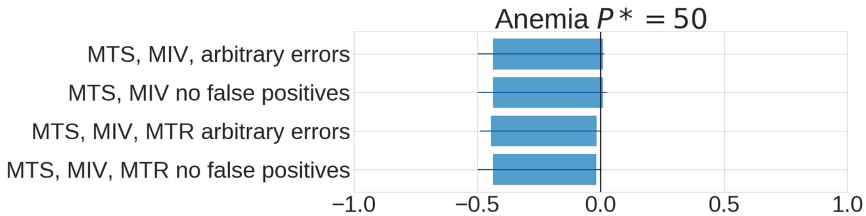

Figure 9: Bounds for the ATE of SNAP on anemia in the NHANES data, allowing for misreporting with a true participation rate of 50%

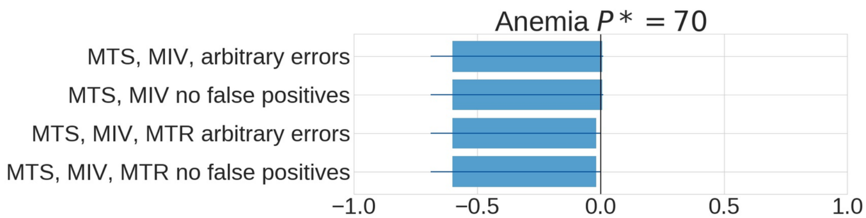

Figure 10: Bounds for the ATE of SNAP on anemia in the NHANES data, allowing for misreporting with a true participation rate of 70%

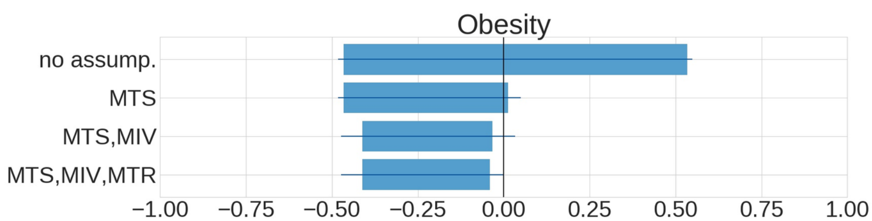

Figure 11: Bounds for the ATE of SNAP on obesity in the NHANES data

of 6% to over 30% points increases in obesity due to SNAP cannot be rejected depending on which assumptions one wishes to maintain. In contrast to the results for anemia and subjective health, assuming no false positives has noticeable identifying power but not of the same magnitude as for food insecurity.

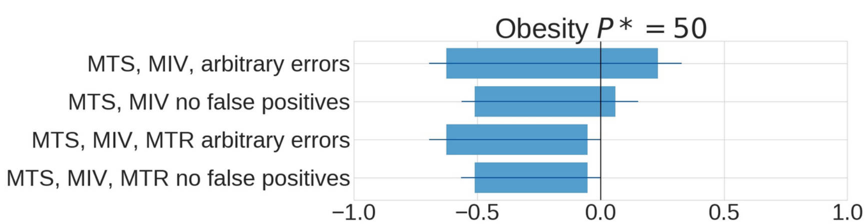

Figure 12: Bounds for the ATE of SNAP on obesity in the NHANES data, allowing for misreporting with a true participation rate of 50%

misreporting

Figure 13: Bounds for the ATE of SNAP on anemia in the NHANES data, allowing for misreporting with a true participation rate of 70%

5 Critical Assessment

Due to the partial identification bounding method, the authors' identification and estimation of the ATE is very clean. The assumptions used to tighten the bounds further are well-motivated. Furthermore, using partial identification allows them to calculate consistent bounds on the ATE that don't require any assumptions, using only frequencies in their data and properties of probabilities in this setting. Such a consistent, assumption-free estimator that is point identified does not exist for observational data.[3] These assumption-free bounds make it possible to transparently show the identifying power of every assumption by reporting how the bounds tighten with each assumption.

However, this approach also has weaknesses:

Firstly, the approach only allows them to bound the ATE. However, it is questionable how interesting the ATE of the program in its current form is to policy makers: The ATE measures the average difference in health between all eligible households receiving food stamps and the abolition of the program. However, even among Republicans 66% support raises in food stamp benefits. (Clement 2017). Thus, the abolition of the program does not seem an interesting quantity to contrast. Instead, policy makers are likely more interested in the effects of adjustments to the program rules, such as changes to the award rules or to the reporting requirements.

Secondly, the approach forces the authors to code all variables as binary. This forces them to give up a lot of important information. This introduces two important weaknesses:

On the one hand, it makes them understate the effect of SNAP. To see why, consider the food insecurity variable: Families answered 18 questions regarding their food insecurity. By coding this variable as binary, they can only report the effect of SNAP on the extensive margin (whether families are food insecure or not). However, all improvements in the intensity of families' food insecurity are disregarded when they do not move families into being food secure.

On the other hand, coding program participation as binary forces them to ignore large differences between participants in the amount of awarded food stamps. To get a feeling for the differences in benefits consider a two person household in 2016: the smallest possible monthly benefit was 16 USD while the largest was above 350 USD (Lauffer 2017, p.103). Using this (non-random) variation in treatment intensity could allow a researcher to derive an ATE per dollar of monthly benefits under reasonable assumptions. With more assumptions it might even be possible to estimate a distribution of marginal treatment effects that could be used to inform the debate on whether and which increases in SNAP benefits are effective in improving children's health outcomes.

Regarding the empirical results, it would have been nice if the authors had included two more sets of results:

Firstly, the CPS data was only used for the analysis without measurement error. Given the substantial underreporting in the CPS (C. and M., 1997) it would have been interesting to see how the bounds they derived to account for misreporting fare in the presence of such a large degree of underreporting.

[3] If we had a randomized control trial, we would be able to get a point-identified consistent estimator of the ATE. However, point-identified estimators in observational data rely on exogeneity assumptions that can rarely be tested. (For example the independence of errors and the outcome variable conditioning on the regressors in OLS.) There are instances where some of these assumptions can be tested (e.g. in the case of over-identified GMM) but never all - as is the case with the bounds developed here.

Secondly, the results including measurement error for obesity were only shown with the two strictest assumptions. However, it is exactly obesity where the MTS[4] and MIV[5] assumptions are the least innocuous. Therefore, they should have also included the results without these two assumptions in the paper.

[4]Children that receive food stamps are less often obese than children who don't

[5]The poorer the child's household the more likely that the child is obese.

References

Marianne P. Bitler, Janet Currie, and John Karl Scholz. WIC Eligibility and Participation. *The Journal of Human Resources*, 38:1139, 2003. doi: 10.2307/3558984. URL https://doi.org/10.2307%2F3558984.

Christopher R. Bollinger and Martin H. David. Modeling Discrete Choice with Response Error: Food Stamp Participation. *Journal of the American Statistical Association*, 92(439):827–835, sep 1997. doi: 10.1080/01621459.1997.10474038. URL https://doi.org/10.1080%2F01621459.1997.10474038.

Bollinger C. and David M. Modeling Discrete Choice With Response Error: Food Stamp Participation. *Journal of the American Statistical Association*, 92:827, 1997.

Scott Clement. Hard-working taxpayers don't suppotr big cuts to food stamps, it turns out, May 2017. URL https://www.washingtonpost.com/news/the-fix/wp/2017/05/25/hard-working-taxpayers-dont-support-big-cuts-to-food-stamps-it-turns-out/?utm_term=.01e90c167132.

James J. Heckman and Edward J. Vytlacil. Chapter 70 Econometric Evaluation of Social Programs Part I: Causal Models, Structural Models and Econometric Policy Evaluation. In *Handbook of Econometrics*, pages 4779–4874. Elsevier, 2007. doi: 10.1016/s1573-4412(07)06070-9. URL https://doi.org/10.1016%2Fs1573-4412%2807%2906070-9.

Sarah Lauffer. Characteristics of Supplemental Nutrition Assistance Program Households: Fiscal Year 2016. Technical Report SNAP-17-CHAR, United States Department of Agriculture, November 2017.

Bruce Meyer, Wallace K. Mok, and James Sullivan. The Under-Reporting of Transfers in Household Surveys: Its Nature and Consequences. Technical report, jul 2009. URL https://doi.org/10.3386%2Fw15181.

Azeem M. Shaikh and Edward J. Vytlacil. Partial Identification in Triangular Systems of Equations With Binary Dependent Variables. *Econometrica*, 79(3):949–955, 2011. doi: 10.3982/ecta9082. URL https://doi.org/10.3982%2Fecta9082.