

WP 2016-11
July 2016



Working Paper

Charles H. Dyson School of Applied Economics and Management
Cornell University, Ithaca, New York 14853-7801 USA

INTRA-HOUSEHOLD INEQUALITY AND OVERALL INEQUALITY

Ravi Kanbur

It is the Policy of Cornell University actively to support equality of educational and employment opportunity. No person shall be denied admission to any educational program or activity or be denied employment on the basis of any legally prohibited discrimination involving, but not limited to, such factors as race, color, creed, religion, national or ethnic origin, sex, age or handicap. The University is committed to the maintenance of affirmative action programs which will assure the continuation of such equality of opportunity.

Intra-Household Inequality and Overall Inequality*

By

Ravi Kanbur

www.kanbur.dyson.cornell.edu

This version: 31 May, 2016

Contents

1. Introduction
 2. Some Basic Analytics
 3. How Much Intra-Household Inequality?
 4. Exercise 1: Growth, Poverty Reduction and Inequality
 5. Exercise 2: Minimum Wage and Poverty
 6. Exercise 3: Targeting of Public Programs
 7. Conclusion
- References

Abstract

Assessing the specific contribution of intra-household inequality to standard measures of overall inequality and poverty is an underdeveloped area of research and policy analysis. However, intra-household inequality is clearly important. It is argued in this paper that neglecting intra-household inequality could lead to (i) an understatement of inequality and an overstatement of the impact of growth on poverty reduction; (ii) a mis-statement of the potential impact of minimum wage policies on poverty; (iii) mis-design of transfer policies to reduce inequality and poverty. Any discussion of inequality in the welfare state cannot afford to ignore intra-household inequality.

* Paper based on presentation to the Conference on The Welfare State and the Fight Against Inequality, Columbia University, 8-9 November, 2015.

1. Introduction

The discourse on inequality in the welfare state and more generally has two strands. One strand focuses on inequality between individuals, and on the poverty of individuals. This is perhaps the fundamental orientation. The standard measures of inequality and poverty are all built up from information on wellbeing at the individual level. But a second strand looks at inequality across broadly defined salient groups. For example, total inequality among individuals is often decomposed into a “between group” and a “within group” component. In this second strand, the grouping is sometimes envisaged as merely instrumental for policy, providing an instrument of intervention; the objective is still inequality between individuals. At other times, however, the group is seen as having normative significance. Now the inequality decomposition has ethical significance over and above the instrumental.

Consider decomposition of inequality across groupings defined by race, gender, ethnicity or caste. Such a decomposition can be used in myriad ways (Kanbur, 2006). In terms of positive analysis, a decomposition is the start of a causal analysis, identifying the key determinants of variation in income among individuals. In effect, it is the non-parametric analog of running a parametric regression with dummy variables for the different groupings, and interaction terms across the groupings.

However, the same decomposition can be turned to normative use. Thus one can develop rules for intervention across groups even if the objective is overall inequality or poverty, if fine differentiation of individuals within a group is very costly. In effect, then, one is using the group identifier as a targeting device (Kanbur and Tuomala, 2016). But one can go further. Following Roemer’s (1998) framework, and as implemented by Paes de Barros et. al. (2009) one can imbue the inequality accounted for by differences across groups with an ethical significance—it can be identified with “inequality of opportunity.”

One grouping which is not usually thought of as such is grouping by household. Each household is a group unto itself, and inequality (and poverty) can be decomposed across these “groups”. There are of course many more groups now than in the case of gender, say, but the basic analytical structure remains the same. We can develop between-group and within-group decompositions, and we can discuss the consequences of policy across and within groups. Policy, in turn, can treat the household grouping as purely instrumental, the ultimate objective being inequality or poverty across individuals, or it can indeed imbue the household per se with special normative significance.

It is this relationship between inequality within and between households, and overall inequality, which is explored in this paper. I hope to show that this approach provides interesting perspectives on inequality and policy, and indeed raises some difficult questions for policy. Section 2 begins with some analytical preliminaries on inequality and poverty decomposition. Section 3 asks how much we know about within-household and between-household inequality. Section 4 is the first of three sections which conduct exercises to illustrate the significance of the intra-household inequality issue. It considers the relationship between growth and poverty reduction through the intra-household lens. Section 5 does the same for the debate on minimum wages, and Section 6 takes up targeting of public expenditure for poverty reduction. Section 7 concludes.

2. Some Basic Analytics

Before proceeding to empirical and policy discussion, I begin with some basic analytics which will be referred to in the rest of the paper. Let there be n individuals each of whom can be allocated to G mutually exclusive and exhaustive groups. Let income or consumption be denoted y_{ig} for individual $i = 1, 2, \dots, n_g$ in group $g = 1, 2, \dots, G$. The population share of group g is written x_g , where the shares obviously sum to 1, and the mean income of group g is written m_g . Overall mean income for the whole population is m . Let overall inequality, a function of all n incomes, be written I . Consider now the case where each individual in a group is given the mean income of that group. We thus have an income distribution with G mass points with income and populations shares given by the G pairs (m_g, x_g) for $g = 1, 2, \dots, G$. This is a distribution in which there is no inequality within each group. What is left is variation due to differences in mean income group across groups. Denote the inequality of this distribution as I_B , the between-group component of overall inequality. Thus the difference between overall inequality and I_B can be called the within-group component of overall inequality, denoted $I_W = I - I_B$.

Obviously, between-group inequality is less than overall inequality. In fact, the distribution where inequality within each group is suppressed by giving each person in group the mean income of that group, is a mean preserving contraction of the overall distribution. Put another way the overall distribution is a mean preserving spread of the distribution which only reflects between group variation. For specific inequality indices, expressions can be derived and amounts calculated for I_W and I_B . A commonly used measure of inequality is the Mean Log Deviation (MLD), which we denote by L . Let MLD within group g be L_g . In this case,

$$\begin{aligned} L &= (1/n) \sum \sum \log(y_{ij}/m) \\ &= L_W + L_B \\ &= \sum x_g L_g + \sum x_g \log(m_g/m) \end{aligned} \quad (1)$$

Thus the within group component of the MLD is simply a weighted sum of the group MLDs.

For poverty, a group decomposition can also be carried out for the FGT class of poverty indices P_α (Foster, Greer Thorbecke, 1984). Let the poverty line be z and the poverty in group g be written $P_{\alpha,g}$

$$\begin{aligned} P_\alpha &= (1/n) \sum \sum [(z-y_{ij})/z]^\alpha \quad \text{for } z \geq y_{ij} \\ &= \sum x_g P_{\alpha,g} \end{aligned} \quad (2)$$

Thus national poverty is a weighted sum of group poverty, the weights being group population shares. The decomposition holds for all values of α . Of particular interest are the cases of $\alpha = 0$, which is the head count ratio measure of poverty; $\alpha = 1$, which is the income gap measure; and $\alpha = 2$, which is the squared gap measure of poverty.

$$\begin{aligned} \text{Notice that } P_\alpha &\text{ is nothing other than the expectation of the function} \\ h(y) &= [(z-y_{ij})/z]^\alpha \quad \text{for } z \geq y_{ij}; \quad 0 \text{ otherwise} \end{aligned} \quad (3)$$

For $\alpha \geq 1$, $h(y)$ is a convex function of y . But since the distribution with inequality within each group suppressed is a mean preserving contraction of the overall distribution, not only inequality but poverty in the overall distribution will be greater than in the between group distribution (where every member of a group is given the mean income of that group). When $\alpha = 0$, the function is neither convex nor concave, so whether poverty is overstated or understated in this way will depend on the detail of the distribution and the location of the poverty line.

The groupings used in the expressions above can be anything, provided they are mutually exclusive and exhaustive. As noted in the introduction, the use of groupings by gender, ethnicity, place of birth etc. is quite common in the literature. In this paper, however, we focus on each household as a group and consider the inequality and poverty consequences of inequality within the household.

3. How Much Intra-Household Inequality?

A large literature has developed in the last three decades on intra-household allocation of resources. This literature has moved in a mutually supportive parallel fashion with a corresponding literature on alternative models of household decision making. The dominant model of household decision making was and to some extent still is the unitary model, where household behavior is represented by the maximization of a single utility function. The household acts as though it had a single budget constraint. Observed variations in consumption across members of the household are then attributed to the budget constraint and the household utility function. Alternatives to the unitary approach include non-cooperative models, and cooperative models where individuals cooperate but bargain over the gains from cooperation. A recent review of this literature is provided in Browning, Chiappori and Weiss (2014).

In all of the above models, there can be inequality within the household. In the unitary model this can be attributed to the degree of inequality aversion in the household utility function. In cooperative models, on the other hand, inequality can be attributed to differences in the fall back options in a Nash bargaining game. And there are many variations on these themes. But a central question is the empirical one of how much intra-household inequality there actually is in a society, and how do we find this out?

Inequality is the variation of some valued attribute across individuals, as in (1). Thus we need to be able to measure this attribute across all individuals, including within a household. The standard instrument, for collecting such information is the household survey. As the name suggests, the survey is administered to households, who have been selected through appropriate stratified random sampling methods, usually drawn from census lists. The survey questionnaire consists of various modules, including those which collect information on individuals with questions on demographic attributes, health, education, anthropometric variables and so on. Some of these variables can then be analyzed for inequality.

A leading example of the use of individual level body mass index (BMI) to assess intra-household inequality is Sahn and Younger (2009, p. S13):

“This paper examines the relationship between level of well-being and inequality at inter-country and intra-household levels, using individuals’ body mass index (BMI) rather than

income as the indicator of well-being. BMI is useful for these purposes because (1) it is measured at the individual rather than household level; (2) it reflects command over food, but also non-food resources that affect health status like sanitary conditions and labour saving technologies; (3) it accounts for caloric consumption relative to needs; (4) it is easily measured; and (5) any measurement error is likely to be random....Perhaps the most striking finding in the paper is that about half of total BMI inequality at the country level is within households. Thus, standard measures of inequality that use household-level data may drastically understate true inequality.” (p. S13).

For the mean log deviation measure of inequality in (1), Sahn and Younger (2009) find that, for 14 household surveys across 7 countries, intra-household inequality accounts for between 55% and 66% of total inequality: “If we did not take into account the ‘within’ household component of total inequality in the seven countries examined, country inequality would have been reduced by more than half, and as much as two-thirds in the case of Cote d’Ivoire.” (p. S32).

A similar point was made in an earlier paper by Haddad and Kanbur (1990), using data on individual level food intake from a small purposely designed household survey in the Philippines. Here individual food intake was first converted to calories and then to a calorie adequacy ratio using individual-specific calorie requirements for the Philippines for different demographic groups. A synthetic distribution was then constructed in which the total calorie intake of a household was allocated pro rata per requirements within the household. This is, in fact, the distribution we would have if we only had food intake information at the household level, which is the case for most national level household surveys. Haddad and Kanbur (1990) then compared actual inequality and poverty (using as poverty line a calorie adequacy ratio of one) with the inequality and poverty in the synthetic distribution:

“Our theoretical analysis suggested that potentially serious errors could be made so far as the *levels* of inequality and poverty are concerned. Empirically, we showed that this is indeed the case – the errors are of the order of 30% or more.” (Haddad and Kanbur, 1990, p. 879).

A more recent approach to measurement of intra-household inequality is to estimate a collective model of intra-household resource allocation, to recover the relevant sharing parameters from the model, and thus get at intra-household inequality in consumption. Thus Lise and Seitz (2011, p. 352) conclude their investigation into UK data as follows:

“...the large dispersion in incomes within the household are highly inconsistent with equal division of consumption. When equal division is relaxed, we estimate that there is substantial inequality within the household. This suggests previous work underestimates the level of individual consumption inequality by between 25% and 50%.”

Overall, then, a picture emerges of significant understatement of inequality (and poverty) when intra-household inequality is ignored. But our standard household survey based measures of consumption inequality and poverty do precisely that. All of the above are special exercises—using either an anthropometric measure of wellbeing, a detailed survey on a small number of households on individual food intake, or a full estimation of a household allocation model.

However, our headline measures of inequality and poverty, especially in developing countries, are not the product of any of these types of exercises. The consumption or expenditure data they collect is at the household level only, and conversion to the individual level essentially involves dividing the household total by the number of (sometime equivalized and sometimes not) individuals in the household. Intra-household inequality is thus suppressed, and overall inequality and poverty is understated by a significant but unknown amount.

4. Exercise 1: Growth, Poverty Reduction and Inequality

How far wrong do we go by ignoring intra-household inequality in our standard national level analysis of inequality and poverty? As noted above, in terms of levels we can be off by at least 25%, and perhaps as much as 65%. But there is another error we can commit in a dynamic setting, which perhaps underlies a disconnect between the trends of official data on poverty and public sentiments on the perceived trend of poverty.

In the early 1990s I was the Director of the World Bank's field office in Ghana, having previously worked on and analyzed poverty and poverty trends in Ghana through the Ghana Living Standards Measurement Survey (Boateng, Ewusi, Kanbur and McKay, 1992). Throughout the late 1980s and the 1990s, official data showed that Ghana had seen poverty reduction since the start of the structural adjustment program in the mid-1980s. Yet this was not the perception on the ground, and I can attest to this, having presented the official statistical view multiple times in multiple venues in the country.

This led me to think about a range of reasons for the disconnect, which I have set out in Kanbur (2001). These included, for example, the practice in official data and in economics more generally, to present the number of poor divided by total population—the head count ratio, or the poverty rate or poverty incidence—rather than the absolute number of poor. I argued that the particularly axiom in Sen (1976) which led to normalization by population size was not necessarily appropriate. In high population growth societies, the absolute number of poor could increase at the same time as the poverty incidence fell. Thus a civil society organization dealing with the poor would see an increase in its activities even as officials pronounced that poverty was falling! I argued for presentation of the absolute number of poor alongside the poverty rate, which is now much more common, and developed analysis of poverty measurement without the relevant Sen axiom (Chakravarty, Kanbur and Mukherjee, 2006).

But there was one reason for the disconnect which I did not fully appreciate at that time, and which I set out much later in Kanbur (2010). And this reason is related to a well-known construct in the literature—the growth elasticity of poverty reduction (Bourguignon, 2003). This is the percentage change in poverty for each percent increase in mean income, this increase happening in distribution neutral terms (i.e. each income increasing at the same rate as the mean). This distribution neutral impact of growth on poverty is one component of decomposing poverty change into a part attributable to growth (change in mean) and a part attributable to distributional change. But what is interesting is that the (distribution neutral) growth elasticity of poverty reduction depends on the level of inequality in the distribution. For the head count ratio, and if the distribution is lognormal, then the elasticity ε can be written as (Bourguignon, 2003):

$$\varepsilon = (1/\sigma) \lambda \{ [\log(z/m)]/\sigma + \sigma/2 \}$$

where z is the poverty line, m is mean income, σ is the standard deviation of log income, and $\lambda\{\}$ is the hazard rate of the standard normal distribution. Thus the growth elasticity is seen to be an increasing function of inequality measured by σ (and a decreasing function of mean income m). Bourguignon (2003) derives similar expressions for other measures of poverty.

The expression above has been an entry point in the literature to discussions of why the impact of growth on poverty reduction varies across countries, which it does empirically. One explanation is that, as shown in the expression above, the poverty reduction impact of distribution neutral growth depends on the level of inequality, because greater inequality means that the incremental fruits of growth flow less to the lower tail of the distribution. The expression above formalizes the intuition for a particular distributional structure. Bourguignon's (2003) Figure 3 gives a feel for the quantitative magnitudes involved: "...the growth elasticity is around 3 if inequality is low - i.e. a Gini coefficient around .3 - but it is only 2 if the Gini coefficient is around the more common value of .4." A large empirical literature has thus grown up linking inequality to poverty reduction in this way.

But now consider the following argument. We know that the official method of calculating poverty using household surveys understates the level of poverty. In a period in which poverty is falling, it will understate its level at the start and at the end of the period. But the official method also understates the level of inequality by at least a factor of a third, which means that it overstates the growth elasticity of poverty reduction by a factor of over a half, from the expression above and Bourguignon's (2003) Figure 3. The variation in σ is not now one across countries, but one between the official distribution and the actual distribution of consumption. Thus the beneficial effects of growth for poverty reduction is overstated by our standard methods and our standard data by as much as 50%.

It is, of course, too much to ask that standard national level household surveys collect individual level consumption data and provide us with the "true" distribution of consumption and from this the actual inequality and poverty as our headline statistics. There are conceptual and empirical issues associated with allocating the consumption value of household public goods—how are we to discern the benefits of a bicycle to each different member of the household? This is in principle doable with a more detailed survey module, but outside the realm of practical feasibility. Even for private consumption items like food where there are no conceptual issues, it would require a much thicker questionnaire to get at individual level food consumption. It would seem that, for developing countries at least, we will have to make do with standard household surveys where consumption and expenditure data is collected at the household level.

What is one to do? The answer is two-fold. First, we should go on doing specialized analyses like the ones discussed at the start of this section, involving non-consumption measures of wellbeing which are easily collected at the individual level, or small scale specialized surveys which collect information on individual level consumption, or estimation of structural models of intra-household allocation. These will continue to update the broad quantitative magnitude of the errors we make when we ignore intra-household inequality in estimating overall inequality and poverty. Second, however, we should carry around us with us a rule of thumb correction for official inequality and poverty measures. We have seen that the error can range from one thirds

to two thirds. We might, therefore, start by using a minimal correction factor of one third for any official inequality and poverty measure to take account of intra-household inequality.

5. Exercise 2: Minimum Wages and Poverty

On the face of it, the discourse on intra-household inequality and that on minimum wages seem to have no connection to each other. The minimum wage discourse is associated with labor markets, the elasticity of labor demand, whether the labor market is competitive or monopsonistic, enforcement of minimum wages, and so on (see, for example, Basu, Chau and Kanbur, 2010; Borat, Kanbur and Mayet, 2012). Typically, protagonists divide into camps which either support minimum wages because of their potential impact on the wages of the working poor, or oppose them because of their possible effects in raising unemployment. The first group argues for the poverty reducing impact of minimum wages; the second group of course highlights the poverty increasing consequences of higher unemployment.

But workers, whether employed or unemployed, do not in general live in isolation. Thus a wage (or lack of it) does not necessarily translate directly into individual consumption. When workers live in households, it is the intra-household allocation process which determines consumption outcomes for them and for non-earners within the household. Intra-household inequality thus plays a key role in determining the inequality and poverty impact of an increase in the minimum wage. The worker composition of households has been used in different ways in the minimum wage debate. In the US, for example, it has been argued that those who would benefit from a minimum wage tend not to live in poverty households (Burkhauser, 2015). Analysis for South Africa shows complex patterns of employment and unemployment (Bhorat et al., 2016, pp 39-40):

“Only 16 percent of working age adults in quintile 1 are employed, while in quintile 5 this percentage rises to 75 percent. This suggests that high rates of unemployment and low rates of employment are at the core of the high levels of observed household poverty levels.It is additionally clear that among poor households, the number of people reliant on each employed person is high.”

The employment composition patterns across households have been prominent in policy debates in South Africa. The first part of the findings above is used in the usual way to suggest that raising the minimum wage will not necessarily have a big impact on poverty since it is unemployment which is a key characteristic of the poor, not employment, and introducing or raising minimum wages may increase unemployment (whether it does so or not is an empirical question; see Borat, Kanbur and Mayet, 2013). But the second part has been used in policy debates, for example by trade union organizations like COSATU in South Africa, to make the following informal argument. Poor households depend on the earnings of a small number of workers. Thus raising the wage of the employed through minimum wage regulation will reduce poverty.

The second of these arguments can be developed as follows in intuitive fashion. Take a competitive labor market with homogeneous labor which initially clears at some low wage. A minimum wage is now introduced which raises the wage but lowers employment and creates unemployment. Suppose that all workers in the country form a giant national household, and that

all income is equally shared between employed and unemployed. For simplicity think of the case where these are the only individuals in society. Then it should be clear that the consumption of each individual goes up when the total wage bill goes up, irrespective of the employment and unemployment levels associated with any wage rate. Further, if the elasticity of labor demand is less than one, which is empirically the case almost universally (in South Africa the conventional number referred to in debates is 0.7), then as the wage rises so does the total wage bill. In this case, therefore, raising the minimum wage will raise consumption of all workers and thus tend to reduce poverty if the previous consumption level was below the poverty line.

Of course if we take the other extreme and assume no income sharing at all between employed and unemployed, then introducing or raising the minimum wage will increase the number of people at very low consumption, and this is to be balanced against the higher consumption of the now smaller number of employed. The net effect on poverty will depend where exactly the poverty line lies in relation to the consumption of the unemployed and that of the employed. If it lies between the two, this will rationalize the view taken on one side of the debate, that higher minimum wages will increase poverty. But the central point is that the degree of income sharing between employed and unemployed turns out to be central to the debate.

The above point is developed formally in Fields and Kanbur (2006), using the FGT class of poverty indices in (2). For the case of the “national household”, clearly the poverty impact of a minimum wage ranges from adverse when there is no income sharing to beneficial when there is perfect income sharing. The analysis is then extended to the case where there is not a national household, but the nation is made up of two person households. Each member of a household can be employed or unemployed with probability given by the national employment and unemployment rates, e and $1-e$ respectively. There are thus three types of households. A fraction e^2 of individuals are in households where both are employed; a fraction $(1-e)^2$ of individuals are in households where both are unemployed; and a fraction $2e(1-e)$ of individuals are in households where one person is employed and the other is unemployed. Given an income sharing rule for members of the household, we can thus derive a national consumption distribution and, given a poverty line, the poverty index. Changes in poverty can then be tracked as the minimum wage changes and the associated national employment rate changes with it.

The overall conclusion of the Fields and Kanbur (2006) analysis is as follows:

“The results of this paper lead us to a more nuanced view about minimum wages than is commonly found in the literature. A higher minimum wage does not necessarily increase poverty because of the unemployment it creates. On the other hand, a higher minimum wage does not necessarily reduce poverty simply because it might increase total labor income, and some of this increased income is shared with the unemployed, either through family sharing or through social sharing. Thus the “standard labor economist’s view” and the “standard trade unionist’s view” are both simplistic. Not only does the truth lie somewhere in between, but it can be characterized precisely in terms of empirically observable parameters.”

And a key parameter is the extent of income sharing within the household.

6. Exercise 3: Targeting of Public Programs

Intra-household issues are quite prominent in the design of transfer programs. The old argument in the UK between whether child benefit should be given through the tax system (essentially through the male paycheck) or through the benefit system (say picked up at the post office by the mother) reflects these concerns. This is the “wallet versus purse” choice faced by policy implementers. In the developing country context, and also in developed countries, the intra-household adjustment following on from nutritional supplements given at school are of great concern. What good is an extra glass of milk at school if the child gets a glass of milk less at home as a result? There is a large literature on intra-household “flypaper effects”—to what extent does the transfer given to an individual “stick” to that individual once intra-household (re)allocation processes have had their sway? Early contributions to the literature include Haddad and Kanbur (1992, 1993). Just two recent papers for developing countries in this spirit are Islam and Hoddinott (2009) and Shi (2012).

These policy issues are linked closely to the literature on unitary versus non-unitary models of the household. If households were unitary, with an integrated budget constraint, it would not matter to whom the resources from the public program were targeted. But there is now considerable evidence against the unitary model, with the implication that who has the transfer matters (see Alderman, Chiappori, Haddad, Hoddinott and Kanbur, 1995; and Browning, Chiappori and Weiss, 2014). The detail of intra-household transfers and adjustments thus need to be understood for a welfare state policy whose objective is to reduce inequality among individuals and the deprivation of individuals. This is not the place to delve into the intricacies of the targeting literature as applied to intra-household issues. Haddad and Kanbur (1993), for example, ask how far wrong we go in targeting for nutritional deficiency if we assume that the only inequality is that between household and assume (erroneously) that there is no intra-household inequality. Their numerical calculations show that there is considerable gain to be had by adding information on intra-household inequality.

Two questions have not, however, been fully explored in this literature. First, keeping the objective of maximizing social welfare defined on individual wellbeing, is intra-household redistribution a more potent instrument for achieving impact on social welfare than inter-household redistribution? If we had costless redistribution instruments but could choose only one—intra-household or inter-household redistribution, which would we choose? The answer depends on which of these two components makes the greater contribution to overall inequality between individuals. The empirical findings discussed in section 3 suggest a split ranging from one third/two thirds to two thirds/one third. It is, then, a country specific empirical matter. But if we take the mid-point of this range, half/half, it would seem that focusing on and eliminating one or the other would make an equal contribution to reducing overall inequality.

However, our redistributive instruments are not perfect, and are not costless. Direct redistribution within a household is not easy to effect. In advanced countries if husband and wife choose to elect for separate taxation then income taxes could, in principle, be used to differentiate within a household. But if most households choose to elect household level taxation then this instrument is better at achieving between household redistribution than within household redistribution. Expenditure instruments such as free school meals could in principle be targeted differentially by age and gender and thus affect intra household inequality, but (i) this

effect will depend on intra-household reallocations after the transfer, and (ii) these instruments will also have inter-household effects. Thus, although we have indirect discussions of this question in the literature, in my view it remains an open area of research.

The second question is as follows. Moving beyond the instrumental, what is the normative legitimacy of a welfare state getting deep into intra-household redistribution issues? At what point should such intervention stop? Notice that similar questions arise in the context of global redistribution. If the objective is global inequality (or poverty), there is clearly a possible instrumental role for transfers to, say, poorer countries. But there are deeper philosophical debates between, for example, Global Rawlsians and Rawls himself, on the nature of moral responsibility towards the poor in non-poor nations (Kanbur and Sumner, 2012).

The issue has a family resemblance to discussions of legitimate and illegitimate realms of redistribution. Thus the literature inspired by Roemer (1998) deems it legitimate to address inequality in outcomes attributable to circumstance, factors outside an individual's control, and effort—factors which an individual controls. For within household inequality of consumption, what is legitimate and what is not in the “equality of opportunity” frame depends on the specific model we have in mind for intra-household allocation. If it is a unitary model then presumably it is all determined by the household utility function, which may be exogenous to everybody except perhaps the patriarch, who would have to be identified. If on the other hand it is a collective model then how do we apply the frame? In the case of adults, on the one hand they choose to participate in the collective enterprise, but on the other hand the distribution of surplus from this enterprise is determined by their outside option, which may be thought of as exogenous to them. For children, of course, we would be right in specifying everything as exogenous.

In my view progressives and egalitarians have not addressed the normative question of the legitimacy of intra-household interventions sufficiently. It is an open area for clarification and development.

7. Conclusion

To conclude briefly, assessing the specific contribution of intra-household inequality to standard measures of overall inequality and poverty is an underdeveloped area for research and policy analysis. However, intra-household inequality is clearly important. Neglecting it could lead to (i) an understatement of inequality and an overstatement of the impact of growth on poverty reduction; (ii) a mis-statement of the potential impact of minimum wage policies on poverty; (iii) mis-design of transfer policies to reduce inequality and poverty. Any discussion of the welfare state thus cannot afford to ignore intra-household inequality.

References

- Alderman, Harold, Pierre-Andre Chiappori, Lawrence Haddad, John Hoddinott and Ravi Kanbur. 1995. "Unitary Versus Collective Models of the Household: Is It Time To Shift The Burden Of Proof?" *World Bank Research Observer*, Vol. 10, No. 1, pp. 1-19.
- Basu, Arnab, Nancy Chau and Ravi Kanbur. 2010. "Turning a Blind Eye: Costly Enforcement, Credible Commitment and Minimum Wage Laws," *Economic Journal*, Volume 120, March, pp. 244-269, 2010.
- Bhorat, Haroon, Ravi Kanbur and Natasha Mayet. 2012. "Minimum Wage Violation in South Africa," *International Labour Review*, Vol. 151, No. 3, pp. 277-287.
- Bhorat, Haroon, Ravi Kanbur and Natasha Mayet. 2013. "The Impact of Sectoral Minimum Wage Laws on Employment, Wages and Hours of Work in South Africa", *IZA Journal of Labor and Development*, Vol. 2, Issue 1, pp. 1-27.
- Bhorat, Haroon, et. al. 2016. "Investigating the Feasibility of a Minimum Wage for South Africa." Development Policy Research Unit, University of Cape Town, http://www.dpru.uct.ac.za/sites/default/files/image_tool/images/36/News_articles/NMW%20Report_final.pdf
- Boateng, E. Oti, K. Ewusi, R. Kanbur and A. McKay. 1992. "A Poverty Profile for Ghana, 1987-88," *Journal of African Economies*, Vol. 1, No. 1, pp. 25-58, 1992.
- Bourguignon, F. (2003): "The growth elasticity of poverty reduction: explaining heterogeneity across countries and time periods," in *Inequality and Growth: Theory and Policy Implications*, ed. by T. S. Eicher and S. J. Turnovsky, Cambridge, MA: MIT Press, Cambridge, MA, Chap. 1, pp. 3-27.
- Brown, Martin, Pierre-Andre Chiappori and Yoram Weiss. 2014. *Economics of the Family*. Cambridge University Press.
- Burkhauser, Richard V. 2015. "The Minimum Wage versus the Earned Income Tax Credit for Reducing Poverty." *IZA World of Labor: Evidence-based Policy Making*: 153. doi: 10.15185/izaw ol.153
- Chakravarty, S., R. Kanbur and D. Mukherjee. 2006. "Population Growth and Poverty Measurement," *Social Choice and Welfare*, Vol. 26, No. 3, pp. 471-483.
- Fields, Gary and Ravi Kanbur. 2007. "Minimum Wages and Poverty with Income Sharing," *Journal of Economic Inequality*, Vol. 5, No. 2, pp. 135-147.
- Foster, J., Greer, J., Thorbecke, E. 1984. A class of decomposable poverty measures. *Econometrica* 52(3), pp. 761-776.

- Haddad, Lawrence and Ravi Kanbur. 1990. "How Serious is the Neglect of Intra-Household Inequality?" *Economic Journal*, Vol. 100, pp. 866-881, September 1990.
- Haddad, Lawrence and Ravi Kanbur. 1992. "Intrahousehold Inequality and the Theory of Targeting." *European Economic Review*, Vol. 36, Issue 2-3, pp. 372-378.
- Haddad, Lawrence and Ravi Kanbur. 1993. "The Value of Intrahousehold Survey Data for Age-Based Nutritional Targeting" , *Annales d'Economie et de Statistique*, No. 29, pp. 65-81.
- Islam, Mahnaz and John Hoddinott. 2009. "Evidence of Intrahousehold Flypaper Effects from a Nutrition Intervention in Rural Guatemala." *Economic Development and Cultural Change*. Volume 57, Number 2, pp. 215-238.
- Kanbur, Ravi. 2001. "Economic Policy, Distribution and Poverty: The Nature of Disagreements," *World Development*, Vol. 29, No. 6, pp. 1083-1094, 2001.
- Kanbur, Ravi. 2006. "The Policy Significance of Inequality Decompositions," *Journal of Economic Inequality*, Vol. 4, No. 3, pp. 367-374, 2006.
- Kanbur, Ravi. 2010. "Globalization, Growth and Distribution: Framing the Questions," in Ravi Kanbur and A. Michael Spence (Editors), Equity in a Globalizing World, The World Bank for the Commission on Growth and Development, pp. 41-70. 132.
- Kanbur, Ravi and Andy Sumner. 2012. "Poor Countries or Poor People? Development Assistance and the New Geography of Global Poverty", *Journal of International Development*, Vol. 24, Issue 6, pp. 686-695.
- Kanbur, Ravi and Matti Tuomala. 2016. "Groupings and the Gains from Targeting", *Research in Economics*, Vol. 70, pp. 53-63, 2016.
- Lise, Jeremy and Shannon Seitz .2011. "Consumption Inequality and Intrahousehold Allocations." *The Review of Economic Studies*, Vol. 78, No. 1, pp. 328-355.
- Paes de Barros R, Ferreira FHG, Vega JRM, Saavedra JC, De Carvalho M, Franco S, Freije-Rodriguez S, Gignoux J. 2009. *Measuring Inequality of Opportunities in Latin America and the Caribbean*. World Bank: Washington DC.
- Roemer JE. 1998. *Equality of Opportunity*. Mass., Harvard University Press: Cambridge.
- Sahn, David and Stephen D. Younger. 2009. "Measuring Intra-Household Inequality: Explorations Using the Body Mass Index," *Health Economics* 18(S1): S13-S36.
- Sen, Amartya K. 1976. "Poverty: an ordinal approach to measurement". *Econometrica* 44: pp. 219-231.

Shi, Xinzheng. 2012. "Does an intra-household flypaper effect exist? Evidence from the educational fee reduction reform in rural China." *Journal of development Economics* 99, No. 2, pp. 459-473.

OTHER A.E.M. WORKING PAPERS

| WP No | Title | Fee (if applicable) | Author(s) |
|---------|--|------------------------|--|
| 2016-10 | Anticipatory Signals of Changes in Corn Demand | | Verteramo Chiu, L., Tomek, W. |
| 2016-09 | Capital Flows, Beliefs, and Capital Controls | | Rarytska, O., Tsyrennikov, V. |
| 2016-08 | Using Unobtrusive Sensors to Measure and Minimize Hawthorne Effects: Evidence from Cookstoves | | Simons, A., Beltramo, T., Blalock, G. and Levine, D. |
| 2016-07 | Economics and Economic Policy | | R. Kanbur |
| 2016-06 | W. Arthur Lewis and the Roots of Ghanaian Economic Policy | | R. Kanbur |
| 2016-05 | Capability, Opportunity, Outcome - -and Equality | | R. Kanbur |
| 2016-04 | Assessment of New York's Pollution Discharge Elimination Permits for CAFO's: A Regional Analysis | | Enahoro, D., Schmit T. and R. Boisvert |
| 2016-03 | Enforcement Matters: The Effective Regulation of Labor | | Ronconi, L. and R. Kanbur |
| 2016-02 | Of Consequentialism, its Critics, and the Critics of its Critics | | Kanbur. R |
| 2016-01 | Specification of spatial-dynamic externalities and implications for strategic behavior in disease control | | Atallah, S., Gomez, M. and J. Conrad |
| 2015-13 | Networked Leaders in the Shadow of the Market - A Chinese Experiment in Allocating Land Conversion Rights | | Chau, N., Qin, Y. and W. Zhang |
| 2015-12 | The Impact of Irrigation Restrictions on Cropland Values in Nebraska | | Savage, J. and J. Ifft |
| 2015-11 | The Distinct Economic Effects of the Ethanol Blend Wall, RIN Prices and Ethanol Price Premium due to the RFS | | de Gorter, H. and D. Drabik |
| 2015-10 | Minimum Wages in Sub-Saharan Africa: A Primer | | Bhorat, H., Kanbur, R. and B. Stanwix |

Paper copies are being replaced by electronic Portable Document Files (PDFs). To request PDFs of AEM publications, write to (be sure to include your e-mail address): Publications, Department of Applied Economics and Management, Warren Hall, Cornell University, Ithaca, NY 14853-7801. If a fee is indicated, please include a check or money order made payable to Cornell University for the amount of your purchase. Visit our Web site (<http://dyson.cornell.edu/research/wp.php>) for a more complete list of recent bulletins.