

1 Introduction

**Stephen P. Jenkins, Arie Kapteyn and
Bernard M.S. van Praag**

The measurement and explanation of household welfare is a multi-faceted undertaking. If one wants to study the well-being of households, one needs to understand household behavior and one needs to devise measures of well-being. The distribution of welfare across households is dispersed in any society, and the additional question arises how to characterize and evaluate this dispersion. Moreover, there are various ways in which one can learn more about the behavior and well-being of households: from pure reflection to broad-based empirical international comparisons; from looking at behavior only to the additional use of an array of subjective indicators of well-being. The chapters in this volume reflect many of, but certainly not all, the facets of household well-being and behavior, and the various approaches one can take to learn more about them.

The choice of approaches and facets is largely inspired by the scientific program adopted by Aldi Hagenaaers in her all-too-brief career. This book reflects her interests and her methodological stance. A beautifully written and balanced review of Aldi Hagenaaers' scientific program is given by Joop Hartog in chapter 2. It is at the same time an apt summary of the plan and the philosophy of this book. Aldi Hagenaaers' work dealt with just about every conceivable economic aspect of household behavior. The starting point may have been a neoclassical one, but it was only a starting point and not a straitjacket. Hartog describes Aldi's work as an attempt to leave the neoclassical citadel in many directions simultaneously. The "citadel" metaphor can be fruitfully developed further: the neoclassical economics citadel is too small to hold all the insights to be acquired about household behavior, and so the fort needs to expand. If we are to make real progress in understanding household behavior and welfare, the city has to grow and gradually the old town walls will become obsolete: interesting relics for tourists or the archaeologist rather than offering true protection against the outside world.

2 Stephen P. Jenkins, Arie Kapteyn and Bernard M.S. van Praag

So what are the directions in which the neoclassical paradigm has to be extended? Some of them are non-controversial and have been dealt with by various authors: dynamics, more than one decision maker in the household, non-linear wage schedules, incomplete information, constraints on working hours, joint production. Others such as preference interdependence across households, or deviations from rationality generate more discussion. Even more controversial than some of the elements of theory that have to be amended are the tools that are used to do it.

Aldi Hagenaaars was one of the members of what is sometimes referred to as the “Leiden School” (economists associated with Bernard van Praag), a group well known for recommending that welfare (or utility, if you prefer) be measured by means of direct questions. By its very nature, this “subjective” approach relates to central questions of social policy such as poverty, income inequality, and welfare, and it generates an interest in the use of household surveys to learn more about personal behavior and welfare. These elements not only characterize Aldi Hagenaaars’ work, they also sum up the approach of the chapters in this book.

Cross-national comparisons of the distribution of welfare

The approach of wanting to learn about the real world not hindered by preconceptions (or at least as few as possible), leaves plenty of room for the exploitation of diverse and rich data sources. Chapter 3, by Aldi herself and two co-authors, Klaas de Vos and Asghar Zaidi, is a prime example. It constitutes a report on a pioneering effort to construct poverty statistics for the European Community (EC). The sheer work of getting together microdata from twelve countries, processing them and solving the numerous inconsistencies that exist between them is mind-boggling. Just for a start, it turned out that data on household incomes in various countries were of too poor quality to be used in the analysis. Hence the authors were forced to use expenditure data instead. This may well prove to have been a blessing in disguise, for if one assumes that households are lifetime consumption smoothers, consumption may be a better reflection of life cycle resources than current income. Not surprisingly, this chapter pays considerable attention to issues of data quality, different definitions of poverty, different equivalence scales to compare the level of resources of different households, etc. The conclusions that one draws about the incidence of poverty in the EC are quite sensitive to the various choices that can be made, and policy makers should be aware of this.

The same point is made by the authors of chapter 4, Tony Atkinson, Karin Gardiner, Valerie Lechêne, and Holly Sutherland. They compare

poverty rates between France and the United Kingdom and investigate the implications of a number of measurement choices. These include the use of expenditure versus income, defining the poverty line as a percentage of the median versus mean income, using persons versus households as the income-receiving unit, measuring income prior to or net of housing costs, using a variety of equivalence scales, and several fractions of median or mean income (40 percent, 50 percent, or 60 percent) to define the poverty line. They show that, for some choices, poverty rates in France and the United Kingdom are virtually identical, and in other cases poverty in France appears to be at least twice as high as in the United Kingdom. The upshot of their analysis is that politicians asking statisticians to calculate poverty rates on the basis of a rather loose definition may get very different answers, depending on how the statisticians fill in the details. This conclusion is strengthened further if one considers alternative poverty measures such as the poverty gap. It is then not difficult to “prove” that either France or the United Kingdom has the “most poverty,” just by choosing the “right” measure and the “right” operationalization of it.

Neither income nor consumption expenditure are ideal statistical measures of material well-being for poverty measurement because both ignore variations across households in preferences for leisure or differences in the level of household production. Assuming that the ultimate concern with poverty derives from concerns about the welfare of households and their members, such neglect may be serious. Chapter 5, by Peter Saunders, Inge O’Connor, and Tim Smeeding, tries to correct for this omission by employing the concept of earnings capacity, originally introduced by Garfinkel and Haveman (1977). The basic idea is to look at the income people would have were everyone to spend the same amount of time in market work, in particular if each person were in full-time employment during the full year (so-called FYFT employment). This is earnings capacity. In order to operationalize such a concept, one has to impute the FYFT earnings for everyone who does not already work full time for the full year. This can be done by regressing log annual earnings on a set of explanatory variables using data for the people observed with FYFT employment. Next the estimated regression coefficients and error variance can be used to predict, and therefore impute, the earnings of everyone else not in FYFT employment.

The method is applied using data for the mid-1980s for five countries: Australia, Canada, West Germany, the Netherlands, and the United States. (The data derive from the Luxembourg Income Study (LIS) and only married couples are considered.) These countries vary quite substantially in terms of unemployment and participation rates, so it comes as

4 Stephen P. Jenkins, Arie Kapteyn and Bernard M.S. van Praag

no surprise that moving from actual earnings to earnings capacity has different effects in different countries. For instance, mean earnings go up by 28 percent in the United States if we move from actual earnings to earnings capacity. In the Netherlands the corresponding amount is 53 percent. The adjustments reduce inequality in all countries concerned, but again the change is much larger in the Netherlands than in the United States or Canada.

Besides the issues concerning analysis of income and earnings distributions across and within countries cited so far, there is one dimension – income mobility – which has not been mentioned at all. In addition to cross-section snapshots on income, one also needs longitudinal movies. In chapter 6, Richard Burkhauser, Douglas Holtz-Eakin, and Stephen Rhody add this additional twist to the study of income distributions, and compare the perspectives derived.

The relevance of mobility for distributional assessments can be brought home by considering a very simple example. Imagine a society consisting of 100 individuals, with an income distribution which is completely fixed over time. Assume furthermore that individuals live for 100 periods, and that incomes are permuted systematically over time in such a way that in each period individuals have a different income. Thus after 100 periods, the income history of every individual differs only in the order in which incomes were received. If, furthermore, the interest rate is zero and capital markets are perfect, everyone has the same lifetime income equal to the mean income in the society. This would be an example of a society with constant cross-sectional income inequality, substantial mobility, and complete lifetime equality of resources. In a somewhat less stylized world one could at least imagine that a society with a high income inequality at any given point in time and high mobility could actually exhibit a distribution of lifetime resources as equal as that of another society with lower cross-sectional income inequality and lower mobility.

Burkhauser, Holtz-Eakin, and Rhody compare income inequality and mobility in Germany and the United States in the second half of the 1980s. They find higher income inequality in the United States than in Germany (which is not unexpected), but surprisingly they do not find very large differences in mobility across the two countries. (The conclusions vary somewhat, depending on the exact definition of “mobility” and the income measure considered.) Why the greater competitiveness of the US labor market did not lead to higher mobility as well as higher inequality is an intriguing finding, deserving of further investigation. Whatever the case, the authors’ results underscore the importance and relevance of analyzing mobility as well as inequality when contemplating social policy.

Subjective approaches

A crucial issue for virtually any anti-poverty policy is how to define “poverty” in the first place. This issue has several aspects, including the choice of a poverty line (identification) and the summarizing of the information about individual material deprivation (aggregation). Rather than telling us what in their opinion the correct, or at least a reasonable, approach should be, Yoram Amiel and Frank Cowell in chapter 7 asked a group of 340 students in six countries about their perceptions of poverty. They showed the respondents various hypothetical income distributions and poverty lines, and then asked them to compare these situations in terms of the severity of poverty associated with these hypothetical situations. In this way, they were able to uncover different dimensions in the perception of poverty – for instance, whether “poverty” is perceived mainly to be a relative or an absolute concept, the difference between a change in poverty implied when all incomes change proportionally or by the same amount, etc. The empirical results are not clear cut, but they are intriguing. It appears, for instance, that slight manipulation of the hypothetical situations offered to the respondents may have quite an effect on judgments about poverty. These subtleties justify extensive further research into what shapes our perception of “poverty.”

Peter Kooreman investigates the explanatory power of subjective information in the modeling of consumer choices in chapter 8. This is another example of the many ways in which one may cross the boundaries of traditional neoclassical revealed preference territory and learn something on the way. Kooreman starts out by building a neoclassical model of consumer choice, in this case relating to the choice of a country in which to spend a vacation. Since different countries provide different characteristics, a utility function defined in terms of vacation characteristics (and other consumption) can be employed to model the choice of destination country as a constrained utility maximization problem. The availability of a data set with respondents’ subjective statements about the importance of various vacation characteristics (e.g. whether there are many possibilities for sightseeing, the quality of food and drink) permits empirical modeling. Although some methodological caveats have to be kept in mind, the overall impression is that the subjective data contribute significantly to an explanation of destination country choices.

Whereas Kooreman uses subjective measures as explanatory variables, Isolde Woittiez and Jules Theeuwes (chapter 9) use a subjective measure of well-being (a score on the Cantril scale) as a dependent variable. They

try to explain variations in self-reported well-being by a host of variables: labor market status, income, education, age, sex, housing variables, variables referring to the respondent's partner, labor market history, and various proxies for preferences for work. There are many reasons why such analyses are of interest. For instance, one may want to look at income maintenance policies, family policies, and the like. The authors concentrate on the effects of labor market status, using the other variables as controls. Generally non-workers (unemployed, disabled, pensioners, non-participants) are less happy than workers, with the exception of the early retired. A closer analysis, using further subjective information, suggests that the main reason for this is the involuntary nature of these non-working states. This also explains why the early retired are about as happy as the working population: freedom to choose breeds happiness.

Summarizing welfare

What is "inequality?" Or, to put it in Gary Fields' terms, "Do inequality measures measure inequality?" Most of the economic literature on inequality looks at income distributions (or distributions of anything else) from a social welfare viewpoint. Famous examples are the inequality measures developed by Atkinson (1970) and Dalton (1920). In chapter 10 Fields abstracts from this, and tries to arrive at reasonable criteria by which one distribution can be judged more equal than another. Alongside the generally accepted Lorenz criteria, Fields develops criteria based on the concepts of "elitism of the rich" and "isolation of the poor." This leads to a class of inequality measures which may be narrowed or broadened depending on which further specific criteria one may find reasonable to impose when considering inequality. The chapter nicely traces out where consensus among observers may be likely and where different people will hold different opinions on what "inequality" really means.

A different approach is taken by Tony Shorrocks in chapter 11. He considers the distribution of "bads" (as opposed to "goods") and introduces what he labels "deprivation profiles." The deprivation profile is related to the distribution function of bads (and therefore also to a transformation of the generalized Lorenz curve). Examples of bads include personal material deprivation, earnings discrimination, and horizontal inequity. Shorrocks' chapter usefully draws attention to common elements in measurement methods used in a variety of contexts, and shows how different distributions of bads may be ranked in terms of the amount of deprivation they entail. He considers a number of basic

desirable properties which any aggregate summary measure should possess and thence defines a class of aggregate deprivation indices (which encompasses the Hagenaars–Dalton family). Shorrocks shows how unanimous deprivation index orderings are equivalent to orderings by deprivation curves, and illustrates his arguments with a specific application to poverty measurement. Shorrocks then applies this apparatus to the study of poverty and the properties of classes of poverty indices. The deprivation profiles are very helpful in investigating circumstances in which different poverty indices will yield similar rankings of poverty across countries or over time.

The robustness of poverty rankings to the choice of summary measure is also James Foster's and Yong Jin's starting point in chapter 12. They consider the influence on measured poverty of both the choice of poverty line and the sensitivity of the poverty index to differences in the distribution of resources below it. The authors concentrate on so-called Dalton utility-gap indices, i.e. indices that are additive functions of the shortfalls between each individual's utility of poverty line income and the utility of actual income. They show that poverty comparisons can be undertaken by considering the associated utility of income distributions and invoking well known theorems on stochastic dominance. One particularly nice example of robustness of the Dalton utility-gap measures is provided by their theorem 2, which says that if we use two different utility functions to operationalize the index and we find for one utility function that a given income distribution exhibits less poverty than another, then this conclusion will remain unchanged if we replace the utility function by one which is more risk averse.

The economies in transition in Eastern Europe are a fascinating laboratory to study how changes in economic structure and social policy affect the distribution of well-being. In her chapter on the changes that took place in the Czech and Slovak republics between 1989 and 1992, Thesia Garner (chapter 13) explores the evolution of income and expenditure inequality during this period. Most western observers would probably guess that inequality would increase substantially when capitalism was introduced. Surprisingly, this does not seem to have happened; certainly not in the Czech republic, and only a little in the Slovak republic. Garner ascribes the slower than expected movement in inequality to a conscious policy, especially in the Czech republic, which reputedly has a strong preference for equality. Combined with the relative prosperity of the Czech republic compared to other economies in transition, the results suggest that economic transition in the Czech republic may be one of the smoothest in existence.

The household, income, and welfare

Describing or characterizing income inequality or the incidence of poverty is one thing; influencing it is another. A prominent policy instrument for influencing the distribution of personal economic well-being is unemployment insurance (UI). In chapter 14, Lars Osberg, Sadettin Erksoy, and Shelley Phipps consider the change in the distribution of well-being associated with UI reforms in Canada. In particular, they compare the 1971 UI regime with the 1994 UI reforms. The 1994 reform substantially reduced the role of the welfare state in income maintenance policies. Such changes not only have incentive effects, e.g. on labor supply, and redistributive effects (the usual focus), but they also affect the income uncertainty people face. The incidence of uncertainty, and hence the impact of the reform, will be different for different groups of people. To evaluate the impact of reforms on economic well-being, one has to simulate behavior under both the old and the new regime, and to evaluate the welfare cost of increased uncertainty. This is exactly what the authors do, using an elaborate micro-simulation model with a database of some 20,000 Canadians (respondents to Statistics Canada's 1984 Survey of Assets and Debts). The labor supply part of the simulation model is based on a Stone–Geary specification for the utility function of individuals, and the same Stone–Geary specification is used to evaluate the welfare loss due to increased uncertainty. The chapter's findings indicate that the UI cutbacks increased inequality and may have decreased average economic well-being.

Although the parameters of the Stone–Geary utility function have been estimated from the behavioral model underlying the simulations, there is one piece of information missing. The labor supply model is static, and hence any monotonic transformation of the utility function will yield identical labor supply, but not identical risk premia. The authors compare the risk aversion implied by their utility specification with estimates from other sources (e.g. the financial literature) and conclude that the Stone–Geary specification employed probably entails too little risk aversion. This implies that their estimate of the welfare loss due to increased uncertainty is a conservative one.

Educational attainment is a primary determinant of the income an individual may expect to enjoy during his or her life. Conversely, decisions to invest in schooling may be expected to be influenced by the extra expected income it generates. In chapter 15, Robert Haveman, Kathryn Wilson, and Barbara Wolfe build a structural model in which an individual's decision to invest in schooling is a function of individual characteristics, family characteristics, neighborhood characteristics, and

the expected income stream associated with the schooling choice made. To properly operationalize such a model requires an extremely rich data set. Haveman, Wilson, and Wolfe use the US Panel Study of Income Dynamics (PSID) and add neighborhood information by matching small area data from the 1970 and 1980 Censuses to the location of the children in the sample. The panel nature of the data set makes it possible, amongst other things, to construct expected income streams conditional on decisions to graduate or not for young people with varying characteristics. These expected income streams feed into a structural model of schooling choice. In terms of significance of estimates and plausibility of results the model appears to work extremely well. Expected income has a significant effect on schooling choice. For instance, the authors estimate that if one were able to increase the future expected income of high school graduates by 10 percent, this would reduce drop out rates by somewhere between 6 and 18 percent. Conversely, income maintenance schemes that gave more support to high school drop outs (either by design, or inadvertently), would increase drop out rates.

One problem intrinsic to any empirical study of welfare across households is the choice and construction of equivalence scales to account for differences in household size and composition (or more generally, “needs”). Chapters 3 and 4, by Aldi Hagenars, Klaas de Vos, and Asghar Zaidi, and by Tony Atkinson, Karen Gardiner, Valerie Lechêne, and Holly Sutherland, draw attention to the sensitivity of empirical conclusions to the choice of equivalence scale. However, rather than having to choose an equivalence scale, one would prefer to let the data, in combination with some well established theory, tell us what equivalence scale to choose. Unfortunately, as the chapter by Richard Blundell (chapter 16) bears out, the story told by theory is largely a negative one: finding unique equivalence scales is impossible, at least on the basis of the data usually available. In essence, the problem is that information on expenditures of households can tell us a fair amount about the costs of children, but not a lot about their benefits. This fundamental identification problem has been known for some time, at least since Pollak’s and Wales’ (1979) paper. Blundell characterizes this identification problem borrowing from his important work with Arthur Lewbel (Blundell and Lewbel, 1991), and reminds us that the identification problem can be solved only by invoking extra information. This information may refer to preferences for different family compositions, subjective information, or it may come from “reasonable” restrictions imposed on preferences by the investigator.

To these possibilities Blundell adds exploitation of information about intertemporal choices. It is well known that, intra-temporally, a utility

function can be identified only up to a monotonic transformation, and since the transformation may depend on family composition, this is the root of the identification problem. However, it is also well known that intertemporal choice restricts the number of transformations of a utility function consistent with observed behavior to a much smaller class. Hence, if one has intertemporal data on consumption and demographics, further headway may be made. This is a neat idea worth pursuing.

The chapter by Blundell is theoretical. By contrast David Johnson's chapter (chapter 17) is empirical, looking at inequality across households and persons. In doing so, Johnson has to address at least two significant issues, namely the choice of equivalence scales to deflate income (or consumption) of different households, and how to account for income inequality within each household. Both issues are practical and fundamental. The fundamental part has to do with an observational problem: generally we do not observe intra-household allocation, nor do we know, as Blundell's chapter shows, how family composition itself influences welfare. One way of handling these problems is to make a range of different assumptions and to explore the extent to which these affect the conclusions of one's analysis, and Johnson does precisely this. He estimates eight different equivalence scales (using the US Consumer Expenditure Survey) and assumes four different sharing rules for consumption within the household. For these various possibilities he then considers four different inequality measures. Not perhaps surprisingly, the estimates of the inequality measures vary considerably across scales and sharing rules. This forcefully drives home the point that we need to know much more about what is going on within the household if we want to have a really reliable basis for social policy.

A similar issue is taken up by Stephen Jenkins and Nigel O'Leary in chapter 18 when studying trends in the income differences between men and women in the United Kingdom between 1971 and 1991. Jenkins and O'Leary consider only two polar sharing rules: the first (equal sharing) assumes that within a family all incomes are shared equally, whereas the second (minimal sharing) assumes that no sharing takes place at all. It seems safe to assume that the truth is somewhere in between these extremes. Observe nonetheless that virtually all income distribution analyses, including official ones, use the equal sharing rule. It is interesting to note that the evidence for a closing of the gender gap is much more clear cut under the minimal sharing than under the equal sharing rule. Since the number of women earning an income in the labor market increased over these two decades, and since women tend to earn more than before on average, their income distribution became more similar to the income distribution for men. The situation under an equal sharing