

Testing a social safety net

Martin Ravallion*, Dominique van de Walle, Madhur Gautam

The World Bank, 1818 H Street, NW, Washington, DC 20433, USA

Received April 1993, revised version received March 1994

Abstract

Standard benefit-incidence analysis does not distinguish policy impacts on persistent poverty from transient poverty. We offer an alternative approach, based on actual and simulated joint distributions of consumption over time, which allows us to distinguish the extent of ‘protection’ against poverty from ‘promotion’ out of poverty. The approach is illustrated by an analysis of the distributional impact of changes in cash benefits introduced to compensate for other policy reforms in Hungary. Cash benefits protected many from poverty, but promoted few out of poverty. The safety net’s impact on poverty was largely due to higher average outlays, rather than improved targeting.

Key words: Poverty; Transfers; Targeting; Mobility; Hungary

JEL classification: I32; I38

1. Introduction

Empirical evidence on the performance of social safety nets is typically *static*; it describes how the incidence of transfers varies according to some measure of the recipient’s current standard of living. That is all cross-sectional surveys allow. This static incidence picture may be quite uninformative about the distributional impacts of policy changes. In many settings, including economies in transition, household living standards are changing over time in often unpredictable ways. The standard incidence

* Corresponding author.

table cannot tell us how much of any reduction in poverty was due to better *protection* of those vulnerable to poverty, versus better performance at *promoting* the poor.¹ The same post-intervention distribution of living standards can be produced in any number of ways; for example, two policies may yield the same number of poor, yet in one case many more fall into poverty, and many escape, than in the other. Clearly, we may be far from neutral to such differences when evaluating a social safety net.

This paper outlines a straightforward approach to assessing dynamic incidence with panel data. We propose measures that distinguish a policy's ability to protect the poor—interpretable as its impact on transient poverty—from its ability to promote the poor—its impact on persistent poverty. We then use this approach to examine the performance of Hungary's social safety net during the late 1980s. This setting is of wide interest for a number of reasons, not least that Hungary has been going through transition from a command-driven to a market-driven economy. Policy reforms during the transition have helped some but hurt others, and the country's system of cash benefits has been used to try to help compensate those likely to be hurt most. We ask how well the safety net performed this function.

The paper is organized as follows. Section 2 discusses some conceptual issues that arise in testing a social safety net, and outlines our chosen approach, and its drawbacks. Section 3 then describes some key features of the setting with bearing on our investigation, while Section 4 discusses the new data set constructed for this study. Section 5 presents our results, while our conclusions are summarized in Section 6.

2. Generic issues in testing a social safety net

2.1. *Measuring 'protection' and 'promotion'*

How should the poverty impacts of the social safety net be quantified? In constructing the usual static incidence picture, or 'poverty profile', households are typically ranked according to some indicator of living standards, and the receipts of various components of social expenditures are compared. Assessing dynamic incidence demands a departure from this method. With panel data, instead of relying on the static-univariate distribution, we can construct the joint distribution over time, in which the panel structure is exploited to show how households moved between welfare groups.²

¹ On this distinction, see Drèze and Sen (1989).

² While there have been a number of empirical studies of aspects of the dynamics of poverty (including Adelman et al., 1985; Bane and Ellwood, 1986; Ravallion, 1988; Ruggles and Williams, 1989; and Chaudhuri and Ravallion, 1994), the difference with our approach is in simulating the effects of safety-net policies on the joint distribution.

In comparing joint distributions—such as with and without policy changes—we will use two tests: how well people are protected from poverty, and how well they are promoted from poverty. To define these, let x denote the indicator of living standards (discussed further below), found in the interval $(0, x^{\max})$. Consider two possible joint distribution functions over dates 1 and 2, namely $F(x_1, x_2)$ and $G(x_1, x_2)$ [i.e. $F(x_1, x_2)$ is the proportion of the population with less than x_1 in period 1, and less than x_2 in period 2, and similarly for $G(x_1, x_2)$]. The corresponding marginal distributions are $F_1(x_1) = F(x_1, x^{\max})$ and $F_2(x_2) = F(x^{\max}, x_2)$, and similarly for G . The poverty line is z , and so the proportion of the population who are poor in period 1 in the F distribution is $F_1(z)$, while a proportion $F_2(z)$ are poor in period 2. By construction, $F_2(z) - F(z, z)$ is the proportion of individuals in the F distribution who are poor in the second period but were not poor in the first. We will say that F protects from poverty better than G if and only if

$$F_2(z) - F(z, z) < G_2(z) - G(z, z).$$

The extent of protection allowed by F will be measured by

$$PROT(z) = G_2(z) - G(z, z) - F_2(z) + F(z, z). \quad (1)$$

Analogously, $F_1(z) - F(z, z)$ of the population were poor in the first period but not the second. F promotes the poor better than G if and only if

$$F_1(z) - F(z, z) > G_1(z) - G(z, z).$$

And the extent of promotion due to F relative to G will be measured by

$$PROM(z) = F_1(z) - F(z, z) - G_1(z) + G(z, z). \quad (2)$$

In all cases considered in this paper the marginal distributions in the first period are identical: $F_1(z) = G_1(z)$, which is simply the pre-intervention distribution. It follows that promotion is equivalent to requiring that $F(z, z) < G(z, z)$, i.e. $PROM$ can be interpreted as a test of whether there is less persistent poverty in the F distribution, the persistently poor being defined as those who were poor in both periods. The residual, $F_2(z) - F(z, z)$, is then interpretable as the amount of transient poverty, which is precisely what $PROT$ tests for. Another implication of identical first-period marginals is that if *both* $PROT$ and $PROM$ are positive, then $F_2(z) < G_2(z)$ (i.e. the incidence of poverty is lower for the F distribution in period 2), though the converse is not true (lower poverty in period 2 is possible with only one of $PROT$ or $PROM$ being positive).

All of these definitions can be generalized to multiple poverty lines. For example, one may ask whether or not $PROT$ holds for all possible z in some interval. If both $PROT$ and $PROM$ are positive, then there must be

first-order stochastic dominance in period 2 over that interval, implying unambiguous poverty comparisons (Atkinson, 1987).

2.2. *Welfare measurement issues*

The choice of welfare indicator may matter to the results obtained using the above methods. One issue is whether the ‘standard of living’ at some date is best measured by commodities actually consumed, rather than *potential* consumption.³ Here we take the former view, though we would note that even if one preferred the latter concept, income would be an imperfect measure of potential consumption, as households will also differ in their liquid wealth, which is rarely known from survey data. The choice between income and consumption can clearly matter in a transition economy since pre-transition wealth (and, probably less so, borrowing) can be used to buffer current living standards to some degree. Two households facing unemployment, one with initial wealth, the other without, will be affected quite differently. It seems likely that consumption will better reflect that difference than income.

Another issue is that households differ in demographic composition and may face different prices at a given date. To deal with this heterogeneity, all consumption expenditures are normalized here by the household-specific and date-specific poverty lines. Thus the welfare comparisons here are based on estimates of ‘welfare ratios’ (Blackorby and Donaldson, 1987). This is only one of a number of possible approaches; alternatives include ‘money metric utility functions’ calibrated to models of consumption/labor supply behavior, including certain approaches to forming demographic scales as special cases.

The credibility of the welfare ratios depends in part on that of the poverty lines used, which should (in theory) be points on the consumer’s cost function corresponding to the poverty line in utility space. In practice, there are serious identification problems in retrieving the cost function from observed demand behavior (see, for example, Pollak, 1991). And the properties of poverty lines—such as the way they adjust for spatial differences in the cost of living, and differences in household size and composition—can have bearing on the conclusions drawn from poverty comparisons (for an overview, see Ravallion, 1994). For example, the choice of equivalence scale can alter how well targeted a policy such as Hungary’s family allowances is to the poor (Jarvis and Micklewright, 1994).

The scales built into the poverty lines can also have implications for the extent of measured mobility. Since the sizes of households in a panel typically change over time, errors in the parameterization of the demo-

³ For recent discussions of the issues, see Atkinson (1991) and Ravallion (1994).

graphic scale could alter the transition probabilities. For example, suppose one used household size (so that welfare is measured by consumption expenditure per person). There are surely scale economies in household consumption; two persons can achieve the same standard of living more cheaply living together than apart. Then true welfare (expenditure per equivalent number of single-person households say) may be constant over time, and yet measured welfare (consumption per actual person) varies. Or the two may move in opposite directions.

The poverty lines that we have used were constructed by the Central Statistical Office (CSO) of Hungary, and we discuss them further in Section 4. While we are confident that the CSO is in a better position than us to decide what poverty lines for Hungary should look like, and we are pessimistic about the prospects of resolving this issue to any greater satisfaction through demand analysis,⁴ one should nonetheless be aware of the potential sensitivity of policy conclusions to the value judgements implicit in the poverty lines used. We comment further on this issue in Section 4.

2.3. Behavioral responses

Whether one uses consumption or income, a common assumption in incidence analysis is that pre-intervention status is revealed by simply subtracting benefits received. This is questionable. Behavioral responses through labor supply, inter-household transfers, and inter-temporal decision-making could greatly alter the incidence result.

The treatment of such responses depends in part on the concept of living standards used. If a household saves part of an increment to social income, and one uses current consumption as the welfare indicator, then one should net out the saved portion. However, if instead one prefers to measure living standards by the opportunity for consumption, then one would treat the saved portion the same way as the consumed portion.⁵

We do not aim to resolve these issues here, but simply to test the sensitivity of results to the choices made. In particular, we will also consider simulations in which only the change in current consumption is valued. In principle, a well-specified and realistic behavioral model could reveal this; in practice, one is not sure what exactly such a model would look like, or how it would be estimated with available data. Here we adopt an ad hoc approach, which is still capable of identifying the key empirical parameters needed for a behavioral incidence analysis. In particular, we use an

⁴ See Browning's (1992) review, and the discussion in Ravallion (1994).

⁵ Similarly, if some of the social income gain is used to 'finance' an increase in leisure—via a change in work effort—how this should be treated will depend in part on whether leisure is valued in assessing the level of living.

econometric model of consumption to estimate the propensity to consume out of social income (*PCSI*), giving the change in consumption expected from a change in social income; a *PCSI* that is positive but less than unity implies that behavioral responses exist, but that net benefits are still positive at the margin. The model is limited, however, in that it does not tell us about the *structure* of those responses, which may involve any of the channels mentioned above. This may matter; one might feel quite differently about a household that saves an increment to social income versus one that works less. These responses will be treated symmetrically here; all that we identify is the net gain to current consumption.

In a cross-section regression of consumption on social income, one would naturally be concerned about omitted variable bias, given that receipts of social income are correlated with a variety of household characteristics, in part through policy design. One could deal with this to some extent by including measured household characteristics as additional regressors. But panel data allow a better—and widely used—option: to exploit the panel structure to estimate a model of consumption in which household fixed effects are eliminated by differencing. That is the course we follow, though we note that this procedure comes at a cost if there is a high degree of noise in the data due to date-specific measurement errors.⁶

3. The setting and policy issues

The transition from a command economy to a market economy poses a number of difficult problems for social policy, not least of which is the issue of how effective the existing social safety net is in preventing an increase in poverty during the transition.⁷ This is more than a concern about the safety net's performance in reaching the persistently poor; there is at least as great a fear that the safety net may be unresponsive to changing household circumstances, and thus relatively ineffective in protecting those who are vulnerable. There is conflicting evidence on whether or not poverty rose in Hungary during the 1980s; poor inactive households (mainly pensioners)

⁶ For example, if two individual- and date-specific variables are given by an individual-specific time mean plus a date-specific white-noise error process, then differencing will entail regressing white noise on white noise with an understandably poor fit; see Deaton (1994). Of course, this is only an example, and it is not an argument against ever differencing the data, as other (familiar) examples can be constructed which would entail equally serious problems for inferences if one does not (see, for example, Hsiao, 1986); rather it speaks to the need for caution in interpreting poor fits in difference regressions.

⁷ For recent reviews of these issues in the Hungarian context, see Jarvis and Micklewright (1994), Fajth and Vita (1992), and World Bank (1992). Evidence on the static incidence of the social safety net is reviewed in Atkinson and Micklewright (1992) and Milanovic (1994).

appear to have seen net gains between the late 1970s and the starting point of this study (1987), while active households have seen no gains over the period, and (some data sources suggest) they may well have experienced rising poverty (Szalai, 1989; Atkinson and Micklewright, 1992). There have been no attempts (to our knowledge) to estimate the contribution of the social safety net to these changes.

During the period covered by our data, the social safety net comprised: (i) employment-related social insurance (pensions, sick pay, family allowances, maternity and child care allowances, and, from 1989, unemployment benefits); (ii) universal benefits (social benefits in kind, including education and health care); and (iii) a limited number of means-tested transfers (social aid and student aid). In recent years, many benefits under (i) have become more widely available due to high labor participation rates and alterations in the rules to expand coverage to non-contributors. The Hungarian social safety net has traditionally been based on the social insurance model (Atkinson and Micklewright, 1992); as a consequence it has few components that restricted eligibility to the poor explicitly. This fact has motivated proposals for 'better targeting' of transfers (World Bank, 1992).

Benefits that do not use a means test can still be well targeted. The indirect indicators of poverty—such as geographic location or family size—built into a scheme, and the incentives the scheme creates for self-selection, as they interact with the behavior of potential recipients, can have a powerful effect on the final distribution of the benefits. It remains an empirical question: How well targeted are cash benefits? These are generic issues in transition economies (compare, for example, Barr, 1992, on Russia), and elsewhere (Ravallion and Datt, 1994, for India).

There were some specific policy changes in Hungary during the period 1987–1989. Tax reforms, including a new personal income tax and VAT, were introduced in 1988.⁸ Several social security and budgetary reforms were also introduced by early 1989. Universal consumption subsidies were cut. The tax and spending reforms are likely to have hurt some groups more than others. Cash benefits were adjusted to protect only those deemed especially vulnerable, notably children and pensioners. These adjustments took the form of increments to family allowances and pensions. Thus, while families with children and pensioners were somewhat compensated for the policy reforms, others, such as wage earners without small children, probably experienced lower real incomes.

⁸ The complexity of taxation and pension-fund contribution rules make it extremely difficult to infer from available data what these variable magnitudes were for 1987, though they are available for 1989. Aside from the general issues raised in the previous section, these changes make estimation of a concept of 'income' which is comparable between 1987 and 1989 problematic.

The cash benefits identified in our data comprise pensions (68% of the total in 1987, 61% in 1989, based on the Household Budget Survey Panel, discussed in Section 4), family allowances (18% in 1987, 23% in 1989), and a number of small items ('social aid', child allowance, sick pay, educational aid, and our imputed housing subsidy); for further details see van de Walle et al. (1994). There was an overall real increase in social expenditures during the period 1987–1989; comparing the aggregate receipts implied by the Household Budget Survey Panel, total real spending increased by 21% (van de Walle et al., 1994). All categories increased, though the gains were proportionately greater for family allowances, which increased by 49% in real terms.

4. Empirical implementation

4.1. *The Household Budget Survey*⁹

Given the considerations of Section 2, and the specific policy changes that occurred during this period (Section 3), there are compelling arguments for basing poverty comparisons on consumption rather than income. We shall use the consumption data collected by the Household Budget Survey (HBS) conducted by the Central Statistical Office (CSO) for two years, 1987 and 1989, converted to constant 1989 prices using the monthly CPI. The surveys, held every two years, follow a sampling procedure in which two-thirds of sampled households are retained for re-sampling from one survey to the next. This HBS feature has here been exploited to create a panel of households, with 5,945 households tracked over the two years.¹⁰

The basic unit of observation in the HBS is the household which may contain more than one family unit. The sample frame is based on the 1980 census, and comprises all Hungarian citizens living in private households in the country (until 1989), excluding households which had a member classified as 'self-employed'.¹¹ The survey also excludes persons living in

⁹ Van de Walle et al. (1994) document further details on the survey (sample frame, sample stratification, interviews, etc.); here we only summarize the salient features relevant to our enquiry.

¹⁰ In theory, of the 12,000 households sampled in each survey date, panel rotation should allow a complete panel of 8000 households. In addition to the usual sample attrition (due to migration, non-response, etc.), the introduction of the category 'self-employed' in the target population in 1989 necessitated a reduction in the re-sampling of the usual two-thirds for that year.

¹¹ In 1989, the HBS was modified to include people belonging to the 'self-employed' class, given the dramatic increase in self-employment activities. Unfortunately, since we want to make use of the panel feature of the surveys, we cannot take advantage of this sample broadening.

institutions (retirement homes, children's homes, etc.), the homeless, Hungarian citizens living abroad and foreign citizens living in Hungary. It is not clear what biases, if any, in our results can be attributed to these restrictions on the sample frame. CSO statisticians designed a detailed set of inverse sampling rates to remove biases in the 1987 round of the panel, though non-random attrition will leave biases in the 1989 round; little can be done about this in our data.

Given the survey technique, we expected that the consumption data in the HBS would be of high quality. Data collection for the HBS was carried out in a three-stage interview process. In the first stage, households were required to maintain a diary for a period of two months in which they recorded daily purchases of consumption items (both quantity and value in current prices), incomes from all sources (except investment income), weekly consumption from own production (both quantity and value in local current prices), household demographic data, data on 'housing conditions', and data on owned plots of land, if any. This two-month diary stage is evenly distributed through the survey year for different households, with one-sixth of the sample keeping diaries in each of the six two-month periods. The second stage consisted of an interview two months after the completion of the diary stage. At this time data were collected on a recall basis on all household income in the previous month and certain medium frequency purchases (e.g. clothing) over the preceding 4 months. The final stage was an interview at the end of the year to collect data, again on a recall basis, on current stocks of consumer durables, expenditures on major consumer durables, construction and real estate activity, and net incomes from agriculture during the preceding year.

In 1987 the diary stage accounted, on average, for two-thirds of all household expenditures; a major proportion of the remaining one-third was accounted for by the second stage, while items collected in the third stage of interviewing made up the remainder. The survey's primary aim was collecting data to analyze consumption and related activities for various socio-economic population groups. It includes exhaustive information on cash benefits. The survey does not contain data on wage rates or hours worked. Other than employment income, the data set contains only qualitative information on earned income in the form of occupation codes.

The consumption data comprise total household expenditure on goods and services, including reported housing expenditures and the value of consumption from own production. However, housing expenditures are missing for about one-third of the sample (it appears that rents for owner-occupied dwellings have not been imputed). They also exhibit extreme variance reflecting the housing market situation in Hungary. Thus the data need to be supplemented to get (imputed) housing expenditures for missing observations, and also to get better estimates for observations possibly suffering from measurement errors. To deal with both problems, we have

replaced reported housing expenditure by estimates from hedonic housing expenditure relationships estimated for each year. Only physical housing attributes are used as regressors in a linear relationship. Regressors include dummy variables for the geographical region to which the household belongs. The 23 regions of Hungary are further subdivided into urban and rural sub-regions. Budapest, which is a large and entirely urban region, is divided into 22 sub-districts. Other variables include dummy variables for whether the dwelling is a house, whether it is government-owned or private, the number of rooms, type of heating, type of bath/flush facilities, flat size, whether it has running water, and whether it has piped gas. Also included are dummy variables for the month of interview to control for unobservable within year structural changes, such as seasonal inflation, and changes in government policies. The details of the estimation procedure and the results are given in van de Walle et al. (1994). On the basis of these new estimates of housing expenditure conditional on the physical characteristics of the dwelling, total household consumption is re-calculated. Subsequent analyses will use this new household consumption estimates. The effects of this procedure on key variables of interest are discussed in van de Walle et al. (1994).

The existence of a large subsidized public housing sector poses a further problem. A dummy variable for government housing was included in the hedonic regressions, and had a (highly) significant negative sign in both years. We assume that this reflects an implicit subsidy through controlled rents in public housing. We thus estimated the housing expenditure for each household as though the dwelling had been obtained on the private rental market. The predicted private-market-equivalent housing expenditure is then the value used to derive total household consumption for each observation. The difference between the private-market-equivalent expenditure and the value obtained by setting the government dummy parameter equal to its estimated value gives the subsidy associated with government-provided housing. In addition to dealing with the missing values, this procedure goes some way toward eliminating measurement error associated with reported housing expenditures.

4.2. The poverty lines

The poverty lines used here were provided by the CSO, and have three components: the cost of a subsistence food basket, 'other expenses', and an allowance for housing costs. Subsistence food expenditures are derived based on a weekly menu for each season costed at average national quarterly prices and aggregated to get the minimum annual food expenditures. These minima are derived for three groups of individuals—children, active adults, and pensioners—and allow for the differential food require-

ments of each group as prescribed by the National Research Institute of Dietetics, Hungary. Next, these households were differentiated into various groups according to their location of residence and demographic characteristics (rural–urban, active–pensioner, and household size and composition). The 1989 HBS is then used to locate households with food expenditures in the range 20% above or below the subsistence food spending for their demographic group (excluding households reporting expenditures which are non-typical of households at the poverty level, i.e. purchases of houses, flats, or cars). The poverty line's other components—'other expenses', i.e. expenses other than food and housing, and 'housing expenses'—are then calculated based on the actual expenditure level of these 'reference' households. See van de Walle et al. (1994) for details of the poverty lines.

The CSO poverty line embodies both scale economies (particularly for housing) and differences in consumption needs between three groups: active adults, children, and pensioners. The equivalence scale embodied in the CSO poverty lines may be critical to inferences on the performance of social expenditures, such as family allowances and pensions.¹² Consider children: the CSO poverty lines have an elasticity of about 0.6 to an increase in the number of children from around zero to three in a family with two adults (van de Walle et al., 1994). This is sufficiently high for it to be true that poorer households (in terms of their welfare ratio) tend to have more children (van de Walle et al., 1994). Naturally, then, a family allowance scheme may seem to be well-targeted to the poor. That is a valid conclusion, as long as one accepts the structure of the CSO poverty lines.

We performed one test of the equivalence scale implicit in the CSO poverty lines. This was based on the Engel method of setting scales, whereby the budget share devoted to food is regressed on total expenditure and a set of variables describing the demographic composition of the household (Deaton and Muellbauer, 1986). Our test entailed regressing the food share on the log of total expenditure, the log of the CSO poverty line, and household size; if the latter were significant, then the Engel method would imply a different set of scales to those implicit in the CSO poverty line. However, household size had no significant effect on the food share controlling for the CSO poverty line as well as household expenditure (van de Walle et al., 1994). This test is not conclusive (given the well-known problems of identifying scales from demand behavior¹³), but it does not suggest that the CSO poverty lines would have to be revised to be consistent

¹² In the context of the family allowance in Hungary, see Jarvis and Micklewright (1994). For a general discussion of how equivalence-scale parameters affect poverty comparisons, see Lanjouw and Ravallion (1994).

¹³ See, for example, Deaton and Muellbauer (1986), Pollak (1991), and Browning's (1992) survey.

Table 1
Welfare ratio distributions

Welfare ratio (% of poverty line)	Cumulative percentage of population			
	(1) 1987	(2) 1989	(3) Mean of (1) and (2)	(4) Mean Welfare ratio ^a
75	3.70	5.29	4.50	3.01
100	17.24	20.50	18.87	15.50
125	38.66	41.94	40.30	36.02
150	57.54	61.50	59.62	56.69
175	71.46	74.40	72.93	71.98
200	81.11	83.14	82.13	82.16
225	87.00	88.50	87.75	88.34
250	91.22	91.80	91.51	92.41

^a Persons ranked by the time-mean welfare ratio.

with the household-size elasticity of a set of scales derived by the Engel method.

4.3. Changes in poverty over the period

The marginal distribution functions of persons ranked by their household welfare ratios for each date are given in Table 1.¹⁴ First-order dominance is indicated, implying an unambiguous increase in consumption poverty; this holds for all poverty lines and poverty measures within a broad class (Atkinson, 1987).¹⁵ Table 1 [column (4)] also gives the marginal distributions based on the two-year mean welfare ratios, i.e. individuals are ranked by the two-year mean of their household welfare ratios.¹⁶ There is less poverty in this distribution than for either year on its own, up to about 150% of the poverty line. And there is less poverty in the distribution based on mean welfare ratios than in the mean of the two marginals based on current-year's welfare ratio up to almost twice the poverty line; thus (for these data) the variability in living standards over time tends to increase measured poverty.¹⁷

¹⁴ The poverty profile and static incidence of cash benefits are described in van de Walle et al. (1994).

¹⁵ This also holds when one considers rural and urban areas separately, and when Budapest is separated from other urban areas. Thus the conclusion that poverty had increased is also robust to measurement error in the poverty line differentials between urban and rural areas.

¹⁶ Using the 1987 household sizes, though the difference using the 1989 household sizes is negligible.

¹⁷ This is an empirical property; on the conditions for it to hold, see Ravallion (1988).

5. Policy simulations

5.1. The base-line joint distribution

Table 2 gives selected points on the base-line joint distribution and corresponding transition matrix over the two dates; each cell gives the percentage of the total population who were in that row's welfare group in 1987 and that column's group in 1989, while the number in square brackets is the transition probability (proportion of each row's total population who were in each column's welfare ratio group in 1989). Thus, for example, 4.18% of people lived in households with a consumption less than the poverty line in 1987 and were between 100% and 125% of the poverty line

Table 2
Base-line joint distributions and transition matrix

	<100	100–125	125–150	150–200	200+	Total 1987 (cumulative)
<100	9.66 ^a [56.07]	4.18 [24.26]	1.67 [9.69]	1.14 [6.62]	0.58 [3.37]	17.23 (17.23) [100.00]
100–125	6.25 [29.18]	6.38 [29.79]	4.10 [19.14]	3.47 [16.20]	1.22 [5.70]	21.42 (38.65) [100.00]
125–150	2.40 [12.71]	5.54 [29.34]	4.78 [25.32]	4.01 [21.24]	2.15 [11.39]	18.88 (57.53) [100.00]
150–200	1.22 [5.18]	3.38 [14.34]	5.66 [24.01]	7.88 [33.43]	5.43 [23.04]	23.57 (81.10) [100.00]
200+	0.42 [2.22]	1.54 [8.15]	3.30 [17.46]	5.61 [29.68]	8.03 [42.49]	18.90 (100.00) [100.00]
Total 1989 (cumulative)	19.94 (19.94)	21.02 (40.96)	19.52 (60.48)	22.11 (82.59)	17.42 (100.00)	100.00

Note: the table gives the percentage of the total population (represented by the panel sample) in the 1987 welfare-ratio group of each row, and the 1989 group of each column. The figure in brackets below each of these percentages is the corresponding 'transition probability', giving the percentage of those in the 1987 group of a given row who are found in the group of each column in 1989. The number in parentheses in the column and row totals are the points on the (marginal) cumulative distributions for each year.

^a Decomposition of persistently poor:

	<80	80–90	90–100
<80	2.66	0.72	0.94
80–90	1.07	0.82	0.75
90–100	0.88	0.94	0.88

in 1989. The column and row totals are simply the marginal welfare-ratio distributions. A footnote to the table also gives a decomposition of the proportion found to be persistently poor, according to the welfare ratios in each year.

There was considerable transient poverty over the period. While 17% of people consumed less than the poverty line in 1987, and 20% in 1989, only 10% were poor at both dates. There was also considerable variability amongst the persistently poor. Still, the people who were poor in 1989 came mainly from those who were consuming less than 150% of the poverty line in 1987, while few of those who escaped poverty between the two dates got farther than 125% of the poverty line.

There are various measures of mobility.¹⁸ A common measure is the correlation coefficient, which is 0.431 between welfare ratios in 1987 and 1989.¹⁹ An alternative measure with some advantages is that proposed by Shorrocks (1978b), based on a comparison of the inequality measures using the two year means with those for each year separately. The Gini index for 1987 and 1989 welfare ratios are 0.227 and 0.229, respectively, while that for the distribution of two-year mean welfare ratios is 0.203. The Shorrocks rigidity index based on these Gini indices is 0.92. However, we know of no comparable estimates for other countries for either of these measures (all of the estimates that we know of are for earnings data; see the survey in Atkinson et al., 1992).

5.2. Simulated distributions

What contribution did the changes in cash benefits over the period make to the joint distribution? To answer this question, we must simulate the counter-factual distributions, without any change in cash benefits. We do this under various assumptions about possible behavioral responses. Initially we assume that pre-reform consumptions are unchanged and that all increments to social incomes are consumed. While this is a natural ‘benchmark’—and is typical of static incidence calculations—later in the paper we consider the implications of relaxing it.

For each simulation we calculate the *PROT* and *PROM* tests described in Subsection 2.1. In the notation of that section, the *F* distribution is that

¹⁸ Shorrocks (1978a) proposes a set of axioms for measuring mobility, and discusses their consistency, and the properties of various measures used in practice. For a recent overview of the issues, see Atkinson et al. (1992).

¹⁹ The OLS regression is (*t*-ratios in parentheses):

$$x_{89} = 0.814 + 0.412 \cdot x_{87} \\ (42.33) \quad (36.83).$$

represented by the base-line distribution. Thus a positive value of *PROT* implies that the actual changes over the period protected the poor relative to the simulated alternative (the *G* distribution), and similarly for *PROM*. We do both tests at the CSO poverty line, and a poverty line set 25% higher.

To first assess the impact of the changes in cash benefits over the period, the joint distribution in Table 3 is simulated under the assumption that cash benefits did not change between the two dates for any household. A comparison of Tables 2 and 3 thus indicates how the changes which actually occurred (as reflected in the actual distribution in Table 2) affected the joint distribution of welfare ratios.

Looking first at the impact on the marginal distribution, we see that there is first-order dominance between the distribution of welfare ratios that we predict would have occurred in 1989 without any changes in cash benefits and that which actually occurred; this can be seen by comparing the cumulative totals in the last row of Table 3 with those in the last row of Table 2. Without the change in cash benefits there would have been higher

Table 3
 No change in cash benefits between 1987 to 1989
PROT(100) = 6.62(10.53); *PROT*(125) = 5.36(7.75); *PROM*(100) = 1.02(2.17); *PROM*(125) = 0.48(0.81)^a

	<100	100–125	125–150	150–200	200+	Actual total, 1987
<100	10.68 ^b	3.17	1.58	1.12	0.68	17.23
100–125	8.02	5.08	3.47	3.48	1.36	21.42
125–150	4.26	4.85	3.92	3.64	2.21	18.88
150–200	3.09	3.97	4.73	6.61	5.18	23.57
200+	1.54	2.15	2.53	5.11	7.58	18.90
Simulated 1989 cumulative	27.58 (27.58)	19.22 (46.80)	16.23 (63.03)	19.96 (82.99)	17.01 (100.00)	100.0

^a z-scores in parentheses; critical values: 1.96 (2.58) at the 5% (1%) level.

^b Decomposition of persistently poor:

	<80	80–90	90–100
<80	3.06	0.75	0.55
80–90	1.87	0.59	0.54
0–100	1.71	0.79	0.83

poverty in 1989 than actually observed; this holds wherever one draws the poverty line, or which poverty measure is used.

We find that, if there had been no change in cash benefits over the period, there would have been an extra 7.6% of the population consuming less than the poverty line by 1989. And the bulk ($PROT(100) = 6.6\%$) of this increment would have been due to non-poor people in 1987 falling into poverty by 1989; from Table 3, it can be seen that 16.9% of the population would have fallen into poverty by 1989 if there had been no change in cash benefits, whereas (from Table 2) the actual percentage was 10.3% falling into poverty. While protecting 6.6% from poverty, the changes in cash benefits only allowed 1.0% to actually escape poverty. Nonetheless, both the *PROT* and *PROM* tests indicate that the actual changes in cash benefits were pro-poor, though only for *PROT* are the differences statistically significant.²⁰

In principle, this positive impact in protecting people from poverty could be due at least in some part to the fact that the average cash benefit increased, rather than to the way the distribution of that increase occurred. The latter can be thought of as the ‘targeting’ of changes in cash benefits (though, as is invariably the case, the changes in distribution presumably reflected both the decisions of participants as well as policy reforms by the government). Thus it is also of interest to ask: What would the outcome have been if the actual increase in mean cash benefit had been equally distributed to all persons?

We give that simulation in Table 4. Here we take the actual increase in aggregate cash benefit, allocate it equally to all persons, and re-calculate the joint distribution and poverty rates. We find that while 19.9% of the population consumed less than the poverty line in 1989, it would have been 22.2% if the increase in cash benefits had been equally distributed. This increase is statistically significant (z score = 3.07). The transitions are also significantly different; while 10.3% of the population fell into poverty by 1989, the proportion would have been 13.0% with equally distributed gains in cash benefits ($z = 4.56$). However, while protection of the poor is evident in this case, there is little difference in the extent to which people escaped poverty.

5.3. Testing for behavioral responses

To test the robustness of these results to possible behavioral responses, we have estimated a model of the first difference of consumption over the two dates as a function of the other variables (listed in the appendix, which gives the means and standard deviations of the differences). In addition to

²⁰ On estimating the standard errors for *PROT* and *PROM*, see van de Walle et al. (1994).

Table 4

An equal gain in cash benefits from 1987 to 1989 set at the mean

 $PROT(100) = 2.68(4.56)$; $PROT(125) = 1.60(2.42)$; $PROM(100) = -0.38(0.77)$; $PROM(125) = -1.28(2.09)^a$

	<100	100–125	125–150	150–200	200+	Actual total, 1987
<100	9.28 ^b	4.15	1.77	1.24	0.78	17.23
100–125	6.22	5.54	3.80	4.32	1.52	21.42
125–150	3.23	4.32	4.54	4.23	2.57	18.88
150–200	2.31	3.31	4.65	7.29	6.02	23.57
200+	1.20	1.70	2.36	5.01	8.63	18.90
Simulated 1989 (cumulative)	2.24 (22.24)	19.02 (41.28)	17.12 (58.40)	22.08 (80.48)	19.54 (100.00)	100.00

^a z-scores in parentheses; critical values: 1.96 (2.58) at the 5% (1%) level.^b Decomposition of persistently poor:

	<80	80–90	90–100
<80	2.57	0.59	0.80
80–90	1.40	0.61	0.62
90–100	1.09	0.83	0.77

cash benefits, we have chosen a set of explanatory variables for changes in the demographic composition of the household, physical and human assets, and occupations, as well as dummy variables for the interview month. However, the key coefficient for our purposes is that on the change in cash benefit; the OLS estimate of the *PCSI* is 0.43 ($t = 10.4$). (The full regression results are given in the appendix.)

Treating social income as exogenous may be questioned. In Hungary, take-up rates are very high; 89% of households in 1987 received some form of social income, while this was true of 94% in 1989. Nonetheless, we also tried a specification in which the 1989 social income was treated as endogenous, with all other right-hand-side variables included in the set of instruments, which also included 1987 values of a number of various other variables (demographic and occupational variables and a dummy variable for sickness) for identification. Depending on the precise set of instruments, our IV estimates of the *PCSI* ranged from 0.35 ($t = 2.9$) to 0.56 ($t = 4.0$). The OLS estimate is about at the middle of this range.

We also tried adding the 1987 values of all right-hand-side variables (in effect relaxing the homogeneity restriction implicit in the first-difference

specification). Some of these were mildly significant, but the key coefficient for our purposes was little affected; the OLS estimate of *PCSI* in the augmented model was 0.49 ($t = 10.7$).

If the *PCSI* is correlated with the level of cash benefits or with the household characteristics which are used to target the cash benefits, then this could bias our results. To test this possibility we re-estimated the *PCSI* from the first-difference model of consumption by stratifying the sample according to whether cash benefits in 1987 were above or below the median; the estimates were 0.41 ($t = 7.04$) and 0.55 ($t = 8.77$), respectively. Thus it is not the case that recipients of large cash benefits tend to have a higher *PCSI*. When we stratified instead by consumption per capita, those above the median had a *PCSI* of 0.41 ($t = 6.37$), while for those below the median it was 0.40 ($t = 8.05$); when stratified by the 1987 welfare ratio the difference was slightly greater: 0.46 ($t = 8.86$) for those below the median, and 0.38 ($t = 6.43$) for those above the median. Poorer households do have a higher propensity to consume out of cash benefits, but the difference is not large.

However, when we stratified by demographic variables, some large differences in the *PCSI* emerged. For households larger than three (the median) the estimated *PCSI* was 0.91 ($t = 10.86$), while for those of size three or less it was -0.06 ($t = 0.12$). On probing further, the difference is found to be correlated strongly with the number of children. For households made up of adults only, the estimate of *PCSI* was -0.09 ($t = 2.17$). For those with one child it was 0.15 ($t = 1.00$). However, for those with two or more children, the *PCSI* was 1.05 ($t = 11.7$). It appears then that families with two or more children tend to consume all of an increment in cash benefits, while others save it. It is not clear why this would happen. We will do simulations with and without this demographic stratification in the *PCSI*.

5.4. Simulated distributions with behavioral responses

Our aim here is to test how important behavioral responses may be to assessments of the performance of the social safety net. The results of the previous subsection suggest quite strong responses; on average, a little more than half of the current gross gain from an increment to cash benefits is dissipated through those responses. However, we do not identify what form those responses take, and there remains the strong possibility that a good deal of an increment to social income is being saved, so that the welfare gains are at another date. Here we are only able to identify impacts on current living standards.

Using estimates of the *PCSI*, we can simulate the joint distributions under alternative assumptions about how cash benefits changed by estimating the change in consumption (and hence the 1989 consumption) for each house-

Table 5

No change in cash benefits between 1987 to 1989 ($PCSI = 0.43$) $PROT(100) = 2.72(4.62)$; $PROT(125) = 2.01(3.10)$; $PROM(100) = 0.58(1.22)$; $PROM(125) = 0.34(0.57)^a$

	<100	100–125	125–150	150–200	200+	Actual total, 1987
<100	10.24 ^b	3.70	1.67	1.05	0.58	17.23
100–125	7.32	5.55	3.95	3.38	1.22	21.42
125–150	3.11	5.29	4.70	3.72	2.06	18.88
150–200	1.84	3.66	5.66	7.20	5.21	23.57
200+	0.73	1.93	3.01	5.38	7.86	18.90
Simulated 1989 (cumulative)	23.24 (23.24)	20.12 (43.36)	18.98 (62.34)	20.73 (83.07)	16.94 (100.00)	100.00

^a z-scores in parentheses; critical values: 1.96 (2.58) at the 5% (1%) level.^b Decomposition of persistently poor:

	<80	80–90	90–100
<80	2.89	0.71	0.73
80–90	1.53	0.65	0.72
90–100	1.15	1.04	0.82

hold for each hypothetical change in social income, leaving all other variables at their data values (including the regression residuals).

The joint distribution in Table 5 is simulated under the assumption that the $PCSI$ is 0.43 for all households and that cash benefits did not change between the two dates for any household. A comparison of Tables 2 and 5 thus indicates how the changes that actually occurred (as reflected in the actual distribution in Table 2) affected both actual poverty incidence (the column totals) and the transitions across groups under this assumption on behavioral responses.

As one would expect, the incorporation of such behavioral responses has noticeably changed the joint distribution (comparing Tables 3 and 5), though a number of the qualitative conclusions are robust. Again we see that there is first-order dominance between the distribution of welfare ratios that we predict would have occurred in 1989 without any changes in cash benefits and that which actually occurred (comparing the cumulative totals in the last row of Table 5 with those in the last row of Table 2). Without the change in cash benefits there would have been higher poverty in 1989 than was actually observed; this holds wherever one draws the poverty line, and

no matter which poverty measure is used. If there had been no change in cash benefits over the period, there would have been an extra 3.3% of the population consuming less than the poverty line by 1989. And virtually all ($PROT(100) = 2.7\%$) of this increment would have been due to non-poor people in 1987 falling into poverty by 1989; from Table 5 it can be seen that 13.0% of the population would have fallen into poverty by 1989 if there had been no change in cash benefits, whereas (from Table 2) the actual percentage was 10.3% falling into poverty. While protecting 2.7% from poverty, the changes in cash benefits only allowed 0.6% to actually escape poverty, and virtually all of these (0.5%) got no further than 125% of the poverty line. Nonetheless, both the *PROT* and *PROM* tests indicate that the actual changes in cash benefits were pro-poor allowing for this behavioral response, though only for *PROT* are the differences statistically significant.

On taking the actual increase in aggregate cash benefit and allocating it equally to all persons (analogously to Table 4) one obtains a result that is very similar to Table 2. This is given in Table 6. For example, while 19.9% of the population consumed less than the poverty line in 1989, we estimate

Table 6

An equal gain in cash benefits from 1987 to 1989 set at the mean ($PCSI = 0.43$)
 $PROT(100) = 0.61(1.08)$; $PROT(125) = -0.16(-0.25)$; $PROM(100) = -0.11(0.23)$;
 $PROM(125) = -0.70(1.15)^a$

	<100	100–125	125–150	150–200	200+	Actual total, 1987
<100	9.55 ^b	4.13	1.68	1.22	0.65	17.23
100–125	6.16	5.93	4.35	3.69	1.28	21.42
125–150	2.65	4.98	4.89	4.19	2.17	18.88
150–200	1.44	3.32	5.40	7.75	5.67	23.57
200+	0.64	1.49	3.16	5.34	8.27	18.90
Simulated 1989 (cumulative)	20.44 (20.44)	19.86 (40.30)	19.48 (59.78)	22.18 (81.96)	18.04 (100.00)	100.00

^a z-scores in parentheses; critical values: 1.96 (2.58) at the 5% (1%) level.

^b Decomposition of persistently poor:

	<80	80–90	90–100
<80	2.67	0.69	0.82
80–90	1.22	0.72	0.68
90–100	0.91	0.94	0.90

that it would have been 20.4% if the increase in cash benefits had been equally distributed. This increase cannot be considered statistically significant ($z = 0.68$). Furthermore, there is no longer first-order dominance; if one used a poverty line set 25% higher, then one would conclude that the equally distributed allocation would have achieved a slightly lower poverty rate, 40.3% instead of 41.0%, though again this difference is not statistically significant ($z = 0.77$). The transitions are also very similar; for example, while 10.3% of the population fell into poverty by 1989, the proportion would have been only slightly higher (10.9%) with equally distributed gains in cash benefits. None of the *PROT* or *PROM* tests is statistically significant. Thus, on introducing these behavioral responses, we find that the safety net's ability to protect the poor was largely attributable to the increase in mean cash benefits; there was little impact (one way or another) from changes in how those gains were targeted.

The results of Subsection 5.3 also suggested that there may be an appreciable difference in behavioral responses between different demographic groups. To see what difference in consumption behavior could have on our assessment of the performance of the social safety net, we re-estimated the joint distributions assuming a *PCSI* of unity for households with two or more children, and zero otherwise. The simulations of the effect of not changing cash benefits were virtually identical to those obtained with a constant *PCSI* (Table 5); detailed results are given in van de Walle (1994). We also did analogous simulations to Tables 4 and 6 under this alternative assumption on behavioral responses, and the results were not appreciably different from Table 2, though with a slightly stronger sign of protection for the poor ($PROT(100) = 0.97\%$, $t = 1.71$); see van de Walle et al. (1994) for details. On the whole, the conclusion that the actual changes in the safety net quite closely approximated the joint distribution of consumption that would have been obtained if the gains had been equally distributed is robust to these alternative assumptions on behavioral responses.

6. Conclusions

In testing the performance of a social safety net, static methods of evaluating benefit incidence have the drawback that no distinction is made between a safety net's ability to protect people from poverty, versus promote people from poverty. We have proposed a straightforward alternative method of quantifying the distributional impacts of a safety net when suitable panel data are available. The transitions of households in and out of poverty over time are simulated under policy changes, with and without data-consistent behavioral responses. Tests are then performed on the

alternative joint distributions to assess whether there is a significant difference in the extent of protection and/or promotion of the poor.

The approach has been implemented on new data for Hungary, 1987–1989, during which period changes in cash benefits were used to help compensate households for a number of the policy reforms implemented, in the context of an economy in transition. Our results suggest that:

(i) There was an increase in consumption-poverty over this period; that claim is robust to the choice of poverty measure or line. There was also considerable transient poverty; roughly half of the persons who were living in poor households in 1989 had not been doing so in 1987. And roughly four out of ten persons who had been poor in 1987 escaped poverty by 1989.

(ii) The gains in social incomes were markedly pro-poor. Without any changes in cash benefits, and ignoring behavioral responses, the poverty rate would have been 6.6 percentage points higher by 1989 than actually observed. This was mainly achieved by preventing households from falling into poverty; far fewer escaped poverty by this means. Thus the changes in the safety net were better at reducing transient poverty than persistent poverty.

(iii) There is evidence of behavioral responses to changes in cash benefits; our estimates suggest that on average about half of an increment in cash benefits received is passed on to current consumption. Nonetheless, the changes in cash benefits were still markedly current-poverty reducing. Incorporating our estimate of the propensity to consume out of cash benefits, we estimate that the poverty rate would have been three percentage points higher by 1989 than actually observed if cash benefits had not changed. Again, this was mainly achieved by preventing households from falling into poverty; far fewer escaped poverty by this means.

(iv) The reduction in transient poverty was due in large part to the increase in mean cash benefits rather than improved targeting. Indeed, allowing for behavioral responses, the rates of poverty, and the transitions into and out of poverty, would have been virtually identical if the same increase in cash benefits over the period had instead been equally distributed.

Acknowledgements

These are the views of the authors, and should not be attributed to the World Bank. We are grateful to staff of the Hungarian Central Statistical Office for providing access to the data used here, and for helping us with our many questions on the data. In particular, we wish to thank Gaspar Fajth, Judit Salamin and Peter Szivos. The journal's editor and referees made a number of helpful comments. In the course of the study we had useful

discussions with Emmanuel Jimenez, including comments on this paper. We are also grateful to seminar participants at the World Bank's Resident Mission, Budapest.

Appendix: Fixed effects model of consumption

Variable	Mean	Std. dev.	Coefficient	t-ratio
Total expenditure	-16303.35	134534.2	-1	
Intercept	1	0	-10199	-2.369**
Total cash benefit	11382.93	44092.33	0.4311	10.382**
Number of male adults aged 19–59	-0.032	0.41	61430	13.871**
Number of female adults aged 19–54	-0.045	0.37	58067	11.277**
Number of males aged 60 and over	0.007	0.24	13016	1.702*
Number of females aged 55 and over	0.025	0.26	27727	3.827**
Number of persons aged 15–18	0.017	0.41	37376	7.204**
Number of persons aged 6–14	-0.053	0.44	31188	6.031**
Number of persons aged 3–5	-0.016	0.31	23609	3.403**
Number of persons aged 0–2	-0.024	0.28	-1202	-0.171
Interview month: March–April [†]	0.164	0.37	-3351	-0.566
Interview month: May–June [†]	0.158	0.37	2326	0.389
Interview month: July–August [†]	0.164	0.37	-21592	-3.644**
Interview month: September–October [†]	0.164	0.37	-7998	-1.350
Interview month: November–December [†]	0.194	0.39	-581	-0.102
Whether own plot of land [†]	0.062	0.42	5714	1.419
Whether own place of dwelling [†]	0.011	0.15	-16001	-1.419
Education of head: elementary school [†]	-0.020	0.29	-1209	-0.139
Education of head: vocational school [†]	0.035	0.26	-8823	-0.844
Education of head: secondary school [†]	-0.010	0.24	19623	1.672*
Education of head: college/graduate [†]	-0.059	0.13	26906	1.657*
Occupation of head: leader/manager [†]	-0.003	0.21	20684	0.671
Occupation of head: white-collar worker [†]	-0.002	0.24	-450	-0.015
Occupation of head: skilled labor [†]	-0.011	0.32	4605	0.156
Occupation of head: unskilled labor [†]	0.006	0.25	8628	0.292
Occupation of head: Self-employed/farmer [†]	0.009	0.09	-19925	-0.590

Notes: 1. All variables except the month of interview are first differences (1989–1987). 2. Observations are weighted by their expansion factors for statistical analysis. 3. Variables marked with a dagger (†) are dummy variables for the respective categories. 4. ** indicates significance at the 5% level and * indicates significance at the 10% level.

References

- Adelman, I., K. Subbarao and P. Vashishtha, 1985, Some dynamic aspects of rural poverty in India, *Economic and Political Weekly* 20, A103–A117.

- Atkinson, A.B., 1987, On the measurement of poverty, *Econometrica* 55, 749–764.
- Atkinson, A.B., 1991, Comparing poverty rates internationally: Lessons from recent studies in developed countries, *World Bank Economic Review* 5, 3–22.
- Atkinson, A.B., F. Bourguignon and C. Morrisson, 1992, Empirical studies of earnings mobility (Harwood Academic Press, Chur, Switzerland).
- Atkinson, A.B. and J. Micklewright, 1992, Economic transformation in Eastern Europe and the distribution of income (Cambridge University Press, Cambridge).
- Bane, M.J. and D.T. Ellwood, 1986, Slipping into and out of poverty, the dynamics of spells, *Journal of Human Resources* 21, 1–23.
- Barr, N., 1992, Income transfers and the social safety net in Russia. Studies of economies in transition, Paper 4 (World Bank, Washington, DC).
- Blackorby, C. and D. Donaldson, 1987, Welfare ratios and distributionally sensitive cost-benefit analysis, *Journal of Public Economics* 34, 265–290.
- Browning, M., 1992, Children and household economic behavior, *Journal of Economic Literature* 30, 1434–1475.
- Chaudhuri, S. and M. Ravallion, 1994, How well do static indicators identify the chronically poor?, *Journal of Public Economics* 53, 367–394.
- Deaton, A., 1994, Data and econometric tools for development economics, in: J. Behrman and T.N. Srinivasan, eds., *Handbook of development economics*, vol. 3 (North-Holland, Amsterdam).
- Deaton, A. and J. Muellbauer, 1986, On measuring child costs: With applications to poor countries, *Journal of Political Economy* 94 720–744.
- Drèze, J. and A. Sen, 1989, *Hunger and public action* (Oxford University Press, Oxford).
- Fajth, G. and L. Vita, 1992, Income trends and social policy developments in Hungary, Paper prepared for the Luxembourg Income Study Meeting, July 1992, Luxembourg Income Study, Luxembourg.
- Hsiao, C., 1986, *Analysis of panel data* (Cambridge University Press, Cambridge).
- Jarvis, S. and J. Micklewright, 1994, The targeting of family allowance in Hungary, in: D. van de Walle and K. Nead, eds., *Public spending and the poor: Theory and evidence* (Baltimore and London, The Johns Hopkins University Press, 1995, forthcoming).
- Lanjouw, P. and M. Ravallion, 1994, Poverty and household size, mimeo, Policy Research Department, World Bank.
- Milanovic, B., 1994, Distributional incidence of cash and in-kind transfers in Eastern Europe and Russia, in: D. van de Walle and K. Nead, eds., *Public spending and the poor: Theory and evidence* (Baltimore and London, The Johns Hopkins University Press, 1995, forthcoming).
- Pollak R.A., 1991, Welfare comparisons and situation comparisons, *Journal of Econometrics* 50, 31–48.
- Ravallion, M., 1988, Expected poverty under risk-induced welfare variability, *The Economic Journal* 98, 1171–1182.
- Ravallion, M., 1994, Poverty comparisons. Fundamentals in pure and applied economics, vol. 56 (Harwood Academic Press, Chur, Switzerland).
- Ravallion, M. and G. Datt, 1994, Is targeting through a work requirement efficient?, in: D. van de Walle and K. Nead, eds., *Public spending and the poor: Theory and evidence* (Baltimore and London, The Johns Hopkins University Press, 1995, forthcoming).
- Ruggles, P. and R. Williams, 1989, Longitudinal measures of poverty: Accounting for income and assets over time, *Review of Income and Wealth* 35, 225–243.
- Shorrocks, A.F., 1978a, The measurement of mobility, *Econometrica* 46, 1013–1024.
- Shorrocks, A.F., 1978b, Income inequality and income mobility, *Journal of Economic Theory* 19, 376–393.

- Szalai, J., 1989, Poverty in Hungary during the period of economic crisis, Background paper for the World Development Report 1990, World Bank, Washington, DC.
- van de Walle, D., M. Ravallion and M. Gautam, 1994, How well does the social safety net work? The incidence of cash benefits in Hungary 1987–89, LSMS Working Paper 102, World Bank, Washington, DC.
- World Bank, 1992, Hungary: Reform of social policy and expenditures. World Bank country study (World Bank, Washington, DC).