JOURNAL OF COMPARATIVE ECONOMICS **26,** 338–357 (1998)
ARTICLE NO. JE981526

Transient Poverty in Postreform Rural China[1]

Jyotsna Jalan and Martin Ravallion*

World Bank, Washington, DC, 20433

Received April 17, 1997; revised March 3, 1998

Jalan, Jyotsna and Ravallion, Martin—Transient Poverty in Postreform Rural China

In a new, relatively high-quality, panel-data set on household consumptions for postreform rural China we find that consumption variability accounts for a large share of observed poverty and is likely to be a severe constraint on efforts to reach the long-term poor. Half of the mean squared poverty gap and over a third of the mean poverty gap is transient, in that it is directly attributable to year-to-year fluctuations in consumption. There is enough transient poverty to treble the cost of eliminating chronic poverty when targeting transfers according to current consumptions and to tilt the balance in favor of untargeted transfers. Anti-poverty policies in China may have to give greater emphasis to the problem of transient poverty. *J. Comp. Econom.*, June 1998, **26**(2), pp. 338–357. World Bank Washington, DC. © 1998 Academic Press

Journal of Economic Literature Classification Numbers: I32, O53, Q12.

1. INTRODUCTION

Some of the poverty observed at any one point in time is bound to be transient, in that it is due to a temporary shortfall in consumption. Let us define ''transient poverty'' as the poverty that can be attributed to intertemporal variability in consumption. The degree of such transient poverty could well be high in underdeveloped rural economies, where incomes are highly dependent on climatic conditions and imperfect credit and insurance arrange-

* Corresponding author.

[1] The authors are with the Development Research Group, World Bank, 1818 H Street NW, Washington, DC, 20433. The authors have had useful discussions on this paper with Harold Alderman, Benu Bidani, Shaohua Chen, Gaurav Datt, Emmanuel Jimenez, Peter Lanjouw, Branko Milanovic, Arijit Sen, Dominique van de Walle, and Michael Walton. They have also benefitted from the comments of the Journal's anonymous referees. The financial support of the World Bank's Research Committee (under RPO 678-69) is gratefully acknowledged. These are the views of the authors and should not be attributed to the World Bank.

0147-5967/98 $25.00

ments leave farm households exposed to income risk. Transient poverty is also likely to be common in economies undergoing structural changes that can have diverse impacts at household level. The poverty that persists in mean consumption over time, we define as ''chronic poverty''; a household whose mean consumption is above the poverty line cannot be chronically poor by this definition although it may still experience transient poverty.

We suggest three main reasons why we want to know how much of observed poverty is transient rather than chronic. First, in assessing overall progress against poverty we may not be indifferent to whether it is transient or not. Consider two countries in which half of the population is poor at each of two dates, but in one country it is exactly the same households who are poor over time, while in the other it is none of the same households. Few observers will view these two extremes the same way. Yet that is what a conventional poverty measure such as the proportion of the population living below the poverty line at one time, does. However, it is unclear how one should weight the two types of poverty in forming an overall assessment. Development policy discussions appear often to give less emphasis to transient poverty than chronic poverty. That position may be contentious; for example, in economies undergoing major structural changes, transient poverty can be a serious policy concern.

Second, distinct policies are often called for to deal with transient poverty. Increasing the human and physical assets of poor people, or the returns to those assets, are thought to be more appropriate to alleviate chronic poverty. Insurance and income-stabilization schemes are seen to be more important policy instruments when poverty is transient (Lipton and Ravallion, 1995). Knowing how much the currently observed level of poverty is transient may thus inform policy choices.

Third, even when the policy objective is to reduce chronic poverty, the existence of transient poverty can influence policy choices. As always, the policies chosen will depend in part on the information available. A long-standing policy issue is how much transfers and public services should be targeted. Variability over time will clearly make current consumption a more noisy indicator of longer-term welfare, and so weaken the case for targeting on the basis of static data. We currently know little about the extent of leakage to the transiently poor from the common modes of targeting on the basis of static data.

All three reasons are relevant to postreform rural China. This is a setting in which one expects to find transient poverty, both because of the variability in agro-climatic conditions and the structural changes occurring as the economy is progressively liberalized. Poverty reduction policies in China have focused on poor-area development programs.[2] The motivation for these programs has empha-

[2] See Leading Group (1988), World Bank (1992), and Jalan and Ravallion (1998a).

sized chronic poverty rather than transient poverty; income-generating investments in local agriculture and rural development are promoted, rather than short-term insurance or state-contingent transfer schemes. Elsewhere we have attempted to estimate the dynamic gains from these programs, and our results indicated that the consumption growth rate in the targeted poor areas was about 1% per year higher than could have been expected without the program (Jalan and Ravallion, 1998a). So these interventions do appear to help alleviate chronic poverty in rural China. Is almost exclusive emphasis on these programs appropriate? In other work, we have found evidence of a significant exposure to income risk in this setting, so that the perfect markets hypothesis can be easily rejected (Jalan and Ravallion, 1998b). If a large share of the poverty at any one date is due to income risk, other policy instruments to smooth income and/or consumption may well be needed.

There is a related issue of how much China should rely on more finely targeted programs. While the evidence available suggests that the areas selected for participation in the poor-area development programs are appreciably poorer than average, the same evidence also indicates that roughly half of the poor are not found in those areas (Jalan and Ravallion, 1998a).[3] Different kinds of policies will be needed to reach these groups, probably entailing finer targeting. The question then arises as to how well this can be done using the sorts of data available for the purpose.

The currently available poverty measures for China (and in many other countries) are quite uninformative about these issues. Longitudinal observations, in which the same households are tracked to form a panel, are needed to distinguish transient poverty. However, the sample surveys on which poverty measures are usually based are single cross sections, in which the sample changes at each date. Past poverty measures for China based on the main national household surveys have also treated the samples as repeated cross sections, even though the main rural survey instrument, the Rural Household Survey (RHS), does not rotate the sample every year. Indeed, there are relatively long periods in which it is possible to form a panel, although that has never to our knowledge been done before.

In this paper, we use a new household-level panel data set formed from the RHS in a large southern and southeastern region of postreform rural China, to address two main questions:

(i) How much observed poverty in rural China is transient, in that it is due to consumption variability?[4]

(ii) How much does the existence of transient poverty in rural China limit

[3] On this point also see Riskin (1994).

[4] A further question of interest is whether the share of poverty that is transient has increased since the reforms. However, we cannot address this question here for lack of prereform data.

effectiveness at targeting people with low average consumption over time using transfers contingent on static data?

The following two sections examine in turn the two questions posed above; in each case we begin by outlining our method in abstract terms before presenting the results for China. Our conclusions are found in Section 4.

2. MEASURING TRANSIENT POVERTY

We measure transient poverty by the contribution to expected poverty of the variability over time in the individual welfare indicator; if the indicator does not change then transient poverty is zero.[5] We do not identify transient poverty as simply crossing the poverty line. Transient poverty is positive for someone who is always poor, but whose consumption varies, for example, due to uninsured income risk. However, this focus does mean that consumption fluctuations entirely above the poverty line are ignored. The effect of variability below the line is determined by the weights built into the poverty measure; here we shall adapt familiar ideas from static poverty measurement to an inter-temporal setting in a fairly natural mannner as discussed below.

We focus on consumption rather than income (although the latter is more commonly used in China) because it is likely to be a better welfare indicator for assessing poverty, particularly when incomes vary over time in reasonably predictable ways, as is invariably the case in a poor rural economy. It would not be particularly interesting in our view to find that there is considerable transient poverty in terms of current incomes in this setting, given that households can to some extent smooth their living standards from income variability, even if far from perfectly.

2.1. Method

To explain our approach, let $(y_{i1}, y_{i2}, \ldots, y_{iD})$ be household i's (positive) consumption stream over D dates and let $P = P(y_{i1}, y_{i2}, \ldots, y_{iD})$ be an intertemporal poverty measure for household i. Quite generally, this measure should reflect both the level of mean consumption over time and its variability around that mean. We define the *chronic poverty* component (C_i) of P as its value if consumption did not vary around the time mean ($\bar{y}_i$):

$$C_i = P(\bar{y}_i, \bar{y}_i, \ldots, \bar{y}_i). \tag{1}$$

We then define the *transient component* (T_i) as the remainder:

[5] Here we follow Ravallion (1988). Also see Rodgers and Rodgers (1993), who discuss the advantages of this approach over alternatives in the literature including those based on the number and length of spells of poverty.

$$T_i = P(y_{i1}, y_{i2}, \ldots, y_{iD}) - P(\bar{y}_i, \bar{y}_i, \ldots, \bar{y}_i). \quad (2)$$

So the intertemporal poverty measure is the sum of the chronic and transient components. Corresponding to each of the household-specific poverty measures there is an aggregate poverty measure across all households, which we denote by dropping the subscripts i.

We impose a number of conditions on the aggregate poverty measure. First, we require that it is both intertemporally and interpersonally additive. It is common to restrict attention to interpersonally additive measures, whereby aggregate poverty is a population-weighted mean of an individual poverty measure.[6] This implies that the measure is ''subgroup consistent'' (Foster and Shorrocks, 1991) in that, if poverty increases in any one subgroup and does not fall in any other, aggregate poverty must increase. We apply the same restriction to the intertemporal poverty measure, so that aggregate poverty for a given household is the expected value over time of a date-specific individual measure, denoted by p_{it}. A possible objection to this assumption is that the extent of household poverty at one date may depend on expenditures at a prior date; e.g., acquiring a bicycle now may make one less poor in the future. This objection is less persuasive if (as is the case in our data) the measure of consumption in a given period includes the imputed value of all commodities consumed in that period, even those purchased previously.

The second set of assumptions concerns the properties of the dated individual poverty measure, which we take to be a function of y_{it}, $p_{it} = p(y_{it})$. A simple example is $p_{it} = 1$ if $y_{it} < 1$ and $p_{it} = 0$ otherwise. The intertemporal poverty measure is then the proportion of dates for which household i falls below the poverty line and the aggregate poverty index is the intertemporal mean of the familiar ''headcount index'' of poverty, given by the proportion of the population falling below the line at each date.

While this is a simple example for $p(y_{it})$, it does not provide a very satisfactory measure of poverty, since the measure tells us nothing about how far below the line the household falls (an analogous concern is expressed concerning the static headcount index, in Sen, 1976). We assume instead that the measure: (i) penalizes losses to the poor and only the poor, in that p is strictly decreasing up to the poverty line, and zero thereafter; (ii) penalizes (or at least does not reward) inequality increases amongst the poor, so p is at least weakly convex in y; and (iii) is continuous at the poverty line. There appears to be broad agreement on the desirability of these properties, with the possible exception of the third.[7] When combined with (i), continuity precludes the

[6] For a survey of the additive measures found in the literature and their properties, see Atkinson (1987).

[7] The popular headcount index, i.e., the proportion of people below the poverty line, is discontinuous at the poverty line. However, most of the more well known ''higher-order'' measures of poverty are continuous (Atkinson, 1987); an exception is the Sen (1976) index, although this

possibility of consumption changes for the least poor (in a neighborhood of the poverty line) being given a higher weight than those amongst the poorest.

The main empirical poverty measure we will use is the squared poverty gap (SPG) index of Foster *et al.* (1984). The SPG for household i is

$$p(y_{it}) = (1 - y_{it})^2 \quad \text{if } y_{it} < 1$$
$$= 0 \quad \text{otherwise,} \tag{3}$$

where y_{it} is normalized by the (possibly household-specific) poverty line and thus takes the value of unity for someone at the poverty line.[8] The aggregate SPG is the household-size weighted mean of $p(y_{it})$ across the whole population. To test robustness to this choice we also consider two alternative measures. The ordinary poverty gap index in which the poverty gap in Eq. (3) is not squared, and so the measure is only weakly convex. The other is the Watts (1968) poverty index, which is simply the population mean log shortfall below the poverty line (counting the nonpoor as having zero shortfall). This index is strictly convex, as is SPG.

2.2. Data

For the purposes of this study, a new panel data set was constructed from the Rural Household Surveys done by China's State Statistical Bureau (SSB). Since 1984, the RHS has been a well-designed budget survey of a random sample of households drawn from a sample frame spanning rural China (including small-medium towns), with unusual effort made to reduce nonsampling errors.[9] Sampled households fill in a diary on daily expenditures and are visited on average every two weeks by an interviewer to check on the diaries and collect other data. There is also an elaborate system of cross-checking at the local level. The consumption data obtained from such an intensive survey process are almost certainly more reliable than those obtained by the far more common cross-sectional surveys in which the consumption data are based on recall at a single interview or possibly with one follow-up interview. For a six year period 1985–1990, the survey was also longitudinal, as it sampled the same households over time. This was done for administrative convenience since local SSB offices were set up in each sampled county. The survey was not intended to be a panel survey,

---

index can be made continuous by a simple re-normalization (Shorrocks, 1995). For further discussions of the case for and against a discontinuity at the poverty line, see Lipton and Ravallion (1995), Shorrocks (1995), Ravallion (1996a), and Bourguignon and Fields (1997). On the specific implications of a discontinuity for measuring transient poverty, see Ravallion (1988).

[8] Almost all poverty measures (including those used here) are homogeneous of degree zero in the vector of individual welfare indicators and the poverty line, so this normalization is possible.

[9] Chen and Ravallion (1996) provide a complete discussion of how the survey was undertaken.

but the panel can still be formed. To avoid spurious transience in the data, quite strong conditions were used to define a panel household.[10]

We constructed measures of chronic and transitory poverty using the panel data over the six-year period 1985–1990 from four contiguous provinces in southern China, namely Guangdong, Guangxi, Guizhou, and Yunnan, having a total population of 176 million in 1990. The region is a fairly good representation of the current regional disparities in rural China. Three of the provinces (Guangxi, Guizhou, and Yunnan) form a region of southwest China that is widely regarded as one of the poorest regions in the country. If we use average rural income as a yardstick, Guizhou is the poorest, while Guangxi and Yunnan are about equally well off on average, although still well below the national rural mean (Chen and Ravallion, 1996). Guangdong, on the other hand, is a relatively rich coastal region. In terms of the rates of income growth there are signs of divergence, with Guangdong having a considerably higher rate of economic growth than any of the other three provinces. Other indicators, e.g., the percentage of illiterate people over 15 years of age, the proportion of sampled households who obtained more than half their income from nonfarm activities and the percentage of households in a minority region confirm the spatial divergence between Guangdong and the other three provinces.

The consumption measure we use is comprehensive in that it includes imputed values for consumption from own production valued at local market prices and it also includes an imputed value of the consumption streams from housing and the inventory of consumer durables (Chen and Ravallion, 1996). Adjustments are also made for expenditures on housing and durable goods. All expenditures have been converted to constant 1985 prices using a province-specific rural consumer price index.

The poverty line is based on a normative food bundle set by the SSB, which assures that average nutritional requirements are met with a diet that is consistent with Chinese tastes. This bundle is valued at province-specific prices. The food component of the poverty line is augmented with an allowance for nonfood goods, consistent with the nonfood spending of those households whose food spending is no more than adequate to afford the food component of the poverty line. Full details on the construction of the poverty lines can be found in Chen and Ravallion (1996).

2.3. Results

In Table 1, we report the household-size-weighted means and interhousehold standard deviations for current consumption and income by year, the

[10] Constructing the panel from the annual RHS survey data proved to be more difficult than expected since the identifiers could not be relied upon. Fortunately, virtually ideal matching variables were available in the financial records, which gave both beginning and end of year balances. The relatively few ties by these criteria could easily be broken using demographic data.

TABLE 1

Summary Data by Year

	1985	1986	1987	1988	1989	1990
Mean consumption (Yuan per person per year, 1985 prices)	318.01 (327.63)	328.39 (354.81)	342.81 (384.39)	344.17 (387.16)	344.20 (419.45)	340.80 (428.11)
Income (Yuan per person per year, 1985 prices)	395.62 (532.42)	427.41 (622.78)	456.12 (669.76)	455.01 (652.01)	452.52 (649.30)	447.72 (649.01)
Correlation coefficient between consumption in current year and mean consumption over 6 years among the transient poor	0.480	0.640	0.719	0.752	0.715	0.636
Rank correlation between consumption in current year and mean consumption over 6 years among the transient poor	0.507	0.641	0.726	0.764	0.719	0.642
Poor in current year and chronically poor (%)	15.32	16.25	15.51	15.88	16.45	16.41
Poor in current year and not chronically poor (%)	13.12	11.24	7.45	6.88	8.85	11.89
Not poor in current year and chronically poor (%)	4.91	3.98	4.72	4.35	3.79	3.82

Note. Standard deviations are in parentheses. Correlation coefficients are significant at 1% level.

interhousehold correlation coefficients between each year's current consumption and time-mean consumption, and a discrete summary of the joint distribution of each date's consumption and time-mean consumption using the poverty line as the cutoff point. On average, between 1985–1990, household consumption and income increased in the sample. Current consumptions are significantly positively correlated with their respective time means. Yet there is appreciable transient poverty in that there are sizable numbers of people relative to the proportions of chronically poor (7–13% of the sample, depending on the year) who are poor in the current year but not chronically poor. A household is deemed to be in chronic poverty if its average consumption over the 6 years is below the poverty line.

A further indication of the extent of variability can be obtained from Table 2 which gives summary statistics on consumption and incomes by groups of households classified according to their time-mean consumption. The intertemporal standard deviation of both consumption and income tends to be lowest for the chronically poor and to rise as mean consumption rises, as does the coefficient of variation (CV) for consumption. There is a significant positive correlation between changes in consumption and changes in income; the correlation tends to be highest for the poor, suggesting that they are less able to protect their consumption from income fluctuations. This fact is also borne out by other research on the same data set indicating that while all households are vulnerable to income risk, the poorest among them are the least able to protect themselves against income shocks (Jalan and Ravallion, 1997a).

Table 3 gives the SPG poverty measures by subgroups defined according to the time-mean consumption as a proportion of the poverty line; results were very similar for the other measures. By construction, the chronic poverty measure is zero for all except those with time-mean consumption below the poverty line. What is more interesting is that 35% of transient poverty is found amongst those below the poverty line on average. However, the extent of transient poverty drops to a negligible amount for households whose time-mean consumption is more than 50% above the poverty line.

A number of other observations can be made about Table 3. Larger households have higher chronic poverty, but there is less difference in their transient poverty; hence, poverty amongst smaller households is more likely to be transient. Relatively high transient poverty amongst small households may well reflect their vulnerability to seasonal labor shortage in this setting where rural labor markets tend to be thin. Both chronic and transient poverty levels decline with the education levels for the head of the household while the proportion of poverty that is transient varies little with the education of the head. Both transient and chronic poverty tend to be higher for households with lower farm yields and less wealth; the proportion of their poverty that is chronic tends also to be lower. These households are less likely to be using

TABLE 2

Summary Data by Time-Mean Consumption Groups

Time mean consumption group	Number of persons in the sample	Consumption			Income			Correlation coefficient between change in consumption over time and change in income
		Mean consumption	Mean of the inter-temporal standard deviation	Mean of the inter-temporal coefficient of variation (%)	Mean income	Mean of the inter-temporal standard deviation	Mean of the inter-temporal coefficient of variation (%)	
Mean consumption (y) $<$ poverty line (z)	7,881	205.73	101.75	49.58	248.56	163.31	65.19	0.393 (0.0001)
$z \leq y < 1.25*z$	9,991	270.90	137.67	50.79	340.64	221.79	64.30	0.409 (0.0001)
$1.25*y \leq y < 1.5*z$	8,594	331.90	171.82	51.77	430.37	276.86	63.27	0.370 (0.0001)
$y \geq 1.5*z$	12,485	492.47	295.70	58.60	666.66	453.01	66.10	0.371 (0.0001)
Full sample	38,951	342.19	188.59	53.26	446.31	296.22	64.83	0.376 (0.0001)

Note. Consumption and income are in Yuan per person per year at 1985 prices. All means are household-size weighted. Figures in parentheses are the p-values under the null hypothesis of no correlation.

TABLE 3

A Profile of Both Transient and Chronic Poverty

Variable	No. of individuals	Transient poverty	Chronic poverty	Total poverty	Percentage of transient poverty in total poverty
Mean consumption (y) $<$ poverty line (z)	7,881	2.00	3.42	5.42	36.90
$z \leq y < 1.25*z$	9,991	0.78	0.00	0.78	100.00
$1.25*z \leq y < 1.5*z$	8,594	0.14	0.00	0.14	100.00
$y \geq 1.5*z$	12,485	0.04	0.00	0.04	100.00
Household size ≤ 2	282	0.79	0.35	1.14	69.30
Household size $= 3$	1,233	0.76	0.67	1.43	53.15
Household size $= 4$	3,996	0.59	0.64	1.23	47.97
Household size $= 5$	7,230	0.54	0.61	1.15	46.96
Household size $= 6$	8,478	0.62	0.53	1.15	53.91
Household size $= 7$	7,700	0.65	0.71	1.36	47.79
Household size > 7	10,032	0.74	0.90	1.64	45.12
Head of hh—illiterate	8,579	0.96	0.99	1.95	49.23
Head of hh—primary school educated	18,351	0.60	0.67	1.27	47.24
Head of hh—secondary school educated	9,287	0.52	0.55	1.07	48.60
Head of hh—high school educated	2,559	0.37	0.35	0.72	51.39
Head of hh—university^{+} educated	175	0.41	0.44	0.85	48.24

Mean yield ≤ 200 kg	10,967	1.07	1.50	2.57	41.63
200 kg < mean yield ≤ 250 kg	6,871	0.71	0.69	1.40	50.71
250 kg < mean yield ≤ 325 kg	11,028	0.46	0.38	0.84	54.76
Mean yield > 325 kg	10,085	0.33	0.16	0.49	67.35
Hh wealth ≤ 2,575 yuan	8,275	1.29	2.31	3.60	35.83
2,575 yuan < hh wealth ≤ 3,900 yuan	9,907	0.78	0.49	1.27	61.42
3,900 yuan < hh wealth ≤ 5,700 yuan	9,638	0.46	0.20	0.66	69.70
HH wealth > 5,700 yuan	11,131	0.20	0.09	0.29	68.97
Standard deviation of hh wealth (std) < 400	8,714	1.00	1.79	2.79	35.84
400 ≤ std < 715	9,678	0.72	0.71	1.43	50.35
715 ≤ std < 1,400	9,963	0.55	0.31	0.86	63.95
Std ≥ 1,400	10,596	0.36	0.13	0.49	73.47
Sample mean	38,951	0.71	0.72	1.43	49.39

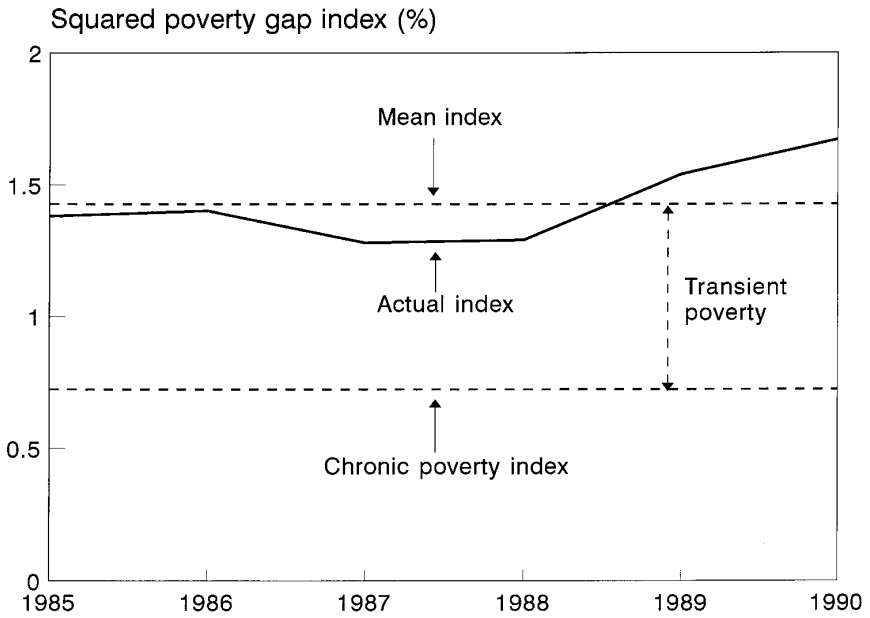

FIGURE 1

modern agricultural techniques that lead to higher average yields but tend also to be more risky. This may also explain why households whose wealth holdings fluctuate more, as indicated by the standard deviation of wealth, have lower poverty but it is more likely to be transient (Table 3).[11]

In Fig. 1, we give the SPG with and without consumption variability. We find that 49% of mean SPG is attributable to variability in consumption. The transient component accounts for a lower share of mean poverty using the poverty gap index and the Watts index, namely 38% and 41%, respectively. However, it is still clear that a large share is transient.

On repeating these calculations for each province separately, we found that the picture is quite different between the relatively well-off province of Guangdong and the other three provinces, but similar amongst the poorer three. The higher average consumption of Guangdong means that poverty is more likely to be transient while the percentages are much lower for the other three provinces. In Guangxi, Guizhou, and Yunnan, 49%, 43%, and 57% (respectively) of the mean SPG can be ascribed to variability in consumption while in Guangdong, it is 84%. The pattern is similar for the other two poverty measures.

[11] Jalan and Ravallion (1997) present multivariate models of both transient and chronic poverty.

3. TRANSIENT POVERTY AS A CONSTRAINT ON REDUCING CHRONIC POVERTY

Since we find considerable transient poverty in this setting, the impact on chronic poverty from a given anti-poverty budget finely targeted to those who appear to be poor according to static data will be diminished. To quantify this loss, we study a stylized policy problem in which we compare outcomes of fine targeting in terms of chronic poverty on the basis of data on current levels of living.[12] This approach does not aim to describe the actual policy problem in this setting with all its constraints. Nor do we subscribe to the view that chronic poverty should be the sole policy concern. Rather, the stylized policy problem studied here is a measurement tool for assessing the extent of transient poverty, where that is judged by how much it constrains attempts to reduce chronic poverty using only cross-sectional data. In short, it provides an alternative way of assessing the extent of consumption variability in the data.

3.1. Method

Consider a set of lump-sum transfers that aim to minimize chronic poverty subject to budget and informational constraints. Let $\Gamma(B)$ denote the minimum level of chronic poverty with a budget B and perfect information about each person's expected consumption and (hence) chronic poverty level. In practice, the information set is incomplete due to transient poverty. Let $\Gamma_t(B)$ denote the minimum level of chronic poverty attainable with the information set restricted to the observed consumptions at date t.

We now confine attention to strictly convex poverty measures that satisfy our other assumptions. These measures penalize inequality amongst the poor. Chronic poverty is minimized by giving the person with lowest time-mean consumption the first allocation from the budget so as to bring that person up to the level of the second poorest. Then both receive the next allocation, and so on.[13] Thus $\Gamma(B)$ can be calculated readily given the distribution of time-mean consumptions.

In calculating $\Gamma_t(B)$, the transfers are based on the observed consumptions at each date although we evaluate them by their impact on chronic poverty. The transfers are allocated the same way (from the poorest up), although this time it is the poorest in terms of *current* consumption. The impact on chronic poverty is then based on the new distribution of time-mean consumptions and is calculated as

[12] Here we follow the approach of Chaudhuri and Ravallion (1994).

[13] It is readily demonstrated that such step-wise targeting will be the allocation that has the largest impact on any additive, decreasing, strictly convex, and continuous poverty measure. For further discussion, see Chaudhuri and Ravallion (1994).

$$\Gamma_t(B) = C(\bar{y}_1 + \tau_{1t}, \ldots, \bar{y}_n + \tau_{nt}), \tag{4}$$

where the transfers τ_{it} for $i = 1, \ldots, n$ minimize the chronic poverty index based on *current* consumptions as given by

$$C(y_{1t} + \tau_{1t}, \ldots, y_{nt} + \tau_{nt}), \tag{5}$$

subject to the additively absorbed public budget:

$$\sum_{i=1}^{n} \tau_{it} = B. \tag{6}$$

By comparing $\Gamma(B)$ with $\Gamma_t(B)$, we can measure directly the extent to which transient poverty reduces the efficacy of transfers based on current consumptions as a means of fighting chronic poverty. Under our restrictions on the class of poverty measures, the functions Γ and Γ_t are strictly decreasing in B and $\Gamma_t(B) \geq \Gamma(B)$ at a given B. The value of $\Gamma_t(B) - \Gamma(B)$ measures the amount of chronic poverty that cannot be eliminated with a budget B using only the data for date t.

We can also calculate the extra budgetary cost of a given impact on chronic poverty as the dual of the above optimization problem. For example, it is common to calculate the budget needed to eliminate poverty with perfect information as given by the aggregate poverty gap; for eliminating chronic poverty that cost is

$$B^* = \sum_{i=1}^{n} \max(1 - \bar{y}_i, 0) \tag{7}$$

at which point $\Gamma(B^*) = 0$. However, when there is latent transient poverty, the poverty gap calculated using cross-sectional data will underestimate the true cost of eliminating chronic poverty, as given by B_t^* such that $\Gamma_t(B_t^*) = 0$.

3.2. Results

Perfect targeting on the basis of the six-year mean consumptions will eliminate chronic poverty at a cost of 6.9 Yuan per person in 1985 prices (Table 4); this represents 1.9% of mean consumption across all provinces and dates. Although this is the minimum cost, it would be impossible to attain in practice. How much more would the cost be under alternative assumptions about the information available? We first compare perfect targeting with a *uniform transfer*, given by the minimum sum needed to eliminate chronic poverty subject to the constraint that all persons (whether poor or not) receive the same amount. This stylized policy option does not require any knowledge about the distribution of consumption other than the value of its lower bound. With a uniform transfer, the cost is 3.6 times the budget under perfect targeting (Table 4).

TABLE 4

Cost of Eliminating Chronic Poverty under Various Assumptions about the Policy-Maker's Information

	Perfect targeting	Uniform transfer	Targeting on the basis of current consumption					
			1985	1986	1987	1988	1989	1990
All four provinces	6.94	25.07	24.88	22.61	21.47	23.29	20.86	22.96
	(100)	(361)	(358)	(326)	(309)	(335)	(300)	(331)
All except Guangdong	9.07	32.52	32.29	29.48	27.98	30.33	27.16	29.92
	(100)	(359)	(356)	(325)	(280)	(303)	(300)	(330)
Guangdong	0.33	0.88	0.72	0.60	0.73	0.69	0.70	0.83
	(100)	(267)	(218)	(181)	(222)	(211)	(211)	(251)
Guangxi	6.30	24.00	18.89	16.99	16.59	18.34	17.02	18.05
	(100)	(379)	(300)	(270)	(263)	(291)	(270)	(286)
Guizhou	14.35	46.23	44.44	41.20	38.12	39.17	35.56	36.96
	(100)	(322)	(310)	(287)	(266)	(273)	(248)	(257)
Yunnan	7.62	26.20	22.21	21.77	22.23	21.83	24.36	23.88
	(100)	(344)	(292)	(286)	(292)	(287)	(320)	(314)

Note. The table gives the expenditure per capita (of the total population) in Yuan at 1985 prices needed to fill exactly the poverty gaps in terms of the six-year mean consumption under alternative assumptions about the information available to the policy maker. Figures in parentheses are percentages of the total budget required to eliminate chronic poverty.

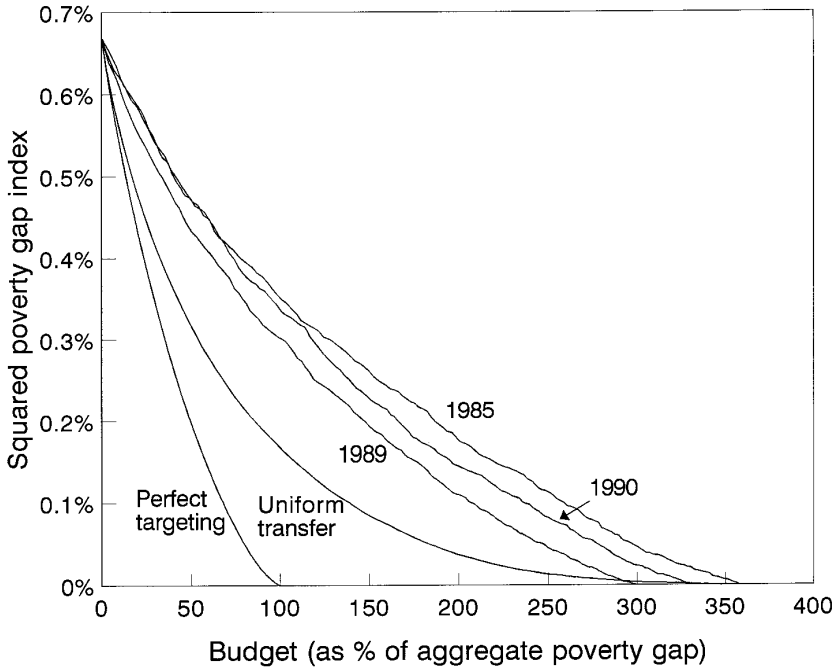

FIGURE 2

Next consider targeting based solely on each year's *current* consumption. To isolate the contribution of transient poverty, we assume that targeting is otherwise optimal in that the poverty gaps are filled exactly. If there was no transient poverty, this cost would be the same as under perfect targeting. However, it turns out that there is so much transient poverty that the cost of eliminating poverty using only current consumptions is much greater than the perfect targeting case. Indeed, the cost is typically only slightly lower than the cost of uniform transfers. There is enough transient poverty to virtually eliminate the cost saving from even optimal targeting on the basis of current consumption.

We repeated these calculations for various budgets. In Fig. 2 we give the measure of chronic poverty (vertical axis) attainable for each aggregate expenditure on transfers (on the horizontal axis). So, for example, starting with any budget on the horizontal axis, the chronic poverty measure on the uniform-transfer curve is obtained when that budget is allocated equally to everyone whether poor or not. The point on the 1990 curve is the measure of chronic poverty obtained if the same budget is allocated according to current consumption in 1990. The point on the perfect-targeting curve is

obtained when the budget is allocated according to the six-year mean consumption. We have left out the curves for 1986–1988 because they are virtually indistinguishable from that for 1989.

Even for the worst year, i.e., 1989, when the budget is less than about 300% of the aggregate poverty gap, uniform transfers dominate allocations based on current consumptions. Judged by the impact on chronic poverty, it would be unambiguously better to share the budget equally than to rely on current consumptions for all except large budgets. We repeated these calculations for the individual provinces, except Guangdong, where there is only a small amount of chronic poverty, and found that the dominance of uniform transfers held for budgets up to 250–300% of the aggregate poverty gap.

It might also be argued that the variability of consumption over time contains a high share of measurement error and that this is why we get these results. However, this is the same sort of data on which actual targeted interventions are often based. Indeed, one would normally expect *more* measurement error in the data available to a policy maker for the purposes of fine targeting, given the incentive for misreporting and the need for more rapid appraisal methods in a comprehensive anti-poverty program as distinct from a sample survey.[14] So fine targeting using static data may well be even less effective in reaching the long-term poor than our simulations suggest.

As a final comment, it is of interest to compare these results with similar calculations in Chaudhuri and Ravallion (1994) for semi-arid areas of rural India using a panel data set of the same length. Targeting transfers on the basis of current consumptions in the Indian data entailed an appreciably larger impact on chronic poverty at any given budget than in our China data. While we have found that uniform transfers dominate for all except high budgets, the Indian data indicated that the cost of a moderate drop in chronic poverty using uniform transfers was two–three times higher than was possible with cross-sectional data. The Indian data indicated far less transient poverty than we have found. In future work we hope to probe this issue.

4. CONCLUSIONS

We find considerable transient poverty in this region of rural China from 1985 to 1990. One-half of the mean squared poverty gap and over one-third of the mean poverty gap is accounted for by year-to-year fluctuations in consumption. About 40% of the transient poverty is found amongst those who are not poor on average, but almost all of this is for households whose mean consumption over time is no more than 50% above the poverty line. Transient poverty is likely to be a significant constraint on the effectiveness

[14] For evidence on the performance of rapid appraisal methods, see Ravallion (1996b).

of targeted anti-poverty policies contingent on current levels of living for reaching the long-term poor. Static consumption data contain considerable noise about long-term poverty status. For example, the full cost of eliminating chronic poverty using a current cross section of consumptions, which is itself a very demanding criterion, is three or four times the poverty gap based on mean consumption over six years. Indeed, targeting on the basis of current consumption has *less* impact on chronic poverty for any given budget than a uniform lump-sum allocation in which the same amount is given to everyone, whether poor or not. A comparison with similar tests for rural India suggests that there is far more transient poverty in our data for rural China.

While our consumption data would appear to be of good quality given the survey methods used, the changes in consumption may well contain more measurement error than do the intertemporal consumption averages. Our measure of transient poverty will reflect this extra measurement error as well as the actual consumption changes. This will lead to over-estimation of the share of poverty that is transient. However, one can conjecture that the static data normally available for fine targeting would be even noisier than our data. This reinforces our conclusion that attempts to reach people with low long-term standards of level of living by fine targeting on the basis of current data may well be disappointing. The same time-varying measurement errors will, of course, also confound attempts to reach the transiently poor.

The degree of transient poverty that we find in these data throws open the question as to whether the current emphasis on fighting chronic poverty in China through poor-area development programs is appropriate. It is far from clear that such programs will be sufficient to eliminate poverty in China. The exposure to uninsured income risk that underlies the high transient poverty will probably persist even within successful program areas. Hence, the many poor in nonprogram areas will not benefit. There is a case for considering more finely targeting programs, although not as a means of fighting chronic poverty but rather as a way of stabilizing incomes by making assistance contingent on adverse events.

REFERENCES

Atkinson, Anthony, ''On the Measurement of Poverty.'' *Econometrica* **55,** 4:749–764, July 1987.

Bourguignon, Francois and Fields, Gary, ''Discontinuous Losses from Poverty, Generalized P_α Measures, and Optimal Transfers to the Poor'', *J. Public Econ.* **63,** 2:155–175, Jan. 1997.

Chaudhuri, Shubham, and Ravallion, Martin, ''How Well do Static Indicators Identify the Chronically Poor?'' *J. Public Econ.* **53,** 3:367–394, March 1994.

Chen, Shaohua, and Ravallion, Martin, ''Data in Transition: Assessing Rural Living Standards in Southern China.'' *China Econ. Rev.* **7,** 1:23–56, Spring 1996.

Foster, James, Greer, Joel, and Thorbecke, Erik, ''A Class of Decomposable Poverty Measures.'' *Econometrica* **52,** 3:761–766, May 1984.

Foster, James, and Shorrocks, A. F., ''Subgroup Consistent Poverty Indices.'' *Econometrica* **59,** 3:687–709, May 1991.

Jalan, Jyotsna, and Ravallion, Martin, ''Modeling Poverty in Rural China.'' *Mimeo.* Development Research Group, The World Bank, Washington D.C.: The World Bank, 1997.

Jalan, Jyotsna, and Ravallion, Martin, ''Are There Dynamic Gains from a Poor-Area Development Program?'' *J. Public Econ.*, **67,** 65–85, Jan. 1998a.

Jalan, Jyotsna, and Ravallion, Martin, ''Are the Poor Less Well Insured? Evidence on Vulnerability to Income Risk in Rural China.'' *J. Devel. Econ.*, forthcoming 1998b.

Leading Group, *Outlines of Economic Development in China's Poor Areas.* Beijing: Office of the Leading Group of Economic Development in Poor Areas Under the State Council, Agricultural Publishing House, 1988.

Lipton, Michael, and Ravallion, Martin, ''Poverty and Policy.'' In Jere Behrman and T. N. Srinivasan, Eds., *Handbook of Development Economics, Volume III.* Amsterdam: North Holland, 1995.

Ravallion, Martin, ''Expected Poverty Under Risk-Induced Welfare Variability.'' *Econ. J.* **98,** 393:1171–1182, December 1988.

Ravallion, Martin, ''Issues in Measuring and Modeling Poverty.'' *Econ. J.* **106,** 438:1328–1343, Sept. 1996a.

Ravallion, Martin, ''How Well Can Method Substitute for Data? Five Experiments in Poverty Analysis.'' *World Bank Res. Observer* **11,** 2:199–221, Aug. 1996b.

Riskin, Carl, ''Chinese Rural Poverty: Marginalized or Dispersed?'' *Amer. Econ. Rev. Papers and Proc.* **84,** 2:281–284, May 1994.

Rodgers, Joan R., and Rodgers, John L., ''Chronic Poverty in the United States.'' *J. Human Res.* **28,** 1:25–54, Winter 1993.

Sen, Amartya, ''Poverty: An Ordinal Approach to Measurement.'' *Econometrica,* **44,** 2:219–231, March 1976.

Shorrocks, Anthony, ''Revisiting the Sen Poverty Index.'' *Econometrica* **63,** 5:1225–1230, Sept. 1995.

Watts, H. W., ''An Economic Definition of Poverty.'' in Daniel P. Moynihan, Ed., *On Understanding Poverty.* New York: Basic Books, 1968.

The World Bank, *China: Strategies for Reducing Poverty in the 1990s.* Washington, D.C.: The World Bank, 1992.