



Columbia University

Department of Economics
Discussion Paper Series

**Smoothing Consumption Under Income Seasonality:
Buffer Stocks vs. Credit Markets**

Shubham Chaudhuri
Christina Paxson

Discussion Paper #:0102-54

Department of Economics
Columbia University
New York, NY 10027

April 2002

Smoothing consumption under income seasonality: buffer stocks vs. credit markets*

Shubham Chaudhuri
Columbia University

Christina Paxson
Princeton University and NBER

December 2001

Abstract

Rural households in many developing economies have incomes that vary seasonally. We explore the implications of this income seasonality for household consumption. We use household-level data from three Indian villages to document seasonal patterns in income and consumption, and to test whether income seasonality produces seasonal consumption variation. Our basic finding is that while there does appear to be some seasonality in consumption patterns, it is much less pronounced than in the case of income, and more surprisingly, that the patterns are quite similar for households with very different seasonal income patterns. While this finding is consistent with well-functioning credit markets, we show, through simulations, that it is also consistent with a simple buffering model of consumption in which cautious households cannot borrow, but can save via the accumulation of assets. We provide evidence that suggests that households rely more on buffering behavior than on credit markets to smooth consumption under income seasonality.

Keywords: consumption smoothing, income seasonality, buffer stocks

JEL codes: E21, O12, Q12

*We thank the International Crop Research Institute for the Semi-Arid Tropics (ICRISAT) for access to the data. Tim Besley, Angus Deaton, Tom Mehen, and Robert Townsend are thanked for comments. An earlier version of this paper was titled: "Consumption smoothing and income seasonality in rural India".

1. Introduction

Rural households in many poor agrarian economies—where rain-dependent crop cultivation is the primary source of household income—have incomes that vary seasonally. And the seasonal variation in income can be quite large. For instance, in the Indian villages we study in this paper, cultivating households receive, on average, about 75% of their annual income in a 3-month period.

What are the consequences of this pronounced income seasonality for consumption? Empirical evidence suggests that household consumption levels also vary seasonally in rural economies, leading in many cases to seasonal variation in nutritional status and health.¹ The presumption, especially in policy circles, has been that the observed consumption seasonality is largely driven by the seasonal variation in income and that the link between the two can be attributed to poorly functioning credit markets. Borrowing constraints caused by credit market imperfections can indeed combine with income seasonality to produce consumption seasonality. If households cannot borrow against future income during the slack season, then consumption levels may well be lower during the period before the harvest. But before we can conclude that the reduction of seasonal consumption variation is an appropriate policy goal, and further, that policy interventions should be directed towards alleviating credit constraints, two important unresolved questions need to be addressed.²

The first has to do with the sources of consumption seasonality. Even in settings where incomes are highly seasonal, income seasonality need not be the source of consumption seasonality. The latter may instead stem from any one or more of a number of sources, among them, seasonal variation in prices, preferences, labor effort, or precautionary savings motives linked to seasonal patterns in the resolution of uncertainty about income. Secondly, though borrowing constraints can, in principle, contribute to consumption seasonality, the extent to which they do so in practice, remains an open question. Our focus in this paper is on this second question. There is some evidence that credit constraints might prevent poor agrarian households from smoothing consumption across years (see, for example, Morduch (1990), Rosenzweig (1988) and Rosenzweig and Wolpin (1993)). However, except for Jacoby & Skoufias (1998), there is little systematic evidence on whether these households are able to smooth consumption across seasons.³

We use household-level panel data from three Indian villages collected by the International Crop Research Institute for the Semi-Arid Tropics (ICRISAT) to investigate whether households are able to smooth consumption across seasons, despite the seasonal variation in

¹For instance, various articles in Chambers et al.(1981) and Sahn (1989) document nutritional deficiencies during the “lean season”, before crops are harvested, in a variety of countries. Paxson (1992) finds evidence of seasonal variation in total consumption expenditure in Thailand. Walker and Ryan (1990) presents evidence on seasonal patterns in nutritional status in a sample of Indian villages, although deficits do not always occur during the lean season.

²See Besley (1994) for a more detailed discussion of this point.

³Jacoby and Skoufias (1998) presents results on income and consumption seasonality using the same data but different methods from the ones we use. They reach similar conclusions regarding the ability of households to smooth consumption across seasons.

A number of other papers have explored the implications of income seasonality for other aspects of household behavior. For instance, Jacoby and Skoufias (1997) examines the impact of income seasonality on school attendance, while Behrman (1988) looks at the impact of seasonality on intra-household allocations. Canagarajah and Pudney (1993) examine the effects of illness and seasonal involuntary unemployment on seasonal consumption.

incomes. In the next section, we document seasonal patterns in income and expenditure, and examine whether household consumption expenditure tracks household income across seasons. We follow the methodology presented in Paxson (1992) and exploit the fact that seasonal income patterns differ across households within a village. We test whether households that have different seasonal income patterns also have different seasonal consumption patterns and, if so, whether seasonal consumption and income patterns are related. The results show little evidence that household consumption tracks income across seasons. Although there are significant seasonal changes in consumption expenditure common to households within villages, the month-to-month variation in consumption is, on the whole, much less pronounced than the month-to-month income variation. More surprisingly, we find that the seasonal patterns in consumption are quite similar for households with very different seasonal income patterns. These findings suggest that the households in our sample are able to and do smooth consumption in the face of seasonal income variation.

While the absence of “tracking” behavior is consistent with the absence of borrowing constraints, it need not imply that credit markets function well within villages, and that precautionary savings motives are unimportant. Recent research on buffer-stock models of saving indicate that even in the complete absence of credit markets, cautious households may accumulate and draw down stocks of physical or financial assets to maintain consumption levels that vary little from year to year (see Deaton (1992)). In these models, large changes in consumption come about only when stocks of assets are drawn down to zero, something that may happen infrequently.

In Section 3 of the paper, we develop and simulate a seasonal version of the buffer-stock model to investigate whether the degree of seasonal consumption smoothing observed in the data is consistent with a buffering model. Our simulations indicate that even in the complete absence of credit markets, buffering behavior can nearly eliminate systematic seasonal consumption patterns. Therefore, our finding that household consumption does not track household income across months does not necessarily imply that credit markets function well within villages. We therefore turn, in Section 4 to other sources of evidence regarding the relative importance of buffer stocks and credit markets in enabling households to smooth consumption under income seasonality. Data constraints limit what we are able to say, but the additional evidence we consider is indicative of buffering behavior and suggests that borrowing constraints may be operative in two of the three villages. Section 5 concludes.

2. Seasonal patterns in income and consumption

This section presents descriptive information on seasonal patterns in income and expenditure for households from three villages in south-central India. The data we use come from the Village-Level Studies (VLS) longitudinal household surveys carried out by the International Crop Research Institute for the Semi-Arid Tropics (ICRISAT), based in Hyderabad, India. We use information from the villages of Aurepalle in the state of Andhra Pradesh, and Shirapur and Kanzara in the state of Maharashtra. In each village, a sample of thirty cultivating and ten landless households was selected. These households were then interviewed frequently- typically every 4 to 6 weeks for eleven years, from 1975 to 1985, and asked questions about socio-demographic characteristics, production decisions, and transactions in markets. In this study we use data from 1976 through 1981, when more detailed transactions data were collected. Walker and Ryan (1990) provide extensive socio-economic profiles of the sample villages, a wealth of institutional, environmental and historical detail,

as well as a comprehensive summary of the many findings of the surveys. Appendix A to this paper provides details on how the different variables we use are measured.

An important feature of these villages is that income derived from crops has pronounced seasonal patterns. This is not surprising, given that these villages are in rain-fed areas where agricultural cultivation depends heavily on monsoon rains. As indicated in Table 1, crop income accounts for a large share of total annual income for farm households, ranging from 47.4% in Kanzara to 61.3% in Shirapur. Irrigation, which could potentially result in less seasonality in profits from crops, is fairly uncommon. Farm households in Aurepalle and Shirapur had, on average, irrigation for thirteen percent of their cropped area over the sample period. Irrigation was even less prevalent in Kanzara, reflecting the fact that Kanzara has rainfall that is higher and less variable than the other areas, see Walker and Ryan (1990, p. 34).i (Kanzara is also richer than the other villages.)

The combination of monsoon-dependent crops and lack of irrigation results in flows of income from crops that are very unevenly distributed over the year, much more so than flows of income from other sources. Figure 1 traces out monthly averages of income from crops, and income from other sources (trade and handicrafts, wages, and livestock.) The seasonal patterns in crop income are consistent with the timing of the monsoon (which occurs between June and September), and the soil conditions and prevalence of irrigation in each of the three villages. In Aurepalle and Kanzara, the main harvest is of crops planted with the onset of the monsoon in May and June ("kharif" season crops), and occurs in late fall. There is a second harvest between March and May. In Aurepalle, this second harvest consists largely of HYV crops, predominantly paddy, grown on irrigated land. In Kanzara it consists of longer-duration kharif crops. The timing of crop income in Shirapur differs from the other villages. There is a single main harvest period, in the first three months of the calendar year. Farms in Shirapur tend to have deep soils that hold the monsoon rainfall. Crops are typically planted after the monsoon is completed, and harvested late (see Walker and Ryan, 1990). Because farm households rely more heavily on crop income than do other households, they also have more seasonally variable total household income. As shown in Table 1, the fraction of income received by farm households in the three highest-income months ranges from 61.3% in Aurepalle to 81.8% in Shirapur. The corresponding fractions for non-farm households are much lower, ranging from 39% in Aurepalle to 46% in Shirapur.

Does the higher degree of seasonality in the incomes of farm households translate into higher seasonality in consumption expenditure? The evidence in Table 1 suggests not. Even for farm households, both food expenditure and expenditure on all nondurable goods are less concentrated (in the top three months) than is income. For example, in Aurepalle, 38% of expenditure by farm households is incurred in the three highest months, whereas 61% of income is earned in the three highest months. Furthermore, the level of concentration of expenditure (in the top three months) is nearly the same for farm and non-farm households within each village.

Further evidence on the timing of income and expenditure for farm and non-farm households is presented in Figure 2. These graphs display the month effects estimated from regressions of the log of monthly expenditure and of monthly income (relative to its average value) on a set of month dummies, which were allowed to vary across farm and non-farm households. Specifically, the equation for income is:

$$A_{imt} = \alpha_m + \alpha_m^F F_{it} + e_{imt} \quad (2.1)$$

where A_{imt} is the ratio of household i 's total income in month m and year t to the household's

average monthly income in year t (i.e. annual income divided by 12.)⁴ This variable is regressed on a set of dummy variables for each month and a set of month dummies interacted with an indicator of farm status F_{it} , which equals 1 if the household has gross cropped area in excess of 0.5 hectares, and equals 0 otherwise. Thus, the parameters α_m measure month effects in the income share for nonfarm households. Month effects in the income share for farm households are measured as $\alpha_m + \alpha_m^F$. The parameters α_m^F represent the difference in the month effects between the farm and non-farm households. An intercept is included, so the first month effect, α_1 , is normalized to equal zero.

The expenditure equation is:

$$\ln(E_{imt}) = X_{it}\beta_0 + \beta^F F_{it} + \beta_m + \beta_m^F F_{it} + u_{imt} \quad (2.2)$$

where $\ln(E_{imt})$ is the logarithm of expenditure of household i in month m in year t (see Paxson (1992) for a derivation of this expenditure equation). Two variants of (2.2) are estimated. In the first, E_{imt} denotes total nondurable expenditure, and in the second it denotes food expenditure. Both measures include the value of food produced and consumed at home. The expenditure measure is regressed on a set of variables X_{it} that are constant within years, including the logarithm of average monthly income, the numbers of males, females, and children in the household at the beginning of the year, a set of year dummies, and a constant. The coefficients β_m denote month effects in expenditure for nonfarm households, and since an intercept is included in the equation the coefficient for the first month (β_1) is normalized to zero. The coefficients β_m^F denote month-specific differences between farm and nonfarm households. Again, since a variable denoting farm status is included, the first of these month/farm interactions is normalized to zero. Thus, the difference between farm and nonfarm expenditure in any month (given X_{it}) equals β^F for month 1 (January), and $\beta^F + \beta_m^F$ for all other months. Table 2 presents the results of F-tests of the significance of the month effects and the month-farm status interactions.

The results in Figure 2 confirm the patterns suggested by Table 1. First, farm and non-farm households have different seasonal income patterns. Income is relatively flat across seasons for non-farm households, and in fact F-tests indicate that the hypothesis that there are no month effects in the incomes of non-farm households cannot be rejected for any village except (perhaps weakly) for Shirapur. The month effects in income for farm households are jointly significant, and the hypothesis of identical month effects across farm and nonfarm households is strongly rejected. Second, although there are typically seasonal patterns in expenditure, these seasonal patterns are for the most part similar across farm and non-farm households. As shown in Figure 2, expenditure is lower during the summer months, and higher during the harvest periods. (Note that the scale for the expenditure graphs is nearly 10 times smaller than for the income graphs.) However, these patterns in expenditure appear for both farm and non-farm households: the post-harvest expenditure boom is not confined to households with high harvest income. This is true for both food and non-durable expenditure. Tests for whether month effects in expenditure are identical across farm and non-farm households cannot be rejected in Aurepalle and Shirapur (see Table 2). The results for Kanzara are slightly different. The hypothesis that the month effects in

⁴Monthly income is divided by average monthly income to remove scale effects. It should also be noted that “year” refers to the agricultural year (June to May) rather than the calendar year, since the data are collected on a crop-year basis. However, month effects in income and expenditure are labelled numerically to correspond to the calendar year, i.e. α_1 denotes the month effect in January, α_2 denotes the month effect in February, etc.

expenditure differ across farm and nonfarm households can be rejected at a confidence level of .019 (for total expenditure) and .005 (for food.) However, the month effects graphed in Figure 2 do not indicate that expenditure tracks income across months. For instance, the expenditure of farm households is low relative to that of nonfarm households between August through November, but this is not a time period during which farm incomes are relatively low.

3. Smoothing consumption under income seasonality: theory

The evidence from the previous section indicates that seasonal patterns in consumption do not track those in income, and that suggests that households are able to smooth consumption in the face of systematic month-to-month fluctuations in income. Such behavior is consistent with the standard permanent income model, with complete credit markets. However, we show in this section that it is also consistent with a buffer-stock model of consumption, in which households cannot borrow at all. Both the permanent income model and the buffer-stock model are well-known in the literature on saving and consumption, but they have been typically used to analyze annual rather than seasonal consumption patterns. Here we explore the implications of the two models for consumption behavior when income is seasonal. We begin with the permanent-income model.

Assume that preferences are quadratic and separable across periods, that the rate of time preference is equal to a fixed interest rate r , and that preferences do not vary across time periods. Then, the change in consumption from period to period can be expressed as:

$$\Delta c_t = c_t - c_{t-1} = \left(\frac{r}{1+r}\right) \sum_{k=0}^{\infty} (1+r)^{-k} [E_t - E_{t-1}] Y_{t+k} \quad (3.1)$$

where c_t denotes consumption in time t and Y_t is non-asset income in time t . E_t is the expectations operator, conditional on information known at time t . In words, the change in consumption between $t-1$ and t is equal to the revision in permanent income between the two periods, where permanent income equals the annuity value of discounted future earnings plus the value of current assets. Equation (3.1) does not explicitly incorporate income seasonality, but since the equation is valid for any income process, it applies equally to situations with seasonal income variation. One can let “ t ” denote a month rather than a year, and allow average income levels (and higher moments of income) to vary systematically across months.

Equation (3.1) implies that since consumption responds only to unexpected innovations in permanent income, deterministic seasonal patterns in income should have no effect on consumption. The expected value of the change in consumption between two months, conditional on information known in the earlier month, is zero. Therefore, (3.1) implies that on average there will be *no* seasonal consumption variation. More general models of consumption—for instance those that allow for seasonal variation in preferences, prices or interest rates, or seasonal patterns in the resolution of income uncertainty (see Chaudhuri(1999) for an example of the latter)—do yield systematic seasonal patterns in consumption. However, even in these more general models, seasonality in income *levels* will not directly translate into consumption seasonality and consumption will still not respond to anticipated income changes.⁵

⁵A large number of papers over the last two decades have used household-level data to empirically investigate whether, contrary to the implications of the basic theory, household consumption does indeed

An alternative to the PIH is the buffer-stock model presented in Deaton (1989,1991). In this model, households are prohibited from borrowing, but can accumulate assets used to buffer consumption from income fluctuations. Deaton adds a precautionary motive for saving by assuming isoelastic preferences, which imply convex marginal utility. In addition, consumers are assumed to be impatient, meaning that their rate of time preference (denoted δ) exceeds the interest rate r . Given these assumptions, consumers save only to buffer their consumption from short-term income fluctuations. Impatience prevents long-term asset accumulation, but caution coupled with borrowing constraints provides incentives to hold a buffer of assets in most periods. Consumption is generally smooth relative to income, but at times declines abruptly in periods when assets are drawn down to zero. We work with this basic model, but make the additional assumption that income is seasonal. We then present simulations of the model, to investigate how credit constraints and precautionary savings motives affect consumption seasonality.

Assume that there are two seasons, 1 and 2. Assets must be nonnegative in both seasons of each year. Infinitely lived consumers starting in season 1 maximize:

$$U_{1t} = \sum_{j=0}^{\infty} E_{jt} [u(c_{1,t+j})(1+\delta)^{2j} + u(c_{2,t+j})(1+\delta)^{2j+1}] \quad (3.2)$$

where δ is the positive rate of time preference, c_{mt} is consumption in season m of year t , and $u(c_{mt})$ is the sub-utility function in season m of year t . The maximization problem for consumers starting in period 2 is identical, but with the season subscripts for consumption switched. Assets evolve across seasons according to:

$$a_{2t} = (1+r)[a_{1t} - c_{1t}] + Y_{2t} \quad (3.3)$$

and

$$a_{1t} = (1+r)[a_{2,t-1} - c_{2,t-1}] + Y_{1t} \quad (3.4)$$

where a_{mt} represents cash-on-hand at the beginning of season m in year t , equal to the sum of assets held over from the previous season (including any interest they have earned), plus Y_{mt} , income earned in season m of year t .

Utility maximization yields the following Euler equations for consumption in each season:

$$u'(c_{1t}) = \max\{u'(a_{1t}), \beta E_{1t}[u'(c_{2t})]\} \quad (3.5)$$

and

$$u'(c_{2t}) = \max\{u'(a_{2t}), \beta E_{2t}[u'(c_{1,t+1})]\} \quad (3.6)$$

where marginal utility is denoted $u'(c_{mt})$, and in what follows we will assume preferences are isoelastic such that $u'(c_{mt})$ equals $c_{mt}^{-\rho}$. The term β equals $[(1+r)/(1+\delta)]$, and impatience implies that $\beta < 1$. Equations (3.5) and (3.6) are standard Euler equations, modified to

respond to predictable income changes. The early results tended to be mixed, perhaps in part due to the difficulties of adequately identifying the predictable component of income changes. More recently, however, a number of papers have cleverly exploited exogenous, institutionally derived sources of variation in the timing and magnitude of household income flows to test this hypothesis. For instance, Souleles (1999) examines the response of household consumption to the receipt of (predictable) income tax refunds, Parker (1999) estimates households' propensity to consume out of predictable increases in take-home pay associated with the reduction in Social Security tax withholdings when the Social Security tax cap is reached, and Shapiro and Slemrod (1995) report the results of a survey in which consumers were asked what they planned to do with the increase in take-home pay that resulted from an executive order issued by President Bush in early 1992 reducing standard tax-withholding rates.

account for the presence of borrowing constraints. Consumption in any period can rise no higher than a_{mt} , and this upper bound on consumption implies a lower bound for marginal utility.

To add content to equations (3.5) and (3.6), it is necessary to make assumptions about the income process in each season. We assume that season 1 is the growing season, during which income is low on average but fairly stable, and that season 2 is the harvest season, during which income has a high mean and a high variance. Specifically:

$$Y_{mt} \sim (\bar{Y}_{mt}, \sigma_m^2), \bar{Y}_1 < \bar{Y}_2, \sigma_1^2 < \sigma_2^2$$

Income in season m has support $[Y_m^L, Y_m^H]$, and we assume that Y_m^L exceeds zero for both seasons. We assume incomes are independently distributed across seasons and years, and that farmers receive no information during the growing season about what harvest season incomes are likely to be.⁶

Given these assumptions, the Euler equations will be solved by a unique set of season-specific consumption functions. Consumption in season m is a function of cash-on-hand in the beginning of that season, and is denoted $c_m(a_m)$. The consumption functions are defined implicitly by the following equations:

$$[c_m(a_m)]^{-\rho} = \max\{a_m^{-\rho}, \beta \int_{Y_n^L}^{Y_n^H} c_n([1+r][a_m - c_m(a_m)] + Y_n)^{-\rho} dF_n(Y_n)\} \quad (3.7)$$

where $F_n(Y_n)$ is the distribution function for Y_n .

We examine how seasonal patterns in the means and variances of income affect the prevalence of borrowing constraints and seasonal consumption patterns. Intuition suggests that consumers will be less likely to face binding borrowing constraints in the high-income harvest period (season 2). Since growing-season income is lower on average than harvest income, farmers have incentives to carry positive assets out of the harvest period. It is shown in Appendix B that a sufficient condition for farmers to never face binding borrowing constraints in season 2 is that $Y_2^L > \beta^{-1/\rho} Y_1^H$. In words, if the worst possible income during the harvest season exceeds the best possible income during the growing season (scaled by a factor of $\beta^{-1/\rho}$, which exceeds 1) then the farmer will always carry positive assets out of the harvest period. Note that this condition is sufficient but not necessary, and it is possible that the farmer will always carry positive assets out of the harvest period even in if $Y_2^L < \beta^{-1/\rho} Y_1^H$. Borrowing constraints will also never bind in the harvest period if $Y_1^L \leq 0$. Households will always carry positive assets out of the harvest season to guard against the possibility that growing season income (and consumption) is zero.⁷

Numerical calculations are required to illustrate further properties of the season-specific consumption functions. We calculated the consumption functions under the following assumptions about income, preferences, and the rate of interest rate.⁸ First, we make the

⁶It is straightforward, in theory, to allow for the accumulation of information about future harvests. Consumption then becomes a function of two state variables, cash-on-hand *and* information held at the beginning of each time period. In practice, adding state variables makes it much more difficult to obtain numerical solutions for the consumption functions, since one must solve for consumption for all possible combinations of the state variables.

⁷In this case, it must be that the household begins life (in season 1) with positive assets, to guarantee positive consumption in the first time period.

⁸The general method for simulating consumption functions is described in Deaton (1989). The analysis in this paper is complicated somewhat by the fact that solutions must be found for two season-specific consumption functions.

simplifying assumption that income in each season follows a Bernoulli distribution, with a 50% chance that realized income is “high” and a 50% chance that it is “low”. Our base-line case (Case 1) is one in which income in the harvest season 2 equals either 50 or 150, and income in the growing season 1 is either 10 or 30, so that on average harvest-season income is five times growing-season income, which is consistent with the evidence from the ICRISAT data. We chose r to equal 0, which would be the appropriate interest rate in a pure stocking model (with no depreciation), and δ to equal .05. The preference parameter ρ is set to 3. We then vary the spread of income in each season, first allowing the coefficient of variation of growing season income to drop from .5 to .25 (Case 2), and then allowing the coefficient of variation of harvest season income to drop from .5 to .25 (Case 3). In all of our calculations, we choose parameters such that $Y_2^L > \beta^{-1/\rho} Y_1^H$, implying that borrowing constraints never bind in the harvest period.

The top panel of Figure 3 graphs consumption in each season against cash-on-hand, for Cases 1 and 3. (The consumption functions for Case 2 are similar.) In each case, harvest season consumption is less than cash-on-hand: borrowing constraints never bind in the harvest season, and farmers always carry positive assets into the following growing season. In the growing season borrowing constraints may bind. The threshold levels of cash-on-hand at which borrowing constraints just bind are shown as vertical lines on the graph. If cash-on-hand in the growing season is less than the threshold level, the farmer consumes all cash on hand in the growing season and enters the following harvest period with no assets. The levels of cash-on-hand at which borrowing constraints just bind are in both cases more than twice average growing-season income.

For each of the three cases, we simulated 5000 “years” of realizations for income, consumption, and assets, and these realizations are used to determine the likelihood that borrowing constraints bind, and the extent of consumption seasonality. The lower panel of Figure 3 graphs 100 “years” of realizations for consumption and cash-on-hand for Cases 1 and 3, and Table 3 presents descriptive statistics on consumption, assets, and borrowing constraints, using the data from all 5000 “years” for all three cases. These results illustrate the following points:

(i) *Borrowing constraints* Borrowing constraints bind infrequently. The infrequency of binding constraints can be seen from the graphs in the lower panel of Figure 3, which show that consumption equals cash-on-hand only occasionally, and from the summary statistics in Table 3, which show that constraints bind only 3% to 10% of the time. The results also indicate that reducing the variance of either harvest-season or growing-season income can increase the likelihood that borrowing constraints bind (contrast Cases 2 and 3 with Case 1). As the variance of income falls, farmers hold lower levels of assets on average, and it is more likely that cash-on-hand in the growing season will fall below the critical level at which borrowing constraints bind. For example, a reduction in the coefficient of variation of harvest-season income from .5 to .25 (i.e. moving from Case 1 to Case 3) reduces average cash-on-hand in the growing season from 158.5 to 92.5, and increases the probability that borrowing constraints bind from 3.3% to 10.1%. An important conclusion to draw from this result is that the prevalence of borrowing constraints is not a good indicator of welfare, since welfare-improving reductions in income variability (holding average income fixed) can result in more frequently binding borrowing constraints.

(ii) *Seasonal consumption patterns* In all three cases, consumption displays systematic seasonal variation. Specifically, consumption falls on average between the harvest and growing season, and rises from the growing season to the harvest. The growth in consumption into the harvest season can be accounted for by the precautionary motive, which tends

to depress consumption in the low-variance growing season relative to the high-variance harvest season, and by borrowing constraints, which at times prevent the household from consuming as much during the growing season as it would if it could borrow. Although consumption displays systematic seasonal variation, the amount of seasonal variation is small, on the order of a 2% change in consumption between seasons (see Table 3).

In summary, the simulation results show that even in the complete absence of credit markets and the presence of substantial income seasonality, buffering behavior may be used to nearly eliminate consumption seasonality. An implication is that our empirical results from the previous section do not necessarily indicate that credit markets work well within these villages, and in fact may be consistent with the buffer-stock model. Our finding that seasonal patterns in consumption are similar for farm and nonfarm households, despite different seasonal income patterns, could simply reflect the fact that farm households smooth consumption through buffering.

4. Buffer stocks vs. credit markets: some suggestive evidence

The theory and simulations discussed in the previous section indicate that the finding that consumption expenditure does not track income across months, despite large seasonal variations in income for some households, does not by itself constitute persuasive evidence that credit markets function well. We therefore turn, in this section to other sources of evidence regarding the relative importance of buffer stocks and credit markets in enabling households to smooth consumption under income seasonality. We first examine direct evidence on the particular mechanisms—transactions in financial markets, both formal and informal, purchases and sales of physical assets, accumulation and decumulation of grain stocks and cash reserves—households rely on to smooth consumption in the face of income seasonality. We then implement a more formal test of the permanent income hypothesis (with complete credit markets). Data constraints limit what we are able to say, but the additional evidence we consider is indicative of buffering behavior and suggests that borrowing constraints may be operative in two of the three villages.

4.1. Direct evidence on the mechanisms of consumption smoothing

On a month to month basis, we do not directly observe the stocks of financial and physical assets, grain and cash held by the households. However, at each interview (roughly about every four to six weeks), households were asked to provide details of all their transactions in financial, asset and grain markets since the previous interview. We use these data to construct two measures of monthly changes in stocks of financial and physical assets. The first is a measure of the net change in financial assets (excluding cash reserves), and is defined as deposits minus withdrawals from banks and other financial institutions, minus net increases in debt (from both formal and informal loan markets.) Changes in cash balances are not recorded, and so are not included in this measure. To obtain our second measure, the “net change in financial and physical assets,” we add to the net change in financial assets, the value of purchases minus sales of land, buildings, equipment, livestock, durable goods, and jewelry. This is not a complete measure of the net change in wealth, since it does not include changes in stocks of grains or cash balances. If the asset transactions data are accurate, then the difference between income and consumption that is not accounted for by the measured net change in financial and physical assets must (by definition) be accounted for by changes in either stocks of cash or grains.

Table 4 provides descriptive statistics on the two measures of changes in physical and financial assets. The top panel of Table 4 presents information on the percentage of observations with month-to-month changes in financial assets or in net financial and physical assets. The percentage of household reporting increases or decreases in net financial assets in the period before the survey (roughly 4 to 6 weeks) ranges from 18% to 41%, with more households reporting decreases than increases.⁹ Farm households appear to be somewhat more active in financial markets than nonfarm households in Shirapur and Kanzara. The percentage of households with changes in net financial and physical assets is higher, but shows the same asymmetry between increases and declines for all villages except Kanzara.¹⁰ Walker and Ryan (1990, Table 7.2) presents information on credit market transactions that is disaggregated by the source of loans, and finds that formal sector loans are relatively rare in Aurepalle and Shirapur, with over 70% of total lending (in rupees) coming from informal sources between 1976-77 and 1984-85. In Kanzara the situation is somewhat different, with formal sector loans accounting for 70% of total lending (by value of loans) and for about 40% of credit market transactions. The bottom panel of Table 6 provides median asset changes conditional on asset changes being non-zero. Farm households have larger median asset changes (either positive or negative) than nonfarm households.

Table 4 indicates that households are fairly active in asset markets. But are transactions in financial and physical asset markets the main mechanism for smoothing consumption? If so, there should be strong month effects in the measured changes in financial and physical assets for farm households. Why? Because the results of the second section suggest that there should be strong month effects in the total savings (i.e., income minus consumption) of farm households. And that is confirmed through regressions of total savings on a set of month effects and month-farm interactions which indicate that there are indeed significant month effects in total savings for both farm and non-farm households and that these month effects differ significantly across the two groups. In particular, one might expect to find that during the growing season farm households either draw down financial assets or take out loans, and then either repay loans or save at harvest time. Walker and Ryan (1990) discuss how smaller farm households often rely on moneylenders to provide growing-season credit, which is repaid at the time of threshing. Furthermore, since farm and nonfarm households have different income patterns but similar consumption patterns, seasonal asset changes should differ across these two groups of households.

To test these propositions we regressed our measure of changes in financial and physical assets on a set of month dummies and a set of month-farm interactions. In all villages, the differences in the month effects in measured asset changes between farm and nonfarm households are jointly insignificant. The conclusion that seasonal consumption smoothing is not accomplished through transactions in financial and physical asset markets is reinforced in Figure 4. There, we plot the monthly averages of total savings and of measured changes in financials and physical assets for farm households in each of the three villages. Graphs

⁹There is, of course, no reason why the percentage of observations that have asset increases must match the percentage of observations with asset declines. However, if it were the case that households take out loans and then gradually repay them over time, one would expect to see the *opposite* pattern. Underreporting of small asset transactions is a possibility.

¹⁰Note that a non-zero net change in financial assets *need not* imply a non-zero net change in the sum of financial and physical assets. Whether it does or not depends on the extent to which changes in financial assets reflect portfolio adjustments rather than ex-post consumption smoothing. For instance, if households were primarily adjusting their portfolios, e.g., drawing down financial assets to finance the purchase of physical assets, the decrease in financial assets would be offset by an increase in physical assets, resulting in no net change in the sum of physical and financial assets.

for nonfarm households are not shown, but indicate much less seasonal variation in saving. As anticipated by the regression results we report above, there appear to be strong seasonal patterns in total savings, but in all three villages, there is little indication that transactions in financial and physical asset markets account for this seasonal variation.

It is of course possible that households underreport transactions in financial and real asset markets. It is also possible that our results are contaminated by outliers which obscure seasonal patterns in the data.¹¹ And so our finding that financial and real asset markets play a limited role in helping households smooth consumption needs to be treated with caution. However, we can provide affirmative evidence that households do rely on accumulation and decumulation of grain stocks to buffer consumption from seasonal income variation. Figure 5 shows monthly averages of the value of crop output and crop sales for farm households in each of the villages. Although the seasonal patterns of crop sales and output are similar, crop sales are more smoothly spread out across months. Overall, this evidence supports the idea that seasonal consumption smoothing is accomplished through changes in stocks of grains and cash.

4.2. Additional tests for complete credit markets

To more formally distinguish the buffer-stock model from the permanent income model, we next turn to an additional implication of the buffering model which provides a clearer alternative against which the permanent income model can be tested. The permanent income model implies that changes in consumption between periods t and $t + 1$ should be orthogonal to any variables whose values are known prior to the realization of income in period $t + 1$: changes in consumption should only reflect new information about permanent income. This is not an implication of the buffer stock model, since borrowing constraints at times prevent consumers from fully adjusting consumption in response to anticipated income changes. Specifically, consumers with low levels of cash-on-hand who want to borrow (but cannot) will experience larger average consumption changes than in periods in which constraints are not binding, and consumption changes will display “excess sensitivity” to anticipated income changes.

This is a well-known difference between the two models and a common way of testing for it has been to implement some version of an excess-sensitivity test, for instance by testing whether changes in consumption observed in household data are correlated with instruments for anticipated income changes. The problem with this test however, is that it only reveals borrowing constraints when such constraints are binding. But as our simulations indicate, given the buffering behavior of households, borrowing constraints rarely bind: under the parameter configurations we use, assets are drawn down to zero only in the growing season, and then too, only between 3% to 10% of the time. We therefore propose, and then implement using the ICRISAT data, a slightly different version of the excess-sensitivity test which addresses this issue.

The test we propose builds on a key insight that emerges from the analysis of the buffering model in Deaton (1991) and is also borne out in our simulations. It is that while buffering behavior appears to delink consumption from income over short intervals so that consumption appears fairly stable even with fluctuating income flows, over longer intervals, consumption and income are more closely linked. What this suggests is that under a buffering model we expect to only observe consumption smoothing over short intervals

¹¹We were unable to apply standard techniques for dealing with outliers (such as looking at seasonal patterns in median asset changes) since the median household in each month reports *no* change in assets.

— what Deaton (1991) terms “high frequency” savings — in contrast, traditional life-cycle models, of which the permanent income model is a special case, imply both high-frequency saving and dissaving, as well as “low frequency” consumption smoothing over longer periods. This is, in principle, a testable difference between the two models if the relevant intervals can be appropriately designated. In environments with pronounced income seasonality a natural distinction is suggested between high frequency consumption smoothing which occurs within an income cycle (i.e. a full harvest-planting cycle) and longer-term consumption smoothing, across income cycles. An obvious way to test for this difference then is to apply the standard excess-sensitivity test first to consumption changes within an income-cycle and then to consumption changes across income-cycles.

To ascertain whether this difference is likely to be empirically discernable we again turn to simulations. For each of the three cases described above, we constructed a series of consumption, income and cash-on-hand variables for 100 households for 10 years (where each “year” contains a growing season and a harvest season) using randomly generated income numbers and the simulated consumption functions shown in Figure 3.¹² These “data” are used to estimate standard Euler equations. We first regress the change in consumption on the lagged value of cash-on-hand, and then regress the change in consumption on the change in income, where the income change is instrumented by lagged assets. Under the null of the PIH, neither lagged assets nor the predicted change in income will affect the change in consumption. With borrowing constraints, we expect that the change in consumption will be negatively correlated with lagged assets, and positively correlated with anticipated income changes. Moreover, we expect that any excess-sensitivity finding will be stronger when the data are aggregated up to a year.

Table 5 shows parameter estimates from these regressions, using simulated data. Columns 1 and 2 contain results when consumption and income changes are measured at the seasonal level. The change in consumption is defined as consumption in the current season minus consumption in the preceding season, and the change in income is defined in a similar manner. Lagged cash-on-hand is the value of cash-on-hand held at the beginning of the preceding season. In columns 3 and 4 the data are aggregated up to the annual level, i.e., over seasons. The change in consumption is therefore annual consumption in the current year minus annual consumption in the past year, and the change in income is defined similarly. Lagged cash-on-hand is defined as cash-on-hand at the beginning of the past year.

As expected, the PIH is rejected using these simulated data. Consumption changes are not orthogonal to lagged cash-on-hand, and the change in consumption is positively related to the predicted change in income. However, the sensitivity of consumption changes to lagged assets and to predicted income changes varies considerably across the three cases, and it is not surprising that the coefficient on predicted income is largest in Case 3, where the frequency that borrowing constraints bind is the highest.

The parameter estimates also depend on whether consumption changes are measured at the seasonal or annual level. The coefficients on lags of cash-on-hand and on predicted incomes are smaller, in absolute value, when consumption changes are measured seasonally (columns 1 and 2) than when the changes are measured over annual intervals (columns 3 and 4). This is, as we suggested above, due to the fact that borrowing constraints bind only in one of the two seasons. When consumption changes are measured at the seasonal level, the

¹²Although we use “year” to refer to the longer-interval containing the growing and harvesting seasons, a more appropriate term might be “crop-cycle” since in two of the sample villages there are actually two “crop-cycles” within an agricultural year.

parameter estimates represent a blend of little excess sensitivity between a harvest season and the following growing season, and higher excess sensitivity between a growing season and the following harvest season, since it is only in this period that borrowing constraints ever bind.

These results correspond well with our earlier discussion. We therefore estimate equations similar to those shown in Table 5 using the ICRISAT data, and using several different time intervals to measure consumption and income changes. These results are presented in Table 6. We regressed consumption changes on income changes separately for each village. In the first column of Table 6, consumption and income changes are measured at the monthly level. The instruments for the change in income include the household's monthly income in the same month of the previous year, and a measure of the area cultivated in the previous month interacted with a set of month dummies. Since all of these variables were known by the household before the measured consumption change took place, they are all valid instruments for income. Under the null of the PIH, the change in income predicted by these variables should not be related to the change in consumption. In column 2 of Table 6, consumption and income changes are measured at the seasonal rather than the monthly level. Consumption and income are first aggregated into two seasons, one spanning May to October (the "kharif" season) and the second spanning November to April (the "rabi" season). For Aurepalle and Kanzara, each of these seasons includes a full cycle of planting and harvesting. Consumption and income changes are measured across seasons, and the instrument list consists of the change in income experienced in the same season in the previous year, as well as measures of cultivated area in each season of the previous year. In column 3, the data are further aggregated up to the annual level, with income and consumption changes measured from year to year. For this last case, the instruments for the income change consist of a vector of lagged changes in assets, including land-holdings, livestock, farm implements, liquid assets, and total wealth. These variables were chosen as instruments because they predict changes in income well, and were known by households before the consumption change took place.

The results provide only modest evidence against the PIH, and the results are very sensitive to the time period over which consumption changes are measured. In all three villages, when consumption changes are measured at the monthly level, predicted income changes are not significantly related to consumption changes. The set of lagged income changes and lagged cultivated area used as instruments are good predictors of monthly income changes (Test 2), but are only weakly related to monthly consumption changes (Test 1). Thus, using monthly data, the PIH can be rejected for no village.

When consumption and income are measured at the seasonal level, the evidence against the PIH is somewhat stronger for two of the villages. In Aurepalle and Shirapur, the IV estimates yield "excess sensitivity" measures of .13 and .09, respectively, and both of these coefficients are significant. Furthermore, in both of these villages the overidentifying restrictions cannot be rejected at the 5% level. For these two villages, it appears that seasonal consumption changes are related to predictable seasonal income changes. The results for Kanzara are different. The set of instruments are good predictors of seasonal income changes, but are unrelated to seasonal consumption changes. As discussed above, Kanzara is the richest of the three villages, and has the most active financial markets, so these results may not be surprising.

The results when consumption and income changes are measured annually are mixed. In Aurepalle, lagged asset changes are related to both annual consumption and income changes (Tests 1 and 2), but consumption changes are not related to predicted income

changes, and the overidentifying restrictions are strongly rejected. Although these results provide evidence against the PIH, the coefficient on the predicted income change is not what would be expected from a buffering model. The annual results for Shirapur and Kanzara do not provide evidence against the PIH. Predicted income changes appear to have no effect on consumption, and lagged asset changes are not highly significant in the reduced form consumption change equations. However, for these two villages lagged asset changes do not predict income changes particularly well, and the results may be explained by a weak set of instruments. Overall, these results are not inconsistent with results other researchers have found using ICRISAT data at the annual level. For example, Morduch (1990), using data on annual food consumption, finds the strongest evidence of borrowing constraints in Aurepalle, some evidence of constraints in Shirapur, and no evidence of constraints in Kanzara.

5. Conclusions

This paper has investigated whether rural Indian households who experience seasonal income variation also experience seasonal consumption variation. The major finding of the first section of the paper is that, at the household level, seasonal consumption and income patterns are largely unrelated. The seasonal consumption patterns of different groups of households are similar despite dramatic differences in the timing of income flows.

One might infer from these results that seasonal credit markets work well within these villages. However, there is both empirical and theoretical evidence that this need not be the case. First, farm households, who experience the largest seasonal income variation, appear to make little use of financial markets to smooth consumption across months. Instead, there is (somewhat sketchy) evidence that farm households accumulate and draw down stocks of cash and grains over the course of the year to smooth consumption. Although this evidence does not prove that specific households face borrowing constraints, it suggests that borrowing constraints may be operative. Second, our simulations of the consumption behavior of cautious households who are barred from borrowing yield consumption patterns that are quite smooth across seasons, with an average increase in consumption from the growing season to the harvest season of only about 2%. Stocking behavior can potentially nearly eliminate systematic seasonal consumption patterns. Third, we find some evidence of “excess sensitivity” of consumption changes to anticipated income changes in two of the three villages. For these two villages, the permanent income hypothesis can be rejected in favor of the buffering model. For these three reasons we would be hesitant to claim that credit markets function perfectly within these villages.

We also find that there are seasonal patterns in consumption common across households within villages. These common seasonal patterns could be due to seasonal patterns in preferences common to all households, or to seasonal patterns in prices (which could in turn be due to seasonal variation in aggregate village income). Without knowing which of these factors is responsible for seasonal consumption variation, it is impossible to determine whether reducing seasonal consumption variation is an appropriate policy goal, and, if so, what policies should be used to achieve a reduction. If seasonal preference variation drives seasonal consumption variation, then policy interventions are clearly unnecessary. If, on the other hand, seasonal consumption variation is due to seasonal price variation, then market interventions may be appropriate, provided there is more seasonal price variation than is efficient. More research is needed on the sources and policy implications of aggregate consumption seasonality.

References

- [1] Behrman, Jere R. (1988). "Nutrition, birth order and seasonality: Intrahousehold allocation among children in rural India", *Journal of Development Economics*, 28(1), February, pp.43-63.
- [2] Besley, T. (1994). "How Do Market Failures Justify Interventions in Rural Credit markets?", *World Bank Research Observer*, 9(1), January, pp.27-47.
- [3] Canagarajah, Sudharshan and Stephen Pudney (1993). "Short-term Seasonal Consumption Behavior, Seasonality, and Labor Market Uncertainty in Rural India," *manuscript*.
- [4] Chambers, Robert, Richard Longhurst, and Arnold Pacey (1981). *Seasonal Dimensions to Rural Poverty*. London: Frances Pinter Limited.
- [5] Chaudhuri, Shubham (1999). "Forward-looking behavior, precautionary savings, and borrowing constraints in a poor agrarian economy: tests using rainfall data," *Columbia University Department of Economics Discussion Paper #9899-10*, August.
- [6] Deaton, Angus (1990). "Savings in Developing Countries: Theory and Review," *World Bank Economic Review* 4 (Proceedings volume).
- [7] Deaton, Angus (1991). "Saving and liquidity constraints," *Econometrica*, 59, 1221-48.
- [8] Deaton, Angus (1992). *Understanding Consumption*, Oxford: Clarendon Press.
- [9] Jacoby, Hanan and Emmanuel Skoufias (1997). "Risk, Financial Markets, and Human Capital in a Developing Country," *Review of Economic Studies*, 64(3), July, pp.311-335.
- [10] Jacoby, Hanan and Emmanuel Skoufias (1998). "Testing Theories of Consumption Behavior Using Information on Aggregate Shocks: Income Seasonality and Rainfall in Rural India," *American Journal of Agricultural Economics*, 80(1), February, pp.1-14.
- [11] Morduch, Jonathan (1990). "Risk, Production and Saving: Theory and Evidence from Indian households," *manuscript*.
- [12] Parker, Jonathan A. (1999). "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, 89(4), September, pp.959-973.
- [13] Paxson, Christina H. (1993). "Consumption and Income Seasonality in Thailand," *Journal of Political Economy*, 101 (1): 39-72.
- [14] Rosenzweig, Mark (1988). "Risk, Implicit Contracts and the Family in Rural Areas of Low-Income Countries," *Economic Journal*, 98(3), December, pp.1148-1170.
- [15] Rosenzweig, Mark and Kenneth Wolpin (1993). "Credit Market Constraints, Consumption Smoothing and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India" *Journal of Political Economy*, 101(2), April, pp.223-244.

- [16] Sahn, David E.(ed.) (1989) *Seasonal Variability in Third World Agriculture: the Consequences for Food Security*, Baltimore: Johns Hopkins University Press.
- [17] Shapiro, Matthew D. and Joel Slemrod (1995). "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding," *American Economic Review*, 85(1), March, pp.274-283.
- [18] Souleles, Nicholas S. (1999). "The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, 89(4), September, pp.947-958.
- [19] Townsend, Robert (1994). "Risk and Insurance in Village India," *Econometrica*, 62: 539-591.
- [20] Walker, Thomas S. and James G. Ryan (1990). *Village and Household Economies in India's Semi-Arid Tropics*. Baltimore: Johns Hopkins University Press.

Appendix A

The data were collected by resident-investigators who lived in the villages during the collection process. The collection of the data was organized into schedules. We use the data from the plot-cultivation schedule, the transactions schedule, and the asset-inventory schedules. The plot-cultivation schedule records the dates and detailed input-output data of all production operations at the plot-level, as well as various plot characteristics, for all operational holdings of the sample households over the entire period of ten years. The schedule was maintained on a monthly basis. The transactions schedule, generally recorded every three to four weeks, catalogs all transactions (i.e., flows of goods, services and cash into or out of the household) reported by the households. Thus, purchases and sales of commodities, loans given or taken, wages paid or received, etc., whether in-kind or in cash, are recorded. In addition an effort was made to record various intra-household transactions. Transactions were classified under accounts such as cultivation, animal-husbandry, labor, financial, non-durable consumption, and durable consumption. Fairly disaggregated item-specific codes were used to identify the goods or services being traded. The coverage of the schedule was however reduced after 1982; specifically, transactions involving non-food or minor food items were excluded. To preserve the comprehensiveness of the consumption data, which we build up from the transactions schedule, we therefore use only the data up to the 1981-82 agricultural year.

The asset-inventory schedules provide a listing of the quantity and value of various assets held by the household, assets such as land, livestock, farm-machinery, grain stocks and jewelry. These schedules were updated only once a year, at the beginning of the crop-year in June. However, information about asset-holdings at the beginning of the 1975-1976 crop-year was only obtained retrospectively in December, 1975. For this reason, but also because the 1975-76 crop-year was the first, and hence, trial year of the survey when a lot of the details were being worked out, we do not use the data from the first year.

Other than the asset-holdings data which were straightforwardly compiled from the relevant schedules, all the variables that we use in the empirical analysis were constructed from the raw data in the primary schedules. Monthly income was calculated using data from the plot-cultivation and transactions schedules. From the plot cultivation schedule we obtained the value of total crop output during the month, as well as the cost of all material inputs and draft and human labor.¹ The difference between the two provided a measure of net income from agriculture. To this we added the net inflows of cash (or cash-equivalent in-kind flows) from animal husbandry, trade and handicrafts-related activities, and off-farm labor which were reported in the transactions schedule for that month. The imputed value of family labor and own bullock time was recorded as a cost in the plot-cultivation schedule by ICRISAT, but was not reported as income in the transactions schedule. We therefore subtracted these imputed costs in arriving at a final figure for the net monthly income of each household.

The estimates of monthly consumption expenditures were constructed entirely from the transactions schedule.² We constructed two separate series, one on food expenditures, and one which included food and other non-durable items. Ceremonial expenditures were excluded from both measures. The expenditure data had to be allocated to specific months since the information on the transactions schedule was collected irregularly. We did this by calculating the time interval between the current and previous survey date, and allocating expenditure to a specific month according to the fraction of the time interval that fell in that month. This procedure for allocating consumption expenditure may make expenditure appear to be somewhat smoother across months than it actually is, and this should be kept

in mind when interpreting the results that follow. However, the time intervals between surveys were generally quite short, with 91.6% of the observed time intervals equal to one month or less, and 99.7% with intervals of two months or less.

Appendix B

We show in this appendix that borrowing constraints will never bind in period 2 if $Y_2^L > \beta^{-1/\rho} Y_1^H$. From (3.7), a borrowing constraint will bind in period 2 for values of cash-on-hand a_2 such that:

$$a_2^{-\rho} > \beta \int_{Y_1^L}^{Y_1^H} c_1(Y_1)^{-\rho} dF_1(Y_1) \quad (\text{A.1})$$

Note that $c_1(Y_1)$ cannot exceed Y_1^H , since Y_1 cannot exceed Y_1^H and consumption cannot exceed cash-on-hand. This implies that:

$$\beta \int_{Y_1^L}^{Y_1^H} c_1(Y_1)^{-\rho} dF_1(Y_1) \geq \beta (Y_1^H)^{-\rho} \quad (\text{A.2})$$

Combining (5.1) and (5.1), a necessary condition for a borrowing constraint to bind in period 2 is that a_2 is such that:

$$a_2^{-\rho} > \beta (Y_1^H)^{-\rho} \iff a_2 < \beta^{-1/\rho} Y_1^H \quad (\text{A.3})$$

The least a_2 that can equal is Y_2^L . It follows that if $Y_2^L > \beta^{-1/\rho} Y_1^H$, it is impossible that (5.1) is true for any value of a_2 , and thus, borrowing constraints will never bind in period 2. Note that this is a sufficient but not necessary condition.

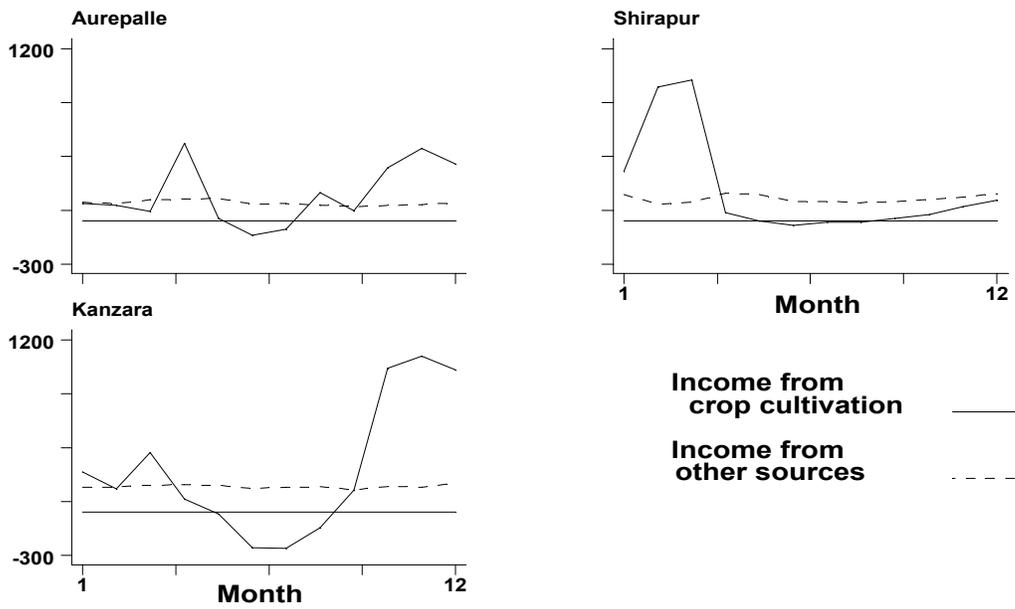


Figure 1
Seasonal patterns in income from different sources

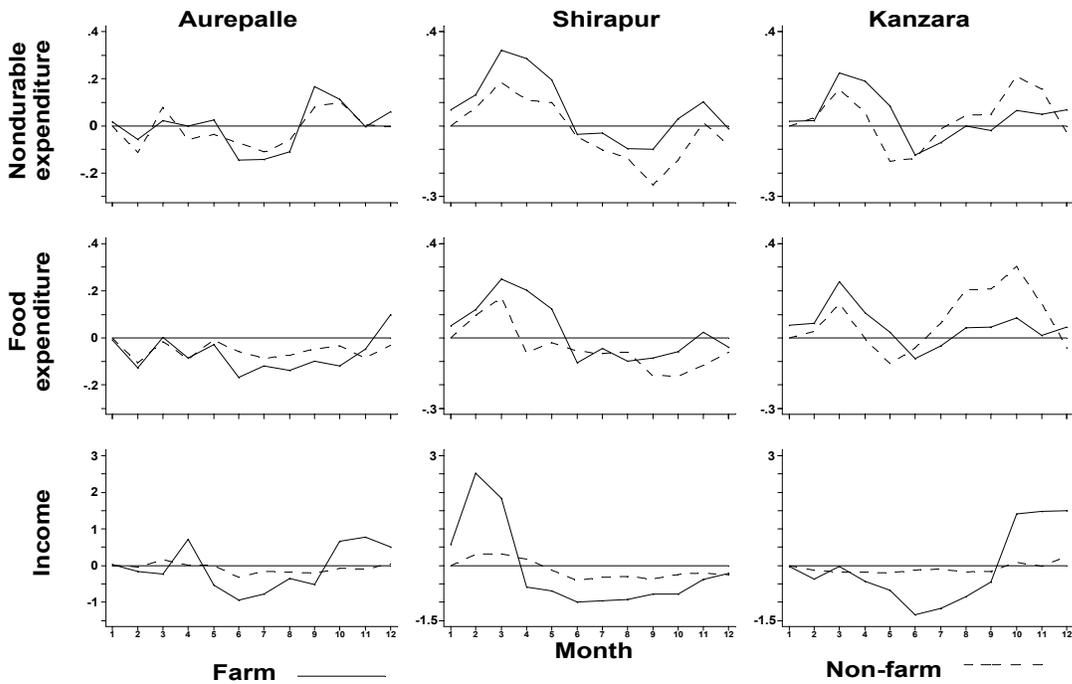
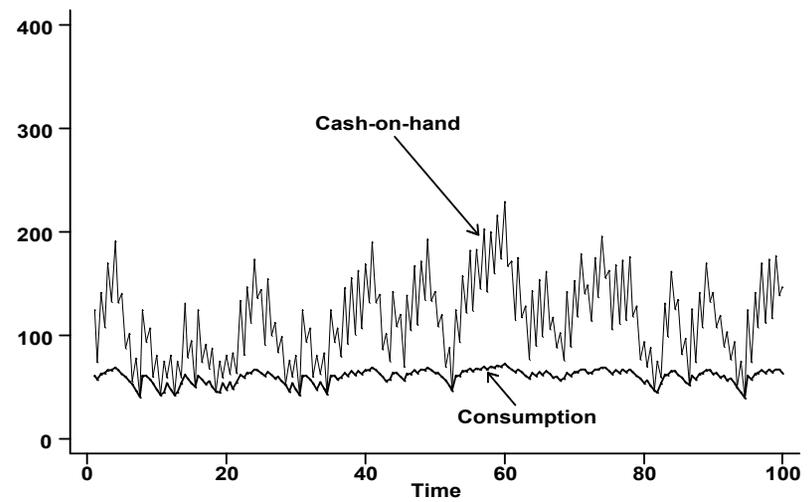
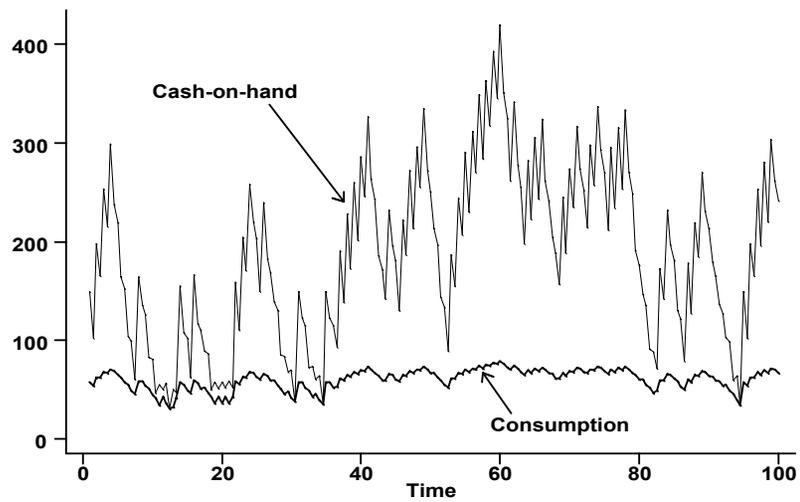
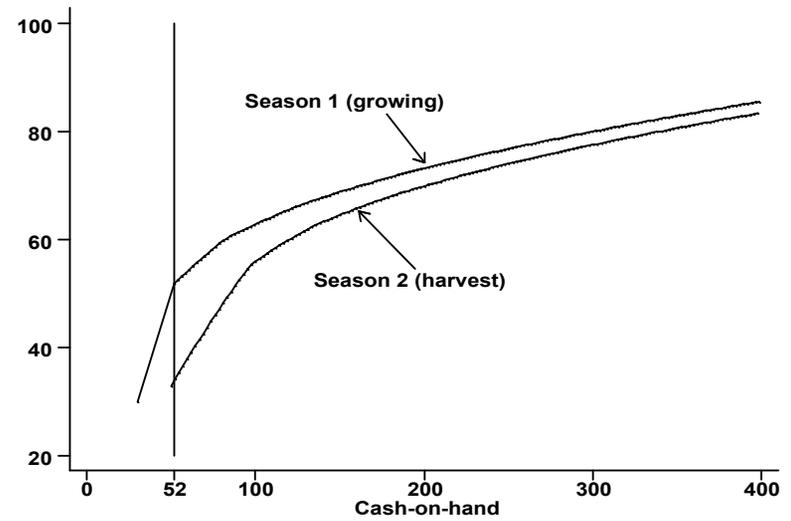
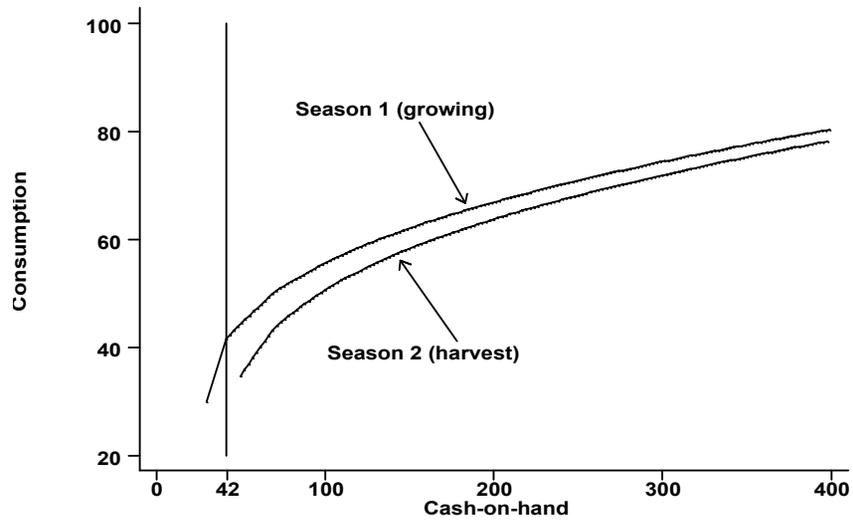


Figure 2
Month effects in income and expenditure



Case 1: Growing season income = 10 or 30
Harvest season income = 50 or 150

Case 3: Growing season income = 10 or 30
Harvest season income = 75 or 125

Figure 3
Simulated consumption functions (top panel) and values of consumption and cash-on-hand for 100 "years" (bottom panel)

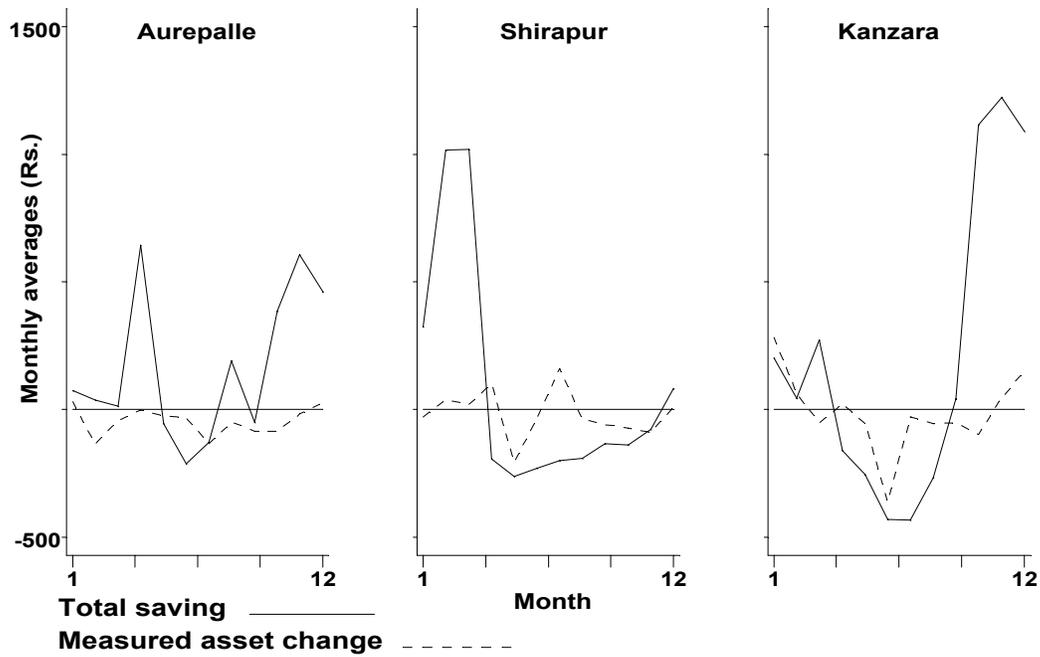


Figure 4
Seasonal patterns in total savings and in measured asset changes

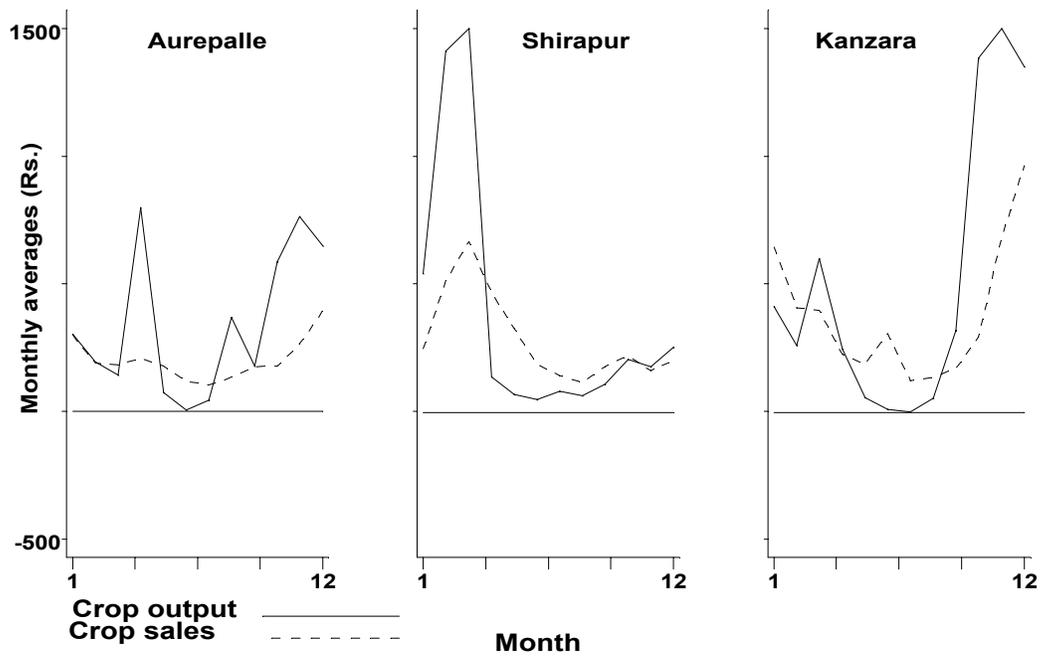


Figure 5
Seasonal patterns in crop output and in crop sales

Table 1
Sample means of income, consumption expenditures and related variables

Village	Aurepalle		Shirapur		Kanzara	
	Non-farm	Farm	Non-farm	Farm	Non-farm	Farm
Range of gross cropped area (ha.)	0 - 0.5	> 0.5	0 - 0.5	> 0.5	0 - 0.5	> 0.5
Monthly income	121.1	385.4	185.0	446.3	184.8	554.0
Gross cropped area (ha.)	0.05	4.48	0.00	6.98	0.03	6.15
Percentage of area irrigated	0.0	13.5	0.0	12.6	0.0	4.4
Percentage of annual income from crop cultivation	1.3	50.4	-1.3	61.3	0.0	47.4
Percentage of annual income received in top 3 months	39.1	61.3	46.3	81.8	41.3	71.5
Monthly expenditure, nondurables	145.4	221.9	242.4	362.1	207.4	350.6
Monthly expenditure, food	105.0	143.7	148.7	222.6	121.4	189.3
Percentage of annual expenditure on food	73.5	68.5	62.6	62.1	60.7	56.1
Percentage of annual nondurable expenditure incurred in top 3 months	35.5	38.4	41.8	37.3	36.3	38.0
Percentage of annual food expenditure incurred in top 3 months	34.8	37.5	39.9	35.7	36.7	38.1
Sample size (household-months)	660	1716	708	1668	552	2040

Notes: Average monthly incomes and expenditures are in 1983 rupees. The percentage of income/expenditure in top 3 months is calculated as the sum of monthly incomes/expenditures in the 3-highest income/expenditure months of each year, divided by annual income/expenditure in that year.

Table 2
Test statistics for month effects in income and expenditure
(F-statistics, p-values in parentheses)

	Aurepalle		Shirapur		Kanzara	
Observations (household-months)	2328		2364		2568	
Income						
Test 1: no month effects, nonfarm	0.80	(0.636)	2.04	(0.021)	0.36	(0.970)
Test 2: no month effects, farm	45.17	(0.000)	98.28	(0.000)	80.42	(0.000)
Test 3: month effects identical	10.64	(0.000)	19.94	(0.000)	13.77	(0.000)
R-square	0.179		0.312		0.259	
Nondurable expenditure						
Test 1: no month effects, nonfarm	1.50	(0.125)	4.46	(0.000)	2.73	(0.002)
Test 2: no month effects, farm	6.62	(0.000)	12.71	(0.000)	8.10	(0.000)
Test 3: month effects identical	0.59	(0.841)	0.56	(0.859)	2.08	(0.019)
R-square	0.451		0.454		0.518	
Food expenditure						
Test 1: no month effects, nonfarm	0.33	(0.980)	2.37	(0.007)	3.12	(0.000)
Test 2: no month effects, farm	3.94	(0.000)	8.77	(0.000)	4.70	(0.000)
Test 3: month effects identical	0.76	(0.678)	1.35	(0.193)	2.45	(0.005)
R-square	0.413		0.444		0.445	

Notes: Test statistics are based on estimates of equations (2.1) and (2.2) in the text. In the top panel, Test 1 is an F-test for whether the month effects α_m in the income variable A_{imt} are jointly insignificant for nonfarm households.

Test 2 is an F-test for jointly insignificant month effects $(\alpha_m + \alpha_m^F)$ for farm households. Test 3 is an F-test for

whether the month effects in income are identical for farm and nonfarm households, i.e, whether α_m^F are jointly insignificant. The second and third panels repeat these test for $\ln(\text{expenditure})$ and $\ln(\text{food expenditure})$. The expenditure equations allow for different intercepts for farm and nonfarm households and the null hypothesis of Test 3 is that the difference between the expenditures of nonfarm households and those of farm households is *constant* across months (rather than equal to 0 for all months). In terms of equation (2.2), the test is that β_m^F are jointly insignificant. The expenditure equations control for $\ln(\text{annual income})$, a set of year dummies, and numbers of males, females, and children in the household. Figure 2 graphs the estimates of the month effects for farm and nonfarm households from these regressions.

Table 3
Descriptive Statistics: Simulated consumption and cash-on-hand values
(Standard errors in parentheses)

	Case 1	Case 2	Case 3
	$Y_1=10$ or 30	$Y_1=15$ or 25	$Y_1=10$ or 30
	$Y_2=50$ or 150	$Y_2=50$ or 150	$Y_2=75$ or 125
Season 1: growing season			
Income: Y_{1t}	19.99 (10.00)	19.99 (5.00)	19.99 (10.00)
Cash-in-hand: a_{1t}	158.50 (72.56)	150.87 (70.26)	92.52 (30.78)
Consumption: c_{1t}	59.59 (9.38)	59.56 (9.31)	59.63 (6.62)
Percentage of times borrowing constrained ($c_{1t} = a_{1t}$)?	0.033	0.051	0.101
Value of c_{1t} at which just constrained	41.58	39.32	51.77
Percentage change in consumption: $\ln c_{1t} - \ln c_{2,t-1}$	-0.0218 (0.047)	-0.0229 (0.033)	-0.0173 (0.057)
Season 2: harvest season			
Income: Y_{2t}	100.50 (50.0)	100.50 (50.0)	100.50 (25.0)
Cash-in-hand: a_{2t}	199.44 (80.96)	191.84 (79.18)	133.17 (35.22)
Consumption: c_{2t}	60.92 (9.75)	60.96 (9.68)	60.64 (6.62)
Percentage change in consumption: $\ln c_{2t} - \ln c_{1t}$	0.0217 (0.135)	0.0228 (0.136)	0.0171 (0.104)

Notes: These numbers are based on 5000 "years" of randomly drawn income values, and on the solutions for consumption functions described in the text. To construct income, we drew 5000 random numbers for each season from a standard normal distribution, and set income equal to its low value if the draw was less than 0 and to its high value if the draw exceeded 0. Each of the three cases uses of the same series of random draws. One hundred "years" of simulated values of consumption and cash-on-hand for Cases 1 and 3 are graphed in Figure 4.

Table 4
Descriptive statistics: asset transactions

Village	Aurepalle		Shirapur		Kanzara	
	Non-farm	Farm	Non-farm	Farm	Non-farm	Farm
Observations (household-months)	660	1716	708	1668	552	2040
Percent of observations with:						
Increase in net financial assets	6.36	7.28	12.29	16.13	6.34	10.88
Decrease in net financial assets	23.33	20.69	21.05	25.06	12.14	15.29
Increase in net financial & physical assets	6.97	9.91	16.53	21.46	11.96	20.29
Decrease in net financial & physical assets	26.52	30.19	25.14	31.00	15.76	18.43
Median values of changes in:						
Financial assets, if change > 0	50	260	50	234	100	312
Financial assets, if change < 0	-70	-150	-65	-200	-75	-215
Financial & physical assets, if change >0	50	117	60	230	47	200
Financial & physical assets, if change < 0	-64	-120	-70	-200	-96	-213

Notes: All information in this table is based on reports of asset transactions collected periodically (every 4 to 6 weeks) during each survey year. The change in net financial assets equals deposits minus withdrawals from banks and other financial institutions, minus net increases in debt. Changes in cash balances are not included. The change in net wealth equals the net change in financial assets, plus purchases minus sales of: land, buildings, equipment, livestock, durable goods, and jewelry. Capital gains and losses are not recorded.

Table 5
Estimates of Euler equations: simulated data

Dependent variable: change in consumption	Seasonal data		Annual data	
	(1)	(2)	(3)	(4)
Case 1				
Lagged(cash-on-hand)	-0.012 (19.20)		-0.046 (19.90)	
Change in income, instrumented with lagged(cash-on-hand)		0.059 (24.00)		0.150 (24.80)
Case 2				
Lagged(cash-on-hand)	-0.012 (18.80)		-0.048 (19.70)	
Change in income, instrumented with lagged(cash-on-hand)		0.058 (23.20)		0.148 (25.10)
Case 3				
Lagged(cash-on-hand)	-0.032 (24.00)		-0.113 (25.90)	
Change in income, instrumented with lagged(cash-on-hand)		0.100 (30.00)		0.262 (23.80)

Notes: These estimates are based on simulations of 10 years of data for 500 households (for each case). In columns 1 and 2, changes in consumption and income are measured as changes from season to season, and lagged cash on hand is cash on hand held at the beginning of the previous season. In columns 3 and 4, consumption and income is aggregated to annual values, and consumption and income changes are measured as changes in annual values from year to year. Lagged cash on hand is measured as cash on hand at the beginning of the preceding year. A dummy variable for the season is included in the regressions shown in column 1 and 2.

Table 6
Estimates of Euler equations: ICRISAT data
 Dependent variable: change in consumption
 (t-statistics in parentheses)

	Month	Season	Year
Aurepalle			
Coefficient on change in income (OLS)	-0.007 (1.13)	0.080 (4.09)	0.010 (0.26)
Coefficient on change in income (IV)	0.007 (0.42)	0.130 (3.29)	-0.090 (0.91)
F-tests (p-values shown)			
Test 1: instruments jointly insignificant in the reduced form consumption equation	0.0368	0.0005	0.0000
Test 2: instruments jointly insignificant in the first-stage income change equation	0.0000	0.0000	0.0000
Test 3: Test of over-identifying restrictions for the IV equation	0.0253	0.0615	0.0000
Shirapur			
Coefficient on change in income (OLS)	0.020 (4.61)	0.050 (2.35)	0.130 (3.23)
Coefficient on change in income (IV)	-0.030 (1.78)	0.090 (2.15)	0.100 (0.84)
F-tests (p-values shown)			
Test 1: instruments jointly insignificant in the reduced form consumption equation	0.8690	0.0587	0.0503
Test 2: instruments jointly insignificant in the first-stage income change equation	0.0000	0.0000	0.0031
Test 3: Test of over-identifying restrictions for the IV equation	0.9800	0.4140	0.0458
Kanzara			
Coefficient on change in income (OLS)	0.004 (0.53)	0.002 (0.14)	0.190 (4.43)
Coefficient on change in income (IV)	0.004 (0.29)	0.040 (1.64)	0.200 (1.44)
F-tests (p-values shown)			
Test 1: instruments jointly insignificant in the reduced form consumption equation	0.0174	0.2836	0.0670
Test 2: instruments jointly insignificant in the first-stage income change equation	0.0000	0.0000	0.0056
Test 3: Test of over-identifying restrictions for the IV equation	0.0110	0.7852	0.0830

Notes: In column 1, the dependent variable is the change in consumption between months. In column 2 it is the change in consumption between seasons, where a season is defined as a six month period running from May to October or from November to April. In column 3, it is the change in consumption between years. The change in income in each regression is defined using the same time frame as the change in consumption. The first row of each panel shows the coefficient (and t-statistic) from an OLS regression of the change in consumption on the change in income. The second row shows the coefficient from an IV regression of the change in consumption on the change in income. In column 1, the instruments consist of the income change from the same month in the previous year and a measure of cultivated area in the previous month interacted with a set of month dummies. In column 2, the instruments consist of the income change from the same season in the previous year, and measures of cultivated area in each season of the previous year. In column 3, the instruments consist of a vector of lagged asset changes, including land-holdings, livestock, implements, liquid assets, and total wealth. The consumption change equations also included a set of month dummies (column 1), season dummies (column 2) and year dummies (column 3). The three test statistics are defined as follows. Test 1 is an F-test for the joint insignificance of the instruments in the reduced form consumption change equation (i.e. the regression of the change in consumption on the time dummies and the set of instruments.) Test 2 is an F-test for the joint insignificance of the instruments in the first-stage income change equation. Test 3 is a test for whether the overidentifying restrictions implicit in the IV estimates of the consumption change equations can be rejected.