**Forward-looking behavior, precautionary savings,
and borrowing constraints in a poor, agrarian
economy: tests using rainfall data**

by

Shubham Chaudhuri, Columbia University

August 1999

Discussion Paper Series No. 9899-10

dp 9899-10
pages: 30

Forward-looking behavior, precautionary savings, and borrowing constraints in a poor, agrarian economy: tests using rainfall data

Shubham Chaudhuri*
Department of Economics
and
School of International and Public Affairs,
Columbia University

August 1999

Abstract

Households in many poor, agrarian economies derive their income from rain-dependent agriculture. This dependence raises the possibility that, from rainfall patterns early on in the crop-cycle, households accumulate information about future cash inflows (e.g., harvests) before the cashflows are realized. Using detailed longitudinal data from three villages in India, this paper explores whether households utilize this information in the ways suggested by modern consumption theories. Two of the central hypotheses suggested by these theories are that households are forward-looking in their consumption behavior and react optimally to the receipt of information; and that in the face of uncertainty about future incomes, households engage in precautionary saving. The basic challenge in implementing tests of these hypotheses has been to find empirical measures of the “news” that households receive, and the uncertainty they face. I show that, in the villages I study, rainfall patterns provide good proxies from which such measures might be constructed. I exploit this fact to test for precautionary saving and forward-looking behavior. I find that precautionary saving and forward-looking behavior do not, by themselves, fully explain the observed consumption patterns. However, when I explicitly incorporate the possibility of binding borrowing constraints into the testing strategy, I obtain fairly strong support for the joint hypothesis of precautionary saving and forward-looking behavior in the presence of borrowing constraints.

*Address: Department of Economics, Columbia University MC-3308, New York, NY 10027. E-mail: sc301@columbia.edu. I would like to thank, without in any way implicating them for any errors that remain, Abhijit Banerjee, Chris Carroll, Angus Deaton, Fumio Hayashi, Glenn Hubbard, Michael Kremer, Annamaria Lusardi, Jonathan Morduch, Chris Paxson and seminar participants at Columbia, Princeton, Brown, Harvard-MIT, and the NBER for providing comments and suggestions. I would also like to thank the International Crop Research Institute for the Semi-Arid Tropics for allowing me access to their data.

1. Introduction

Modern consumption theories emphasize the role of information accumulation and uncertainty in determining consumption patterns. Two of the central hypotheses suggested by these theories are first, that households are forward-looking in their consumption behavior and react immediately to the receipt of information about future income; and second, that in the face of uncertainty about future incomes, households engage in precautionary saving.¹ Whether households do in fact behave in the ways suggested, and the extent to which they do so, has implications for a number of macroeconomic policy issues, from the likely efficacy of various fiscal measures to the determinants of savings rates. A large number of papers over the last two decades have used household-level data to explore the empirical validity of these hypotheses.² The primary challenge has been to find empirical measures of the “news” that households receive, and the “uncertainty” they face. With the sort of household-level data that are usually available for the U.S. and other developed economies, convincing proxies have been hard to find.

The basic idea of this paper is that in poor, agrarian economies, where households derive their incomes from rain-dependent agricultural cultivation, good proxies may be constructed from data on rainfall patterns. Agricultural cultivation is a lengthy time-consuming process, drawn out over several months, and especially in rainfed areas, is subject to a variety of rainfall-related weather shocks at each stage of the crop-cycle. The dependence on rainfall raises the possibility that from rainfall patterns early on in the crop cycle, households may accumulate information about future cashflows (e.g., harvests) months before the cashflows are actually realized. For instance, the late arrival of monsoon rains (in June) might imply (everything else equal) poorer harvests (and hence, lower incomes) in December. If so, the late arrival of the monsoon conveys news (in this case, bad) about future incomes that forward-looking households ought to respond to, by adjusting their consumptions in June itself. Moreover, because, by definition, the arrival of news implies the resolution of at least some uncertainty, intertemporal and cross-sectional variation in the timing and volume of rainfall should capture some of the uncertainty faced by agrarian households. Thus, data on rainfall patterns, in principle, provide good—good, in the sense that they are observable, exogenous and vary over time and locations—empirical proxies for the “news” that households receive and the “uncertainty” they face. I implement this idea using longitudinal household data and rainfall data from three villages in India. I show that rainfall patterns do indeed provide good proxies for measures of news and uncertainty and exploit this fact to test for precautionary saving and forward-looking behavior.

Beyond providing a potentially clean test, this exercise is of independent interest. The incomes of agrarian households are notoriously volatile because of the significant year-to-

¹Strictly speaking, households that display a precautionary motive for saving increase their savings in response to increased uncertainty about their future *consumption* prospects. Except in some pathological cases, uncertainty about future income will translate directly into uncertainty about future consumption. Note also that risk-aversion, by itself, does not imply precautionary saving. Kimball (1990) explores the distinction between risk-aversion and precautionary saving behavior.

²Browning and Lusardi (1996) provide a very useful survey of this literature.

year fluctuations in rainfall in many parts of the world. Such volatility can, in the absence of insurance markets or other risk-sharing arrangements, impose significant welfare costs on these households and this has long been a source of concern for policymakers. A large literature has emerged over the last two decades examining the coping mechanisms poor, agrarian households use to deal with this volatility.³ This paper is also a contribution to that literature.

In the next section, I document the seasonality and volatility of incomes in the three villages. I also show that, in these villages, rainfall patterns provide information about future cash inflows before the cashflows are realized. Section 3 lays out the theoretical implications of informational accumulation for consumption behavior. The empirical strategy I pursue is described in Section 4. In Section 5 I present and discuss the results of the tests of forward-looking behavior. The results and discussion of the tests of precautionary saving are in Section 6. I find that precautionary saving and forward-looking behavior do not, by themselves, fully explain the observed consumption patterns. But when the testing strategy is modified to take into account the possibility of borrowing constraints, the modified tests provide fairly strong evidence in favor of the joint hypothesis of forward-looking behavior and precautionary saving in the presence of borrowing constraints. Section 7 concludes.

2. Income seasonality and the accumulation of information

2.1. The data and the environment

The data I use come from the Village-Level-Studies (VLS) longitudinal household surveys carried out by the International Crop Research Institute for the Semi-Arid Tropics, based in Hyderabad, India. Started in 1975, the VLS surveys initially covered six villages in three broad agro-climatic zones within the semi-arid region of India (i.e., south-central India). Three of the villages were dropped after 1980 at which time the survey was extended to four new villages. Routine data collection was discontinued in all villages after 1985. Daily rainfall data were, however, collected up to 1990. The cross-sectional coverage of the surveys is limited—forty households in each village. But the surveys provide detailed data on practically all aspects of production, consumption and asset-holdings of the sample households. 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.

I use the data from the three villages—Aurepalle in the state of Andhra Pradesh, and Shirapur and Kanzara in the state of Maharashtra—that were covered continuously between 1975 and 1985. However, due to concerns about reliability, I do not use the data from the first and last years (1975-76 and 1984-85) of the survey. Also, because the coverage of consumption expenditures was reduced after 1981-82, I only use the consumption data from the six years 1976-77 to 1981-82. Finally, I exclude any landless households from my sample. This leaves about thirty or so households in each of the three villages.

The primary source of income in all three villages is agricultural cultivation, which, because the villages are in rainfed areas where irrigation is relatively rare, is dependent on

³The survey chapter by Besley (1995) and the symposium in the *Journal of Economic Perspectives*, Case (1995) provide nice overviews of this literature.

the southwest monsoon rains. The agricultural crop cycle is therefore governed by rainfall patterns that are highly seasonal. The monsoon rains typically begin in early June and recede by mid-October. Figure (1) plots the average (over the period 1975-1990) monthly rainfall in each of the three villages. In Aurepalle and Kanzara, the main crop is planted with the onset of the monsoon. Harvesting begins in late September for short-duration crops and continues till the following March for longer-duration crops. Some households plant a second crop during the winter months that is harvested between March and May. In Shirapur, crops are typically planted after the monsoon and harvested in the first few months of the following calendar year.

I next establish some basic properties of household income in these villages where income is defined as the sum of profits from agricultural cultivation and net inflows from other activities including off-farm agricultural labor, animal husbandry and trades and handicrafts. For this purpose, and for the subsequent analysis, I divide the crop year into three seasons—the monsoon season (M) from June to September, the winter season (W) from October to January, and the summer season (S) from February to May. Table (1) presents an analysis of variance of seasonal income. Household and agricultural year effects are evident, but the main source of variation is seasonal. Figure (2) plots the season effects in income in the three villages. Incomes are highest during the main harvest period, which for Aurepalle and Kanzara is during the winter, and for Shirapur, during the summer. Incomes are lowest in all three villages during the monsoon.

The pronounced seasonality of income is accompanied in these villages by significant volatility in incomes from year to year. Table (2) displays the cross-sectional average, by season and village, of the (intertemporal) coefficient of variation of income for each household over the eight years of the panel. The average coefficient of variation ranges from 0.45 for monsoon season income in Kanzara to 1.27 for monsoon income in Shirapur. Incomes are clearly quite volatile from year to year.⁴ The dependence of agricultural incomes on rainfall, however, raises the possibility that, *within the crop year*, the volatility is partially anticipated—i.e., that, from rainfall patterns early on in the crop cycle, households are able to accumulate information about future cashflows (e.g., harvests) months before the cashflows are actually realized.⁵

2.2. Do rainfall-related weather shocks convey information about future income?

To explore this possibility I begin by positing the following process for the income of a household h in season j of crop-year t :⁶

$$Y_{htj} = X_{ht}\alpha^j + K_{ht}\beta^j + \sum_{i=1}^j R_{ti}\gamma_i^j + \sum_{i=1}^j (K_{ht} \otimes R_{ti})\delta_i^j + \sum_{i=1}^{j-1} \phi_i^j e_{hti} + e_{htj} \quad (2.1)$$

⁴Note, from the first column of Table (2), that the coefficient of variation of annual income is lower than that of income in any one season, which suggests that within the agricultural year, seasonal incomes tend to be negatively correlated. This may reflect variation in the timing of income receipts from year to year, or it may be due, as Kochar (1995) suggests, to household-level efforts to smooth *income* ex-post through, for instance, adjustments in labor supply. Still, even the volatility of annual income is high.

⁵For instance, the late arrival of the monsoon rains in June may signal poorer harvests during the winter season.

⁶From here on I use $\{1,2,3\}$ and $\{M, W, S\}$ interchangeably to index seasons.

Here X_{ht} and K_{ht} are vectors of the household's characteristics *at the beginning of crop-year t* , characteristics that capture sources of variation in Y_{htj} (across households and crop-years) that are unrelated to the gradual resolution of uncertainty over the course of the crop-cycle. X_{ht} includes the household's total wealth, liquid assets, the age of the household head, the schooling of the household head, and the number of adults in the household, while K_{ht} includes the land area owned by the household, the proportion of this area that is irrigated, and the number of bullocks owned by the household. R_{ti} is a vector of rainfall-related weather *shocks* that are realized (*and are hence assumed to enter the household's information set*) in season i , where $i \leq j$. The key assumption implicit in (2.1) is that R_{ti} potentially affects income, not only in season i , but also in later seasons. $(K_{ht} \otimes R_{ti})$ is the vector of interactions between K_{ht} and R_{ti} . Lastly, e_{hti} represents an unobservable (to us) non-weather-related idiosyncratic shock—e.g. ill health, localized crop damage—that is realized in season i but, again, potentially affects income in later seasons as well.

Several points need to be made about (2.1). First, note that this specification of the seasonal income process captures the potential accumulation of information over the course of the crop-year. This can be seen by rewriting (2.1) as:

$$Y_{htj} = X_{ht}\alpha^j + K_{ht}\beta^j + \sum_{i=1}^j \eta_{hti}^j \quad (2.2)$$

where, with the normalization, $\phi_j^j = 1$:

$$\eta_{hti}^j = R_{ti}\gamma_i^j + (K_{ht} \otimes R_{ti})\delta_i^j + \phi_i^j e_{hti} \quad (2.3)$$

A household's income in season j of year t is therefore posited to be the sum of a benchmark level— $X_{ht}\alpha^j + K_{ht}\beta^j$ —determined by its characteristics at the beginning of the year, and a series of shocks— η_{hti}^j , $i = 1, \dots, j$ —information about which the household accumulates in the seasons leading up to and including season j . In other words, η_{hti}^j represents the news that the household receives in season i , where $i \leq j$, about its income in season j . The idea here is that a household begins the crop-year with a base information set $\{X_{ht}, K_{ht}\}$, which is supplemented in each season i with information obtained as various rainfall-related (R_{ti}) and idiosyncratic shocks (e_{hti}) are realized. For instance, the onset date of the monsoon (whether it's early or late) provides the household with additional information that it uses to update its forecasts of its income during the monsoon and later seasons. Table (3) lists all the variables that I include in the base information set, and in each of the three subsequent seasonal information sets.

Of course, for this interpretation (of η_{hti}^j as news) to be valid, the rainfall-related variables, R_{ti} , must truly represent “shocks” in the sense that they are unanticipated. That is, if $\Omega_{ht,i-1}$ denotes the household's information set in season $i - 1$, it must be that $E[R_{ti} \mid \Omega_{ht,i-1}] = 0$, where $E[\cdot \mid \Omega_{ht,i-1}]$ denotes the expectations operator conditional on $\Omega_{ht,i-1}$. I confirm that that is indeed the case by testing whether the rainfall-related variables are predictable from past weather patterns. The tests indicate the absence of both inter-year and inter-seasonal (i.e., intra-year) correlations.

The second point to note about (2.1) is that it is a semi-reduced form representation of the income process. The incomes of agrarian households, such as those in my sample, are derived primarily from agricultural cultivation. Therefore, underlying the income process is an agricultural production function that links the various input use decisions that the

household makes to the harvests that are realized. Insofar as I do not directly control for these input choices, (2.1) is a *reduced* form representation of the income process. Note, however, that this omission is deliberate because I wish to estimate the *total* effect of weather-related shocks on income and that includes any effects that result from alterations in input use in response to weather-related shocks. The specification is a *semi-reduced* form one in that I do condition on K_{ht} .⁷ The inclusion of K_{ht} serves two purposes. First, by interacting K_{ht} with the rainfall-related variables R_{ti} I am able to construct *household-specific* proxies for the exogenous shocks. And second, by including K_{ht} , I control for the possibility that, because of factor and asset market imperfections, a household's production decisions may be influenced by its ownership of productive assets.⁸

The third and final point about the specification is that it assumes that the weather-related shocks to income accumulate *additively* over the course of the crop-year. The significance of this somewhat restrictive assumption (and some evidence in favor of it) is discussed in the next section.

To test whether rainfall-related weather shocks that are realized in a particular season contain any news about (i.e., have any power in predicting) incomes in that and later seasons, I separately estimate (2.1) for each village *and season*. I deal with the fact that e_{hti} is not directly observable by estimating the three seasonal income processes implied by (2.1) sequentially. I begin by estimating the monsoon season income process, which poses no problems since e_{htM} enters as a standard disturbance term. I then include the estimated residual $\hat{e}_{htM}$ as a proxy for e_{htM} in estimating the winter season income process. Finally I estimate the summer season income process using $\hat{e}_{htM}$ and $\hat{e}_{htW}$ as proxies for e_{htM} and e_{htW} .

Table (4) reports the results I obtained. The first two columns of each panel show the results for the monsoon-season income process, the next two, for the winter-season income process, and the final two columns, for the summer season income process. Shown are the p -values from the F -tests for the joint significance of the variables— R_{ti} , $(K_{ht} \otimes R_{ti})$, and where applicable, $\hat{e}_{hti}$ —that serve as proxies for news about current and future income. The first, third, and fifth columns indicate the significance of the entire set of proxies. The remaining columns reveal whether the household-specific proxies are significant. The two sets of results are similar in every case but one.

Except in Aurepalle, monsoon-season rainfall surprises do not capture well the variations in monsoon season income—see the first row and first two columns of each panel. This primarily reflects the fact that in Shirapur and Kanzara such cashflows are rare. Monsoon-season rainfall surprises do, however, predict summer cash inflows in all three villages—fifth and sixth columns of the first row of each panel—and, in Aurepalle and Shirapur, provide good proxies for shocks to winter cashflows as well—third and fourth columns of the first row of each panel. Winter rainfall surprises are everywhere significant—though only marginally so in Kanzara—in explaining variations in winter cash inflows (third and fourth columns of the second row of each panel); and except in Kanzara, are also significant in predicting

⁷In other words, (2.1), is a specification of the conditional or restricted-profit function.

⁸Note that K_{ht} denotes the productive assets *owned* by the household *entering* crop-year t rather than the assets actually allocated to production in crop-year t . The distinction is important both in theory and in the data. For instance, the amount of land *owned* by a household entering a crop-year does not reflect any information received during the crop-year whereas the amount of land *cultivated* by the household may.

summer income. Fluctuations in rainfall, of which there is little in all three villages during the summer months, do nevertheless seem to explain variations in summer cashflows in Aurepalle and Shirapur.

Overall, the results from Table (4) suggest that in the three villages in my sample, rainfall-related weather shocks do contain news about current and future incomes, and hence, provide good proxies from which measures of the “news” that households receive, and the “uncertainty” they face, might be constructed. The idea of using rainfall data to construct proxies for shocks to income in agrarian settings is certainly not a new one. Paxson(1992) uses rainfall data to construct estimates of transitory shocks to income, while Jacoby and Skoufias(1997, 1998) use rainfall measures as proxies for unanticipated changes in income.⁹ Binswanger and Rosenzweig(1993) use measures of weather risk constructed from rainfall data to examine the influence of weather risk and wealth on the portfolio (productive and other assets) choices of agrarian households. The use I make of rainfall-related variables builds upon these previous papers. First, unlike Paxson(1992) I have the data to allow for household-specific weights on the aggregate shock (through the interaction terms $(K_{ht} \otimes R_{tk})$). Secondly, unlike Jacoby and Skoufias(1997, 1998), I use the rainfall data to construct proxies not only for shocks to *current* income but also for *news about future income*. And finally, whereas Rosenzweig and Binswanger(1993) looks at how weather risk influences *ex-ante* investment decisions, I focus on how, conditional on the portfolio choices they make, households adjust their consumption decisions *ex-post*, in response to weather risk and weather-related shocks.

3. Implications for consumption behavior

In this section I detail the implications of information accumulation over the course of the crop-year for the consumption behavior of households. To most easily highlight these implications I begin by going through a simple example in which I assume that income is exogenous, and for which a closed-form solution is available. I then highlight some potential complications that arise because, in practice, the incomes (more correctly, profits) of agrarian households depend on production decisions that they, themselves, make over the course of the crop year, where these decisions are presumably influenced by any new information that becomes available.

3.1. A simple example

Assume, as in (2.2) above, that the income that household h receives in season j of crop-year t is given by:

$$Y_{htj} = X_{ht}\alpha^j + K_{ht}\beta^j + \sum_{i=1}^j \eta_{hti}^j$$

where X_{ht} and K_{ht} are vectors of household characteristics, α^j and β^j are season-specific weighting factors, and η_{hti}^j is defined as:

$$\eta_{hti}^j \equiv E[Y_{htj} \mid \Omega_{hti}] - E[Y_{htj} \mid \Omega_{ht,i-1}]$$

⁹Wolpin(1982) was among the first papers to use rainfall data in studying consumption behavior, though he uses the data in a slightly different way. He exploits spatial variation in the moments of the rainfall distribution (estimated from historical data) to instrument permanent income.

where $E[.]$ denotes the expectations operator and $\Omega_{ht,i-1}$ denotes the information set of household h in season $i - 1$ of year t . η_{hti}^j therefore represents the revision, in season $i \leq j$, of the household's expectations about income in season j ; or in other words, η_{hti}^j is the *news* about season j income that the household receives in season $i \leq j$. Note that implicit in this specification of the income process is the assumption that once we control for the benchmark characteristics of the household, information about seasonal incomes in year t accumulates only within the year.

From the definition of η_{hti}^j it follows that:

$$\begin{aligned} E[\eta_{hti}^j \mid \Omega_{ht,i-1}] &= 0 & \forall h, \forall i \leq j \\ E[\eta_{hti}^k \eta_{htj}^k \mid \Omega_{ht,i-1}] &= 0 & \forall h, \forall i < j \leq k \end{aligned}$$

However, it is possible, and in fact quite likely, that the news, η_{hti}^j , that the household receives in season i about season j income may be correlated with the news, η_{hti}^k , it receives in the same season i about income in season k where $i < j \leq k$. Thus, I assume that:

$$E[\eta_{hti}^j \eta_{hti}^k \mid \Omega_{ht,i-1}] = \sigma_{hi}^{jk} \quad \forall h, \forall i < j \leq k$$

where σ_{hi}^{jk} is potentially non-zero and may be either negative or positive (except when $j = k$, in which case it has to be positive).

Beginning in season k of crop-year s , infinitely-lived households maximize the expected value of the intertemporally separable utility function:

$$E_{sk} \left(\sum_{j=k}^3 \beta^{j-k} v(C_{hsj}) + \sum_{t=s}^{\infty} \sum_{j=1}^3 \beta^{3(t-s)-k+j} v(C_{htj}) \right) \quad (3.1)$$

subject to the sequence of seasonal budget constraints:

$$C_{htj} + A_{htj} = Y_{htj} + (1 + r)A_{ht,j-1} \quad (3.2)$$

Each crop-year t has three seasons, indexed by j , and $v(.)$ denotes the within-season utility function. C_{htj} is household h 's consumption level in season j of crop-year t (denoted (t, j)), and $\beta < 1$ is its subjective discount factor. A_{htj} is the household's asset holdings at the end of (t, j) , and r is the season-to-season rate of return on assets. The Euler equation for this problem is the usual one, i.e.,:

$$v'(C_{ht,j-1}) = \beta(1 + r)E[v'(C_{htj}) \mid \Omega_{ht,j-1}] \quad (3.3)$$

To proceed further in characterizing the consumption process, additional assumptions need to be made. I assume that the household's preferences are described by the constant absolute risk aversion (CARA) utility function:

$$v(C_{htj}) = -\frac{1}{\theta} e^{-\theta C_{htj}}$$

and that the various components of news that are received in season i are jointly normally distributed:

$$\eta_{hti} \sim N(0, \Sigma_{hi})$$

In addition, to save on notation, I make the otherwise inessential assumption that:

$$\beta = \frac{1}{1+r}$$

With these assumptions, the results of Caballero(1990) can be adapted to this seasonal context to generate the following closed-form solution for consumption changes:

$$\Delta C_{htj} \equiv (C_{htj} - C_{ht,j-1}) = \frac{r}{1+r} \sum_{k=j}^J (1+r)^{j-k} \eta_{htj}^k + \frac{1}{2} \theta V[C_{htj} | \Omega_{ht,j-1}] \quad (3.4)$$

Here, $V[C_{htj} | \Omega_{ht,j-1}]$ is the conditional variance of one-period-ahead consumption, which, under our assumptions about the income process, is given by:

$$V[C_{htj} | \Omega_{ht,j-1}] = \left(\frac{r}{1+r} \right)^2 \left(\sum_{k=j}^J \sum_{l=j}^J (1+r)^{2j-k-l} \sigma_{hj}^{kl} \right) \quad (3.5)$$

Two key implications of (3.4) are the focus of this paper. The first is that households should be forward-looking in the sense that household consumption should *respond immediately to news about future income*. This can be seen by rewriting the innovation in period (t, j) consumption as:

$$\begin{aligned} \xi_{htj} &\equiv C_{htj} - E[C_{htj} | \Omega_{ht,j-1}] \\ &= \frac{r}{1+r} \eta_{htj}^j + \frac{r}{1+r} \sum_{k=j+1}^J (1+r)^{j-k} \eta_{htj}^k \end{aligned}$$

The first term on the right hand side represents the influence of the shock to current income; the second term captures the adjustment called for because of the news received in (t, j) about incomes in later seasons.

The second implication is that households should display a precautionary motive. To see this, take the conditional expectation of (3.4), which yields:

$$\begin{aligned} E[\Delta C_{htj} | \Omega_{ht,j-1}] &= \frac{1}{2} \theta V[C_{htj} | \Omega_{ht,j-1}] \\ &= \frac{1}{2} \theta V\left[\frac{r}{1+r} \sum_{k=j}^J (1+r)^{j-k} \eta_{htj}^k | \Omega_{ht,j-1}\right] \end{aligned} \quad (3.6)$$

In other words, households should depress their consumption (i.e., engage in precautionary saving) in $(t, j-1)$ because they are uncertain about the news, η_{htj}^k , they will receive in (t, j) about incomes in season j and in later seasons k . What this implies is that, *on average*, consumption should grow more going into periods in which a lot of uncertainty is resolved (i.e., information is received), as measured by the conditional variance of the news that is received.

Precautionary savings motives are usually stated in terms of households engaging in rainy day saving in the face of uncertainty about future incomes. As (3.6) makes clear, a more correct statement would be that the extent of precautionary saving in $(t, j-1)$ is driven by the uncertainty about the news that will be received in (t, j) , captured here by the *variance of the news*. In general, this will differ from the unconditional variance (or the

variance conditional on the base information set) of the income that is received in (t, j) , and this distinction becomes especially important when constructing empirical proxies for the uncertainty that households face. For instance, in this example, the variance of (t, j) income is:

$$V(Y_{htj}) = \sum_{i=1}^j \sigma_{hi}^{jj}$$

which, clearly, is not the same as the expression in (3.5). Finally, note that it is, therefore, possible in this setting for the seasonal pattern of consumption to differ quite markedly from the seasonal pattern of income since the seasonal pattern of consumption is driven by the *seasonal patterns in the resolution of uncertainty*.

3.2. Potential complications

The example I outline above assumes a simple exogenous income process in which the shocks to income accumulate additively over the course of the crop-year and the distribution of these shocks is fixed over time. The assumption of exogeneity may be questioned on the grounds that the incomes of agrarian households clearly depend on input choices that the households themselves make over the course of the crop-year. That is to say, a more general and realistic representation of the income process would be one where the income of a household h in season j of crop-year t is given by:

$$Y_{htj} = F^j(X_{ht}, K_{ht}, Z_{ht,i \leq j}^j, \varepsilon_{ht,i \leq j}^j) - \sum_{k=j}^3 p_j^z Z_{htj}^k$$

Here $F^j(\cdot)$ represents the crop revenues realized in (t, j) . Crop revenues are a function of:

$$Z_{ht,i \leq j}^j \equiv (Z_{ht1}^j, \dots, Z_{htj}^j)$$

where Z_{hti}^j denotes the vector of inputs applied by the household in season i (where $i \leq j$) on plots that are harvested in season j . Crop revenues also depend on:

$$\varepsilon_{ht,i \leq j}^j \equiv (\varepsilon_{ht1}^j, \dots, \varepsilon_{htj}^j)$$

where ε_{hti}^j denotes an exogenous *shock*—i.e., a random variable whose expectation, conditional on lagged information, is zero—that directly affects the harvest in season j but is realized in season $i \leq j$. Finally, $\sum_{k=j}^3 p_j^z Z_{htj}^k$ denotes cultivation expenditures in (t, j) ; these include both expenditures on inputs applied towards the harvest in (t, j) and those on inputs for crops to be harvested at a later date.

In each season, based on the information available to it at that point, and given its earlier decisions, each household decides on its consumption level, C_{htj} and the levels of various inputs, $Z_{htj}^{k \geq j} \equiv (Z_{htj}^j, \dots, Z_{htj}^3)$.

The question then is how this more realistic representation of the income process alters the implications for consumption behavior that were outlined in the simple example. To get a sense of the complications that can arise, consider a household's decision, in period (t, j) , regarding the level of input use on plots that are to be harvested in $(t, j + 1)$. If, as seems quite plausible, the marginal product of the input depends on ε_{htj}^{j+1} , the shock that is realized in (t, j) , the level of input use will clearly depend on ε_{htj}^{j+1} . And, by influencing the

level of input use, the realization of ε_{htj}^{j+1} will, therefore, provide information on the *level* of the harvest that can be expected in $(t, j + 1)$. This type of information accumulation is captured by the simple process assumed thus far. But if the marginal product of an input applied in (t, j) also depends on the shock that is realized in $(t, j + 1)$ —whether it does or not depends on the separability properties of $F^j(\cdot)$ —the level of input use in (t, j) will, in addition, influence the *variance* of the harvest in $(t, j + 1)$. Thus, in this setting, the realization of a weather-related shock, ε_{htj}^k , in period (t, j) may provide news not only about the level of income but also news about the variance of income in future seasons; in other words, the income process may be *conditionally heteroskedastic*.

Conceptually, this added complication does not pose any problems, because, as Caballero (1990) shows, even with a conditionally heteroskedastic income process, the basic conclusions about how households should react to either news or the presence of uncertainty remain. But the empirical specification of the seasonal income process I have assumed so far, (2.1), only allows weather-related shocks to accumulate additively, which rules out the possibility of conditional heteroskedasticity. To correct this, I extended the specification to include interactions of the rainfall-related shocks that are realized in different seasons, or in other words, to allow the shocks to accumulate multiplicatively. However, possibly because the inclusion of so many additional terms reduces the already limited degrees of freedom, none of the interaction terms were significant. Thus, the empirical strategy I outline in the next section, and the results I report in the following two sections, are based on estimates of the simple additive income process.

4. Empirical implementation

While the assumption of CARA preferences proved useful in highlighting the implications of information accumulation for consumption behavior, in empirically testing these implications, I adopt the constant relative risk aversion (CRRA) specification for the utility function. I do so because much of the more recent microeconomic consumption literature is based on this specification, which implies:

$$v'(C_{htj}) = C_{htj}^{-\rho} \quad (4.1)$$

Let r denote the net seasonal return on a risk-free asset to which all households are assumed to have access. With this assumption and (4.1), a second-order Taylor expansion of the Euler equation (3.3) yields:

$$\Delta \ln C_{htj} = \alpha + \frac{1}{\rho} \ln \beta [1 + r_{j-1}] + \frac{1}{2} \rho V(\ln C_{htj} \mid \Omega_{ht,j-1}) + \xi_{htj} \quad (4.2)$$

where:

$$\Delta \ln C_{htj} = \ln C_{htj} - \ln C_{ht,j-1}$$

α is a constant containing higher-order terms, $V(\ln C_{htj} \mid \Omega_{ht,j-1})$ is the conditional variance of the log of one-period-ahead consumption, and ξ_{htj} is the adjustment to consumption that the household makes in period (t, j) . ξ_{htj} is defined as:

$$\xi_{htj} \equiv \ln C_{htj} - E(\ln C_{htj} \mid \Omega_{ht,j-1}) \quad (4.3)$$

and therefore satisfies:

$$E(\xi_{htj} \mid \Omega_{ht,,j-1}) = 0 \quad (4.4)$$

The basic challenge in empirically implementing tests of forward looking behavior and precautionary saving based on equation (4.2) is that $V(\xi_{htj} \mid \Omega_{ht,,j-1})$ and ξ_{htj} are not directly observable. Appropriate proxies for the two variables have, therefore, to be constructed.

With the CRRA specification there is no closed form solution available for the consumption process. I, therefore, proceed based on the following approximation, which is motivated by the CARA example:

$$\begin{aligned} \xi_{htj} &\approx [Y_{htj} - E(Y_{htj} \mid \Omega_{ht,,j-1})] \\ &\quad + \sum_{k=j+1}^J [E(Y_{htk} \mid \Omega_{htj}) - E(Y_{htk} \mid \Omega_{ht,,j-1})] \\ &\approx \eta_{htj}^j + \sum_{k=j+1}^J \eta_{htj}^k \end{aligned} \quad (4.5)$$

In other words, I assume that the consumption innovation can be approximated by the unweighted sum of η_{htj}^j , the shock to season j income that is realized in season j and the news, η_{htj}^k , received in season j about incomes in seasons $k > j$. The season-specific estimates of (2.1) then provide natural empirical proxies for the news that households receive. For instance, as an empirical proxy for η_{htj}^k , where $k > j$, I use:

$$\begin{aligned} \hat{\eta}_{htj}^k &\equiv \hat{E}(Y_{htk} \mid \Omega_{htj}) - \hat{E}(Y_{htk} \mid \Omega_{ht,,j-1}) \\ &= R_{tj} \hat{\gamma}_j^k + (K_{ht} \otimes R_{tj}) \hat{\delta}_j^k + \hat{\phi}_j^k \hat{e}_{htj} \end{aligned} \quad (4.6)$$

The other unobservable variable in equation (4.2) is $V(\xi_{htj} \mid \Omega_{ht,,j-1})$. Since I proxy ξ_{htj} with:

$$\hat{\xi}_{htj} \equiv \sum_{k=j}^3 \hat{\eta}_{htj}^k \quad (4.7)$$

the obvious proxy for $V(\xi_{htj} \mid \Omega_{ht,,j-1})$ is the sample *conditional* variance of (4.7), which I denote from now on as $\hat{V}(\hat{\xi}_{htj} \mid \Omega_{ht,,j-1})$:

$$\begin{aligned} \hat{V}(\hat{\xi}_{htj} \mid \Omega_{ht,,j-1}) &= \hat{V}(\sum_{k=j}^J \hat{\eta}_{htj}^k \mid \Omega_{ht,,j-1}) \\ &= \sum_k \hat{V}(\hat{\eta}_{htj}^k \mid \Omega_{ht,,j-1}) + \sum_k \sum_{l>k} \widehat{C\hat{V}}(\hat{\eta}_{htj}^k, \hat{\eta}_{htj}^l \mid \Omega_{ht,,j-1}) \end{aligned}$$

In constructing the sample variances and covariances that make up the variance of (4.7), I explicitly recognize the information structure of the model. Thus, for instance, the sample conditional variance of $\hat{\eta}_{htj}^k$ is constructed as:

$$\hat{V}(\hat{\eta}_{htj}^k \mid \Omega_{ht,,j-1}) \equiv [\hat{\gamma}_j^k + K_{ht} \hat{\delta}_j^k]^2 \hat{V}(R_j) + [\hat{\phi}_j^k]^2 \frac{1}{T-1} \sum_s \hat{e}_{hsj}^2 \quad (4.8)$$

where:

$$\hat{V}(R_j) \equiv \frac{1}{14} \sum_{s=1975}^{s=1990} R_{sj}^2$$

and, for purposes of notational ease, I have ignored the fact that R_{tj} and K_{ht} are both vectors.

It is worth clarifying the way in which $\hat{V}(\hat{\xi}_{htj} \mid \Omega_{ht,j-1})$ differs from the empirical measures of income uncertainty used in past studies using micro-data. A typical proxy for income uncertainty is the household-specific sample variance of the residuals from a longitudinal regression of income on a set of covariates that capture any cross-sectional variation. The problem with such a measure in this context is twofold. First, it is an unconditional variance from a time-series perspective. If the covariates that capture cross-sectional variation, K_{ht} in this setting, vary over time, this variation would be incorrectly included in the measure. Second, and this is the more serious problem, to the extent that households accumulate information over the course of the crop-cycle, the measure would overstate the uncertainty faced by the household in the preceding period. By the same token, if households receive news about future cashflows, the variance of *realized* net cashflows underestimates the uncertainty that is resolved in any given period. What enables me to avoid these problems, as is clear from the construction of $\hat{V}(\hat{\xi}_{htj} \mid \Omega_{ht,j-1})$, is the fact that the rainfall data provide proxies for the unanticipated intertemporally varying income shocks.

5. Does household consumption respond to news about future income?

Forward-looking behavior on the part of households is implicit in all modern theories of consumption. A basic implication of forward-looking consumption behavior is that if a household receives news about its *future* income—for instance, that it will be higher than the household had earlier expected—the household should adjust its *current* consumption accordingly in an effort to smooth consumption over time. This implication has been tested in a couple of different ways.

The most widely used approach has been to test the exclusion restriction, $\gamma = 0$, implied by forward looking behavior in an equation such as:

$$\Delta \ln C_{htj} = \beta_0 + \gamma Y_{ht,j-1} + u_{htj}$$

where $Y_{ht,j-1}$ is lagged-income (or any variable assumed to be in the household's information set at time $(t, j - 1)$), or an equation in which Y_{htj} is decomposed into an anticipated, $\hat{E}[Y_{htj} \mid \Omega_{ht,j-1}]$, and an unanticipated, $\hat{\eta}_{htj}^j$, component:

$$\Delta \ln C_{htj} = \beta_0 + \beta_1 \hat{\eta}_{htj}^j + \gamma \hat{E}[Y_{htj} \mid \Omega_{ht,j-1}] + u_{htj}$$

These *excess-sensitivity* tests, which provide an indirect test—indirect in that they look for evidence *inconsistent* with forward-looking behavior—have been largely motivated by the fact that whereas empirical proxies for the news that households receive about future income are difficult to come by, variables, such as lagged income, upon which the excess-sensitivity tests may be based, are more readily observable.

A different approach, proposed initially by Campbell (1987), has been implemented using household-level data by Deaton (1992b), Udry (1995), Alessie and Lusardi (1997) among others. These papers provide direct evidence of forward-looking behavior by testing whether households save in anticipation of negative shocks to future income. Related are the *excess smoothness* tests implemented in papers such as Campbell and Deaton (1989), which

posit a particular income process and test whether consumption responds to innovations in income to the extent implied by the theory.

The fact that I am able to construct empirical proxies for the news that households receive about future income permits me to adopt an approach closer in spirit to the second one described above.¹⁰ The basic equation I estimate is:

$$\Delta \ln C_{htj} = \beta_0 + \beta_1 Y_{htj} + \beta_2 \left[\sum_{k=j+1}^3 \hat{\eta}_{htj}^k \right] + u_{htj} \quad (5.1)$$

which is an estimable version of equation (4.2), where I have, for the present, omitted the conditional variance term, and have, following the approximation in (4.5), decomposed ξ_{htj} , the consumption innovation, into two terms. The second term:

$$\sum_{k=j+1}^3 \hat{\eta}_{htj}^k \quad (5.2)$$

corresponds exactly to the second term in (4.5), and is a proxy for the sum of the news that the household receives in (t, j) about net cash inflows in future seasons.¹¹ The coefficient, β_2 , on this term therefore provides an estimate of the responsiveness of current consumption to news about *future* cash inflows. If households are forward-looking and decrease(increase) their consumption in (t, j) upon receiving news that their future cash inflows are likely to be lower(higher) than earlier expected, β_2 should be positive.

The first term, Y_{htj} , serves as a proxy for the first term in (4.5); β_1 therefore captures, albeit imprecisely, the response of consumption to news about current realized cash inflows. I include Y_{htj} rather than a proxy for just the unanticipated component, as (4.5) suggests, for purely empirical reasons. Rainfall-related surprises in season j affect both season j income as well as future incomes, and this is reflected in the construction of the proxies for news about current and future income. My main concern was that the inclusion of a noisy measure of the unanticipated shock to season j income might yield a spuriously significant estimate of β_2 —spurious in the sense that it might be really capturing the effect of current income on current consumption—because the measures of the news about future incomes are constructed *from the same underlying data on season j rainfall surprises*. The only way to avoid this possibility was to include current income, Y_{htj} , in the equation. Unless households are fully insured against income risk, or Y_{htj} is fully anticipated, β_1 should be positive.

Table (5a) presents the results from estimation of (5.1). The first column reports the estimates obtained from the full sample. The results provide weak evidence that households do engage in forward-looking consumption smoothing behavior. The estimate of β_2 is positive and, with a t -statistic of 1.703, is marginally significant (at the 10% level).¹² At the average consumption level during the monsoon season, Rs.1260, the point estimate of

¹⁰To follow the second approach *exactly*, I would need to observe savings on a *seasonal* basis. However, the ICRIAT data only record changes in asset-holdings from year to year, and thus only provide an *independent* measure of *annual* savings. A measure of seasonal savings can, of course, be imputed from the available seasonal consumption and income data. But doing so would, if incomes are measured with error, raise concerns about spurious correlations between income and savings.

¹¹The proxies are constructed according to (4.6).

¹²All the statistics reported in Table (6) are based on standard errors that have been corrected for the presence of generated regressors in the estimated equation.

0.0000091 implies a marginal propensity to consume out of expected future income of 0.0115 ($=0.0000091 \times 1260$). In other words, if rainfall patterns in the monsoon signal that cash inflows in the remaining two seasons are likely to be Rs.2752 higher than earlier expected, households increase their monsoon season consumption by Rs.31 ($=0.0115 \times 2752$).¹³ Without a closed form expression for consumption growth, I cannot say whether this response matches that predicted by theory. But a rough calculation based on the permanent income model, for which closed form expressions are available, suggests that the estimated response is low. Under a strict permanent income model, and assuming a stationary income process, the marginal propensity to consume out of innovations to one-period-ahead income is given by $\frac{r}{(1+r)^2}$ where r is the interest rate. The estimated marginal propensity to consume, when matched to this expression, yields an implied interest rate of just over 1%. This is much lower than the interest rates observed in the villages in my sample, which are sometimes as high as 30% from season to season.

The low estimate of β_2 might be explained by a number of factors. The first is the fact that it was obtained using an estimate of news about future cash inflows that was not discounted (see (5.2)). I therefore re-estimated (5.1), replacing (5.2) with a discounted sum calculated under different assumptions about the level as well as the extent of seasonal variation in the interest rate r . With interest rates in the range generally reported in the three villages, the estimates of β_2 increase very slightly relative to the base case reported in the first column of Table (5a), but still remain low. I do not therefore report these estimates.

A second possible explanation for the low estimate of β_2 is that the estimate of news I use is a poor proxy for the “theoretically correct” measure of news. Note that I am not referring here to the attenuation bias in small samples from the use of generated regressors—though that too may be a problem—rather, my concern is that the approximation, (4.5), that I used to motivate the empirical specification may be inappropriate. There is, however, little that I can do to address this concern.

I turn therefore to a third possible explanation, and that is the influence of binding borrowing constraints. Even if forward-looking households wish to smooth consumption in response to news of future income, with binding borrowing constraints, their capacity to do so may be limited. And this would dampen the response of consumption to news about future income.

I explore this possibility in two ways. First, I split the sample on the basis of observed wealth at the beginning of the panel and re-estimate equation (5.1) separately on each of the two sub-samples. The basic idea behind the approach is this. Survey and anecdotal evidence from the sample villages suggest that borrowing opportunities are largely shaped by wealth, with the wealthier cultivating households enjoying what effectively resembles an open line of credit from local moneylenders, and in some instances, greater access to loans from institutional sources. If the presence of binding borrowing constraints does partially explain the weak response to news about future cashflows observed in the sample as a whole, given the greater access to credit apparently enjoyed by wealthier households, the estimated response should be stronger among these households than it is among the poorer households.

¹³For the sample as a whole, Rs.2752 represents 32% of the average net cash inflows during the winter and summer seasons combined; 0.32 is the cross-sectional average of the intertemporal coefficient of variation of combined net cash inflows in these two seasons.

On the other hand, if the estimated response, though low, is spurious—for instance, because the estimate of news about future cashflows is constructed from variables that also influence cashflows in (t, j) , it may simply be picking up the response of consumption to concurrent realized cash inflows¹⁴—the estimates of β_2 should not differ significantly in the two sub-samples.

This ‘split-sample’ approach closely resembles that pursued by, among others, Zeldes(1989), and using these same data, Morduch(1990). But while those papers use this approach to investigate whether borrowing constraints help explain the excess-sensitivity of consumption growth to lagged income, I use it to clarify the interpretation of what is essentially an excess-smoothness finding. The distinction is important. To see why, consider the finding of Zeldes(1989) that, upon splitting his sample on the basis of wealth, the excess-sensitivity result disappears for wealthy households, but remains for poor households, for whom consumption growth remains significantly negatively correlated with lagged income. Using an argument similar to the one laid out above, Zeldes (1989) suggests that this finding is indicative of borrowing constraints. But as Carroll(1997) very persuasively argues, this finding is, in principle, also consistent with precautionary saving behavior in the absence of borrowing constraints.¹⁵ The main advantage of the test I propose is that a finding that wealthy households respond more to news about future income cannot be explained in terms of precautionary savings motives. In principle, therefore, the split-sample approach applied in this context provides a cleaner test of borrowing constraints.

The results reported in the second and third columns of Table (5a) provide support for the hypothesis of borrowing constraints. The estimate of β_2 for the sub-sample of wealthy households is positive and significant. At the mean monsoon consumption level for wealthy households, the point estimate, 0.0000145—which, while still low, is nearly 50% larger than the estimate for the sample as a whole—implies a marginal propensity to consume of 0.025. For the sample of poor households, on the other hand, the estimate of β_2 is actually negative, though statistically insignificant. Moreover, the consumption of poor households is much more sensitive to realized cashflows (which are at least partially anticipated) than is that of wealthy households. This too suggests the presence of borrowing constraints since it may reflect excess sensitivity to anticipated cash inflows. I also test whether the apparent differences in the estimated coefficients across the two sub-samples are statistically significant. The second panel of Table (5a) reports the relevant F -test statistics and associated p -values. The difference in the estimates of β_1 is statistically significant (p -value of 0.029). However, given the large standard error on the estimate of β_2 for the sub-sample of poor households, I cannot reject the hypothesis that the estimates are equal across the two sub-samples.

Splitting the sample on the basis of observed wealth offers a way of testing for forward looking behavior in the presence of borrowing constraints. It does not, however, exploit a

¹⁴Though, as noted above, the presence of Y_{htj} in the estimated equation makes this unlikely, it is nevertheless possible, especially if Y_{htj} is grossly mismeasured.

¹⁵The intuition behind this claim is this. With precautionary savings motives, consumption growth will be positively correlated with the variance of consumption. The variance of consumption will in turn be higher for poorer households and households with low incomes, who have fewer assets to buffer their consumption from income shocks. Hence, consumption growth will be negatively correlated with lagged income, and especially so for poorer households.

key aspect of optimal consumption behavior under borrowing constraints, which is that the response to news about future cashflows is potentially *asymmetric*. The intuition behind this is simple. Borrowing constraints, when they bind, do so in only one direction—they prevent households from consuming as much as they would like in a given period. Without restrictions on savings, and these seem unlikely in any setting, households are always free to adjust their consumption downward in response to bad news. Therefore, in the presence of borrowing constraints, the (downward) response of consumption to bad news should be much stronger than the (upward) adjustment in the face of good news. This suggests a testing strategy based on splitting the sample into one where households have received good news and another where they have received bad news.¹⁶ Without closed form solutions to the optimal consumption policy functions, the specific criterion used to split the sample is necessarily ad-hoc. I chose to split the sample based on the sign of the estimate of news about future income and re-estimated equation (5.1) separately for each of the sub-samples.

The first column of Table (5b) reports the results from this exercise. These results also provide evidence in favor of both forward-looking behavior and the hypothesis of borrowing constraints. The response of household consumption to good news (i.e., $\Sigma_{k>j} \hat{\eta}_{htj}^k > 0$) is statistically insignificant; on the other hand, the results indicate that when households receive bad news—i.e., revise downwards their expectations of future cash inflows—they also reduce their current consumption.¹⁷ This can be seen from the point estimate of β_2 for the sub-sample of observations where $\Sigma_{k>j} \hat{\eta}_{htj}^k < 0$; the estimate, 0.0000214, is, with a t -statistic of 2.578, statistically significant, and is more than double the estimate for the combined sample (0.0000091), and is larger also than the estimate for the sample of wealthy households. Moreover, the difference in the estimates of β_2 from the two sub-samples is statistically significant (see the bottom row of the first column).

The two approaches I adopt above in splitting the sample are not mutually exclusive and a natural next step is to combine the two. I therefore re-estimated equation (5.1) separately on the sub-samples split by wealth, but allowed both β_1 and β_2 to vary in each sub-sample based on whether the households received good news or bad news. The second and third columns of Table (5b) reports the results I obtained. Again, the results are consistent with both forward-looking behavior and borrowing constraints in that even poor households do seem to adjust their consumption downwards in the face of bad news, while the consumption of wealthy households is responsive to both good and bad news. At this level of disaggregation, though, the estimated responses are not significant.

¹⁶ Again, there is a precedent for this strategy. It is similar to that adopted by Altonji and Siow(1987), but that paper implements it to clarify an excess sensitivity finding while I use it to modify an excess-smoothness test. This difference is reflected in the different predictions about the direction of the asymmetry. In Altonji and Siow(1987), the presence of liquidity constraints implies that consumption should grow more in response to positive anticipated income changes than it falls in response to negative anticipated income changes.

¹⁷ These results could, in principle, also be consistent with precautionary savings motives without borrowing constraints. Carroll and Kimball(1996) show that for a large class of utility functions, the presence of a precautionary motive implies that consumption is a concave function of wealth. That in turn would imply that the increase in consumption from an increase in wealth should be smaller (in absolute value) than the decrease in consumption from an equivalent (in absolute value) decrease in wealth. The asymmetric response of consumption to good and bad news may reflect this concavity.

6. Do households engage in precautionary saving?

To test for precautionary saving I estimate:

$$\Delta \ln C_{htj} = \beta_0 + \beta_1 Y_{htj} + \beta_2 [\Sigma_{k>j} \hat{\eta}_{htj}^k] + \beta_3 \ln \hat{V}(\hat{\xi}_{htj} \mid \Omega_{ht,j-1}) + u_{htj} \quad (6.1)$$

which is simply equation (5.1) with an added variable—an estimate of the conditional variance of the news. If households engage in precautionary saving, which implies that consumption growth should be greatest going into periods where the most uncertainty is resolved, and for households who face the most uncertainty, β_3 should be positive and significant. A variant of this equation has provided the basis for most prior tests of precautionary saving. The precautionary savings parameter, β_3 , has, in most of these studies, been identified solely from cross-sectional variation in the measure of uncertainty, which itself has generally been constructed from time-series data on individual incomes. But this raises the possibility that these measures will be correlated with unobserved cross-sectional variation in attitudes towards risk—for instance, relatively more risk-averse individuals may choose professions that offer relatively more stable income streams. As Carroll(1997) demonstrates, this possibility reduces the usefulness of an equation such as (6.1) as a basis for tests of precautionary saving, except in situations where there is an observable and exogenous source of risk that varies across households and/or over time periods. The basic message of this paper has been that in a poor, agrarian economy, rainfall variation provides just such a source. And this is reflected in the fact that there are three sources of variation in $\ln \hat{V}(\hat{\xi}_{htj} \mid \Omega_{ht,j-1})$ —cross-sectional, year-to-year, and inter-seasonal—all of which are plausibly orthogonal to variations in risk preferences.

The first column of Table (6a) reports the results from estimation of equation (6.1) on the combined sample of households. The estimates of β_1 and β_2 are similar to those reported earlier. However, the estimate of β_3 , which should be positive and significant if households engage in precautionary savings, is neither.

As before, a possible explanation for this result may be the presence of borrowing constraints. But note that whereas the test of forward looking behavior based on (5.1) is complicated by the possibility of a binding borrowing constraint in (t, j) , the test of precautionary saving is complicated by the possibility of binding borrowing constraints in $(t, j - 1)$. This is because, as Deaton(1991) shows, in the presence of borrowing constraints, the marginal utility of consumption, and hence consumption itself, follows a threshold autoregression. A binding borrowing constraint in $(t, j - 1)$ delinks consumption in (t, j) from consumption in $(t, j - 1)$. Equation (4.2) is, then, no longer an appropriate representation of consumption growth.

Despite this difference, splitting the sample on the basis of observed wealth again offers a way of testing for both borrowing constraints and precautionary saving. I, therefore, re-estimated equation (6.1) separately on the sub-samples of wealthy and poor households. The results are reported in the second and third columns of Table (6a). For the sub-sample of poor households, the estimate of β_3 continues to be statistically insignificant, though it is now negative. Among wealthy households, however, the response of consumption to future uncertainty is consistent with precautionary saving behavior. The point estimate of β_3 is statistically significant, and at 0.0422, is quite large. The sample standard deviation of $\ln \hat{V}(\hat{\xi}_{htj} \mid \Omega_{ht,j-1})$ among wealthy households is 1.199 in the winter season. A one

standard deviation increase in this proxy for uncertainty about consumption during the winter season therefore leads to a 5.0% reduction in consumption during the monsoon ($1.199 \times 0.0422 = 0.050$). This represents a fairly large adjustment given that, for these households, average consumption growth between the monsoon and winter seasons is 15%.

Table (6b) report the results I obtained from estimation of the following equation in which I have added $\ln C_{ht,j-1}$ to (6.1):

$$\Delta \ln C_{htj} = \beta_0 + \beta_1 Y_{htj} + \beta_2 [\Sigma_{k>j} \hat{\eta}_{htj}^k] + \beta_3 \ln \hat{V}(\hat{\xi}_{htj} \mid \Omega_{ht,j-1}) + \beta_4 \ln C_{ht,j-1} + u_{htj} \quad (6.2)$$

The motivation for this specification, which is admittedly ad-hoc, comes from a consideration of the Euler equation (3.3), modified to incorporate a binding borrowing constraint in period $(t, j - 1)$:

$$v'(C_{ht,j-1}) = \beta[1 + r_{j-1}]E(v'(C_{htj}) \mid \Omega_{ht,j-1}) + \mu_{ht,j-1} \quad (6.3)$$

Here, $\mu_{ht,j-1} > 0$ is the Lagrange multiplier associated with the borrowing constraint in $(t, j - 1)$. Assuming a CRRA utility specification, a Taylor expansion of (6.3) analogous to that used to derive (4.2) yields:

$$\Delta \ln C_{htj} = \alpha + \frac{1}{\rho} \ln \beta[1 + r_{j-1}] + \frac{1}{2} \rho V(\ln C_{htj} \mid \Omega_{ht,j-1}) - \frac{1}{\rho} \phi(\ln C_{ht,j-1}) + \xi_{htj} \quad (6.4)$$

where:

$$\phi(\ln C_{ht,j-1}) \equiv \ln(1 - \frac{\mu_{ht,j-1}}{C_{ht,j-1}})$$

When the borrowing constraint is not binding, $\phi(\ln C_{ht,j-1}) = 0$, and (6.4) reduces to (4.2). If it is binding, however, (6.4), and not (4.2), is the correct representation of consumption growth and the test of precautionary saving should be based on the former equation. The problem with doing so is that $\phi(\ln C_{ht,j-1})$ is unobservable. When empirically implementing (6.4), I, therefore, substitute $\ln C_{ht,j-1}$ in its place— $\phi(\ln C_{ht,j-1})$ is weakly monotonically related to $\ln C_{ht,j-1}$ —and this leads to the specification, (6.2). Intuitively, with a binding borrowing constraint, the initial empirical specification, (6.1), which is based on (4.2), is overdifferentiated, and to correct for this, we need to introduce a lagged dependent variable into the equation.

The inclusion of $\ln C_{ht,j-1}$ as an added regressor has a marked effect on the estimate of β_3 for both the sample as a whole as well as for the sub-sample of poor households. For the sample as a whole, the estimate of β_3 goes from a statistically insignificant 0.0007, under the base specification, to 0.0512 with a t-statistic of 6.657 under (6.2) (see the first column of Table (6b)). The estimate for the sub-sample of poor households goes from -0.0147 to 0.0293. A one standard deviation (1.316) increase in uncertainty about winter season consumption therefore depresses consumption in the monsoon season by 3.9%; this, too, represents a fairly large adjustment given that, on average, consumption grows by 8.0% from the monsoon season to the winter season for these households. For the wealthy households, precautionary saving continues to be significant though the estimate of β_3 is slightly lower than that obtained under the base specification. The difference in the estimates of β_3 from the two sub-samples, which was statistically significant (with a p-value of 0.021) under the base specification, is no longer so.

The estimates of β_4 , the coefficient on lagged consumption, are also of some interest since, in theory, they can provide some indication of the importance of borrowing constraints. For each of the three samples, the coefficient on lagged consumption is negative and highly significant, which is consistent with the hypothesis of borrowing constraints. I do not, however, place too much weight on these estimates because they likely suffer from attenuation bias, given the possibility that consumptions may be measured with error.¹⁸

7. Conclusion

Households in poor, agrarian economies often derive their income primarily from rain-dependent agriculture. This dependence raises the possibility that, from rainfall patterns early on in the crop-cycle, households accumulate information about future cash inflows (e.g., harvests) before the cashflows are realized. Using detailed longitudinal data from three villages in India, this paper explored whether households utilize this information in the ways suggested by modern consumption theories.

Two hypotheses suggested by these theories are that households are forward-looking in their consumption behavior and react optimally to the receipt of information; and that in the face of uncertainty about future incomes, households engage in precautionary saving. The basic challenge in implementing tests of these hypotheses has been to find empirical measures of the “news” that households receive, and the uncertainty they face. I show that, in the villages I study, rainfall patterns provide good proxies from which such measures might be construct. I exploit this fact to test for precautionary saving and forward-looking behavior.

I find that precautionary saving and forward-looking behavior do not, by themselves, fully explain the observed consumption patterns. However, when I explicitly incorporate the possibility of binding borrowing constraints into the testing strategy, I obtain fairly strong support for the joint hypothesis of precautionary saving and forward-looking behavior in the presence of borrowing constraints.

¹⁸Ravallion and Chaudhuri(1997) provide evidence that suggests that the consumption data from these three villages are systematically mismeasured because of under-reporting of consumption of own-produced cereals and grains, especially by the larger cultivating households.

References

- [1] Alessie, R. and A. Lusardi (1997). "Saving and income smoothing: evidence from panel data," *European Economic Review*, 41(7), July, pp. 1251-1279.
- [2] Altonji, J.G. and A. Siow (1987). "Testing the response of consumption to income changes with (noisy) panel data," *Quarterly Journal of Economics*, 102(2), May, pp. 293-328.
- [3] Besley, T. (1995). "Savings, credit and insurance," in Behrman, J. and T.N. Srinivasan, eds., *Handbook of Development Economics, Volume 3A* (Amsterdam: Elsevier Science, North Holland), pp. 2123-2207.
- [4] Browning, M. and A. Lusardi (1996). "Household saving: micro theories and micro facts," *Journal of Economic Literature*, 34(4), December, pp. 1797-1855.
- [5] Caballero, R.J. (1990). "Consumption puzzles and precautionary savings," *Journal of Monetary Economics*, 25(1), January, pp. 113-136.
- [6] Campbell, J.Y. (1987). "Does saving anticipate declining labor income? An alternative test of the permanent income hypothesis," *Econometrica*, 55, pp. 1249-1273.
- [7] Campbell, J.Y. and A.S. Deaton (1989). "Why is consumption so smooth?," *Review of Economic Studies*, 56(3), July, pp. 357-373.
- [8] Carroll, C.D. (1997). "Buffer-stock saving and the life cycle/permanent income hypothesis," *Quarterly Journal of Economics*, 112(1), February, pp. 1-55.
- [9] Carroll, C.D. and M.S. Kimball (1996). "On the concavity of the consumption function," *Econometrica*, 64(4), July, pp. 981-992.
- [10] Case, A. (1995). "Symposium on consumption smoothing in developing countries," *Journal of Economic Perspectives*, 9(3), Summer, pp. 81-82.
- [11] Deaton, A.S. (1991). "Saving and liquidity constraints," *Econometrica*, 59(5), September, pp. 1221-1248.
- [12] _____ (1992a). *Understanding Consumption* (New York: Oxford University Press).
- [13] _____ (1992b). "Saving and income smoothing in Cote d'Ivoire," *Journal of African Economies*, 1(1), pp. 1-24.
- [14] Jacoby, H. and E. Skoufias (1997). "Risk, financial markets, and human capital in a developing country," *Review of Economic Studies*, 64(3), July, pp. 311-335.
- [15] _____ (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.
- [16] Kimball, M. (1990). "Precautionary saving in the small and in the large," *Econometrica*, 58(1), January, pp. 53-73.

- [17] Kochar, A. (1995), "Explaining household vulnerability to idiosyncratic income shocks," *American Economic Review, Papers and Proceedings*, vol.85, no.5.
- [18] Morduch, J. (1990). "Risk, production, and savings: theory and evidence from Indian households," *mimeo*, Harvard University.
- [19] Paxson, C. (1992). "Using weather variability to estimate the response of savings to transitory income in Thailand," *American Economic Review*, vol.82, no.1.
- [20] Ravallion, M. and S. Chaudhuri (1997). "Risk and insurance in village India: a comment", *Econometrica*, 65(1), January, pp. 171-184.
- [21] Rosenzweig, M.R. and H.P. Binswanger (1993). "Wealth, weather risk and the composition and profitability of agricultural investments," *Economic Journal*, 103(1), January, pp. 56-78.
- [22] Udry, C. (1995). "Risk and saving in northern Nigeria," *American Economic Review*, 85(5), December, pp. 1287-1300.
- [23] Walker, T. S., and J. G. Ryan (1990). *Village and Household Economies in India's Semi-Arid Tropics*, (Baltimore: Johns Hopkins University Press).
- [24] Wolpin, K.I. (1982). "A new test of the permanent income hypothesis: the impact of weather on the income and consumption of farm households in India," *International Economic Review*, 23(3), October, pp. 583-594.
- [25] Zeldes, S. (1989). "Consumption and liquidity constraints: an empirical investigation," *Journal of Political Economy*, 97(2), April, pp. 305-346.

Figure 1: Average monthly rainfall (1975-1990)

Figure 2: Seasonal patterns in income

Table (1)

Analysis of variance of seasonal income		
Effects	<i>F</i> -statistic	<i>p</i> -value
Household	10.57	.0000
Agricultural year	7.74	.0000
Season	187.93	.0000
Village×agricultural year	1.29	.2071
Agricultural year×season	0.52	.9223
Village×season	80.73	.0000
Village×season×agricultural year	1.14	.2780

Table (2)

Cross-sectional averages of the intertemporal coefficient of variation of income				
Village	Annual income	Monsoon	Winter	Summer
Aurepalle	0.39	0.80	0.49	0.49
Shirapur	0.37	1.27	0.47	0.47
Kanzara	0.33	0.45	0.69	0.70

Table (3)

Variables	
X_{ht} =	wealth, liquid assets, age of head, schooling of head, adults
K_{ht} =	owned-land area, proportion irrigated, bullocks
R_{tM} =	monsoon onset date, monsoon rainfall, dry spells
R_{tW} =	October rainfall, November–January rainfall
R_{tS} =	summer rainfall
Information sets	
Base =	$\{X_{ht}, K_{ht}\}$
Monsoon =	$\text{Base} \cup R_{tM} \cup (K_{ht} \otimes R_{tM})$
Winter =	$\text{Monsoon} \cup R_{tW} \cup (K_{ht} \otimes R_{tW})$
Summer =	$\text{Winter} \cup R_{tS} \cup (K_{ht} \otimes R_{tS})$

Table (4)

Estimated equation:

$$Y_{htj} = X_{ht}\alpha^j + K_{ht}\beta^j + \sum_{k=1}^j R_{tk}\gamma_k^j + \sum_{k=1}^j (K_{ht} \otimes R_{tk})\delta_k^j + \sum_{k=1}^{j-1} \phi_k^j \hat{e}_{htk} + e_{htj}$$

p-values from *F*-tests of:

$$(1) H_0 : \gamma_k^j = \delta_k^j = \phi_k^j = 0$$

$$(2) H_0 : \delta_k^j = \phi_k^j = 0$$

Information set	Income					
	Monsoon		Winter		Summer	
	(1)	(2)	(1)	(2)	(1)	(2)
Aurepalle						
Monsoon	0.0774	0.0446	0.0172	0.0807	0.0000	0.0000
Winter			0.0000	0.0006	0.0005	0.0015
Summer					0.0016	0.0006
Shirapur						
Monsoon	0.4533	0.4201	0.0000	0.0000	0.0000	0.0000
Winter			0.0004	0.0004	0.0022	0.0123
Summer					0.0368	0.1470
Kanzara						
Monsoon	0.5293	0.5200	0.1147	0.1310	0.0000	0.0000
Winter			0.0869	0.0897	0.3363	0.6579
Summer					0.2732	0.5478

Table (5a)

Dependent variable: $\Delta \ln C_{htj}$	Sample		
	Combined	Wealthy	Poor
Coefficient estimates (t -statistic):			
$Y_{htj} \equiv \text{current income} : \hat{\beta}_1$	0.0117 (3.656)	0.0101 (2.962)	0.0379 (3.107)
$\Sigma_{k>j} \hat{\eta}_{htj}^k \equiv \text{estimate of news about future income} : \hat{\beta}_2$	0.0091 (1.703)	0.0144 (2.179)	-0.0006 (-0.076)
F -test statistics (p -value):			
$H_0 : \hat{\beta}_1(\text{wealthy}) = \hat{\beta}_1(\text{poor})$		4.795 (0.0288)	
$H_0 : \hat{\beta}_2(\text{wealthy}) = \hat{\beta}_2(\text{poor})$		2.557 (0.110)	

Table (5b)

Dependent variable: $\Delta \ln C_{htj}$	Sample		
	Combined	Wealthy	Poor
Coefficient estimates (t -statistic):			
$Y_{htj} \times \mathcal{I}(\Sigma_{k>j} \hat{\eta}_{htj}^k > 0) : \hat{\beta}_1^+$	0.0133 (2.995)	0.0116 (2.222)	0.0407 (2.846)
$Y_{htj} \times \mathcal{I}(\Sigma_{k>j} \hat{\eta}_{htj}^k < 0) : \hat{\beta}_1^-$	0.0099 (2.248)	0.0100 (2.287)	0.0369 (2.929)
$\Sigma_{k>j} \hat{\eta}_{htj}^k \times \mathcal{I}(\Sigma_{k>j} \hat{\eta}_{htj}^k > 0) : \hat{\beta}_2^+$	-0.0044 (-0.508)	0.0132 (0.971)	-0.0190 (-1.890)
$\Sigma_{k>j} \hat{\eta}_{htj}^k \times \mathcal{I}(\Sigma_{k>j} \hat{\eta}_{htj}^k < 0) : \hat{\beta}_2^-$	0.0214 (2.569)	0.0167 (1.532)	0.0206 (1.212)
F -test statistics (p -value):			
$H_0 : \hat{\beta}_1^+ = \hat{\beta}_1^-$	0.349 (0.555)	0.065 (0.799)	0.067 (0.796)
$H_0 : \hat{\beta}_2^+ = \hat{\beta}_2^-$	4.220 (0.040)	0.031 (0.861)	3.805 (0.051)

Notes: $\mathcal{I}(\cdot)$ is the indicator function; coefficient estimates multiplied by 1000.

Table (6a)

Dependent variable: $\Delta \ln C_{htj}$	Sample		
	Combined	Wealthy	Poor
Coefficient estimates (t -statistic):			
$Y_{htj} \equiv \text{current income}^1 : \hat{\beta}_1$	0.0117 (3.493)	0.0083 (2.584)	0.0414 (3.209)
$\Sigma_{k>j} \hat{\eta}_{htj}^k \equiv \text{estimate of news}^1$ about future income : $\hat{\beta}_2$	0.0092 (1.705)	0.0143 (2.190)	0.0007 (0.089)
$\ln \hat{V}_{htj} \equiv \log(\text{estimated}$ variance of news in $(t, j) : \hat{\beta}_3$	0.0007 (0.079)	0.0422 (1.978)	-0.0147 (-1.296)
F -test statistics (p -value):			
$H_0 : \hat{\beta}_3(\text{wealthy}) = \hat{\beta}_3(\text{poor})$		5.325 (0.021)	

Table (6b)

Dependent variable: $\Delta \ln C_{htj}$	Sample		
	Combined	Wealthy	Poor
Coefficient estimates (t -statistic):			
$Y_{htj} \equiv \text{current income}^1 : \hat{\beta}_1$	0.0151 (5.471)	0.0116 (4.143)	0.0505 (4.469)
$\Sigma_{k>j} \hat{\eta}_{htj}^k \equiv \text{estimate of news}^1$ about future income : $\hat{\beta}_2$	0.0090 (1.875)	0.0112 (1.842)	0.0024 (0.334)
$\ln \hat{V}_{htj} \equiv \log(\text{estimated}$ variance of news in $(t, j) : \hat{\beta}_3$	0.0512 (6.657)	0.0392 (2.193)	0.0293 (2.862)
$\ln C_{ht,j-1} : \hat{\beta}_4$	-0.3799 (-16.165)	-0.4541 (-9.948)	-0.3670 (-12.919)
F -test statistics (p -value):			
$H_0 : \hat{\beta}_3(\text{wealthy}) = \hat{\beta}_3(\text{poor})$		0.242 (0.623)	

Notes:¹ coefficient estimates multiplied by 1000

1998-1999 Discussion Paper Series

Department of Economics
Columbia University
1022 International Affairs Bldg.
420 West 118th Street
New York, N.Y., 10027

The following papers are published in the 1998-99 Columbia University Discussion Paper series which runs from early November to October 31 of the following year (Academic Year).

The following is the Columbia University, Economics Department's website where viewers may access more information about the Discussion Paper series.

[Http://www.columbia.edu/cu/economics/](http://www.columbia.edu/cu/economics/)

All 1997-98, 1998-99 and future papers will be accessible on-line at the following website:

[Http://www.ssrn.com](http://www.ssrn.com)

Discussions papers that are not available on-line maybe obtained by written request to the Economics Department. Please check the Ordering Papers section for details.

Copy Requests

To order 1996-97 and prior discussion papers, please write to the Discussion Paper Coordinator at the address below, along with a check for the appropriate amount, made payable to: Department of Economics, Columbia University. Please be sure to indicate the discussion paper number or the particular series in your written request. *Orders cannot be processed without payment*, and they cannot be taken over the phone, or by fax or email.

Past discussion papers (1996-97 and prior papers) are available for purchase in *U.S. dollars only*, at the cost of:

US: per paper \$4.00 / per series: \$140.00
Canada: per paper \$5.00 / per series: \$150.00
Overseas: per paper \$7.00 / per series \$185.00

1998-99 Discussion Papers
Economics Department, Columbia University

11/9/99

Series no.	Title	Authors
9899-01	Propensity Score Matching Methods for Non-experimental Causal Studies	Dehejia, R.
9899-02	Institutional Solutions to the Principal-Agent Problem in African Health Care	Leonard, K. Leonard, D.
9899-03	The Economics of Repeated Extortion	Choi, J.P. Thum, M.
9899-04	To Pay or Not to Pay: Managerial Decision making and Wage Withholding in Russia	Desai, P. Idson, T.
9899-05	Wage Arrears, Poverty, and Family Survival Strategies in Russia	Desai, P. Idson, T.
9899-06	Averaged Periodogram Spectral Estimation with Long Memory Conditional Heteroscedasticity	Henry, M.
9899-07	Household Production, the Bundling of Services and Degradation, and Non-monotonic Environmental Engel Curves	Chaudhuri, S. Pfaff, A.
9899-08	A Carbon Sequestration Supply Function And Development of Feasible Clean Development Mechanism Rules for Tropical Forest Carbon Sinks	Kerr, S. Pfaff, A.
9899-09	Asymmetric Information and the Role of NGOs in African Health Care	Leonard, K. Leonard, D.
9899-10	Forward-looking Behavior, precautionary savings, and borrowing constraints in a poor, agrarian economy: test using rainfall data	Chaudhuri, S.