



INTERNATIONAL POLICY CENTER
Gerald R. Ford School of Public Policy
University of Michigan

IPC Working Paper Series Number 68

Why is Mobility in India so Low? Social Insurance, Inequality and Growth

Kaivan Munshi
Mark Rosenzweig

December, 2007

Why is Mobility in India so Low? Social Insurance, Inequality, and Growth *

Kaivan Munshi[†]

Mark Rosenzweig[‡]

December 2007

Abstract

This paper examines the hypothesis that the persistence of low spatial and marital mobility in rural India, despite increased growth rates and rising inequality in recent years, is due to the existence of sub-caste networks that provide mutual insurance to their members. Unique panel data providing information on caste loans, marriage, and migration are used to link caste networks to household and aggregate mobility. Wealthier households within their sub-castes are more likely to both migrate and inter-marry, suggesting that they are not being adequately compensated by the network in their role as net lenders. Conversely, among households with the same wealth, those in higher-wealth caste networks are more likely to obtain loans and are less likely to be mobile, providing direct evidence that the networks restrict mobility. At the aggregate level, the networks appear to have coped successfully with the rising inequality within sub-castes that accompanied the Green Revolution. The results suggest that caste networks will continue to smooth consumption in rural India for the foreseeable future, as they have for centuries, unless alternative credit mechanisms of comparable quality become available.

*We received helpful comments from Andrew Foster, Rachel Kranton, Ethan Ligon and seminar participants at Arizona, Chicago, Essex, Georgetown, Harvard, IDEI, ITAM, LEA-INRA, LSE, UCLA, and NBER. Alaka Holla provided excellent research assistance. Research support from NICHD and the National Science Foundation is gratefully acknowledged.

[†]Brown University and NBER

[‡]Yale University

1 Introduction

Increased mobility is the hallmark of a developing economy. Although individuals might be tied to the land they are born on and the occupations that they inherit from their parents in a traditional economy, the emergence of the market allows individuals to seek out jobs and locations that are best suited to their talents and abilities. Among developing countries, India stands out for its remarkably low levels of occupational and geographic mobility. Munshi and Rosenzweig (2006), for example, show how caste-based labor market networks have locked entire groups of individuals into narrow occupational categories for generations. India lags behind other countries with similar size and levels of economic development in terms of geographical mobility as well.¹ Figure 1 plots the percent of the adult population living in the city, and the change in this percentage over the 1975-2000 period, for four large developing countries: Indonesia, China, India, and Nigeria (UNDP 2002). Urbanization in all four countries was low to begin with in 1975 but India falls far behind the rest by 2000.

Data from the Indian census indicates that just one-fifth of the growth in the urban population from 1991 to 2001, which we have seen to be relatively low to begin with, can be attributed to migration. Despite the restructuring of the Indian economy during the 1990's, the proportion of individuals in the population that changed residence in the decade preceding the 1991 and 2001 census rounds was roughly constant, and among these migrants less than a third were men seeking jobs. Consistent with these national trends, a representative sample of rural Indian households in 1982 and 1999 that we use for much of the analysis in this paper indicates that in rural areas migration rates of men out of their origin villages are low and actually declined, from 10 percent in 1982 to 6 percent in 1999.² Indeed, it is standard practice for researchers to ignore migration in empirical studies based in rural India, although a coherent explanation for such immobility rooted in the fundamental features of the local economy is lacking.³

Low rates of migration are not the only indicators of immobility in India. The basic marriage

¹By geographic mobility we mean a permanent change in residence. There is data indicating that temporary migration, in which residence remains the same but workers move for the purpose of temporary employment, has increased in India.

²Women have traditionally migrated outside the village to marry in India. In our data, of the rural women marrying between 1982 and 1999, more than 88 percent had left their origin village by 1999, and marriage is almost always the reason for this exit. Along these lines, the 2001 census indicates that movement due to marriage accounts for roughly 45 percent of all migration in India, while employment, business, and the movement of entire families accounts for just 39 percent of migration (similar statistics are obtained in the 1991 round). We will consequently focus on male out-migration when measuring spatial mobility in this paper.

³The assumption that the rural population is essentially immobile has been made in studies of local governance in rural India (Banerjee et al., 2005), the determinants of rural schooling (Foster and Rosenzweig, 1995), and the effects of rural industrialization (Foster and Rosenzweig, 2005).

rule in Hindu society is that no individual is permitted to marry outside the sub-caste or *jati*. Social mobility will be severely restricted by this rule because individuals are forced to match within a very narrow pool. Social mobility, as measured by inter-caste marriage, continues to be low in rural India despite the economic changes within and across castes that have taken place over the past decades. Figure 2 reports rates of marriage outside the *jati* for the children and siblings of all household heads over the period 1950-1999, based on retrospective information collected in the 1999 round of the rural survey. We see that out-marriage in 1999 is less than 6 percent of all marriages, and this rate has remained stable for 50 years in rural areas.

Why is mobility in India so low? Many *ad hoc* explanations are available; for example, one explanation for the historically low rural-urban migration in India in the 1970's and 1980's is that opportunities in the rural areas expanded with the increase in agricultural productivity that accompanied the Green Revolution, and so the push out of the rural areas that drives migration in other economies may have been absent. However, over the past 15 years or more Indian growth rates, inclusive of the non-agricultural sector, have been high by any standard and male migration and out-marriage continue to be low, at least in rural areas. Similarly, it could be argued that individuals continue to marry within their *jatis* simply because they have a strong preference for partners with the same background and characteristics. However, this cannot explain why out-marriage has not increased despite the increase in within-*jati* inequality that we document below. Although alternative explanations may be available for the patterns that we have described, none of them can explain all of these phenomena and all are silent on which households do become mobile.

The particular (unified) explanation for both low migration and low out-marriage that we propose in this paper is that rural *jati*-based networks, which have been active in smoothing consumption for centuries in the absence of well functioning markets, restrict mobility. Once the individual has married outside the *jati* or migrated outside the village, he is less susceptible to the sanctions that are imposed on those who fail to honor their network obligations. This individual will consequently be excluded from the mutual insurance arrangement in equilibrium (see Greif 1993 for a similar argument in a different context). Individuals who out-marry or migrate thus lose the services of the caste networks, which dampens mobility when alternative risk-sharing arrangements of comparable quality are unavailable.

We use in this paper newly-available survey data describing the population of rural India over the past three decades that identifies the *jatis* of the immediate relatives of household heads and their

spouses and provides detailed information on loans and transfers to (i) examine the hypothesis that caste networks providing mutual insurance play an important role in limiting mobility and (ii) assess the prospects for both the decay of these networks and for increased mobility as economic growth proceeds. A direct test of the hypothesis that rural households forego mobility in return for superior insurance is that those who leave networks are less insured. However, any attempt to estimate the loss of insurance due to out-marriage or migration must take account of the fact that both insurance and mobility are endogenously determined. In our view there are no credible instruments for marriage or migration that would identify their effects on insurability. Our strategy instead is to identify those households who would be least affected by a loss in network services and those networks that are of relatively poor quality. We then proceed to show that it is precisely those households and the members of those vulnerable networks that have the greatest propensity to be mobile.

We first show, using data from the 1982 and 1999 rounds of the rural survey, that caste loans are more important than bank loans or moneylender loans in smoothing consumption and in particular for meeting contingencies such as illness and marriage that impose infrequent but very large costs on households. We also show that caste loans are also received on more favorable terms, with respect to both interest rates and collateral requirements, than alternative sources of finance. There is a large literature on credit markets in developing countries that has primarily focused on the interaction between traditional local moneylenders and formal banks. More recently, attention has shifted to micro-finance arrangements. This literature, however, has ignored informal network loans, which we will see are an important source of credit in rural India.

How well do these caste networks function? Based on Townsend's (1994) work on risk-sharing in rural India, many studies have implemented a test of full risk-sharing in which a key implication is that household consumption should be completely determined by aggregate consumption in the group around which the mutual insurance is organized. Previous contributions to this literature that are situated in rural India have treated the village as the social unit, whereas we argue in this paper that the *jati*, which extends beyond village boundaries, is the relevant unit around which the network is organized.⁴ We consequently implement a modified Townsend test, using a national panel sample of rural households over a three-year period, 1969-71, to assess if household consumption co-moves strongly with aggregate *jati* consumption, net of village consumption. We find this to be the case, and

⁴An exception is Morduch (2004) who considers sub-caste groupings within villages as mutual-insurance networks. Given the data used, however, he could not fully implement a model incorporating caste networks, which extend beyond villages.

also show that alternative measures of aggregate consumption, at the level of the district, the broad caste category, or occupation, are not correlated with household consumption.

Although a fair amount of consumption smoothing is sustained at the level of the village and the *jati*, our version of the Townsend test rejects full risk-sharing, matching the results from many previous studies (see, for example, Townsend 1994, Ligon 1998, and Fafchamps and Lund 2000). This leads us to consider a model of mutual insurance with limited commitment as a more accurate descriptor of the functioning of the caste network. In this model, a household that receives a positive income shock in a given time-period will transfer resources to one or more members of the network who received a negative shock in that period, in return for which (state-contingent) transfers will flow in the opposite direction for some periods in the future (Ligon, Thomas and Worrall 2002). This model suggests that risk-sharing will not be perfect, as we find, and that flows of resources across households will be in the form of quasi-loans rather than pure gifts or transfers. A comparison of caste loans versus pure transfers, by value, indicates that this is indeed the case in our data. There is a large empirical literature on mutual insurance in developing economies that has focussed almost exclusively on pure gifts or transfers as the mechanism through which consumption smoothing occurs (Attanasio and Albarran 2003 is a notable exception). As with the literature on credit, what is missing from the insurance literature is the role played by informal network loans in smoothing consumption. Apart from documenting this role, the analysis in the paper also generates predictions from the limited commitment model linking the caste networks that support these loans to mobility.

The key predictions of the model are that (i) higher quality networks, where quality is measured by average wealth in the *jati*, are associated with lower mobility, and (ii) wealthier households *within* their *jatis* are more likely to exit their networks and, therefore, to be more mobile. To derive these predictions we extend the standard limited commitment model of mutual insurance to allow for wealth inequality among the participants. The insurance arrangement with limited commitment is shown to be more difficult to sustain when there is inequality within the *jati*, for example, if some households receive positive shocks more often than others. These wealthier households end up being lenders more often than borrowers and unless the compensatory transfers that flow back to them increase in magnitude they will end up subsidizing the network. Thresholds on minimum consumption, enforced by redistributive social norms, could prevent such asymmetric transfers from being implemented, in which case growing wealth inequality within *jatis* could lead to the wealthiest households within their *jatis* being pushed past their participation constraints. In this framework, the wealthiest households

would unambiguously have the greatest propensity to out-marry and migrate since they are least affected by the loss in network services.⁵ By a symmetric argument, holding household wealth constant, an increase in average *jati* wealth makes the household a net borrower and should, therefore, lower its propensity to out-marry and migrate.

To test these predictions, we need a source of exogenous variation in wealth inequality within the *jati*. The Indian Green Revolution, which began in the late 1960's, was an important force increasing inequality within *jatis* that were historically quite homogeneous. Although the Green Revolution substantially increased agricultural productivity and farm incomes, all growers did not gain access to this superior technology simultaneously. Some regions were better suited to the early High Yielding Varieties (HYVs) of seeds than others, and although cross-breeding with local varieties ultimately allowed the new technology to be adopted throughout the country, those areas that had a head start ended up with a steeper trajectory than those that followed. This spatial variation in wealth in the aftermath of the Green Revolution increased inequality within *jatis*, which typically span a wide area. Indeed, comparison of Gini coefficients of the rural wealth distribution in 1982 and 1999 indicate that within-*jati* inequality rose by 13.7 percent. In contrast, within-village wealth inequality rose by only 6.8 percent over the same time-period.

We exploit the timing of HYV seed availability as the exogenous source of variation that determines changes in wealth within and between *jatis* in the empirical analysis. We find, consistent with the model, that the caste-loan position (loans-in minus loans-out) of a household is decreasing in its own wealth but increasing in overall *jati* wealth. Own wealth also affects loans received from banks and moneylenders, but aggregate *jati* wealth does not. The fact that the wealthiest individuals within the *jati* are now net lenders does not imply that they will exit the network, unless it is the case that the compensatory transfers (implicit interest rates) in the mutual insurance arrangement fail to adjust sufficiently. The data indicate that that relatively wealthy individuals do receive caste loans at lower rates and disburse loans at higher interest rates. The strong redistributive norms that have historically been in place in these communities make it unlikely, however, that these households will be compensated completely in their new role as net lenders. And, indeed, we find that conditional on average *jati* wealth, the propensity to be mobile is increasing in household wealth. Conversely, among households with the same wealth, those belonging to *jatis* with lower average wealth are significantly

⁵Previous research on individual participation in collective institutions (Banerjee and Newman 1998, La Ferrara 2002) suggests that the relationship between relative wealth and exit is ambiguous. For the mutual insurance arrangements that we examine in this paper, however, exit occurs unambiguously at the top of the distribution.

more likely to out-marry and migrate.

Apart from establishing a link between caste networks and household mobility, the analysis also connects network viability and wealth inequality to aggregate growth and mobility. The theoretical framework and the empirical results indicate that when caste networks are active, increases in aggregate wealth brought about by economic growth, with no accompanying increase in within-network inequality, have little effect on mobility. What matters for changes in mobility is not even (exogenous) changes in inequality in the general population, but rather inequality within the *jati*. Our estimates indicate that increasing inequality by shifting wealth from the bottom to the top of the wealth distribution would actually *lower* overall exit; households at the top of the distribution would be more likely to exit but households at the bottom of the distribution would be even more likely to stay. It then follows that the increase in within-*jati* inequality in the aftermath of the Green Revolution might actually have reduced mobility rates, that were low to begin with, even further. Low mobility has negative implications for growth, but the resilience of the caste networks in the face of substantial increases in inequality suggests that they will continue to smooth consumption in rural India in the foreseeable future, as they have for centuries, unless alternative credit mechanisms of comparable quality become available.

The paper is organized in six sections. The next section establishes that caste loans are an important and preferred source of finance for smoothing consumption and meeting contingencies. Section 3 implements the modified Townsend-test to provide evidence that *jati* networks play an important role in smoothing consumption in rural India. Section 4 extends the model of mutual insurance with limited commitment to identify the effect of an increase in wealth inequality within the *jati* on the loan position and the implicit interest rate faced by households at different positions in the wealth distribution. The identity of those households that might want to exit can then be established immediately. Section 5 verifies the implications from the preceding section and then extends the analysis to derive the relationship between networks, inequality, and aggregate mobility. Section 6 concludes.

2 Sources of Finance and the Costs of Contingencies in Rural India

In this section we show that loans from caste members are important and preferred mechanisms through which consumption is smoothed in rural India. We also show that the comparative advantage of the caste loans over alternative sources of finance has been maintained over time and document

the frequency and magnitudes of the health and marital consumption shocks that are experienced by rural households. The evidence is based on a panel survey of rural Indian households covering the period 1982 through 1999. The baseline survey is the 1982 Rural Economic Development Survey (REDS) carried out by the National Council of Applied Economic Research (NCAER) in 1981-82 in 259 villages located in 16 states (the major states except Assam). The sample of 4,979 households is meant to be representative of all rural households in those states. Subsequently, all households in the 1982 survey (with the exception of those residing in Jammu and Kashmir) in which at least one member remained in the village were resurveyed in 1999.

A key feature of both surveys is that information on the source and purpose is provided for every loan that was outstanding at the beginning of the reference period or obtained during the reference period. Although the 1982 and 1999 survey instruments were designed for the most part to permit analysis across the two time periods, some sections did not coincide precisely. For example, the classification of activities that loans are used for is much coarser in 1999 and, in particular, consumption expenses do not appear as a separate category. Because an important role of the caste networks, and the quasi-loans that they provide, is to smooth consumption, we restrict our description of loans by source and by purpose to the 1982 survey.

The 1982 survey data indicate that rural Indian households do have access to credit, despite the well documented credit market imperfections in that country. 17 percent of households in the 1982 survey round reported credit outstanding in the year prior to the survey and average outstanding debt (conditional on holding any debt) was as large as 54 percent of the household's annual income. The households received credit from three main sources: banks, moneylenders, and the caste network. As noted, caste networks have received little attention in both the credit literature and the consumption smoothing literature. Nevertheless, one-third of households with credit outstanding reported that they had received caste loans and these loans are important by value as well. The 1982 survey data indicate that of the 1,423 loans recorded for the survey households those from caste members made up 12.3 percent of all loans in value, approximately equal to the amount households obtained from moneylenders (12.2 percent). Bank loans were 46.3 percent of all loans. Table 1 reports the proportion of loan value both by source and purpose. As can be seen, caste and moneylender loans are also similar in that they are disproportionately used to cover consumption expenses and for meeting contingencies such as illness and marriage. For example, although loans from caste members were 12 percent of all loans in value, they were 23 and 43 percent, respectively, of the value of all consumption and

contingency loans.⁶ Similarly, loans from money lenders were 47 and 27 percent of all consumption and contingency loans. In contrast, 53 percent of loans for operating expenses were from banks, compared with six and two percent from caste members and moneylenders. And, banks supplied 26 percent of all investment loans, compared with 17 percent from caste members as well as moneylenders.

Loans for contingencies such as marriage and illness, which are disproportionately supplied by the caste network, tend to be large but infrequent. Although the 1982 survey round included details of loans and transfers that were received by the households in the sample, information on marriage and medical expenditures must be obtained from the 1999 round. 5.8 percent of the households in that round reported that they incurred expenditures on marriage ceremonies and celebrations in the year prior to the survey. Among those households, around half reported additional dowry expenditures, presumably because one or more girls were married out in that year.⁷ In addition, 9 percent of households in the 1999 round reported that a member had received treatment at a hospital or nursing home in the past year, with adult members receiving treatment in two-thirds of these households.

Thus, as many as 15 percent of households in a given year may incur unusual expenditures associated with marriage or major illnesses. How burdensome are these contingencies? On average, a single contingency loan accounts for over 59 percent of the receiving household's annual income, which is substantially larger than the corresponding statistic for investment loans (47 percent). These large loan sizes reflect the considerable expenditures associated with marriage and serious illness. For example, dowries averaged Rs. 6,767 for female members of the surveyed households who married in the 1981-83 period.⁸ Based on marriage expenditures in the year prior to the 1999 survey, we estimate that dowries are twice as large as expenditures on ceremonies and celebrations. If this ratio has remained constant over time, then total expenses for marriages around 1982 that involved a dowry would have been around Rs. 10 thousand. These figures imply that average marriage expenditures (including dowry) could have been as high as twice a household's annual income. This is a substantial sum of money for households in poor agrarian economies. At the same time, the average value of contingency

⁶Caldwell, Reddy and Caldwell (1986) surveyed nine villages in South India after a two-year drought and found that nearly half (46%) of the sampled households had taken consumption loans during the drought. The sources of these loans (by value) were government banks (18%), moneylenders, landlord, employer (28%), relatives and members of the same caste community (54%), emphasizing the importance of caste loans for smoothing consumption.

⁷The 1999 round of the survey collected retrospective dowry information for all marriages of the immediate relatives of the household heads and their spouses. Almost all marriages were associated with a dowry expenditure - a payment from the bride's family to the groom's family. See Anderson (2004) for a discussion of the role of dowry in India.

⁸This statistic is computed for the sisters and the sisters-in-law of the 1999 household heads (who also appeared in the 1982 round of the survey) who married in the 1981-83 period. Expenditures on ceremonies and celebrations are only available for marriages that took place in the year prior to the 1999 survey.

loans received from the caste network in the year prior to the 1982 survey was Rs. 4,660, which is 95 percent of the receiving households' annual income in that year. Thus, we see that caste loans contribute substantially to meeting the expenses associated with marital events. Detailed information on medical expenditures, by condition, was not collected in any of the survey rounds. Nevertheless, we expect these expenditures to be substantial as well, particularly for serious conditions where the household is most likely to seek expensive private treatment (Luke and Munshi 2007). The loss in household income when an adult member is afflicted with a serious long-term condition can also be substantial and must be taken into account when computing the economic impact of a health shock.

Although expenditures on marriage and health may be infrequent, failure to secure credit when these expenditures do arise can have catastrophic consequences. Gifts and transfers between close relatives could be used, in principle, to meet these expenses. These sources of finance are about as frequent as caste loans in the 1982 data, but they are less than half the size of loans in value, and in practice caste loans appear to be the primary mechanism through which rural Indian households cover major contingencies when they arise.

We argue in this paper that caste networks restrict mobility because comparable arrangements are unavailable, particularly for meeting contingencies. Without separating loans by purpose, Table 2 also shows that loan terms - the average interest rate, the proportion of zero-interest loans, and the proportion of loans requiring collateral - are more favorable for caste loans on average. The statistics, weighted by the value of the loans, are computed separately for the 1982 and the 1999 rounds, allowing us to examine any changes in the term structure of the loans over time. Statistics reported for 1982 are based on the 1,423 loans that were used to compute the statistics in Table 1. Statistics reported for 1999 are based on the 1,687 loans obtained by the sampled households in that year, or still outstanding in that year.

Table 2 shows that for both 1982 and 1999 caste loans have (statistically significant) lower interest rates than either bank or moneylender loans in both years. A substantial fraction of the caste loans are also zero-interest, consistent with the patterns reported elsewhere for informal quasi-loans (for example, Fafchamps and Lund 2000). Not only are bank interest rates higher than the average interest rates charged by caste members (15 versus 11 percent in 1982 and 10 versus 8 percent in 1999), but most caste loans also do not require collateral (84 percent in 1982 and 98 percent in 1999). In contrast, almost half of bank loans in 1982 and over 83 percent of bank loans in 1999 required some collateral. As is well known, moneylender loans often do not require collateral, but the average

interest rate charged by moneylenders is much higher than that charged by caste members - 17 versus 11 percent in 1982 and 31 versus 8 percent in 1999.

Tables 1 and 2 establish that loans from caste members are important for smoothing consumption, particularly for contingencies that make large expenditure demands, and continue to be advantageous to borrowers compared with loans from the two major alternative sources of finance in rural India. The analysis that follows will formally test the role of caste networks and the loans that they provide in smoothing consumption.

3 Caste Networks and Consumption Smoothing

In his pioneering study of risk and insurance in village India, Townsend (1994) derives a simple test to assess whether households are fully insured. Following Morduch (2004) and Bardhan and Udry (1999), the set of Pareto optimal consumption allocations with full risk-sharing can be obtained as the solution to the central planner's problem of maximizing a social welfare function

$$W = \sum_{t=0}^T \delta^t \sum_{s=1}^S \pi_s \sum_{i=1}^N \lambda_i U(C_{it}^s)$$

where $\delta \in [0, 1)$ is a discount factor, π_s is the probability of state s occurring, λ_i is individual i 's welfare weight, and C_{it}^s is his consumption allocation in period t and state s , subject to the constraint that total consumption in that period and state should not exceed total income, $\sum_i C_{it}^s = \sum_i y_{it}^s$. The risk-averse individual's utility function $U(C_{it}^s)$ has the usual properties and the implicit assumption underlying the resource constraint is that there is no storage and no savings.

Combining the first-order conditions obtained for any two individuals i and j from this constrained maximization problem, full risk-sharing implies the following well known condition:

$$\frac{U'(C_{it}^s)}{U'(C_{jt}^s)} = \frac{\lambda_j}{\lambda_i}.$$

The ratio of marginal utilities for any two individuals will be constant across all states of the world at all points in time. Assuming CRRA preferences, taking logs, summing over all j and then dividing by N , the number of individuals in the mutual insurance arrangement, we arrive at Townsend's regression specification:

$$\log(C_{it}^s) = \frac{1}{N} \sum_{j=1}^N \log(C_{jt}^s) + \left[\frac{1}{\gamma} \left(\log \lambda_i - \frac{1}{N} \sum_{j=1}^N \log \lambda_j \right) \right]$$

where γ is the coefficient of relative risk aversion. This condition should hold in each time period, in any state of the world, and so Townsend's test of full risk-sharing can be easily implemented if panel data over successive years are available:

$$\log(C_{it}) = \alpha \log(y_{it}) + \beta \log(C_t) + f_i$$

where $\log(C_t)$ measures average log-consumption among the participants in the insurance arrangement and the additional variable that is introduced, y_{it} , is the individual's income in period t . For a relatively short panel, it is reasonable to assume that the welfare weights remain constant over time and so the fixed effect f_i will collect all the terms in square brackets above. With full risk-sharing, the individual's consumption in any state of the world will be determined by aggregate consumption ($\beta > 0$), but will be independent of his income ($\alpha = 0$). For the special case with CRRA preferences, $\beta = 1$ as above.

Townsend implements the test of full-insurance by assuming that mutual insurance is organized at the level of the village. Although some risk-sharing mechanisms, notably the local bank and the moneylender, will no doubt operate at this level, we are interested in the role that caste networks play, net of these mechanisms. We consequently investigate whether individual consumption tracks aggregate caste consumption, net of village consumption, using a panel data set covering three successive crop years 1968-69, 1969-70, and 1970-71 at the onset of the Indian Green Revolution.

To implement this modified Townsend test, we need information on each household's *jati* affiliation. The 1982 and 1999 rural surveys that we use for much of the analysis in this paper followed an earlier three-year longitudinal survey, also conducted by the NCAER, over the 1969-71 period. This survey covered 4,118 households in the 17 major states of India and was designed to be representative of the entire rural population of the country in those years. The 1982 survey built on the longitudinal study, adding households where necessary to construct a sample that was representative of the rural population at that later time, while the 1999 survey attempted to track all households in the 1982 round, including those that had partitioned. Detailed *jati* information was not collected in the 1969-71 survey or the 1982 follow-up, but this deficiency was rectified in the 1999 survey round. It is consequently possible to assign *jatis* to those households in the 1969-71 panel who were re-surveyed in 1982 and subsequently in 1999. The test of full-risk sharing, over the 1969-71 period, is consequently restricted to the 1,181 households for which *jati* affiliation is available and which belong to *jatis* with at least 10 sampled households. This sample subset is not a random sample of the 1969-71 households.

However, all time-invariant household characteristics (including the welfare weight) are subsumed in the household fixed effect when implementing the Townsend test.

Table 3, Column 1, begins with Townsend's specification, including village consumption and own income as regressors. Village consumption is measured as the average of log consumption in each village-year. The coefficient on village consumption is 0.73, the coefficient on log income is 0.17. Although there appears to be a fair amount of consumption smoothing, full risk-sharing is rejected - the hypotheses that the own income coefficient is zero and the village consumption coefficient is one are both rejected at the 5 percent level.⁹ Townsend and numerous studies that have followed arrive at essentially the same conclusion. The analysis that follows will explore the additional role played by the *jati* in smoothing consumption.

When the caste system was first established, individuals born into a *jati* were locked into the traditional occupation assigned to it over their lifetimes. These restrictions on occupational mobility gradually weakened over time and today some degree of occupational heterogeneity will exist in any *jati*. What maintains the *jati*'s salience (and its ability to support networks serving many different roles) is the rule of marital endogamy, which we have noted continues to be maintained.

What is the appropriate geographical domain in which intra-*jati* marriages occur? If we take the idea that each *jati* was originally defined by an occupation seriously, then a *jati* could potentially span the entire country. This is not the case in practice, however, because India's many languages create natural social and spatial boundaries. For example, consider the case of the Patils, a cultivator caste from the Marathi-speaking part of the country, and the Patels, also a cultivator caste, but from the adjacent Gujarati-speaking area. The Patils and the Patels have the same traditional occupation and hold comparable positions in the caste hierarchy; judging from their names, these groups clearly served the same economic role in the distant past. Nevertheless, Patils and Patels do not inter-marry, simply because they speak different languages. Modern Indian states are conveniently organized along linguistic lines and so the *jati* statistic that we use in the paper will be constructed within each state.

Table 3, Column 2 includes the average of log consumption in each *jati*-year as an additional regressor to assess the role that caste networks play in smoothing consumption. The coefficient on own income is hardly affected by the inclusion of this additional regressor. However, the coefficient on village consumption does decline and the coefficient on *jati* consumption is positive and significant. Despite the significance of the *jati* coefficient, there are two reasons why this coefficient may be

⁹These estimates are not substantially changed when the sample includes all surveyed households.

underestimated in Column 2. First, the *jati* statistic is based on a sample of households drawn, in turn, from a sample of villages in each state. The measurement error in this statistic will consequently be much greater than the corresponding error in village consumption, which is based on a sample of households in each village. Second, caste loans are used disproportionately to meet relatively infrequent contingencies. The Townsend-test effectively measures the household’s ability to smooth small but frequent income shocks for which the caste network is less useful. Nevertheless, the results in Column 2 indicate that household consumption does co-move significantly with aggregate consumption in the household’s *jati*, consistent with the important role played by caste loans in Table 1.

One concern with the result reported above is that *jati* consumption may simply proxy for unobserved determinants of consumption that are common across households in a geographical area that is larger than the village; for example, a single bank will typically serve multiple villages. To rule out this possibility, the average of log consumption in each district-year is included as an additional regressor in Column 3. Although this variable does have a positive effect on household consumption, the *jati* consumption coefficient remains very stable (a similar result is obtained for state-level consumption).

Our view that social networks in rural India are organized at the level of the endogamous sub-caste is based on the idea that the marriage ties linking members of a *jati* improve information flows and reduce the probability that any individual will renege on his network obligations. An alternative view of the *jati* consumption effects that we obtain is that this variable simply proxies for unobserved socioeconomic characteristics that directly determine consumption and are common to households at the same level in the social hierarchy. Many sub-castes occupy the same position in this hierarchy. We construct an aggregate caste consumption statistic based on information provided in the survey on the household’s social position, at the level of the state-year, and include this variable as an additional regressor in Table 3, Column 4. Aggregate caste-hierarchy consumption does not covary with individual household consumption, whereas the aggregate *jati* consumption coefficient retains its statistical significance.

Although the members of a *jati* may no longer be employed in the same occupation, economic activity continues to be highly correlated within the group for historical reasons. Table 3, Column 5 consequently includes average consumption in the household’s broad occupational category, computed at the level of the state-year, as an additional regressor to allow for the possibility that *jati* consumption is proxying for unobserved occupation-specific shocks.¹⁰ The coefficient on this variable is statistically

¹⁰There are nine occupational categories provided in the data. We used the occupation of the household head to

significant, but the *jati* coefficient once again remains significant and very stable. Thus, it is not unobserved geographic clustering or common socioeconomic characteristics across households that matters for consumption smoothing, but something instead that is specific to the sub-caste or *jati*. We have assumed that a typical *jati* spans a state. Although each major regional language is associated with a single Indian state, multiple states are Hindi-speaking. As a final robustness check, we drop Hindi-speaking states, across which marriages could conceivably take place, in Table 3, Column 6. Matching the stability of the estimates across all the specifications we have experimented with, the coefficients with this reduced sample of households remain very similar to what we obtained with the full sample in Column 2. Although the domain of the social networks cannot be observed directly, the results in Table 3 lend support for the view that they are organized at the level of the *jati* in rural India.

4 Mutual Insurance with Inequality

The statistics reported in Tables 1-2 indicate that caste loans are an important and preferred mechanism for smoothing consumption. The results reported in Table 3 suggest that networks organized at the level of the *jati* facilitate the flow of these loans, but that full risk-sharing is not achieved. In the discussion that follows we describe a model of mutual insurance with limited commitment in which quasi-loans rather than reciprocal transfers emerge as the optimal consumption-smoothing mechanism in equilibrium, consistent with the data. The model is subsequently extended to link the caste network to individual and aggregate mobility. We begin by studying the relationship between wealth inequality within the *jati*, which we expect to have increased with the Green Revolution, on the pattern of loans and the (implicit) interest rate. This allows us to predict which individuals might have been the first to exit and which networks might have been most vulnerable in the aftermath of this exogenous change.

4.1 Caste Loans as Mutual Insurance

For simplicity, consider a mutual insurance arrangement with two individuals and two payoffs: high (H) and low (L). Payoffs are independent across individuals and over time. The probability that individual 1 receives the high payoff and individual 2 receives the low payoff is denoted by P_{HL} . The probabilities of the remaining states of the world occurring are denoted by P_{LH} , P_{LL} , and P_{HH} .

 compute this statistic. Not surprisingly, the majority of households were engaged exclusively in agriculture.

respectively. All four probabilities will, of course, sum up to one. To begin with, assume that $P_{HL} = P_{LH}$, which implies that both individuals are equally wealthy. With a perfect insurance arrangement, these risk averse individuals will consume at a level of $(H + L)/2$ in any period with unequal payoffs, and so will be strictly better off than they would be in autarky. Consumption levels are exactly the same for the two individuals in any state of the world, which implies that the ratio of their marginal utilities will be constant (equal to one in this special case) across all states, satisfying the condition for full risk-sharing derived earlier.

However, perfect insurance might not always be implementable. The individual's incentive to deviate is greatest when he receives H in a given period and his partner receives L . For the special case without commitment, as analyzed by Coate and Ravallion (1993), the individual will weigh the gain from deviating in that period, $H - (H + L)/2$, against the future loss in insurance, assuming that both individuals return permanently to autarky once either deviates. Social sanctions help deter such deviations, but it will often be the case that only partial insurance can be sustained.

Partial insurance without commitment is characterized by a transfer that is strictly less than $(H - L)/2$ in states with unequal payoffs. While the ratio of marginal utilities in states with equal payoffs continues to be one, this is evidently no longer the case in states with unequal payoffs, violating the full risk-sharing condition. Transfers will be reduced as little as possible below $(H - L)/2$, up to the point where the high-payoff individual's participation constraint just binds, but the two individuals will nevertheless often end up consuming at very different levels in states of the world with unequal payoffs.

Ligon, Thoman and Worrall (2002) describe how a higher level of insurance can be sustained with *limited* commitment: under this constrained-efficient arrangement the individual who receives H in a given period t and makes a transfer to his partner who received L will receive compensatory transfers in return that maintain the same ratio of marginal utilities (or as close as possible to that ratio) in all subsequent periods with equal payoffs (L, L or H, H).¹¹ The process starts afresh when unequal payoffs (H, L or L, H) are once again obtained. Although the arrangement with limited commitment may dominate an arrangement without commitment, the individual who receives H will still consume at a higher level than the individual who receives L , which is why transfers must flow in the opposite direction in all subsequent periods with equal payoffs.

¹¹If the ratio of marginal utilities in the initial state with unequal payoffs is set so high that the individual subsequently making the compensatory transfers would prefer to exit the arrangement in either of the states with equal payoffs, then this ratio will be adjusted in that state so that his participation constraint just binds.

Mutual insurance with limited commitment can be characterized as a series of quasi-loans connecting members of the network, whereas full insurance and imperfect insurance without commitment are associated with the flow of gifts or pure transfers between members. The rejection of full risk-sharing and the dominance of caste loans in our data is consistent with the existence of a constrained efficient risk-sharing arrangement in rural networks. The analysis that follows explores how this arrangement would respond to a increase in inequality within the *jati*.

4.2 Wealth Inequality and Mutual Insurance

New High Yielding Varieties (HYVs) of wheat and rice were introduced throughout the developing world in the 1960s, dramatically increasing farm incomes. Certain areas of rural India were better suited to the early HYVs than others and so were quicker to benefit from the new technology. Although the development of hybrid varieties tailored to local growing conditions ultimately made this superior technology available throughout the country, the early start in some areas gave rise to persistent spatial wealth inequality. *Jatis* span a wide area within a state, which implies that wealth inequality would have grown within the caste networks, with some members fortuitously benefitting more from the new technology than others.

To understand the effect of this increase in wealth inequality on the pattern of transfers within the mutual insurance arrangement, we write out a more formal version of the limited commitment model. Let V_l be the net present value to an individual – the lender – who has just received H , while his partner received L , from participating in the arrangement. Let V_b be the corresponding net present value for that individual when he is a borrower, receiving a payoff L , while his partner receives H . We normalize so that the value from deviating is zero. V_l, V_b thus represent the net gain from participation over deviation.

With limited commitment, the lender who has just received H , while his partner received L , will remain in the lending regime, receiving compensatory transfers in return, as long as equal payoffs (H, H or L, L) are obtained. Let $U^l(n)$ be the (certainty equivalent) utility that the lender derives from all lending regimes of length n . Assuming that individual 1 is the lender, V_l can be expressed as

$$V_l = \sum_{n=1}^{\infty} P_n \left(U^l(n) + \delta^n [qV_l + (1 - q)V_b] \right) \quad (1)$$

where P_n is the probability that the current lending regime will persist for n periods, δ is the discount factor, and $q = P_{HL}/(P_{HL} + P_{LH})$ is the probability that individual 1 will enter a fresh

lending regime when the current regime is completed. Payoffs are independent over time and so the expected sequence of payoffs is exactly the same following the H, L state or the L, H state. By symmetry, V_b for individual 1 when he is a borrower can thus be expressed as

$$V_b = \sum_{n=1}^{\infty} P_n \left(U^b(n) + \delta^n [qV_l + (1-q)V_b] \right), \quad (2)$$

where $U^b(n)$ is the (certainty equivalent) utility that the individual derives from all borrowing regimes of length n .

Adding equation (1) and equation (2) above,

$$qV_l + (1-q)V_b = \frac{q \sum_n P_n U^l(n) + (1-q) \sum_n P_n U^b(n)}{1 - \sum_n P_n \delta^n}.$$

Substituting this expression in the V_l equation (1), we finally obtain

$$V_l = \sum_n P_n U^l(n) + \frac{\sum_n P_n \delta^n}{1 - \sum_n P_n \delta^n} \left[q \sum_n P_n U^l(n) + (1-q) \sum_n P_n U^b(n) \right]. \quad (3)$$

Although the lender receives compensatory transfers with the limited commitment arrangement, he is still worse off in any lending regime than he would be in autarky, $U^l(n) < 0$. It is the anticipated benefits of insurance in the future, when he is a borrower, that discourage him from deviating; $U^b(n) > 0$. Given q and δ , the transfers in the lending and the borrowing regimes will adjust in equilibrium such that the individual's participation constraint just binds when he receives H and his partner receives L and he makes his initial transfer. Because we have normalized so that the value of deviating is zero, this implies that $V_l = 0$ in the equation above.

Although V_l in equation (1) was derived for individual 1, the corresponding equation for individual 2 requires only that we redefine q as $P_{LH}/(P_{HL} + P_{LH})$. Up to this point we have assumed that both participants in the insurance arrangement receive the same payoffs in the high and low state, H, L , and have the same probability of receiving the high payoff, $P_{HL} = P_{LH}$, which implies $q = 0.5$ for both participants. To generate a mean-wealth preserving increase in inequality within this arrangement, we could either allow the payoffs or the probability of success to diverge across the two individuals. We first model the mean-wealth preserving increase in wealth inequality as an increase in P_{HL} accompanied by a compensating decrease in P_{LH} because this provides us with unambiguous comparative statics. q increases for the now wealthier individual 1, while q decreases for individual 2.

Holding constant P_{HH} , P_{LL} , and the transfers that were in place prior to the change in q , it follows that $\sum_n P_n U^l(n)$, $\sum_n P_n U^b(n)$, and $\sum_n P_n \delta^n$ will be unchanged. Since $\sum_n P_n U^l(n) < 0$ and $\sum_n P_n U^b(n) > 0$, it is then evident from equation (3) above that an increase in q for the wealthier individual 1 will lead to a decline in V_l , violating his participation constraint.

To bring V_l up to zero once again, the following changes in the pattern of transfers are required. First, the transfer made by the wealthier individual when he receives H and his partner receives L will decline in value. This implies that the transfers that flow in the opposite direction during the remainder of that lending regime must increase in value to maintain the new ratio of marginal utilities.¹² The net effect of these changes in the pattern of transfers will be to increase $U^l(n)$. Second, the transfer received by the wealthier individual when he receives L and his partner receives H will now increase. The transfers that he returns during the remainder of that borrowing regime will decline to maintain the ratio of marginal utilities, and the net effect will be to once again increase $U^b(n)$.

Because the wealthier individual now provides a smaller transfer in the H, L state and receives larger compensatory transfers for the remainder of that lending regime, the implicit interest rate on the loans that he provides must go up. By the analogous argument, the implicit interest rate on the loans that he receives must go down. However, the change in his loan position is ambiguous. Because q has increased for him, he is more likely to be in a lending regime than a borrowing regime. But the size of the loans that he gives out is now smaller, while the size of the loans that he receives is larger. We expect that the change in loan size will be dominated by the first-order change in the probability of being a lender (q), and numerical solutions to the model (not reported) indicate that this is indeed the case.

Thus the model implies that conditional on average wealth in the network, an increase in the household's wealth should lower the interest rate on the loans that it receives, increase the interest rate on the loans that it disburses, and decrease its loan position (loans received minus loans given out). Wealthier households might demand less caste loans because they have the collateral to access bank loans. Conversely, they might have a greater demand for caste loans by virtue of their superior investment opportunities. This implies that the overall effect of household wealth on the loan position will be ambiguous. However, conditional on household wealth, an increase in *jati* wealth will unambiguously increase the household's loan position when caste networks are active.

¹²The implicit assumption here that the borrower's participation constraint does not bind when making these compensatory transfers. If the constraint does bind, then the initial transfer from individual 1 to individual 2 will decline even further.

If we modeled a mean-wealth preserving increase in inequality by an increase in the high payoff in the H,L state for individual 1 and a matching decline in the high payoff in the L,H state for individual 2 (assuming now that q is the same for both individuals), then the effect on the loan position and the interest rate is ambiguous but still potentially consistent with the results above. Holding constant the transfers that were in place with equality, the now wealthier individual 1 gets to keep more in the H,L state than he did before, which pushes him below his participation constraint. At the same time, this risk averse individual benefits less from insurance (over autarky) since he is wealthier, and so the net effect on his participation constraint and, by extension, on the loan position and interest rates in the new regime with inequality, is ambiguous. For the less wealthy individual 2, the forces described above work in the opposite direction, but the net effect will still be ambiguous. Note that such ambiguity does not arise with the alternative formulation of the comparative statics – increasing P_{HL} and decreasing P_{LH} – described above. We normalized so that the payoff from autarky was zero, before and after the increase in inequality, for simplicity. In fact, the gain from insurance declines for individual 1 once inequality is introduced. Because this individual needed additional compensation in any case, this additional force would only reinforce the result derived above. By the same argument, allowing for an increase in the benefit of insurance over autarky would only reinforce the result that the less wealthy individual is willing to give and receive loans at less favorable terms to preserve the integrity of the system.

4.3 Wealth Inequality and Network Stability

The preceding discussion indicates that the network will maintain its stability when faced with changes in wealth positions as long as the pattern of transfers or quasi-loans is sufficiently responsive. In particular, households that become on average better off must receive more favorable terms on loans, with poorer households experiencing a deterioration in loan terms. It is possible, however, that social pressures could prevent such changes from being implemented in practice, with the wealthy increasingly subsidizing poorer households. One motivation for such redistribution would be to ensure that all members of the community remain above a nutrition threshold, which might be an efficiency enhancing policy in a subsistence economy (Polanyi 1957, Gersovitz 1983, Atkeson and Ogaki 1996). Caste loans were also used disproportionately for contingencies such as illness and marriage in Section 2 and we have noted that failure to receive adequate credit when such contingencies arise can have serious consequences. The community might in that case be forced to relax the stipulation that

borrowers fully repay their contingency loans in subsequent states of the world with equal payoffs.

There is an extensive anthropological literature that describes the often substantial redistribution of wealth across households in traditional agrarian economies.¹³ This redistribution was enforced by social norms that sanctioned wealthy individuals who failed to honor their customary obligations and accorded high status to those that did (Scott 1976). If such norms are resilient and prevent changes in the pattern of transfers even as inequality within the network grows then the “tax” on the wealthiest individuals will increase. This could result in exit from the network by the wealthy, once inequality crosses a threshold level. Platteau (1997), for example, documents such patterns of exit from cooperative arrangements among Senegalese fishermen and in a Nairobi slum. Given the increases in intra-caste inequality that accompanied the Green Revolution it is entirely possible that the wealthiest members of the *jati* would have ended up subsidizing the rest of the network. Since the cost of out-marriage and migration would then be lower for them, we would expect such households to be most mobile, *ceteris paribus*.

The key assumption underlying the preceding argument is that those individuals who out-marry or migrate lose the services of the network. We provided an intuitive explanation for why this should be the case in the Introduction, but the theoretical framework allows us to derive this result more formally. We did not introduce social sanctions in the characterization of mutual insurance above, but if deviation is accompanied by a punishment S , then this is equivalent to adding S on the right hand side of the V_i expression in equation (1) and equation (3). Recall that V_i describes the net gain from remaining in the insurance arrangement over deviating, at the onset of a lending regime.

The social sanction S is lowered for the individual who has married outside his *jati* because only the individual himself and his birth relatives, but not his affines (relatives by marriage), can be punished when he deviates. S is lowered for the individual migrant because while his relatives can be punished when he deviates, it is more difficult for the community to reach him. Members of the sub-caste are located throughout the state and so the migrant could potentially be monitored by his network even after leaving the village. In practice, however, *jatis* tend to be concentrated in a relatively small number of villages. There are on average 42.7 *jatis* in a state but only 6.7 *jatis* per village in our sample (which consists of about 40 households per village). A migrant who leaves his village to maximize his wage gain is thus unlikely to encounter other members of his *jati* where he settles. The network’s

¹³Scott (1976) is the classic reference in the literature on the “moral economy,” but see also Popkin (1979) for an opposing view.

ability to punish an individual, S , is consequently lowered for the individual who has married outside his *jati* or migrated, which implies that $V_l < 0$ for him at the level of insurance (lending) that can be sustained by other members of the network. Because this particular individual's ability to provide insurance is relatively limited, the other members of the network will avoid partnering with him in equilibrium.

The model implies that if the pattern of transfers within the mutual insurance arrangement does not adjust sufficiently as inequality grows, then the wealthiest households within the *jati* will end up marrying outside the *jati* or migrating. Conditional on *jati* wealth, an increase in household wealth will increase the propensity to exit the network. Conversely, conditional on household wealth, an increase in *jati* wealth makes the household a net borrower, shifting marriage and migration decisions in the opposite direction.

Variation in (absolute) household wealth will influence marriage and migration decisions independently of the network mechanism. Wealthier households will likely do better on the “open” marriage market, outside the *jati*. This will reinforce the (conditional) household wealth effect on out-marriage derived above. Wealth will also directly affect the ability to finance migration and the opportunity costs of leaving. Wealthier households might possess the resources that are needed to compete successfully in an urban environment, but those households might also have more to lose by leaving. Thus the overall effect of household wealth on migration is ambiguous. Note, however, that conditional on household wealth, an increase in *jati* wealth will unambiguously *decrease* both out-marriage and migration when redistributive norms are in place. Households belonging to wealthier *jatis* are less likely to exit, emphasizing the role of the caste networks in restricting mobility in India.

Apart from identifying which households are most likely to exit, the model also allows us to analyze the effect of an increase in inequality within the *jati* on overall mobility. Consider a transfer from the poor to the rich in a *jati*. This would increase the propensity of the rich to exit, whereas the poor would be even more likely to stay. The impact of a mean-wealth preserving increase in inequality on mobility is consequently ambiguous. Later we will establish empirically that inequality actually reduces mobility, reinforcing the low mobility that historically existed in India.

5 Empirical Analysis

5.1 Loan Terms by Wealth

An implication of the limited commitment model is that relatively wealthy households within a network give and receive caste loans at rates that are more favorable to them. The data on loans from the 1982 and 1999 surveys are consistent with this. We classified households within each *jati* into two wealth categories – low and high – using median wealth *within* the *jati* in the relevant survey round as the cut-off. As can be seen in Table 4, Columns 1-2, wealthier households within the *jati* receive loans at interest rates that are almost two percentage points lower than the interest rates for loans obtained by the less wealthy. The wealthy also disburse loans at interest rates that are over 1.5 percentage points higher than those loans given out by the less wealthy. In contrast, interest rates for loans received from banks are identical for low- and high-wealth households within the *jati* in Columns 3-4. And interest rates on moneylender loans are actually higher for wealthier households in Columns 5-6, consistent with the commonly held view that moneylenders have local monopoly power and price discriminate.

Although the gaps in the lending and receiving interest rates associated with caste loans between low- and high-wealth households are substantial, because of sample size these differences are not statistically significant. Later we will present evidence that relatively wealthy households have a greater propensity to migrate and out-marry, indicating that interest rates do not adjust sufficiently and that the mutual insurance arrangement is not flexible enough to deter exit.

The statistics in Table 4 do not provide estimates of how an exogenous increase in a household's relative wealth (within the *jati*) changes the interest rate on caste loans or how an exogenous change in mobility (out-marriage or migration) affects the household's access to mutual insurance. Given that only a fraction of households receive caste loans, any analysis of interest rates is based on a selected sample and so we do not attempt to identify wealth effects on interest rates. Instead we look at how changes in household and *jati*-level wealth affect the household's caste-loan position to assess the implications of the mutual insurance model with inequality. With respect to mobility, we will study how household wealth and *jati* wealth jointly affect out-marriage and migration. If mobility leads to a loss in network services, then those (relatively wealthy) households who benefit the least from the network when redistributive norms are in place, should be the first to exit. Conversely, *jatis* with low average wealth should be the most vulnerable to such exit.

5.2 Specification and Identification of Wealth Effects

To identify the effects of changes in household and *jati*-level wealth on a household's caste-loan position we will estimate a regression of the form

$$D_{it} = \alpha W_{it} + \beta \bar{W}_t + f_i + \epsilon_{it}, \quad (4)$$

where D_{it} measures household i 's loan position in period t (1982, 1999), W_{it} is its wealth in that period, \bar{W}_t is average wealth in the *jati*, f_i is a fixed effect, and ϵ_{it} collects all other unobserved determinants of D_{it} . Apart from household wealth and *jati* wealth, the household's caste-loan position will depend on other sources of finance (banks and moneylenders), as well as its demand for loans. This demand will depend, in turn, on the household's investment opportunities and its preferences for saving versus consumption. Some determinants of the loan position, such as the household's propensity to save, are time invariant and will be subsumed in the fixed effect f_i . Because two rounds of data are available, we can difference over time to estimate an equivalent regression of the form

$$\Delta D_{it} = \alpha \Delta W_{it} + \beta \Delta \bar{W}_t + \Delta \epsilon_{it}. \quad (5)$$

But changes in investment opportunities or access to finance over time would affect changes in household wealth and (possibly) *jati* wealth, as well as changes in the loan position. For example, as documented by Burgess and Pande (2004), there was a substantial increase in bank coverage in rural areas over the survey interval. The availability of more formal finance would clearly affect investments and wealth accumulation, while at the same time reducing the importance of networks for consumption insurance, although we see that bank loans are less important for this purpose. We will include a variable indicating the presence of a bank in the village in all of the specifications, but changes in access to informal finance and investment opportunities are difficult to observe and measure. Our solution to this identification problem is to instrument for ΔW_{it} and $\Delta \bar{W}_t$. Valid instruments would determine wealth accumulation and, by extension, changes in wealth over time, without being correlated with *changes* in these unobserved variables.

In rural India, wealth accumulation in households has four main sources - growth in the value of fixed assets due to changes in productivity; increases in asset accumulation, such as investment in irrigation; asset sales and purchases; and household division. To eliminate the latter, we started with the 1982 households for whom, based on the 1999 information, we could identify the *jati* of the head

and then aggregated the wealth of any and all of the households in 1999 that split-off from the 1982 households. Thus we have a balanced sample of 3,441 households in each of two years. To increase the precision of *jati*-level aggregates, we eliminated all households in *jatis* with less than 10 surveyed households, resulting in a balanced two-year panel of 2,341 households.¹⁴

As previously discussed, the availability of High Yielding Varieties of wheat and rice in the late 1960s substantially increased farm incomes in India and thus the value of land, particularly irrigated land. However, all areas of the country did not benefit immediately from the new technology. The early HYVs, particularly the rice HYVs, were unsuitable for cultivation in many areas, and it was only by cross-breeding with local varieties that the new technology could be adopted throughout the country (see Munshi 2004 for details). This process was completed by the early 1980s, and so all the households in our sample had access to HYVs in both the 1982 and the 1999 survey rounds. Nevertheless, differences in the timing of initial HYV adoption would have initiated distinct wealth trajectories, across households and *jatis*, which we exploit in the empirical analysis.

The Intensive Agricultural Advanced District Program (IAADP) was introduced in selected districts, typically one per state, in the late 1960s to increase the spread of the new technology. This program was placed in areas that were anticipated to be particularly suited to HYV adoption, and households in the IAADP districts were provided with an assured supply of credit and fertilizers. We consequently include a binary variable indicating whether the household was located in an IAADP village, as well as an indicator for whether any household in its village adopted HYV in 1971, as measures of the timing of HYV adoption.

High Yielding Varieties require expensive inputs such as irrigation and fertilizer, which only wealthy households (or *jatis*) can afford. The amount of land (acres) historically inherited by the household heads in the 1982 survey round would thus have determined the timing of HYV adoption and, by extension, the household's subsequent wealth trajectory. The ability to invest in expensive inputs would in general depend on both wealth and the availability of bank credit, and so a binary variable indicating whether a bank was present in the village in 1971 is included in the set of instruments as well. Because the instruments must predict changes in household wealth as well as *jati* wealth, we complete the set of instruments by including *jati*-level averages of inherited land and the presence of HYV in the village in 1971.

The first stage regression, which includes the instruments discussed above, as well as the change

¹⁴The statistics in Table 4 are computed with the same balanced panel.

in the presence of a village bank over time, which appears as an independent regressor in the second stage, is reported in Appendix Table A1. Inherited land, both at the household and the *jati* level, and the presence of HYV in the village in 1971 (at the level of the *jati*) are significant predictors of changes in household wealth in Column 1. Inherited land and the presence of HYV in the village in 1971, at the level of the *jati* alone, as well as the IAADP indicator, are significant predictors of changes in *jati* wealth in Column 2. The F-statistic testing the joint significance of the excluded instruments is sufficiently large in the first-stage regressions (the p-values are well below 0.05).

To increase the power of the first stage, as a way to improve the precision of the second-stage estimates, we also separated inherited land into irrigated and unirrigated land. The coefficients with this augmented specification in Appendix Table A1, Columns 3-4 are qualitatively similar to what we obtained in Columns 1-2, but the F-statistics are now substantially larger. We will estimate the marriage and migration regressions with both sets of instruments, and while the estimated wealth effects are very stable across the two specifications, they do become more precise (and significant at the 5 percent level) with the full set of instruments.

While there appears to be sufficient power in the instruments that we have chosen we also consider the possibility that the instruments might fail to satisfy the exclusion restriction. Areas that adopted HYV early had very different characteristics from areas that adopted later, and some of these characteristics could, in principle, have been associated with changes in non-farm opportunities or access to finance outside the caste network over time. By the same reasoning, households with greater inherited wealth in 1982 could have been endowed with characteristics such as education or initiative that are associated with *changes* in opportunities or resources in a dynamic economy. Later we will estimate equation (5) with out-marriage and migration as dependent variables. The exclusion restriction in that case is that the timing of HYV adoption should not determine *changes* in the propensity to out-marry or migrate, except through changes in household and *jati* wealth. We have many more instruments than endogenous variables, and so we will later carry out tests of the overidentifying restrictions to verify the validity of these instruments. We will see that these tests do have bite, failing to pass with some of the specifications that we report.

5.3 Descriptive Statistics

Table 5, Panel A presents for 1982 and 1999 the average loan position for the panel households and the value of caste loans, measured by the value of all caste loans received in the survey year plus caste

loans outstanding in that year, in 1982 Rupees. Table 5, Panel A also reports the average value of bank loans and moneylender loans. The level of bank loans is substantially larger than the level of caste loans, while the level of moneylender and caste loans are comparable, matching the patterns reported in Table 1 for the 1982 round. The relatively low average value for the caste loans underestimates their importance since a large fraction of households will report zero loan outstanding in any given year. Conditional on holding a caste loan, the average value of such a loan in 1982 is as high as Rs. 1,220, and we noted earlier that average caste loans for contingencies are four times as large as that amount (roughly equal to the household's annual income). The importance of the caste network in providing credit appears to have been stable over time. Bank loans and moneylender loans are also stable over time, and in general access to finance appears to have changed very little over a relatively long 20-year period.

Table 5, Panel B reports the incidence of marriage outside the *jati* and migration from the village of birth, our two measures of mobility. The measure for out-marriage, constructed from the 1999 marriage histories, is whether any child of the household head married someone who was not a member of the head's *jati* in the 10-year period prior to the survey. The measure of migration is whether any male aged 20-30 at the time of the survey and residing in the household 10 years prior to the survey date had left the village permanently by the survey date.

As can be seen in the table, out-marriage continues to be infrequent in rural India. And the level of male migration actually declines from 1982 to 1999 (the difference over the two survey rounds is significant at the 5 percent level). This is not due to lack of growth - panel C of the table reports average household wealth in the sample for the two survey years. Wealth per-household increased four-fold over the 1982-99 period in real terms, which is a substantial change over what is essentially a single generation. *Jati* wealth, which is computed as the average over the sampled households in each *jati* in each survey year, tracks the household wealth statistic by construction. Finally, consistent with the government program to increase rural bank access over the period, the data indicate that the proportion of households with a bank in the village increased from 0.19 in 1982 to 0.36 in 1999, emphasizing the need to include this variable in the regressions that follow.

5.4 Loan Estimates

The model of mutual insurance predicts that conditional on *jati* wealth, an increase in the household's wealth should make it a net lender. By a symmetric argument, conditional on the household's wealth,

an increase in network partners' (*jati*) wealth should make the household a net borrower. With loans-in minus loans-out as the dependent variable, this implies that $\alpha < 0$ and $\beta > 0$ in equation (5).

Table 6, Column 1 reports instrumental variable estimates of this regression, with the restricted set of instruments reported in Appendix Table A1, Columns 1-2. Consistent with the framework the coefficient on own wealth is negative, while the coefficient on *jati* wealth is positive. Both coefficients are significant at the 5 percent level. Column 2 replaces the net loan position with loans-in as the dependent variable. As expected, the same pattern of (statistically significant) coefficients is obtained.

Table 6, Columns 3-4 include bank loans and moneylender loans as the dependent variables. The theoretical model assumes that the household either participates in the mutual insurance arrangement or quits the network and finds finance elsewhere. In that case, exit from the caste network will be associated with an increased demand for bank or moneylender loans. If bank and moneylender loans are substitutes for caste loans in this way then the coefficient on own wealth will be positive and the coefficient on *jati* wealth will be negative in the bank and moneylender equations. However, we could imagine instead that households obtain capital from different sources simultaneously. The importance of any source of finance will then be determined, in part, by the interest rate on its loans relative to the other sources. We know that the interest rate on caste loans is declining, and hence the demand for those loans is increasing, in household wealth (conditional on *jati* wealth). The demand for bank loans and moneylender loans must then go in the opposite direction and so the coefficient on own wealth will be negative, while the coefficient on *jati* wealth will be positive. Additionally, banks or moneylenders could view borrowers (of given own wealth) from *jatis* that are more wealthy as being more credit worthy so that bank/moneylender loans and caste loans are complements.¹⁵ In that case the *jati* wealth coefficients would display the same signs in all three loan equations. Finally, a household's loan position will also depend on its absolute wealth. Most bank loans require collateral, which a wealthier household is better positioned to provide. Recall, also, from Table 3 that wealthier households paid much higher interest rates on moneylender loans, which would lower their demand for that source of finance.

In general the sign of the own wealth and the *jati* wealth coefficient is ambiguous in both the bank and the moneylender loan regressions. Not surprisingly, there is no particular pattern to those coefficients in Columns 3-4, in contrast to what we obtained in Columns 1-2, although bank loans look

¹⁵The 1999 data indicate that 2.3 percent of caste loans are obtained to repay loans from other sources; neither bank nor moneylender loans are used for this purpose.

more like complements to caste loans than do moneylender loans. However, the *jati* wealth coefficient is insignificant in both Column 3 and Column 4. This contrasts with the strong relationship between *jati* wealth and caste loans, suggesting that the *jati*-level variable is not just picking up an aggregate demand for loans. Rather, the set of loan results by source are consistent with the view that caste networks are able to mobilize a particular type of credit for their members.

The effect of changes in household and *jati* wealth on the household's propensity to exit the network, as a consequence of changes in its caste loan position, are difficult to interpret from Table 6. Although the effect of these wealth changes on the household's average loan position may be small, the effect on the small fraction of households that did hold caste loans in the year preceding the survey will be much larger. All households will demand contingency loans, albeit infrequently, from their caste networks. These loans are utilized to cover large negative income shocks and a risk-averse household would be very sensitive to even small changes in their size. The estimates in the caste loan regression are most useful in establishing that the coefficients on household and *jati* wealth have the expected patterns of signs implied by the limited-commitment framework. To assess the impact of these changes in caste credit and the terms of the caste loans on mobility, we turn next to the marriage and migration regressions.

5.5 Migration and Marriage Estimates

Table 7 reports estimates of household and *jati*-level wealth effects on network exit measured by migration and out-marriage. We noted earlier that the coefficient on household wealth is difficult to interpret unambiguously in these regressions because an absolute increase in wealth (independent of *jati* wealth) could directly affect marriage and migration decisions other than through network effects. Nevertheless, we see that the coefficient on own wealth is positive across Columns 1-4 in Table 7. This finding is consistent with transfers within the *jati* not being sufficiently responsive to changes in inequality to deter the wealthier households from exiting. The estimated *jati* wealth effects are most useful in establishing a role for the caste networks in determining individual mobility, and here we see that the coefficient on *jati* wealth is negative across all four columns. Combined with the caste loan estimates in Table 6, the results indicate that households who are more (less) likely to be net borrowers because they are members of wealthier *jatis* are also those households who less (more) likely to exit.

The magnitude of the effect of a change in average *jati* wealth on mobility depends, as indicated

in the model, on the extent to which loan terms can respond in response to changes in the *jati* wealth distribution. If terms are completely flexible, then changes in *jati* wealth characteristics will have no effect on mobility. As seen in Table 5, average wealth in the sample increased from 5 to 20 thousand Rupees between 1982 and 1999. The point estimates in Table 7, Columns 1-4 indicate that if *jati* wealth increased by that amount but own household wealth remained the same, the propensity to out-marry would decrease by 1.4 percentage points (20 percent) and out-migration would decline by from 1 to 1.4 percentage points (10 to 14 percent). These effects are non-trivial but relatively small, indicating that while the rural networks are not able to fully-adjust loan terms, these terms are flexible.¹⁶

The point estimates from Columns 3 and 4 can also be used to assess the impact of overall economic growth on mobility. In particular, we can use the estimates to assess the effect of an increase in average wealth on out-marriage and migration. Because economic growth increases both own wealth and *jati* wealth, the net effect of wealth change is obtained by adding the own wealth and *jati* wealth coefficients. The summed point estimates indicate that the fourfold increase in wealth between 1982 and 1999 described above would have actually reduced out-marriage, by 0.4 percentage points and only increased migration by 0.75 percentage points. Thus the substantial increase in average rural wealth that occurred in rural India would have had little effect on out-marriage or migration, consistent with the view that it is inequality within the *jati* that matters for network stability.

Although the networks may have been resilient to changes in average wealth, individuals towards the top of the wealth distribution nevertheless appear to have had a greater propensity to exit, by out-marrying and migrating outside the village. The median of the standard deviation of wealth within *jatis* is Rupees 13,318 in the 1999 sample. The point estimates in Columns 3-4 indicate that increasing the wealth of a household by this amount, with no change in average *jati* wealth would have increased the probabilities that a member of the household marries someone outside the *jati* by 0.8 percentage points (a 9 percent increase, based on the 1999 level) and migrates from the village by 1.9 percentage points (a 32 percent increase, based on the 1999 level). Those individuals that the theoretical framework identifies as having the greatest propensity to exit the network (and who, as indicated in Table 6 receive less caste loans) are indeed substantially more likely to out-marry and

¹⁶A couple of coefficients in Columns 1-2, estimated with the restricted set of instruments, are not significant at the 5 percent level, but they become significant with the full set of instruments in Columns 3-4. The F-statistic on the over-identification test is not surprisingly larger in Columns 3-4, particularly with out-marriage as the dependent variable. Notice, however, that the point estimates are very stable across the alternative specifications.

migrate, reinforcing the connection between exit and these variables.

What about the effects on mobility due to rising inequality within *jatis*? The regression specification in Columns 1-4 does not permit an assessment of the effect of a mean-preserving change in inequality, implemented by transferring wealth from the bottom of the distribution to the top of the distribution, because own wealth effects are assumed to be linear (symmetric for low and high wealth individuals). It is evident that transfers of this sort would have no net effects on exit. The theoretical framework does not, however, imply that wealth effects are symmetric; this is just an assumption that we adopt for convenience in Columns 1-4. We consequently test the symmetry assumption in Columns 5-6 by interacting own wealth with a binary variable, which indicates whether the inherited wealth of the household head in the 1982 survey round lies above the average inherited wealth in his *jati*.¹⁷

The estimates of the own wealth-above mean interaction coefficients in Columns 5-6 indicate that the response of out-marriage and migration to wealth change is larger (although not statistically significant) for households below the average wealth in their *jati*. The point estimates indicate that if 13,318 Rupees (the median of the standard deviation of wealth within *jatis*) were transferred from a household with wealth below the mean to a household above the mean, out-marriage would have declined by only 1/2 a percentage point. However, migration would have declined more substantially, by 2.8 percentage points. The increase in inequality associated with the Green Revolution would thus have actually dampened out-migration, which was low to begin with, even further.

Finally, although the traditional caste networks have evidently been robust to economic growth and changes in inequality in the short run, our estimates do indicate that exit is more likely among the wealthy within the network, consistent with the mutual insurance model. This selective exit will over time worsen the average wealth of the network. The estimates of the *jati* wealth effect in Columns 3-4 permit an assessment of the effect of this decay in the network on future mobility. The average *jati* wealth in the 1999 round was Rupees 20,445. Picking the *jati* whose average wealth was closest to this number, we discarded the top 10 percent of households from its wealth distribution and re-computed the average wealth. The average wealth in that *jati* substantially declined, to Rupees 11,593. The associated 8,852 Rupee decline in *jati* wealth, however, would have increased both out-marriage and migration by only 0.8 percentage points. Neither network decay effect is especially large, which suggests that the caste networks might remain firmly in place, and that mobility will continue

¹⁷The above-mean indicator is interacted with the full set of instruments used in Columns 3-4 to complete the set of instruments in Columns 5-6. The above-mean indicator does not appear in the second stage because it is subsumed by the household fixed effect.

to remain low in the future, unless comparable alternative credit arrangements become available.¹⁸

6 Conclusion

In this paper we have examined the hypothesis that the persistence of low spatial and marital mobility in rural India, despite increased growth rates and rising inequality in recent years, is due to the existence of caste networks that provide mutual insurance to their members. Unique panel data identifying sub-caste membership and providing information on caste loans, marriage, and migration are used to link caste networks to household and aggregate mobility. We show that sub-caste networks are particularly important in providing assistance in the form of low-interest loans with no collateral requirements in times when households incur unusually large expenditures associated with marriage and illness. Moreover, consistent with a limited-commitment mutual insurance model the data indicate that wealthier households within sub-caste networks pay even lower interest rates on caste loans and charge higher interest rates to same-caste borrowers compared with less wealthy members of the same sub-caste. These households are nevertheless more likely to both migrate and inter-marry, suggesting that they are not being adequately compensated by the network in their role as net lenders. Conversely, among households with the same wealth, those in higher-wealth caste networks are more likely to obtain loans and are less likely to be mobile, providing direct evidence that the networks restrict mobility.

Although our analysis connects caste networks to mobility, it does not indicate the extent to which mobility would increase if an alternative credit arrangement of comparable quality became available. A comparison of marriage trends in urban and rural India provides us with a sense of how large the consequent increase in mobility could be. Figure 3 reports out-marriage in Bombay city, based on a survey of school children conducted by the authors in 2001. As in rural areas, in the 1970's and 1980's out-marriage in Bombay was very rare. However, in contrast with the stability that we observed in Figure 2 for rural households, out-marriage increased steeply over time in Bombay, particularly in the 1990's, starting at 2 percent in the early 1970's and reaching as high as 12 percent by the late 1990's.¹⁹

Munshi and Rosenzweig (2006) describe how caste-based labor market networks historically assisted their members in finding well-paying jobs in particular occupations in Bombay. They also

¹⁸Alternative arrangements may not have emerged precisely because the caste networks are so effective. Kranton (1996) provides a formal model that describes this tension between market and non-market institutions.

¹⁹To be comparable with Figure 2, the statistics in Figure 3 are computed using the parents and siblings of the sampled school children.

document restrictions on mobility that were imposed by those networks, and the low out-marriage in the 1970's and 1980's is consistent with the presence of such restrictions. However, the urban caste-based job networks quickly became irrelevant with the restructuring of the Indian economy in the 1990's and the availability of new economic opportunities outside of the traditional caste occupations. This is reflected in the steep increase in out-marriage in the post-1990 period. In contrast, the principal role of rural caste networks in providing consumption insurance was not eroded by the recent developments in the Indian economy. The urban experience suggests, however, that if the rural caste networks ceased to be salient because of the availability of new, superior forms of credit, mobility could surge in rural India in the future as well. Although such mobility would be growth-enhancing, it is important to recognize that alternative risk-smoothing mechanisms would not necessarily provide the subsidized insurance to the poor that is a key feature of the caste network.

References

- [1] Anderson, Siwan. 2004. Dowry and Property Rights. University of British Columbia.
- [2] Atkeson, Andrew and M. Ogaki. 1996. Wealth Varying Intertemporal Elasticities of Substitution: Evidence from Panel and Aggregate Data. *Journal of Monetary Economics*. 38, no. 3, pp. 507-534.
- [3] Attanazio, Orazio and Pedro Albarran. 2003. Limited commitment and crowding out of private transfers: evidence from Mexican villages. *Economic Journal*. 113, 486.
- [4] Banerjee, Abhijit and Andrew Newman. 1998. Information, the Dual Economy, and Development. *Review of Economic Studies*. 65, pp. 631-653.
- [5] Banerjee, Abhijit, Lakshmi Iyer and Rohini Somanathan. 2005. Public Action for Public Goods. MIT.
- [6] Bardhan, Pranab and Christopher Udry. 1999. *Development Microeconomics*. Oxford: Oxford University Press.
- [7] Burgess, Robin and Rohini Pande. 2005. Do Rural Banks Matter? Evidence from the Indian Social Banking Experiment. *American Economic Review*. June.
- [8] Caldwell, John C., P.H. Reddy, Pat Caldwell. 1986. Periodic High Risk as a Cause of Fertility Decline in a Changing Rural Environment: Survival Strategies in the 1980-1983 South Indian Drought. *Economic Development and Cultural Change*. pp.677-701.
- [9] Coate, Stephen and Martin Ravallion. 1993. Reciprocity without Commitment: Characterization and Performance of Informal Insurance Arrangements. *Journal of Development Economics*. 40, pp. 1-24.
- [10] Fafchamps, Marcel and Susan Lund. 2000. Risk-Sharing Networks in Rural Philippines, mimeo, Oxford University.
- [11] Foster, Andrew D. and Mark R. Rosenzweig. 1996. Technical Change and Human Capital Returns and Investments: Evidence from the Green Revolution. *American Economic Review*. September.
- [12] Foster, Andrew D. and Mark R. Rosenzweig. 2005. Agricultural Development, Industrialization and Rural Inequality, mimeo, Harvard.

- [13] Gersovitz, Mark. 1983. Savings and Nutrition at Low Incomes. *Journal of Political Economy*. 91, 5. pp. 841-55.
- [14] Greif, Avner. 1993. Contract Enforceability and Economic Institutions in Early Trade: the Maghribi Traders' Coalition. *American Economic Review*. 83, 3, pp. 525-48.
- [15] Kranton, Rachel. 1996. Reciprocal Exchange: A Self-Sustaining System. *American Economic Review*. 86,4, pp. 830-851.
- [16] La Ferrara, Eliana. 2002. Inequality and Group Participation: Theory and Evidence from Rural Tanzania. *Journal of Public Economics*. 85, pp. 235-273.
- [17] Ligon, Ethan. 1998. Risk-Sharing and Information in Village Economies. *Review of Economic Studies*. 65, pp. 847-864.
- [18] Ligon, Ethan, Jonathan P. Thomas and Tim Worrall. 2002. Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies. *Review of Economic Studies*. 69, pp. 209-244.
- [19] Luke, Nancy and Kaivan Munshi. 2007. Social Affiliation and the Demand for Health Services: Caste and Child Health in South India. *Journal of Development Economics*. 82(2), pp. 256-279.
- [20] Morduch, Jonathan. 2004. Consumption Smoothing Across Space: Testing Theories of Risk-Sharing in the ICRISAT Study Region of South India. in Dercon, Stefan, ed., *Insurance Against Poverty*. Oxford Press.
- [21] Munshi, Kaivan. 2004. Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution. *Journal of Development Economics*. 73, 1, pp. 185-215.
- [22] Munshi, Kaivan and Mark Rosenzweig. 2006. Traditional Institutions Meet the Modern World: Caste, Gender and Schooling Choice in a Globalizing Economy. *American Economic Review*. 96(4): 1225-1252.
- [23] Platteau, Jean-Philippe. 1997. Mutual Insurance as an Elusive Concept in Traditional Rural Communities. *Journal of Development Studies*. 33, 6, pp. 764-796.
- [24] Polanyi, Karl. 1957. *The Great Transformation*. Boston: Beacon Press.

- [25] Popkin, Samuel L. 1979. *The Rational Peasant: The Political Economy of Rural Society in Vietnam*. Berkeley: University of California Press.
- [26] Scott, James C. 1976. *The Moral Economy of the Peasant: Rebellion and Subsistence in Southeast Asia*. New Haven: Yale University Press.
- [27] Townsend, Robert. 1994. Risk and Insurance in Village India. *Econometrica*. 62, 3, pp. 171-84.
- [28] UNDP. 2002. *Deepening Democracy in a Fragmented World*. Human Development Report.

Figure 1.
Change in Percent Urbanized, by Country, 1975-2000




Figure 2
Rates of Out-Marriage (Hindus), by Decade, Rural India 1950-1999 (N=31,529)




Figure 3
Rates of Out-Marriage (Hindus), by Quinquennia, Mumbai 1970-2002 (N=5,406)




Table 1: Loans Received by Source and Purpose, 1982

Loan source:	Caste (1)	Bank (2)	Moneylender (3)	Other (4)
Total loan value (%)	12.33	46.30	12.19	29.18
Loan value by purpose (%):				
Investment	17.07	26.47	16.83	39.63
Operating expenses	6.08	53.47	1.82	38.63
Contingencies	42.61	20.56	27.48	9.35
Consumption	23.11	15.08	47.42	14.39

Note: Statistics are computed using the 1,423 loans received by the sampled households in the 1982 survey round.

Loan value is computed as the percentage of all loans (by value) received for that purpose that is provided by that source.

Loan values sum up to 100 across the four sources in each row.

Investment includes land, house, business, etc.

Operating expenses are for agricultural production.

Contingencies include marriage, illness, etc.

Table 2: Terms of Loans, by Source and Year

Year:	1982			1999		
	Caste	Bank	Moneylender	Caste	Bank	Moneylender
Source:	(1)	(2)	(3)	(4)	(5)	(6)
Interest rate	10.70 (0.50)	14.88 (0.47)	16.99 (0.42)	8.23 (0.91)	10.16 (0.23)	30.63 (2.30)
Percentage zero-interest loans	34.87	0.27	2.84	59.78	0.17	15.07
Percentage loans requiring collateral	16.23	48.95	18.99	1.31	83.21	24.78

Note: Statistics are computed using the 1,423 loans received by the sampled households in the 1982 round and the 1,687 loans received in the 1999 round.

Statistics are weighted by the value of the loan.

Standard errors in parentheses.

Table 3: Tests of Full Risk-Sharing

Dependent variable:	log own-consumption					
	(1)	(2)	(3)	(4)	(5)	(6)
Log own-income	0.174 (0.033)	0.171 (0.033)	0.172 (0.033)	0.168 (0.034)	0.169 (0.033)	0.137 (0.032)
Village log-consumption	0.725 (0.033)	0.635 (0.042)	0.576 (0.047)	0.638 (0.039)	0.624 (0.045)	0.669 (0.045)
Jati log-consumption	--	0.232 (0.031)	0.216 (0.031)	0.239 (0.034)	0.209 (0.029)	0.247 (0.033)
District log-consumption	--	--	0.095 (0.048)	--	--	--
Caste log-consumption	--	--	--	-0.024 (0.065)	--	--
Occupation log-consumption	--	--	--	--	(0.123) (0.073)	--
R-squared	0.882	0.884	0.884	0.883	0.884	0.867
Number of observations	3,543	3,543	3,543	3,387	3,543	2,184

Note: regressions use three years of data 1969-71 for each household.

All regressions include household fixed effects.

Standard errors in parentheses are clustered at the state-year level.

Only jatis with more than 10 households in the sample are included in the regressions.

Caste in column 4 measures broad hierarchical category and the corresponding statistic is computed in each state-year.

Occupation in column 5 measures broad occupational category and the corresponding statistic is computed in each state-year.

Column 6 excludes Hindi speaking states (Bihar, Madhya Pradesh, Rajasthan, Uttar Pradesh).

Table 4: Interest Rates by Source and Household Wealth

Loan source: Wealth category:	Caste		Bank		Moneylender	
	High	Low	High	Low	High	Low
	(1)	(2)	(3)	(4)	(5)	(6)
Borrowing	10.08 (0.83)	11.98 (0.69)	12.09 (0.25)	12.08 (0.26)	28.78 (2.04)	14.22 (0.66)
Lending	10.83 (1.82)	9.20 (2.16)	--	--	--	--

Note: the household is the unit of observation.

Interest rates at the level of the household are weighted by household sampling weights to compute reported statistics.

Interest rates are computed by pooling loans in 1982 and 1999.

Standard errors are in parentheses.

The cut-off separating low and high wealth is the median wealth level within the jati in each year.

The hypothesis that the mean interest rate for low and high wealth households is equal cannot be rejected at the 5 percent level except for borrowing from moneylender.

Table 5: Descriptive Statistics, Panel Sample

Year:	1982	1999
	(1)	(2)
Panel A: Loan Value by Source		
Caste loans-in minus loans-out	44.21 (31.55)	41.34 (13.83)
Caste loans	71.42 (11.43)	81.72 (10.78)
Bank loans	393.96 (89.54)	235.39 (35.03)
Moneylender loans	47.77 (7.61)	46.13 (10.42)
Panel B: Marriage and Migration		
Out marriage	0.07 (0.01)	0.09 (0.01)
Migration	0.10 (0.01)	0.06 (0.01)
Panel C: Wealth and Access to Banks		
Household wealth	4831.91 (163.98)	20311.48 (1408.72)
Jati wealth	4609.11 (81.18)	20103.21 (1182.78)
Bank in village	0.19	0.36

Standard errors in parentheses. All statistics are computed using sample weights.

Statistics are computed using households in the 1982-1999 panel.

Statistics computed using jatis with at least 10 households in sample and households with heads at least age 35 in 1982.

Table 6: FE-IV Loan Estimates

Dependent variable: Loan source:	Loan position		Loans in	
	Caste		Bank	Moneylender
	(1)	(2)	(3)	(4)
Household wealth	-0.014 (0.006)	-0.009 (0.004)	-0.060 (0.031)	0.004 (0.004)
Jati wealth	0.006 (0.003)	0.007 (0.003)	0.020 (0.026)	-0.006 (0.004)
Bank in village	17.273 (98.082)	141.654 (68.575)	453.976 (340.690)	31.762 (96.710)
Constant	-9.982 (77.736)	-73.021 (75.098)	219.995 (124.592)	12.409 (68.736)
F statistic (over-id test)	0.88	1.46	1.39	2.14
p-value	0.53	0.26	0.28	0.11
Number of observations	2,094	2,094	2,094	2,094

Standard errors in parentheses are robust to clustering at the state level.

Loan position is measured as loans in minus loans out.

Instruments include inherited land, initial HYV adoption (IAADP and HYV adoption in the village in 1971), bank in 1971.

Inherited land and HYV in the village in 1971 are computed at household and jati level.

Sample restricted to jatis with at least 10 households in sample and households with heads at least age 35 in 1982.

Table 7: FE-IV Out-Marriage and Migration Estimates

Instrument set: Specification: Dependent variable:	Restricted		Full			
	Symmetric wealth effects		Asymmetric wealth effects			
	Out marriage	Migration	Out marriage	Migration	Out marriage	Migration
	(1)	(2)	(3)	(4)	(5)	(6)
Own wealth x 10 ⁻⁶	0.62 (0.34)	1.24 (0.55)	0.63 (0.31)	1.41 (0.57)	0.92 (0.77)	3.29 (1.36)
Own wealth*above-mean x 10 ⁻⁶	--	--	--	--	-0.39 (0.44)	-1.20 (0.71)
Jati wealth x 10 ⁻⁶	-0.93 (0.36)	-0.64 (0.48)	-0.88 (0.34)	-0.91 (0.45)	-0.98 (0.43)	-2.06 (1.25)
Bank in village x 10 ⁻²	-0.70 (0.85)	-0.25 (1.71)	-0.69 (0.84)	-0.35 (1.73)	-0.85 (1.04)	-1.19 (2.44)
Constant x 10 ⁻²	2.94 (0.83)	3.03 (2.14)	2.81 (0.86)	3.14 (2.10)	2.84 (0.84)	2.30 (2.08)
F statistic (over-id test)	0.56	0.25	2.23	1.07	0.69	0.40
p-value	0.76	0.95	0.09	0.44	0.81	0.98
Number of observations	896	925	896	925	896	925

Standard errors in parentheses are robust to clustering at the state level.

Instruments include inherited land, initial HYV adoption (IAADP and HYV adoption in the village in 1971), bank in 1971.

Inherited land and HYV in the village in 1971 are computed at household and jati level.

Full set of instruments separates irrigated and unirrigated inherited land.

Regressions with asymmetric wealth effects interact own wealth with a binary variable indicating whether inherited wealth in 1982 was above jati mean.

Additional instruments in Columns 5-6 include full set interacted with binary variable indicating whether inherited wealth in 1982 was above jati mean.

Sample restricted to jatis with at least 10 households in sample and households with heads at least age 35 in 1982.

Table A1: First Stage Estimates

Dependent variable:	Household wealth	Jati wealth	Household	Jati wealth
	change	change	wealth change	change
	(1)	(2)	(3)	(4)
Inherited land	13.84 (2.56)	0.02 (1.47)		
Inherited land (jati average)	47.98 (15.56)	77.81 (25.09)		
Inherited unirrigated land	--	--	14.66 (4.20)	-0.44 (1.77)
Inherited irrigated land	--	--	13.63 (6.13)	3.61 (6.61)
Inherited unirrigated land (jati average)	--	--	26.27 (9.91)	55.32 (19.13)
Inherited irrigated land (jati average)	--	--	87.04 (14.92)	117.48 (45.97)
HYV in the village in 1971 x 10 ³	1.66 (2.80)	-1.85 (1.73)	1.09 (2.61)	-2.78 (1.81)
HYV in the village in 1971 x 10 ³ (jati avg.)	18.36 (7.44)	29.96 (11.92)	14.74 (5.92)	26.35 (10.77)
IAADP district x 10 ³	5.72 (3.84)	11.56 (4.89)	3.42 (3.30)	8.92 (4.22)
Village bank in 1971 x 10 ³	-0.33 (2.71)	-2.91 (2.98)	-0.65 (3.29)	-3.33 (2.89)
Bank change (1982-1999) x 10 ³	-0.27 (3.79)	-5.00 (4.20)	-1.49 (3.56)	-6.20 (4.69)
F statistic	7.79	3.24	32.97	3.68
p-value	0.0008	0.0328	0.0000	0.0146
R-squared	0.087	0.198	0.100	0.219
Number of observations	2094	2094	2,094	2,094

Standard errors in parentheses are robust to clustering at the state level.

Dependent variables are computed as the change between 1982 and 1999.

All variables in the regression are excluded from the second stage except bank change (1982-99).

Regressions restricted to jatis with at least 10 households in sample and households with heads at least age 35 in 1982.