Chapter 15

Insurance

15.1. Basic concepts

In this chapter, we take up the study of *insurance*. The main ideas that we seek to develop can be illustrated with an example that involves just two people.

Suppose that Asaf and Sharif are two farmers. We begin by supposing that they produce the same crop and that they use the same amounts of land and other inputs. Suppose that the harvest can take one of two values: $2,000 if all works out well and $1,000 if there is some damage to the crop. Assume that the chances of either of these events occurring is 1/2, and that this probability is the same for both Asaf and Sharif.

If Asaf and Sharif do not know each other, then they will try to do something about this uncertainty on their own. You have seen several examples by now of risk aversion and how this gives rise to the desire to smooth consumption, and we need not repeat the details here. Suffice it to say that both Asaf and Sharif would like to even out their consumption. They would like to run down their wealth a bit to make up for the shortfall when the crop is bad and stock some of the excess when the crop is good.

Using one's own wealth to smooth uncertain shocks in income is called *self-insurance*. Essentially, individuals with the wealth to do this are protecting themselves against uncertainties that they cannot control. Self-insurance can work through several channels. Stocks of cash or accumulated savings in banks can be run down (or added to) for the purpose. The same can be done with grain stocks, although holding such stocks can be costly because grain is not perfectly durable. However, this may be a preferred form of savings if rural banks are few and far between or if there are restrictions on the rapid liquidation of savings. Other assets may be run down or accumulated as well. Livestock and jewelry are two such assets, although these may be sold as a last resort. The box on self-insurance in Indian villages using bullocks will give you a flavor of how self-insurance works in practice.

Self-Insurance and Bullocks

Possibly the most important mechanism for consumption smoothing, other than market credit and intra- or interfamily lending, is the sale and purchase of assets. Farmers hold many different forms of wealth, including land, jewelry, currency, capital goods (such as pump sets), animals (such as bullocks), and stocks of food grain. It may be possible to use such assets as *buffer stocks*; that is, families accumulate stocks in periods of relative affluence and deplete these reserves to finance consumption expenditures in more difficult times. Do rural households engage in substantial buying and selling of assets for this purpose? A study by Rosenzweig and Wolpin (1993), based on the ICRISAT villages in India, suggested that for certain forms of assets, this is indeed the case.

Immobile capital—land and buildings—constitute a major part of farmers' wealth in this study, accounting for approximately 85% of total wealth. However, it is a widely observed fact that although the market for land *rentals* is very active in these regions, the asset market for land is highly inactive. Of all rural households surveyed in 1970–71 by the National Council of Applied Economic Research in India, only 1.5% undertook any kind of land sale, which suggests that this method is rarely used to finance consumption in times of need. In the non-real-estate wealth category, financial assets (stocks, bonds, etc.) have a very small share (less than 5% even for large farmers), and thus play an insignificant role. Average crop inventories held over the year accounts for about a quarter of this wealth; however, it varies a lot across harvest and nonharvest seasons, and is probably used more to smooth consumption over the year rather than across years. A sizable portion (about 19% on average) of nonland wealth is held in the form of jewelry, but the data indicate that buying and selling jewelry, too, is minimal. By far the largest component of nonland wealth is bullocks (representing 50% of the wealth for small farmers, over 33% for midsize ones, and about 27% for large farmers), and evidence is strong that farmers vary their ownership of bullocks as a primary instrument to smooth consumption.

The asset and rental markets for bullocks have exactly the opposite characteristics compared to those for land. Although there is an extremely well organized, regionally integrated market for bullocks,[1] short-term bullock leases are extremely uncommon (see Chapter 12 for further discussion). The nonexistence of the market for bullock power (as opposed to the market for bullocks) implies that there can be substantial productivity gains for a farmer who owns an animal or two; thus, from a purely productivity viewpoint, turnover in bullock ownership should be expected to be low.[2] However, a ten-year survey revealed that 86% of households were involved in at least one transaction in bullocks, indicating that many of these were probably "distress sales," motivated to meet consumption requirements.

[1] In India, large centralized bullock "fairs" are held in many regions at specific times of the year.

[2] Part of the potential efficiency losses arising from the absence of free flow in bullock power services is mitigated by the active rental market in land. Rosenzweig and Wolpin estimated that a farmer who did not own a bullock was 63% more likely to lease out land than a farmer who owned at least one.

A careful analysis of the data shows that this is indeed the case. Due to the presence of well integrated bullock markets, bullock prices are largely immune to village-specific production shocks, which is precisely the factor against which farmers seek insurance. This is highlighted by the fact that over 60% of bullock sales were made to buyers outside the village, with 10% going to buyers located more than 20 kilometers away.[3] The statistical results show that the probability of a bullock purchase increases significantly when income is high, and the probability of a sale decreases. Moreover, other things held constant, those people holding higher stocks of bullocks (or even pump sets) are less likely to make a purchase in the future, suggesting that on the average, just as with any other buffer stock, farmers try to maintain a target level of the asset. This pattern is strongly borne out in data on medium and small farmers. For large farmers, the fit is much weaker, suggesting that this group has much greater access to credit and other conventional instruments and does not have to rely heavily on asset sales as a means to achieve consumption targets.

Does the widespread incidence of distress sales (coupled with the absence of rental markets for draft animals) hamper efficiency and lower average agricultural output? Rosenzweig and Wolpin felt that it does. At the prevailing market prices for bullocks and other inputs, their profit function estimates show that the optimal number of bullocks owned by midsized farmers is about two. However, the average size of the bullock stock is 0.94. Through numerical simulations of their econometric model, they estimated that the provision of actuarially fair weather insurance does not increase farmer welfare substantially. The weather is only one source of risk, and farmers already have some insurance through limited credit. On the other hand, the creation of supplementary assured sources of income (say through a rural job creation policy) raises bullock stocks closer to optimal levels and thereby increases agricultural efficiency.[4]

The second major form of smoothing is taking recourse to credit. We have already discussed this in great detail and there is nothing that we need add here.

[3] This observation is unlikely to hold with much generality across countries, especially in those areas where intervillage integration is relatively small. For example, Czukas, Fafchamps, and Udry [1997] test whether livestock is used as a consumption smoothing device by households in Burkina Faso. Unlike the observations for India, the authors find that the effects of random shocks on livestock sales are indeed small, suggesting both that markets are local and that there are a fair number of alternative smoothing options.

[4] We need to be a little careful, though, in interpreting these results and drawing the conclusion that a serious "underinvestment" in bullocks exists. After all, the *aggregrate* supply of bullocks must be quite inflexible at least in the short run, and can increase only to the extent that more favorable conditions lead to greater breeding. The problem is that profit and optimality calculations are performed under the *prevailing* market price for bullocks: if consumption smoothing motives are removed through the provision of alternative sources of insurance, the demand and supply conditions in this market change entirely, leading to entirely new equilibrium prices and optimality calculations. However, the mere stabilization of bullock stocks leads to some increase in output and welfare, even if the *average* size of these holdings does not increase.

What I would like to focus on here is a third source of income smoothing: that attained through mutual insurance between economic agents. As the discussion continues, you will see that this source of smoothing is, after a point, difficult to distinguish from credit, but for the moment, we postpone these complications.

To develop this idea, bring Asaf and Sharif together. We can think of four possible outcomes: (1) both Asaf and Sharif produce the good output; (2) both produce the bad output, and [(3) and (4)] one produces high while the other produces low. In the first two cases there is little that Asaf and Sharif can do to help each other; they are in identical positions. In the last two cases, they can agree to share their output equally. What this boils down to is a promise that Asaf will pay Sharif $500 when he produces high and Sharif produces low, and Sharif will pay Asaf $500 when the opposite occurs.

The following observations are to be noted. First, *if the chances of producing high output are independent across Asaf and Sharif*, elementary probability tells us that the chance of each of these four cases is 1/4. Without any consumption smoothing, Asaf's consumption in the four situations is (2,000, 1,000, 2,000, 1,000); a similar statement holds for Sharif. With the mutual insurance scheme in place, this consumption stream changes to (2,000, 1,000, 1,500, 1,500). This reflects the fact that in the first two states there is nothing to be done in the way of mutual insurance, whereas in the last two cases income is equalized. Does this provide an improvement in expected utility? It certainly does, if Asaf and Sharif are risk-averse. You can see this in exactly the same way as we did in previous chapters. Mutual insurance gives rise to a mutual gain for Asaf and Sharif.

Second, observe that *at any date*, if the high–low scenario occurs, either Asaf will have to pay Sharif or Sharif will have to pay Asaf, but not both. Mutual insurance demands that ex post, after the outcome is revealed, one party make a unilateral transfer to the other. Moreover, the slate is wiped clean after this transfer. Next year, the *same* person may be asked to pay up again. The point is that a past transfer from one person to the other does not carry a historical burden for the recipient. This is what makes a scheme of pure insurance conceptually distinct from credit. A credit scheme ties past transfers to the expected pattern of transfers in the future. An insurance scheme in its purest form does not.

Third, observe that the power of mutual insurance depends greatly on the *degree of correlation* between the outputs of Asaf and Sharif.[5] Suppose, for instance, that high or low outputs are determined entirely by the state of the weather, and not by the possibility of localized pest damage. If Asaf and Sharif are in the same village, the weather will affect them equally, and cases (3) and (4) will *never* occur. Both farmers will always be in the same position

[5] See Appendix 2 at the end of the book for a detailed discussion of this concept.

and no mutual insurance is possible. Weakening this extreme example a bit, if both outputs are highly positively correlated, there will be situations in which the high–low scenario occurs, but this will happen with small probability. There *are* gains to be had from mutual insurance, but the magnitude of the gains is small ex ante.

Reversing this argument, imagine a situation in which the fortunes of Asaf and Sharif are perfectly negatively correlated. It may be hard to conceive of such a situation in reality, but treat it as a benchmark for a moment. In this case, the high–low scenario will *always* occur and the gains from insurance are huge. Thus it appears that insurance groups are most likely to be formed among those people whose fortunes are as negatively correlated as possible. There is some truth to this statement, but there are also reasons to avoid such extremes (even if they exist), which we shall come to later.

The general point is that for mutual insurance to work, there must be a certain absence of positive correlation between the fortunes of the participating economic agents. This is one reason why crop insurance is difficult: agriculture does present large correlations because of the weather. Having said this, though, it is only fair to recognize that there are situations that *are* idiosyncratic across various economic agents (illness, festivities, local crop damage), and these can benefit from insurance along the lines described in the preceding text.

Finally, suppose that there are not just two people, but many replicas of Asaf and Sharif. Assume again that their fortunes are all independent. Now ask yourself what the *average* income of this large string of people will be in any given year. This problem is identical to a coin-tossing problem: toss a coin several times, and each time write down 1,000 if heads comes up and 2,000 if tails comes up. The average of this string will be close to 1,500 for a long string of tosses.[6] For exactly the same reason, the average income will be $1,500 (identify each independent coin toss with a person), provided that the number of farmers is large.

Now suppose each of these farmers follows the rule of contributing $500 to a common pot if their income is $2,000 and claiming $500 from the common pot if their income is $1,000. If the number of farmers is large enough, then it follows from the argument in the previous paragraph that in any one year, this rule is *viable*, in the sense that the surplus or deficit in the pot will be negligibly small. Again it is true that in any *one* year, a farmer will be a net contributor to or a net recipient from the common pot, but note that each farmer succeeds, in this way, in smoothing his income *completely*. In the two-farmer case, this would have been possible only if the two incomes were perfectly and negatively correlated, so that Asaf gets $2,000 whenever Sharif gets $1,000, and vice versa. However, such a state of affairs is unlikely. With

[6] This is what we know of in probability theory as the (strong) *law of large numbers*.

a large number of farmers, independence of fortunes is enough to get full smoothing.

It is important to note that these insurance schemes are rarely constituted in some formal way. The box (to come) on reciprocal credit–insurance arrangements in Nigeria will give you some idea of how these transactions are carried out. Underlying the transactions is some social norm of reciprocity—one that pushes you to help out a community member in need. The transfer you make might be in cash or might be in kind (such as reciprocal exchanges of labor on farming plots). As we have seen on more than one occasion in this book, social norms cannot survive for long if the economic incentives that hold such norms begin to wither away, and so despite our acceptance of reciprocity and mutual insurance as a *social* institution, we are looking for the economic underpinnings that permit such an institution to survive.

15.2. The perfect insurance model

15.2.1. Theory

The example of the previous section allows us to formulate a model of perfect insurance. Suppose that a village is populated by a large number of farmers. For simplicity, let us assume that the farmers are all identical, though this is in no way necessary at a conceptual level. For each farmer, describe his income at each date by the equation

$$Y = A + \epsilon + \theta. \tag{15.1}$$

The income equation has three components. Think of the first (A) as each farmer's *average* income (this is a number such as \$1,500 in the example of the previous section). The second component (ϵ) is a random shock that may have the same distribution across farmers, but affects each one *independently*. For instance, the crop may be eaten by animals, or destroyed by bugs, or there may be inappropriate application of water or fertilizer. The shock may represent a shortfall of income due to a sudden spike in consumption needs; for instance, a part of income may be paid off for illness (in this case think of Y as net income after such necessary expenditures). It is a local or *idiosyncratic* source of uncertainty, which might happen to every farmer but happens in an independent fashion. Finally, the third component (θ) captures all the common or *aggregate* variation in the village. This is a source of uncertainty whose realization in any year affects *all* farmers in the same way. The best way to think of this is to imagine that these are common income variations stemming from the weather (the state of the monsoon, for instance).

Because we regard A as the average income, we should think of the random terms ϵ and θ as having mean value equal to zero.

Now, a little thought will convince you that if there are a large number of farmers, all the idiosyncratic variation embodied in ϵ can be insured away in exactly the same manner as in the previous section. Farmers simply pay their realized value of ϵ into a common fund. Because ϵ is sometimes positive and sometimes negative, this amounts to saying that some farmers make positive contributions, whereas others stand to receive payouts from the common fund. This way, all the idiosyncratic variation in ϵ can be ironed out. With the insurance fund in place, a typical farmer's *insured income* is now given by

$$\bar{Y} = A + \theta. \tag{15.2}$$

Compare (15.1) and (15.2). The income stream $\bar{Y}$ carries less risk than the original stream Y, because we have removed the idiosyncratic component. Average income is still A, of course. It follows that if farmers are risk-averse, they prefer this system of mutual insurance.

What about the aggregate shock θ? Can't we pool this in the same way? The answer is no, and the reason is that the value or realization of θ is the *same* for all the farmers in the village, so there is no insurance possible here. Neglecting self-insurance and credit for the moment, the village as a whole cannot mutually insure itself against a random event that affects everybody in the same way.

Hence it is not the case that under the perfect insurance scheme insured incomes don't fluctuate at all: they do, but the source of fluctuation is the aggregate uncertainty in the system that cannot be insured away with the use of a common pool of funds. This line of reasoning also suggests that one important determinant of mutual insurance is the relative significance of idiosyncratic to aggregate risk.

15.2.2. Testing the theory

The model that we described also yields a natural test for the existence of perfect insurance in a community such as an entire village. Controlling for movements in *aggregate* village consumption (which act as an indicator of aggregate uninsurable risk), fluctuations in individual household income, or, indeed, idiosyncratic shocks such as unemployment should *not* have any effect on the consumption of that household.

As we have seen on more than one occasion, regression analysis is a good tool for such a test.[7] Use the average consumption of the community and household income as independent variables. If possible, also include consumption-side shocks (such as unemployment) as independent variables. The dependent variable is household consumption. Roughly speaking, if

[7] See Appendix 2 at the end of the book for an introduction to regressions.

the theory of perfect insurance is right, the estimated coefficients on all the household-specific variables should be zero, and that on group average consumption should be unity. Individual consumption should track group consumption more closely than it tracks individual income.

Townsend [1993] explored this implication using three villages from the ICRISAT study, which should by now be familiar to you from repeated references in the last few chapters. Recall that the ICRISAT villages, located as they are in the semi-arid tropics of southern India, all display a high degree of risk. These risks stem primarily from crop cultivation, but spill over to labor markets, so that a landless laborer who earns wages faces a substantial amount of risk as well. The three villages studied were Aurapalle, Shirapur, and Kanzara. Household consumption data are available annually for a sample of forty households in each of these three villages for the period 1975–84. There is substantial variation in household income, and if this variation is taken as a measure of household risk, then idiosyncratic components are very much in evidence. The best way to see this is to compare the *differences* between household income and the average income of the sample, for different households. Subtracting the sample average is a bit like removing the aggregate risk. Then, if we interpret what's left as idiosyncratic risk, it is certainly the case that there is a lot of it: this residual appears to fluctuate independently across households (see Townsend [1993, Figure 1]). This means that the three ICRISAT villages present a potentially fertile case for testing the perfect insurance hypothesis.[8]

A major problem arises, however, in the construction of an appropriate regression equation. We have household-level consumption, and measures of household income. What do we use for average consumption? The answer to this question depends on whom we suspect to be the relevant group that is engaged in mutual insurance. One possibility is that the *entire* village is engaged in mutual insurance, as one group so to speak. In that case we include some measure of village average consumption. However, all we have are data from the sample!

One option is to use the sample consumption average as a proxy for the village average and hope that by the law of large numbers this gives us a good estimate of the village average. However, this runs us into another sort of problem: even if we perform the regressions household by household, if the sample is large, the *average* estimated coefficient on group consumption must be unity! This is not a sign that the perfect insurance model passes with flying colors, but just that individual consumptions in the sample must sum to the sample average (Deaton [1990]).

[8] At the same time, be aware that there is substantial measurement error in estimating household incomes. This sort of error invariably biases estimated coefficients on household income variables toward zero. Some of this error can be avoided by averaging the income data over households with the same crop/plot technologies, but this remains a serious issue. Measurement error can also be allowed for (to some extent) by statistical correction. See Townsend [1993, Section 6] for more details.

One way to get around this problem is to run the regression household by household (using the sample average exclusive of the household in question) and see whether *each* of the estimated coefficients on group average (and not just the average of these coefficients) is close to a value of 1. The results of this exercise are mixed.[9] Depending on the length of the time period and the choice of consumption categories in the regressions (see footnote 9), the coefficients are sometimes, but not always, close to unity. A formal statistical test of the hypothesis that the coefficients are indeed 1, is not rejected using the full set of data, but the power of this test deteriorates with other variants of the regression.[10]

The observation that the coefficients on household income are indeed close to zero is more robust and suggests that variations in household income do not significantly affect household consumption. This outcome suggests that quite a bit of smoothing is taking place, but whether it is due to insurance per se or some other form of smoothing (such as self-insurance or credit) is difficult to say.

Another way to avoid the problem of missing aggregate data on consumption is to consider *differences* between household consumption and the sample average of households. This is tantamount to subtracting group averages from both sides of the regression, equation, so that the right-hand side now contains only individual terms, such as household income of various categories, and dummy variables for sickness or unemployment. Now, to test whether group consumption moves individual consumption more than individual income does is simply a matter of whether or not the coefficients on the individual variables are zero. In contrast to the earlier regressions, these coefficients are now significantly different from zero, but their magnitudes are small.[11] For instance, the effect of household income on individual consumption seems to be no more than 14%, which is pretty small. This is like saying that a one dollar increase in contemporary household income increases consumption by no more than fourteen cents. Moreover, events such as sickness (leading to loss of work) are insignificant in the determination of household consumption.

The preceding findings suggest that although we cannot be sure that mutual insurance is at work here, there is substantial smoothing accomplished by ICRISAT households through some means or other. Informal credit markets probably account for a large chunk of this. There is separate evidence

[9] A particularly nasty problem is that data collection between the years 1982–84 dropped some consumption categories. To be sure, this has the spurious effect of making individual consumption move with household consumption. To sidestep this problem, we can use a restricted set of consumption categories throughout (such as grain consumption) or perform the regression by leaving out these last two years. The results concerning the comovement of household and group consumptions are predictably weaker.

[10] For a discussion of hypothesis testing, see Appendix 2 at the end of the book.

[11] See Appendix 2 at the end of the book for a detailed discussion of the difference between *significance* and *magnitude*.

that ICRISAT households smooth fluctuating income through credit transactions (see Cain [1981] and Walker et al. [1983]). Other mechanisms might include grain inventories (Paxson and Chaudhuri [1994]), holdings of cash (Lim and Townsend [1994]), and even the sale and purchase of livestock (Rosenzweig and Wolpin [1993] and the box on self-insurance and bullocks).

Note, however, that among the growing number of empirical studies conducted on this question, the ICRISAT villages do relatively well in exhibiting a great deal of consumption smoothing. Townsend [1995] reports on similar regressions carried out on Thai data. The results seem to convincingly reject the presumption of perfect risk sharing, although in other studies on Thailand, a significant degree of smoothing cannot be ruled out (Paxson [1992] used data on Thai rice farmers to make this point). Similarly, Deaton's [1994] study of the Côte d'Ivoire suggests that although there is comovement of household consumption, full insurance can be convincingly rejected. Household income does significantly influence household consumption. Townsend [1995] provides a discussion of these and other studies.

Even among the ICRISAT villages, the ability to smooth may vary significantly across households. Using a somewhat different test for insurance than the one described here, Morduch [1995] found substantial evidence of consumption smoothing among relatively better-off farmers, but not so for small farmers and landless laborers (Townsend's work contains similar observations as well). Morduch also made the point that insurance might look good simply because households try to smooth their *income* streams so that the remaining fluctuations can be absorbed by the available consumption smoothing mechanisms. For instance, some households may forego the cultivation of a crop with high expected return simply because the yield is more risky and cannot be smoothed through insurance. In this case, the household chooses a "safer" crop and its income fluctuations are therefore smaller and consequently easier to smooth.[12] The point is that there is a hidden cost (which is the expected value sacrified by choosing a safer crop). Thus, Morduch [1995, p. 108] concludes, "for less well off households in the ICRISAT villages, production choices will be made with an eye to reducing the likelihood that the shocks will happen in the first place."

15.3. Limits to insurance: Information

The general failure of the perfect insurance model to fit the facts suggests that there are limits to the ability of households to insure one another, even if the income shocks are idiosyncratic. Where do these limits come from? The

[12] The same sort of observation is true for the *way* in which a given crop is produced. Rosenzweig and Binswanger [1993] investigated this for the ICRISAT villages; see also the earlier study by Bliss and Stern [1982] on fertilizer choice in the northern Indian village of Palanpur.

answer is based on the same principles that informed our study of credit markets: limited information and limited enforceability.

15.3.1. Limited information about the final outcome

Think of two sorts of informational problems. One is that a person may ask for or expect transfers from his community while providing them with deliberately wrong or misleading information regarding his economic state: he might lie about the size of his harvest, he might lie about illness in the family, or, at the very least, he might be in a *position* to lie, in the sense that his community members do not have information to verify that the claimed occurrences indeed occurred.

Now this is certainly a possibility if we extravagantly expand the notion of what a community means. It is unlikely (even if speedy electronic transfers of funds were possible between villages) that a farmer in one extreme of the state of Bahia would be engaged in insurance with another unknown farmer in the other extreme of that state.[13] Even if such insurance were mutually beneficial, it is very difficult to verify the underlying circumstances on which that insurance must be based. So it's certainly true that this sort of informational barrier kicks in and precludes insurance over very large anonymous groups or spatial distance, even if such insurance could, in principle, be profitable for all concerned.

However, does this sort of informational problem occur within the village? It is hard to say, and the answer must surely vary across villages. One of the great strengths of traditional society is that they are highly endowed with *social capital*. Much of social capital is simply information flow. *A* knows if *B*'s harvest has been bad or if *B*'s father-in-law really has food poisoning. Social capital provides a fund of socially available information that permits a community to interact in ways that would not be possible in an anonymous society. Insurance is one such way, and there are others. We already discussed social capital in the context of credit in Section 14.3.7 and earlier in the context of permanent labor markets in Chapter 13. In villages where mobility is on the increase, this sort of knowledge begins to die away, reducing the scope of insurance networks. Likewise, a village that is segmented into subgroups (say, along caste lines or perhaps even along economic hierarchies such as landowners and the landless) may be unable to develop insurance schemes that cut across these lines of segmentation, simply because the flow of information is restricted. All of these theories suggest the need to extend empirical work to define more precisely what we mean by the community, and to the extent that our notion of community fails to capture these finer

[13] This does not prevent organized financial institutions from indirectly linking them, but that is another matter.

points (mainly because of lack of data), we might fail to capture the existence of insurance "subnetworks" simply because we are looking for one giant network, which doesn't exist.

15.3.2. Limited information about what led to the outcome

The second kind of informational problem is more subtle and, as we will see, does not altogether rule out the possibility of insurance. This is the problem of *moral hazard*—the possibility that the occurrence of some insurable event can be influenced by the unobservable actions of the individual. We've already encountered moral hazard in our discussion of land markets (Chapter 12) and in our discussion of credit rationing (Chapter 14). Clearly, moral hazard matters for insurance as well.

Consider mutual insurance of harvests. Everyone knows that there are idiosyncratic shocks that can affect the state of the harvest on a particular plot. Water, fertilizer, or pesticides may have been badly applied (inadvertently), or there may be damage to the crop caused by events outside the farmer's control, or (and this is the source of the problem), the crop may be bad because the farmer *deliberately* skimped on the use of these inputs. The size of the harvest may be visible for all to see (so that the first informational problem is irrelevant), but *why* the harvest is what it is requires information of a different kind altogether. The problem, of course, is that in the presence of full insurance, the incentive to deliberately underapply inputs is higher.

To see this in its simplest form, consider a community of a large number of identical farmers. Assume that all uncertainty is idiosyncratic. Just as in the example of Asaf and Sharif, we suppose that there are two possible outputs, which we denote by L (for "low") and H (for "high"). The twist in the story is that for each of the farmers there are *two* possible probabilities of producing the high output: call them p and q, and suppose that $p > q$. The idea is that with greater care and better application of inputs, each farmer can produce H with probability p, but if he underapplies inputs, the probability falls to q. Denote the extra cost of diligent effort and input use by C (this is expressed in utility units). Finally, let u be the utility function that describes how net outputs are translated into units of utility for each of the farmers.

Begin by analyzing the farmer's problem without insurance. If he diligently farms the land, his expected utility is $pu(H) + (1-p)u(L)$, but he incurs the extra cost C up front, so his *net* expected utility is given by the expression

$$pu(H) + (1-p)u(L) - C.$$

On the other hand if he does not apply effort, the probability of the high output falls to q, although he saves on the extra cost, so that his expected

utility is now

$$qu(H) + (1-q)u(L).$$

Assume that in isolation, each farmer would like to put in the extra effort; that the probability gain of the high output makes it worthwhile to do so. This amounts to assuming that

$$pu(H) + (1-p)u(L) - C > qu(H) + (1-q)u(L). \tag{15.3}$$

Now with insurance we can do even better. We already saw in several examples that if the farmer is risk-averse, then the utility of the expected output exceeds the expected utility of output,

$$u(pH + (1-p)L) > pu(H) + (1-p)u(L),$$

so there are clearly gains to be had from insurance. Moreover, with a large number of farmers, it is possible to set up an insurance scheme that (approximately) guarantees this for every farmer. To see how, suppose that every farmer with high output pays the sum $(1-p)(H-L)$ into a common pool, while every farmer who produces the low output receives a payout of $p(H-L)$. Do the necessary simple algebra to verify that with this scheme, each farmer gets the average value of output, $pH + (1-p)L$, *for sure*. Now all we have to do is to make sure that the scheme is feasible: that total payouts balance total receipts. To see this, note that with a large number of farmers, approximately a fraction p of them get the high output, so that the total inflow into the pool is (approximately) $p(1-p)(H-L)$. At the same time, approximately $(1-p)$ of the farmers produce the low output, so that the total outflow from the pool is $(1-p)p(H-L)$. Inflows and outflows match, so we are in business.

However, there is a flaw in this argument. We *are* in business if we can somehow guarantee that each farmer continues to put in the high level of effort. But it is not possible to do so under this insurance scheme if the extra effort cannot be verified by outside parties. If each farmer receives $pH + (1-p)L$ *regardless* of whether his output is high or low, then why incur the extra effort to raise the probability that output will be high? Thus, each farmer has an incentive to slack off under the perfect insurance scheme, and, of course, if all farmers think this way, the scheme cannot survive.

Notice that the problem is the same as that encountered in the principal–agent model. There is a trade-off between the provision of insurance and the provision of incentives. If the community wants to maintain effort at the higher level, it will have to offer incomplete insurance so that individual incentives are not destroyed. How much insurance can it safely offer?

To solve this problem, suppose that X is a farmer's net consumption when output is high and Y is his net consumption when output is low, in any insurance scheme. The requirement that the scheme be approximately viable is equivalent to the equation

$$pX + (1-p)Y = pH + (1-p)L. \tag{15.4}$$

If X and Y are too close together, people will slack off. To see how close they can be, take a second look at equation (15.3), and substitute X for H and Y for L. Note that high effort will be chosen as long as

$$pu(X) + (1-p)u(Y) - C \geq qu(X) + (1-q)u(Y)$$

or equivalently, rearranging these terms, as long as

$$(p-q)[u(X) - u(Y)] \geq C.$$

Thus the *closest* that X and Y can get without upsetting the desired incentives is given by the equation

$$(p-q)[u(X) - u(Y)] = C. \tag{15.5}$$

Equations (15.4) and (15.5) are two equations in two unknowns. Use them to solve for X and Y. This is the *second-best insurance scheme*—the best one can do when incentives pose a constraint.

Note that in any solution, it must be the case that $X > Y$ [simply examine (15.5)]. This means that under the second-best scheme, *individual consumption must move with individual income*. Such a finding is consistent with the data, although to empirically test this model against others requires finer analysis.[14]

To complete this model, consider a numerical example just to see how much insurance is restricted by these considerations. Let $H = 2{,}000$ and $L = 1{,}000$ just as in the Asaf–Sharif example. Let $p = 3/4$ and $q = 1/4$. Suppose that the utility function is given by $u(x) = \sqrt{x}$. Finally, let the additional cost of diligence assume the value 2 in utility units. Check that the condition (15.3) is comfortably satisfied, so that in isolation each farmer will indeed put in the higher level of effort. Now we solve for the second-best insurance scheme, described by the pair of numbers X and Y that solve (15.4) and (15.5). These two equations reduce to

$$\tfrac{3}{4}X + \tfrac{1}{4}Y = 1{,}750$$

[14] Ligon [1993] pursues this line of research.

and

$$\tfrac{1}{2}[\sqrt{X} - \sqrt{Y}] = 2.$$

Solving, we see that X is approximately \$1,832 and Y is approximately \$1,505. Thus farmers with good harvests pay out around \$170 each and farmers with bad harvests receive around \$500 each. More insurance would certainly be desirable but cannot be achieved, because the constraints on incentives would be violated.

Informational constraints pose a real problem to effective insurance. Groups with better access to the information of their members are generally in a better position to provide mutual insurance. This is also true of groups that are composed of members who feel some altruism toward one another. In this case, the incentive constraints are somewhat loosened, because each member internalizes (perhaps partially) the undesirable implications of his deviation for other group members. Hence, it will not be surprising to see that insurance arrangements between members of extended families or of the same kin group[15] are relatively common.

The flip side of this argument is that diversification possibilities within the extended family may be relatively limited. This will be the case if all members participate in a similar activity and are therefore exposed to the same set of aggregate shocks. We expect, then, that families will attempt to diversify their activities or spatial locations to lower the correlation between the incomes of family members. Rosenzweig and Stark (1989) found a tendency for marriage and kinship ties to be established across families from different geographical regions, perhaps to maximize insurance possibilities while retaining (or creating) the altruistic ties that are needed to counterbalance the absence of information. In the same spirit, Paulsen [1995] studied remittances from migrants in cities to their families in the Thai countryside.

15.4. Limits to insurance: Enforcement

We now turn to a different limit on insurance, that has to do with the *enforcement* of insurance arrangements.[16] Recall that mutual insurance schemes are rarely contracts that are written down in black and white and routinely backed or enforced by the law. Rather they are informal arrangements, set in the context of a social norm that encourages reciprocity and punishes deviations through social sanctions such as ostracization or public rebuke. This, too, is a form of *social* capital.

[15] For studies of reciprocal arrangements within kin groups, see La Ferrara [1997].

[16] For a more formal treatment of some of the issues in this section, see Coate and Ravallion [1993], Fafchamps [1996], Kletzer and Wright [1995], and Thomas and Worrall [1994].

Whether a village's endowment of social capital is enough to permit a system of mutual insurance to function is a question that can only be settled through detailed empirical work, but it is useful, anyway, to discuss the various issues that arise with the help of a simple model.

15.4.1. Enforcement-based limits to perfect insurance

Suppose, once again, that there are a large number of farmers. The simplest case is one in which all uncertainty is idiosyncratic, so let's begin with it. Assume that there are two outputs, H and L, and that the high output is produced with probability p. Because we have already discussed the moral hazard problem, we take the probability of the high output to be fixed in what follows (equivalently, effort levels are perfectly monitored).

Consider perfect insurance. We already saw what this gives each person. With a large number of farmers, each farmer receives the constant amount $pH + (1-p)L$ regardless of what he produces. Let this "middle amount" be denoted by M. We already discussed the possibility that such a scheme can be undermined because of moral hazard, so this is not an issue we consider here.

Now picture in your mind a situation where this scheme is in place and transfers are made according to it year after year. Consider a particular year in which a particular farmer produces a high output. In that year he *could* receive a utility of $u(H)$ if he pocketed this output in its entirety, but the insurance scheme dictates that he give up some amount to the common pool, so that his utility under the scheme is really $u(M)$, which is obviously smaller. We can use this information to construct a measure of the *gain* enjoyed by a high-output individual if he decides to ignore an existing insurance arrangement and simply consume the output himself. Denoting the gain by G, we have

$$G = u(H) - u(M).$$

This is a measure of the incentive *not* to conform to the arrangement. Of course, this is only one side of the story. If the farmer deviates in this fashion, he will presumably be excluded from future access to insurance.[17] There is also the question of social sanctions, as we already discussed. All of this adds up to a future *loss* if he fails to conform today.

Just as we noticed for credit markets (Chapter 14) and (permanent) labor markets (Chapter 13), the size of the loss from future exclusion depends on

[17] Informational considerations are once again relevant here. Depending on how easily the news of his "crime" spreads through the community, this individual may find it more or less easy to become part of another group. We already discussed such issues in the context of credit and labor markets. They are probably less relevant in the case of insurance, involving as it does a larger group of people at a time.

the farmer's farsightedness. In each future date he loses the amount $u(M) - [pu(H) + (1-p)u(L)]$, which is just the difference between his insured utility and his stand-alone expected utility. If he has a mental horizon of N periods, the total loss is just this expression multiplied by N, the number of dates for which he thinks ahead. To this add a loss S, which is a measure of the utility damage imposed by social sanctions (such as a public scolding in the main village square). The total loss (L) from a deviation is thus given as

$$L = N\{u(M) - [pu(H) + (1-p)u(L)]\} + S.$$

Now put together the gains and losses to check if perfect insurance is indeed viable. For viability, we must have $L \geq G$, and this gives rise to the fundamental *enforcement constraint*

$$N\{u(M) - [pu(H) + (1-p)u(L)]\} + S \geq u(H) - u(M). \tag{15.6}$$

This looks like an ugly equation (and it is), but it gives us some insight into the nature of the enforcement constraint.

First, look at the role of S, which is the punitive power of social sanctions. Clearly, the larger is S, the more likely it is that (15.6) holds. Second, the enforcement constraint is also more likely to hold if people value future interactions to a greater degree. This notion is reflected in a larger N. In part, a larger value of N may be interpreted as greater foresight or greater patience on the part of the farmer, but it can also be interpreted as a measure of greater faith in the durability of the scheme itself. If mutual insurance is expected to continue in the indefinite future, the loss from a single deviation is perceived to be large, and so the deviation is less likely to occur. This also brings out a "bootstrap" characteristic of such schemes (which appears in other situations as well, such as cooperative credit). The perceived instability of the insurance scheme may be self-fulfilling: N is driven down, (15.6) fails, and the scheme falls apart simply because it is expected to. Social institutions are very much like the outcomes of coordination games, and multiple equilibria are surely possible.

Third, is it the case that the greater the scope of insurance the more likely it is that the scheme will survive? This sounds pretty obvious but actually (15.6) tells us that it isn't generally true. Insurance is more beneficial when stand-alone outputs are more dispersed; in this example, when there are large differences between L and H. To be sure, this raises the perceived losses from not participating in the scheme, but it also means that more has to be contributed when the output is good [recall that the contribution is $(1-p)(H-L)$, which increases with dispersion]. This means that there is more to be gained from a deviation. The net effect can go either way.

15.4.2. Enforcement and imperfect insurance

What if equation (15.6) fails to hold? This means that *perfect* insurance is unattainable. We are then in a situation analogous to that discussed in the case of informational failure, where, nevertheless, some more limited form of insurance was possible. We can follow in the steps of that analysis to see what happens here. Once again, denote by X the consumption of a high-output farmer and by Y the consumption of a low-output farmer. For the scheme to be viable, the "insurance budget" must balance; this is just equation (15.4) which we reproduce here:

$$pX + (1-p)Y = pH + (1-p)L = M. \tag{15.7}$$

Next, the enforcement constraint must be satisfied. This amounts to the statement that the gain from a deviation must not outweigh the future losses. The one-time gain is $u(H) - u(X)$. The loss is

$$N\{[pu(X) + (1-p)u(Y)] - [pu(H) + (1-p)u(L)]\} + S,$$

so the enforcement constraint may be rewritten as

$$N\{[pu(X) + (1-p)u(Y)] - [pu(H) + (1-p)u(L)]\} + S \geq u(H) - u(X). \tag{15.8}$$

It will help to rewrite this in a way so that the Xs and Ys are all on one side. Transposing terms, we obtain the equivalent condition

$$u(X) + N\{pu(X) + (1-p)u(Y)\} \geq u(H) + N\{pu(H) + (1-p)u(L)\} - S. \tag{15.9}$$

We are after the second-best insurance scheme. We already know that everybody wants as much consumption smoothing as possible (subject to the enforcement constraint), so we are looking for the "least dispersed" values of X and Y that are compatible with both (15.7) and (15.9). There is an easy way to do this graphically. Figure 15.1 plots how the left-hand side of the enforcement constraint (15.9) changes as we alter X, starting from the perfect insurance level [which is, of course, $pH + (1-p)L$] and moving up to the no-insurance level H. Implicitly, the value of Y is altered, so that the viability condition (15.7) holds throughout this exercise. Thus Y *falls* from the perfect insurance level [also $pH + (1-p)L$] to the no-insurance level L (but this change is not explicitly carried in the diagram).

What happens to the left-hand side of (15.9)? *At* the perfect insurance point $X = pH + (1-p)L$, the value of this expression must rise. The reason

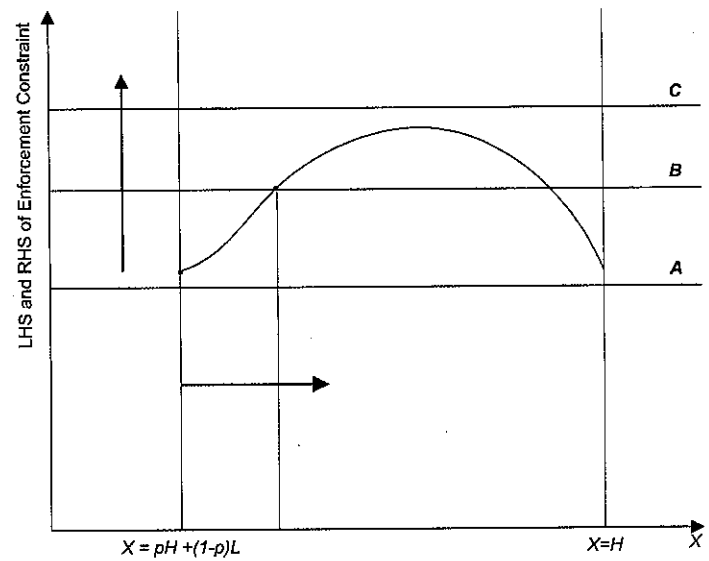

Figure 15.1. The enforcement constraint and imperfect insurance.

is that the term $pu(X) + (1-p)u(Y)$ is approximately unchanging for small variations in X, at the perfect insurance point (because it attains a maximum there).[18] However, $u(X)$ increases with X, so the combined expression $u(X) + N\{pu(X) + (1-p)u(Y)\}$ must rise with X.

What does this mean in words? It means that if we move away from the first-best insurance level, enforcement becomes easier. This is because the deviation gain falls, while the losses from future denial of insurance do not fall as fast, at least initially.

After a while, this trend reverses itself. The fall in $pu(X) + (1-p)u(Y)$ starts to become significant for higher values of X, and ultimately this dominates the gain in $u(X)$, so the combined expression reaches a maximum and then turns down. It is possible that by this time, X has already reached the no-insurance point H so that, effectively, there is no turning down of the combined expression. Which particular case occurs depends on the parameters of the problem and is of no great significance to us, so we also put in a declining segment in Figure 15.1 for completeness.

The horizontal line in Figure 15.1 represents the right-hand side of the enforcement constraint (15.9). It is flat because it is independent of the value of X. Note that its height depends on various parameters, among them the

[18] This and the observations to follow can be easily verified by taking derivatives of the relevant expression with respect to X, accounting at the same time for the comovement of Y. I omit the details.

value of social sanctions. The lower the line, the greater the range of insurance schemes that satisfy the enforcement constraint.

We draw three instances of this horizontal segment. For the values represented by line A, perfect insurance is enforcable and there is no need to look at other schemes. Now imagine that social norms in this economy begin to wither away, so that the value of social sanctions declines. This has the effect of shifting the horizontal line *upward* to the line marked B. At this level, perfect insurance falls apart. The second-best insurance scheme is achieved by a higher value of X (and consequently a lower value of Y). This means that individual consumption has to move with individual income. The lower the value of social sanctions, the more will be this variability. Finally, there comes a point (shown by the line marked C) at which *no* insurance is possible at all.

Thus the lack of enforceability limits insurance and creates second-best insurance structures in which people must face, to some extent, the variation in their own individual fates. This model may be modified in many different ways, to yield different insights.

Self-insurance and mutual insurance

First, we can interact self-insurance or credit with mutual insurance. Following a deviation from mutual insurance, it may be possible to self-insure to some extent. This means that punishments are effectively lower, and the tendency for mutual insurance schemes to fall apart is thereby accentuated. The greater the possibilities for self-insurance, the more likely it is that mutual insurance schemes will begin to crumble. At one level this is only natural: as alternatives appear, informal substitutes can die out without much difficulty. But it is important to appreciate that the process exhibits a built-in inefficiency. The value of a deviation from mutual insurance is the *current* (high-output) return plus the future return from self-insurance. If this value is larger than the returns from mutual insurance and if there are no additional social sanctions, the mutual scheme will collapse. However, note that *at* the collapse point, the mutual scheme generates larger value than self-insurance does, provided that the power of social sanctions is small.[19] In this sense the transition is inefficient.

Correlated shocks

We also can extend the model to include aggregate (correlated) shocks to output (see Coate and Ravallion [1993]). It is easy enough to write down

[19] The current payoff from a deviation, if a farmer has produced the high output, is always larger than the return that any insurance scheme can provide. It follows that if S is small, the mutual insurance scheme fails to be enforceable at a point when it still does better than self-insurance. You may want to formally prove this by first adapting the enforcement constraint (15.6) to include a weaker self-insurance scheme following the deviation. Put $S = 0$. Now examine the point at which mutual insurance fails the enforcement constraint.

an enforcement constraint for this problem. Such a constraint has an inte esting feature. If the entire economy is in an aggregate depressed state, th farmers who do *relatively* well in that state will have a greater incentive t deviate from any given mutual insurance scheme. This follows from dimir ishing marginal utility: when total output is lower, the value of *not* makin some given contribution to an insurance fund must rise. At the same tim future expectations of gains from insurance are unchanged (if the shocks a independently distributed over time).

Consequently, it becomes more difficult to satisfy the enforcement co straint when times are bad for the economy as a whole, and so insuranc weakens in such times. However, we must be careful: contrast this ph nomenon with one in which there is a permanent improvement in the l of one (or some) members of the group. Faced with a permanent improv ment, they too will not contribute to mutual insurance. Thus an aggrega state of economic depression or a selective state of well-being can both co tribute to the destruction of enforceable insurance opportunities.

The blurring of credit and insurance

Our third (and last) extension is much more subtle than the other two an has profound consequences for the conceptual distinction between credit an insurance. Recall that the distinguishing feature of insurance is a certain lac of regard for history. It is possible (though perhaps unlikely) that Asaf ma have a bad output three years in a row, and Sharif may have enjoyed a hig output for those three years. Our insurance scheme dictates that Sharif mu make contributions to Asaf for those three years. Moreover, once the thre years are over, the slate is wiped clean again. Thus pure insurance has a un lateral air to it: transfers do not carry an obligation to repay. Indeed, strang though it may sound, this is a necessary property of the perfect insuranc scheme in the model that we described. It's not that such schemes are e tirely missing in the world we live in: for instance, medical insurance has th property—but it is harder to imagine informal schemes with this property.

It turns out that the enforcement constraint throws some interesting lig on this issue. It is true, as we already mentioned, that the first-best schem has a "history-independence" property. However, what if enforcement co straints prevent implementation of the first-best scheme? This is precisely t situation that we have considered so far, yet in formulating our second-be scheme, we stuck firmly to pure insurance and did not allow past transfe to influence the future. It turns out that if we drop this requirement, we ca do even better (Fafchamps [1996]; Kletzer and Wright [1995]; Thomas an Worrall [1994]).

To see this better scheme, return to the formulation of the second-be arrangement pursued in equations (15.7)–(15.9) and Figure 15.1. Recall t

basic idea in this construction: high-output individuals may be reluctant to contribute to the insurance scheme because they stand to gain by keeping their entire output. They balance these gains against the losses of not being able to participate in the scheme in the future (there are also social sanctions, but we neglect these here as they are not needed to make the main point). The way we previously solved for X—the consumption of the high-output person—was by balancing these two conflicting forces in a way that was consonant with the maximum possible degree of smoothing. In doing so, we neglected another source of balance; that is, *the scheme could be designed to increase, for a high-output person, the future returns from participation in the insurance scheme.* This raises the losses from future nonparticipation and thus loosens the enforcement constraint and permits a greater degree of smoothing today.

This means that the production of high output today raises a person's share of the pie in tomorrow's scheme and likewise lowers the portion for a low-output person who receives a handout from the scheme today. However, such an adjustment doesn't affect the low-output farmer's enforcement constraint: he is *receiving* money from the system, so there is no constraint to worry about in this case.[20]

Of course, this state of affairs is not permanent: today's high-output person may be tomorrow's low-output producer, and then he will slip down a notch as far as the subsequent share in the insurance fund is concerned. Likewise, today's low-output producer may recover his "status" by producing high in the future.

The algebra of this formulation is complicated and we will not pursue the details here, but you should be able to see how the basic idea works, and—more importantly—its implications. By linking an individual's payouts in the past to current claims in the insurance scheme, the scheme ceases to be one of pure insurance and acquires some of the characteristics of credit. Your current entitlement is higher if you have paid out more in the past, which sounds, to some extent, like you are entitled to a "repayment" of some of the money that you contributed yesterday. However, this is an entitlement only *on average*. It may be that your compatriot does badly again, and then the insurance aspect of the scheme will do battle with the credit aspect: you may be "required" to forgive or roll over some of the past "debt" and even advance further "loans," so it's more complicated than pure credit. Still more interesting is the interpretation of repayment from your compatriot depending on *your* economic state today: he will be required to pay you more if your current economic state is bad, relative to when it is good.

This seemingly esoteric mixed-credit insurance scheme brings us to the frontiers of current research in this area. Is it all esoterica in reality? The

[20] This is true provided we are careful not to reduce his participation utility from tomorrow onward *below* the level he can attain on his own.

answer, surprisingly enough, is no: as empirical investigations unfold, we see that a great deal of informal financial interaction in developing countries is characterized by norms of both reciprocity and obligation. The former demands repeated support if circumstances require it, so it has the flavor of insurance. The latter demands some accounting of the past, so it has the flavor of credit. We end this chapter with a box that describes a pioneering empirical study in this area.

Credit as Insurance: A Case Study from Northern Nigeria

How can markets and economic institutions adapt to missing information and problems of enforcement? We have seen how these constraints can limit reciprocal transactions and force a blend of credit and insurance to develop. A study of rural credit markets in northern Nigeria by Udry [1993, 1994] provides a rich empirical illustration of institutional adaptation.

Nigeria is predominantly Islamic: over 90% of the population is Muslim. Islamic law strongly condemns the practice of charging interest for loans. Therefore, usury has always been taboo in traditional Islamic societies. Fixed repayment periods are also prohibited: "And if the debtor is in difficulty, then (there should be) postponement to a time of ease" says the *Koran*. (2: 280; quoted in Udry [1993]). At the same time, Nigerian agriculture is subject to seasonal fluctuations and random shocks (droughts, crop disease, etc.). Thus income flows are naturally highly uncertain and irregular. The need to borrow and lend in times of distress is therefore paramount. Given the stricture against charging interest, which should strongly dampen the incentive to lend, how do economic institutions deal with the problem of consumption smoothing? Is there any significant credit market alive in rural Nigeria? Udry's study, based on data collected from four randomly selected villages near the city of Zaria in the Kaduna state of northern Nigeria, shows that the answer is very much in the affirmative and that credit implicitly takes the form of interhousehold insurance against income shortfalls.

Of the 400 sample households, only 10% reported *not* having engaged in any borrowing or lending activity over the period of the survey. Indeed, 50% of households *both* borrowed and lent at various points during this interval. A noticeable feature of these loans is that they were almost always restricted within tightly knit communities, whose members knew about each other's conditions, needs, and activities. Ninety-seven percent (by value) of the 808 informal-sector loans recorded were either between residents of the same village or between relatives. Almost all such loans were remarkably informal in character: there was no written contract or witness, nor was a rate of interest or repayment date specified, implicitly or explicitly. Loans seem to be advanced purely on the basis of trust.

What is interesting about the Nigerian credit market is that the amount of money paid back depends not only on the borrower's financial situation at the time of repayment, *but also on that of the lender*. Udry's econometric analysis confirmed that the importance of the latter factor in determining the implicit rate of interest is statisti-

cally significant. This led him to conclude that the loans were not, in fact, equivalent to "equity shares" in the borrower's farm or business, but mixed devices with components of insurance and mutual risk pooling embedded in them.

It is not surprising, therefore, that the implicit realized rate of interest, calculated from loaned and repaid amounts, showed large variations; however, the median nominal monthly rate of return was about zero. Given that the monthly rate of inflation was a steep 3.7%, the average loan clearly earned a significantly negative real rate of return. This further strengthens the interpretation that the loans were really forms of mutual insurance, based on need and reciprocity. However, they were not unilateral transfers either, as a pure insurance would dictate. A transfer today carries *some* obligation to repay tomorrow.

Theory suggests that a mixed credit–insurance contract is symptomatic of enforcement problems. A household that receives a loan may decline to repay later or may repay only a meager amount. Likewise, households experiencing relative prosperity may refuse to help out their neighbors who are in need. How does the community keep default or delinquency at bay? It seems that discipline is maintained by the pressure of community opinion and the threat of censure or ostracization—instruments that can be more powerful than mere pecuniary penalties. Note, moreover, that an impersonal market, by the very fact of being *impersonal*, does not have these instruments at its disposal. What an impersonal market does share with this informal institution is another source of punishment: temporary or permanent exclusion from future participation in reciprocal transactions.

In the few cases of dispute and perceived default, a complaint is made to a community leader or village head by the lender. If, after review of the situation, the leader holds the borrower to be guilty, the punishment imposed is usually a verbal admonition and, in extreme cases, a threat to make the matter public (leading to shame and possible sanctions from the community). The latter threat is considered so severe and damaging that in his survey, Udry found not a single instance of it.

Other issues of moral hazard are likely to be secondary in the Nigerian context because of the ease of information flow. Because lending is restricted within closely knit communities, borrowers usually find it impossible to misrepresent their true financial situation or to feign crop damage, illness, or any other misfortune.

A glance at default figures in the sample is revealing. Of the 808 loans, only 47 (about 5%) were in default as perceived by the lenders themselves. Astonishingly, only three of these cases pertained to borrowers who experienced negative shocks to their income after the loan was taken. This must mean that borrowers who do not repay because of bad shocks are generally not *perceived* as having defaulted. This goes to show that it is part of the norm that the repayment burden is reduced or excused for those borrowers who run into financial trouble: proof again of the insurance aspect of these transactions.

Let's end with some observations on intervillage credit transactions. Note that a shock to income can be of two types: an *idiosyncratic* shock, which affects a single household (e.g., an illness in the family, flooding or insect damage to the family's plot, etc.), and an *aggregate* shock, which can affect an entire village or region (e.g., rainfall shortage, the flooding of a river, etc.). Because most borrowing and lending is restricted to small communities and villages, it seems that the rural credit market should be efficient in pooling risk and providing insurance against the first kind of shock, but not against the second. Udry estimated that in his sample, about half (58%) of the variation in farm yields was caused by the latter kind of shocks, so this limitation must be serious.

However, although there is practically no direct outside lending (the informational disadvantage of outsiders and the inability to charge fixed, nonstate contingent interest rates in violation of Islamic law are surely the reasons), flows of funds between villages are not entirely absent. In northern Nigeria, the Hausa tradition of long-distance trading has given rise to a network of merchants. With widespread outside connections and intimate personal knowledge of people in their own resident villages, these merchants act as pipelines for outside credit to enter the village, while still providing loans on terms that do not violate the accepted standards of the community. Intervillage insurance, though not as prevalent as intravillage arrangements, does not appear to be altogether absent.

15.5. Summary

This chapter studied *insurance*: the provision of funds to smooth out unexpected variations in income or in consumption needs. We began with an example to show how insurance can be profitable, and used this example to distinguish between *self-insurance*, which is the smoothing of consumption using one's own assets, and *mutual insurance*, which requires interaction across a number of individuals. We observed that insurance possibilities are maximal when the fortunes of individuals are *negatively correlated* (I am doing well, on average, when you are not, and vice versa). Although negative correlation may be difficult to come by in practice, a weaker condition—the *independence* of incomes—also yields considerable scope for insurance. Insurance possibilities do diminish, however, as the fortunes of individuals become more and more *positively correlated* and vanish altogether if incomes and needs move entirely in tandem across individuals. Finally, we observed that the larger the group, the greater the possibilities of mutual insurance.

We then introduced a benchmark model, which we call the *perfect insurance* scenario. In this scenario there are a large number of people who face an aggregate shock that is common to all of them (say, the weather) and the additional possibility of idiosyncratic variations in income (or consumption needs). We showed that it is optimal to pool all the idiosyncratic variation and that it will be impossible to insure against any of the aggregate variation. Hence, the solution to the perfect insurance problem yields a testable hypothesis: if the model is good, then we should see that individual (or household-level) consumption is unaffected by *individual* incomes, but it should move in perfect correlation with the aggregate consumption of the insurance group (the entire village, let us say).

We turned to regression analysis to test this implication. Use household consumption as the dependent variable and "group consumption," household income, and other household observables (such as a spell of unemployment) as independent variables. Roughly speaking, if the theory of perfect insurance is right, the estimated coefficients on all the household-specific variables should be zero and the coefficient on group average consumption should be unity. We noted several problems with this approach that largely stemmed from missing data: for instance, what do we use for group consumption when all we have are data on a sample of households? We then discussed several empirical studies. The consensus is that although the coefficient on household income is generally small (certainly way smaller than unity), we are pretty far from perfect insurance. Idiosyncratic shocks that affect household income do have a significant effect on the consumption of that household.

This conclusion motivated a study of models with *imperfect* insurance. We identified two broad factors that might impose constraints on insurance: *limited information* and *limited enforcement*.

The informational problem is one of *moral hazard*: group members may not be sure that a particular member has suffered an income loss or a consumption spike; that is, it might be difficult or impossible to verify the outcome, on which insurance transfers are predicated. At a more subtle level, it may be possible to verify the outcome, but not what led to it: did a farmer suffer a bad harvest because he was really or truly unlucky or did he apply inadequate levels of effort to cultivation? With these uncertainties, it might be foolhardy to attempt perfect insurance: there would be huge efficiency losses in production because individuals would then have the incentive to apply inadequate effort to income-generating activities. We studied the imperfect insurance problem under these constraints and described a "second-best" insurance arrangement, suitably constrained by the limited information.

One implication of this model is that individuals run into a trade-off between diversification and moral hazard. Diversified, uncorrelated incomes across group members are easiest to achieve when members are all in different occupations or they are in the same occupation, perhaps, but are geographically separated, which is hard to accomplish within a village. On the other hand, an extension of insurance schemes across villages runs straight into the moral hazard problem, because it is much harder to verify outcomes and inputs. We discussed how this trade-off can be ameliorated in practice.

The second problem has to do with enforcement. A farmer with a good harvest may be required to make a transfer to the common insurance fund. Why would he do so? One answer is that if he does not comply, he will face social sanctions from the village. He might also be excluded from future participation in the insurance scheme. Such sanctions and exclusions have costs, but they must be traded off against the benefits from retaining extra income today. This notion gives rise to another sort of constrained problem: find the best insurance scheme consistent with the requirement that all farmers will pay up when times are good. We analyzed the properties of this scheme.

There are several implications of the enforcement-constrained model. One is that increased opportunities for self-insurance tend to destroy schemes of mutual insurance, because such opportunities lower the punishment value of future exclusion. Individuals are more likely to deviate from an ongoing scheme. Likewise, the loss of social ties in the village weakens the power of social disapproval: this, too, will cause insurance schemes to fall apart.

Another implication of the enforcement-constrained model is that it generally predicts the blurring of credit and insurance. We observed that pure insurance is "history-free": the fact that I make a transfer today has no implications for what might happen tomorrow. However, when pure insurance is ruled out by problems of enforcement, a second-best compromise might be to make my future returns from the scheme positively related to transfers that I make today. The reason is that this eases the no-deviation constraint and allows for better insurance overall. The linking of future returns to current payouts can be intrepreted as a form of credit.

We ended the chapter with a discussion of rural markets in Nigeria, where this sort of mixed scheme can be seen in practice.

Exercises

■ (1) Discuss the main problems with agricultural crop insurance. Pay particular attention to the problems of correlation and information. Using these concepts, comment on the possibility of a two-tier system in which a giant insurance company makes insurance payouts to the village as a whole, while a village council decides on how to allocate the payouts to individual residents.

■ (2) Suppose that two persons, A and B, can each earn a high income (H) or a low income (L). The following table describes the probability of each of the four events. Make sure that you understand what each cell means: in particular, you should note that $q + s + 2r = 1$.

	H for A	L for A
H for B	q	r
L for B	r	s

(a) If each person expects H with probability p, and this is independent of what happens to the other person, describe q, r, and s in terms of p.

(b) Now assume that the fortunes of the two persons are correlated. Using the variables in the table, how would you capture the idea that the incomes of *A* and *B* might be *positively* correlated? Negatively correlated? How would the extremes of perfect correlation be expressed?

(c) Which of these cases do you find most conducive for mutual insurance between *A* and *B*?

■ (3) Suppose that we find that temporary income shocks (such as those due to weather fluctuations) have small effects on the consumption of a household. Then this is *prima facie* evidence that households can smooth such shocks, but the smoothing may arise from self-insurance, credit, or mutual insurance. If you have data on consumption and income for the household over several periods of time, discuss how you might use the intertemporal data to distinguish between these three sources of smoothing.

■ (4) Suppose that a household builds up stocks of assets, such as bullocks, land, jewelry, and cash. When faced with an adverse income shock in a particular year, the household will use these assets to smooth consumption. Discuss the factors that influence the choice of assets used in smoothing.

■ (5) Consider two economies: one with a high population density, so that villages are clumped together and intervillage transportation costs are low, and another in which villages are spread apart and largely isolated from one another. In which economy would you expect bullocks to be used for consumption smoothing?

■ (6) Consider two groups of risk-averse farmers. One group practices mutual insurance, but the other does not. In the questions to follow, base your arguments on the assumption that a menu of production techniques and crops are available for cultivation, and that farmers choose from this menu.

(a) If farmers can collectively choose their crops/techniques, argue that the insuring group is likely to display a higher average rate of return on its farming investments.

(b) On the other hand, argue that an existing system of insurance might dampen the *individual* incentive to adopt newly available methods or crops with a higher average rate of return.

■ (7) Quite apart from considerations of altruism, discuss why insurance among members of an extended family is more likely than insurance among a group of strangers of the same size and occupational composition.

■ (8) We have seen that moral hazard restricts the scope of insurance: if unobserved effort must be expended to produce efficient levels of output, then

such effort may not be forthcoming in the presence of insurance. Review the relevant arguments in this chapter. Now, (i) discuss if poor farming groups are more or less likely to be prone to this problem than rich farming groups and (ii) show that a mixture of group credit and insurance (rather than pure insurance alone) may help avoid the moral hazard problem.

■ (9) Review the material on enforcement constraints in insurance. Why might higher mobility (of people) lead to a breakdown in mutual insurance?