Poor but rational?

Esther Duflo[*]

January 2003

Modern development economics emerged with the realization that poverty changes the set of options available to individuals. Poverty thus affects behavior, even if the decision maker is "neo-classical": unboundedly rational, forward-looking, and internally consistent. The "homo economicus" at the core of neo-classical economics ("calculating, unemotional maximizer", (Mullainathan and Thaler, 2000) would behave differently if he was poor than if he was rich. Asset market failures and preferences towards risk are sufficient to explain why asset ownership matters, why worthwhile transactions and investments may not always take place, and why the poor may remain poor as a result. The initial theoretical advances[1] opened a new empirical agenda to mainstream economists.

Prior to these advances, the debate had been revolving around the phrase "poor but efficient", popularized by Ted Schultz (1964). According to this, the poor certainly have bad lives but there is nothing special about them; they just do the best they can under the difficult circumstances life has placed them in; their fields are as productive as they can be (Tax, 1953), they just cannot be very productive. Rejecting (or accepting) the hypothesis of "poor but efficient", meant rejecting (or accepting) all the postulates of neo-classical economics.

When the theoretical work made it clear that being poor meant being cut off from many opportunities that were available to others, the task of empirical economics shifted to providing evidence for market inefficiencies, and the potential of economic policies to alleviate them. A new paradigm, "poor but neo-classical" (but not necessarily efficient) helped define an empirical agenda and structure a vision of the world, even though it often remained implicit in empirical work. While the poor (and the rich) are all perfectly rational, the markets left to themselves may not produce an efficient outcome. In turn, many of its predictions have been substantiated by the data. But there are also some fundamental facts for which this view of the world does not account. Using two classic examples, which have been very fertile ground for research in development economics -- insurance and agricultural investment -- I will try to explore how far this agenda led us, and what remains out of its reach.

[1] Among others: Loury (1981), Banerjee and Newman (1991), Banerjee and Newman (1993), Banerjee and Newman (1994), Galor and Zeira (1993), Aghion and Bolton (1997), and Piketty (1997).

1. Insurance

The poor, it is commonly acknowledged, face a very risky environment: the weather is uncertain, crops fail for all sorts of reasons, prices are volatile, illness strikes often, etc. Because they are close to subsistence, risk is also particularly painful to the poor. This makes insurance both more valuable and also easier to implement, since the very threat of cutting someone's insurance if he is caught cheating should be very powerful. Moreover, at least some of the poor live in a close-knit environment (the "village"), where information flows easily, and people have the possibility of exerting sanctions against each other if they are found to be abusing the system. This should alleviate some of the common problems that insurers traditionally face: How to distinguish bona fide claims from forgery? How to prevent people who know they are insured from taking unreasonable risks, now that they are not going to pay for it? How to avoid a situation where only people who know they may encounter problems in the future sign up for insurance?

In a village, part of the risk is common to all families: If there is a really severe drought, it affects everyone. But part of the risk is specific to the circumstances of specific households: for example, someone's cow may die. This has made the village institution a fertile testing ground for one of the "poor but efficient" hypothesis: Within a village, the poor should be able to insure each other against the part of the risk that is common across households.

Townsend (1994) made this point in a very influential article. Using detailed data from several Indian villages, he argued that the incomes of different families within a village have ups and downs at very different times. This creates substantial scope for insurance, and Townsend argued that advantage was fully taken of this possibility. The consumption of all families within a village move very close together: When someone has a bad year, everybody in the village suffers a little bit, and the affected family's consumption does not fall behind that of others.

This all goes to show, the article argued, that the village institution is fully efficient. The article generated a very lively controversy, and the question of how well rural households cope with risk and insure each other was a focus of much of the research in development economics in the following years.

Taking stock after ten years, it appears relatively clearly that the claim made in the seminal Townsend paper gave somewhat too much credit to the village institution. Subsequent work by Townsend himself, (Townsend, 1995a; Townsend, 1995b), as well as others, recognized the incompleteness of the insurance provided in the village, and the variety of insurance arrangements across villages.

There are several reasons why insurance may be imperfect even in the village economy. First, villagers may be able to hide part of their output from others. Second, if efforts cannot be perfectly monitored, perfect insurance might result in disastrous outcomes, if household members stop working in anticipation of being bailed out if and when their

output is low. The villagers will thus need to balance the need for insurance with the necessity of giving people incentive to work, and as a result, will insure each other only partially. Ligon (1998) argued that the data analyzed by Townsend (1994) is actually consistent with this model.

These two explanations, however, do not do justice to Townsend's central intuition, that the village institution should precisely be better than modern institutions in dealing with these dimensions. A third explanation draws from the specificity of the village institution. Agreements are not backed legally and, therefore, cannot rely on external pressure to be honored. In other words, people cannot be coerced to stay in the system: if someone is not happy with the transfers they have to make at any point in time, they can decide to walk away and operate on their own (Coate and Ravallion 1993). Worse still, several villagers may decide to walk away together and form their own insurance group (Genicot and Ray, 2003). Individuals who are enjoying a good year will then continue to participate only if the transfers they have to make today are not greater than the value of the insurance in the future. This limits the extent of insurance that can be provided in the village; in particular, someone who has done particularly well in a given year may be offered the option of receiving higher net transfers from the common pool in all subsequent years. Insurance starts looking much more like credit, since future claims are linked to today's contributions to the pool. Udry (1990) finds evidence of this in Nigerian villages, where debt repayments are contingent on how well both the creditors and the debtors did in the previous year. Udry rejects the hypothesis that villagers are fully insured.

All these arguments suggest that if individuals have good information on what others are doing (so that they cannot shirk or make false claims), and have a strong reason to stay together, they should be insuring each other. One group which seems to satisfy these conditions is the family: its members know each other, expect to stay together, and should therefore be able to achieve an efficient outcome, at least within themselves. Yet, insurance seems less than perfect in the family. The private consumption of household members in Southern Ghana seems to be completely unrelated to the income of their partners (Goldstein, 2000). Of course, this could be because, in reality, household members are effective at hiding income from each other: This seems likely, since when asked directly about their partner's income, household members seem to know very little about it. Since an individual's private consumption is mostly made of goods that can be consumed out of the house (beverages, meals taken out, transportation, kola nuts, etc.), it is plausible that the individual is consuming part of the income on the road between the market and the home, before it reaches the house and can be put into the common pool. If family members do not observe each others' actions, they cannot fully insure each other.

Troubles do not stop here, however. In joint work with Chris Udry (Duflo and Udry, 2004), we study whether household members in Cote d'Ivoire are able to insure each other against shocks that all of them can observe. Households grow different crops, which react differently to the same rainfall: For example, men tend to grow tree crops, which are sensitive more to the previous year's rain than to this year's rain. Women grow vegetables, which are sensitive to current rainfall. Thus, variation in quarterly rainfall is a

predictor of variation in each individual's income: A good year for the women in the household can be a bad year for the men in the household. We then ask the following question: If a particular year is good for the women of the household in a village, do we observe a shift in the type of goods that are consumed in the household in that year, relative to a good year for the men? This is exactly what we find: If the year is good for women in a particular household, more money is spent on food expenditures than in good years for men. In fact, none of the unexpected increase in the income from the male cash crops (coffee and cocoa) is spent on extra food; all of it is spent on private goods (clothing, alcohol, and tobacco). This is particularly striking, since none of the information we use in our analysis is unknown to the household members: We use only the information on the weather which everybody in the household can observe. The household members seem not to insure each other against variation in income that they can perfectly observe. Furthermore, the household members do not seem to have an option of just quitting the household, since they are linked by an intricate web of exchange (the women prepare food, the men bring in more income, etc.), so that it is somewhat difficult to take as given the explanation that intra-household insurance arrangements are limited by the threat of walking away. And finally, household members have strong sanctions at their disposal, which should help them enforce efficient transfers.

The study on Cote d'Ivoire adds one more piece of evidence which further undermines the hope of finding an explanation for household behavior in a "poor but neo-classical" framework. There is one crop in Cote d'Ivoire, yam, which is traditionally grown only by men (some specific operation cannot be accomplished by women, and property rights on yam fields and yam crops are clearly attributed to males). However, a strong social norm limits the legitimate uses of the proceeds from yams (Meillassoux, 1965). Yams are supposed to be used for feeding household members and for taking care of children. We therefore treated yams as a separate group, and examined how expenditure on various goods reacted to variation in yam income predicted by rainfall. Good years for yams are indeed associated with more expenditure on food, both purchased and eaten at home. Expenditures on private goods (alcohol, tobacco, ornamental clothing and jewelry) do not respond at all to increases or decreases in yam income. Expenditure on education responds only to changes in yam income, not to changes in female or male income. Men seem to treat income coming from different sources differently, and do not use the windfall from one source to compensate for any shortfall from the other source.

The household therefore seems to keep separate "mental accounts" (Thaler, 1994), treating different types of income differently. The separation of accounts goes beyond the failure to put together money that "belongs" to different people: Different sources of income are allocated to different uses depending on their origin. The fact that these accounts respond differently to observable shocks in income is difficult to reconcile with imperfect observability, moral hazard, or limits on self-enforcing insurance schemes. The complexity of intra-household sharing arrangements seems to resist explanations based only on information and incentives. Understanding why these norms and arrangements emerge (which problems they solve) and how they are sustained requires a deeper understanding of decision-making of individuals and groups. It will also shed some light

on how important economic decisions are made, and what constraints are put on the household.

2. Agricultural Investment

The classic articulation of the "poor but efficient" hypothesis for agricultural households is in the book by Schultz (1964) which argues that poor peasants are on the productive efficiency frontier, citing notably the study by Tax (1953), <u>Penny Capitalism</u>, which studies peasants in Guatemala.

If agricultural production is efficient, the investment, effort, and production on the land should not depend on who is tilling it: Whoever is working on the land should extract the maximum from the land, and these profits should then be shared. The impact of tenancy arrangements on agricultural investment and productivity suggests that the story is more complicated. For instance, Shaban (1987) showed that in India a given farmer works 40% more and uses 20% more fertilizer on his own land than on land that he is sharecropping. Even after accounting for intrinsic differences in productivity of different plots of land, agricultural productivity is 30% higher on land farmed by the owner. Tenancy arrangements are clearly inefficient.

The notion of limited liability provides a possible unified framework to explain why sharecropping arrangements arise, and why they are inefficient (Banerjee, Gertler and Ghatak, 2002). If there is a limit on how miserable someone can be (e.g., in no circumstance can someone's last bowl of food be taken away), a tenant has to be protected in bad years, and thus his payment cannot be fully dependent on how well he did. Limited liability also explains why those who do not have land cannot borrow to purchase it, and thus why there are tenants and landlords in the first place. It implies that land or wealth redistribution would increase investments and productivity. The poor are different because they are desperate (Banerjee, 2001): Having nothing to lose, they cannot be made fully responsible for their actions. They cannot thus be given the same opportunities as others, and this explains the persistence of poverty.

In this world, productivity should be maximal on owner-occupied land, at least when the necessary investments are not larger than the maximum an individual can borrow. Since land can serve as collateral, maximum permissible borrowing should be related to the value of the land. Very large investments (such as digging a new irrigation well) may not take place, even if they would be profitable eventually for a single farmer,[2] but the choice of seeds, the use of fertilizer, the use of bullocks, etc., should be efficient.

This intuition seems at odds with a number of facts, sharing the feature that the technology employed on owner-occupied farms seems to be far from the most efficient one. For example, Goldstein and Udry (1999) show that the rate of returns to growing pineapples greatly exceeds that of growing any of the other crops that are traditionally

---

[2] In practice, very large investments will often benefit more than one farmer, so one other source of inefficiency is that it will be difficult to get everyone to agree on what to build and who should pay for it.

planted by farmers: The rate of return for growing pineapple is 531%, while it is only 21% for growing traditional crops. Yet, only 24.4% of male farmers, and almost none of the female farmers, grow pineapple. This is surprising, since pineapple growing requires a moderate initial investment, which could easily be covered by a loan secured by the land owned by these farmers.

It is true that any new investment requires taking some risk that the poor may be wary of facing. For example, in India, Morduch (1993 and 1995) documents that poor farmers are less likely than richer farmers to use the high yielding varieties of seeds. One possible explanation is that farmers are reluctant to borrow, perhaps for fear of losing their land if the investment did not pay off. They will prefer a safe strategy with a low return to a riskier strategy with a higher return. It is reasonable to think that the poorer someone is, the more he dislikes taking risks. In that case, the poor (and sometimes even more so the near-poor, who have more to lose) will tend to be more conservative than the rich, and more likely to fail to undertake efficient investments. In particular, small owners may be very reluctant to use their land to secure a mortgage. This is poverty as vulnerability (Banerjee, 2001). Again, this means that the poor will tend to remain poor, while those who start with enough wealth will be able to accumulate more of it. According to Morduch, this accounts for the discrepancies in hybrid seed use across rich and poor farmers: Those who have access to assets which can safeguard their consumption are more likely to use fertilizers and hybrid seeds.[3]

There remains many facts that resist either of these views of the world. To clarify the discussion, I will focus on one example that I have been studying for a few years in collaboration with Michael Kremer and Jonathan Robinson--fertilizer use in Western Kenya.

Maize is one the main staple cereals in the region, and most farmers grow corn on their own fields, mainly for home consumption. There are still periodic episodes of food shortages in this region, so increasing agricultural productivity is critical. In experimental farms, chemical fertilizers have been shown to generate very large increases in yield, and are recommended by the agricultural ministry as part of a package including hybrid seeds, fertilizer at planting, and fertilizer at "top dressing" (when the plant is knee-high). While commercial maize farmers in other regions of Kenya all use this package, in Busia (the region we study) less than 20% of farmers use fertilizer at any point in time.

Prima facie, neither limited liability nor risk aversion seem capable of explaining such a low level of fertilizer use. Fertilizer comes in small packages, so it could be used on a very small part of the farmer's land. This would not require a large investment, so it seems implausible that the inability to borrow should be a major factor limiting investment. For example, it costs 8 shillings, less than the price of 1 kilogram of maize, to apply fertilizer on an area of 30 square meters. This is the average production of only 10 square meters of maize in a relatively bad season. Even if there is no credit market, the

---

[3] Note that the necessity for the poor to adopt this conservative behavior will disappear if the poor can be insured. In this world, asset redistribution is not the only way to increase efficiency and investments.

farmer could always save during one season to start to use fertilizer the next season. The size of the investment could be chosen so as to maintain the acceptable level of risk.

One obvious possibility is that fertilizer, while very effective in experimental farms and for commercial farmers who have access to mechanized agricultural tools, is actually not profitable for these farmers. To investigate this, International Child Support (ICS), an NGO with a long-standing presence in the area, conducted small field trials on farmers' lands. During five consecutive seasons, ICS asked a group of farmers (randomly selected from the list of parents in several schools) to select two or three small contiguous plots of land (each measuring 30 or 60 square meters). A field officer then randomly selected one plot to be a control plot, and the others to be subjected to different "treatments" (fertilizer at top dressing only, or the full "package" of hybrid seeds, fertilizer at planting and top dressing). Because the recommendations on how much fertilizer to use were conflicting, ICS also varied the quantity of fertilizer we applied on a given plot from season to season. The field officer monitored fertilizer application and measured yields in treatment and control plots, but otherwise did not provide any other inputs or instructions, so that these results should reflect the profitability of using a given quantity of fertilizer in regular conditions.

The average (annualized) rate of return of applying a small quantity of fertilizer at top dressing is very large: 231% on average over 6 seasons. With larger quantities on the same sized plot, the average rate of return goes down. The experiments are too coarse to tell us what is the optimal quantity of fertilizer to use , but it seems clear that it is positive. The full package recommended by the ministry of agriculture, however, has low or negative returns (-86% over 2 seasons); the reason is that, in real conditions, hybrid seeds often do not germinate and need to be replaced. In this case, the initial investment in fertilizer at planting is lost. The experiments thus taught us that it is clearly efficient for these farmers to use some fertilizer on their land. However, they also showed that, even though the technology has been around for a long time, what to use and how to use it are not necessarily effectively conveyed to the farmers.

This may open an avenue to rescue the hypothesis that lack of access to credit prevents fertilizer use. Learning about a technology requires some experimentation, which needs effort, and presumably some losses in the first years (Foster and Rosenzweig, 1995). If the farmer knows that after this initial investment, he will only be able to use fertilizer on a very small scale for a long period of time (because he cannot have access to the funds to do it on a large scale), he might rightly consider it worthless to do the initial investment to master the technology.

The field trials allow us to directly evaluate the strength of this hypothesis: In essence, they constituted a particularly intensive form of agricultural extension. After the harvest, an ICS field officer visited each farmer, and discussed their experience as well as the experience of everybody who participated in the trials with them. The field officer went through detailed profit calculations, both using the farmer's own data and using the average numbers. The farmer thus learnt both how to use fertilizer and how profitable it really was: ICS gave the farmer a chance to learn how to use fertilizer at no cost to them.

If learning how to use fertilizer was the main barrier to adoption, we might thus expect that farmers participating in the trials would have higher fertilizer use than those who did not. Since farmers were randomly selected from parents of a few schools, the other parents of the same schools form a natural comparison group. We thus monitored the adoption of fertilizers by all farmers who had participated in ICS trials for the following seasons.

Indeed, fertilizer adoption increased in this group. On average, while 20% of the farmers who did not participate in the trial adopted fertilizer in a given season, 37% of those who did participate were using fertilizer.

The explanation in terms of fixed cost of learning may thus have some bite, but fails to explain the entire phenomenon. First, more than 60% of the farmers continue not to use fertilizer. And secondly, the effect of having participated seems to diminish over time: After three seasons, only 29% of farmers who had initially participated in the trial use fertilizer.

We thus have to turn to other explanations. A possibility, which sounds plausible in view of the farmers' overwhelming explanation of why they are not using fertilizer (98% of them say they do not use fertilizer because they have no money) is that farmers find it very difficult to save even the small amount of money necessary to purchase fertilizer. We also observe that farmers seem reluctant to use fertilizer on a very small plot: Perhaps, there is a psychological cost to deviating from "business as usual". In this case, farmers need to be able to save (or borrow) a reasonable amount of money to start using fertilizer.

To test the hypothesis that the lack of savings opportunities prevents the farmers from using as much fertilizer as they would want, ICS offered farmers the option to purchase fertilizer on their behalf immediately after the harvest, when maize is relatively plentiful. ICS would either deliver the fertilizer right away, or store it until the time of top dressing. The farmer could pay either with cash, or with the equivalent value in maize.

This opportunity was offered to farmers in six different seasons, and proved very popular. In the first year, 12 out of the 54 farmers to whom ICS had offered this opportunity purchased fertilizer in advance. The increase in fertilizer adoption occasioned by the scheme was the same as the number of people who elected to purchase fertilizer under the scheme. This suggests that all those who elected to take up the scheme were new adopters. By offering farmers a way to save for just a few weeks, ICS doubled the rate of fertilizer usage.

In the following years, the scheme was offered to farmers who had participated in the agricultural trials. Among those, the take-up of the scheme varied at least 30%. It was higher when farmers were asked to make a decision on the spot, and reached 80% among farmers who had sold a fraction of their maize crop to ICS for some other purpose. This simple commitment device led to a very sharp increase in adoption rate, over and above the effect of learning alone.

We are far from understanding fully why this program was so successful. Did farmers value the opportunity to commit their money to be used for fertilizer purchase because they knew that if they kept the money they would be tempted to use it on other things before the time for purchasing and using the fertilizer occurs (Laibson, 1991)? Are they particularly keen to exchange cash for fertilizer because maize is not likely to disappear quite as rapidly? Are they protecting the cash against themselves (they know now that they need to tie their hands), or are they worried about their families or their neighbors? Were farmers just lured by fertilizer as one possible purchase, and would they have been just as happy with a radio or some alcohol? Were farmers particularly aware of the value of fertilizer, after we had just weighed and valued the output of the two plots, and would this have lost its salience if they had waited longer?

Many further experiments will be necessary to disentangle these possibilities, and, with ICS, we are currently designing and implementing some of them. What is clear, however, is that the inefficiently low level of fertilizer investment in this region cannot easily be accounted for in the "poor but neo-classical" paradigm. It is definitely the task of economists to understand this fundamental production decision. It will require a deeper understanding of the determinants of decision-making in a very poor environment than the "homo economicus" framework has left space for.

3. Towards a theory of behavior in poverty

The "poor but neo-classical" research program systematically explored the structural constraints that poverty imposes on the decision-making of an unboundedly rational individual in an environment where information is incomplete. The two examples we just covered have shown that many important facts can be accounted for in this framework, but that other facts are resistant.

A natural tendency would then be to turn to new hypotheses, incorporating insights of psychology to better understand economic decisions: This has been the research agenda of "behavioral economics" (Mullainathan and Thaler, 2000). Behavioral economists have argued that three main ways in which humans deviate from the standard economic model need to be incorporated into economic reasoning. Their ability to analyze information, compute, and remember is limited. Their willpower is also bounded: They do not always make choices that are in their best interest in the long run. Finally, they are not purely self-interested. Developing countries could then be seen as ideal testing grounds for some of these theories, since program evaluation experiments and field experiments can be combined, as in the example of the fertilizer program. There may be more to learn about human behavior from the choices made by Kenyan farmers confronted with a real choice than from those made by American undergraduates in laboratory conditions. There are a few examples of "real world" experiments that test specific hypotheses in a developed country, such as the "Save More Tomorrow" Program (Thaler and Benartzi, 2004), where individuals were offered the possibility of signing up to a program where a part of their future salary increases would be saved. There are also natural experiments, such as

changes in default rules for enrollment in voluntary retirement savings plans (Madrian and Shea, 2002), or variation in the debt structure of a financial aid package with the same monetary value (Field, 2002). Real or natural experiments, however, are rare in developed countries; it might be easier to collaborate with NGOs to offer programs with different rules, and different characteristics, designed to test specific behavioral hypotheses. This could be a fruitful avenue of collaboration between practitioners and academics, since it may help practitioners to design programs that effectively help improve an individual's options.

Faithfully applying the theories designed for developed countries to the analysis of the decisions of the poor in developing countries would, however, be making the same mistake as the proponents of the "poor but efficient" framework, and failing to recognize the central insight of the "poor but neo-classical" line of research. Trying to reduce the behavior of a Kenyan farmer who does not use fertilizer and that of an American employee to the same model may be as fruitless as trying to convince oneself that Guatemalan farmers are on the efficiency frontier. Being poor almost certainly affects the way people think and decide. Perhaps when choices involve the subsistence of one's family, trade-offs are distorted in different ways than when the question is how much money one will enjoy at retirement. Pressure by extended family members or neighbors is also stronger when they are at risk of starvation. Or, decision making may be influenced by stress.

What is needed is a theory of how poverty influences decision-making, not only by affecting the constraints, but by changing the decision-making process itself. That theory can then guide a new round of empirical research, both observational and experimental.

References

Aghion, Philippe, and Patrick Bolton (1997) "A trickle-down theory of growth and development with debt overhang," *Review of Economic Studies*, 64(2), 151--72.

Banerjee, Abhijit V. (2001) "The two poverties," *Nordic Journal of Political Economy*, 26(2), 129--141.

Banerjee, Abhijit V., Paul L. Gertler, and Maitreesh Ghatak (2002) "Empowerment and efficiency: The economics of tenancy reform," *Journal of Political Economy*, 110(2), 239--280.

Banerjee, Abhijit V., and Andrew Newman (1991) "Risk bearing and the theory of income distribution," *Review of Economic Studies*, 58(2), 211--235.

Banerjee, Abhijit V., and Andrew Newman (1993) "Occupational choice and the process of development," *Journal of Political Economy*, 101(2), 274--298.

Banerjee, Abhijit V., and Andrew Newman (1994) "Poverty, incentives, and development," *American Economic Review*, 84(2), 211--15.

Coate, Stephen, and Martin Ravallion (1993) "Reciprococity without commitment: Characterization and performance of informal insurance arrangements," *Journal of Development Economics*, 40, 1--24.

Duflo, Esther, and Christopher Udry (2004) "Intrahousehold resource allocation in Côte D'ivoire: Social norms, separate accounts and consumption choices," NBER Working Paper #10498.

Field, Erica (2002) "Educational debt burden and career choice: Evidence from a financial aid experiment," Princeton University Industrial Relations Section Working Paper #469.

Foster, Andrew D., and Mark R. Rosenzweig (1995) "Learning by doing and learning from others: Human capital and technical change in agriculture," *Journal of Political Economy*, 103(6), 1176--1209.

Galor, Oded, and Joseph Zeira (1993) "Income distribution and macroeconomics," *Review of Economic Studies*, 60(1), 35--52.

Genicot, Garance, and Debraj Ray (2003) "Group formation in risk-sharing arrangements," *Review of Economic Studies*, 70(1), 87-113.

Goldstein, Markus Paul (2000) "Intrahousehold allocation and farming in Southern Ghana," Ph.D. Dissertation, University of California, Berkeley.

Goldstein, Markus, and Christopher Udry (1999) "Agricultural innovation and resource management in Ghana," Mimeo, Yale University; Final Report to IFPRI under MP17.

Laibson, David (1991) "Golden eggs and hyperbolic discounting," *Quarterly Journal of Economics*, 62, 443--477.

Ligon, Ethan (1998) "Risk sharing and information in village economics," *The Review of Economic Studies*, 65(4), 847--864.

Loury, Glenn C. (1981) "Intergenerational transfers and the distribution of earnings," *Econometrica*, 49(4), 843--867.

Madrian, Brigitte, and Dennis F. Shea (2002) "Preaching to the converted and converting those taught: Financial education in the workplace," Mimeo, Graduate School of Business, University of Chicago.

Meillassoux, Claude (1965) Anthropologie Economique des Gouros de Cote d'Ivoire
        (Paris: F. Maspero).

Morduch, Jonathan (1993) "Risk production and saving: Theory and evidence from
        Indian households," Mimeo, Harvard University.

Morduch, Jonathan (1995) "Income smoothing and consumption smoothing," *Journal of
        Economic Perspectives*, 9(3), 103--114.

Mullainathan, Sendhil, and Richard Thaler (2000) "Behavioral economics," NBER
        Working Paper 7948, October.

Piketty, Thomas (1997) "The dynamics of the wealth distribution and the interest rate
        with credit rationing," *The Review of Economic Studies*, 64(2), 173--189.

Schultz, Theodore W. (1964) Transforming Traditional Agriculture (New Haven,
        London: Yale University Press).

Shaban, Radwan (1987) "Testing between competing models of sharecropping," *Journal
        of Political Economy*, 95(5), 893--920.

Tax, Sol (1953) Penny Capitalism: A Guatemalan Indian economy. (Washington, DC:
        Smithsonian Institution Institute of Social Anthropology, Publication No. 16).

Thaler, Richard H. (1994) "Psychology and savings policies," *AEA Papers and
        Proceedings* 84(2), 186--192.

Thaler, Richard H., and Shlomo Benartzi (2004) "Save more tomorrow: Using behavioral
        economics to increase employee saving," *Journal of Political Economy*, 112(1),
        S--164--187.

Townsend, Robert (1994) "Risk and insurance in village India," *Econometrica*, 62(4),
        539--591.

Townsend, Robert (1995a) "Consumption insurance: An evaluation of risk-bearing
        systems in low-income economies," *Journal of Economic Perspectives*, 9(3), 83--
        102.

Townsend, Robert (1995b) "Financial systems in Northern Thai villages," *Quarterly
        Journal of Economics*, 110(4), 1011--1046.

Udry, Christopher (1990) "Credit markets in Northern Nigeria: Credit as insurance in a
        rural economy," *World Bank Economic Review*, 4(3), 251--69.