Microeconomic Theory
Lecture 12: Moral Hazard and Sharecropping

Marcel Fafchamps
Oxford University

Readings: Varian, Chapter 25.

Optional reading:

Bernard Salanié, *The Economics of Contracts*, 1998, MIT Press.

J.E. Stiglitz, "The Causes and Consequences of Dependence of Quality on Price", *Journal of Economic Literature,* 25:1-48, 1987.

Pranab Bardhan, *The Economic Theory of Agrarian Institutions,* Oxford UP

Debraj Ray, *Development Economics*, Princeton University Press

1. Imperfect information in economics

The last two or three decades have witnessed an explosion of theory work on imperfect information issues. Early work on contracts and information asymmetries goes back to the early 1970's with Akerlof's paper on adverse selection (QJE 1970), Stiglitz' paper on moral hazard (REStud 1974) and Spence's book on signaling (1974). Research on these issues only truly picked up in the 1980's, with several other papers by Stiglitz and rigorous efforts at generalizing the moral hazard and adverse selection models.

In part because Stiglitz' 1974 paper was motivated as an analysis of sharecropping, a new generation of development economists rushed into this new research area and began applying the imperfect information toolbox to all kinds of agrarian structure issues. Several of these young economists were originating from India (Bardhan, Basu, Braverman, etc) where agrarian institutions were a hot policy topic, especially in and around West Bengal from where many of them were coming. This resulted in an explosion of applied theory papers purporting to represent one or another feature of rural India. Unfortunately, little empirical work was undertaken at the time. Moreover, many of the models that were proposed could not be tested anyway because they required data on information asymmetries which is difficult if not impossible to collect (virtually by design). So for a while the applied theory literature went into free-wheeling mode, a state of affairs that is still reflected in existing development economics textbooks. My perception is that interest in sharecropping and asymmetric information per se has now waned. But there remains a very strong focus on institutional issues, which are perceived as fundamental to the development process.

Akerlof, Stiglitz, and Spence pioneered a new way of looking at the role of information in economics. In so doing, they were only revisiting issues initially raised by Hayek and Coase in the 1920's and 1930's – but long forgotten. As Hayek (1945) originally pointed out, incomplete information is the norm rather than the exception: workers know better what they are doing than their boss, producers know better the quality of their products than consumers, borrowers know better their chances of repaying a loan than lenders, etc. Hayek was a fervent believer in the market system and used these ideas to indict centrally planned economies and to extol the virtues of the market mechanism. Today, however, the theory of incomplete information focuses primarily on how imperfect information distorts market outcomes. Examples of market outcomes that can result from imperfect information include:

- information processing: large organizations (firms, government agencies, international aid organizations) spend time and resources circulating information (memos, reports, meetings, accounts); this is necessary to ensure coordination. Circulation of information requires literacy and numeracy skills (to read and write reports and hold accounts). Without these skills, countries cannot have effective large organizations.
- search: agents spend time and resources finding out about trading opportunities, e.g., job search, unemployment.
- statistical discrimination: e.g., lenders or employers rely on observable characteristics (race, gender, place of residence) to infer hidden traits.
- signaling: e.g., agents invest in non-productive activities to signal their type (education, dress code, religion, speech).
- rationing: e.g., certain borrowers are unable to get a loan at any price.
- market failure: e.g., certain types of exchange do not take place because quality/effort cannot be assessed or because search is too costly, e.g., insufficient delegation of authority, overworked bosses.
- market segmentation: e.g., merchant banks lend to corporations at a low rate; commercial banks lend to small firms at a higher rate.
- advertising: resources spent to inform the consumer about trade opportunities.
- intermediaries: agents specialize in facilitating search, disseminating information, e.g., banks, brokers, job agencies, traders.
- third-party certification: e.g., reference letters, audits are examples of non-productive activities that only serve to circulate information.
- networking: e.g., agents cultivate friends and acquaintances as sources of information on job openings, product quality, market opportunities.
- long-term contracts: e.g., once a reliable and competent employee has been found, employer tries to keep him/her.
- incentive contracts: e.g., contracts that seek to provide incentives to work hard and perform well, as opposed to straight sales.

The issues listed above are at the core of industrial organization and labor economics. Most of the emphasis of information theory has been on the last point, the so-called theory of contracts and two of its most analyzed variants, the hidden action model (often referred to as moral hazard) and the hidden information model (often referred to as adverse selection). These two types of contracts have received an inordinate amount of attention, possibly because they have long resisted efforts to derive simple intuitive results from them. For a while they were like the holy Grail of contract theory; now they have fallen somewhat by the wayside. Still, they remain the required entry point into an enormous literature that, for the most part, casts itself in relation to hidden action or hidden information models. For this reason, penetrating the development literature on agrarian institutions is much easier if one is armed with a good understanding of these two basic models. Learning about these two models saves a lot time.

Contract theory is organized around a basic set-up called the principal-agent model. This model, which analyzes the strategic interactions between two stylized economic agents, is quite unusual because it assumes a basic asymmetry between them. The agent is the usual utility/profit maximizer of previous lectures, but the behavior of the principal is quite different. First, the principal is assumed to move before the agent. Second the principal decides what contract to offer to the agent. Third, the principal acts strategically in the sense that he anticipates the reaction of the agent and chooses the contract accordingly. In other words, when deciding what kind of contract to offer to, say, an employee, the principal figures how hard the employee will work under each contract. (For those familiar to imperfect competition models, the set-up is not dissimilar to the Stackelberg leader model: the principal leads and the agent follows.) In practice, this is equivalent to assuming that the principal computes the reaction function of the agent to all possible contracts and chooses the contract that maximizes the principal's utility. Of course, the agent always has the right to walk away from a contract; after all, economics is about voluntary exchange. The principal thus also has to worry about giving enough that the agent wishes to enter into a contract. Voluntary participation enters into the model in the form of a voluntary participation or individual rationality constraint. The reaction function calculation is called the incentive compatibility constraint in this literature.

Most contract models are partial equilibrium models: they focus on the relationship between two individuals in isolation from the rest of the economy. Because contract models are mostly interested in the qualitative features of a particular situation, their results do in general carry through even when agents and principals compete with each other. In his Chapter 25, Varian considers one specific case in which entry into the principal's category is free, but the number of agents is fixed. With free entry, principals' profits must be driven to zero. Imposing the zero profit condition on top of the model can thus identify a particular equilibrium. Another possibility, not discussed by Varian, is when there is free entry in the agent's category, but fixed in the principal's category. This type of equilibrium can be analyzed simply by setting the voluntary participation constraint so that agents derive zero net profit from participation. We will revisit these issues further in details below.

2. The full information principal-agent model

The full information model serves as a useful benchmark. It can be written as follows:

$$\max_{b,s(.)} x(b) - s(x(b)) \text{ subject to} \tag{2.1}$$

$$s(x(b)) - c(b) \geq \bar{u} \text{ (voluntary participation constraint)} \tag{2.2}$$

$$s(x(b)) - c(b) \geq s(x(z)) - c(z)\ \forall z \quad \text{(incentive compatibility constraint)} \tag{2.3}$$

where z is an arbitrary action by the agent, b is the action that the principal tries to promote, $s(.)$ is the contract/incentive schemes that the principal proposes to the agent, $x(z)$ is the output of action z, $c(z)$ is the cost of taking action z for the agent, and $\bar{u}$ is the reservation utility of the agent, that is, the utility that the agent could guarantee himself if he were to refuse the contract proposed by the principal. The principal's utility is equal to the output $x(b)$ minus what is paid to the agent, $s(x(b))$. The agent's utility is what he gets from the principal $s(x(b))$ minus the cost of taking the action $c(b)$. The voluntary participation constraint requires that the agent gets at least as much from the contract than what he could guarantee himself by refusing to contract. The incentive compatibility constraint ensures that, if b is the action stipulated by the principal, it is also the action that yields the highest payoff to the agent.

Note that the principal chooses not just the value of a particular variable, as in previous optimization models; he also chooses an entire function $s(.)$. This makes principal-agent problems much harder than the kind of optimization problems we have analyzed so far. Without any restriction imposed on the form of the function, it is very difficult in general to figure what the function $s(.)$ should look like. In this case, the contract $s(.)$ is a function of the agent's action. This is only possible if the principal can actually observe the action of the agent. This is why the above model is called the full information model: the principal knows everything there is to know, from the action of the agent to his preferences $c(b)$ and reservation utility $\bar{u}$.

The solution to the above model turns out to be extremely simple. In this case, it is better to rely on intuition rather than math to solve the problem. First note that since the principal observes the agent's action, he can stipulate the action b and set the agent's payoff such that any other action will result in an arbitrarily large negative payoff. Given this payoff structure, the agent is bound to choose the action suggested by the principal (incentive compatibility constraint satisfied). This solves for the shape of function $s(.)$, i.e., a set reward if the agent follows the order of the principal, and an large (possibly infinite) penalty if the agent deviates. Note that the function $s(.)$ is non-linear and discontinuous. In fact, it is quite a bizarre function that would have been very hard to derive algebraically. The value of $s(b)$ is fixed by the voluntary participation constraint, taking as given that the agent will act as told. Since there is no reason for the principal to give more to the agent than is strictly necessary, the voluntary participation constraint is binding exactly and $s(x(b)) = \bar{u} + c(b)$. This means that the principal maximizes:

$$\max_b \; x(b) - c(b) - \bar{u}$$

The solution is the first best outcome $x'(b^*) = c'(b^*)$, marginal revenue equals marginal cost, where b^* denotes the first best level of output.

The conclusion is that with full information, the outcome of the principal-agent model is first best. (Beware: we have not proved this for more general, i.e., non-additive utility functions for the agent and principal. See Salanié for details.)

2.1. Contract restrictions and Marshallian inefficiency

In the model above we have not imposed any restriction on the contract function. It is useful to examine what would happen if we restricted $s(.)$ to take a specific form. We first examine what happen if we impose the condition that:

$$s(x) = \alpha x$$

If we limit our attention to contracts of this form, the principal-agent problem takes the form:

$$\max_{b,\alpha} \; (1-\alpha)x(b) \text{ subject to} \tag{2.4}$$

$$\alpha x(b) - c(b) \geq \bar{u} \text{ (voluntary participation constraint)} \tag{2.5}$$

$$b = \arg\max_b \alpha x(b) - c(b) \text{ (incentive compatibility constraint)} \tag{2.6}$$

where we have written the incentive compatibility contrainst in another equivalent form. The first order condition for the agent is:

$$\alpha x'(b^{**}) = c'(b^{**})$$

from which we immediately see that $b^{**} < b^*$ if $\alpha < 1$. This is easy to show graphically: $c'(b)$ typically is flat or curves up while $x'(b)$ falls down – at least at optimum point.[1] Consequently, multiplying the $x'(b)$ curve by $\alpha < 1$ shifts the curve down and yields an intersection at a lower level of b. Of course, the principal could achieve first best by setting $\alpha = 1$, but then the principal would not get paid, which is obviously not optimal from the principal's point of view. This result is a cornerstone of the literature on incentives. It was first noted by Marshall and for this reason is called Marshallian inefficiency. It is the reason why Marshall condemned sharecropping as an archaic and inefficient contract.

It is interesting to see what happens if we generalize ever so slightly the contract to allow for a fixed transfer/side payment F in addition to the share α:

$$s(x) = \alpha x - F$$

The difference appears minor yet it has dramatic consequences on the solution. The reason is that the principal has one additional degree of freedom. This enables him to give sufficient incentives to the agent to exert effort while capturing the surplus with the side payment F. Formally, the principal-agent problem takes the form:

$$\max_{b,\alpha,F} (1-\alpha)x(b) + F \text{ subject to} \tag{2.7}$$

$$\alpha x(b) - c(b) - F \geq \bar{u} \text{ (voluntary participation constraint)} \tag{2.8}$$

$$b = \arg\max_b \alpha x(b) - c(b) - F \text{ (incentive compatibility constraint)} \tag{2.9}$$

From our earlier example, we already know that setting $\alpha = 1$ achieves first best. We also know that first best maximizes the surplus or gain from trade. The question then is whether the principal can capture all the surplus via F. If he can, then setting $\alpha = 1$ is optimal since it generates the largest surplus. We want the principal to capture all the surplus, i.e. that:

$$\begin{aligned} \alpha x(b) - c(b) - F &= \bar{u} \\ F &= \alpha x(b) - c(b) - \bar{u} \end{aligned}$$

Since there is no reason for the principal not to take all the surplus irrespective of the value of α, we therefore conclude that $\alpha = 1$ and $F = x(b^*) - c(b^*) - \bar{u}$ are the solution in this case.

What is remarkable is that when we allow for the side payment we again obtain first best even though we have not used the non-linear, discontinuous contract of the previous sub-section. An immediate corollary is that there exist many contracts that can deliver the same outcome – hence making the search for 'the' optimum contract difficult since it is not unique. We also see that first best is achieved here but the contract does not resemble sharecropping. Rather, it resembles a fixed rental: the agent pays a fixed price F but keeps all output. In the jargon of this literature, the agent is called the 'residual claimant' on output. Incentive theory often insists that to provide maximum incentive, an agent should be made residual claimant. One problem with this approach is that the agent must be able to shoulder the risk associated with possible fluctuations in output. If the agent cannot, a lower share of α and a lower payment F may be required to insure the agent against bad states of nature. This is basically the idea behind the hidden action model. To this we now turn.

[1] If the two curves intersect at a point where $x'(b)$ is increasing, the second order condition is violated and the intersection describes the minimum, not the maximum.

3. Imperfect information

3.1. The hidden action model

The hidden action model differs from the above in that the principal can no longer observe the action of the agent. Of course, if output is a deterministic consequence of the agent's action, this is not a problem: the agent's action can be perfectly inferred from observing output. The model in section 2 applies. The situation becomes interesting only if output is stochastic. Assume that there are n possible output levels $(x_1, ..., x_n)$. These output levels occur with probabilities that depend on the agent's action. To keep things simple, assume that the agent can only take two actions, a or b and let's denote the probability of output i if action j is taken as π_{ij}. Write $s(x_i)$ as s_i. Supposing that action b is taken, the principal-agent problem now is:

$$\max_{s_i} \sum_{i=1}^{n} (x_i - s_i)\pi_{ib} \text{ subject to}$$

$$\begin{aligned}
\sum_{i=1}^{n} u(s_i)\pi_{ib} - c_b &\geq \sum_{i=1}^{n} u(s_i)\pi_{ia} - c_a \text{ (incentive compatibility constraint)} \\
\sum_{i=1}^{n} u(s_i)\pi_{ib} - c_b &\geq \bar{u} \text{ (voluntary participation constraint)}
\end{aligned}$$

One important feature to recognize in the above model is that problems arise principally when the agent is risk averse. To see why, suppose the contrary, i.e., that the agent is risk neutral. It is then possible for the principal to charge a fixed price and let the agent be residual claimant of the product of his effort, i.e., $s_i = x_i - F$. Faced with this contract, the agent chooses the action that equates marginal return to marginal cost and first best is achieved. The level of F is determined by the voluntary participation constraint, i.e., $F = \sum_{i=1}^{n} x_i\pi_{ib} - c_b - \bar{u}$. In other words, the principal extracts all the surplus generated by the agent, $\sum_{i=1}^{n} x_i\pi_{ib} - c_b$, except for reservation utility $\bar{u}$. Since first best maximizes joint surplus and the principal can extract all the surplus, this solution is optimal. Incentive problems thus arise because the agent is unwilling or unable to assume all the risk generated by his actions. If the agent could be made residual claimant of the fruits of his action, incentive problems disappear.

In the general case with a risk averse agent, it turns out that the above optimization problem is incredibly hard to solve. This is true in spite of the fact that we have already made numerous simplifying assumptions: finite number of outcomes, two actions, all outcomes reachable from the two actions, simple preferences, no equilibrium conditions, etc. The fact is that there is no general result regarding this model.

In order to obtain a result, it is necessary to make additional assumptions either regarding the form of the contract itself (i.e., restricting one's attention to specific functional forms such as fixed share of output to agent, with or without side payment), the probabilities π_{ij}, and/or the preferences $u(s_i)$. In particular, it is incorrect to say that it is in general in the interest of the principal to provide incentives for the agent to exert the costly, high productivity action. This is surprising since the phrase 'moral hazard' is often referred to situations where economists would suggest the presence of an 'incentive' problem, meaning, that the agent must provided adequate incentives to apply proper care. In this sense, the rigorous literature on moral hazard is much less conclusive than casual economic speak would lead us to believe. The literature is nicely summarized in Salanié's book.

3.2. An application of moral hazard to sharecropping

One example of interest to us is that of sharecropping. Since Stiglitz (1974), moral hazard is thought to be a major rationale behind the existence of sharecropping. The way this is shown is to take the hidden action model above and impose a particular functional form to the payment schedule s_i such that $s_i = \alpha x_i + \beta$. In other words, compensation of the agent is limited to contracts that stipulate a fixed payment β (positive or negative) and a payment $(1-\alpha)x_i$ proportional to output. This functional form includes not only sharecropping but also fixed rent (negative β, and $\alpha = 1$) and fixed wage (positive β and $\alpha = 0$). With this simplification (and a few others...), it is possible to show that sharecropping is optimal when the landlord (principal) is risk neutral and the tenant (agent) is risk averse. In addition, a special case exists where the tenant's share α decreases as his aversion to risk increases (a monotonicity result). When the tenant is risk neutral, fixed rental is optimal.

Singh ("Theories of Sharecropping", in *The Economic Theory of Agrarian Institutions*, Pranab Bardhan ed., OUP, 1989) presents a nice summary of the moral hazard sharecropping model. In Singh's notation, the landlord's problem is written:

$$\max_{\alpha,\beta,L} E_\theta[(1-\alpha)\theta Q(L) - \beta] \text{ subject to}$$

$$\begin{aligned} E_\theta[U(\alpha\theta Q(L) + \beta)] - L &\geq K \\ E_\theta[U'(\alpha\theta Q(L) + \beta)\alpha\theta Q'(L)] - 1 &= 0 \end{aligned}$$

where θ is a random output shock with $E[\theta] = 0$, L is labor, $Q(L)$ is output, α is the share of output going to the tenant, β is a fixed payment (which can be positive or negative), and K is the tenant's reservation utility – i.e., the utility the tenant could achieve if he were to refuse the contract offered by the landlord. The first constraint is the voluntary participation constraint and the second is the incentive constraint, expressed in the form of the tenant's first order condition.[2] Note that we have assumed that the landlord is risk neutral and that the tenant's utility is additively separable in labor and consumption.

Solving the landlord problem works in two steps. We first note that, since utility is increasing in consumption, the landlord can always push the tenant all the way down to his reservation utility by setting β high enough. Consequently, we only need to consider the case when the participation constraint is binding. The two constraints can thus be solved for values of L and β as a function of α and K. Let's write these relationships $L(\alpha, K)$ and $\beta(\alpha, K)$. Substituting into the landlord's objective function, we obtain a first order condition for α:

$$\begin{aligned} -Q + (1-\alpha)Q'L_\alpha - \beta_\alpha &= 0 \\ \alpha &= 1 - \frac{Q + \beta_\alpha}{Q'L_\alpha} \end{aligned}$$

By totally differentiating the voluntary participation constraint, we get:[3]

$$\beta_\alpha = \frac{d\beta}{d\alpha} = -Q\frac{E[U'\theta]}{E[U']}$$

[2] The conditions under which the first order condition can be used in this way is discussed in Salanie. Basically, the tenant's optimization problem must be convex and have a single interior solution for this approach to work.

[3] This is what Singh does, but strictly speaking it is incorrect. The correct approach would be to totally differentiate the two constraints with respect to L, α, and β, and solve the resulting (linear) system in terms of $d\beta/d\alpha$. This would allow L to vary as α varies. Here we implicitly assume that L_α is small and can be ignored for the sake of obtaining $d\beta/d\alpha$.

which is negative and less than Q in magnitude if the tenant is risk averse. Hence $Q + \beta_\alpha > 0$. Whether α is less than 1 depends on the sign of L_α. Intuitively, one would indeed expect that effort increases with the incentive to work α so that L_α. If this is the case, $\alpha < 1$. However, it is difficult to establish $L_\alpha > 0$ in general. The reason, I believe, is that α also has an effect on income and higher income reduces the incentive to work, as the tenant substitute leisure for consumption. Singh claims that, if $U''' < 0$, then $L_\alpha > 0$.

Many researchers have noted the predictions of this model are at first glance hard to reconcile with the evidence. The model predicts that the share varies with the circumstances of the landlord and tenant and consequently should vary across tenancy contracts. In practice, sharecropping contracts often taken set values, such as one half or one third. Also, the model does not generate 'mass points' around fixed rental and fixed wage contracts; such contracts are predicted to arise only with Lebesgue measure 0.

Other versions of this model exist. One of them forces the side-payment to be 0. In this case, $0 < \alpha < 1$ always, but this model cannot explain why the side-payment is not allowed (since it would increase the payoff to the landlord) and without side payment, the model does not have fixed rental and fixed wages as possible outcomes. Ultimately, the hidden action model is not a convincing model of sharecropping. It has proved much more useful in thinking about managerial incentives in firms.