

Technological Design and Moral Hazard

Edward Simpson Prescott

The classic moral hazard model studies the problem of how a principal should provide incentives to an agent who operates a project for him. In this model, the principal only observes the realized output and not the agent's effort. Consequently, the agent must be induced to work hard with compensation that depends on performance. Because many contracts explicitly tie rewards and punishments to performance, the model is a workhorse of modern economics, with applications to insurance contracts, employee and executive pay, sharecropping contracts, corporate finance, and bank regulation, to name just a few.

In the moral hazard model the exact dependence of compensation on performance depends on the relationship between the agent's input and the project's output. Most analysis takes this relationship, or technology, as given; that is, something that cannot be modified by the principal.

There are many situations, however, where the principal has some control over this technology. For example, a principal can design a production process so some outcomes are more likely than others when certain inputs are applied. A production line can be designed so that if sufficient care is not supplied, it will break down. In agriculture, the fertilizer, the type of seed, and other inputs all effect the stochastic properties of production. Debt contracts frequently include loan covenants that put restrictions on the activities of a borrower, such as working capital requirements.¹ Financial regulation works similarly. Banks have limits on their activities. For example, a bank cannot lend more than a set fraction of its assets to a single borrower. Similarly, money market mutual funds are limited to investing in short-term, safe, commercial paper, and as a

■ The views expressed in this article do not necessarily represent those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

¹ See Black, Miller, and Posner (1978) for more information on loan covenants as well as connections with bank regulation.

consequence they have a very different risk profile than banks, even though the money market liabilities are close substitutes to some bank liabilities.

In each of the above examples, the principal has some choice over the functional relationship between the agent's actions and his output. In the agricultural case, the connection is through use of inputs. For debt contracts, loan covenants are used to keep a borrower away from potentially dangerous conditions. In the money market and bank regulation examples, the investment restrictions reduce the variance of returns.

The purpose of this article is to work out some of the implications of this line of thought.² Only some are explored because there are many different dimensions along which the technology could be changed. Consequently, the analysis is necessarily limited and mainly exploratory. Still, it emphasizes the principles at work and demonstrates why this margin of choice is potentially important.

Most of the issues are illustrated with two examples in which the principal can adjust the technology. The first example gives the principal wide latitude in determining the technology and starkly illustrates how powerful this margin may be. It also demonstrates that this margin strongly affects the optimal contractual form. The second example limits the principal to choosing between only two technologies, but it demonstrates that the principal may be willing to choose a less productive technology, as measured by expected output, because it reduces the incentive problem. In this example, the limited choice among technologies is motivated by the interaction between a principal's decision and inferences made from financial market prices. In particular, the principal's decision to liquidate the firm alters the informativeness of market prices.

1. THE MODEL

In the basic moral hazard model, the agent chooses an action a that combines with a random shock to produce output q . In this article the principal has some control over how a interacts with the randomized shock to produce output. The choice made by the principal is called the *technological* choice and is indexed by i . The relationship between the inputs and the output is described by the conditional probability distribution function $p_i(q|a)$. For simplicity, q is assumed to take on only a finite number of values. Of course, $\sum_q p_i(q|a) = 1$. Finally, based on the output, the principal pays the agent his consumption c .

²The moral hazard literature touches on this issue, but the implications and importance of this idea may not be fully appreciated, since the results are scattered across different applications. One important application of this idea is in the task assignment model of Holmstrom and Milgrom (1991) and Itoh (1991). They study how to assign workers to tasks, which in turn affects the technology faced by an agent. One paper, however, that explicitly considers the principal's choice of technology is Lehnert, Ligon, and Townsend (1999).

Preferences

The agent cares about his consumption and his effort. His utility function is $U(c) - V(a)$, with U strictly concave and V increasing in a . The principal is risk neutral so he only cares about the project's surplus, that is, $q - c$.

The principal offers the agent a contract that consists of three items: the principal's choice of technology, what action the agent takes, and how the agent is paid as a function of the output. Formally,

Definition 1 *A contract is a technological index i , a recommended action a , and an output-dependent compensation schedule $c(q)$.*

The agent has an outside opportunity that gives him \bar{U} units of utility. For this problem, this means that the contract has to give him at least that amount of expected utility before he will agree to work for the principal. Therefore, a feasible contract must satisfy

$$\sum_q p_i(q|a)U(c(q)) - V(a) \geq \bar{U}. \quad (1)$$

The point of the moral hazard problem is to generate nontrivial dependence of consumption on output. But from what has been described to this point, there is little reason to expect such a nontrivial dependence. For example, if the agent is risk-averse, that is, U is strictly concave, then an optimal contract fully insures the agent against variations in output, paying the agent a constant wage with the principal absorbing all the risk in output.

Dependence of consumption on output is generated by assuming that the agent's action is *private information*; that is, the principal does not observe it. Consequently, the principal must set up the compensation schedule $c(q)$ to *induce* the agent to take the recommended action. Inducement means here that given $c(q)$ it is in the agent's best interest to take the recommended action. More formally, if the principal wants the agent to take a , then the contract must satisfy the following constraint:

$$\sum_q p_i(q|a)U(c(q)) - V(a) \geq \sum_q p_i(q|\hat{a})U(c(q)) - V(\hat{a}), \quad \forall \hat{a}. \quad (2)$$

A contract is called *feasible* if it satisfies constraints (1) and (2).³ The principal chooses a feasible contract that maximizes his utility. We can find such a contract by solving the following constrained maximization program:

³ In mathematics, a *program* refers to the problem of choosing an object that maximizes (or minimizes) an objective function subject to that object satisfying a set of constraints.

Moral Hazard Program

$$\max_{i,a,c(q)} \sum_q p_i(q|a)(q - c(q))$$

s.t. (1) and (2).

Analysis

To keep the analysis simple, assume that there are only two actions, a_l and a_h , with $a_l < a_h$. The latter action gives the principal more expected output but gives the agent more disutility. Also assume that in the optimal contract the agent is supposed to take a_h . In this case, there is only one incentive constraint. It is

$$\sum_q p_i(q|a_h)U(c(q)) - V(a_h) \geq \sum_q p_i(q|a_l)U(c(q)) - V(a_l). \quad (3)$$

Taking the first-order conditions to the program, with (2) replaced by (3), gives

$$-p_i(q|a_h) + \lambda U'(c(q))p_i(q|a_h) + \mu(p_i(q|a_h) - p_i(q|a_l))U'(c(q)) = 0,$$

where λ is the Lagrangian multiplier on constraint (1) and μ is the multiplier on (3). Simplifying gives

$$\forall q, \quad \frac{1}{U'(c(q))} = \lambda + \mu \left(1 - \frac{p_i(q|a_l)}{p_i(q|a_h)}\right). \quad (4)$$

Equation (4) describes the relationship between $c(q)$ and the parameters of the problem. Each Lagrangian multiplier is nonnegative and will be positive if its corresponding constraint binds. These variables affect consumption but not as much as the *likelihood ratio* $\frac{p_i(q|a_l)}{p_i(q|a_h)}$ does.

The likelihood ratio determines how c changes with q . If it decreases with q then consumption increases with q . Inspection of (3) reveals why. When this ratio is low, a high level of consumption rewards the agent relatively more for taking a_h than for taking a_l . Conversely when this ratio is high, a low level of consumption punishes the agent relatively more for taking a_l than for taking a_h . The likelihood ratio determines when the principal should use the carrot and when he should use the stick.

Conditions under which the likelihood ratio is decreasing in q include the normal distribution and some others. (For more information see Hart and Holmström [1987] and Jewitt [1988].) Still, most distributions do not satisfy this monotone likelihood property. This lack of robustness has always been a concern for this class of models because most contracts are monotonic as well as being simpler than those that solve the moral hazard program. For example, many sharecropping contracts are linear in output, while salesman

are often paid a fixed wage plus a percentage of sales, sometimes with a bonus for hitting performance targets.⁴

By choosing i , the principal is essentially choosing these ratios and, in doing so, directly affects the severity of the incentive constraints. The above analysis suggests that the principal will want to make this ratio high for some outputs in order to use the stick and low for others in order to provide the carrot. If, as was argued earlier, the principal has some control over the properties of the technology, then this will strongly affect technological design as well as compensation schedules. The following examples are designed to explore this idea.

2. A SIMPLE EXAMPLE

This example examines the extreme case where the principal can *only* control the probability distribution of output for the low action. The only restriction on these probabilities is that the chosen distribution still needs to produce the same expected output. In this problem, it is assumed that the solution is such that the principal wants to implement the high action.

Each choice of the technology index i corresponds to a choice of the entire probability distribution over q given a_l . For this reason, it is convenient to drop explicit reference to i and just let the principal choose the entire function $p(q|a_l)$.

The programming problem for this example is

$$\max_{p(q|a_l) \geq 0, c(q)} \sum_q p(q|a_h)(q - c(q))$$

subject to the participation constraint

$$\sum_q p(q|a_h)U(c(q)) - V(a_h) \geq \bar{U}, \quad (5)$$

the incentive constraint

$$\sum_q p(q|a_h)U(c(q)) - V(a_h) \geq \sum_q p(q|a_l)U(c(q)) - V(a_l), \quad (6)$$

and the constraints on technology

$$\sum_q p(q|a_l) = 1, \text{ and} \quad (7)$$

$$\sum_q p(q|a_l)q = \bar{q}, \quad (8)$$

⁴ See Townsend and Mueller (1998), however, for a description of sharecropping contracts that are formally linear but in practice are much more complicated.

where \bar{q} is the expected output amount that the distribution must produce.

The first set of first-order conditions is identical to (4). It is

$$\forall q, \frac{1}{U'(c(q))} = \lambda + \mu \left(1 - \frac{p(q|a_l)}{p(q|a_h)}\right). \quad (9)$$

The second set of incentive constraints comes from taking the derivative with respect to $p(q|a_l)$. Letting η be the Lagrangian multiplier on (7) and ν the multiplier on (8), these conditions are

$$\forall q, -U(c(q))\mu + \eta + \nu q \leq 0, \quad (= 0 \text{ if } p(q|a_l) > 0). \quad (10)$$

There is not necessarily an interior solution to this problem, so it is possible that (10) holds at an inequality.⁵

Equation (10) implies that $c(q)$ is monotonic over q such that $p(q|a_l) > 0$. Whether it is increasing or decreasing depends on the sign of ν . It is important to note that this does not mean that $c(q)$ is monotonically increasing everywhere. The result only applies for outputs for which $p(q|a_l)$ is chosen to be strictly positive. For outputs in which $p(q|a_l) = 0$, equation (9) implies that

$$\frac{1}{U'(c(q))} = \lambda + \mu. \quad (11)$$

Thus, consumption is the same for all of these values of output. Comparing (11) with (9) and noting that $\mu > 0$ reveals that consumption is higher for values of q that satisfy $p(q|a_l) = 0$ than for values that satisfy $p(q|a_l) > 0$.

These results are summarized by the following proposition:

Proposition 1 *The optimal contract is characterized by the following features: i) Consumption is a constant for all values of q such that $p(q|a_l) = 0$; ii) Consumption is monotonically increasing or decreasing in q over values of q such that $p(q|a_l) > 0$; iii) Consumption levels for q such that $p(q|a_l) = 0$ are higher than consumption levels for q such that $p(q|a_l) > 0$.*

At first glance, the monotonicity result is appealing because many contracts are monotonic. But since monotonicity only applies to outputs for which $p(q|a_l) > 0$, the degree of monotonicity will depend on the range of values of output with this property. As the following analysis demonstrates, there only needs to be two such outputs.

Proposition 2 *Let $Q = \{q | p(q|a_l) > 0\}$. There exists a solution in which there are no more than two outputs with $q \in Q$.*

Proof: The variables $p(q|a_l)$ only affect the right-hand side of the incentive constraint, (6), and the constraints on the technological design, (7) and

⁵ Notice that it has been implicitly assumed that $c(q)$ will be an interior solution.

(8). The lower the value of the right-hand side of (6) the less binding the incentive constraint will be and the better the consumption schedule that can be implemented. Consequently, for a given consumption schedule a solution will solve the following program for the principal,

$$\min_{\forall q \in Q, p(q|a_l) \geq 0} \sum_{q \in Q} p(q|a_l) U(c(q))$$

subject to

$$\sum_{q \in Q} p(q|a_l) = 1, \quad (12)$$

$$\sum_{q \in Q} p(q|a_l) q = \bar{q}. \quad (13)$$

This program is a linear program. If a solution exists to a linear program, which is true by assumption here, then a basic feasible solution exists. A basic feasible solution is one in which the number of non-zero valued variables is less than or equal to the number of constraints. In this problem there are only two constraints so there is a solution where all but two variables are necessarily zero. So if Q had more than two elements, their probabilities would be zero and they would not be in Q . **Q.E.D.**

The $p(q|a_l)$ function is set to make the utility from taking the low action as low as possible. Consequently, the program puts as much weight as possible on the lowest values of $c(q)$. Proposition 2 shows that there needs to be only two such points, which helps us characterize the compensation schedule. Since only two outputs are in Q , $p(q|a_l) = 0$ for all other values of q . So by (9) the value of consumption for these outputs is a constant. The compensation schedule in this problem is a wage except for the one or two outputs for which $p(q|a_l) > 0$. Which one or two outputs will be chosen cannot be determined without solving for the multipliers and all the other variables.

The goal of this exercise is to demonstrate the striking effect that technological choice by the principal can have on the properties of an optimal contract. Still, the optimal contract with its punishments on two levels of intermediate outputs does not look like contracts used in practice. Indeed, the analysis suggests that the incentive problem will be relatively minor since the low levels of consumption are only paid infrequently for the two outputs with $p(q|a_l) > 0$. Furthermore, the contract looks a lot more like a wage contract than the pay for performance contracts the theory was designed to describe. Fortunately, as the next section demonstrates, adding a small modification to the problem generates an optimal contract that is much more appealing on “realism” grounds. In particular, it will be monotonically increasing and relatively simple.

A Monotonicity Extension

Assume now that the agent can costlessly destroy output and that the principal does not know if he destroyed any output. All the principal observes is what is left. The agent does not consume any of the destroyed output. This assumption adds another source of private information to the problem; one that is easy to analyze. The ability to destroy output requires that the compensation schedule $c(q)$ be weakly monotonically increasing in output.⁶ Otherwise, the agent could destroy some output and claim the higher consumption. If there are n possible output realizations and q_j refers to the j th output, then this assumption requires adding the following constraints to the program.

$$\text{for } j = 1, \dots, n - 1, \quad c(q_j) \leq c(q_{j+1}). \quad (14)$$


An analysis of the amended program is not too different from that of the earlier program. The first-order conditions analogous to (9) only have some additional multipliers for the monotonicity constraints. The remaining first-order conditions are the same as (10).

The addition of the monotonicity constraints prevents the solution to the earlier program from being optimal. If the agent produced the one or two outputs that correspond to $p(q|a_l) > 0$ —the outputs with the lowest consumption—he could simply destroy some of the output and receive a higher level of consumption that goes with his new lower output.

The optimal contract to the program with the monotonicity constraints retains some of the same features as the solution to the earlier program. There are still only one or two outputs for which $p(q|a_l) > 0$. If there are two such outputs, they split the contract into three distinct regions: one region less than the lower of the two outputs, a middle region between the two outputs, and a third region above the higher output. In the lower range consumption is a constant and $p(q|a_l) = 0$, in the second region consumption is also a constant but higher than the consumption of the first point as well as higher than the consumption in the first region, and in the third region consumption is yet again a constant but at an even higher level. The contract is a step function with three steps. It resembles a contract with a wage and two levels of bonuses (or, a contract with a wage and one bonus level and one lower wage level for poor performance, etc.). Figure 1 illustrates. The point is that the contract keeps the desired monotonicity property and is relatively simple.

Proposition 3 *An optimal contract to the program with the monotonicity constraint is characterized by a compensation schedule that is: i) monotonic; ii) characterized by the three regions described above; and iii) consumption for*

⁶ Adding this source of private information to the standard model is enough to generate monotonic consumption but the optimal contract can still be complicated with lots of contingencies. As will be shown, the combination of technological choice and monotonicity also simplifies the contract in important ways.

Figure 1 Optimal Compensation Schedule

Notes: Example of what an optimal compensation-sharing rule might be. It is broken into three regions with a linear portion in each region. Consumption is weakly monotonically increasing.

each of two outputs that separate the two regions is equal to the consumption in one of the adjacent regions.

Proof: *i)* Follows directly from the monotonicity constraints (14). To prove *ii)*, use the same argument as before to argue that there are only two output levels for which $p(q|a_l) > 0$. These points are the boundaries. Below, between, and above, there is full insurance within each region because if there was not, consumption could be smoothed, which would deliver the risk averse agent the same utility at lower cost to the principal and not affect incentives. Finally, for *iii)*, if consumption of either of these two points was not equal to consumption in one of the adjacent regions, then consumption in the regions could be made closer together without altering the agent's utility. Analogous to the argument in *ii)*, this is a less expensive way for the principal to provide utility to a risk-averse agent. **Q.E.D.**

There are two lessons to this example. First, the modification generates contracts that are appealing on observational grounds. Second, the principal

Table 1 Probability Distributions of Output and Expected Output for Each Action

	0	1	2	$E(q a)$
a_l	1/3	1/3	1/3	1
a_h	ε	1/3	2/3- ε	5/3 - 2 ε

Notes: The parameter ε is nonnegative.

will try and design the technology so that if the agent slacks off (takes a_l), certain outputs will be very likely. In particular, he wants off-equilibrium probability distributions to be as revealing as possible when the agent slacks off.

3. LIQUIDATION EXAMPLE

The next example demonstrates that because of incentives, sometimes the principal prefers a less-productive technology, as measured by expected output. The reason for this counterintuitive result is that sometimes a less-productive technology alters the likelihood ratios in such a way that the incentive constraint is relaxed enough to outweigh the loss in output.

There are three outputs, $q_1 = 0$, $q_2 = 1$, and $q_3 = 2$. As before, there are only two actions, a_l and a_h . The production function is described in Table 1. The exercise is to assume that $\varepsilon \geq 0$ and then to vary it to illustrate how that affects the solution. Expected output for the high action is higher than that of the low output for any $\varepsilon < 1/3$.

The literal description of this problem is different from the standard model. Mathematically, however, it will be identical to the moral hazard program. The description is useful because it better motivates the example.

Now, assume that q represents an intermediate valuation of the project's long-term prospects. The principal does *not* observe q . There is a market, however, that trades securities based on the long-term value of the project. The market observes q and prices its securities accordingly. Alternatively, market participants have varying sources of information that are combined and communicated, however imperfectly, through the market price. Importantly, the principal observes the market price and makes an inference about the true q from it.

So far, the problem is no different than that of the standard model; the principal does not observe q directly but he can infer it from the market price of the security. Now, however, the principal has the option of liquidating the

Table 2 Probability Distributions of Output and Expected Output for Each Action

	0	1	2	$E(q a)$
a_l	0	$2/3$	$1/3$	$4/3$
a_h	0	$1/3 + \varepsilon$	$2/3 - \varepsilon$	$5/3 - \varepsilon$

Notes: Principal liquidates whenever the traded security indicates that $q = 0$ or $q = 1$.

project right after q is created (and observed and traded upon by the market). If he liquidates, the value of the project becomes one.

Markets, as always, are forward looking. In this context, this means that the market takes into account the effect of the principal's liquidation strategy on the value of the project. For example, if the strategy is to liquidate the project when the market price indicates that $q = 0$ or $q = 1$, then the market will trade the security at a price that indicates that $q = 1$. Indeed, under this liquidation strategy the security would never trade at a price of zero!⁷

The problem for the principal here is to decide on the best liquidation strategy. If he does not liquidate, the technology is the one described in Table 1. If the principal liquidates when the market price indicates $q = 0$ or $q = 1$, then the principal has essentially chosen the probability distributions to be those described in Table 2. No other feasible liquidation strategy is preferable, so the other ones are not explicitly considered.

In the liquidation case, $q = 1$ is not literally the amount produced since the agent may have produced $q = 0$, but liquidating turns it into an output level of one. Because the principal chooses whether to liquidate, the principal is essentially choosing between the probability distribution in Table 1 and the one in Table 2. Thus, the program has been mapped into the mathematical structure of the moral hazard program. Furthermore, for $\varepsilon > 0$ expected output in Table 1 is less than that in Table 2 for both actions. It is in this sense that the first technology is technologically inferior to the second technology. Yet, as we will shortly see, the first technology is sometimes superior when incentive considerations are taken into account.

The two problems can be compared by merely contrasting the incentive constraints. As before, assume that the principal wants to implement a_h . The no-liquidation incentive constraint, i.e., the one from choosing the technology

⁷ Related, a liquidation strategy of liquidating only when $q = 0$ would create an equilibrium existence problem. Under this strategy, $q = 0$ would never be observed because the price would be one. But if the price is one, the supervisor would never liquidate! Liquidating when the market price is zero or one avoids this circularity.

in Table 1, is

$$(1/3 - \varepsilon)(U(c(q_3)) - U(c(q_1))) \geq V(a_h) - V(a_l). \quad (15)$$

The liquidation incentive constraint, i.e., the one from choosing the technology in Table 2, is

$$(1/3 - \varepsilon)(U(c(q_3)) - U(c(q_2))) \geq V(a_h) - V(a_l). \quad (16)$$

The only difference between the two constraints is the replacement of $U(c(q_1))$ in (15) with $U(c(q_2))$ in (16). This should not be surprising. Output q_1 is not produced in the liquidation case, so it is not a factor in that case.

The striking feature of this example is that for small enough values of ε the principal prefers the no-liquidation technology even though the liquidation technology produces a higher expected output (for either action). The best way to see this is to analyze the limiting case where $\varepsilon = 0$. Consider the contract where the principal chooses the no-liquidation technology, sets $c(q_1)$ equal to its minimum level and sets $c(q_2) = c(q_3)$. Assuming that $c(q_1)$ can be set low enough so that the incentive constraint (15) holds, then this solution provides full insurance. Indeed, the incentive constraint (15) does not bind. Because the agent chooses a_h , the principal has not given any output up by not liquidating and the low consumption for producing $c(q_1)$ is a very powerful way of preventing the agent from choosing a_l .⁸ In contrast, if the principal liquidated the project with this consumption schedule, the agent would take a_l because he would never suffer the penalty from producing q_1 .

As ε gradually increases from zero, the principal starts foregoing output by not liquidating. Still, for small values of ε the incentive effect of setting $c(q_1)$ to a low value outweighs the loss in output as well as any cost to the agent from producing q_1 . (As ε grows, the optimal contract will no longer provide constant consumption over q_1 and q_2 , and $c(q_1)$ will increase.) Consequently, the “inferior” no-liquidation technology is preferred to the liquidation technology for incentive reasons. Of course, as ε gets large enough, the output effect will dominate the incentive effect and only then will the principal prefer the liquidation strategy.

4. CONCLUSION

This article worked through two examples to illustrate the importance of technological design on moral hazard. The first example gave the principal wide latitude in designing the probability distribution. It illustrated the mechanics of the approach and demonstrated that large effects on optimal compensation schedules were possible. The second example studied a problem in which the

⁸The likelihood ratio is infinite in this case. What this is indicating here is that consumption is not an interior solution and, in this case, is set to its lower bound.

liquidation strategy affected the informativeness of output. It demonstrated that sometimes the principal was willing to forgo output in return for a more informative distribution of output.

The main conceptual difficulty in these examples is determining how much latitude to give the principal in setting the probability distributions. What choice to offer the principal will depend on the application. The second example, with its problem of inferring true output from a market price, had a natural way of limiting the principal's control over the technology. Other applications will suggest different dimensions to this choice. Regardless of the application, what the analysis makes clear is that the technological design dimension to the moral hazard problem is an important one. It affects the surplus for the principal and the shape of the compensation schedule.

REFERENCES

- Black, Fischer, Merton H. Miller, and Richard A. Posner. 1978. "An Approach to the Regulation of Bank Holding Companies." *Journal of Business* 51 (3): 379–412.
- Hart, Oliver and Bengt Holmström. 1987. "The Theory of Contracts." In *Advances in Economic Theory, Fifth World Congress*. Cambridge, England: Cambridge University Press.
- Holmström, Bengt and Paul Milgrom. 1991. "Multitask Principal-Agent Analysis." *Journal of Law, Economics & Organization* 7: 24–52.
- Itoh, Hideshi. 1991. "Incentives to Help in Multi-Agent Situations." *Econometrica* 59 (3): 611–37.
- Jewitt, Ian. 1988. "Justifying the First-Order Approach to Principal-Agent Problems." *Econometrica* 56 (5): 1177–90.
- Lehnert, Andreas, Ethan Ligon, and Robert M. Townsend. 1999. "Liquidity Constraints and Incentive Constraints." *Macroeconomic Dynamics* (1): 1–47.
- Townsend, Robert M. and Rolf A. E. Mueller. 1998. "Mechanism Design and Village Economies: From Credit to Tenancy to Cropping Groups." *Review of Economic Design* 1 (1): 119–72.