

A Primer on Moral-Hazard Models

Edward S. Prescott

Moral Hazard — *The effect of insurance on insureds' behavior.*

Moral hazard, a long-time concern in the insurance industry, is increasingly being recognized as a concern in the regulation of banking and other financial industries. A classic example of its possible perverse effects is the selling of a fire insurance contract to a group of uninsured individuals. If the premiums are based on the actuarial data of this group's loss experiences, then the contract will be unprofitable. The reason for this loss is that with the introduction of fire insurance, insured people take fewer precautions than before against fires, raising losses above historical levels. It is this adverse effect of insurance on people's behavior that is moral hazard, and it is because of these adverse effects insurance contracts frequently contain clauses that attempt to minimize this behavior such as deductibles and copayments.

Ever since moral-hazard models were formalized mathematically in the early 1970s, applications of the models have burgeoned.¹ The models have been applied to just about any field where contractual relationships play an important role. In development economics, they have been used to study agricultural sharecropping contracts. In corporate finance, they have been used to study capital structure and executive compensation. In labor economics, they have been used to study employee compensation.

Moral-hazard analysis plays an important role in the theories of bank regulation and, more generally, financial regulation. For instance, government guarantees of bank deposits, be they explicit or implicit, reduce the incentive for depositors to monitor their banks. This lack of monitoring can create

■ The author would like to thank Andreas Hornstein, Brent Hueth, Jeff Lacker, Wenli Li, and John Weinberg for helpful comments. The views expressed in this article do not necessarily represent the views of the Federal Reserve Bank of Richmond or the Federal Reserve System.

¹ Some of the earliest work on this topic was done by James Mirrlees, one of the 1996 Nobel Laureates in Economics. See Dixit and Besley (1997) for a review of his contributions.

incentives for financial institutions to take on excessive amounts of risk. One striking example of the perverse risk-taking effects caused by poor regulatory design is the U.S. savings and loan crisis in the 1980s.² Moral hazard may also be a problem when financial institutions are large enough to be “too big to fail.” Some commentators fear that too-big-to-fail institutions have an incentive to take on excessive risk because these institutions (and even their “uninsured” creditors) will be bailed out in the event of a failure.³

Moral hazard is also assigned a prominent role in some analyses of causes of the recent Asian financial crisis. In these analyses, lax domestic financial regulation led to excessive risk-taking by Asian financial institutions. Furthermore, as argued by Calomiris (1998), the expectation of International Monetary Fund bailouts for developing country banking sectors gave foreign investors an incentive to finance risky activities.

Moral-hazard models are studied by analyzing constrained maximization programs, an important class of optimization problems. Though some of these programs are easy to study, the moral-hazard class is a particularly difficult one to analyze. Consequently, an extensive literature has developed that provides conditions that simplify the program. Unfortunately, these conditions are unappealingly restrictive.

To avoid these simplifying assumptions, we present another approach to analyzing moral-hazard models whereby we compute solutions to numerical examples. There are two advantages to this approach. First, it can be used to study problems that are not amenable to analytical methods. Indeed, the methods for computing numerical examples in this article succeed in cases in which the standard analytical simplification does not apply. The second advantage to computing solutions is that it gives one the ability to answer quantitative questions. For instance, the effect of deposit insurance on bank risk-taking is an inherently quantitative question. So is the question of whether a smaller amount of deposit insurance would result in a more efficient mix of risk and insurance. If one does not compute examples, the answers to these questions, as well as others, cannot be quantified.

In this article, we use linear programming as the computational technique for solving moral-hazard programs. Linear programming has been widely applied in management science, operations research, engineering, industry, and economics. Much is known about this class of programs, and practical algorithms for computing solutions to them have existed since the 1940s. Today,

² See Kareken (1983) for an early warning about this crisis. See also Benston and Kaufman (1998), who not only discuss this episode but also the effectiveness of The Federal Deposit Insurance Corporation Improvement Act of 1991 (FDICIA), the reform that resulted from the episode.

³ For a recent statement of this view, plus a proposed solution, see Feldman and Rolnick (1998).

there are numerous commercial and publicly available software codes that implement them.

Developing the moral-hazard program as a linear program requires making two departures from its standard formulation. One change is minor but the other change, allowing randomization or lotteries in the contract, is not. The implications of making this latter change will be discussed later. It is important, however, to understand that the change is not merely a technical assumption. Rather, it is one that can be justified on economic grounds.

As we review the moral-hazard model in the following paragraphs, we discuss the role of private information and describe the moral-hazard program as it is usually seen in the literature. This section is self-contained and can be read as such if the reader is interested only in the basic intuition of moral-hazard problems.

Using the standard formulation as a starting point, the next section introduces lotteries. We discuss the economic reasons for using them and present the linear program. (Two appendices spell out the linear program in more detail: Appendix A contains a short review of linear programs, and Appendix B derives the linear programming representation of the moral-hazard program.) Next, we compute the solution to a bank regulation example followed by some concluding comments. Finally, Appendix C contains a list of the papers that formulate private-information problems other than the moral-hazard problem as linear programs.

1. THE MORAL-HAZARD PROBLEM

The moral-hazard problem usually is formulated in terms of a contract between a principal and an agent who “works” for him. The principal and the agent can be people or institutions. With regards to agricultural sharecropping, the principal is the landowner and the agent is the tenant. With regards to banking, the principal is the bank regulator and the agent is the bank. With regards to executive compensation issues, the principal represents the collective interests of the shareholders while the agent is the chief executive officer.

In the moral-hazard problem, the agent works on a project for the principal. The amount of work the agent performs affects the probability distribution of the project’s return. The problem is that the principal cannot monitor the agent’s work, so the agent’s effort is *private information*; that is, it is observed only by the agent himself.⁴ In some models, the agent’s amount of effort is not observed. In other models, precisely how the task is performed is not observed. Our model captures either specification, so we will refer generally to the agent taking an action, be it a level of effort or a specific task. In agricultural sharecropping arrangements, the action is the amount of effort the tenant applies to working

⁴ Actions that are private information are also sometimes called hidden.

the sharecropped land, and the output is the crop yield. In bank regulation, the action is the risk-return profile of the bank's investments, and the output is the bank's return. In the executive compensation example, the action is the amount of resources aimed towards increasing the corporation's profits rather than the CEO's own personal satisfaction (for example, expanding the size of the corporation or avoiding painful decisions like layoffs), and output is the corporation's profit. In each example, the agent takes actions, unobserved by the principal, that affect output.⁵

Moral-hazard models are normally developed so that there is a conflict between the agent and the principal over the action the agent should take. For example, the agent might prefer a low-effort action because he dislikes hard work, while the principal might prefer a high-effort action because it increases the expected output of a project. This conflict in and of itself does not cause moral hazard but it does so when combined with private information on the agent's action.

To understand the role of private information, it is helpful to first consider the opposite situation, that is, when both the principal and the agent observe the action taken (commonly called the full-information case). In this case the principal and the agent could simply make a contract that fixes the level of action the agent should take. The exact action they would agree upon depends on many factors, such as the outside opportunities of the agent's labor, the outside opportunities of the principal's investment project, and the amount of compensation the agent receives. Nevertheless, once the two parties agree upon an action, they can enforce it because both parties observe it.⁶

Returning to the private-information assumption, one would see that the principal and the agent could still write a contract specifying the agent's action. But in this case, how could the contract be enforced? After all, if the principal cannot observe the agent's action, how can he make sure the agent took it? Although the principal could ask, there is no guarantee that the agent would respond truthfully; the agent can always reply that he fulfilled his end of the contract, whether he did or not.

While the principal cannot make the agent work a specified action directly, he may be able to induce the agent to take the desired action using the

⁵ Not allowing the principal to observe the action at all is, admittedly, an extreme assumption. Landowners can inspect the fields, regulators can monitor asset quality. Still, these measures are not perfect. It is expensive to audit and difficult to interpret signals. The basic moral-hazard problem laid out here should be considered a starting point for detailed analysis of any particular contractual arrangement. Adding features like auditing to the moral-hazard problem can be, and has been, done. We will return to this issue later.

⁶ The private-information literature usually assumes that if both parties observe a variable, then they can write enforceable contracts on it. Clearly, contracting is not so simple in practice. The largest literature that tries to address this problem is the one on incomplete contracts. See Hart (1995).


information he does have, namely, the project's output. In moral-hazard problems, one usually assumes that the output is publicly observed, that is, seen by both the principal and the agent. It follows that the only device the principal can use to encourage the agent is to make compensation depend on output. The principal, of course, wants to do this in such a way that the agent willingly takes the contracted action. The idea is that actions can be contracted upon, but *only* if they are consistent with the agent's incentives as determined by the compensation schedule. As we will see, compensation schedules can be effective at inducing actions but not without a cost.

It might be helpful to return to the agricultural example. Think of a landowner who lets a tenant farm a piece of his land. If the landowner spends his time on other activities, such as farming other pieces of land he owns, then he cannot monitor the tenant's efforts on the plot. He does not observe how carefully the tenant weeds the plot. Nor does he observe how many hours a day the tenant works the plot. All he can do is infer from the plot's yield the care and effort the tenant put into working the land. But this inference is far from perfect. If the crop fails, it could be because the tenant did not properly apply himself. But, it could also be the result of disease, insect infestation, insufficient rain, or a host of other factors beyond the tenant's control. Disentangling the effects of factors under the agent's control (his effort) and factors not under his control (the weather) is the essence of the moral-hazard problem because it leads to a trade-off between insurance and incentives. The contract should insure the tenant against events beyond his control, but it should also provide him an incentive to do what he is supposed to. Often, these are conflicting goals.

Formal Development of the Model

There are five sequential stages to the model. First, the principal and the agent agree to a contract. Second, as part of this contract, the principal recommends an action and then the agent decides whether or not to take it. These two steps are separated so that it is clear that the agent is choosing his action at that point. Then, the output is realized and, finally, the principal compensates the agent. Figure 1 summarizes these steps in a timeline.

Figure 1 Timeline of Moral-Hazard Problem



Environment

There are three variables that matter in this problem. First, there is the action that the agent takes. We identify an action by a and restrict it to lie in a set A , where A is an interval. (For example, A could be the interval between 0 and 1.) The second variable is the output, which we call q . Possible outputs lie in the set Q . For simplicity, we assume that there are only a finite number of elements in this set. The final variable is consumption c , which is restricted to lie in the set C . Like A , it is an interval.

The output q is determined by the agent's choice of action and a random shock that occurs after the agent has taken his action. The principal observes the output but not the shock or the agent's action. The idea is that he cannot infer from the output alone how hard the agent worked. It is most convenient to drop explicit references to the shock and instead describe the relationship between the action and the output by the conditional probability $p(q|a)$. This function is the probability distribution of output given the action. Because it is a probability distribution, $\sum_q p(q|a) = 1$ for each a . Finally, for simplicity we assume at this time that each output is possible for each action, that is, $p(q|a) > 0, \forall a \in A, q \in Q$.

Preferences

The agent cares about his consumption and his effort. We write his utility function as $u(c, a)$. For the moment, we do not make any assumptions on the form this function takes. The principal only cares about the project's surplus, that is, $q - c$. Depending on the model, the surplus may be negative, that is, the principal pays the agent out of his own funds. The principal's utility function is $w(q - c)$.

Deterministic Contracts

In the standard formulation of the moral-hazard problem, the model is solved for the optimal *deterministic contract*. A contract consists of the action the agent is supposed to take and the compensation schedule, that is, consumption as a function of output. The term "deterministic" refers to an assumed property of the contract, namely, that no randomization is allowed in the contract's terms. What we mean precisely by randomization is described in the next section.

Definition 1 *A deterministic contract is a recommended action a and an output dependent compensation schedule $c(q)$.*

The Approach

Our goal is to find one of the best feasible contracts that satisfies some criterion. Economists do this by solving a constrained-maximization program.

These programs consist of an objective function and a set of constraints. The objective function ranks alternative contracts according to some criterion. The constraints describe the set of contracts that are feasible. In this problem, the constrained-maximization program represents the problem facing a principal who is trying to determine the best feasible contracts to give the agent.

Objective Function

In moral-hazard problems, it is usually assumed that the principal owns the project and designs the contract. The objective function then is the principal's utility. It is written

$$\sum_q p(q|a)w(q - c(q)).$$

Participation Constraint

The first constraint is the participation constraint; it is also sometimes called an individual rationality constraint. The variable U represents the value of outside opportunities to the agent, opportunities that are not usually explicitly modeled. The constraint represents the idea that the principal-agent relationship does not exist in a vacuum. Since the agent has other activities that he can do, he will only sign the contract if it is at least as good as the best of these outside opportunities. Though we do not explore the issue in this article, the level of U can have a strong effect on the optimal contract. The higher U is the less surplus the principal will be able to get from the project. The participation constraint is

$$\sum_q p(q|a)u(c(q), a) \geq U. \quad (1)$$

Incentive Constraints

Incentive constraints are the formal method of accounting for private information. To see how moral hazard restricts the set of feasible contracts, consider the following problem. Assume that the principal is risk-neutral, that is, $w(q - c) = q - c$, and that the agent's preferences are separable $u(c, a) = U(c) - V(a)$. We assume that $U(c)$ is concave ($U'(c) > 0$, $U''(c) < 0$) indicating that the agent dislikes riskiness in consumption. We also assume that $V(a)$ is nonnegative and increasing in the action, indicating that the agent prefers a lower action to a higher one. (The action can be thought of as the level of effort here.)

Now consider the following compensation schedule, $c(q) = \bar{c}$, that is, consumption is independent of the realized output. We assume that this $c(q)$, together with the assigned a , satisfy (1). Under this contract, the principal bears all of the risk. His consumption (the surplus) moves one-for-one with

fluctuations in the output. In contrast, the agent's consumption is unaffected by fluctuations in output because he receives a fixed payment. In fact, it is easy enough to show that this contract is the optimal one if the only constraint is (1). The first-order conditions on consumption are

$$\forall q, \frac{1}{U'(c(q))} = \lambda, \quad (2)$$

where λ is the Lagrangian multiplier on the participation constraint. Only a contract with constant consumption satisfies these constraints.

Now consider what happens if this contract is chosen (and the contract does not assign the lowest action to the agent), but the agent's action is private information. Because the principal does not observe the agent's action, he cannot make him work the contracted action. The agent will take the action that is in his best interest, which is defined by *his* maximization problem, that is, $\max_a \sum_q p(q|a)[U(\bar{c}) - V(a)]$. But because consumption does not depend on output, the agent's action does not affect his consumption! Consequently, he takes the action that gives him the least disutility rather than the action recommended by the contract. Thus, the example contract, despite its desirable insurance properties, is not feasible.

Economists capture the effect of private information by using *incentive* or *incentive-compatibility* constraints. These constraints are simply a way of recognizing that any feasible contract will have to be compatible with the agent's incentives. For a contract $(a, c(q))$ to be incentive compatible, it must satisfy

$$a \text{ solves } \max_{\tilde{a}} \sum_q p(q|\tilde{a})u(c(q), \tilde{a}). \quad (3)$$

This constraint just says that the agent will take the action that is in his own best interest as determined by the compensation schedule $c(q)$.

Another way to write this constraint, one which will be more convenient later, is to express it by a direct pairwise comparison between taking the recommended action and taking all other actions. This representation is

$$\sum_q p(q|a)u(c(q), a) \geq \sum_q p(q|\hat{a})u(c(q), \hat{a}), \quad \forall \hat{a} \in A. \quad (4)$$

As before, this constraint says action a must be optimal from the agent's perspective. Notice that there are many constraints.

We can now proceed to the constrained-maximization program. To repeat, the problem is to find one of the best feasible deterministic contracts. For a contract to be feasible, it must satisfy both the participation and the incentive-compatibility constraints. An optimal deterministic contract is a solution to the program below.

Program with Deterministic Contracts

$$\max_{a, c(q)} \sum_q p(q|a)w(q - c(q))$$

s.t. (1) and (4).

Properties of Solutions

This program is surprisingly difficult to analyze. The problem is that if A is a continuum, then there are a *large* number (a continuum) of incentive constraints. To put these constraints into a manageable form, researchers normally try to substitute the first-order condition (FOC) from the agent's problem (3) for (4). When the agent's preferences are separable, that is, $u(c, a) = U(c) - V(a)$, the first-order condition is $\sum_q p_a(q|a)U(c(q)) = V_a(a)$. The programming problem with this constraint is much easier to analyze than the programming problem with (4).

Unfortunately, this "first-order approach" is not valid in general. First-order conditions are only sufficient for describing an optimum of a concave function. There is no guarantee, however, that at the optimal $c(q)$ the function $\sum_q p(q|a)u(c(q), a)$ is concave in the action.

What researchers have attempted to do is to find conditions that make this substitution valid. These conditions, however, are restrictive. In particular, the most commonly used assumptions are that the agent is risk-averse, his preferences are separable, the principal is risk-neutral, and the technology $p(q|a)$ satisfy the monotone likelihood ratio property (MLRP) and the convexity of the distribution function condition (CDFC).⁷

The MLRP condition is $\frac{p_a(q|a)}{p(q|a)}$ increasing in q . This condition guarantees that output is increasing stochastically in effort, that is, higher output is more likely for higher effort than lower efforts. If $P(q|a)$ is the cumulative distribution function, then CDFC is

$$P(q|\alpha a + (1 - \alpha)a') \leq \alpha P(q|a) + (1 - \alpha)P(q|a'), \quad \forall a, a' \in A, \quad \forall \alpha \in (0, 1),$$

and for all q . This condition provides a form of diminishing returns in effort. Both of these conditions are rather restrictive, and many natural technological specifications do not satisfy them both. For example, the distribution $q = a + \theta$, where θ is normally distributed, satisfies MLRP but does not satisfy CDFC.

When the first-order approach is valid, the optimal contract can be fairly well characterized. For example, if the previous assumptions on preferences and

⁷ For more details than are presented here, see Hart and Holmström (1987). Rogerson (1985) contains an important early proof of this result. Jewitt (1988) extends these results for the case of a risk-neutral principal, and Alvi (1997) extends these results for nonseparable preferences.

technology holds, then the first-order condition to the constrained-maximization problem (not the agent's subproblem) is

$$\frac{1}{U'(c(q))} = \lambda + \mu \frac{p_a(q|a)}{p(q|a)}.$$

The variables λ and μ are the Lagrangian multipliers on the participation constraint (1) and the incentive constraint (FOC version), respectively. Compare this with (2), the FOC for the full-information program. The only difference is the term $\mu \frac{p_a(q|a)}{p(q|a)}$, which is the effect on the solution from adding the incentive constraints. Now, because of private information, consumption does depend on output. In particular, one can easily verify that when MLRP is satisfied, the term $\frac{p_a(q|a)}{p(q|a)}$ is increasing in q . Because the Lagrangian multipliers are nonnegative in this problem, this property implies that consumption is increasing in output. Unfortunately, most technologies do not satisfy MLRP. Consequently, optimal consumption-sharing rules need not be monotonic.

An Example

The following executive compensation problem illustrates what a "typical" compensation schedule may look like.⁸ In this example, the owners of a bank (the principal) must devise a compensation schedule for their chief executive officer (the agent). The CEO is risk-averse and would rather apply his effort to activities that do not necessarily increase bank profits. To simplify the model, we assume that an increase in his action gives the bank higher expected profits but gives the CEO lower utility. Furthermore, we assume that the technology satisfies the MLRP.⁹


The solid line in Figure 2 shows the optimal compensation schedule as a function of the bank's profits. Compensation increases with bank profits as one would expect for a technology that satisfies MLRP. To illustrate the role of private information, we also report the optimal full-information compensation schedule, calculated from the solution to the full-information program. (That program is the private-information program without the incentive constraints.) Under full information, the risk-averse executive bears no risk.¹⁰ He is fully insured against all fluctuations and the risk-neutral bank absorbs all fluctuations in output.

Our comparison of the compensation schedules reveals one cost of private information. For incentive reasons, some insurance must be sacrificed. There

⁸ There is a literature that examines executive compensation using moral-hazard models. See Jensen and Murphy (1990) and Haubrich (1994).

⁹ The technology does not satisfy CDFC. We computed this example using the techniques described later in the article.

¹⁰ Any variation in the schedule is due to numerical approximation.

Figure 2 Optimal Compensation Schedules

is another cost, however, not illustrated in Figure 2: the program implements a lower action with private information than with full information. Apparently, in this example, it is too expensive to implement the optimal full-information action when there is private information.

2. CONTRACTS WITH LOTTERIES

In our review of the moral-hazard problem we restricted the contract space to deterministic contracts. In many cases, however, randomization in the terms of the contract may improve welfare.¹¹ This section covers these issues and explains how lotteries are a necessary component in developing the linear program.

There are two types of randomization that one may place in the contract. The first type involves making the recommended action a random. Instead of choosing a single action a , the principal chooses a probability distribution

¹¹ With randomization one can show by *Revelation Principle*-type arguments that no communication game between the principal and the agent can improve upon the direct mechanism considered in this problem. See Myerson (1982).

over all of the possible recommended actions. We will call this probability function $\pi(a)$. The second type of randomization is contained in the compensation schedule. Instead of choosing $c(q)$, the principal chooses $\pi(c|q, a)$, that is, a probability distribution of consumption given each realized output and each recommended action with positive probability. It is important to note that contracts with lotteries do not preclude deterministic contracts. Deterministic contracts are still feasible, but they are simply degenerate lotteries over the relevant sets.

Definition 2 *A contract with lotteries is a probability distribution over recommended actions, $\pi(a)$, and a probability distribution over consumption as a function of the output and the recommended action, $\pi(c|q, a)$.*

Economic Role of Lotteries

There has been little research on the role of lotteries in moral-hazard problems. What we do know is that under certain strong conditions, randomization is undesirable, and under certain weaker conditions, it is desirable. In general, analytical results concerning randomization appear to be quite difficult to derive. We will not attempt to summarize the results but will point out a few apparent ones. The reader interested in more detail can examine Arnott and Stiglitz (1988) and the citations contained within.

If the agent is risk-averse and his utility function is separable, and if the principal is risk-neutral, then we know that consumption lotteries are *not* optimal. To see this result, imagine a contract with a nondegenerate consumption lottery. Replace the consumption lottery with a deterministic compensation schedule that leaves the agent's utility unchanged for any realization of q and a . Because of concavity of the utility function, the schedule uses less consumption in expectation without violating the participation and incentive constraints. The reduced expected payment can be returned to the principal, increasing his utility. With nonseparable preferences, however, such similar conditions are much harder to provide. See Arnott and Stiglitz (1988) and the citations contained within.

Cases with optimal nondegenerate consumption lotteries require either a nonconvexity or nonseparability in preferences. If the agent is a risk-seeker, consumption lotteries (for reasons having nothing to do with incentives) will be valuable. Consumption lotteries may also be desirable if the agent's action affects his risk-aversion. Cole (1989), in a different model, provides one such example.

Action lotteries have been less systematically studied than consumption lotteries. Arnott and Stiglitz (1988) demonstrate that action lotteries will occur if in the space of contracts with deterministic actions, the principal's expected utility is nonconcave in the agent's expected utility. Unfortunately, conditions

under which this situation occurs are far from obvious. Lehnert (1998) and Prescott (1998) contain examples with action lotteries (or something similar).

To summarize, analytical results on the optimality of nondegenerate lotteries are difficult to derive. It seems then that computation of examples should play an important role in examining this issue.

Concerns with Lotteries

One common criticism of lotteries is that they are typically not found in the explicit terms of contracts. While this may be true, the complicated terms of optimal deterministic contracts are also typically not found in the explicit terms of contracts. This criticism then is not one aimed at the use of lotteries per se but at contract theory in general.

One response to the criticism is that explicit terms in contracts are not necessarily an accurate guide to its true terms. Frequently, there are implicit terms or contingencies in contracts. For example, Townsend and Mueller (1998) contains examples of implicit contingencies in sharecropping contracts.

Another response to the criticism is that lotteries may be indistinguishable from other state-contingent transfers and may represent unmodeled transactions. Cole and Prescott (1997) demonstrate that ex ante randomization, the sort defined by the action lotteries, need not be implemented through individual contracts but, instead, can be implemented through ex ante gambling that generates wealth inequality. In a developing economy context Lehnert (1998) argues that ex ante randomization may be implemented by financial intermediaries like Rotating Savings and Credit Associations (ROSCAs). ROSCAs frequently use lotteries to determine the order in which members receive funds. In both papers, there is a continuum of agents so the ex ante lotteries represent equilibrium *fractions* of the population. Interpreting the lotteries as equilibrium fractions of the population is a particularly intuitive interpretation that we will return to later in our example. It is also possible that the consumption lotteries may represent state-contingent transfers tied to exogenous events, transfers that one might interpret as part of a normal financial contract.

In some arrangements, lotteries are an explicit part of the economic mechanism. For example, the Internal Revenue Service randomly audits tax returns. Some indivisible goods or duties are assigned by lottery. For example, citizens are assigned to jury duty by lot. No doubt in the case of employment decisions sometimes the difference between a successful job candidate and an unsuccessful one is simply luck.

In my view, if we are to take the discipline of optimizing seriously, we should not restrict the contract space unless there is an economic argument for excluding these types of contracts.¹² This issue is not merely a technical

¹² See Boyd and Smith (1994), who argue that the gains from randomized auditing in resolving firm bankruptcy are miniscule.

one. In Lehnert (1998), policy prescriptions differ between an economy with deterministic contracts and one with lotteries. If for a particular problem one decides that the evidence does not support lotteries, then one needs to choose exogenous parameters so that the optimal contract is deterministic. The theory then provides an additional layer of requirements on the class of models we work with.

Making the Moral-Hazard Problem a Linear Program

We now show how to develop the moral-hazard program with lotteries as a linear program. This article does not deal with the topic of linear programming in any detail. There are plenty of well-written references on the topic. It does, however, contain a brief review of linear programs in Appendix A, summarized in the following paragraph.

A linear program is a constrained-maximization program that satisfies the following conditions:

1. The objective function and the constraints are linear;
2. There is a finite number of variables; and
3. There is a finite number of constraints.

By adding lotteries, condition 1 is satisfied. Lotteries have the further advantage that they are nonnegative by definition so the choice variables in the program are nonnegative, a requirement of the standard form of linear programming. (See Appendix A.) To satisfy conditions 2 and 3 another assumption is necessary, namely, that the sets C and A , in addition to the set Q , each contain a finite number of elements. For example, if A has two elements $\{a_1, a_2\}$, then, with respect to the recommended action, the principal chooses two variables, $\pi(a_1)$ and $\pi(a_2)$, the probability of a_1 and the probability of a_2 . In general, the elements in these sets—often called the grids—can be made arbitrarily large, providing an approximation to the continuum case.

The derivation of the linear program is contained in Appendix B, but to summarize here, the deterministic contract $(a, c(q))$ is replaced with the contract $(\pi(a), \pi(c|q, a))$. Next, the objective function and constraints are algebraically manipulated so that the choice variable is the probability distribution $\pi(c, q, a)$. The term $\pi(c, q, a)$ is the unconditional probability distribution of each possible consumption, output, and action triplet.

Usually, the new choice variable is the probability distribution of each point in the grid $P = C \times Q \times A$, though sometimes P is a subset of the grid. For expositional ease, we assume the former in this section. Since C , Q , and A all have finite numbers of elements, P has a finite number of points. For example, if $C = \{c_1, c_2, \dots, c_l\}$, $Q = \{q_1, q_2, \dots, q_m\}$, and $A = \{a_1, a_2, \dots, a_n\}$, then P has lmn elements. Each element is indexed by a (c_i, q_j, a_k) vector. One way to list all the elements is

$$\begin{aligned}
& (c_1, q_1, a_1), \\
& (c_2, q_1, a_1), \\
& \dots \\
& (c_n, q_1, a_1), \\
& (c_1, q_2, a_1), \\
& (c_2, q_2, a_1), \\
& \dots \\
& (c_l, q_m, a_n).
\end{aligned}$$

The choice variable $\pi(c, q, a)$ is an lmn -dimensional vector. The value of $\pi(c_i, q_j, a_k)$ is the probability of the (c_i, q_j, a_k) triplet occurring.

Moral-Hazard Program with Lotteries

$$\begin{aligned}
& \max_{\pi} \sum_{c,q,a} \pi(c, q, a)w(q - c) \\
& \text{s.t.} \quad \sum_{c,q,a} \pi(c, q, a)u(c, a) \geq U, \tag{5}
\end{aligned}$$

$$\sum_{c,q} \pi(c, q, a)u(c, a) \geq \sum_{c,q} \pi(c, q, a) \frac{p(q|\hat{a})}{p(q|a)} u(c, \hat{a}), \quad \forall a, (\hat{a} \neq a) \in A \times A, \tag{6}$$

$$\forall \bar{q}, \bar{a}, \quad \sum_c \pi(c, \bar{q}, \bar{a}) = p(\bar{q}|\bar{a}) \sum_{c,q} \pi(c, q, \bar{a}), \tag{7}$$

$$\sum_{c,q,a} \pi(c, q, a) = 1, \text{ and } \forall c, q, a, \quad \pi(c, q, a) \geq 0. \tag{8}$$

This program is the same as the deterministic program discussed earlier but includes lotteries. Indeed, except for the addition of two constraints, the programs are identical in structure. Again, the objective function is the principal's utility. Constraint (5) is the participation constraint. The incentive constraints are (6). On the right-hand side of (6) is the likelihood ratio $\frac{p(q|\hat{a})}{p(q|a)}$. This ratio is very important in private-information models because it influences the ability of the principal to reward and punish the agent. For example, if the ratio is low for some q , then high compensation would reward an agent for taking the recommended action more than it would reward him for taking the deviating action \hat{a} . Similarly, if the ratio is high for some q , then low compensation would punish an agent who takes the deviating action \hat{a} more than it punishes him for taking the recommended action a .

The last two sets of constraints are specific to problems with lotteries. They are not analogous to any constraints in the deterministic formulation.

Constraints (7) are the technology, or Mother Nature, constraints. As Appendix B describes in more detail, these constraints are added to the program so that despite choosing $\pi(c, q, a)$, the principal really only chooses the contract $(\pi(a), \pi(c|q, a))$; that is, the constraints ensure that a feasible $\pi(c, q, a)$ is consistent with $p(q|a)$. The final constraint (8) ensures that $\pi(c, q, a)$ is a lottery, i.e., it sums to one and each element is nonnegative.

It is straightforward to verify that this problem is a linear program. Lotteries deliver linearity of the objective function and the constraints. The grids then ensure that $\pi(c, q, a)$ is a finite-dimensional vector and that there are a finite number of constraints.

Computation

In general, the two limitations to computing linear programs are computer speed and memory. Roughly, the larger the dimensions of the problem, the more memory is required and the longer the problem takes to solve. Unfortunately, as the sizes of the grids increase, these problems quickly grow in size. For example, if l , m , and n are the number of elements in C , Q , and A , respectively, then the number of variables in the problem is lmn . The number of constraints grows similarly. There is one participation constraint, $n(n-1)$ incentive constraints, nm technology constraints, and one probability measure constraint. For example, if there are 10 actions, 20 outputs, and 50 consumptions, then the linear program has 10,000 variables and 292 constraints.

The size of the linear program that can be computed depends on the user's software and hardware. Although recent advances in computer technology have greatly increased the size of programs that may be computed, size limits may still arise. If size limits bind, the programmer may be forced to restrict the number of elements in the grids. Whether this is a serious issue for the results depends on the problem. For some problems, grids are the natural specification, such as when there are indivisibilities. For many problems, the grid does not need to be very fine.¹³

With commercial code or high-quality public domain code written in a compiled language, the biggest limitation on computation is probably memory. When this problem arises, the Dantzig-Wolfe algorithm can be used to com-

¹³ The discussion of the grid brings up the issue of *grid lotteries*. Grid lotteries are lotteries over adjacent points in the grid, be they consumption or action lotteries. They frequently appear when computing these problems, so it is important to understand that they have no economic content. They only reflect the approximation of a continuum inherent in the grids. For example, the executive compensation schedule presented in Figure 2 was computed as a linear program and the optimal solution contained some grid lotteries. This situation is most obvious in the full-information contract, where the compensation schedule is not quite horizontal. This deviation would disappear if the grid was made successively finer. Indeed, one strategy for computing these problems is to solve them first with a relatively coarse grid and then add grid points to areas where there is positive weight in the solution.

pute a moral-hazard program. This algorithm solves linear programs that have a special constraint structure, a structure that the moral-hazard program satisfies. Prescott (1998) solves several moral-hazard programs using this algorithm. He finds that for given memory limitations substantially larger problems can be computed using this method.

3. A BANK REGULATION EXAMPLE

In this example, we study a moral-hazard problem where the principal is a bank regulator and the agent is a bank. The bank is funded by insured deposits. Because depositors are insured, they are unconcerned about the bank's return. For simplicity, we drop all reference to them from the problem. Since depositors have no incentive to monitor, it is the job of the bank regulator to devise a regulatory regime.

The bank can engage in either an opaque investment strategy or a transparent one. The opaque strategy consists of investing in assets that are difficult for outsiders to evaluate and are potentially risky. Examples of opaque strategies include business loans or complicated derivative contracts. The transparent strategy consists of investing in safe assets that are easy for outsiders to evaluate. Examples could include holdings of Treasury bills or other money market assets. We assume that the bank may engage in only one investment strategy. It cannot split its assets between an opaque and a transparent strategy. Also, for simplicity we assume that the bank has a fixed deposit base so it cannot choose the amount of funds to invest.

The information assumptions on the bank's actions are slightly different from the standard formulation, although except for a minor modification to the incentive constraints, the program will not change. The strategy the bank engages in is public information, that is, both the bank and the regulator know whether the bank invests in opaque or transparent assets. But if the bank engages in the opaque strategy, there is some private information. Specifically, we allow the bank to choose the riskiness of its opaque strategy without letting the regulator observe this choice. In contrast, if the transparent strategy is chosen, we assume that there is no private information. Finally, the bank's return (net of payments to depositors) is public information.

This problem is designed to study two issues about bank regulation. First, the model can exhibit coexistence of multiple regulatory regimes, one for each possible investment strategy adopted by the bank. Second, it can be used to study the role of fines in mitigating excessive risk-taking. Fines and other penalties are an essential part of the precommitment approach, a recent proposal for regulating banks' trading accounts.¹⁴

¹⁴ See Kupiec and O'Brien (1995), Prescott (1997), and Marshall and Venkataraman (1998). This approach advocates letting banks choose their capital level but fining them if losses exceed

If the bank engages in the opaque investment strategy, we assume that the bank chooses only the variance of its returns. Unlike most problems, the bank's choice of action has no effect on the mean of returns. Also unlike most problems, the bank's utility is not directly affected by the action chosen. To create a conflict of interest between the regulator and the bank, we first assume that the bank has limited liability. Because of its limited liability the bank prefers a high variance return as its losses are limited to zero.¹⁵ Our second assumption that will guarantee a conflict of interest is to assume that there is a social loss to bankruptcy, an event defined as a negative return by the bank. The loss creates a dislike for variance on the part of the regulator who we assume is benevolent and wishes to minimize social losses. Our setup creates a trade-off between risk-seeking behavior by the bank (driven by the limited liability) and the social cost from poor realizations of the bank's investment. The difference between the trade-off here and the one described earlier in the executive compensation problem is that in the compensation problem there was a trade-off between insuring the executive and giving him incentives to take the desired action.

The regulator has two devices available to influence the bank's actions. First, it can regulate the investment strategy of the bank. The regulator by fiat can decide which lines of investment that the bank may engage in. Second, the regulator may impose return-dependent fines, though the regulator's ability to do this is limited.

We limit regulators' ability to levy fines in two ways. First, with limited liability, a fine cannot exceed the bank's return. Second, the amount of the fine that the regulator may impose in any state must lie between zero and an upper bound. Reasons for these bounds are not modeled, but they could exist for political reasons. For example, it might be difficult politically to make explicit payments (negative fines) to banks. Furthermore, we assume that fines are not merely a transfer; instead, we assume that a fraction of collected fines represent a social loss to society. Consequently, not only does the regulator want to prevent bankruptcy, but it also prefers not to impose fines.

Setting up the Linear Program

In this problem it is more convenient to put the level of fines in the grid rather than the bank's profit (consumption). Bank profits are easy enough to calculate, as they are the difference between the return and the fine, that is, $c = q - f$.

Grid

The bank can engage in either of two investment strategies. If the bank engages

capital.

¹⁵ Limited liability makes the bank a risk-seeker over certain ranges of returns.

in the transparent strategy, then we assume that there is only one action the bank can take. We write this set as $A_{tr} = \{0.0\}$. If the bank engages in the opaque investment strategy, its set of feasible actions is $A_{op} = \{0.4, 0.7, 1.0, 1.3\}$. The value of the action corresponds to the standard deviation of the project's return. To be consistent with previous notation we can write $A = (A_{tr}, A_{op}) = \{0.0, 0.4, 0.7, 1.0, 1.3\}$. Because the investment strategy is public information, the incentive constraints only apply to actions in the set A_{op} .¹⁶

It is easiest if we make the return grid depend on the investment strategy. We will make the extreme assumption about the transparent strategy that there is no variance in its return. In particular, we assume that $Q_{tr} = \{0.5\}$ and that $Q_{op} = \{-0.3, -0.25, -0.2, \dots, 1.7\}$.

Because both the principal's and the agent's utilities in fines are linear (they are both risk-neutral), it is necessary to have only two points in the fine grid, an upper and a lower bound. Lotteries over the bounds can obtain any intermediate level of fine without being an approximation. Indeed, the weighted sum of the bounds will be interpreted as the actual level of the fine.

By assumption the lower bound on fines is zero. But because of limited liability, the upper bound on fines must be less than the bank's return. The fine must also be less than the exogenously set upper bound on fines, which is 0.12 in this example. We capture these constraints by writing the fine grid as return dependent. It is $F(q) = \{0\}$ if $q < 0$ and $F(q) = \{0, \min\{q, 0.12\}\}$ if $q > 0$. The return and action grid for this problem is thus $(Q_{tr} \times A_{tr})$ followed by $(Q_{op} \times A_{op})$, that is, $\{(0.5, 0.0), (-0.3, 0.4), (-0.25, 0.4), \dots, (1.7, 1.3)\}$. The entire grid is then created by appending the appropriate $F(q)$ for each return, action pair.

Preferences

The bank is risk-neutral with limited liability so its preferences are $u(f, q, a) = \max\{q - f, 0\}$. Notice that the bank's utility is not affected, at least not directly, by the action it chooses. The regulator only cares about the social costs of bankruptcy and fines. Bankruptcy is defined as negative return, that is, if $q < 0$. The regulator's preferences are

$$w(f, q) = \begin{cases} -1 & \text{if } q < 0, \\ -0.5f & \text{otherwise.} \end{cases}$$

These preferences assume that it is costly to resolve bankruptcies and to collect

¹⁶ Here is the minor change from the standard formulation that we referred to earlier. Keeping the notation of the standard formulation, the incentive constraints are now

$$\sum_{c,q} \pi(c, q, a) u(c, a) \geq \sum_{c,q} \pi(c, q, a) \frac{p(q|\hat{a})}{p(q|a)} u(c, \hat{a}), \quad \forall a, (\hat{a} \neq a) \in A_{op} \times A_{op}.$$

Notice that there are no incentive constraints if the bank adopts the transparent strategy.

finer. By maximizing this utility function, the regulator is minimizing the social costs of bankruptcy and fines.¹⁷

Technology

The probability distribution of the bank's return when it adopts the transparent strategy is predetermined by our assumption that there is no variance in the return. The probability of $q = 0.5$ when $a = 0.0$ is 1.0. If the bank adopts the opaque strategy, then its action affects the variance of its returns. Let $f(q|a)$ denote a normal distribution with mean -0.7 and standard deviation a evaluated at q . The conditional probability of each output for the opaque strategy is

$$p(q|a) = \frac{f(q|a)}{\sum_{Q_{op}} f(q|a)}.$$

The denominator normalizes the function to sum to one, which we need to do because of the grid. The technology is probably better understood by considering Figure 3. Notice that the higher the standard deviation (the higher actions), the higher the probability of bankruptcy (the area under each curve for $q < 0$).

In this problem, the regulator dislikes bankruptcy and fines because of their social cost. Without limited liability, the bank would not care which action it takes since each has the same mean. With limited liability, however, the bank receives no disutility from negative returns, so it prefers high-variance actions. This preference is the source of the conflict between the bank and the regulator. The only tools available to the regulator in this problem are its limited ability to impose fines on the bank and its ability to order the bank to engage in either of the investment strategies.


Solution

The linear program was solved for the case where the bank receives $U = 0.60$ utils. For this parameterization, the optimal action for the bank is a lottery over the transparent and opaque strategies. The bank is assigned action $a = 0.0$ (the transparent investment strategy) with probability 0.47, and it is assigned action $a = 0.4$ under the opaque investment strategy with probability 0.53.

When the bank is assigned the transparent strategy, no fines are imposed because there are no incentive constraints for this strategy. When the bank is assigned the opaque strategy, however, the regulator imposes fines to prevent

¹⁷ Arguably, one would want to add a constraint that the regulator recovers enough resources in fines to compensate depositors in the event the bank loses money, that is, when returns are negative. In the interest of keeping the example as close as possible to the structure of the basic moral-hazard program, we left out this constraint. However, in the solution to the example computed later, fines are sufficiently high that in expectation they cover expected losses (even assuming that half of the fines are lost as part of the fine collection process).

Figure 3 Probability of Each Output Given Each Action for the Opaque Strategy



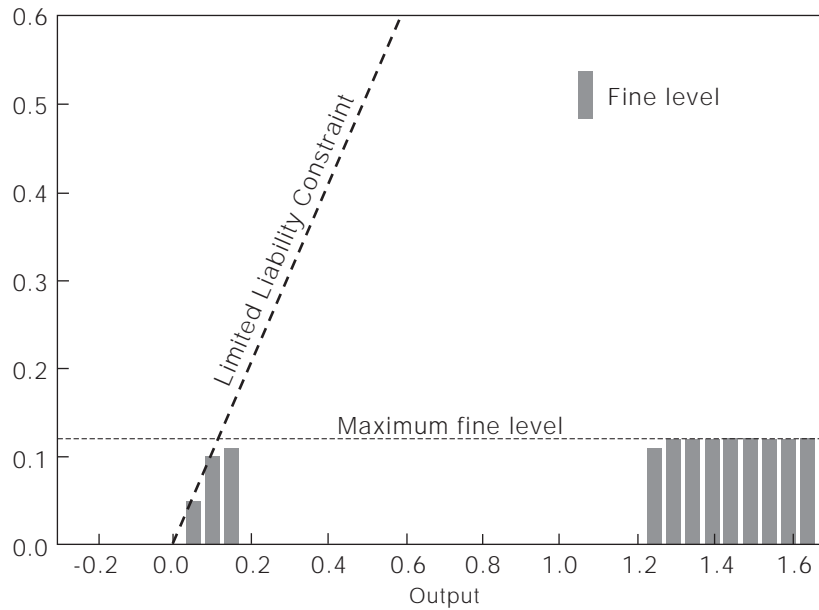
the bank from taking the higher risk actions. Figure 4 shows the optimal fine schedule if the bank is assigned the opaque strategy.

Because the regulator wants to minimize the amount of fines imposed, it wants to limit their use to when they are most effective. They are most effective if imposed on returns that are likely if the bank takes the high-variance action and unlikely if the bank takes the low-variance action (like it is supposed to). In this problem, low and high returns satisfy this condition. With limited liability, however, fines can hardly be imposed on the bank when it realizes a low return. In contrast, although the limited liability constraint is not binding for high returns, the fine ceiling binds instead.¹⁸

Interpretation

We interpret the investment-strategy lottery as representing the coexistence of two different regulatory regimes. To see this, imagine that instead of there

¹⁸One more point to note about this example. The binding incentive constraint in this example is not the one corresponding to the downward adjacent action (the discrete analog to the FOC on the agent's subproblem); instead it is the one corresponding to the highest variance action. Consequently, the first-order approach would not work in this example.

Figure 4 Optimal Fine Schedule

being one bank, there are a large number, or a continuum, of them. Each of these banks is identical to the one in the example and is treated identically, at least ex ante. Mathematically, the linear program is unchanged but now the lotteries represent equilibrium fractions of the population of banks. Under this interpretation, 47 percent of the banking sector would be engaged in transparent investment strategies and would not be subject to any fines. These banks resemble narrow banks. The remaining 53 percent of the banking sector would be engaged in the opaque investment but would be subject to a very strong regulatory regime.

Parallels to these results exist in U.S. financial regulation. Thrifts, credit unions, banks, and mutual funds all face different limitations on their investment strategies. For example, thrifts must hold a certain fraction of assets in housing and consumer lending. Credit unions are limited in their commercial lending. Money market mutual funds limit investment to money market instruments. Furthermore, all of these institutions have their own regulators and set of rules.

The investment strategy lottery is an example of an action lottery. As we discussed earlier, these lotteries can occur if there is a nonconvexity in utility space. In this example, banks get higher utility (and the regulator less utility) if the bank takes the opaque strategy than if it takes the risky strategy. The

lottery ensures that the bank gets its reservation utility in expectation. For different levels of reservation utility it is possible that all banks may engage in deterministic strategies.

In this example, banks receive higher utility from taking the opaque investment, so why would any bank take the transparent strategy? For the lottery to be implemented there needs to be some kind of control over the number of banks that may engage in the opaque investment strategy. In financial regulation, chartering may serve such a role. Many financial institutions must receive a charter to operate. Regulators' control of charters can limit the fraction of banks that can engage in certain strategies. A less command-and-control method would be to price deposit insurance so that banks pay the social costs of their investment strategy. Transparent banks would pay no deposit premiums, while opaque banks would pay enough to cover the expected social costs they cause, net their expected fine payments.

Another important feature of this example is the important role that fines play in mitigating risk-taking behavior. Like the action lottery, the optimal fine schedule (or at least a portion of it) parallels existing regulatory practice. The fines imposed for low returns are suggestive of the early closure rules under FDICIA. Under FDICIA, regulators can close banks when their capital drops to a low enough level. There is no direct parallel, however, for the high return fines. If we take the extreme view that existing regulatory rules cannot be improved upon, then we might say that the technology (the $p(q|a)$) in the example is incorrectly specified. And no doubt the technology is incorrectly specified; it was chosen to illustrate lotteries and the role of fines, not to match data. Another source of misspecification is that the model is missing crucial elements that would preclude the imposition of such fines. For example, well-run and innovative organizations should generate higher returns than other banks and such activities should not be discouraged.

Still, the high-return fines are suggestive of an interesting modification to this model. Dye (1986) adds to the moral-hazard model by also allowing the principal to verify ex post the agent's action but only at a cost. The problem for the principal is to determine the returns (if any) under which he should spend resources to verify the agent's action. For the parameters in our example, we conjecture that it would be optimal for the regulator to verify whenever there are high returns (the same returns under which it is optimal to impose fines in the present example). Indeed, the optimal regulatory regime would probably entail a combination of costly ex post verification with fines.¹⁹

¹⁹ Selective costly ex post verification is used in fair lending examinations by the Federal Reserve. Presently, one portion of the Federal Reserve's fair lending enforcement procedures starts with a statistical analysis of Home Mortgage Disclosure Act (HMDA) data. This data, which banks and other financial institutions are required by law to collect and report, contains information on each loan applicant. The information includes variables such as the loan applicant's race, the loan

4. CONCLUDING COMMENTS

The strength of computation lies in its ability to answer quantitative questions and its ability to investigate models that are difficult to analyze. The moral-hazard problem is one such model, of particular importance to many problems in economics, including financial regulation. Still, the basic model is limiting in many ways. Compensation is the only device available to the principal but in practice other mechanisms are important. For example, bank regulators receive reports, they monitor, and they observe signals. Many papers have investigated variants on the moral-hazard model, though not many have used the methods developed in this article. Appendix C lists the few such papers that use linear programming to solve examples. Since there are considerable difficulties in analyzing the basic moral-hazard model, one would expect that these difficulties would still exist in variants on the model and, therefore, numerical methods should be increasingly valuable as an analytical tool. With continued rapid advances in computer hardware and software, computation should become an increasingly effective way to study moral-hazard and other private-information problems.

APPENDIX A LINEAR PROGRAMMING

A linear program written in standard form is

$$\begin{array}{ll} \max & \mathbf{c}\mathbf{x} \\ & \mathbf{x} \geq 0 \\ \text{s.t.} & \mathbf{A}\mathbf{x} = \mathbf{b}, \end{array}$$

where \mathbf{c} is a $(1 \times n)$ vector, the choice variable \mathbf{x} is an $(n \times 1)$ vector, \mathbf{b} is an $(m \times 1)$ vector, and \mathbf{A} is an $(m \times n)$ matrix, often called the coefficient matrix and *not* related to the set of actions A .

For a more complete description, see any linear programming textbook such as Luenberger (1973) or Bertsimas and Tsitsiklis (1997). The important points to note are that the objective function and the constraints are linear in \mathbf{x} , that \mathbf{x} is a finite-dimensional vector, and that there are a finite number of

size, and the bank's decision. Some other information is reported but overall only a rather limited set of variables is collected. (Applicant credit history, for example, is not reported in the HMDA data.) If statistical analysis of an institution's HMDA data reveals a disparity in loan acceptances by race, then the analysis proceeds to a deeper and more costly review. For larger institutions, the next step in the analysis is to draw a sample of loan files from which examiners collect a much more detailed set of information on loan applicants than is available in the HMDA data. This detailed data set is then statistically analyzed, and if race appears to be statistically significant, then white-minority matches are generated and examiners look at the actual loan files to perform the deepest (and costliest) review. For more information see Avery, Beeson, and Calem (1997).

constraints. Neither the equality constraints nor the nonnegative values of \mathbf{x} are critical features. Problems with inequality constraints or variables that may be negative can be easily converted to standard form.

Two classes of linear programming algorithms are presently in use. The most common are simplex-based routines. Simplex-based algorithms move along the frontier of the constraint set until an optimum is reached. The simplex algorithm was developed in the 1940s by Dantzig and has proven to be an efficient method in practice for computing solutions to linear programs. More recently, interior point algorithms have been developed. These algorithms move through the interior of the constraint set. It is commonly reported that the simplex algorithm is faster for small problems, but for large, sparse (lots of zeros in the constraint matrix) problems interior point methods can be very effective.

APPENDIX B DERIVATION OF THE LINEAR PROGRAM

With lotteries, the choice variables are $\pi(a)$ and $\pi(c|q, a)$. When placed into the standard program, the program becomes

$$\begin{aligned} \max_{\pi(a), \pi(c|q, a)} \quad & \sum_{c, q, a} \pi(c|q, a) p(q|a) \pi(a) w(q - c) \\ \text{s.t.} \quad & \sum_{c, q, a} \pi(c|q, a) p(q|a) \pi(a) u(c, a) \geq U, \end{aligned} \quad (9)$$

$$\begin{aligned} \forall a \ni \pi(a) > 0, \quad & \sum_{c, q} \pi(c|q, a) p(q|a) u(c, a) \geq \\ & \sum_{c, q} \pi(c|q, a) p(q|\hat{a}) u(c, \hat{a}), \forall \hat{a} \in A, \end{aligned} \quad (10)$$

$$\forall q, \forall a \ni \pi(a) > 0, \quad \sum_c \pi(c|q, a) = 1, \pi(c|q, a) \geq 0, \quad (11)$$

$$\sum_a \pi(a) = 1, \text{ and } \pi(a) \geq 0. \quad (12)$$

This program is not a linear program because neither the objective function nor the participation constraint is linear.

To make it a linear program, we use the identity

$$\pi(c, q, a) = \pi(c|q, a) p(q|a) \pi(a) \quad (13)$$

and make the joint distribution $\pi(c, q, a)$ our choice variable.

Technology Constraints

By choosing the joint probabilities $\pi(c, q, a)$, the principal is choosing the conditional probabilities $\pi(c|q, a)$ and the unconditional probability $\pi(a)$, as is consistent with the formulation of the problem. The conditional distribution on the technology $p(q|a)$, however, is exogenous. We can keep this distribution exogenous by adding the following constraints

$$\forall \bar{q}, \bar{a}, \sum_c \pi(c, \bar{q}, \bar{a}) = p(\bar{q}|\bar{a}) \sum_{c,q} \pi(c, q, \bar{a}). \quad (14)$$

If a joint distribution $\pi(c, q, a)$ that satisfies (14) is chosen, then the principal has only implicitly chosen the contractual terms $\pi(a)$ and $\pi(c|q, a)$.

Incentive Constraints

The incentive constraints guarantee that the agent always takes the recommended action. Thus, the constraints are for any action recommended with positive probability, so for all a such that $\pi(a) > 0$,

$$\sum_{c,q} \pi(c|q, a)p(q|a)u(c, a) \geq \sum_{c,q} \pi(c|q, a)p(q|\hat{a})u(c, \hat{a}), \quad \forall \hat{a} \in A. \quad (15)$$

The left-hand side of (15) gives the utility the agent receives if he takes the recommended action. The right-hand side gives the utility the agent receives if he takes any other action. On the right-hand side, the deviating action \hat{a} enters the utility function and affects the probability distribution of output. It does not affect the compensation schedule because the principal uses the recommended action in the compensation schedule.

To make these constraints linear in $\pi(c, q, a)$, we first define $\pi(c, q|a) = \pi(c|q, a)p(q|a)$ as the conditional probability of the consumption-output pair (c, q) given that action a is recommended. Next, we make the substitution $\pi(c|q, a) = (\pi(c, q|a))/(p(q|a))$ into (15) to obtain

$$\sum_{c,q} \pi(c, q|a)u(c, a) \geq \sum_{c,q} \pi(c, q|a) \frac{p(q|\hat{a})}{p(q|a)} u(c, \hat{a}), \quad \forall \hat{a} \in A. \quad (16)$$

The term $\pi(c, q|a) \frac{p(q|\hat{a})}{p(q|a)}$ is the probability that the agent receives the pair (c, q) given that a was recommended but the agent instead takes action \hat{a} .

The final step in making the incentive constraints linear is to multiply both sides of (16) by the unconditional probability distribution $\pi(a)$ in order to express them in terms of $\pi(c, q, a)$. Rather than writing out the incentive constraints only for $\pi(a) > 0$, we write them as

$$\sum_{c,q} \pi(c, q, a)u(c, a) \geq \sum_{c,q} \pi(c, q, a) \frac{p(q|\hat{a})}{p(q|a)} u(c, \hat{a}), \quad \forall a, (\hat{a} \neq a) \in A \times A. \quad (17)$$

This set of constraints applies not only to actions assigned positive probability but also those assigned zero probability. Notice that the constraints hold trivially for actions such that $\pi(a) = 0$. This is quite convenient because normally we do not know beforehand which actions are assigned positive probability. Finally, the number of incentive constraints is finite because the number of elements in A is finite.

Probability Constraints

The last set of constraints ensures that $\pi(c, q, a)$ is a probability measure.

$$\sum_{c,q,a} \pi(c, q, a) = 1, \text{ and } \forall c, q, a, \pi(c, q, a) \geq 0. \quad (18)$$

Moral-Hazard Program with Lotteries

$$\begin{aligned} & \max_{\pi} \sum_{c,q,a} \pi(c, q, a)w(q - c) \\ \text{s.t. } & \sum_{c,q,a} \pi(c, q, a)u(c, a) \geq U, \end{aligned} \quad (19)$$

(14), (17), and (18).

Equation (19) is the participation constraint. This program is a linear program because it has a finite number of linear constraints, a linear objective function, and a finite number of variables.

APPENDIX C RELATED LITERATURE

The moral-hazard problem is not the only private-information problem that may be formulated in terms of a linear program. Myerson (1982) contains a general environment that incorporates several types of private-information problems. Prescott and Townsend (1984), Townsend (1987a), and Townsend (1993) formulate several private-information models as linear programs. Below we list a number of private-information models and papers that set them up as linear programs.

Costly State Verification

In this model, the private information is the agent's income. The agent sends a report on his income which the principal may audit at a cost. This model has been heavily used in the auditing, macroeconomics, and finance literatures. Townsend (1988) formulates this problem as a linear program. He also shows

that the problem with the restriction that auditing be a deterministic function of the report is a mixed integer linear program. See also Boyd and Smith (1994).

Moral Hazard with a Public Input

This model is just like the moral-hazard problem except that there is also a publicly observed input into production. See Lehnert (1998) or Lehnert, Ligon, and Townsend (forthcoming 1999).

Hidden Information

In this model, a shock to income or preferences is hidden information. Several well-known models such as the Mirrlees (1971) optimal tax problem or the Mussa and Rosen (1978) monopolist problem can be worked into the linear programming formulation. Prescott and Townsend (1984) have formulated a related insurance problem as a linear program.

Hidden Information with Moral Hazard

This problem combines a shock to preferences followed by a moral-hazard problem. See Myerson (1982) or Prescott (1996).

Repeated Private Information

This problem repeats the static private-information problem over multiple periods. Phelan and Townsend (1991), Lehnert, Ligon, and Townsend (forthcoming 1999), Ligon, Thomas, and Worrall (forthcoming 1999), and Yeltekin (1998) analyze variations on the problem using a combination of linear programming methods and dynamic programming.

Limited Communication

Some problems with private information and limited communication have been formulated as linear programs. See Townsend (1987b), Townsend (1989), and Prescott (1996).

Multi-Agent Problems

Problems where multiple agents have private information can be formulated in this way. See, for example, Townsend (1993), Prescott and Townsend (1999), and Yeltekin (1998).

Limited Commitment

Lacker (1989) formulates a two-agent limited commitment with costly enforcement as a linear program. Ligon, Thomas, and Worrall (forthcoming 1999) also contain limited commitment features.

REFERENCES

- Alvi, Eskander. "First-Order Approach to Principal-Agent Problems: A Generalization," *Geneva Papers on Risk and Insurance Theory*, vol. 22 (June 1997), pp. 59–65.
- Arnott, Richard, and Joseph E. Stiglitz. "Randomization with Asymmetric Information," *RAND Journal of Economics*, vol. 19 (Autumn 1988), pp. 344–62.
- Avery, Robert B., Patricia E. Beeson, and Paul S. Calem. "Using HMDA Data as a Regulatory Screen for Fair Lending Compliance," *Journal of Financial Services Research*, vol. 11 (April 1997), pp. 9–42.
- Benston, George J., and George G. Kaufman. "Deposit Insurance Reform in the FDIC Improvement Act: The Experience to Date," Federal Reserve Bank of Chicago *Economic Perspectives*, vol. 22 (Second Quarter 1998), pp. 2–20.
- Bertsimas, Dimitris, and John N. Tsitsiklis. *Introduction to Linear Optimization*. Belmont, Mass.: Athena Scientific, 1997.
- Boyd, John H., and Bruce D. Smith. "How Good are Standard Debt Contracts? Stochastic versus Nonstochastic Monitoring in a Costly State Verification Environment," *Journal of Business*, vol. 67 (October 1994), pp. 539–61.
- Calomiris, Charles W. "The IMF's Imprudent Role as Lender of Last Resort," *Cato Journal*, vol. 17 (Winter 1998), pp. 275–94.
- Cole, Harold Linh. "Comment: General Competitive Analysis in an Economy with Private Information," *International Economic Review*, vol. 30 (February 1989), pp. 249–52.
- , and Edward C. Prescott. "Valuation Equilibrium with Clubs," *Journal of Economic Theory*, vol. 74 (May 1997), pp. 19–39.
- Dixit, Avinosh, and Timothy Besley. "James Mirrlees' Contributions to the Theory of Information and Incentives," *Scandinavian Journal of Economics*, vol. 99 (June 1997), pp. 207–35.
- Dye, Ronald A. "Optimal Monitoring Policies in Agencies," *Rand Journal of Economics*, vol. 17 (Autumn 1986), pp. 339–50.
- Feldman, Ron J., and Arthur J. Rolnick. "Fixing FDICIA: A Plan to Address the Too-Big-To-Fail Problem," Federal Reserve Bank of Minneapolis *The Region*, vol. 12 (March 1998), pp. 3–22.
- Hansen, Gary D. "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, vol. 16 (November 1985), pp. 309–27.
- Hart, Oliver. *Firms, Contracts, and Financial Structure*. Oxford: Clarendon Press, 1995.

- _____, and Bengt Holmström. "The Theory of Contracts," in Truman F. Bewley, ed., *Advances in Economic Theory: Fifth World Congress*. Cambridge: Cambridge University Press, 1987.
- Haubrich, Joseph G. "Risk Aversion, Performance Pay, and the Principal-Agent Problem," *Journal of Political Economy*, vol. 102 (April 1994), pp. 258–76.
- Jensen, Michael C., and Kevin J. Murphy. "Performance Pay and Top-Management Incentives," *Journal of Political Economy*, vol. 98 (April 1990), pp. 225–64.
- Jewitt, Ian. "Justifying the First-Order Approach to Principal-Agent Problems," *Econometrica*, vol. 56 (September 1988), pp. 1177–90.
- Kareken, John H. "Deposit Insurance Reform or Deregulation Is the Cart, Not the Horse," Federal Reserve Bank of Minneapolis *Quarterly Review*, vol. 7 (Spring 1983), pp. 1–9.
- Kupiec, Paul H., and James O'Brien. "A Pre-commitment Approach to Capital Requirements for Market Risk," Working Paper 95–36. Washington: Board of Governors of the Federal Reserve System, July 1995.
- Lacker, Jeffrey M. "Limited Commitment and Costly Enforcement," Working Paper 90–2. Richmond: Federal Reserve Bank of Richmond, December 1989.
- Lehnert, Andreas. "Asset Pooling, Credit Rationing, and Growth," Finance and Economics Discussion Series 1998–52. Washington: Board of Governors, December 1998.
- _____, Ethan Ligon, and Robert M. Townsend. "Liquidity Constraints and Incentive Constraints," *Macroeconomic Dynamics* (forthcoming 1999).
- Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. "Mutual Insurance, Individual Savings and Limited Commitment," *Review of Economic Dynamics* (forthcoming 1999).
- Luenberger, David G. *Introduction to Linear and Nonlinear Programming*. Reading, Mass.: Addison-Wesley Publishing Company, Inc., 1973.
- Marshall, David, and Subu Venkataraman. "Bank Capital Standards for Market Risk: A Welfare Analysis," *European Finance Review*, vol. 2 (1998), pp. 125–57.
- Mirrlees, James A. "An Exploration in the Theory of Optimal Taxation," *Review of Economic Studies*, vol. 38 (July 1971), pp. 175–208.
- Mussa, Michael, and Sherwin Rosen. "Monopoly and Product Quality," *Journal of Economic Theory*, vol. 18 (August 1978), pp. 301–17.
- Myerson, Roger B. "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, vol. 10 (June 1982), pp. 67–81.

- Phelan, Christopher, and Robert M. Townsend. "Computing Multi-Period, Information-Constrained Optima," *Review of Economic Studies*, vol. 58 (October 1991), pp. 853–81.
- Prescott, Edward C., and Robert M. Townsend. "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," *Econometrica*, vol. 52 (January 1984), pp. 21–54.
- Prescott, Edward S. "Computing Moral-Hazard Problems Using the Dantzig-Wolfe Decomposition Algorithm," Working Paper 98–6. Richmond: Federal Reserve Bank of Richmond, June 1998.
- _____. "The Pre-Commitment Approach in a Model of Regulatory Banking Capital," Federal Reserve Bank of Richmond *Economic Quarterly*, vol. 83 (Winter 1997), pp. 23–50.
- _____. "Communication in Private Information Models: Theory and an Application to Agriculture." Manuscript. Richmond, December 1996.
- _____, and Robert M. Townsend. "Boundaries and Connectedness of Economic Organizations." Manuscript. Richmond, February 1999.
- Rogerson, Richard. "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, vol. 21 (January 1988), pp. 3–16.
- Rogerson, William P. "The First-Order Approach to Principal-Agent Problems," *Econometrica*, vol. 53 (November 1985), pp. 1357–67.
- Townsend, Robert M. *The Medieval Village Economy: A Study of the Pareto Mapping in General Equilibrium Models*. Princeton, N.J.: Princeton University Press, 1993.
- _____. "Currency and Credit in a Private Information Economy," *Journal of Political Economy*, vol. 97 (December 1989), pp. 1323–44.
- _____. "Information Constrained Insurance: The Revelation Principle Extended," *Journal of Monetary Economics*, vol. 21 (March/May 1988), pp. 411–50.
- _____. "Arrow-Debreu Programs as Microfoundations of Macroeconomics," in Truman F. Bewley, ed., *Advances in Economic Theory: Fifth World Congress*. Cambridge: Cambridge University Press, 1987a.
- _____. "Economic Organization with Limited Communication," *American Economic Review*, vol. 77 (December 1987b), pp. 954–71.
- _____, and Rolf A. E. Mueller. "Mechanism Design and Village Economies: From Credit to Tenancy to Cropping Groups," *Review of Economic Dynamics*, vol. 1 (January 1998), pp. 119–72.
- Yeltekin, Sevin. "Dynamic Principal-Multiple Agent Contracts." Manuscript. Stanford University, December 1998.