

The Form of Incentive Contracts: Agency with Moral Hazard, Risk Neutrality and Limited Liability *

Joaquín Poblete[†], Daniel Spulber[‡]

December 2011

Abstract

The analysis obtains a complete characterization of the optimal incentive contract in an agency model with moral hazard, risk neutrality, and limited liability. We introduce a "critical ratio" that is a function of the agent's effort and the random state of the world in the state-space setting. The critical ratio indicates the returns to providing the agent with incentives for effort in each state of the world. The agency contract takes the form of debt (a capped bonus) when the critical ratio is increasing (decreasing) in the state of the world. The critical ratio's properties replace the more restrictive assumptions commonly used in agency models, including the monotone likelihood ratio property (MLRP). The critical ratio has a simple decomposition: it is the product of the hazard rate of the state times the marginal rate of technical substitution between effort and the state, so its properties can be readily evaluated in most applied models in economics and finance. The critical ratio has a corresponding reduced-form property, which we term the decreasing hazard rate in effort property (DHREP). The critical ratio approach also yields insights into agency with adverse selection.

Key Words: Agency, Incentives, Contract, Moral Hazard, Debt, Limited Liability.

*We thank the coeditor David Martimort and two anonymous referees for their helpful and insightful comments that greatly improved the paper. We also thank Sandeep Baliga, Marco Ottaviani, Yuk-fai Fong, Yogmin Chen, Tom Gresik and Alessandro Pavan for interesting comments. Poblete is grateful to the Searle Center for Law, Regulation and Economic Growth for research support. Spulber is grateful for the support from a research grant from the Ewing Marion Kauffman Foundation on Entrepreneurship.

[†]London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom E-mail: j.j.poblete-lavanchy@lse.ac.uk

[‡]Elinor Hobbs Distinguished Professor of International Business, Management & Strategy, Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL, 60208. E-mail: jems@kellogg.northwestern.edu

1 Introduction

Debt and equity financing generate different types of agency costs by affecting performance incentives for managers and entrepreneurs. Jensen and Meckling (1976) argue that for a given project, equity financing gives managers incentives to shirk or to obtain perks from the firm, while debt financing provides incentives for more efficient performance.¹ Building on the insight that debt is better than equity for reducing moral hazard, Innes (1990) shows that, in risky environments, debt-style contracts are optimal for inducing agent effort when there is limited liability for both a risk-neutral principal and a risk neutral agent. These fundamental results raise the critical question: what are the general conditions under which debt financing and debt-style agency contracts are optimal? In this paper, we consider limited-liability, moral hazard contracts with risk-neutral agents and obtain highly general and natural conditions under which the optimal contract takes the form of debt. We fully characterize the optimal agency contract, and also give conditions under which the optimal contract takes the form of either a capped bonus or equity.

To address the design of the optimal agency contract, we introduce a *critical ratio* that provides an intuitive and easily-derived condition for characterizing the optimal contract. Intuitively, if we increase the marginal payment to the agent in the neighborhood of a particular outcome, this has benefits and costs for the principal. The benefit is that the agent's incentives to exert effort increase, and the cost is that the payments to the agent increase. *The critical ratio is the benefit from increasing marginal incentives for agent effort divided by the overall contract costs from increasing those incentives, so that the critical ratio measures the returns to providing incentives for effort in each state of the world.* The critical ratio therefore suggests how the parties should design a contract so as to motivate agent effort efficiently. Using the state-space or parametric representation of uncertainty, the critical ratio is a function of the state variable and the agent's effort and summarizes the relevant information about the probability distribution of the state variable and the form of the

¹ On the other hand, Jensen and Meckling (1976) point out that debt financing adversely affects the manager's choice of projects and entails monitoring and bankruptcy costs, arguing that the tradeoff between these incentive effects and the moral hazard effects of equity helps to explain why firms choose a capital structure that combines debt and equity financing.

production technology.² Applying the critical ratio, we show that the optimal contract is such that the sensitivity of the contract to the outcome is zero when the critical ratio is below a cutoff value and this sensitivity is one when the critical ratio is above that cutoff.

Our main contribution is a general characterization of the optimal contract for the principal-agent model with moral hazard, limited liability, and a principal and an agent who are risk-neutral. The critical ratio approach yields a complete classification of optimal agency contracts. When the critical ratio is increasing in the random state, the agency contract should associate greater incentives for effort with higher realizations of the state. Therefore, when the critical ratio is increasing in the random state variable, it follows that the optimal contract is debt, because the agent then retains the full marginal product of effort when the outcome of the project clears a threshold equal to the face value of the debt. Conversely, when the critical ratio is decreasing in the random state, the agency contract should associate greater incentives for effort with lower realizations of the state. It follows that when the critical ratio is decreasing in the state, the optimal contract is a capped bonus because the agent then retains the full marginal product of effort when the outcome of the project is less than the cap. When the critical ratio is constant in the random state, the parties are indifferent between contracts with the same payment at the lowest outcome and equity-style linear agency contracts are optimal, but so are many other contracts in this case.

The main results of the paper are as follows. First, we provide a complete characterization of the optimal agency contract and show existence when the critical ratio is increasing, decreasing, or constant. Our analysis strengthens considerably the characterization of optimal agency contracts by introducing assumptions that are much less restrictive than traditional assumptions such as the Monotone Likelihood Ratio Property (MLRP) and the Convex Distribution Function Condition (CDFC). In particular, we show that MLRP implies that the critical ratio is increasing in the state variable but an increasing critical ratio does not imply MLRP.

² In the state-space or parametric representation, the outcome of a task performed by the agent depends on the effort of the agent and the realization of a state variable. The principal only observes the outcome of the project, so that the outcome is a signal of the unobservable agent effort and the state. There is common knowledge regarding the probability distribution of the state variable, the form of the production technology, the agent's preferences, and the agent's cost of effort. The contract is a payment to the agent that depends only on the commonly-observed outcome.

Second, we relate our results for the state-space (or parametric) representation to the more standard reduced-form representation of agency.³ We show that the critical ratio is increasing in the state variable if and only if the reduced-form distribution over outcomes contingent on the agent's effort satisfies a Decreasing Hazard Rate in Effort Property (DHREP). This property states that the hazard rate of the outcome distribution should be decreasing in the agent's effort level. DHREP implies that the optimal agency contract takes the form of debt.

Third, we argue that it is very easy to determine whether the critical ratio is increasing, decreasing, or constant for a wide range of applied problems in economics and finance. We apply a state-space or parametric approach that specifies the production function that associates the state variable and the agent's effort with the outcome and also specifies the probability distribution function of the state variable. Then, we obtain a simple multiplicative decomposition of the critical ratio, which equals the hazard rate of the state variable multiplied by the marginal rate of technical substitution (MRTS) between the agent's effort and the state variable. Because of this decomposition, the properties of the critical ratio (increasing, decreasing, or constant in the state) depend on the properties of the hazard rate of the distribution of the state variable and the MRTS of the production function. Thus, an advantage of using the state-space approach is that it corresponds to many models in economics and finance with random states that affect preferences, initial endowments, technologies, information, exogenous public policies, and states of nature. The state-space approach to uncertainty also is useful because it is consistent with empirical analysis in economics and finance.

Fourth, we obtain a comparison between models of agency with moral hazard and models of agency with adverse selection. This is because the state variable in our state-space representation of the moral hazard model corresponds to private information in the adverse selection model.⁴ We show that in the local incentive compatibility condition with adverse selection, the distortion in the agent's choice of effort in comparison to the full information optimum is equal to the reciprocal of our critical ratio. We also show that the sum across agent's types of the reciprocal of our critical

³ The state-space approach in agency originates in Spence and Zeckhauser (1971), Ross (1973), and Harris and Raviv (1976) who apply a first-order approach with risk-averse agents. The reduced-form approach to agency originates with Mirrlees (1974, 1976).

⁴ There is a vast literature on agency models with adverse selection. For surveys and analysis, see Laffont and Tirole (1993), Laffont and Martimort (2002), and Bolton and Dewatripont (2005).

ratio exactly equals information rents in the corresponding adverse selection model. We find that our increasing critical ratio condition differs in an interesting way from the Spence-Mirrlees single-crossing condition that is standard in adverse selection models.⁵ When the agent's effort and the agent's type (or respectively, the state variable) are complements, concavity of the production technology in effort and the agent's type imply that the increasing MRTS in our setting and the Spence-Mirrlees single-crossing condition are equivalent. Although the agent's type (or respectively the state variable) are unobservable, it is desirable in this case to reward effort for higher types, because the higher-type agent's effort is more productive.

The optimality of debt-style contracts in our model is important because it shows that such contracts perform better under uncertainty than other contractual forms, including linear contracts such as equity-type sharing rules and more complex payment schedules based on outcomes. Our analysis confirms the argument of Jensen and Meckling (1976) who observe that debt strengthens incentives relative to equity because the entrepreneur keeps the whole marginal return to effort in states where the firm is solvent. Our analysis also confirms the optimality of risky debt contracts that was shown by Innes (1990) and Holmstrom and Tirole (1997); see also the extended discussion of Innes in Bolton and Dewatripont (2004). We strengthen the analysis of Innes (1990) by substantially generalizing the assumptions; in particular we do not assume MLRP. See also Jewitt et al. (2008), who consider the effects of imposing minimum and maximum limits on payments that represent limited liability and other constraints and find that payment limits can induce option-type contracts. In contrast to the present work, they assume that the agent is risk averse and impose various assumptions including MLRP.

Our analysis also is important because the properties of optimal contracts help to explain financial arrangements and contractual forms that are widely used in practice. With an increasing critical ratio, which we believe should be the most common case, the contract is debt because of limited liability and the need to shift incentives for performance to higher realizations of the state. Debt financing is widely used by firms, entrepreneurs, and households, including bank loans, credit arrangements, and corporate bonds. Debt-style contracts correspond to financial assets including

⁵ On the Spence-Mirrlees condition, see Mirrlees (1971) and Spence (1974).

securities and bonds whose features resemble options.⁶ Debt-style incentive contracts are widely used in compensation to provide incentives for employees and managers and call options are commonly used in executive compensation. Debt-style incentives also are used in public and private procurement and in most types of insurance contracts.

A decreasing critical ratio implies that the agency contract is a capped bonus because of limited liability and the need to shift incentives for performance to lower realizations of the state. In practice, bonus contracts often feature a cap on earnings (Healy, 1985, Arya et al. 2007). Jensen (2003) recommends linear bonus incentives, although without setting targets or imposing caps. Arya et al (2007) find that capping bonuses in compensation plans remains a puzzle and they suggest that bonus caps can help align incentives between owners and managers. According to our model, this type of contract should be observed in cases when effort and the state are substitutes, so that the agent's effort is particularly productive for low realizations of the state. Our results suggest that debt may be a more common contractual form because debt is consistent with the combination of normal inputs in production and many standard probability densities of the state.

The paper is organized as follows. Section 2 considers the basic model of agency and presents our main assumptions. Section 3 introduces the critical ratio and examines its implications for agency contracts. Section 4 derives the optimal contract and shows how the form of the contract depends on the properties of the critical ratio. Section 5 discusses the connection between our critical ratio and the Mirrlees reduced-form approach in moral hazard problems and the Spence-Mirrlees condition in adverse selection models. Section 6 concludes.

2 The Basic Model of Agency

Consider two risk-neutral economic actors who enter into a contract. The actor designated as the principal is an investor who provides financial capital and the actor designated as the agent is an entrepreneur who devotes effort to establishing and operating a firm. The project also can represent other situations: the agent performs a service under authority delegated by the principal, the agent

⁶ See Cox and Rubinstein (1985) on financial options and Dixit and Pindyck (1994) on real options.

is an employee who produces a good within a firm owned by the principal, or the agent is a producer who provides a good under a contract from the principal.

The principal provides the production technology for the project and the agent supplies productive effort, $a \geq 0$. The agent's action, a , also represents the cost of effort to the agent. The production technology given by

$$\Pi = \Pi(\theta, a), \quad (1)$$

where the outcome, Π , depends on the state, θ , and the agent's effort, a . The outcome represents an economic output, the agent's effort is an input, and the state of nature, although exogenous, also can be viewed as an input. The principal observes the outcome, Π , but cannot observe either the agent's effort, a , or the state, θ . The agent chooses the action, a , before the realization of the state, θ . The outcome Π can represent revenue or profit resulting from the realization of the state and the agent's effort. In many applications, the outcome Π is characterized as a signal of the underlying state and effort.⁷ The state, θ , is a random variable with density function $f(\theta)$, cumulative distribution function $F(\theta)$, and support $[0, \bar{\theta}]$, where $\bar{\theta}$ may be infinite. In a later section, we consider the relationship between the state-space formulation of the model and the standard reduced-form representation.

The state-space formulation allows specification of the properties of the production technology. Assume that $\Pi(\theta, a)$ is twice differentiable in a and θ , and let $\Pi_\theta(\theta, a) = \partial\Pi(\theta, a)/\partial\theta$ and $\Pi_a(\theta, a) = \partial\Pi(\theta, a)/\partial a$.

Assumption 1 *The production technology, $\Pi(\theta, a)$, is increasing in θ .*

This assumption allows us to define $\hat{\theta} = \hat{\theta}(\Pi, a)$ as the state that satisfies

$$\Pi(\hat{\theta}(\Pi, a), a) = \Pi, \quad (2)$$

for an effort level a , and a realization of output, Π .

Next, we require that the production technology exhibit diminishing marginal returns to effort.

⁷ See Spence and Zeckhauser (1971) and Conlon (2009).

Assumption 2 *The production technology, $\Pi(\theta, a)$, is increasing and weakly concave in a for $\theta > 0$. At the lowest state, 0, the outcome, $\Pi(0, a)$, is constant in a .*

The concavity assumption is not needed for our main results and only assumed to simplify the exposition in Corollary 2. Because the lowest outcome is constant in effort, the support for a distribution of outcomes starts at $\Pi(0, a)$.

The contract between the principal and the agent is fully described by a payment from the principal to the agent, w . The contract is based only on the outcome, Π , because the agent's choice of effort, a , and the realization of the random variable, θ , are not observable to the principal,

$$w = w(\Pi). \quad (3)$$

For ease of presentation and without loss of generality, the fixed cost of the project is normalized to zero. The entrepreneur's wealth, or initial endowment, is normalized to zero. The entrepreneur's liability constraint implies that the payment to the agent must be nonnegative, $w \geq 0$. The entrepreneur's limited liability prevents the entrepreneur from self-financing the firm or, equivalently, rules out the principal selling the task to the agent and achieving the first-best effort level. Note that the first-best effort level, a^{FB} solves $\max_a \int_0^{\bar{\theta}} \Pi(\theta, a) f(\theta) d\theta - a$. Let $\lim_{a \rightarrow \infty} \int_0^{\bar{\theta}} \Pi(\theta, a) f(\theta) d\theta - a < 0$ so that the first-best effort is finite.

For any realization of Π , the principal's net benefit from the project is equal to the outcome minus the payment to the agent,

$$v(w(\cdot), \Pi) = \Pi - w(\Pi). \quad (4)$$

The principal's expected net benefit given the form of the outcome function and the distribution of the random variable, θ , equals

$$V(w(\cdot), a) = \int_0^{\bar{\theta}} [\Pi(\theta, a) - w(\Pi(\theta, a))] f(\theta) d\theta \quad (5)$$

The entrepreneur has an opportunity cost $u_0 > 0$. The entrepreneur's net benefit is given by the payment net of the cost of effort,

$$u(w(\cdot), a, \Pi) = w(\Pi) - a - u_0. \quad (6)$$

Given the contract, w , the agent chooses effort to maximize his expected net benefit,

$$U(w(\cdot), a) = \int_0^{\bar{\theta}} w(\Pi(\theta, a)) f(\theta) d\theta - a - u_0. \quad (7)$$

The agent's effort is said to satisfy incentive compatibility if it maximizes his expected net benefit, $a \in \arg \max_x U(w(\cdot), x)$. The contract is individually rational for the agent if his expected net benefit is nonnegative, $U(w(\cdot), a) \geq 0$.

The optimal contract maximizes the expected net benefit of the principal, V , subject to incentive compatibility for the agent and individual rationality for the agent. Following Innes (1990), we require the net benefit of the principal and that of the agent to be nondecreasing in the outcomes Π . As Innes explains, making the principal's net benefit, $v(w(\cdot), \Pi)$, non-decreasing in Π rules out situations in which the principal wants to subvert the contract. This prevents the principal from sabotaging the project to avoid making payments to the agent after the agent has committed effort. The requirement that the principal's net benefit is non-decreasing in Π also rules out the situation in which the agent borrows money to inflate the outcome of the project and thereby increase his returns based on reported performance. The requirement that the agent's net benefit, $u(w(\cdot), a, \Pi)$, is non-decreasing in Π prevents the agent from sabotaging the task to expropriate the principal.

We define a feasible contract based on the requirements that net benefits are nondecreasing in outcomes and the agent has limited liability.

Definition 1 *A contract w is feasible if (a) the principal's net benefit and the agent's net benefit, $v(w(\cdot), \Pi)$ and $u(w(\cdot), a, \Pi)$, are non-decreasing in the outcome, Π , and (b) the payment to the agent is non-negative, $w \geq 0$.*

Requiring the agent's net benefit and the principal's net benefit to be non-decreasing in the outcome implies $\Pi' - w(\Pi') \leq \Pi'' - w(\Pi'')$, and $w(\Pi') \leq w(\Pi'')$ for outcomes $\Pi' \leq \Pi''$. Therefore, $0 \leq w(\Pi'') - w(\Pi') \leq \Pi'' - \Pi'$, which, taking limits, implies that $0 \leq w'(\Pi) \leq 1$, where $w'(\Pi) = dw(\Pi)/d\Pi$ and, in particular, $w(\Pi)$ is continuous and differentiable almost everywhere. It will be shown that the principal's optimization problem is linear in the slope of the contract so that constraints on the slope are necessary for the problem to be well defined.

The principal's problem of choosing an optimal contract subject to feasibility restrictions can be stated as follows,

$$\max_{w,a} \int_0^{\bar{\theta}} [\Pi(\theta, a) - w(\Pi(\theta, a))] f(\theta) d\theta, \quad (8)$$

subject to

$$a \in \arg \max_x U(w, x), \quad (9)$$

$$U(w, a) \geq 0, \quad (10)$$

$$0 \leq w'(\Pi) \leq 1, \text{ for all } \Pi, \quad (11)$$

$$w(\Pi) \geq 0, \text{ for all } \Pi. \quad (12)$$

The first constraint is the agent's incentive compatibility condition for the agent's choice of effort, a , and the second constraint is the agent's individual rationality condition. The third set of constraints results from feasibility and limits the slope of the contract. The constraint $w(\Pi) \geq 0$ represent the agent's limited liability.

Finally, to make sure that the problem has a solution, we make the following assumption

Assumption 3 *There exists a feasible contract w and an effort level a such that the agent's incentive compatibility and individual rationality constraints hold.*

The assumption will hold whenever the cost of the project and the agent's outside opportunity, u_0 , are small enough.

3 The Critical Ratio

This section provides an intuitive discussion of optimal contracts. The general idea is that contracts have two roles; to provide the agent with incentives for performance and to compensate the principal for his investment. The optimal contract is the one that combines these roles most effectively. The trade-off between these two roles is measured by the critical ratio, which is a benefit-cost ratio at a particular outcome. The benefit of increasing the slope of the contract, $w'(\Pi)$, by a small amount $\delta(w')$, in the small interval $[\Pi(\theta, a), \Pi(\theta + d\theta, a)]$, is $[f(\theta)\Pi_a(\theta, a)d\theta]\delta(w')$. This is the increase in marginal return to effort at the state associated with the outcome Π , weighted by the

probability that the outcome occurs given the agent's effort, a . The cost of increasing the slope of the contract is that the principal must pay more at all realizations higher than θ , so that the cost is $[(1 - F(\theta))\Pi_\theta(\theta, a)d\theta]\delta(w')$. The critical ratio in (13) below is the benefit of providing incentives for effort divided by the cost of providing those incentives so that the critical ratio, at a given state, is the expected return to providing incentives for effort in that state.

The critical ratio has a simple multiplicative decomposition that will be useful in economics and finance applications, as will be seen in a later section. By viewing the agent's effort, a , and the state, θ , as inputs in the production of the economic output, Π , we can refer to the ratio of the marginal product of effort to the marginal product of the state, $\frac{\Pi_a(\theta, a)}{\Pi_\theta(\theta, a)}$, as the marginal rate of technical substitution (MRTS) of agent effort for the state. This is the amount that agent effort must be reduced when the state is increased so as to keep the output constant. The ratio $\frac{f(\theta)}{1-F(\theta)}$ is the hazard rate of the state. Then, the critical ratio has the following decomposition.

Definition 2 *The critical ratio $\rho(\theta, a)$ is given by the product of the hazard rate of the state and the MRTS of the agent's effort for the state,*

$$\rho(\theta, a) = \frac{f(\theta)}{1 - F(\theta)} \frac{\Pi_a(\theta, a)}{\Pi_\theta(\theta, a)}. \quad (13)$$

We show that the optimal agency contract with moral hazard depends on the form of the critical ratio. The ratio measures the expected return to providing incentives in each state by determining the relation between incentives and the expected payoff when the slope of the contract, w' , changes. If the critical ratio is higher in some state than in others, then that state provides relatively greater incentives per unit of compensation and therefore that state is more efficient at inducing agent effort. An efficient contract should therefore have a higher slope in those states with a higher critical ratio. From the principal's perspective, among all contracts that implement the same effort level, the optimal contract is the one that gives the agent the smallest possible compensation. Since our assumptions require that the slope of the contract is not greater than one and not smaller than zero, the optimal contract from the principal's standpoint is a contract with a slope equal to one in states with high critical ratios and a slope equal to zero in states with low critical ratios.

We define *binary contracts* as those contracts whose slope is equal to one when the critical ratio is

above a certain threshold and equal to zero when the critical ratio is below that threshold. For a threshold $\lambda > 0$ and the agent's action, a , the set of binary contracts is given by

$$W_\lambda(a) = \{w : w'(\Pi(\theta, a)) = 0 \text{ if } \rho(\theta, a) < \lambda \text{ and } w'(\Pi(\theta, a)) = 1 \text{ if } \rho(\theta, a) > \lambda\}. \quad (14)$$

Let $w_\lambda(\Pi(\theta, a), a)$ designate an element of the set of contracts $W_\lambda(a)$.⁸ To see why there is a set of binary contracts, note for example that if $\rho(\theta, a)$ is constant and equal to λ , then all contracts are binary. The class of binary contracts is used to determine the design of the optimal contract.

4 The Optimal Contract

This section examines the optimal principal-agent contract and characterizes the form of the contract. The optimal contract between the principal and the agent involves choosing the payment function and the effort level that maximize the principal's expected net benefit over the set of feasible contracts. We first define a regularity condition for the critical ratio.

Definition 3 *The critical ratio $\rho(\theta, a)$ is said to be regular if $\rho(\theta, a) \geq \rho(\theta', a)$ implies that $\rho(\theta, a') \geq \rho(\theta', a')$ for every θ, a, θ' , and a' .*

The regularity condition holds if the critical ratio is increasing, decreasing, or constant in θ . We show later that the regularity condition holds for a wide range of models in economics and finance.

We now present our main result, which shows that when the critical ratio is regular, all optimal contracts are binary. The proof is based on establishing a single-crossing property of the agent's expected net benefit.

Proposition 1: *If the critical ratio is regular, then all optimal contracts are binary. That is they take the form $w_\lambda(\Pi; a)$, for some λ .*

Proof: Suppose to the contrary that the optimal contract, w , is not binary and implements the agent's effort level \hat{a} . Consider a binary contract, $w_\lambda(\Pi, \hat{a})$, that yields the same net benefits for the

⁸ With some abuse of notation, we use $w'_\lambda(\Pi(\theta, a); a)$ as the derivative of $w_\lambda(\Pi(\theta, a); a)$ with respect to the output, Π .

agent at that effort level, $U(w_{\hat{\lambda}}, \hat{a}) = U(w, \hat{a})$ and that gives the agent the same payment when the state is 0, $w(\Pi(0, \hat{a})) = w_{\hat{\lambda}}(\Pi(0, \hat{a}); \hat{a})$. For a proof that such a binary contract always exists, see Lemma 1 in the appendix. From the definition of $U(w, a)$, it follows that for any contract, w , and any effort, a ,

$$\frac{dU(w, a)}{da} = \int_0^{\bar{\theta}} w'(\Pi(\theta, a)) \Pi_a(\theta, a) f(\theta) d\theta - 1. \quad (15)$$

And therefore

$$\frac{dU(w_{\hat{\lambda}}, \hat{a})}{da} - \frac{dU(w, \hat{a})}{da} = \int_0^{\bar{\theta}} [w'_{\hat{\lambda}}(\Pi(\theta, \hat{a}); \hat{a}) - w'(\Pi(\theta, \hat{a}))] \Pi_a(\theta, \hat{a}) f(\theta) d\theta \quad (16)$$

Applying integration by parts to the definition of $U(w, a)$ implies that

$$U(w, a) = w(\Pi(0, a)) + \int_0^{\bar{\theta}} w'(\Pi(\theta, a)) \Pi_{\theta}(\theta, a) (1 - F(\theta)) d\theta - a - u_0. \quad (17)$$

Noting that $U(w_{\hat{\lambda}}, \hat{a}) = U(w, \hat{a})$, and $w(\Pi(0, \hat{a})) = w_{\hat{\lambda}}(\Pi(0, \hat{a}), \hat{a})$ implies

$$U(w_{\hat{\lambda}}, \hat{a}) - U(w, \hat{a}) = \int_0^{\bar{\theta}} [w'_{\hat{\lambda}}(\Pi(\theta, \hat{a}); \hat{a}) - w'(\Pi(\theta, \hat{a}))] \Pi_{\theta}(\theta, \hat{a}) (1 - F(\theta)) d\theta = 0$$

Subtracting $\hat{\lambda} [U(w_{\hat{\lambda}}, \hat{a}) - U(w, \hat{a})]$ from (16) gives

$$\frac{dU(w_{\hat{\lambda}}, \hat{a})}{da} - \frac{dU(w, \hat{a})}{da} = \int_0^{\bar{\theta}} [w'_{\hat{\lambda}}(\Pi(\theta, \hat{a}); \hat{a}) - w'(\Pi(\theta, \hat{a}))] [\rho(\theta, \hat{a}) - \hat{\lambda}] \Pi_{\theta}(\theta, \hat{a}) (1 - F(\theta)) d\theta. \quad (18)$$

The expression in (18) is positive by the definition of the contract $w_{\hat{\lambda}}$, the fact that w is not binary, and the feasibility constraint $0 \leq w'(\Pi(\theta, a); a) \leq 1$. These three conditions imply that $w'_{\hat{\lambda}}(\Pi(\theta, \hat{a}); \hat{a}) \geq w'(\Pi(\theta, \hat{a}))$ whenever $[\rho(\theta, \hat{a}) - \hat{\lambda}]$ is positive, and $w'_{\hat{\lambda}}(\Pi(\theta, \hat{a}); \hat{a}) \leq w'(\Pi(\theta, \hat{a}))$ whenever $[\rho(\theta, \hat{a}) - \hat{\lambda}]$ is negative.⁹ Therefore, we know that $\frac{dU(w_{\hat{\lambda}}, \hat{a})}{da} > \frac{dU(w, \hat{a})}{da}$ at \hat{a} .

We next show that $U(w_{\hat{\lambda}}, \cdot)$ and $U(w, \cdot)$ cross exactly once at \hat{a} . For this, consider a second effort level $\tilde{a} \neq \hat{a}$ with $U(w, \tilde{a}) = U(w_{\hat{\lambda}}, \tilde{a})$. Regularity of the critical ratio implies that at the effort level \tilde{a} , $w_{\hat{\lambda}}(\Pi(\tilde{a}, \theta); \tilde{a})$ is still a binary contract, so that there exists a $\tilde{\lambda}$ (different from $\hat{\lambda}$) such that

⁹ Since $w(\cdot)$ is not binary the inequalities must hold as strict inequalities in a set of positive measure.

$w'_\lambda(\Pi; \tilde{a}) = 0$ if $\rho(\theta, \tilde{a}) < \tilde{\lambda}$ and $w'_\lambda(\Pi; \tilde{a}) = 1$ if $\rho(\theta, \tilde{a}) > \tilde{\lambda}$. Since by assumption $U(w, \tilde{a}) = U(w_\lambda, \tilde{a})$, and by construction $w(\Pi(0, \tilde{a}); \tilde{a}) = w_\lambda(\Pi(0, \tilde{a}); \tilde{a})$, the same argument as in (18) gives

$$\frac{dU(w_\lambda, \tilde{a})}{da} - \frac{dU(w, \tilde{a})}{da} = \int_0^{\bar{\theta}} [w'_\lambda(\Pi(\theta, \tilde{a}); \tilde{a}) - w'(\Pi(\theta, \tilde{a}); \tilde{a})][\rho(\theta, \tilde{a}) - \tilde{\lambda}]\Pi_\theta(\theta, \tilde{a})(1 - F(\theta))d\theta. \quad (19)$$

Given the definition of $\tilde{\lambda}$, this expression is positive when evaluated at \tilde{a} . Therefore, whenever a contract gives the same net benefits as the binary contract, $U(w, a) = U(w_\lambda, a)$, we have that $\frac{dU(w_\lambda, a)}{da} > \frac{dU(w, a)}{da}$. It is well-known that if two functions satisfy the property that whenever they are equal, the slope of one of them is larger, then they cross at most once. The single-crossing property implies that $U(w_\lambda, a) = U(w, a)$ can only hold at \hat{a} . The single-crossing property of U has several implications. (i) If $a \in \arg \max_x U(w, x)$ and $a' \in \arg \max_x U(w_\lambda, x)$, then $a' > a$. (ii) Letting

$$U(w) = \max_a \int_0^{\bar{\theta}} w(\Pi(\theta, a))f(\theta)d\theta - a - u_0 = \int_0^{\bar{\theta}} w(\Pi(\theta, \hat{a}))f(\theta)d\theta - \hat{a} - u_0, \quad (20)$$

single-crossing at \hat{a} implies $U(w_\lambda) \geq U(w)$. (iii) If $\hat{a} \in \arg \max_x U(w, x)$ and $a' \in \arg \max_x U(w_\lambda, x)$, letting $V(w) = V(w, a)$ and $V(w_\lambda) = V(w_\lambda, a')$ implies that $V(w_\lambda) > V(w)$. This is because from (i), $a' > \hat{a}$, which implies that $V(w, \hat{a}) = V(w_\lambda, \hat{a}) < V(w_\lambda, a')$, where the last inequality uses $\Pi(\theta, a)$ increasing in a and $w'_\lambda \leq 1$ and strictly less than 1 in a set of positive measure.

From (ii) and (iii), the contract w_λ generates more net benefits for the principal than does the contract w , and satisfies the agent's individual rationality constraint. This is a contradiction. This implies that if there exists an optimal contract, it must belong to the class of binary contracts. \square

Proposition 1 shows how limited liability and nondecreasing net benefits for the principal and the agent smooth the form of the optimal contract. Consider instead a moral hazard model of agency with risk neutral parties without limited liability on the agent side or constraints on transfers to the principal. In that model, the incentive weight must be arbitrarily large on a measure zero set of outcomes, which would create existence and characterization problems. In contrast, the feasibility requirements in the agency model presented here and in Innes (1990) impose upper and lower bounds on transfers. These bounds lead to smoothing such that contractual incentives place weight on a positive measure set of outcomes. Then, the form of the critical ratio, which represents the returns to providing incentives for effort to the agent in each state, determines the form of the optimal contract.

When the critical ratio $\rho(\theta, a)$ is *increasing* in θ , Proposition 1 implies that an optimal contract must take the form of debt. With an increasing critical ratio, the optimal agency contract associates greater incentives for effort with higher outcomes. In our setting, an increasing critical ratio means that the optimal contract design sets the largest possible slope on the highest realization of the outcome, then sets the largest possible slope on the next highest realization of the outcome, and so on. The process continues until the relevant participation constraint is binding or the benefits of providing incentives are outweighed by the costs. As a consequence, the contract generates a threshold for outcomes such that the agent receives all of the output in excess of the threshold.¹⁰ Therefore, an increasing critical ratio leads to an optimal contract in the form of debt.

The face value of the debt specified by the optimal contract is a cut-off level, R , such that the agent receives no payment when the outcome is below the cut-off and the agent receives the difference between the outcome and the cut-off value otherwise. That is, when the agency contract takes the form of debt, the agent receives

$$w(\Pi) = \max\{\Pi - R, 0\}, \quad (21)$$

for some $R \geq 0$. In this case, the principal receives the outcome net of the payment, which is the minimum of the outcome and the cut-off level,

$$\Pi - w(\Pi) = \min\{\Pi, R\}. \quad (22)$$

With a debt-style contract, the principal can be viewed as the lender and the agent as the borrower. The agent is the residual claimant, receiving returns, however, only when the debt has been paid, with the payment to the agent equal to the outcome net of the debt.

Debt-style incentive contracts also can be viewed as call options for which the agent is the buyer and the principal is the seller of the option. Under a call option contract of the type used in financial markets, the buyer has the right but not the obligation to buy a commodity or financial asset from the seller at a particular time and at a strike price R . The buyer will choose to exercise a call option when the market price rises above the strike price specified in the contract. The buyer then benefits by earning the difference between the market price and the strike price. If we interpret the outcome

¹⁰ We thank a referee for suggesting this intuition.

Π as the market price of a security and the cut-off level R as the strike price, the contract gives the agent the value of the call on its expiration date. The agent as the owner of the call option has an incentive to increase effort so as to increase the returns obtained when the market price exceeds the strike price. Call options based on a firm's stock often are used in executive compensation because they give managers an incentive to increase the price of the firm's stock above the strike price.

When the critical ratio $\rho(\theta, a)$ is *decreasing* in θ , Proposition 1 implies that an optimal contract must be a capped bonus. The bonus is capped because a decreasing critical ratio implies that it is desirable to associate more incentives for performance with lower realizations of the outcome. Again, limited liability and nondecreasing net benefits for the agent and the principal constrain the weight that can be put on any realization of the outcome. The optimal contract design sets the largest possible slope for the contract on the lowest realization of the outcome, then sets the largest possible slope on the next lowest realization of the outcome, and so on. The optimal contract generates a threshold for outcomes such that the agent receives all of the output below the threshold and a payment equal to the cap when the outcome is above the threshold. Therefore, a decreasing critical ratio leads to an optimal contract that is a capped bonus.

When the critical ratio is decreasing in θ , the optimal contract specifies a cap on the bonus, R . Then, the agent receives the minimum of the outcome and the bonus cap,

$$w(\Pi) = \min\{\Pi, R\}, \quad (23)$$

for some $R \geq 0$. In this case, the principal's net return can be viewed as a debt to the agent,

$$\Pi - w(\Pi) = \max\{\Pi - R, 0\}. \quad (24)$$

With a capped bonus contract, the principal is the residual claimant, receiving returns only when the cap is exceeded, with the principal's earnings equal to the outcome net of the critical value. Also, the capped bonus contract can be viewed as the principal owning a call option on the outcome with the bonus cap as the strike price.

Propositions 1 and 2 show that the problem is more closely related to the critical ratio than it is to the likelihood ratio originally considered by Innes (1990). The critical ratio appears directly

from the integration by parts, while the likelihood ratio is only indirectly related to the principal's problem, requiring additional effort to relate it to the contract. This allows Proposition 1 to be simpler and more general than the result in Innes (1990). This is discussed more fully in the next section.

We now show that an optimal contract exists.

Corollary 1: *If the critical ratio $\rho(\theta, a)$ is increasing (decreasing) in θ for every effort level a , an optimal contract exists and takes the form of debt (a capped bonus).*

By Proposition 1, when the critical ratio $\rho(\theta, a)$ is increasing (decreasing) in θ , an optimal contract will always take the form of debt (a capped bonus). To prove the corollary, it only remains to show that among all debt (capped bonus) contracts, there exists an optimal one. This is done in the appendix. We show in a later section that an increasing critical ratio is a considerably weaker condition than MLRP, so this result extends the analysis in Innes (1990).

Next consider optimal contracts when the critical ratio $\rho(\theta, a)$ is *constant* in the state θ . Proposition 2 shows that any two contracts that implement the same effort level and have the same payment at the outcome $\Pi(0, a)$ provide the same expected benefits to the principal.

Proposition 2: *If the critical ratio $\rho(\theta, a)$ is constant in θ at a , the principal and the agent are indifferent between any two contracts that implement the effort level a , as long as $w(\Pi(0, a))$ is the same in both contracts.*

Proof: Integrating by parts gives

$$U(w, a) = w(\Pi(0, a)) + \int_0^{\bar{\theta}} w'(\Pi(\theta, a)) \Pi_\theta(\theta, a) (1 - F(\theta)) d\theta - a - u_0.$$

Dividing and multiplying by $\Pi_a(\theta, a)f(\theta)$ within the integral and recalling that $\rho(\theta, a)$ is constant implies that

$$U(w, a) = w(\Pi(0, a)) + 1/\rho \int_0^{\bar{\theta}} w'(\Pi(\theta, a)) \Pi_a(\theta, a) f(\theta) d\theta - a - u_0.$$

Any contract that implements the effort level a must satisfy the first-order condition

$$\int_0^{\bar{\theta}} w'(\Pi(\theta, a)) \Pi_a(\theta, a) f(\theta) d\theta = 1.$$

Replacing the first-order condition in the expected utility implies

$$U(w, a) = w(\Pi(0, a)) + 1/\rho - a - u_0,$$

which establishes that the agent's expected utility is independent of the contract as long as the contracts have the same payment at the lowest outcome, $w(\Pi(0, a))$. Because the principal and agent are both risk neutral, this also shows that the principal's expected utility is independent of the contract \square .

Intuitively, when $\rho(\theta, a)$ is constant, all states are equally efficient at providing incentives and therefore it is irrelevant which states are used to provide incentives. It can be checked easily that in this situation all contracts are binary.

Define linear contracts as having payments that are proportional to the outcome and $w(\Pi(0, a)) \geq 0$. Note that when the critical ratio is constant in θ , equity-style linear contracts satisfy the binary contracts definition. When the critical ratio is constant, the existence of an optimal linear contract follows from Proposition 2. Linear contracts are often used in practice, taking the form of sharecropping in agriculture, piece rates in manufacturing, cost sharing in procurement, sales commissions, and other types of proportional reward sharing between principals and agents. Linear contracts are featured often in studies of agency in economics and finance. Of course the present analysis suggests that linear contracts are rarely actually optimal.

Corollary 2: *If the critical ratio $\rho(\theta, a)$ is constant in θ for every effort level a , there always exists an optimal linear contract.*

Existence holds because assumption 2 states that production is concave in effort, and therefore linear contracts implement any effort level in $[0, a^{FB}]$, so that we can restrict attention to linear contracts without loss of optimality. Debt contracts also are optimal so that, with a critical ratio that is constant in θ , the principal and agent are indifferent between debt and equity contracts.

5 Discussion

In agency models with moral hazard, the principal observes the outcome that results from the agent's effort and the random state. The standard approach in the literature examines the probability distribution over outcomes induced by the agent's effort. Our discussion thus far has emphasized the state-space representation, although the results can be obtained using the Mirrlees reduced-form approach. Conlon (2009) provides a detailed analysis of the relation between the two representations. This section shows that our assumptions are considerably weaker than the standard assumptions in the moral hazard agency literature. Also, this section compares our assumptions with those in agency models with adverse selection.

5.1 Comparison with the Standard Moral Hazard Agency Model

The model can be expressed in terms of the reduced-form distribution as is now standard in the moral hazard literature. Observe first that for the outcome Π to be less than or equal to an outcome $\bar{\Pi}$, it must be the case that the state θ is less than or equal to $\hat{\theta}(\bar{\Pi}, a)$. Let $G(\Pi, a)$ be the induced distribution for the outcome Π , given action a ,

$$G(\Pi, a) = F(\hat{\theta}(\Pi, a)).$$

Define the following ratio using the reduced-form distribution,

$$\gamma(\Pi, a) = -\frac{G_a(\Pi, a)}{1 - G(\Pi, a)}. \quad (25)$$

This expression exactly equals the critical ratio when evaluated at the production function,

$$\rho(\theta, a) = \gamma(\Pi(\theta, a), a). \quad (26)$$

We introduce the following property of the reduced-form distribution.¹¹

¹¹ This definition is based on the property of probability distributions for which the hazard rate is decreasing in a parameter of the distribution (see Ross (1983)).

Definition 4 *The reduced-form distribution of the outcome, $G(\Pi, a)$, satisfies the decreasing hazard rate in effort property (DHREP) when the hazard rate for output in the reduced form is decreasing in the agent's effort, a , $\frac{\partial}{\partial a} \frac{g(\Pi, a)}{1 - G(\Pi, a)} < 0$ for all Π .*

It follows that the critical ratio, $\rho(\theta, a)$, is increasing in the state, θ , for all a , if and only if the reduced-form distribution of the outcome, $G(\Pi, a)$, satisfies DHREP. To see why, notice that the production function $\Pi(\theta, a)$ is increasing in the state and

$$\gamma_{\Pi}(\Pi(\theta, a), a) = -\frac{\partial}{\partial a} \frac{g(\Pi, a)}{1 - G(\Pi, a)}. \quad (27)$$

Therefore, DHREP in the reduced-form setting is equivalent to an increasing critical ratio in the state-space setting. Debt therefore is an optimal contract if the reduced-form distribution of output satisfies DHREP and if the production function satisfies our Assumptions 1 and 2.

The reduced-form and state-space representations are equivalent. The simplicity of Proposition 1, as compared to the more complicated analysis in Innes (1990), does not follow from our use of the state-space representation, but rather from the increasing critical ratio condition, or equivalently the DHREP. The form of the incentive contract in Proposition 1 can be obtained readily using the Mirrlees reduced-form and DHREP. To characterize the optimal agency contract using the Mirrlees representation, maximize $\partial E[w(\tilde{\Pi})]/\partial a = \int_0^{\infty} w(\Pi)g(\Pi, a)d\Pi = -\int_0^{\infty} G_a(\Pi, a)w'(\Pi)d\Pi = \int_0^{\infty} w'(\Pi)\gamma(\Pi, a)[1 - G(\Pi, a)]d\Pi$ holding constant $E[w(\tilde{\Pi})] = \int_0^{\infty} w(\Pi)g(\Pi, a)d\Pi = \int_0^{\infty} w'(\Pi)[1 - G(\Pi, a)]d\Pi$.

The increasing critical ratio property is less restrictive than the standard MLRP assumption.¹²

The MLRP condition can be stated as

$$\frac{\partial}{\partial \Pi} \left[\frac{g_a(\Pi, a)}{g(\Pi, a)} \right] > 0.$$

Whenever MLRP holds for a parameter, a , of the reduced form distribution, the hazard rate of the distribution is decreasing in that parameter.¹³ Therefore, MLRP implies DHREP, or equivalently,

¹² MLRP is used widely in moral hazard agency models, see for example Mirrlees (1976), Holmstrom (1979), Rogerson (1985), Jewitt (1988), Laffont and Martimort (2002), and Bolton and Dewatripont (2005).

¹³ See Ross (1983) and Shaked and Shanthikumar (2007). A related application of MLRP is given in Kim (1997); see also Park (1995). To see why, write the effect of effort on the reduced-form distribution as

MLRP implies that the critical ratio is increasing in θ . The converse is not true. This establishes that the critical ratio, rather than MLRP, determines the form of the optimal contract.

Proposition 3: (i) *MLRP implies that the critical ratio $\rho(\theta, a)$ is increasing in the state θ .* (ii) *MLRP is not necessary for the critical ratio $\rho(\theta, a)$ to be increasing in the state θ .*

The proof is given in the appendix. The proposition was suggested by a referee. DHREP, or an increasing critical ratio, implies that the reduced-form distribution satisfies first-order stochastic dominance (FOSD) in the agent's effort, $G_a(\Pi, a) < 0$.¹⁴ So, DHREP is a weaker condition than MLRP but a stronger condition than FOSD and provides exactly the amount of structure needed for the present agency problem.

As noted previously, our analysis does not require any regularity assumptions, such as CDFC, $G_{aa}(\Pi, a) \geq 0$.¹⁵ Such sufficient conditions are very difficult to satisfy in practice.¹⁶ Because our

follows,

$$G_a(\Pi, a) = - \int_{\Pi}^{\bar{\Pi}} \frac{g_a(\pi, a)}{g(\pi, a)} g(\pi, a) d\pi.$$

MLRP implies that

$$- \int_{\Pi}^{\bar{\Pi}} \frac{g_a(\pi, a)}{g(\pi, a)} g(\pi, a) d\pi < - \frac{g_a(\Pi, a)}{g(\Pi, a)} \int_{\Pi}^{\bar{\Pi}} g(\pi, a) d\pi = - \frac{g_a(\Pi, a)}{g(\Pi, a)} [1 - G(\Pi, a)].$$

Therefore, $G_a(\Pi, a)/[1 - G(\Pi, a)] + g_a(\Pi, a)/g(\Pi, a) < 0$, which holds if and only if DHREP holds.

¹⁴ If the hazard rate of a distribution is decreasing in a parameter of the distribution, the distribution satisfies first-order stochastic dominance (see Ross, 1983). To see why, notice that $\frac{\partial \ln(1-G(\Pi, a))}{\partial \Pi} = \frac{-g(\Pi, a)}{1-G(\Pi, a)}$, and therefore

$$G(\Pi, a) = 1 - \exp \left(- \int_0^{\Pi} \frac{g(\pi, a)}{1 - G(\pi, a)} d\pi \right).$$

Differentiating, we obtain

$$G_a(\Pi, a) = \left(\int_0^{\Pi} \frac{\partial}{\partial a} \frac{g(\pi, a)}{1 - G(\pi, a)} d\pi \right) \exp \left(- \int_0^{\Pi} \frac{g(\pi, a)}{1 - G(\pi, a)} d\pi \right) < 0.$$

¹⁵ Innes (1990) makes an assumption that is weaker than CDFC. It is possible however to prove Innes's result without this regularity assumption.

¹⁶ The MLRP and CDFC conditions in combination greatly restrict the class of distributions and production functions. In particular, Jewitt (1988) observes that few distributions satisfy both the MLRP and CDFC conditions. One distribution was provided by Rogerson (1985) (attributed to Steve Matthews) and later two classes of differentiable examples were provided by Licalzi and Spaeter (2003).

assumptions do not require either MLRP or a regularity assumption similar to CDFC, our analysis shows that debt contracts are optimal for a wide class of production technologies and distributions of the state, or correspondingly, for a wide class of reduced-form distributions, as is discussed in the next section.

5.2 *The Properties of the Critical Ratio*

We now show that we can verify the properties of the critical ratio for a large class of standard probability distributions and production technologies. The state-space representation of uncertainty and the multiplicative decomposition of the critical ratio allow us to consider the applicability of specific probability distributions for the state and specific functional forms for the production technology. Because the critical ratio is the product of the hazard rate of the state and the MRTS, it is sufficient to consider the properties of the two components separately. For example, the critical ratio is increasing if one component is increasing and the other component is either increasing or constant.

The hazard rate, $\frac{f(\theta)}{1-F(\theta)}$, measures the likelihood that the state has a particular value contingent on the state being greater than or equal to that value. Assumptions on the hazard rate are commonly used in adverse selection models such as auctions and nonlinear pricing. Hazard-rate conditions often are applied in studies of reliability and lend themselves to empirical testing and calibration.¹⁷ Applied models in economics and finance often make assumptions on the form of the probability distributions that correspond to increasing, decreasing, or constant hazard rates.

The hazard rate is either constant, increasing, or decreasing for many common probability distributions. For example, the uniform distribution defined on a bounded interval has a hazard rate that is increasing in θ . The exponential distribution, $f(\theta) = \varphi \exp(-\varphi\theta)$, $\varphi > 0$, defined on $[0, \infty)$, has a constant hazard rate. For $\alpha > 1$ ($\alpha < 1$), the gamma distribution, $f(\theta) = \varphi(\varphi\theta)^{\alpha-1} \exp(-\varphi\theta)/\Gamma(\alpha)$, on $[0, \bar{\theta}]$, $\varphi > 0$, has an increasing (decreasing) hazard rate. Also, for $\alpha > 1$ ($\alpha < 1$), the Weibull distribution $f(\theta) = \varphi\alpha\theta^{\alpha-1} \exp(-\varphi\theta^\alpha)$, on $[0, \bar{\theta}]$, $\varphi > 0$, has an increasing (decreasing) hazard rate.

¹⁷ See for example Hall and Van Keilegom (2005).

(see Barlow et al., 1963). The Pareto distribution $f(\theta) = \varphi[\theta_0^\varphi/(\theta)^{\varphi+1}]$, $\varphi > 0$, defined on $[\theta_0, \infty)$, has a hazard rate that is decreasing in θ . The Pareto distribution, or power-law distribution, is widely used in finance and in network economics; see Barabási and Albert (1999) and Gabaix et al (2003).

The MRTS, $\frac{\Pi_a(\theta, a)}{\Pi_\theta(\theta, a)}$, is a common expression in economic models that is used to indicate the rate at which an input can be substituted for another input in production. The MRTS also plays an important role in agency models with adverse selection, as will be discussed in a later section. The MRTS is increasing in the state if effort and the state are complements, $\Pi_{a\theta} > 0$, and the function $\Pi(\theta, a)$ is concave in the state, $\Pi_{\theta\theta} \leq 0$,

$$\frac{d}{d\theta} \frac{\Pi_a(\theta, a)}{\Pi_\theta(\theta, a)} = \frac{\Pi_{\theta a}(\theta, a)\Pi_\theta(\theta, a) - \Pi_{\theta\theta}(\theta, a)\Pi_a(\theta, a)}{(\Pi_\theta(\theta, a))^2} > 0.$$

The complementarity condition is implied by the increasing differences condition that is used in monotone comparative statics analysis, see Topkis (1998) and Milgrom and Shannon (1994). Interpreting output as consumer utility, the ratio $\frac{\Pi_a(\theta, a)}{\Pi_\theta(\theta, a)}$ indicates the consumer's marginal rate of substitution (MRS).

The MRTS is increasing in θ for a given effort a , $\frac{\partial}{\partial \theta} \frac{\Pi_a(\theta, a)}{\Pi_\theta(\theta, a)} > 0$, if and only if the agent's unobservable effort, a , is a *normal* input. Many production functions satisfy input normality. These include: the Cobb-Douglas production function, $\Pi(\theta, a) = \zeta\theta^\beta a^\gamma$, for any positive ζ , γ , and β ; any increasing function H of the product of the agent's effort and the state, $\Pi(\theta, a) = H(\theta a)$; and the additive production function $\Pi(\theta, a) = A(\theta) + H(\theta)Q(a)$, where $H(0) = 0$ and A , H are increasing and concave in θ and Q is increasing. In Stiglitz' (1974) classic model of sharecropping, there is a multiplicative production function, $\Pi(\theta, a) = \theta Q(a)$, where $Q(a)$ is increasing, so that effort is a normal input. Such multiplicative outcome functions are used in many types of economic models in which the random state variable represents prices, technology parameters, or natural states such as weather and demographic effects.

The MRTS is constant in θ for the additive production function $\Pi(\theta, a) = A\theta + Q(a)$, where Q is an increasing function and $A > 0$, and for any increasing function H of the sum of the agent's effort and the state, $\Pi(\theta, a) = H(\theta + a)$. The analysis then must be adjusted for $\Pi(0, a)$ nonconstant in

a. Either of these production functions combined with an exponentially distributed hazard rate are sufficient for the set of optimal contracts to contain linear contracts.

5.3 Comparison with the Adverse Selection Agency Model

The state-space formulation that we use allows for an interesting comparison between the moral hazard agency model and the adverse selection agency model. This is because there is a close correspondence between the *state variable* in the moral hazard model and the *agent's type* in the adverse selection model. Spence and Zeckhauser (1971) and others have also noted this connection. In comparing moral hazard with adverse selection, it is important to observe the centrality of hidden actions in the moral hazard model and hidden information in the adverse selection model. In the moral hazard model of agency, the principal seeks to induce the agent to devote more effort and in the adverse selection model of agency, the principal seeks to induce the agent to reveal private information. In this way, incentives are connected to hidden action in the moral hazard model and to hidden information in the adverse selection model.

The main insight from comparing the two models is the following. In the adverse selection model, the cost of inducing revelation of hidden information in a given state equals the reciprocal of the critical ratio. In the moral hazard model, the reciprocal of the critical ratio is the cost of inducing incentives in a given state.

Consider an adverse selection model of agency that is similar to the moral hazard model considered here.¹⁸ Optimal contract design with a risk-neutral principal and agent and limited liability has been studied as an adverse selection problem by Sappington (1983) in a discrete-types framework. We generalize his setting to allow a continuum of agent types θ and let our Assumptions 1 and 2 hold. Invert the production function, $\Pi = \Pi(\theta, a)$, to obtain the cost of effort needed to achieve output Π in state θ , $a = C(\Pi, \theta)$. Applying the revelation principle, the agency contract can be expressed as a direct revelation mechanism in which the principal makes a transfer $t(\theta)$ to the agent

¹⁸ The optimal agency contract with adverse selection may resemble debt under some conditions. Gale and Hellwig show that the optimal contract with adverse selection is debt based on revelation of information arguments, see also Myers and Majluf (1984). Stiglitz and Weiss (1981) point out that in their setting, incentive schedules may dominate debt contracts when there is adverse selection.

and receives the output $\pi(\theta)$, and the agent's action is reporting his type θ to the principal.

Normality of inputs has related effects in the moral hazard and adverse selection models of agency. In the moral hazard model, hidden effort must be normal, and in the adverse selection model, hidden information must be normal. Recall that in the present moral hazard model, the increasing MRTS condition, $\frac{\partial}{\partial \theta} \frac{\Pi_a(\theta, a)}{\Pi_\theta(\theta, a)} > 0$ for a given effort a , holds if and only if the agent's unobservable effort, a , is a normal input. In the adverse selection model, the Spence-Mirrlees single-crossing condition, which reduces to $\frac{\partial}{\partial a} \frac{\Pi_\theta(\theta, a)}{\Pi_a(\theta, a)} > 0$ for a given type θ , holds if and only if the agent's unobservable type, θ , is a normal input. Normality is associated with the agent's action in each model: effort in the moral hazard model and reported type in the adverse selection model.

Complementarity of inputs has similar effects in both the adverse selection and the moral hazard models. With adverse selection (moral hazard), complementarity of effort and the agent's type (state variable) means that the marginal return to effort is greater in better states in either model.¹⁹ Although the agent's type, or the state variable, are unobservable, the principal provides greater incentives for information revelation to higher-type agents in the adverse selection model and greater incentives for effort in better states in the moral hazard model. The Spence-Mirrlees single-crossing condition holds if the agent's effort and his type are complements, $\Pi_{\theta a}(\theta, a) > 0$, and the production technology is weakly concave in effort, $\Pi_{aa}(\theta, a) \leq 0$. Recall that in our framework, the increasing MRTS condition holds if the agent's effort and the state are complements, $\Pi_{\theta a}(\theta, a) > 0$, and the production technology is weakly concave in the state, $\Pi_{\theta\theta}(\theta, a) \leq 0$. The Spence-Mirrlees single-crossing condition implies that the reciprocal of the critical ratio $\rho(\theta, a)$ is increasing in the agent's effort. This in turn implies that

$$-C_{\Pi\theta}(\Pi, \theta) = \frac{1}{\Pi_a(\theta, a)} \frac{\partial}{\partial a} \frac{\Pi_\theta(\theta, a)}{\Pi_a(\theta, a)} > 0, \quad (28)$$

so the marginal cost of effort is lower for higher-type agents.

The hazard rate of the state or the agent's type, $\frac{f(\theta)}{1-F(\theta)}$, has a somewhat similar technical role in

¹⁹ Some of the examples given in the earlier discussion satisfy both the increasing MRTS condition and the Spence-Mirrlees condition. These include the Cobb-Douglas technology, the multiplicative production technology, $\Pi(\theta, a) = \theta Q(a)$, and the additive technology, $\Pi(\theta, a) = A(\theta) + H(\theta)Q(a)$, where A , H and Q are increasing and concave functions.

moral hazard and adverse selection. In the adverse selection setting, the hazard rate represents the trade-off between providing incentives to the agent of type θ and the cost of providing payments to agents with higher types. A standard assumption in the adverse selection literature used to obtain a separating equilibrium is that the hazard rate is increasing in θ .²⁰ In the moral hazard setting, the hazard rate shows the trade-off between providing the agent with incentives when the state of nature turns out to be θ and having payments that are greater in higher states. In the moral hazard model, a nondecreasing hazard rate, along with an increasing MRTS, yields an increasing critical ratio, so debt contracts are optimal.

Consider the derivation of the optimal contract with adverse selection. For ease of presentation, let $u_0 = 0$. With adverse selection, the optimal mechanism chooses the function $\pi(\cdot)$ to maximize the principal's benefit subject to the agent's individual rationality (IR) participation constraints, $t(\theta) - C(\pi(\theta), \theta) \geq 0$, and the incentive compatibility (IC) constraints, $t(\theta) - C(\pi(\theta), \theta) \geq t(\hat{\theta}) - C(\pi(\hat{\theta}), \theta)$, for all $\theta, \hat{\theta}$ in $[0, \bar{\theta}]$, where $\hat{\theta}$ is the agent's report of his type. The principal's expected benefit is $V(\pi, t) = \int_0^{\bar{\theta}} [\pi(\theta) - t(\theta)] f(\theta) d\theta$.

By standard arguments due to Mirrlees (1971) (see Bolton and Dewatripont, 2005), the optimal contract maximizes the principal's benefit subject to a local IC constraint, a participation constraint for the lowest-type agent, and a monotonicity constraint on the agent's action,

$$\max_{\pi, t} \int_0^{\bar{\theta}} [\pi(\theta) - t(\theta)] f(\theta) d\theta,$$

subject to

$$t'(\theta) = C_\pi(\pi(\theta), \theta) \pi'(\theta), \tag{29}$$

²⁰ Consider the nonlinear pricing study by Maskin and Riley (1984). Their assumptions (1984, Proposition 4) generate second-order conditions for the firm's choice of a nonlinear pricing schedule with adverse selection. The consumer's utility function can be interpreted as a production function, $\Pi = \Pi(\theta, a)$, where the agent's type is a utility parameter, θ , and the agent's action is consumption, a . Their assumptions, while much more extensive, include the following restrictions, $\Pi_a(\theta, a) > 0$, $\Pi_{aa}(\theta, a) < 0$, $\Pi_{a\theta}(\theta, a) > 0$, $\Pi_\theta(\theta, a) > 0$, and $\Pi_{\theta\theta}(\theta, a) < 0$. Maskin and Riley's (1984) assumptions imply that the marginal rate of substitution is increasing in the type. To obtain a fully-separating equilibrium, Maskin and Riley (1984, equation 23) further assume that the hazard rate of agent types, $f(\theta)/(1-F(\theta))$, does not decline too rapidly by requiring that the distribution of agent types satisfies $2(f(\theta))^2 + (1-F(\theta))f'(\theta) > 0$. This is a weaker condition than requiring that the hazard rate of agent types be increasing, that is $(f(\theta))^2 + (1-F(\theta))f'(\theta) > 0$. Their assumptions, with strengthening to obtain an increasing hazard rate for the shock, imply that the critical ratio $\rho(\theta, a)$ in our setting is increasing in θ .

$$t(0) - C(\pi(0), 0) \geq 0, \quad (30)$$

$$\pi'(\theta) \geq 0. \quad (31)$$

It can be shown that the lowest-type participation constraint is binding.

Applying standard arguments, given local incentive compatibility and the participation constraint for the lowest-type agent, the principal's objective can be restated as

$$V(\pi, t) = \int_0^{\bar{\theta}} [\pi(\theta) + C_\theta(\pi(\theta), \theta) \frac{1 - F(\theta)}{f(\theta)} - C(\pi(\theta), \theta)] f(\theta) d\theta. \quad (32)$$

The first-order condition for the principal's choice of $\pi(\cdot)$ is then

$$C_\pi(\pi(\theta), \theta) - C_{\theta\pi}(\pi(\theta), \theta) \frac{1 - F(\theta)}{f(\theta)} = 1. \quad (33)$$

When the local first-order condition holds, applying the Spence-Mirrlees single-crossing condition implies that the marginal product of effort is greater than the marginal cost of effort, since $C_{\theta\pi}(\pi(\theta), \theta) < 0$. Effort is therefore inefficient.

Noting that $C_\pi(\pi(\theta), \theta) = 1/\Pi_a(\theta, a(\theta))$ and substituting for $C_{\theta\pi}(\pi(\theta), \theta)$ using (28), the first order condition (33) can be written as

$$\Pi_a(\theta, a(\theta)) - 1 = \left[\frac{\partial}{\partial a} \frac{1}{\rho(\theta, a)} \right]_{a=a(\theta)}. \quad (34)$$

In the adverse selection model, where the principal is inducing revelation of the agent's type, the reciprocal of the critical ratio is increasing in the agent's effort by the Spence-Mirrlees condition, so the right hand side of (34) is positive. This contrasts with our setting, where the principal is inducing agent effort, and the critical ratio is increasing in the state.

The critical ratio can also be used to understand the slope of the contract in the adverse selection model. The principal's transfer payment to the agent can be expressed as a contract, $t(\theta) = w(\pi(\theta))$. Then, the slope of the contract equals $t'(\theta) = w'(\pi(\theta))\pi'(\theta)$. Applying the agent's local incentive compatibility condition (29) and the first-order condition from the principal's problem in the form

of (34) implies

$$w'(\pi(\theta)) = C_\pi(\pi(\theta), \theta) = \frac{1}{\Pi_a(\theta, a(\theta))} = \frac{1}{\left[\frac{\partial}{\partial a} \frac{1}{\rho(\theta, a)} \right]_{a=a(\theta)} - 1}. \quad (35)$$

Finally, the critical ratio can be used to represent information rents in the adverse selection model.

From (31) and $C_\theta(\pi(\theta), \theta) = -\frac{\Pi_\theta(\theta, a(\theta))}{\Pi_a(\theta, a(\theta))}$, a type- θ agent has information rents,

$$t(\theta) - C(\pi(\theta), \theta) = \int_0^\theta -C_\theta(\pi(z), z) dz = \int_0^\theta \frac{\Pi_\theta(z, a(z))}{\Pi_a(z, a(z))} dz. \quad (36)$$

Information rents increase with an agent's type. Summing information rents across agent types and integrating by parts gives expected information rents,

$$\int_0^{\bar{\theta}} \int_0^\theta \frac{\Pi_\theta(z, a(z))}{\Pi_a(z, a(z))} dz f(\theta) d\theta = \int_0^{\bar{\theta}} \frac{1}{\rho(\theta, a(\theta))} f(\theta) d\theta. \quad (37)$$

So, expected information rents in agency with adverse selection equal the sum of the reciprocal of our critical ratio. The reciprocal of the critical ratio in the adverse selection model is therefore the cost of inducing the agent to reveal hidden information.

Compare (37) with the agent's expected net benefits in the moral hazard agency model. When the optimal contract is debt with $R > 0$ and $w(\Pi(0, a); a) = 0$, from (17), the agent receives expected net benefits of

$$U(w, a) = \int_{\hat{\theta}(R, a)}^{\bar{\theta}} \frac{1}{\rho(\theta, a)} \Pi_a(\theta, a) f(\theta) d\theta - a. \quad (38)$$

The reciprocal of the critical ratio indicates the principal's cost of inducing the agent to provide hidden effort at each state.

6 Conclusion

Optimal agency contracts provide a pattern of incentive payments to agents that in general differs from a linear transformation of the underlying outcome. This recalls the common observation in finance that options generate a pattern of returns that could not be obtained with the underlying stock (Cox and Rubinstein, 1985, p. 45). When the critical ratio is increasing in the random state, higher states are better for providing incentives, so that the optimal contract is debt. When the critical ratio is strictly decreasing in the random state, lower states are more efficient at providing

incentives and the optimal contract is a capped bonus. When the critical ratio is constant in the random state, all states are equally efficient at providing incentives and only then are linear equity-style contracts efficient, but so are many nonlinear contracts.

The result that the optimal contract takes the form of debt when the critical ratio is increasing has far-reaching implications. In contrast to complex performance schedules, debt contracts are easy to design and apply and have standardized legal formats, thus mitigating transaction costs. The optimality and simplicity of debt-style contracts helps to explain the use of performance targets and rewards in a wide variety of economic situations. Debt-style contracts are optimal for sharecropping, procurement, insurance, employee performance, and regulatory incentives. Moreover, debt-style contracts are at the heart of financial contracts and performance rewards for entrepreneurs and managers. Debt contracts have important properties that facilitate market pricing and trading. Financial markets have extensive experience in pricing and trading debt, options, and similar assets.

Our framework allows for tractable economic analysis that indicates how the form of optimal agency contracts depends on the properties of the underlying uncertainty and production technology. This approach should help researchers to derive optimal contracts for a very wide range of applications in economics and finance. The increasing critical ratio requirement, or equivalently DHREP, which we introduce in the reduced-form setting, is less restrictive than the standard MLRP. Additionally, our analysis helps to identify connections between moral hazard and adverse selection models of agency. Our approach to agency with moral hazard should be highly useful in characterizing the many extensions of the basic agency model.

7 References

References

Arya, A., J. Glover, and B. Mittendorf, 2007, "Aligning Incentives by Capping Bonuses," Chapter 8 in R. Antle, F. Gjesdal, and P. J. Liang, eds., *Essays in Accounting Theory in Honour of Joel S. Demski*, Part II, New York, Springer, pp. 169-182.

Barabási, A.-L. and R. Albert, 1999, "Emergence of Scaling in Random Networks," *Science*, 286 (5439), October 15, pp. 509-512, DOI: 10.1126/science

- Barlow, R. E. , A. . Marshall, F. Proschan, 1963, "Properties of Probability Distributions with Monotone Hazard Rate," *Annals of Mathematical Statistics*, 34 (2), June, pp. 375-389.
- Bolton, P. and M. Dewatripont, 2004, *Contract Theory*, Cambridge, MA: MIT Press.
- Conlon, J., 2009, "Two New Conditions Supporting the First-Order Approach to Multisignal Principal-Agent Problems," *Econometrica*, 77, pp 249-278.
- Cox, J. C. and M. Rubinstein, 1985, *Options Markets*, Englewood Cliffs, N.J.: Prentice Hall.
- Dixit, A. K. and R. S. Pindyck, 1994, *Investment under Uncertainty*, Princeton: Princeton University Press.
- Diamond, D. W., 1994, "Financial Intermediation and Delegated Monitoring," *Review of Economic Studies*, 51, pp.393-414.
- Gavaix, X., P. Gopikrishnan, V. Plerou, and H. E. Stanley, 2003, "A Theory of Power-Law Distributions in Financial Market Fluctuations," *Nature*, 423, May 15, www.nature.com/nature, pp. 267-270.
- Gale, D. and M. Hellwig, 1985, "Incentive-Compatible Debt Contracts: The One-Period Problem," *Review of Economic Studies*, Vol. 52, No. 4, October, pp. 647-663.
- Hall, P. and I. Van Keilegom, 2005, "Testing for Monotone Increasing Hazard Rate," *Annals of Statistics*, 33 (3), pp. 1109-1137.
- Harris, M. and A. Raviv, 1979, "Optimal Incentive Contracts with Imperfect Information", *Journal of Economic Theory*, 20 (2), pp. 231-259
- Healy, P. M., 1985, "The Effect of Bonus Schemes on Accounting Decisions," *Journal of Accounting and Economics*, 7, pp. 85-107.
- Holmstrom, B., 1979, "Moral Hazard and Observability," *Bell Journal of Economics*, 19, pp. 74-91.
- Holmstrom, B. and J. Tirole, 1997, "Financial Intermediation, Loanable Funds, and the Real Sector," *Quarterly Journal of Economics*, 112, August, pp. 663-691.
- Holmstrom, B. and L. Weiss, 1985, "Managerial Incentives, Investment and Aggregate Implications: Scale Effects", *Review of Economic Studies*, 52, July, pp. 403-425.
- Innes, R. D., 1990, "Limited Liability and Incentives Contracting with Ex-ante Action Choices," *Journal of Economic Theory*, 52, pp.45-67.
- Jensen, M., 2003, "Paying People to Lie: The Truth about the Budgeting Process," *European Financial Management*, 9, pp. 379-406.
- Jensen, M., and W. Meckling, 1976, "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure," *Journal of Financial Economics*, 3, pp. 305-360.
- Jewitt I., 1988, "Justifying the First-Order Approach to Principal-Agent Problems," *Econometrica*, 56, pp. 1177-1190.
- Jewitt, I, O. Kadan, and J. M. Swinkels, 2008, "Moral Hazard with Bounded Payments." " *Journal of Economic Theory*, 143 (1), November, pp. 59-82.
- Kim, S. K., "Limited Liability and Bonus Contracts," *Journal of Economics & Management Strategy*, 1997, 6 (4), pp. 899-913.

- Laffont, J.-J. and D. Martimort, 2002, *The Theory of Incentives: The Principal-Agent Model*, Princeton: Princeton University Press.
- Laffont, J.-J. and J. Tirole, 1993, *A Theory of Incentives in Procurement and Regulation*, Cambridge, MAL MIT Press.
- Licalzi, M. and S. Spaeter, 2003, "Distributions for the First-Order Approach to Principal-Agent Problems," *Economic Theory*, 21, pp. 167-73
- Maskin, E. and J. Riley, 1984, "Monopoly with Incomplete Information," *RAND Journal of Economics*, Vol. 15, No. 2, Summer, pp. 171-196.
- Milgrom, P. and C. Shannon, 1994, "Monotone Comparative Statics," *Econometrica*, 62 (1), January, pp. 157-180.
- Mirrlees, J. A., 1971, "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38 (2), April, pp. 175-208.
- Mirrlees, J. A., 1974, "Notes on Welfare Economics, Information, and Uncertainty," in Balch, McFadden, and Wu, eds., *Essays on Economic Behavior under Uncertainty*, Amsterdam: North Holland Publishing Co.,
- Mirrlees, J. A., 1976, "The Optimal Structure of Incentives and Authority within an Organization," *Bell Journal of Economics*, 7 (1), Spring , pp. 105-131.
- Myers, S.C., 1984, "The Capital Structure Puzzle," *Journal of Finance*, 39, pp. 575-592
- Myers, S.C. and N.S. Majluf, 1984, "Corporate Financing and Investment Decisions When Firms Have Information the Investors Do Not Have", *Journal of Financial Economics*, 13, pp. 187-221.
- Park, E., 1995, "Incentive Contracting Under Limited Liability" *Journal of Economics and Management Strategy*, 4, pp. 477-490.
- Rogerson W. P., 1985, "The First-Order Approach to Principal-Agent Problems", *Econometrica*, 53, pp. 1357-1367.
- Ross, S., 1973, "The Economic Theory of Agency: The Principal's Problem." *The American Economic Review*, 63, pp. 134-139.
- Ross, S., 1983, *Stochastic Processes*, New York: Wiley.
- Sappington, D., 1983, "Limited Liability Contracts Between Principal and Agent," *Journal of Economic Theory*, 29, pp. 1-21.
- Shaked, M. and J. G. Shanthikumar, 2007, *Stochastic Orders*, New York: Springer.
- Spence, A. M., 1974, Competitive and Optimal Responses to Signals: An Analysis of Efficiency and Distribution, *Journal of Economic Theory*, 7, pp. 296-332.
- Spence, A. M. and R. Zeckhauser, 1971, "Insurance, Information, and Individual Action." *American Economic Review*, 61 , pp. 380-387.
- Stiglitz, J. E. 1974, "Incentives and Risk Sharing in Sharecropping," *Review of Economic Studies*, 41 (2), 219-255.
- Stiglitz, J. E. and A. Weiss, 1981, "Credit Rationing in Markets with Incomplete Information," *American Economic Review*, 71 (3), pp 393-410.

8 Appendix

Lemma 1: *For any feasible contract, w , there always exist a binary contract, w_λ , such that $U(w_\lambda, \hat{a}) = U(w, \hat{a})$ and $w(\Pi(0, \hat{a}); \hat{a}) = w_\lambda(\Pi(0, \hat{a}); \hat{a})$.*

Proof: Let $S_\lambda^+(\hat{a}) = \{\theta : \rho(\theta, \hat{a}) > \lambda\}$ and $S_\lambda^0(\hat{a}) = \{\theta : \rho(\theta, \hat{a}) = \lambda\}$. Define

$$\chi_\lambda(\Pi, \hat{a}) = \begin{cases} 1 & \text{if } \hat{\theta}(\Pi, \hat{a}) \in S_\lambda^+, \\ 0 & \text{otherwise,} \end{cases}$$

$$\psi_\lambda(\Pi, \hat{a}) = \begin{cases} 1 & \text{if } \hat{\theta}(\Pi, \hat{a}) \in S_\lambda^0, \\ 0 & \text{otherwise.} \end{cases}$$

Finally, define

$$w_{\lambda,b}(\Pi; \hat{a}) = \int_0^\Pi [\chi_\lambda(t) + b\psi_\lambda(t)] dt + w(\Pi(0, \hat{a}); \hat{a}).$$

The contract $w_{\lambda,b}$ with $b \in [0, 1]$, is a binary with a slope of b when $\rho(\theta, a) = \lambda$. This contract has the following properties:

(a) $U[w_{\lambda,b}(\Pi; \hat{a}); \hat{a}]$ is continuous and nondecreasing in b . This follows from the fact that $w_{\lambda,b}$ is continuous in b .

(b) $\lim_{\varepsilon \rightarrow 0} U(w_{\lambda-\varepsilon,0}) = U(w_{\lambda,1})$. To see this, fix γ and pick ε small enough so that $\int_{\lambda-\varepsilon}^\lambda \rho(\theta, \hat{a}) d\theta < \gamma$. This can be done because there can be at most $1/\gamma$ points with a mass greater than γ . Define $\zeta_{\lambda,\gamma}(\Pi) = \{1 \text{ if } \rho(\hat{\theta}[\Pi, \hat{a}], \hat{a}) \in (\lambda - \varepsilon, \lambda), 0 \text{ otherwise}\}$. Then, it follows that $w_{\lambda-\varepsilon,0}(\Pi; \hat{a}) = \int_0^\Pi [\chi_{\lambda-\varepsilon,0}(t, \hat{a})] dt = \int_0^\Pi [\chi_{\lambda,0}(t, \hat{a}) + \psi_\lambda(t, \hat{a}) + \zeta_{\lambda,\gamma}(t, \hat{a})] dt = w_{\lambda,1}(\Pi) + \int_0^\Pi \zeta_{\lambda,\gamma}(t, \hat{a}) dt$. Finally notice that $\int_0^\Pi \zeta_{\lambda,\gamma}(t) dt < \gamma$. Therefore $\lim_{\varepsilon \rightarrow 0} w_{\lambda-\varepsilon,0}(\Pi; \hat{a}) = w_{\lambda,1}(\Pi; \hat{a})$ and by the continuity of $U(\cdot)$, $\lim_{\varepsilon \rightarrow 0} U(w_{\lambda-\varepsilon,0}[\cdot]) = U(w_{\lambda,1}[\cdot])$.

To prove the result, suppose to the contrary that such a binary contract does not exist. Then there must exist a λ such that $U(w_{\lambda,0}, \hat{a}) < U(w, \hat{a})$ but $U(w_{\lambda-\varepsilon,0}, \hat{a}) > U(w, \hat{a}) \forall \varepsilon > 0$. Notice now that, by property (b), $U(w_{\lambda,1}, \hat{a}) > U(w, \hat{a})$. Finally, using property (a) and the intermediate value theorem, there exists a value b such that $U(w_{\lambda,b}, \hat{a}) = U(w, \hat{a})$, which is a contradiction. \square

Proof of Corollary 1: A feasible debt contract exists by Lemma 1. The expected utility of the agent with a debt contract that has face value R is given by

$$U(R, a) = \int_0^{\bar{\theta}} \max\{\Pi(\theta, a) - R, 0\} f(\theta) d\theta - a - u_0. \quad (39)$$

The function $U(R, a)$ is continuous and decreasing in R and it has increasing differences in $\{a, -R\}$. Let $A(R) = \{a : a \in \arg \max_x U(R, x) \text{ and } a \geq 0\}$. It is easy to check that the set $A(R)$ is closed. Let $a(R)$ be the largest element in $A(R)$. By standard monotone comparative statics arguments, $a(R)$ is a decreasing function. Because it is decreasing and bounded it has at most countable many discontinuities.

The expected utility of the principal can be written as

$$V(R, a) = \int_0^{\bar{\theta}} \min\{\Pi(\theta, a), R\} f(\theta) d\theta. \quad (40)$$

Let $V^H(R) = V(R, a(R))$. This function is discontinuous only if $a(R)$ is discontinuous. Because $a(R)$ is decreasing, $V^H(R)$ is also decreasing at any discontinuity point.

Define the function $\tilde{V}(R)$ as follows:

$$\tilde{V}(R) = \sup_{\tilde{R} \leq R} V^H(\tilde{R}).$$

$\tilde{V}(R)$ is clearly a nondecreasing and continuous function, since $V^H(\tilde{R})$ is decreasing at any discontinuity point. Consider now the following modified problem:

$$\max_R \tilde{V}(R) \text{ subject to } U(R) \geq 0.$$

The problem has a solution because we are maximizing a continuous function and the set of debt contracts that satisfy $U(R) \geq 0$ is closed. We can always restrict attention to a compact set of debt contracts with $R \leq \bar{R}$ for some \bar{R} positive because $U(R)$ converges to $-u_0$ as R grows, and therefore high levels of debt will not satisfy the agent's participation constraint. Let the solution to this problem be \tilde{V}^* . Let R^* be the lowest R such that $\tilde{V}(R) = \tilde{V}^*$. Note that R^* is well defined because \tilde{V} is a continuous function.

We claim that a contract with face value R^* can implement \tilde{V}^* . To see why, remember that $\tilde{V}(R) =$

$\sup_{R \leq \tilde{R}} V^H(\tilde{R})$. So, there exists a sequence $a^n \in A(R^* - \varepsilon^n)$, with $\lim_{n \rightarrow \infty} \varepsilon^n = 0$ and $\lim_{n \rightarrow \infty} a^n = a^*$ such that $\lim_{n \rightarrow \infty} V(R^* - \varepsilon^n, a^n) = \tilde{V}^*$. If R^* cannot implement \tilde{V}^* , it must be that $a^* \notin A(R^*)$, which means that there exists \tilde{a} such that $U(R^*, \tilde{a}) > U(R^*, a^*)$. But, since U is continuous in both arguments by picking n sufficiently large this also means that $U(\hat{R} - \varepsilon^n, \tilde{a}) > U(\hat{R} - \varepsilon^n, a^n)$, which is a contradiction. Finally we claim that the contract with face value R^* , is optimal. Suppose not. Then, there exists R such that $V(R) > V(R^*)$ and $U(R) \geq 0$. But, this implies that $\tilde{V}(R) \geq V(R) > \tilde{V}(R^*)$, which contradicts the optimality of R^* in the modified problem.

The proof of existence of an optimal capped bonus is analogous. \square

Proof of Proposition 3. (i) As already noted, MLRP implies DHREP. (ii) Consider, for example, the reduced-form probability density function

$$g(\Pi, a) = \begin{cases} \frac{1+a(1-\Pi)}{1+a} & \text{if } \Pi < 1, \\ \frac{1+a-\Pi}{(1+a)a} & \text{if } 1 < \Pi < 1+a. \end{cases}$$

To check that the reduced-form distribution does not satisfy MLRP, notice that for $\Pi < 1$,

$$\frac{g_a(\Pi, a)}{g(\Pi, a)} = -\frac{\Pi}{[1 - a(1 - \Pi)](1 + a)},$$

which is decreasing in Π . Note, however, that for $\Pi > 1$, MLRP does hold.

To see that the reduced-form distribution satisfies *DHREP*, observe that, for $\Pi \leq 1$, $G(\Pi, a) = \Pi - \frac{a}{2+2a}\Pi^2$. So, both $g(\Pi, a)$ and $G(\Pi, a)$ are decreasing in a . Then, $\frac{\partial}{\partial a} \frac{g}{1-G}$ is also decreasing in a . For $\Pi > 1$, *DHREP* also holds since *MLRP* holds. Intuitively, for $\Pi < 1$, the slope of the reduced-form density, $g(\Pi, a)$, becomes more negative as a increases. Thus, conditional on $\Pi < 1$, it follows that as a increases, lower values of Π become more likely relative to higher values of Π . This violates MLRP. On the other hand, as a increases, more of the mass of the distribution is shifted to the right of $\Pi = 1$, which means that both the probability density function and the cumulative distribution function are falling for $\Pi < 1$. Therefore, the hazard rate is decreasing in a for all $\Pi < 1$, and it also is decreasing in a for $\Pi > 1$. \square