1998

UCD Department of Agricultural Economics

Contracts 1896

WORKING PAPER SERIES

University of California, Davis

Department of Agricultural & Resource Economics

Input Control and Common Risk: Addressing Agent Heterogenity and Risk Aversion in the Presence of Moral Hazard

by

Rachael E. Goodhue

Working Paper No. 98-1

Acknowledgements: This paper is based on my Ph.D. thesis. Benjamin Hermalin, Jeff Perloff, Gordon Rausser, Leo Simon and David Zilberman provided many helpful comments and suggestions during the dissertation-writing process. Richard Sexton and seminar participants at the University of California at Davis also provided useful feedback on this work. An earlier version of this work appeared as Working Paper No. 816, Department of Agricultural and Resource Economics, University of California at Berkeley. Bob Lyons and Leo Simon provided assistance with BRS, the computer simulation environment in which simulations for this paper were performed. Financial support was provided by the National Science Foundation and the Poultry Science Association. The author is a member of the Giannini Foundation of Agricultural Economics.

ABSTRACT:

Examining empirical regularities regarding contract provisions leads to a pair of puzzles regarding these contracts: why do principals control inputs, and why do they use a statistically insufficient estimator to calculate compensation for heterogeneous agents? Rather than attributing such provisions to real world imprecisions and imperfections, this paper provides an agency theoretic framework that explains these puzzles. The principal controls inputs in order to reduce the information rents paid to agents and increase profits. Forcing agents to bear additional income risk through the use of an imprecise estimator can reduce the cost of providing incentive compatibility more than it increases the cost of compensating the agents for the added income risk. These theoretical findings are illustrated with examples from agricultural production contracts.

1. Introduction

This paper addresses two questions regarding the structure of principal-agent relationships: why does the principal control non-labor inputs, and how does the principal select his measure of the realization of common uncertainty when relative compensation is used to protect agents against common risk? Such contract provisions are often regarded as resulting from real world imprecisions and imperfections, such as liquidity constraints and the costs of obtaining information, rather than being grounded in incentive design.

These provisions are observable in actual contracts. In agriculture, for example, many production contracts between farmers and food processors and other principals provide farmers with necessary inputs required for production. Fruit and vegetable processors and shippers often dictate the time of planting and harvest, and specify other decisions, such as the variety grown and which fertilizers and other chemicals may be used. Broiler processors control the two important variable inputs used in the production of chickens: the chicks and feed. Processors specify what capital equipment growers must use. Outside of agriculture, many franchising contracts require franchisees to purchase inputs from the franchisor or approved suppliers. Executive compensation contracts are an example of relative compensation; some of these contracts include clauses which reward executives based on their company's performance relative to the industry average. In agriculture, broiler processors compensate their farmers in part based on their performance relative to the average, so that farmers' exposure to common risk is reduced.

My examination of these contract provisions establishes incentive-based motivations for their adoption when agents are heterogeneous, and, for relative compensation, risk averse. Proposition 1 establishes an incentive-based reason for a principal to control inputs that has not been identified previously in the agency theory literature: by controlling inputs the principal can reduce the information rents she must pay to high productivity agents. The principal must offer an incentive-compatible contract menu to heterogeneous agents of unknown ability in the presence of model hazard. When agents of different types would select a different input-effort combination to produce a given amount of output, she can reduce the information rents paid to highly productive agents by controlling the amount of input utilized by each agent.

Heterogeneity and information rents also play a role in the principal's choice of an informationally inefficient relative compensation measure. Holmstrom (1982) showed that a principal will maximize profits by removing the maximum amount of uncertainty from homogeneous agents' payment streams. Proposition 2 confirms that this result may be extended to heterogeneous agents of known types. Proposition 3 establishes that the use of the simple average of output does not remove the maximum amount of uncertainty from agents' payments when types are unknown. Proposition 4 suggests that under some conditions the principal may maximize profits by forcing agents to bear more than the minimum amount of risk, due to interactions between income risk and information rents.

In contrast to this paper, which examines the principal's control of non-labor inputs as a means of minimizing the information rents paid to agents, the agency theory literature generally has assumed that there is no substitutability between labor and other inputs when considering the principal's incentive problem. Khalil and Lawarree (1995) examine a related neglected issue, which is residual claimancy as a source of information rents. They find that when input and output monitoring are both feasible and equally costly to the principal, she will prefer monitoring labor when she is the residual claimant and monitoring output when the agent is the residual claimant. In related work, Maskin and Riley (1985) compare effort to output monitoring when the agent is the residual claimant, and find that output monitoring is more desirable, since high-ability agents will exert more effort when their marginal incentives are not distorted.

For a single risk-averse agent or for multiple risk-averse homogeneous agents, welfare is directly related to the quality of the principal's information regarding contractual outcomes (Grossman and Hart 1983, Holmstrom 1982). In these frameworks, the agent's action does not affect his attitude toward income risk. Consequently, reducing risk in the agent's payment stream increases total welfare, so it is never profitable to add a random element to the agent's income stream. In this paper, in contrast, the agent's risk attitude is dependent on his action. With heterogeneous agents, these factors may lead to profits increasing with the variance of agents' payments. Maskin and Riley (1984) derive a related result in the context of auctions with heterogeneous risk-averse buyers. They find that it is optimal for the seller to utilize an auction with lump-sum fines and subsidies that increases the riskiness of income for buyers with low reservation prices and reduces the riskiness of income for buyers with high reservation prices. Their analysis differs from this one in two important ways. Maskin and Riley allow the risk faced by buyers to be a control variable for the seller, who must then induce them to reveal their types truthfully. Thus, the principal can completely insure the agents if she desires. Here, in contrast, the risk faced by agents is a function of exogenous uncertainty as well as the principal's choices. The exogenous uncertainty places a positive lower bound on the income risk faced by agents regardless of the principal's decisions. That is, the principal can not completely insure the agents. Further, the principal faces a combined hidden information hidden action problem, whereas Maskin and Riley consider only a hidden information problem.

2. Modeling Input Control

In order to focus on the effects of input control, I begin by assuming away risk aversion and uncertainty. Agents' types and input use are unobservable to others. The only way that the principal can control input usage is to administer inputs directly. The problem faced by the principal is that agents will not truthfully reveal their private information unless the principal gives them incentives to do so.

Agents are one of two types: high ability h with probability p and low ability l with probability $q = 1 - p$. These probabilities are known to the principal and the agents. Production by an agent of type i is a function of the agent's type-specific productivity parameter t_i, effort a_i and an input

Q_i that may be provided by the principal or the agent at a constant unit cost of c. Output for type i is assumed to have constant returns to scale and is specified as follows:

$$x_i = Q_i^{\rho} a_i^{\phi} t_i \tag{1}$$

The principal maximizes profits. The agents maximize utility, which is separable in income and effort. Agents face an increasing marginal disutility of effort. The agents' utility is defined as follows:

$$U = y_i - d(a_i) \tag{2}$$

where y_i is income. Exerting effort is costly for the agents, so that $d' > 0$; and each additional unit of effort is increasingly costly, so that $d'' > 0$. Attention is restricted to compensation functions that are monotonically increasing in output, x_i. This assumption eliminates the possibility of perverse compensation functions that provide lower compensation to agents producting higher levels of output. While such functions are theoretically possible, they do not appear to be significant empirically. Given this assumption, I can then restrict attention to affine compensation functions without any further loss of generality, since utility is linear in income.[1] Accordingly, income is defined as the sum of output multiplied by the per-unit compensation plus a lump-sum transfer, or

$$y_i = W_i x_i + T_i \tag{3}$$

[1] To see this, consider the compensation function, $w(x)$, and any two distinct points in the domain of this compensation function, call them x_l and x_h. Since $w(x)$ is monotonically increasing, $w(x_l) \neq w(x_h)$. Agents of either type have a discrete choice between the two points on the compensation function. Their decision is influenced by the total compensation they receive at each point and the marginal value of the compensation paid at each of the two relevant output levels. Let $w'(x_l) = W_l$ and $w'(x_h) = W_h$. For any marginal value of compensation, W_i, I can define a T_i such that $w(x_i) = W_i x_i + T_i$. Thus, the important features of the compensation function are fully captured by an affine function.

Under this compensation function and the other specifications above, utility will be strictly concave in effort (or output) for an agent.

The timing of the model is as follows: the principal offers a menu of contracts, agents announce their types by choosing a contract, production occurs, output is observed, and agents are compensated.

3. Profit Maximization: Input Control

In general, a principal may choose to supply non-labor inputs to the agent, or may allow the agent to determine input levels. In the case of broiler production, chicks and feed are necessary inputs. The processor provides both inputs to growers, and chooses the amount supplied. Similarly, although growers are responsible for purchasing and maintaining capital equipment they must comply with processor specifications, including equipment upgrades.

For analytical convenience I assume a single composite input.[2] When agents' types are unknown, the principal's need to provide incentive compatible contracts drives a wedge between her cost of a specified x_l, x_h pair and the agents' utility-maximizing (and production cost minimizing) production decisions. When the principal controls the input, she can reduce this wedge. Control over the input reduces the cost of maintaining incentive compatibility for high productivity agents.[3]

To provide insight into the nature of the informational advantage of input control, I begin by analyzing the principal's choice of contracts in the first best case when types are known. When types are known, the principal's ability or inability to control the input does not affect her choice of contract. When types are unknown, Proposition 1 shows that the principal increases her profits

[2] Note that this assumption implies no loss of generality if the production process requires that the non-labor inputs must be used in fixed proportions.

[3] In contrast to the situation modeled here, the principal may allow the agent to choose the input level and will then compensate the agent for some share of the input cost. These arrangements may be chosen because the principal's knowledge of production is less accurate than the agent's, or production knowledge is relatively expensive for the principal to collect. In either case, losses due to information rents are relatively small compared to the cost of reducing them. Here I assume that all parties have the same information regarding production and profit-maximizing input use given the type parameter.

by assigning the input, which reduces the cost of maintaining incentive compatibility for the high ability agent.

3.1. **Known types.** The case of known types is the theoretical benchmark for evaluating the hidden information case.[4] When the principal can restrict heterogeneous agents to their reservation utility levels, she can increase profits by designing different contracts for the different types. When there is no hidden information, however, the principal does not gain from controlling the input *in addition* to specifying the output level. The principal and the agent have identical production costs, and there are no information costs for the principal so both parties will choose the input combination that minimizes production costs, the "neoclassical" allocation, for any output level.

Fact: *When types are known, the principal's ability or inability to control the input does not affect her profits or choice of contracts.*

Observe that the fact implies that the principal will desire to vary input assignments across heterogeneous agents.

3.2. **Unknown types.** In the first-best case, the agent and the principal face the same production costs. Whoever controls the input will select the neoclassical production cost-minimizing ratio of effort and the input. In the hidden information case, information costs drive a wedge between the cost of x_l for the low-ability agent and for the principal. The agent faces the same costs of production as in the first-best case, but the principal faces information costs in addition to production costs. When types are unknown, the principal's choice of inputs is affected by these incentive considerations. In particular, for fixed x_l the principal provides proportionately more of the input relative to agent effort than she does in the first-best case, in order to reduce the high ability agent's information rents, which are a function of the gap between the effort a low ability

[4] It also describes the principal's choice of contract if she is dealing with homogeneous agents of known type in the absence of moral hazard. (The case of a single type of unknown ability is not considered.) Consequently, comparing the outcomes of cases with known and unknown types provides comparisons of the contract outcomes not only for these two case but also for the cases of homogeneous versus heterogeneous agents.

agent must exert to produce x_l and the effort a high ability agent must exert to produce x_l. Hidden information increases the cost of low-agent effort relative to the cost of the input for every (x_l, x_h) pair. When the principal assigns the input, she can respond to this change in her relative input prices by adjusting the input mix away from effort and toward the input for the low-ability contract. When the agent controls the input, the principal cannot adjust her input mix.

There are two effects that increase profits in the second-best case when the principal rather than the agent controls the input. First, the principal lowers the information rent obtained by the high-ability agent for any given x_l by controlling the amount of the input used. Second, since the information rent is lowered for every possible x_l, the x_l chosen for the optimal contract increases, so there is a reinforcing scale effect. Consequently, profits under hidden information are lower when the agent controls the input. These observations are summarized in the following proposition:

Proposition 1. *When types are unknown, the principal increases profits by assigning input levels according to type rather than allowing agents to choose their own input levels.*

The information rent and scale effects may be observed in Figure 1. The figure is plotted in output-total cost space. The two curves $TC(a_l^{***}, Q_l^{***})$ and $TC(a_h^{***}, Q_h^{***})$ represent the total cost of the cost-minimizing, or neoclassical, combination of effort and the input for low and high ability agents, respectively. x_l^* and x_h^* are the principal's first best, or full information, choices of output for the two types. The marginal cost of agent effort and the input are equal to their marginal revenue product for each agent type. Under full information, the principal and the agent both choose the production cost-minimizing input combination, $a_i^{***}(x_i^*), Q_i^{***}(x_i^*)$. The total cost curves are lower envelopes of cost curves which hold the amount of the input fixed and then vary effort to obtain different amounts of the output for each type. Some of these iso-input cost curves are depicted in the graph. Note that these curves are tangent to the total cost curve at the first-best output levels for both types.

In the first-best case, the principal maximizes profits taking only production costs into account. In the second-best case, the principal must consider both information and production costs. When the agents control the input, the principal chooses her profit-maximizing contracts knowing that the agents will choose the production cost-minimizing combination of effort and the input. This will result in a pair of output choices such as x_h^* and x_l^a. The principal's choice of output for the low type is distorted downward due to information costs. In order to induce the high ability agent to truthfully reveal his type, the principal must pay him information rents equal to the amount of utility he would gain from pretending to be a low ability agent. The monetary value of these rents when the agent controls the input is simply equal to the vertical distance CG, since a lying high ability agent will choose the production cost-minimizing combination of effort and the input.

When the principal controls the input, she specifies how much of the input will be provided to an agent who announces a given type. This limits the high ability agent's scope to profit from lying about his type, since he will not be able to choose his least costly input and effort combination. Consider what would happen if the principal could specify that an agent announcing himself to be low ability must use input level Q_l^{***} to produce output level x_l^a. Then a lying high ability agent would be forced to use $Q_l^{***}(x_l^a)$ to produce x_l^a. This would cause the high ability agent to produce on his associated iso-input line, reducing his information rents to the amount CE. Obviously, x_l^a will no longer be the profit-maximizing output choice for the principal.

When the principal controls the input, her choice for the high ability agent will remain at x_h^*, using the production cost-minimizing input amount. Her choice for the low ability agent will be at a point such as x_l^p. There are two things to notice about this point: first, that the principal's ability to reduce information rents by controlling the input allows her to increase the output required of a low ability agent relative to when the agents control the input; and second, that the principal will choose a different effort-input combination for x_l^p than the production cost-minimizing combination.

By slightly increasing the ratio of input to effort, the principal obtains a first-order decrease in information rents (DF) while bearing only a second-order increase in production costs (AB).

The results in this section highlight a novel dimension of the input control component of the principal-agent relationship. Existing explanations stress aspects such as differential access to credit, different costs for obtaining the input and agent liquidity constraints (Lehnert, Ligon and Townsend 1997).[5] By contrast, this solution establishes an incentive-based reason why principals may choose to supply their agents with inputs. The principal can reduce information rents and increase profits by controlling the input.

This result provides insight into the organization of the broiler industry not offered by existing explanations of input control and its role in industry development. The contractor's provision of variable inputs is sometimes viewed as a historical coincidence; early contractors were feed companies or other creditors. When chicken farmers could not pay their debts from raising a flock, the creditor would purchase another flock and supply the farmer with feed and other inputs. The returns from this additional flock would be credited against the farmer's debt. In this situation, farmers already owned the necessary capital equipment. According to the historical accident view, ongoing contracts have simply continued this division of inputs. Knoeber (1989) credits processor input control with encouraging innovation, since the processor then received the benefit of the innovation in inputs. Since processors operate on a much larger scale than individual growers, they had greater incentives to innovate. In this view, input control contributed to the rapid technical progress of the broiler industry. Knoeber and Thurman (1995) examine the risk properties of broiler contracts on a per flock basis. They find that the processor's provision of chicks and feed significantly shifts price risk from growers to the processor. This paper highlights the incentive effects of input *control*, rather than focusing on the effects of input *provision*. Input control extends to situations where the principal specifies the inputs to be used, as well as cases where the principal

[5] Agricultural cooperatives, for example, are formed in part to provide individual farmers with the advantage of large-scale input purchases and output sales enjoyed by processors and other large-scale entities.

actually provides the input. Thus, in the broiler industry input control considerations can explain incentives underlying both the processor's provision of chicks and feed and his specification of the capital equipment that growers must use.

This distinction is a particularly important one in agricultural production contracts. Goodhue and Rausser (1998) and Goodhue (1997) apply a more general version of this result, where exerting control over inputs and the production process is costly for the principal. A comparison of contract provisions and practices in the fresh and processed tomato industries supports the generalized version of this theoretical finding. Fresh tomatoes are a much more lucrative crop, on a per unit basis, and are much more costly to produce. Further, the effects of input decisions on the quality and value of the final product are less easily measured at the time the shipper takes possession of the tomatoes than is the case for processing tomatoes.

These differences lead to the theoretical prediction that the principal will exercise more control over the production process for fresh tomatoes than for processed tomatoes. An examination of representative production contracts for the two crops shows that this is indeed the case. The fresh tomato packer chooses the variety planted, while the processor and grower jointly choose the variety. The processor specifies only that the grower must adhere to relevant pesticide laws, while the packer includes a best practices clause as well. The packer monitors the crop while it is in the field, and may require the grower to modify his growing practices based on crop monitoring observations. The processor, in contrast, monitors fields only to aid in planning the processing season. The shipper exercises a greater degree of control over the harvesting process than the processor does.

These examples show that input control is an important consideration in agricultural production contracts. Further, the evidence suggests that incentive effects are an important consideration when principals design contracts.

4. Optimal Sharing Rules with Common Uncertainty

Modeling common and idiosyncratic production uncertainty in conjunction with risk-averse agents provides a theoretical context for evaluating the risk-sharing properties of another common contract feature; relative agent compensation. Holmstrom (1982) demonstrates that if agents are homogeneous then each agent's inputs, output and the average output may be used to calculate an optimal sharing rule. This section is an extension of Holmstrom's work. It investigates and compares the minimum amount of information required for the principal to implement an optimal sharing rule for heterogeneous agents depending on whether agents' types are known or unknown. An optimal sharing rule is defined as one that will induce agents to provide the efficient amount of effort.

To establish a base for evaluating relative contracts in the presence of production risk and agent heterogeneity, the following results are derived. First, in Proposition 2 if agents are heterogeneous but types are known to the principal, then an optimal sharing rule may be calculated on based on the information contained in the output, type and input use for each agent and average output. Second, in Proposition 3 if agents are heterogeneous and the principal does not know types then a sharing rule based only on an individual's own inputs and output and the average level of inputs and output sacrifices information relative to a sharing rule based on the data contained in the entire vector of individual inputs and outputs.

In order to evaluate the role of production risk, the modeling framework will be modified in the following ways: First, agents are assumed to be risk-averse in income. Second, production now has a stochastic component, due to two unobservable shocks. One shock, η, is common to all agents, and the other shock, ϵ_i, is idiosyncratic to agent i. All shocks are independent of each other, and have means of zero. The introduction of these unobservable shocks means that even if types are known, actions can no longer be inferred from output, so that the principal now faces a hidden action problem in conjunction with the hidden information problem she faced previously. With these modifications, the utility of agent i and expected output for agent i have the following form:

$$U_i = u(y_i) - d(a_i) \tag{4}$$

$$\mathbb{E}[x_i] = a_i^{\phi} Q_i^{\rho} t_i \tag{5}$$

where $u' > 0$, $u'' < 0$, $d' > 0$, and $d'' > 0$. The actual realization of x_i depends on the realization of the common and idiosyncratic production uncertainty random variables. Under these assumptions, expected utility is strictly concave in output for an agent.

The principal faces the following econometric problem: she must use her observed data, agents' outputs, to estimate individual agents' effort levels. In the presence of common uncertainty, the output of all other agents conveys information about the effort of any given agent by providing additional information regarding the realization of the level of common uncertainty.

Holmstrom (1982) establishes that sufficient statistics for the vector of observables may be used to specify sharing rules for homogeneous agents that weakly Pareto dominate sharing rules based on the observables themselves, and that, if a sharing rule is based on a set of signals that are not a sufficient statistic for the observables, then it will be Pareto dominated by a sharing rule based on the observables or a sufficient statistic for them. Hence, in order to show that a vector of signals, γ, may be the basis for an optimal sharing rule, it is enough to show that γ is a sufficient statistic for the vector of agents' outputs.[6]

[6] This paper uses Holmstrom's definition of sufficiency which is based on the statistical decision theory sufficiency condition: "A function $T_i(y)$ is said to be *sufficient for y with respect to a_i* if there exist functions $h_i(.) \geq 0$, $p_i(.) \geq 0$ such that: $g(y, a) = h_i(y, a_{-i}) p_i(T_i(y), a)$ for all y and a in the support of g." (Holmstrom, 1982, p.330-1.)

Proposition 2. *If agents' types and the amount of the input provided to each agent are known, then average output is a sufficient statistic for the vector of individual outputs and an optimal sharing rule for each agent may be implemented based on his own output and average output.*

The proposition is a simple variation on Holmstrom's proposition 8, so it is stated without proof.

The proposition below establishes that Holmstrom's result regarding the sufficiency of an individual's own output and average output for the entire vector of outputs when calculating the optimal sharing rule for that agent does not extend to the case where agents are heterogeneous and their types are unknown. When agents' types are unknown, the average output contains less information about the realization of the common uncertainty than the entire vector of outputs. Essentially, the actual level of output for each agent provides information about his type, which in turn increases the informativeness of the estimate of η. This information is lost when the average is used as the basis for the estimator, lowering the estimator's precision.

Proposition 3. *If each agent may be one of two types (high and low) and each agent's type is unknown to the principal, the estimator of η computed using average output has lower precision than the estimator of η computed using the entire vector of observations on output, so that the average output is not a sufficient statistic for the vector of individual outputs.*

Proof. The principal estimates the realization of the common uncertainty using the outputs produced by n agents, who may each be one of two unobservable types. For each individual i, the distribution of η is drawn from the distributions associated with the two possible agent types with probability p for the high type, $0 < p < 1$, and probability $q = 1 - p$ for the low type. For each individual the estimate of the realized η is the difference between actual realized output and expected output:

$$\hat{\eta}_i = x_i - (q a_l^{\phi} Q_l^{\rho} t_l + p a_h^{\phi} Q_h^{\rho} t_h) \tag{6}$$

Given the entire sample of n observations, $\hat{\eta}$ is the maximum likelihood estimator, which has the following mean and variance:

$$\mathbb{E}[\hat{\eta}] = \bar{x} - (qa_l^{\phi}Q_l^{\rho}t_l + pa_h^{\phi}Q_h^{\rho}t_h) \tag{7}$$

$$var(\hat{\eta}) = \frac{(q^2+p^2)\sigma_\epsilon^2}{n} \tag{8}$$

where $\bar{x}$ is the average realized output. The precision of the sample mean is the inverse of its variance, or $\frac{n}{(q^2+p^2)\sigma_\epsilon^2}$.

We will compare the precision of $\hat{\eta}$ with the precision of the estimator of η obtained when the sample mean of output, $\bar{x}$, is used, denoted $\hat{\hat{\eta}}$. The density of $\hat{\hat{\eta}}$ has the following form:

$$f(\hat{\hat{\eta}}|\mu,\sigma^2) = \sum_{m=0}^{n} g(m|n,p)\frac{1}{(2\pi)^{0.5}\sigma_\epsilon}exp(-\frac{1}{2}\frac{(\bar{x} - \frac{n-m}{n}a_lQ_l^{\rho}t_l + \frac{m}{n}a_hQ_h^{\rho}t_h - \mu)^2}{\frac{\sigma_\epsilon^2}{n}}) \tag{9}$$

where $g(m|n,p)$ is a binomial distribution and x represents the number of high ability agents out of a total sample size of n. The variance of $\hat{\hat{\eta}}$ is $\frac{\sigma_\epsilon^2}{n}$.

Comparing this to the variance of the estimator of η when the information from the entire sample is used, we see that since $0 < p < 1$

$$\frac{\sigma_\epsilon^2}{n} > \frac{(q^2+p^2)\sigma_\epsilon^2}{n} \tag{10}$$

The variance of $\hat{\hat{\eta}}$ is higher than the variance of $\hat{\eta}$. Equivalently, the precision of the estimator based on the sample average is lower than the precision of the estimator based on individual observations.

□

In the case of broiler contracts, the bonus formula uses the unadjusted average as the basis of comparison, and does not account for an individual's type. The resulting estimator of the common uncertainty is not the maximum precision estimator available to the processor for heterogeneous growers.[7] The minimum variance estimator would utilize information regarding individuals' types to remove the maximum amount of common uncertainty. Similarly, using an unadjusted industry average as a executive performance benchmark does not fully utilize information available regarding differences in firm performances over time.

Why do these principals use an inefficient estimator? If agents are risk averse, principals could reduce their risk-compensating payments to agents by improving the precision of their estimator

[7] The estimator would have maximum precision if growers were homogeneous. Goodhue (1997) uses actual contract outcomes for a representative broiler industry contract to test and confirm empirically that the processor indeed employs an insufficient estimator for the realization of the common production uncertainty when calculating grower compensation.

of the common uncertainty. The next section considers various reasons why principals may choose this inefficient estimator. It demonstrates that the profit-maximizing estimator may not be the highest-precision estimator.

5. Common Uncertainty and Information Rents

There are a number of reasons that a principal may not use the most precise estimator of the realization of the common uncertainty. First, it is possible that agents are not risk-averse. Second, it is possible that agents are not significantly heterogeneous in a way that affects their productivity. A third possibility is that the principal's choice of a relative compensation measure may be due to the need to collect information; it might be prohibitively expensive to learn about agents' types, or the information may need to be collected over time. This explanation is unsatisfactory in the context of the broiler industry; processors have information regarding grower abilities, which they use when assigning flocks (Goodhue 1997).

The remainder of this section examines another solution: it may not be profit-maximizing for the principal to remove the maximum amount of uncertainty from agents' payment streams. When agents are heterogeneous and risk averse, reducing risk has two effects: risk compensation payments decline, and information rents paid to highly productive growers increase. The interactions of these two effects will distort the optimal levels of expected output. Intuitively, heterogeneity and risk interact in this way because when agents are heterogeneous the principal must induce them not only to take an action but to take the desired action given their type. When it is very costly to satisfy incentive compatibility constraints relative to the cost of satisfying individual rationality constraints, the benefits of reduced risk payments to all agents may be outweighed by the cost of increased information rents to high ability agents.

In this framework, the principal chooses her profit-maximizing contract menu facing three sorts of costs: neoclassical production costs, or the cost of inducing a risk-neutral agent of known type to produce a given level of output; incentive costs, or the distortionary effects of hidden information combined with stochastic output; and the cost of risk, or the risk premiums that risk-averse agents must receive to be induced to undertake a given action. Under some conditions, the profit-maximizing contract menu may use a lower-precision estimator than the best-available estimator of the realization of the common uncertainty.

Intuitively, this occurs when the benefits of reducing the information rents paid to the high ability agents outweigh the costs of compensating all agents for the increased income risk they must bear. When agents face increased income risk, both a low ability agent and a defecting high ability agent produce less for a fixed contract $\{W_l, T_l\}$. Consequently, the principal can reduce T_h in the corresponding high ability contract, since information rents are lower. However, x_l and x_h are both lower, reducing total revenues and profits. The new equilibrium contract will sacrifice some of this reduction in information rents in order to increase production and revenues, even though this implies an increase in agent compensation due to the increased risk. Under the higher risk contract pair, increasing the marginal incentive provided to the low ability agent is less costly in terms of its effect on incentive compatibility for the high ability agent than it is under the low risk contract pair. Under some conditions, the net effect of the ensuing adjustments in wages, transfers and desired output will be increased profits.

Proposition 4. *The estimator that maximizes the principal's profits may differ from the statistically efficient estimator.*

The proposition states that there are cases where the information rent-reducing effects of increased uncertainty outweigh its risk compensation-increasing effects on the principal's total profits. Table 1 and Table 2 provide examples where profits increase and decrease with increased variance of the principal's estimator of the common uncertainty, respectively.[8]

[8] For convenience, the simulation abstracts from considerations related to input use. The following functional forms and parameter values were used to produce the table: product price = 897.47294 , $p = 0.2851$, $n = 45$, $\bar{U} = 5$, $\sigma_\eta = 0.8450$, $\sigma_\epsilon = 8.0857$, $U(y_i) = (y_i)^\alpha - \frac{\sigma_y^2(1-\alpha)}{2 * y_i}$ where $y_i = W_i x_i + T_i$, $v(a_i) = a_i^\mu$, $x_i = a_i t_i$, $\alpha = 0.855$, $\mu = 1.295$ and $t_l = 1$.

TABLE 1. Profits Increasing with $\sigma_{\hat{\eta}}^2$

Var.	First Best High Variance Estimator	First Best Low Variance Estimator	Second Best High Variance Estimator	Second Best Low Variance Estimator
$\sigma_{\hat{\eta}}^2$	0.0190	0.0112	0.0190	0.0112
π	18260879.2442	18260896.2468	17727473.2347	17727471.7771
x_h	2449.8136	2449.8137	2449.0960	2449.0852
w_h	897.4375	897.4375	897.4350	897.4283
T_h	-1380905.6494	-1380905.7478	-1379110.6052	-1379094.3924
y_h	817649.0121	817649.0083	818793.9424	818784.0495
σ_{yh}^2	7320725.5326	7315253.3460	7320224.8702	7314647.3810
$c.v.(y)_h$	0.0033	0.0033	0.0033	0.0033
U_h	5.0000	5.0000	69.7930	69.7863
x_l	30.5346	30.5351	1.1580	1.1579
w_l	830.8586	830.9127	4.6412	4.6409
T_l	-14858.3574	-14860.4932	8.6837	8.6812
y_l	10511.5421	10511.5222	14.0584	14.0548
σ_{yl}^2	5595006.9256	5590394.9352	174.7576	174.5663
$c.v.(y)_l$	0.2250	0.2249	0.9403	0.9401
U_l	5.0000	5.0000	5.0000	5.0000
x_{hl}			91.0948	91.0907
y_{hl}			431.4718	431.4204
σ_{yhl}^2			1344.0824	1344.7767
$c.v.(y)_{hl}$			0.0850	0.0850
U_{hl}			69.7930	69.7863
$t_h = 15.5124$, $\mu = 1.855$				

The first thing to note about the examples is that x_h is distorted downward from its first-best level in each second-best case. With risk-averse agents, their marginal incentives depend on total income, so transfers affect production incentives. Total utility depends on the variance of income, which is a function of the components of piece-rate income, x_i and W_i, and total income. Due to

TABLE 2. Profits Decreasing with $\sigma^2_{\hat{\eta}}$

Var.	First Best High Variance Estimator	First Best Low Variance Estimator	Second Best High Variance Estimator	Second Best Low Variance Estimator
$\sigma^2_{\hat{\eta}}$	0.0190	0.0112	0.0190	0.0112
π	12985867.3068	12985889.2828	12526766.8400	12526767.0869
x_h	1706.5217	1706.5217	1705.8775	1705.8775
w_h	897.4567	897.4568	897.4567	897.4567
T_h	-975416.9748	-975417.0432	-973778.8422	-973778.8424
y_h	556112.4084	556112.4049	557172.3556	557172.3016
σ^2_{yh}	6912672.7639	6906810.0475	6912381.9106	6906518.2332
$c.v.(y)_h$	0.0047	0.0047	0.0047	0.0047
U_h	5.0000	5.0000	69.7002	69.6998
x_l	26.3075	26.3082	1.1706	1.1706
w_l	860.4825	860.5170	5.3471	5.3472
T_l	-13680.0688	-13681.6229	8.7093	8.7062
y_l	8957.0553	8956.9989	14.9687	14.9656
σ^2_{yl}	6001053.8436	5995807.0568	231.9100	231.6910
$c.v.(y)_l$	0.2735	0.2734	1.0174	1.0171
U_l	5.0000	5.0000	5.0000	5.0000
x_{hl}			77.4962	77.4963
y_{hl}			423.0925	423.0917
σ^2_{yhl}			1107.1392	1107.6037
$c.v.(y)_{hl}$			0.0786	0.0787
U_{hl}			69.7002	69.6998
$t_h = 14.0124$, $\mu = 1.90$				

these interdependencies, it may be less costly to adjust x_h and W_h from the first-best levels than to further adjust a_l, W_l, T_h and T_l, as is the case in these examples. Note that the distortion in x_h is relatively small compared to the distortion in x_l. [9]

[9] The direction of these scale adjustments is not monotonically linked to the direction of the change in profits.

Two parameter values differ between Table 1 and Table 2: the type parameter for the high ability agent, and the disutility of effort parameter, which is constant across types. In Table 1 there is a wider ability gap between the two types and a lower disutility of effort than in Table 2. Given the values of the other parameters, the decrease in the disutility of action increases the amount produced for a given marginal incentive by all agents, but the increase is larger for high ability agents than low ability agents. The increase in the high ability agent's type parameter increases this gap. These effects influence the net outcome for the principal when she increases the variance of her estimator of the common uncertainty.

What is driving the increase in profits with the increased variance of the common uncertainty estimator in Table 1? A wider ability gap and a lower disutility of effort increases the benefit to the principal of increasing the marginal incentives for both types. That is, it costs less to induce each type to produce the same output with higher income risk when the disutility of effort is lower. With a larger ability parameter for the high ability agent, the benefit of increasing his marginal incentive is even larger in terms of its effects on revenues. Comparing the outcomes under the equilibrium contracts, the coefficient of variation of income for a defecting high ability agent is a larger share of the coefficient of variation for a low ability agent in Table 1. In other words, the changes in the variance of the estimator has a larger effect on the utility of a defecting high ability agent than it does on the utility of a low ability agent. Thus the increase in risk lowers information rents relatively more under the parameters in Table 1. Again comparing the two examples, note that in Table 1 the adjustment to the increased risk involves increasing the equilibrium levels of output for both high and low ability agents. The cost of maintaining incentive compatibility was reduced enough by the increased risk borne by agents to lead to increased total output in equilibrium. This scale effect contributed to increased profits.

The difference between agent ability levels and the disutility of effort are not the only parameters that affect the impact of increased variance of the principal's estimate of the common uncertainty on profits. The share of high ability agents in the total agent population and risk aversion parameters are also important. The more high ability agents there are, the greater the relative importance of information rents. The net effect of changes in risk aversion is indeterminate, and depends on other parameters. Under the functional forms specified for the example, increased risk aversion increases the cost of compensating agents for any level of risk and for any change in the risk borne for a given level of output. On the other hand, it also reduces the information rents obtained by a defecting high ability agent for any contract.

As mentioned in the previous section, a sample of contract outcomes under a representative broiler contract showed that broiler processors do not use the best available estimator of the common uncertainty when compensating growers (Goodhue 1997). This section illustrates that it may actually increase processors' profits to use a higher-variance estimator instead.

6. Conclusion

Two empirically important contract features, the principal's control of inputs and choice of a relative compensation benchmark, may be motivated by agent incentive considerations in the presence of moral hazard, agent heterogeneity and risk aversion, rather than simply being theoretically imperfect real world strategies. In particular, when agents are heterogeneous and risk averse, the principal's choice of a relative compensation measure affects the high-ability agent's information rents, in addition to affecting the risk compensation that must be paid to all agents. If the information rent effect is significantly important, the principal may maximize profits by not using the highest-precision estimator of the realization of common uncertainty. This finding provides an explanation for broiler processor's use of a statistically inefficient estimator of the realization of common uncertainty in their calculation of grower compensation.

The analysis demonstrated that the principal reduces high ability agents' information rents by controlling inputs. An extension of this result is to introduce a cost of input control for the principal. The costlier it is to control the input, relative to the information rent reduction, the less likely it is that the principal will choose to control the input. A comparison of the degree of control exerted by principals in representative fresh and processed tomato contracts verifies this prediction empirically. These findings regarding empirically important contract features provide ways to analyze organizational strategies more completely within an agency theoretic framework.

References

Goodhue, Rachael E., "Agricultural Complementarities and Coordination: Modeling Value Differentiation and Production Contracting." PhD dissertation, University of California at Berkeley, Berkeley, California 1997.

——— **and Gordon C. Rausser**, "Product Quality and Off-Farm Control: Fresh and Processed Tomatoes," February 1998. Working Paper.

Grossman, Sanford J. and Oliver D. Hart, "An Analysis of the Principal-Agent Problem," *Econometrica*, January 1983, *51* (1), 7–45.

Holmstrom, Bengt, "Moral Hazard in Teams," *Bell Journal of Economics*, Autumn 1982, *13* (2), 324–340.

Khalil, Fahad and Jacques Lawarree, "Input versus Output Monitoring: Who is the Residual Claimant?," *Journal of Economic Theory*, June 1995, *66* (1), 139–157.

Knoeber, Charles R., "A Real Game of Chicken: Contracts, Tournaments, and the Production of Broilers," *Journal of Law, Economics and Organization*, 1989.

——— **and Walter N. Thurman**, "'Don't Count Your Chickens...': Risk and Risk Shifting in the Broiler Industry," *American Journal of Agricultural Economics*, 1995.

Lehnert, Andreas, Ethan Ligon, and Robert M. Townsend, "Liquidity Constraints and Incentive Contracts," November 24 1997. Working Paper.

Maskin, Eric and John Riley, "Optimal Auctions with Risk-Averse Buyers," *Econometrica*, November 1984, *52* (6), 1473–1518.

——— **and** ———, "Input versus Output Incentive Schemes," *Journal of Public Economics*, October 1985, *28* (1), 1–23.

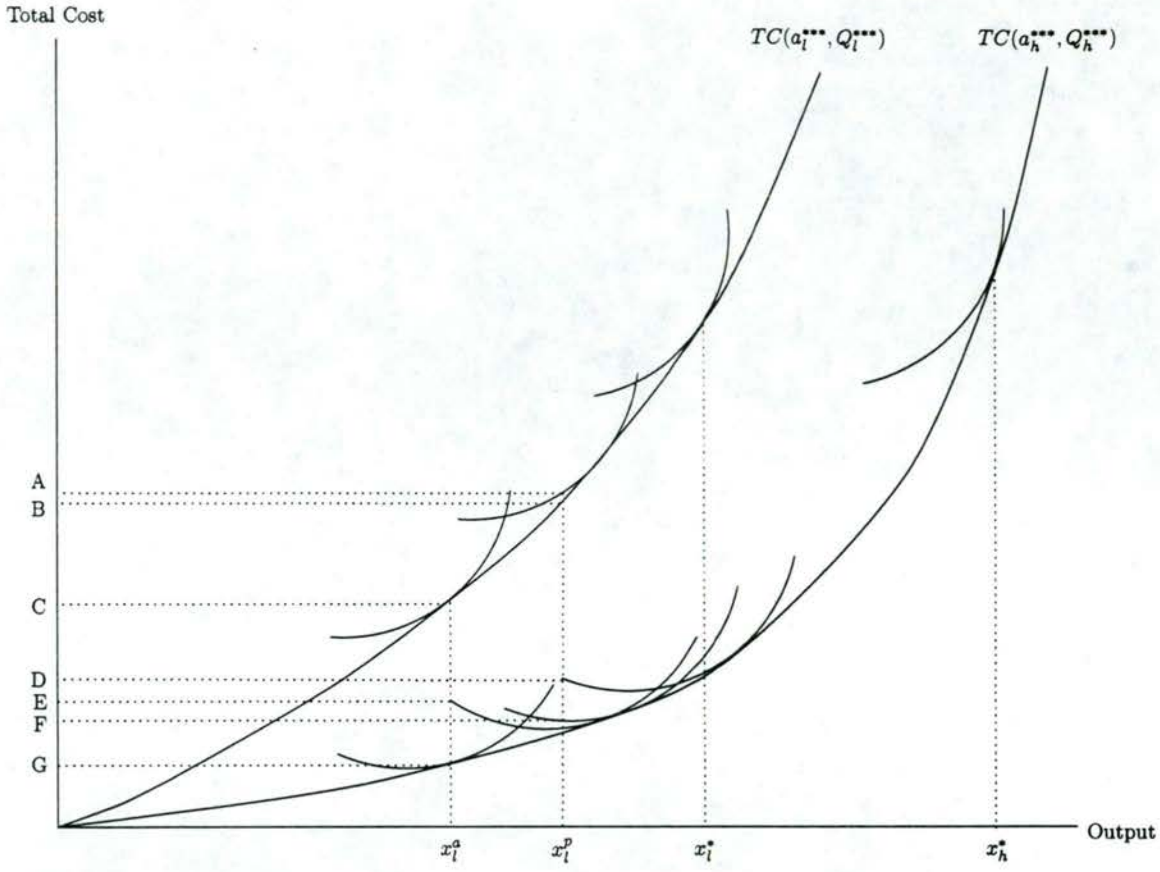

FIGURE 1. The Effects of Input Assignment