Optimal Incentives under Moral Hazard and Heterogeneous Agents: Evidence from Production Contracts Data

Pierre Dubois
University of Toulouse (INRA, IDEI) and CEPR
and
Tomislav Vukina
Department of Agricultural and Resource Economics
North Carolina State University

Paper prepared for presentation at the XIth Congress of European Association of Agricultural Economists - EAAE, August 24-27, 2005
Copenhagen, Denmark

OPTIMAL INCENTIVES UNDER MORAL HAZARD AND HETEROGENOUS AGENTS: EVIDENCE FROM PRODUCTIONS CONTRACTS DATA

Abstract: In this paper we develop an analytical framework for the estimation of the structural model parameters of an incentive contract under moral hazard with heterogeneous agents. Using micro level data on swine production contract settlements, we confirm that contract farmers are heterogenous with respect to their risk aversion and that this heterogeneity affects the principal's allocation of production inputs across farmers. Assuming that contracts are optimal, we obtain estimates of a lower and an upper bound of agents' reservation utilities. We show that farmers with higher risk aversion have lower outside opportunities and hence lower reservation utilities.

Keywords: Contracting, Heterogenous Agents, Moral Hazard.

JEL Classification: D82, L24, Q12, K32, L51.

Introduction

In many business environments, including agriculture, economic agents often interact with each other repeatedly and business is conducted using a series of short-term contracts. The use of contracts to vertically coordinate the production and marketing of agricultural commodities has become common practice in many agricultural sectors including livestock, fruits and vegetables, tobacco, etc. To solve the asymmetric information problems between processors (principals) and independent farmers (agents), the majority of contracts use high powered incentives schemes to compensate farmers. Another interesting characteristic of many production contracts is that all agents contracting with the same principal are operating under formally identical contract provisions (Levy and Vukina, 2002). However, explicitly uniform contracts may not necessarily guarantee that all agents are treated equally. When the principal and agents contract repeatedly, an explicitly uniform but incomplete contract leaves a possibility for the principal to treat agents differently after learning about their types (abilities, risk aversions, costs of effort, etc.). Typically, these contracts specify a general payment formula that expresses the agent's reward as a function of his performance but in which the base payment and the incentive power of the contract depend on the provision of some inputs by the principal. Introducing the choice of these strategic variables as part of the contract design, the principal is able to change what appears to be a uniform contract into individualized contracts tailored to fit agents' preferences or characteristics.

The objective of this paper is to study this contract design problem, to present a method that would allow us to identify and estimate the structural parameters of the moral hazard model, and to test predictions aimed at assessing the empirical reliability of the model. In order to identify the heterogeneity among agents, we assume that they have different risk aversion attitudes and that their preferences are observed by the principal. In the empirical part of the paper we use the panel data containing individual settlements of production contracts. The data originates from a company that contracts the production of live hogs with independent farmers. Our analysis explains an apparent anomaly frequently observed in many agricultural contracts which manifests itself in the principal's use of seemingly uniform contracts for the purposes of governing the relationships with heterogeneous agents.

The literature concerned with empirically testing contract theory related to this paper follows two distinct approaches. One line of research takes contracts as given and model the behavior of the principal and the agents under the observed contractual terms without using any optimality argument about the contract design. For example, interesting studies in labor economics of Paarsch and Shearer (2000; 2004) use the observations on incentive contracts and longitudinal individual outputs in order to estimate how effort responds to incentives provided by piece rate contracts. They do not study the optimal design of these incentive contracts and do not use the identifying power of the contract optimality. However, in some sense, part of their study is related to ours because they also make an assumption about the contract design and use it to identify the heterogeneity of agents with respect to their cost of effort.

However, their assumption is not optimality-specific, since they assume that the employer cannot discriminate between workers according to their observable cost of effort but rather that the contract is designed to satisfy at least the participation constraint of the less able worker. Other papers within the same paradigm include, for example, Dubois and Vukina (2004); and Abbring et al. (2003). They take advantage of the fact that they can observe the actual contracts and perhaps some changes in the contract forms, which enable them to test various implications of moral hazard.

The other line of research in empirical testing of the contract theory takes the reverse perspective and assumes that contracts are optimal. Then, it derives predictions about the determinants of some observed contract parameters and test those predictions empirically. This approach is often used when one does not observe all of the exact contractual terms agreed upon between a principal and an agent. A good example of this approach is the empirical work on sharecropping contracts where the goal is usually to test between alternative theories of contract design, for example between transaction cost versus risk sharing explanations (Allen and Lueck, 1994; Dubois, 2002; Ackerberg and Botticini, 2002).

Our paper goes further than most empirical contract theory papers in the sense that it presents the combination of the above two approaches. In particular, we first empirically check several testable implications of contract theory without assuming the contract optimality thanks to the fact that we observe the true contract parameters. Being able to observe all relevant contract stipulations allows us to model the agent's behavior in a way that is consistent with the assumption that contracts are either optimal or suboptimal. After modelling the agent's behavior, we look at the principal's decisions and contract design, and assuming that contracts are optimal, we derive implications of the principal's optimal decisions. In particular, using the identifying assumption that contracts are optimal, we are able to obtain estimates of a lower and an upper bound of agents' reservation utilities. We confirm that contract farmers are heterogenous with respect to their risk aversion parameters and that this heterogeneity affects the principal's allocation of production inputs across farmers. We also show that farmers with higher risk aversion have lower outside opportunities because lower reservation utilities.

Industry description and data

Swine production in the United States is characterized by an increasing presence of vertically integrated firms (called integrators) that contract the production (grow-out) of hogs with independent farmers. The contract production is dominated by large national companies (Smithfield Foods, Premium Standard Farms, etc.,) that run their businesses through smaller profit centers that issue contracts, supply inputs and slaughter finished animals.

A production contract is an agreement between an integrator company and a farmer (grower) that binds the farmer to specific production practices. Different stages of production of animals are typically covered by different contracts and farmers generally specialize in the production of animals under one contract. The most frequently observed contracts in the swine industry are single production stage contracts such as farrowing contracts, nursery contracts and especially finishing contracts. All production contracts have two main components: one is the division of responsibility for providing inputs, and the other is the method used to determine grower compensation. Growers provide land, housing facilities, utilities (electricity and water) and labor and are also responsible for manure management and disposal of dead animals. An integrator company provides animals, feed, medications and services of field men. Companies also own and operate feed mills and processing plants and provide transportation of feed and live animals. When it comes to specifying integrator's responsibilities for providing inputs, the terms of the contract are intentionally vague. The integrator decides on the volume of production both in terms of the rotations of batches on a given farm as well as the number (density) and weight of incoming animals (feeder pigs) inside the house. A typical scheme for compensating growers in finishing contracts is based on a base plus bonus payment per pound of gain (live weight) transferred, where a bonus payment reflects some efficiency measure such as feed conversion.

The data set used in this study is an unbalanced panel from Martin (1997). It contains a sample of contract settlement data for individual growers who contracted the finishing stage of hog production with an integrator in North Carolina. The data set spans the period between December 1985 and April 1993, for a total of 802 observations. Each observation represents one contract realization, i.e., the payment received and the grower performance associated with one batch of animals delivered to the integrator's processing plant. There are 122 growers in the data set and the number of observations per grower ranges from 2 to 37.

The size of the grow-out operation (the number of finishing houses) varies across growers between one and five houses. All houses under contract have approximately the same capacity. The median density of a house is 1,226 hogs per house and the mean density is 1,234 hogs per house. The contract coverage varies across farms and time. Sometimes one contract will cover multiple houses on a given farm, other times each house will be covered by a separate contract. In cases when multiple houses are covered by one contract, the grower payment is calculated by treating all houses as one unit. The coverage of the contract is determined by the timing of the placement and genetic composition of feeder pigs. The animals covered by the same contract have to be placed on a given farm at the same time and have to have similar genetic characteristics. The average length of the production cycle is approximately 19 weeks. Counting one additional week for the necessary cleanup gives a maximum of 2.6 batches of finished hogs per house per year.

The particular finishing contract that generated the data is fairly representative for the industry as a whole. The contract requires that growers furnish fully equipped housing facilities and that they follow the management and husbandry practices specified by the integrator. The contract guarantees the grower a minimum of 7 batches of feeder pigs and is automatically renewed unless cancelled in writing. The integrator provides the grower with feeder pigs, feed, medication, veterinary services and services of the field personnel. The quality of all inputs as well as the time of placement of feeder pigs and shipment of grown animals are exclusively under control of the integrator.

The compensation to grower i for the batch of hogs under contract t, as the payment for husbandry services and the housing facilities rental, is calculated on a per pound of gain basis with bonuses earned on a per head basis. The bonus is based on the difference between the individual grower's feed conversion, expressed as pounds of feed divided by pounds of gain $\frac{F_{it}}{q_{it}}$, and a standard feed conversion ratio ϕ. If the grower's ratio is above the standard, he receives no bonus and simply earns the base piece rate α multiplied by the total pounds gained q_{it}. If the grower's ratio is below the standard ratio, the difference is multiplied by a constant β to determine the per head bonus rate. The total bonus payment is then determined by multiplying the bonus rate by the number of pigs marketed, where the marketed pigs $(1 - m_{it})H_{it}$ are those feeder pigs that survived the fattening process and m_{it} is the animal mortality rate. Mathematically, the exact formula for the total compensation is:

$$R_{it} = \alpha q_{it} + \max[0, \beta(\phi - \frac{F_{it}}{q_{it}})(1 - m_{it})H_{it}] \tag{1}$$

During the period covered by the data set some parameters of the payment mechanism (1) have changed. The base piece rate varied with the type of feeder pigs placed on a grower farm. For commingled feeder pigs $\alpha = 0.0315$, whereas for integrator's own nursery feeder pigs $\alpha = 0.0275$. Also, as a result of technological progress in nutrition and housing design, the feed conversion standard was lowered from $\phi = 3.50$ to $\phi = 3.35$. However, after the lower feed conversion standard was introduced, the higher standard of 3.50 remained in effect for commingled pigs. Consequently, we have three different payment schemes: $(\alpha = 0.0315, \phi = 3.50)$, $(\alpha = 0.0275, \phi = 3.50)$ and $(\alpha = 0.0275, \phi = 3.35)$. All observed feed conversion ratios are below the benchmark feed conversion (ϕ), so the truncation of the bonus payment at zero can be harmlessly ignored and the payment scheme simplified as

$$R_{it} = \alpha q_{it} + \beta(\phi - \frac{F_{it}}{q_{it}})(1 - m_{it})H_{it}. \tag{2}$$

In addition to individual grower contract settlement data, the proposed methodology requires the integrator-level price data for the inputs and the output. However, such data is not available. Instead we use the regional market prices for feed, feeder pigs and finished hogs, also obtained from Martin (1997). The feed prices are quarterly figures for the Appalachian region, the feeder pig prices are monthly observations for North Carolina and the market prices for finished hogs are monthly prices received by North Carolina farmers for barrows and gilts.

Model

We model the integrator-grower relationship in a principal-agent framework. The timing of the contractual game played between the principal and the agent is as follows. The principal (integrator) proposes the contract to the agents (growers) on a *take-it-or-leave-it* basis. The contract specifies the division of responsibilities for providing inputs and the payment formula. The integrator is required to provide animals (feeder pigs) and feed and the grower is required to provide housing for animals and labor (exert effort). After the grower observes the payment formula and the number and the weight of incoming feeder pigs supplied by the integrator, she accepts or reject the contract. A grower that accepted the contract then exerts effort.

The tasks performed by the grower are not perfectly observable by the integrator, who therefore faces a moral hazard problem in the delegation of production tasks. The incentives to the grower to behave according to the principal's objective are provided through the payment scheme which always includes a particular type of bonus (premium) mechanism. In our data, the bonus depends on a perfectly observable and verifiable performance measure which is the feed conversion ratio. The agent's payment (2) can then be written as a linear function of the performance measure, i.e. the feed conversion ratio $f_{it} = \frac{F_{it}}{q_{it}}$, such that

$$R_{it} = \tilde{\alpha}_{it} - \tilde{\beta}_{it} \left(f_{it} - \phi \right) \tag{3}$$

where the fixed component ($\tilde{\alpha}_{it}$) and the slope ($\tilde{\beta}_{it}$) of this linear function depend on some parameters as

$$\begin{aligned} \tilde{\alpha}_{it} &= \tilde{\alpha}_{it} \left(\kappa_{0it}, H_{it} \right) = \alpha q_{it} = \alpha \left[\kappa_{it} \left(1 - m_{it} \right) - \kappa_{0it} \right] H_{it} && (4) \\ \tilde{\beta}_{it} &= \tilde{\beta}_{it} \left(H_{it} \right) = \beta (1 - m_{it}) H_{it} && (5) \end{aligned}$$

with κ_{it} being the weight of outgoing finished hogs, κ_{0it} the weight of incoming feeder pigs, and H_{it} the heads of animals placed on the farm. When the principal proposes the contract to the agent, he proposes the payment scheme (3) where parameters $\tilde{\alpha}_{it}$ and $\tilde{\beta}_{it}$ are known. Thus, at the time the agent has to accept or reject the contract, the contractual payment consists of a fixed payment $\tilde{\alpha}_{it}$, and a premium part which is tied to the performance $(\phi - f_{it})$ with the known incentive power $\tilde{\beta}_{it}$. After accepting the contract, the agent exerts effort. While some additional random shocks may affect the weight and mortality of animals, we assume that the parameters of this affine function are fixed at the time the grower chooses his effort and that the only source of risk comes form the performance in terms of feed conversion.

The assumption that the parameters $\tilde{\alpha}_{it}$ and $\tilde{\beta}_{it}$ depend on conditions and variables known and observed by the grower when he chooses his effort is reasonably realistic. First, the grower always observes the number H_{it} and the weight κ_{0it} of feeder pigs when they arrive on the farm. The grower also knows that the pigs are grown until they reach their target weight κ_{it}. Finally, the grower can accurately judge the mortality rate m_{it} by observing the genetic make-up and the overall condition of feeder pigs delivered to the farm and the density at which they are stocked. Empirically, we see that there is actually very little variation in mortality rates given H_{it} and a bit more variation in the weight of finished animals.

Agent's behavior

We assume that grower i's preferences over revenue R_{it} and effort e_{it} at period t are described by the utility function $U_i(R_{it} - C(e_{it}))$ which is public and observed by the principal. $C(.)$ is a positive increasing function implying that effort is costly. We assume that growers exhibit constant absolute risk aversion such that $U_i(R_{it} - C(e_{it})) = \frac{-1}{\theta_i} \exp -\theta_i(R_{it} - C(e_{it}))$ where θ_i is the absolute risk aversion parameter and that the stochastic revenue is normally distributed. Under these assumptions, grower i's expected utility can be expressed as an increasing concave function of a mean-variance criterion (which corresponds to the certainty equivalent value of revenue) and his maximization problem can be written as:

$$\max_{e_{it}} W_i(R_{it}, e_{it}) = ER_{it} - \frac{\theta_i}{2} VarR_{it} - C(e_{it}) \tag{6}$$

where agent's revenue R_{it} comes from the payment (3) received from the integrator and the coefficient $\theta_i > 0$ measures the absolute risk aversion.

First, let's specify how the observed outcome stochastically depends on the unobservable grower effort and assume that

$$f_{it}(e_{it}) - \phi = (\lambda - e_{it})u_{it} \tag{7}$$

where λ reflects some fixed ability parameter of grower i, e_{it} is the costly effort which improves (reduces) the feed conversion ratio, and u_{it} is an i.i.d. (across growers and periods) production shock with mean 1 and variance σ^2. This specification shows that a unit of effort is worth one unit of feed conversion ratio which gets transformed into revenue through $\tilde{\beta}_{it}$. Since the cost of effort is monetary, it must be in the same units as revenue, hence we specify

$$C(e_{it}) = \gamma \tilde{\beta}_{it} e_{it}$$

where $0 < \gamma < 1$.

Next, using (3) and (7) we can write the agent's certainty equivalent of net revenue as

$$W_i(R_{it}, e_{it}) = \tilde{\alpha}_{it} - \tilde{\beta}_{it} [Ef_{it} - \phi] - \frac{\theta_i}{2} \tilde{\beta}_{it}^2 Var[f_{it}] - \gamma \tilde{\beta}_{it} e_{it} \tag{8}$$

and the first order condition for the maximization problem in (6) becomes

$$-\frac{\partial}{\partial e_{it}} Ef_{it} - \frac{\theta_i}{2} \tilde{\beta}_{it} \frac{\partial}{\partial e_{it}} Var[f_{it}] = \gamma$$

Given (7), it is clear that

$$\begin{aligned} Ef_{it} - \phi &= \lambda - e_{it} \\ Var[f_{it}] &= (\lambda - e_{it})^2 \sigma^2 \end{aligned}$$

which gives the following expression for the optimal effort level:

$$e_{it}^* = \frac{1 - \gamma}{\sigma^2 \theta_i \tilde{\beta}_{it}} + \lambda. \tag{9}$$

As standard in incentive problems, equation (9) reveals that more risk averse growers, i.e. those with higher θ_i, exert lower equilibrium effort, and also that stronger incentives power $(-\tilde{\beta}_{it})$ increases effort. Notice also that our specification implies that optimal effort is only affected by the incentives power of the contract $(-\tilde{\beta}_{it})$ and not by the constant part of the payment $\widetilde{\alpha}_{it}$. This result has an important consequence for the equilibrium strategy that the integrator (principal) would pursue when it comes to deciding how many feeder pigs to allocate to each grower (agent) according to his risk aversion.

Principal's choices

Now, we model the principal's behavior taking into account the agent's optimal response. We assume that the principal is risk neutral and maximizes the expected profit per grower.

The integrator's profit function is given by:

$$\pi_{it} = pQ_{it} - w_F F_{it} - R_{it}\left(H_{it}, \kappa_{0it}\right) - w_H\left(\kappa_{0it}\right) H_{it} \tag{10}$$

where p is the market price of hogs, $Q_{it} = \kappa_{it}\left(1 - m_{it}\right) H_{it}$ is the total live weight removed from the grower's farm, $R_{it}\left(H_{it}, \kappa_{0it}\right)$ is grower payment, w_F is the market price of feed and $w_H\left(\kappa_{0it}\right)$ is the market price of feeder pigs of weight κ_{0it}.

By deciding how many feeder pigs (H_{it}) of which weight (κ_{0it}) to place on a grower's farm, the principal can vary the contract parameters $\tilde{\alpha}_{it}$ and $\tilde{\beta}_{it}$. As mentioned before, the contracts between the integrator and all agents have the same structure (summarized by the payment scheme (2)), but the allocation of integrator-supplied inputs among growers of different characteristics is not stipulated in the contract and the integrator can choose them unilaterally. Within the class of contractual payments that are observed in the data, varying the quantity and quality of production inputs across growers allows the integrator to use his bargaining power in designing individual incentive contracts for each grower. Notice that in modelling the principal's behavior one can either use a constrained optimality argument by saying that the principal has to choose within the class of payment functions that are empirically observed. This approach will generate some prediction about the principal's "constrained" optimal choices. Alternatively, one can also argue that principals are not legally constrained to use any particular form of payments to agents and therefore those payment schemes that are observed are in fact optimal. Then, one can use the identifying assumption that the observed contracts are actually optimal which would enable the identification of an additional heterogeneity dimension across agents.

As required by the theory, optimal contracts should depend on agent's preferences and her outside opportunities. In particular the incentive power of the contract in such a moral hazard environment should depend on the particular trade-off between risk sharing and incentives that depends on the agent's preferences, whereas the fixed component of the contractual payment should depend on the agent's reservation utility. As agent's preferences (in particular risk aversion) and reservation utilities (depending on outside options and preferences) are likely to be heterogenous, we expect that the principal will tailor particular incentive contracts according to the agent's types.

The problem faced by the integrator is to choose the contract parameters in order to maximize its profit under the incentive compatibility and individual rationality constraints of the agent. Assuming that reservation utilities are not time-varying, this can be formally described as follows:

$$\max_{H_{it}, \kappa_{0it}} E\pi_{it} = E[pQ_{it} - w_F F_{it} - R_{it}\left(H_{it}, \kappa_{0it}, e_{it}\right) - w_H\left(\kappa_{0it}\right) H_{it}]$$

subject to

$$\begin{aligned} EU_i(R_{it} - C(e_{it}^*)|\kappa_{0it}, H_{it}) &\geq \overline{U}_i \\ e_{it}^* &= \arg\max_{e_{it}} EU_i(R_{it}\left(H_{it}, \kappa_{0it}\right) - C(e_{it})|\kappa_{0it}, H_{it}) \end{aligned}$$

where $\overline{U}_i$ is the reservation utility of agent i, and where we denote $R_{it}\left(H_{it}, \kappa_{0it}, e_{it}\right) = \tilde{\alpha}_{it}\left(\kappa_{0it}, H_{it}\right) - \tilde{\beta}_{it}\left(H_{it}\right)\left(f_{it}(e_{it}) - \phi\right)$.

Using the certainty equivalent of the agent's utility, the principal's maximization program is thus equivalent to

$$\max_{H_{it}, \kappa_{0it}} E\left[\pi_{it}\left(H_{it}, \kappa_{0it}\right)\right] = E[pQ_{it} - w_F F_{it} - R_{it}\left(H_{it}, \kappa_{0it}, e_{it}\right) - w_H\left(\kappa_{0it}\right) H_{it}]$$

subject to

$$
\begin{aligned}
W_i(R_{it}(H_{it},\kappa_{0it}), e^*_{it}) &\geq \overline{W}_i \\
e^*_{it} &= \arg\max_{e_{it}} W_i(R_{it}(H_{it},\kappa_{0it}), e_{it})
\end{aligned}
$$

where $\overline{W}_i = U_i^{-1}(\overline{U}_i)$ and the function $W_i(.)$ is defined as in (8). Now, one can incorporate the incentive constraint in the profit function of the principal by replacing the effort level by its optimal value. Then the maximization problem of the principal becomes:

$$\max_{H_{it},\kappa_{0it}} E\pi^*_{it}(H_{it},\kappa_{0it}) = E[pQ_{it} - w_F F_{it}(e^*_{it}) - R_{it}(H_{it},\kappa_{0it}, e^*_{it}) - w_H(\kappa_{0it}) H_{it}]$$

subject to

$$W_i(R_{it}(H_{it},\kappa_{0it}), e^*_{it}) \geq \overline{W}_i$$

where π^*_{it} is the profit function incorporating the incentive constraint.

Assuming the contracts are optimal, one does not have to solve the above principal's problem to determine the equilibrium contractual terms $(H^*_{it},\kappa^*_{0it})$ as a function of observed variables because one can use the directly observed values $(H^*_{it},\kappa^*_{0it})$. However, one has to determine whether the participation constraint is binding or not. If the principal can only choose κ_{0it} and H_{it} to maximize profit and if he has to use the payment formula in (2), then there is no reason for the participation constraint to be binding. Actually, one can see that the choice of κ_{0it} and H_{it} moves the parameters of the linear payment $\tilde{\alpha_{it}}$ and $\tilde{\beta_{it}}$ the same way as the principal could do by choosing $\tilde{\alpha_{it}}$ and $\tilde{\beta_{it}}$ directly, but unlike in the standard principal-agent models, κ_{0it} and H_{it} also change some other part of the principal's profit function. However, if manipulating the choice variables makes the participation constraint not binding, the principal can easily make it binding by adding or subtracting a fixed transfer T_{it} to the agent's revenue R_{it}. Adding such a constant does not change the incentive constraint (as shown by the expression for the optimal effort (9)), thus the principal can perform the maximization program by incorporating only the incentive constraint and then ask for a fixed transfer from the agent in case the participation constraint is not binding.

Therefore, since we exactly observe the contract agreed between the principal and the agent, and the fact that the observed contract is assumed optimal, one can deduce that the solution to

$$(H^*_{it},\kappa^*_{0it},T^*_{it}) = \underset{\substack{H_{it},\kappa_{0it},T_{it} \\ s.t.\ W_{it}(R_{it}(H_{it},\kappa_{0it})-T_{it},e^*_{it})\geq\overline{W}_i}}{\arg\max} \{E\pi^*_{it}(H_{it},\kappa_{0it}) + T_{it}\}$$

is such that

$$T^*_{it} = 0 \text{ and } (H^*_{it},\kappa^*_{0it}) = \underset{H_{it},\kappa_{0it}}{\arg\max}\{E\pi^*_{it}(H_{it},\kappa_{0it})\}$$

and that participation constraint is binding

$$W_{it}(R_{it}(H^*_{it},\kappa^*_{0it}), e^*_{it}) = \overline{W}_i. \tag{11}$$

Therefore, in the following sections, we will first characterize the solution to the principal's maximization program $\max_{H_{it},\kappa_{0it}} \{E\pi^*_{it}(H_{it},\kappa_{0it})\}$, and then exploit the binding participation constraint (11) obtained under the assumption that contracts are optimal in order to derive testable implications.

In order to characterize the principal's maximization program $\max_{H_{it},\kappa_{0it}} \{E\pi_{it}(H_{it},\kappa_{0it})\}$, we need to examine the functional forms of the per head cost function for feeder pigs $w_H(\kappa_0)$ and the mortality function $m_{it}(H)$. Towards this objective, we introduce two assumptions:

- Assumption 1: $w_H(.)$ is increasing convex.

Assumption 1 is likely to be satisfied because of the fact that feed conversion rapidly worsen (increases) and therefore the feeding costs progressively increase as animals grow larger (heavier).

The second assumption that we need to introduce deals with the specific form of the mortality function. It is intuitively obvious that the number of animals placed on a grower's farm cannot be infinite given that the housing facilities are of finite size. The mortality rate will be increasing and necessarily approaching 100% when H approaches infinity. This implies that profits will obtain at a maximum for $H < \infty$. To simplify, we assume that the mortality rate function $m_{it}(H_{it})$ is such that the profit function has a unique maximum $(H_{it}^*(\theta_i), \kappa_{0it}^*(\theta_i))$ and do the following assumption:

- Assumption 2: $m_{it}(H_{it})$ is increasing concave with $m''(1-m) + 2m'^2 \geq 0$ and $2m' + m''H > 0$.

If we label $H_{it}^s = (1 - m_{it}(H_{it}))\, H_{it}$ the number of animals shipped (i.e., the number that survive the fattening process), notice that the condition $2m' + m''H > 0$ is simply equivalent to $H_{it}^{s\prime\prime}(H_{it}) < 0$ that is the number of animals survived $H_{it}^s(H_{it})$ is a concave function of the number of animals placed H_{it}. For example, this assumption is satisfied on $[0, 2\eta]$ with the mortality rate function

$$m_{it}(H_{it}) = 1 - \exp -\frac{H_{it}}{\eta}; \text{ with } \eta > 0. \tag{12}$$

Now we are in the position to state the following two results:

Proposition 1: The optimal decisions $(H_{it}^*(\theta_i), \kappa_{0it}^*(\theta_i))$ made by the integrator are such that $\frac{\partial \kappa_{0it}^*}{\partial \theta_i}(\theta_i)$ is positive if and only if the elasticity of survived animals with respect to risk aversion $\frac{\partial \ln H_{it}^{s*}}{\partial \ln \theta_i} = \frac{\partial \ln\left[(1-m_{it}(H_{it}^*))H_{it}^*\right]}{\partial \ln \theta_i}$ is larger than -1 i.e.

$$\frac{\partial \kappa_{0it}^*}{\partial \theta_i} > 0 \Leftrightarrow \frac{\partial \ln H_{it}^{s*}}{\partial \ln \theta_i} > -1$$

Proof: Available upon request. □

Proposition 2: If the following conditions are satisfied:

$$\begin{aligned} p - \phi w_F + \alpha &> 0 \\ \phi w_F - \alpha - w_H'(\kappa_{0it}) &< 0 \\ \frac{\partial \ln \kappa_{0it}^*}{\partial \ln \theta_i} + \frac{\partial \ln}{\partial \ln \theta_i}\left(\frac{\partial}{\partial H_{it}^*}\left(\frac{1}{1 - m_{it}(H_{it}^*)}\right)\right) &< 1 \end{aligned}$$

then, the optimal decisions $(H_{it}^*(\theta_i), \kappa_{0it}^*(\theta_i))$ made by the integrator are such that

$$\frac{\partial H_{it}^*}{\partial \theta_i} < 0 \Rightarrow \frac{\partial \kappa_{0it}^*}{\partial \theta_i} > 0.$$

Proof: Available upon request. □

Now, let's look at the expression of the certainty equivalent measure of agent's utility $W_{it}(R_{it}(H_{it}, \kappa_{0it}), e_{it}^*)$ and replace e_{it}^* by its analytical expression from (9). This yields

$$\begin{aligned} W_{it}(R_{it}(H_{it}, \kappa_{0it}), e_{it}^*) &= \tilde{\alpha}_{it} - \tilde{\beta}_{it}\left[Ef_{it} - \phi\right] - \frac{\theta_i}{2}\tilde{\beta}_{it}^2 Var\left[f_{it}\right] - \gamma \tilde{\beta}_{it} e_{it}^* \\ &= \tilde{\alpha}_{it} + \frac{(1-\gamma)^2}{2\sigma^2\theta_i} - \gamma\tilde{\beta}_{it}\lambda \end{aligned}$$

Referring back to expressions for contract parameters (4) and (5), the measurement error in the weight of animals at the end of the production period implies that $\tilde{\alpha}_{it}$ is observed with an error but not $\tilde{\beta}_{it}$. Let's assume that κ_{it} is thus measured with an error ε_{it} that is supposed to be uncorrelated with κ_{0it}, H_{it}, and independent across observations. The observed weight of finished animals is therefore $\widetilde{\kappa_{it}} = \kappa_{it} + \varepsilon_{it}$ and then the observed variable $\tilde{\tilde{\alpha}}_{it}$ becomes

$$\begin{aligned} \tilde{\tilde{\alpha}}_{it} &= \alpha \left[\widetilde{\kappa_{it}} \left(1 - m_{it}\right) - \kappa_{0it} \right] H_{it} = \alpha \left[\kappa_{it} \left(1 - m_{it}\right) - \kappa_{0it} \right] H_{it} + \alpha \left(1 - m_{it}\right) H_{it} \varepsilon_{it} \\ &= \tilde{\alpha}_{it} + \varsigma_{it} \end{aligned}$$

where $\varsigma_{it} = \alpha \left(1 - m_{it}\right) H_{it} \varepsilon_{it}$. Using this results, it follows that

$$W_{it}(R_{it}\left(H_{it}^{*}, \kappa_{0it}^{*}\right), e_{it}^{*}) = \tilde{\tilde{\alpha}}_{it}^{*} + \frac{\left(1-\gamma\right)^{2}}{2\sigma^{2}\theta_{i}} - \gamma \tilde{\beta}_{it}^{*} \lambda - \varsigma_{it}.$$

Taking into account the fact that the participation constraint is binding (11), and changing the notation for the observed $\tilde{\tilde{\alpha}}_{it}^{*}$ into $\tilde{\alpha}_{it}^{*}$ for simplicity, we obtain that

$$\tilde{\alpha}_{it}^{*} = \Omega_{i} + \gamma \lambda \tilde{\beta}_{it}^{*} + \varsigma_{it} \tag{13}$$

where $\Omega_i = \overline{W}_i - \frac{(1-\gamma)^2}{2\sigma^2\theta_i}$ and $E\left(\varsigma_{it} | \kappa_{0it}^{*}, H_{it}^{*}, \Omega_i\right) = \alpha \left(1 - m_{it}\right) H_{it} E\left(\varepsilon_{it} | \kappa_{0it}^{*}, H_{it}^{*}, \Omega_i\right) = 0$. Since our data provide observations on $\tilde{\alpha}_{it}^{*}$, κ_{0it}, H_{it}^{*} and $\tilde{\beta}_{it}^{*}$, we can state the following result:

Proposition 3: The agent's reservation utility is a weighted sum of Ω_i identified from (13) and Ψ_i (with unknown weight γ)

$$\overline{W}_i = \Omega_i + \left(1 - \gamma\right) \Psi_i$$

where $\Psi_i = \frac{1-\gamma}{2\sigma^2\theta_i}$ will be identified from performance data using expression (16).

• The lower bound $\overline{W}_{\inf}^{i}$ and the upper bound $\overline{W}_{\sup}^{i}$ on the reservation utility of agent i, $\overline{W}_i$, are identified as

$$\Omega_i = \overline{W}_{\inf}^{i} \leq \overline{W}_i \leq \overline{W}_{\sup}^{i} = \Omega_i + \Psi_i. \tag{14}$$

• If $\Omega_i\left(\theta_i\right)$ is non increasing in θ_i, then one can reject that $\overline{W}_i\left(\theta_i\right)$ is increasing in θ_i (even weakly).

• The parameter $\gamma\lambda$ is identified.

Proof: Available upon request. □

Proposition 3 shows that the assumption that contracts are optimal allows the identification of the lower and the upper bound for the reservation utility of agents. Then, one can test the model restriction $\gamma\lambda > 0$, and explore the correlation between Ω_i and θ_i, as well as the relationship between $\overline{W}_{\inf}^{i}$, $\overline{W}_{\sup}^{i}$ and θ_i.

Identification and Estimation Results

Using the panel data described before, we can now estimate the structural model we developed so far. Substituting (9) in (7) yields the formula for the difference between the benchmark feed conversion ϕ and the equilibrium feed conversion

$$\phi - f_{it}^{*} = \frac{1-\gamma}{\tilde{\beta}_{it}\sigma^{2}\theta_{i}} u_{it} \tag{15}$$

which by taking logs gives the following equation

$$\ln((\phi - f_{it})\,\tilde{\beta}_{it}) = \ln(\frac{1-\gamma}{\sigma^2 \theta_i}) + \ln(u_{it}). \tag{16}$$

The individual level parameters θ_i in (16) can be estimated with a linear regression including growers fixed effects. Notice, however, that θ_i's are identified only up to scale since $\ln(\frac{1-\gamma}{\sigma^2}) - \ln(\theta_i) = \ln(k\frac{1-\gamma}{\sigma^2}) - \ln(k\theta_i)$ for any $k > 0$. Nevertheless, once the estimates of θ_i are known, one can test for the heterogeneity of risk aversions across growers.

The estimation of (16) shows that the unexplained variance accounts for around 50% of the total variance. An F test that all $\ln(\theta_i)$ are equal strongly rejects the homogeneity of growers with respect to their risk aversion ($F(121, 680) = 5.34$). The distribution of risk aversion parameters θ_i is characterized by the fact that the median risk aversion is 43% higher than the value of the 25^{th} percentile of the distribution and 21% lower than the value of the 75^{th} percentile of the distribution. These measures are independent of the scale of coefficients and show substantial heterogeneity across growers regarding their risk aversion.

Performance

Our next objective is to test whether the theoretical implications of the model are consistent with the data. We first check whether the sufficient conditions on the mortality function $m_{it}(H_{it})$ that we introduced in Assumption 1 are satisfied. The data does not allow us to estimate function $m(.)$ and its first and second derivatives non-parametrically because the sample size is not large enough for such a demanding estimation but one can use the parametric form (12) for mortality from which it follows that

$$H_{it} = -\eta \ln(1 - m_{it})$$

and then estimate parameter η by least squares. The results show that $\widehat{\eta} = 26,300$ (with the standard error of 445) and the functional fit is quite good with $R^2 = 79\%$. When estimating η's that vary across feeder pigs type, the R^2 goes up to 85% while the estimates of η are $26,000$ (s.e. 638); $27,300$ (s.e. 724); and $15,100$ (s.e. 708). Notice that for the mortality function in (12), the assumption that led to our Proposition, i.e., $2m' + m''H > 0$ is satisfied if $H < 2\eta$. Since the observed values of H_{it} are between 1,100 and 1,500 per house, this condition is easily satisfied.

Next, using the structural estimates of risk aversion parameters θ_i we want to test the main propositions of the paper. First, we want to test whether the integrator supplies more feeder pigs to less risk averse growers by looking at the relationship between H_{it} and θ_i. First, a non-parametric test of independence between H_{it}, or the average over contracts of H_{it} for grower i, and θ_i shows that independence is strongly rejected. The Spearman rank correlation coefficient is negative and strongly significant. Next, a non-parametric estimate of $E(H_{it} \mid \theta_i)$ obtained by using a standard kernel regression method (shown in Figure) clearly indicates that $E(H_{it} \mid \theta_i)$ is a strictly decreasing function of θ_i, and so does a linear regression model (whose results are not reported here).

Next, although the scale of risk aversion is not identified, the elasticity of the number of animals placement with respect to risk aversion is uniquely identified. A non parametric estimation of $E(\ln H_{it} \mid \ln\theta_i)$ shows that the function is linear and the linear regression gives the estimate $\frac{\partial E(\ln H_{it} \mid \ln\theta_i)}{\partial \ln\theta_i} = -0.84$ with a robust standard error of 0.02. This result shows that a 10% increase in absolute risk aversion results in a 8.4% decrease in the number of animals that the integrator would place on the grower's farm. Based on Proposition 2, this result suggest that the weight of feeder pigs should increase with growers' risk aversion. The result is confirmed by by looking at the elasticity of survived animals with respect to risk aversion,

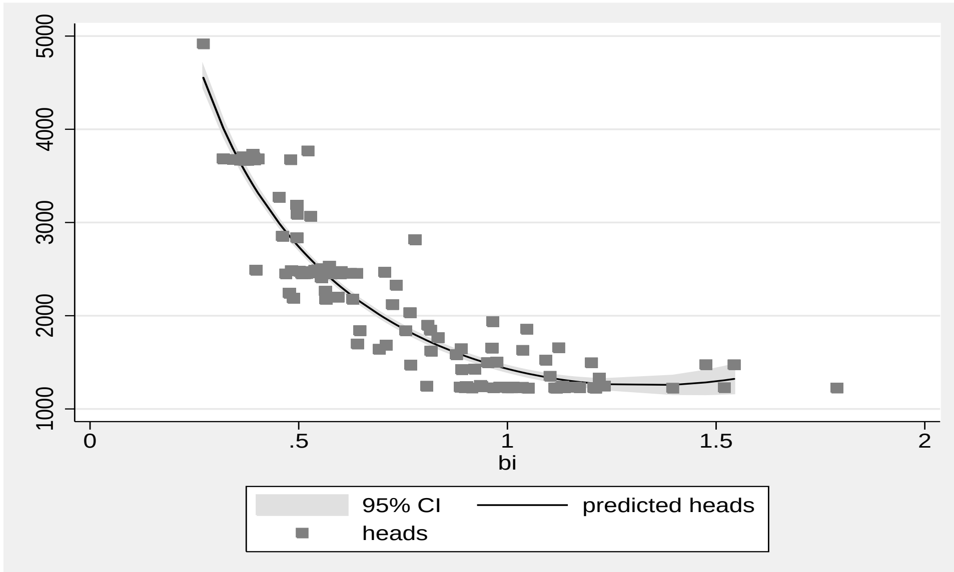

Figure 1. Non parametric estimate of $E(H_{it} \mid \theta_i)$.

i.e., $\frac{\partial E\left(\ln H_{it}^s \mid \ln \theta_i\right)}{\partial \ln \theta_i} = -0.85(0.02) > -1$, which based on Proposition 1, says that the weight of the incoming feeder pigs that the integrator places on a grower's farm would increase with risk aversion if and only if the elasticity of survived animals with respect to θ_i is greater than -1. In fact the results show that $\widehat{\frac{\partial \ln \kappa_{0it}}{\partial \ln \theta_i}} = 0.04(0.01)$. A non-parametric estimate of the weight of incoming feeder pigs conditional on the risk aversion parameter shown in Figure clearly indicates that $E(\kappa_0 \mid \theta_i)$ is an increasing function of θ_i.

Cost of moral hazard

The welfare cost of moral hazard emanates from the fact that contract growers are risk averse and face uncertain income streams. The volatility of income constitutes a direct real cost to growers and can be thought of as the cost of moral hazard in the sense that without moral hazard, integrators could pay growers constant wages to compensate them for their effort in case effort were observable and verifiable. However, obtaining the exact welfare estimates of the cost of moral hazard is impossible because the marginal cost of effort (γ) and the absolute risk aversion coefficient are not identified (θ_i is identified only up to scale). Nevertheless, it is interesting to look at the relationship between the mean and the variance of growers' revenues and their risk aversion parameters. First, 60% of the variance of total payments to growers R_{it} is explained by the *between-growers variance*. Second, a linear regression shows a significant negative relationship between the *within-grower variance* (estimated for each grower along the time dimension of the panel data) and risk aversion. Also, the mean payment is significantly decreasing with risk aversion. The grower level variability of income is such that the average standard deviation is \$3,960 with a median of \$2,856. The above results point out that the cost of moral hazard to growers is substantial.

Moreover, it is important to note that the costs of asymmetric information arise not only from the fact that part of the performance risk (in terms of feed conversion) has to be borne by growers (because they have to be given the correct incentives to perform), but also from the fact that the integrator allocates different number of animals to different growers according to their risk aversions. We anticipate that more risk averse growers would have lower revenues because,

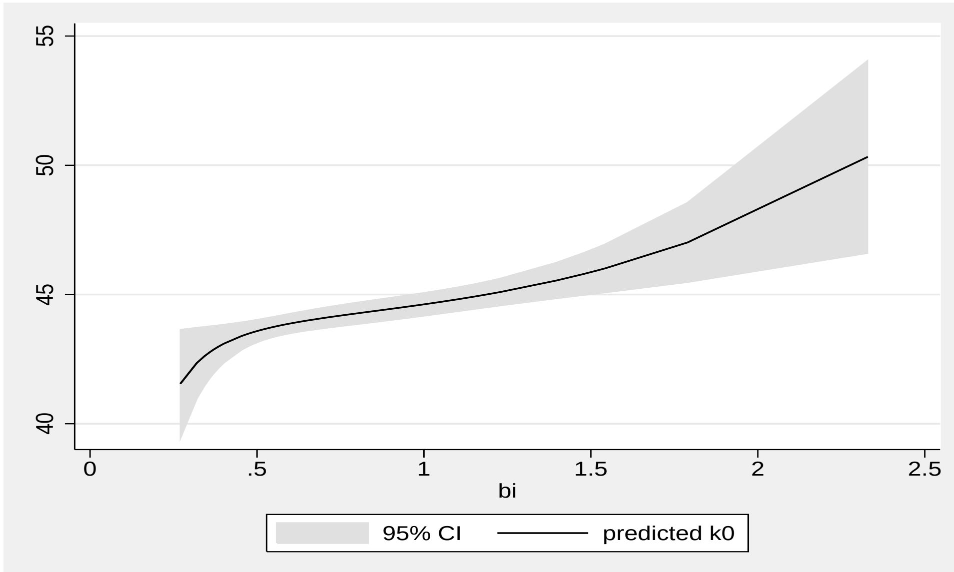

Figure 2. Non parametric estimate of $E(\kappa_0 \mid \theta_i)$.

ceteris paribus, they perform worse in terms of the feed conversion ratio (which reduces their bonus payment), but also because they receive fewer animals compared to the less risk averse growers.

Notice however that the relationship between grower risk aversion and his expected revenue is theoretically ambiguous. Looking at the equilibrium effort equation (9), it follows that the optimal effort decreases with higher risk aversion but also with $\tilde{\beta}$ and hence H_{it}. Therefore, since more risk averse growers receive fewer animals, the overall comparative statics effect of risk aversion on the optimal effort and hence on the expected revenue is undetermined.

The empirical results show that the revenues of more risk-averse growers are less volatile but, also, on average lower. Table 1 shows the average of the means and standard deviations of each grower's revenue R_{it} for different percentiles of the distribution of θ_i. Except for the 50-60 percentiles of the distribution, the relationship shows a negative link between the mean and the variance of grower revenue and risk aversion. This empirical result shows that the net effect of risk aversion on revenue is negative. This net effect is a combination of the indirect effect of risk aversion on the equilibrium values of H and κ_0 via the fixed component and the incentive power of the payment and the direct effect of risk aversion on performance through effort provision.

Table 1: Risk Aversion and Revenue.

% Distribution of θ_i	Mean R_{it} (in US$)	Standard Deviation R_{it}
0-10%	32 709	6 491
10-20%	25 087	5 914
20-30%	23 623	3 969
30-40%	21 227	3 195
40-50%	17 947	2 197
50-60%	18 408	5 971
60-70%	12 906	2 570
70-80%	12 651	3 164
80-90%	11 466	1 999
90-100%	10 995	1 949

Heterogeneity in reservation utilities

Now, we look at the reservation utility of growers (agents) by estimating equation (13), that is

$$\tilde{\alpha}_{it}^{*} = \Omega_i + \gamma\lambda\tilde{\beta}_{it}^{*} + \varsigma_{it}.$$

Using generalized least squares, we obtain consistent estimates of $\{\Omega_i\}_{i=1,..,I}$ and of $\widehat{\gamma\lambda} = 0.025$ (0.0038). This estimate shows that $\gamma\lambda$ is significantly different from zero and positive, indirectly confirming the validity of the model. Recall that both the cost of effort and the ability parameters need to be positive, so their estimated product being positive does not reject the model. A test that all the Ω_i are zero strongly rejects the null hypothesis (F(122, 679)=1.27, p-value=0.0371).

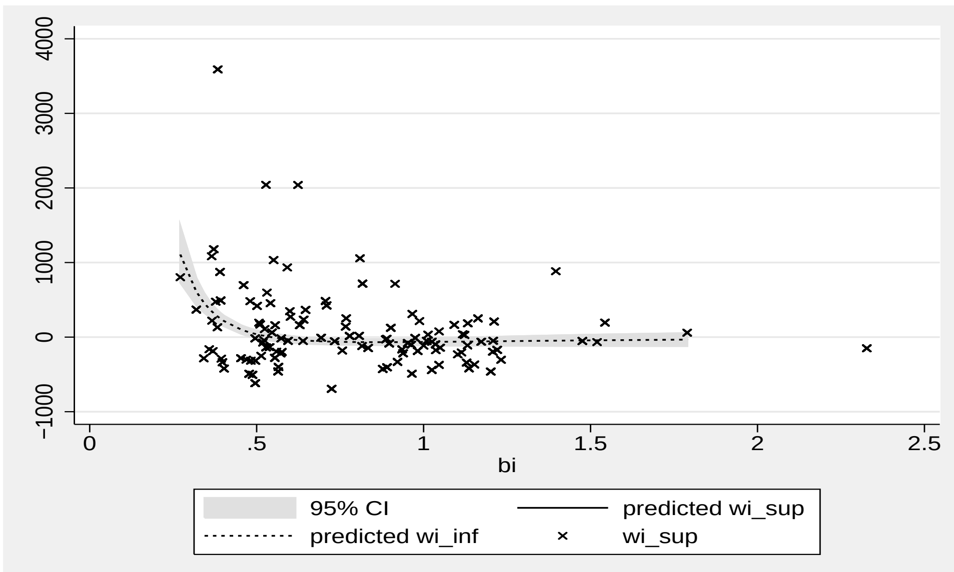

Figure 3. Nonparametric estimate of $E(\overline{W}_{\inf}^{i} \mid \theta_i)$ and $E(\overline{W}_{\sup}^{i} \mid \theta_i)$.

With the obtained estimates we look at the relationship between Ω_i and θ_i. A non parametric estimate of the relationship shows that it is clearly decreasing. Since Ω_i consists of two components: the reservation utility $\overline{W}_i$ and $-(1-\gamma)\Psi_i$ which is increasing in θ_i, it follows that the reservation utility $\overline{W}_i$ has to be decreasing in θ_i. This result implies that agents with higher risk aversion have lower outside opportunities because lower reservation utilities. Figure shows a non parametric estimate of the upper and lower bound estimates of the reservation utility.

Finally, notice that if we considered that agents take into account some risk in κ_{it}, then we should have modified the agent's certainty equivalent revenue by adding the mean-variance value of this additional risk. The shock on the total live weight obtained at the end of the growing period can be denoted η_{it}. Assuming that this random shock is of mean zero and constant variance across agents, the agent's certainly equivalent revenue would become

$$W_{it}(R_{it}\left(H_{it}^{*}, \kappa_{0it}^{*}\right), e_{it}^{*}) = \tilde{\alpha}_{it}^{*} + \frac{(1-\gamma)^2}{2\sigma^2\theta_i} - \gamma\tilde{\beta}_{it}^{*}\lambda - \theta_i var\left[\eta_{it}\right]$$

One can show easily that in this new model, Ω_i would become $\Omega_i = \overline{W}_i - \frac{(1-\gamma)^2}{2\sigma^2\theta_i} + \theta_i var\left[\eta_{it}\right]$. Although, this approach would weaken the possibility to identify the agents' reservation utilities (because the absolute value of θ_i is not identified), the additional term $\theta_i var\left[\eta_{it}\right]$ being increasing in θ_i, would reinforce the fact that $\overline{W}_i$ has to be decreasing in θ_i when Ω_i decreases with θ_i, which has been empirically confirmed.

Conclusion

In this paper we studied the question of optimal contracting under moral hazard when agents have heterogenous preferences. In this case, heterogeneity calls for individually designed contracts, which stands in sharp contrast to what we frequently observe in the economy. The examples of principals using apparently uniform contracts when dealing with heterogenous agents are found in many agricultural sectors, particularly in livestock production contracts such as broilers, turkeys, and hogs. The main elements of all agricultural production contracts are the payment mechanism and the division of responsibilities for providing inputs. The payment mechanism consists almost always of a variable piece rate with bonuses for the efficient use of the principal's supplied inputs and is always the same for all agents. However, contracts never specify the quantity and quality of integrator supplied inputs for each grower. We show that the observed contracts are only formally uniform, and that the principals are using their discretion when it comes to matching inputs with agents of different preferences (risk aversion). Using this variation in contract variables, the principal in fact manages to design the individualized contracts tailored to fit the individual grower's preferences or characteristics.

The paper has two conceptually distinct parts. In the first part we develop an analytical framework for the econometric estimation of the degree of risk aversion of contract producers and carry out its empirical estimation using the individual growers performance data from the swine industry. We found that contract farmers are heterogenous with respect to their risk aversion parameters and that this heterogeneity affects the principal's allocation of production inputs across farmers. The main characteristic of this part of the paper is that it takes the observed contract as given and model the behavior of the agents under the observed contractual terms without using any optimality argument about the contract design.

The obtained results are then used to look at the cost of moral hazard associated with growers' risk aversion. We show that the costs of asymmetric information arise not only from the fact that part of the performance risk has to be borne by growers (because they have to be given the correct incentives to perform), but also from the fact that the integrator allocates different number of animals to different growers according to their risk aversions. More risk averse growers will have lower expected revenues because on average they perform worse, but also because they receive fewer animals compared to the less risk averse growers. These results were confirmed in a variety of different empirical tests.

In the second part of the paper, we look at the principal's decisions and contract design, and assuming that contracts are optimal, we derive the implications of the principal's optimal decisions. Compared to other papers on applied contract theory, this part of the paper stands out in that we use both the assumption of optimality of the contract and the fact that the contract payments are accurately observed in our data. Using the optimality of contracts as an identifying restriction, we were able to obtain estimates of bounds on agents' reservation utilities. We show that farmers with higher risk aversion have lower outside opportunities and hence lower reservation utilities.

References

Abbring, J.H., P.A. Chiappori, J.J. Heckman and J. Pinquet (2003). Adverse selection and moral hazard in insurance: Can dynamic data help to distinguish? *Journal of the European Economic Association* 1: 512-521.

Ackerberg, D.A. and M. Botticini (2002). Endogenous matching and the empirical determinants of contract form. *Journal of Political Economy* 110(3): 564-591.

Allen, D.W. and D. Lueck (1992). Contract choice in modern agriculture: Cash rent versus cropshare. *Journal of Law and Economics* 35(October): 397-426.

Dubois, P. (2002). Moral hazard, land fertility and sharecropping in a rural area of the Philippines. *Journal of Development Economics* 68(1): 35-64.

Dubois P. and T. Vukina (2004). Grower risk aversion and the cost of moral hazard in livestock production contracts. *American Journal of Agricultural Economics* 86 (August): 835-841.

Levy A. and T. Vukina (2002). Optimal linear contracts with heterogeneous agents. *European Review of Agricultural Economics* 29: 205-217.

Martin, L. L. (1997). Production contracts, risk shifting, and relative performance payments in the pork industry. *Journal of Agricultural and Applied Economics* 29 (December): 267-278.

Paarsch, H. J. and B. Shearer (2000). Piece rates, fixed wages, and incentive effects: Statistical evidence from payroll records. *International Economic Review* 41(1): 59-92.

Paarsch, H. J. and B. Shearer (2004). Male-female productivity differentials: the role of ability and incentives. Mimeo.