AGRICULTURAL MARKETS AS RELATIVE PERFORMANCE EVALUATION

BRENT HUETH AND ETHAN LIGON

Abstract. Relative performance schemes (e.g., tournaments) are commonly used in markets for a variety of livestock and processing commodities, while explicit versions of these schemes are rarely (if ever) used in markets for fresh fruits and vegetables and specialty grains. We show how contracts for each of these latter commodities do in fact provide—via a payment mechanism which crucially depends on this existence of a market price—relative performance incentives. In such an environment, compensation that is offered as some increasing function of market revenue implements a relative performance scheme by making each grower's payment an increasing function of his own output, but a decreasing function of others' output.

1. Introduction

Holmström (1982) models the problem facing a group of agents, each of whom produces an observable quantity of some good by taking an unobservable action. A rather robust conclusion from this line of inquiry is that when the production of one agent isn't independent of the production of another, then the compensation awarded the first agent should depend not only on his own production, but also on the output of others. In particular, if agents' production is positively correlated (because of some common shock, for example) then an agent's compensation should be larger when other agents' output is smaller, a phenomenon which Holmstrom terms "relative performance evaluation."

Relative performance evaluation ought to be particularly important for agricultural producers and the various intermediaries who purchase their output. These producers face considerable common risk from aggregate production shocks, and so one would expect that

Date: November 19, 1999.
This research has been sponsored in part by NRI Competitive Grants Program/USDA award #98-35400-6092, but the views expressed in this article are not endorsed by the sponsoring agency.

relative performance evaluation could improve efficiency. In fact, certain types of agricultural intermediaries do use explicit relative performance schemes, as predicted by theory. For example, Knoeber & Thurman (1989) analyze the tournaments (a particular form of relative performace incentives) used by "integrators" in the procurement of live chickens. Tsoulouhas & Vukina (1999) note that similar relative performance schemes are used to procure a wide variety of processing commodities (e.g., livestock, canned fruits and vegetables). In these commodities, relative performance can be used to shield growers against various sorts of common (across producers for a given intermediary) weather, pest, disease, and other production risks.

Despite the importance of relative performance evaluation for some agricultural commodities, for others the explicit use of relative performance seems to rarely (if ever) be observed—leading examples include fresh fruits and vegetables. However, a defining feature of this class of commodities seems to be that growers under contract receive payment which is based on the price of the commodity in downstream markets, which is in turn determined by aggregate supply and demand. Often the payment made to the grower takes the form of a share of total revenue derived from the sale of the grower's produce.

The dependence of an individual producer's compensation on aggregate supply and demand conditions when contracts don't feature relative performance incentives is interesting, since we'd expect an efficient contract to make compensation depend only on variables which convey information about the unobservable actions taken by the grower (Sheldon 1996). Why is the compensation awarded in many agricultural contracts based on current market revenues, even when the current market price used to compute revenue carries little or no information about a particular grower's unobservable actions? In this paper we argue that this seeming puzzle has a simple explanation: compensation based on current market revenue is in fact a form of relative performance evaluation when variation in prices is due to supply shocks, as tends to be the case in markets for agricultural commodities.

To give an intuition for the sort of mechanism we have in mind, consider an adverse weather shock which affects the mean output of all growers of some fresh-market commodity. When common shocks have a similar influence on the output of many growers, relative performance incentives amount to rewarding a particular grower for how he's done relative to other growers; thus, his compensation should be an increasing function of his own output, but a decreasing function of aggregate output. Since, *ceteris paribus*, low aggregate output translates into high prices, giving the grower some increasing function of revenues as compensation satisfies both of these requirements.

In what follows, we examine the extent to which actual contract structures like those described above are consistent with the "optimal" relative performance scheme predicted by theory. In Section 2 we begin by discussing two sorts of real-world contracts we've observed, neither of which has any explicit relative performance incentives, but each of which exposes the producer to aggregate price risk. To understand these contracts, allowing for the possibility of market based relative performance evaluation, we need a model in which output, prices, and grower's compensations are all jointly determined; in short, we need a model of general equilibrium with private information. We construct such a model in Section 3. Though rather simple and specified at the level of primitive preferences, technologies, and endowments, the model delivers a great deal, including a set of equilibrium contracts governing the provision of a finite number of commodities, a non-trivial equilibrium mechanism governing the matching of growers with intermediaries, and well-defined supply and demand schedules indexed by the supply shocks which buffet the economy. The equilibrium contracts in this environment are arbitrary mappings from the state of the economy into a family of probability distributions over allocations—in particular we do not *impose* any relationship between grower compensation, aggregate prices and individual output. We show that the equilibrium contracts will include relative performance incentives, and that one simple way of implementing these contracts is to make compensations depend on price. In Section 4 we display a particular economy from the class of general equilibrium models considered in Section 3. The example economy has only one consumption good, but allows for uncertainty. In equilibrium, the example features partial insurance against idiosyncratic risk and contracts with relative performance evaluation; further, the relative performance incentives in the contract can be provided simply by making compensation contingent on price via revenue. Section 5 concludes.

2. Observed Contracts

Agricultural contracts vary considerably according the commodity being contracted for, but to date very little of this variation has been the subject of research. This may be surprising, given the voluminous research on contracts in other fields of economics (Chambers 1999). In an effort to catalog some of this variation, we describe two forms of contracts, each of which seems to be used for a fairly wide range of commodities. Figure 1 displays stylized versions of these two forms of agricultural contract. The first, which we call a *broker contract*, offers compensation $w_b(p, y) = (1-\tau)py$, where y is net output, τ is a parameter, and p is the market price at or near the time of delivery. The broker contract is common in markets for many fresh fruit and vegetable commodities (Dimitri & Lichtenberg 1999). In these markets, various types of intermediaries sell growers' output on commission to downstream buyers at negotiated prices which depend on current market conditions.

In the top panel of Figure 1, we plot the compensation schedule $w_b(p, y)$ as a function of output for two different values of p—one corresponding to a "good supply shock" (low price) and the other corresponding to a "bad supply shock" (high price)—and as a function of revenue for each shock. The broker contract offers growers compensation which is strictly increasing in output and revenue, and provides only limited protection against aggregate price risk. Also, for a given shock, growers face higher-powered incentives in adverse supply conditions (marginal compensation as a function of output is greater when the aggregate shock is bad and prices are high).

In the bottom panel of Figure 1, we display a *market premium contract*. Here, growers receive some premium δ above the market price so that total compensation is given by $w_p(p, y) = (\delta + p)y$. Market premium contracts are common when the output y is grown under similar circumstances (e.g., manner of planting and cultivation) and in the same geographic region as some standardized commodity with a well defined market. Examples include specialty (e.g., high oil) and seed corn contracts, both of which which offer producers

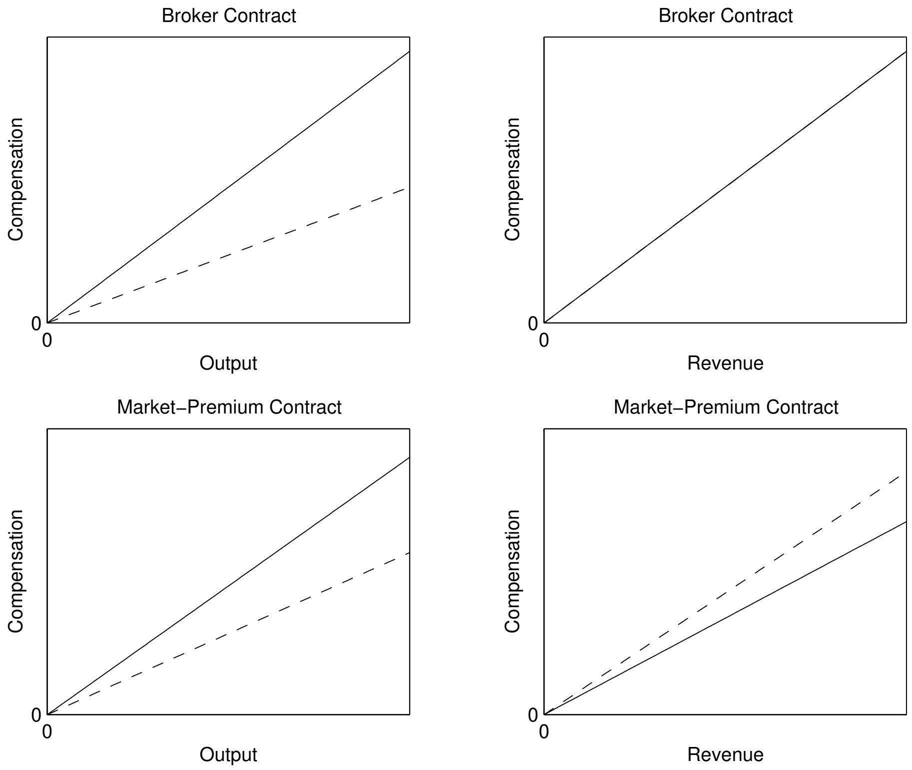

FIGURE 1. Two Stylized Contracts. This figure displays the payment mechanisms offered in two contracts as functions of output and revenue for two different prices. Compensation in the good production state (low price) is indicated by the dashed lines; compensation in the bad production state (high price) is indicated by the solid line.

an "above the board" payment mechanism that provides a premium above the futures market price for commodity corn (Hamilton 1995). The market premium contract exhibits very similar features to the broker contract: both can be written as some function $w(p, y) = g(y) + h(py)$. Also, note an interesting feature of the market premium contract: although compensation depends on a price, the price which is featured in the contract is *not* the market price of the commodity (e.g., seed corn) under contract, but rather the price of a completely different good (commodity corn).

Neither contract in Figure 1 appears to use a relative performance scheme. This is a puzzle since there is considerable scope for such a scheme when producers face common production uncertainty. There is no obvious reason that the same types of output or quality based relative performance schemes mentioned earlier could not be employed here. Moreover, price risk could be reduced by offering a fixed unit price, or in situations where it's feasible, by making payment depend on the difference between the price received for a grower's particular lot and some average market price.[1] Instead, we observe compensation that depends on realized production and on some aggregate market price that may or may not be directly related to sale of a grower's output in some downstream market. More precisely, compensation usually depends directly on revenue, where again the price used to calculate revenue may or may not be directly associated with sale of a producer's output in a downstream market.[2]

One possible explanation for dependence of grower compensation on revenue is that each grower's produce may fetch a different price because of idiosyncratic variation in quality. If a grower's unobserved action affects quality, then price (via revenue) provides a valuable signal regarding the grower's action, and some dependence of compensation on price will actually be a feature of the optimal contract. Hueth & Ligon (1999) investigate this hypothesis for a specific type of contract used in California's fresh-market tomato industry. However, explaining why compensation might depend on price does not explain the apparent absence of relative performance evaluation in fresh market contracts. That is, while moral hazard in the provision of quality may lead to dependence of compensation on price, it should only do so through *idiosyncratic* variation in price. As noted above, considerable market risk could be filtered away by simply making compensation depend on deviations between the price fetched for an individual producer's lot and the average price across a group of producers, or some average market price report. To our knowledge, this type of 'filtering' never occurs. Moreover, in the case of the market premium contract, compensation depends on a price

[1]In the case of fresh fruits and vegetables, the information infrastructure necessary to make this work already exists. The "identity preservation" required to keep track of the price actually received for a grower's crop is ubiquitous, and is already used in the majority of fresh market contracts. The Department of Agriculture reports daily average market prices for many fresh market commodities in various location across the United States. Alternatively, an intermediary could use the average price received on sales for his specific growers.

[2]However, we do not want to convey the notion that this is the dominant or only important feature of agricultural contracts. In fact, such contracts are normally quite non-linear and complicated schedules that include provisions such as a base payment (independent of production and price), quality incentives, short term credit, mixed ownership of productive assets, and so on. Nevertheless, as indicated in the stylized contracts presented above, an important element of most all the contracts we have encountered is that gross revenue often forms the basis for compensation. For further discussion of contracting mechanisms in a cross section of California's fruit and vegetable industry, see Hueth, Ligon, Wolf, & Wu (1999) and Wolf, Hueth, & Ligon (1999).

that is entirely independent of a single grower's actions. In this case, calculating revenue based on the current market price would seem only to *increase* risk to the producer, relative to a compensation rule based on, for example, an expected market price.

3. The Model

In the previous section we've argued informally that exposing producers to aggregate price risk is a way of implicitly providing relative performance incentives. For the argument to hold, we actually need very little. Common supply shocks need to affect producers' output in a similar way (such a shock can't be good for one producer's output and bad for another's), and demand curves need to be downward sloping (so that an increase in aggregate supply results in a decrease in price). However, though aggregate price risk provides *some* relative performance incentives, it is by no means apparent that an *optimal* relative performance evaluation scheme would look anything at all like the contracts pictured in Figure 1. Moreover, in an economy with private information, it's not entirely obvious that downward-sloping demand curves will necessarily obtain, even for normal goods, since in general the producers' compensation may be a non-linear, non-monotone function of output [see, for example, the example of Grossman & Hart (1983)]; it's possible to construct examples in which income effects dominate substitution effects.

In this section and the next, we support our argument that market prices can play an important role in the provision of relative performance incentives by defining the competitive equilibrium for an economy with private information, and then by constructing a competitive equilibrium for a simple example economy. This example economy provides support for our main contention, since in fact demand curves *are* downward sloping in that example, but also allows us to greatly extend our claim: in our example economy, not only does the use of market prices provide relative performance incentives, but the real-world broker contract described in the previous section can be shown to provide the *optimal* relative performance incentives for a particular, plausible specification of preferences and technologies.

3.1. **Preferences and Technology.** There are m commodities in the economy; quantities of each these denoted by y_i, $i = 1, 2, \ldots, m$. We assume there to be a unit mass of identical agents. Each agent is endowed with a unit of time, which is divided between labor and leisure (leisure *not* being among the m commodities listed above). At the beginning of the period, each agent makes a decision regarding how much labor a_i to devote to the production of some good i, and we we assume that agents produce at most a single commodity.

For simplicity, we suppose that labor is the only input to production for any good, and that each agents' production is conditionally independent of other agents' production. However, there is some common unobserved shock ϵ_i specific to the production of each commodity. The common production technology for each commodity i is described by a distribution function $F^i(y|a_i, \epsilon_i)$, with a_i the labor devoted to production of the ith commodity. We assume that larger shocks produce more output, in the sense of first-order stochastic dominance, or that for any $\epsilon' > \epsilon$, $F^i(y|a_i, \epsilon) > F^i(y|a_i, \epsilon')$ for all y and a_i; further, we assume that the support of the distribution function F^i doesn't depend on a_i (though it may depend on ϵ).

Let the labor supply of an agent who produces commodity i by working a units be described by the probability distribution $\Sigma^i(a)$. Then the aggregate supply of commodity i given supply shock ϵ_i will be $\bar{y}_i = \sigma_i \int y dF^i(y|a, \epsilon_i) d\Sigma^i(a)$, where σ_i is the fraction of agents who produce commodity i.[3] The production shocks have a joint distribution function $G(\vec{\epsilon})$, where $\vec{\epsilon} = (\epsilon_1, \epsilon_2, \ldots, \epsilon_m)$, the marginal distribution of the ith shock being denoted $G^i(\epsilon)$.

Each agent's preferences are represented by a utility function which is assumed to take the addilog form, with

$$U(c_1, c_2, \ldots, c_m, \ell) = \ell + \sum_{i=1}^{m} \alpha_i \log(c_i).$$

Here ℓ denotes the agent's leisure, equal to $1 - a$, where a is the labor the agent devotes to production, and the preference parameters α_i are assumed (without loss of generality) to be positive and to sum to one. Each agent chooses some amount of labor a, receives some income w, and then, taking prices $\vec{p} = (p_1, p_2, \ldots, p_m)$ as given, allocates this income to the

[3] For a justification of the point that "probabilities equal proportions" for agents on a continuum, see Kirman (1981) or Judd (1985).

purchase of a bundle of commodities in such a way as to maximize his utility, solving

$$V(w, \vec{p}) = \max_{\{c_i\}} \sum_{i=1}^{m} \alpha_i \log(c_i)$$

subject to the budget constraint $\sum_{i=1}^{m} p_i c_i \leq w$. Note that this defines the function $V(w, \vec{p})$; this is like an indirect utility function, except that we've left aside for now the agent's choice of labor. Since we can add a constant to the agent's utility function without affecting the implied preference ordering, we can rewrite the agent's utility function as $V(w, \vec{p}) - a$. Finally, because of the simple addilog form of the utility function, it's easy to show that $V(w, \vec{p}) = \log(w) + \sum_{i=1}^{m} \alpha_i \log(\frac{\alpha_i}{p_i})$.

Given our assumption of common addilog preferences, the aggregate demand function for the ith commodity can be written as some $\bar{c}_i = \frac{\alpha_i}{p_i}\bar{w}$, where p_i is the spot price of the ith commodity, and where $\bar{w}$ is aggregate income. Note that $\bar{w}$ depends on supply shocks via the influence of such shocks on the agents' income; thus, both supply and demand schedules 'shift' according to what set of shocks $\vec{\epsilon}$ is realized.

3.2. **Firms.** Although each agent is able to engage in the production and marketing on his own, the uncertainty associated with production and the fact that agents are risk averse implies that each might prefer to write a contract with a firm which, by pooling risks across many agents, can reduce the uncertainty faced by any individual agent. We suppose there to be many identical firms, each of which markets a single commodity. The form of the arrangement between firm and producer is that the producer gives all output of commodity i to the firm, which realizes revenue $p_i y_i$ from marketing the producer's output, and compensates the producer by making a contingent transfer w_i. This transfer may depend on the agent's output, of course, but it may also depend on the supply shock ϵ_i, since stochastic dominance with respect to ϵ_i combined with observing the aggregate supply of commodity i implies that the firm can perfectly infer the supply shock ϵ_i.

Any firm which markets commodity i seeks to maximize expected profits by designing a contract which minimizes the costs of inducing a suitable number of agents to (i) sign the contract, and (ii) devote an optimal amount of labor to the production of the ith commodity.

When agents' labor can be observed, this problem is relatively simple. Because of the linearity of the firm's objective and the conditional independence of production, the problem of designing a contract for all agents can be solved simply by considering the problem of designing a contract for one agent. Suppose that a particular agent has a reservation utility u; that is, he has an expected utility of u if he *doesn't* agree to market his output through the firm. To induce the agent to sign the contract, the firm must offer at least this same level of utility.

A contract for the production of commodity i for an agent with reservation utility u takes the form of a mapping and a probability distribution. The first object maps realizations of agent's output y_i and supply shock ϵ_i into a probability distribution over the agent's compensation w_i, and is written as some $\Phi^i(w|y_i, \epsilon_i)$. The second object is the probability distribution over the labor a_i demanded by the firm, $\Sigma^i(a)$; accordingly, a contract is a pair (Φ^i, Σ^i). Both the firm and the agent take as given a price functional $p(\vec{\epsilon})$ which maps the m-vector of shocks into a vector of prices, with $p(\vec{\epsilon})$ denoting the price of the ith commodity. Each firm which produces good i seeks to maximize expected profits, choosing contracts so as to solve

$$\Pi^i(u) = \max_{(\Phi,\Sigma)} \int \left[p_i(\vec{\epsilon})y - w\right] d\Phi(w|y, \epsilon_i) dF^i(y|a, \epsilon_i) dG(\vec{\epsilon}) d\Sigma(a) \tag{1}$$

subject to satisfying the participation constraint

$$\int \left[V(w, p(\vec{\epsilon})) - a\right] d\Phi(w|y, \epsilon_i) dF^i(y|a, \epsilon_i) dG(\vec{\epsilon}) d\Sigma(a) \geq u. \tag{2}$$

Since $V(w, \vec{p})$ is a concave function of w, it's easy to see that the solution to this problem is to insure the agent against variation in output y_i by providing some constant compensation w. In particular, notice that the agent's compensation doesn't depend on the aggregate shocks $\vec{\epsilon}$, and thus doesn't depend on aggregate output or prices.

Matters become more interesting when we admit the possibility that for some production technologies it may not be possible to monitor labor. In this case, revelation principle arguments (Harris & Townsend 1981, Myerson 1982) imply that the problem facing the firm is similar to the problem given above, but with the addition of an incentive compatibility constraint of the form

$$\Sigma \in \underset{\hat{\Sigma}}{\operatorname{argmax}} \int [V(w, p(\vec{\epsilon})) - a] \, d\Phi(w|y, \epsilon_i) dF^i(y|a, \epsilon_i) dG(\vec{\epsilon}) d\hat{\Sigma}(a) \tag{3}$$

This simply requires that it be in the agent's best interests to undertake any labor request Σ made by the firm.

3.3. **Equilibrium.** Recall that when designing the optimal contract, each firm takes as given the price functional $p(\vec{\epsilon})$. Of course, prices are determined by aggregate supply and demand, which in turn depend on the contracts designed by the firms. We need to find an equilibrium; first we provide a definition.

Definition 1. A competitive equilibrium is a price functional $p(\vec{\epsilon})$, an assignment of agents to different industries $\{\sigma_i\}_{i=1}^m$, and a set of contracts $\{(\Phi^i, \Sigma^i)\}_{i=1}^m$ such that

1. Given the price functional $p(\vec{\epsilon})$, contracts for each commodity i solve the firm's problem of maximizing expected profits (1) subject to the participation constraints (2) and incentive compatibility constraints (3).
2. Given the price functional $p(\vec{\epsilon})$, the optimal contracts (Φ^i, Σ^i), and the assignment of agents σ_i, expected profits in every industry i are zero.
3. Given the optimal contracts (Φ^i, Σ^i), and the assignment of agents σ_i, the price functional must clear all commodity markets for all shocks $\vec{\epsilon}$.
4. All agents must be assigned to produce some commodity, so that $\sum_{i=1}^m \sigma_i = 1$.

The first condition is a requirement that firms solve their problem correctly given prices, and the second condition just ensures that each firm earns zero expected profits in doing so. Note that our use of agents' indirect utility functions and the incentive compatibility constraints means that condition (1) *also* implies that agents maximize their expected utilities, taking as given prices and contracts.

The third condition is simply the usual market clearing condition, which equates aggregate demand and supply given any realized supply shock, or

$$\sigma_i \int y dF^i(y|a,\epsilon_i) d\Sigma^i(a) = \frac{\alpha_i}{p_i(\vec{\epsilon})} \sum_{j=1}^{m} \int w d\Phi^j(w|y,\epsilon_j) dF^j(y|a,\epsilon_j) d\Sigma^j(a),\ i = 1,\dots,m,$$

Notice that the spot markets here are only for the m commodities, and takes as given the assignment of labor. Clearing in the labor market is implied by conditions (4) (all agents do something) and condition (2) (no firm makes positive expected profit).

The existence of a competitive equilibrium satisfying Definition 1 follows from arguments advanced by Prescott & Townsend (1984). Although not an entirely transparent feature of this definition it should be clear that in general, equilibrium contracts will depend on the equilibrium price which obtains for each commodity.[4] This feature captures the main point of our paper, which is that "market" price can be used by firms to deliver relative performance incentives to producers when variation in this price is due to aggregate supply shocks. In the following section, we examine this feature further by determining (up to a linear transformation) equilibrium contracts, prices, and consumptions for a simple example economy.

4. An Example

Our description of the equilibrium contracts given in the previous section is in terms of distribution functions (Φ^i, Σ^i). However, the observed contracts described in Section 2 appear to make producer's compensation a deterministic function of output and prices. Fortunately, in many cases of interest the optimal distributions (Φ^i, Σ^i) will turn out to be degenerate, so that both the action taken by each producer and the producer's compensation can be regarded as deterministic functions.

In this section we provide a simple example economy of the sort described in Section 3. We completely characterize production of one of the m goods in this economy, and show that

[4] Compensation depends on the (observable) supply shocks $\vec{\epsilon}$, and for each vector of shocks there is a vector of equilibrium prices $p(\vec{\epsilon})$. Thus, conditioning compensation on the supply shocks is equivalent to conditioning compensation on price.

the equilibrium contract for this particular good looks very much like the "broker contract" described in the introduction.

The economic environment is as described in Section 3, with two additional assumptions. First, we make the simplifying assumption that aggregate income for all agents in the economy, $\bar{w}$, is a constant. This does not imply, of course, that the aggregate income of producers of a particular commodity is constant; we have in mind the case of agricultural commodities in the US, where variation in the incomes of, say, lettuce farmers may be large, but is too small a part of the economy to have an appreciable effect on aggregate GDP.

Second, we specify a particular distribution function for the production of the ith commodity. In particular, let the supply shock ϵ_i take on values greater than or equal to one, and let

$$f^i(y|a,\epsilon_i) = \begin{cases} \frac{y^{\epsilon_i - 1}}{\Gamma(\epsilon_i)} \frac{1}{a} \exp(-ya^{1/\epsilon_i}) & \text{for } y > 0 \\ 0 & \text{elsewhere.} \end{cases} \tag{4}$$

This particular distribution belongs to the family of "gamma-type" distributions. Two properties of this distribution are worthy of note. First, expected output given (a, ϵ_i) is

$$\mathrm{E}(y|a,\epsilon_i) = \epsilon_i a^{1/\epsilon_i},$$

so that expected output is an increasing, (weakly) concave function of a. Second, the likelihood ratio for (4) is given by

$$\frac{f_a^i(y|a,\epsilon_i)}{f(y|a,\epsilon_i)} = \frac{1}{a}\big(\frac{y}{\mathrm{E}(y|a,\epsilon_i)} - 1\big),$$

where we interpret the a which appears in this formula as the action recommended by the intermediary. Heuristically (but inaccurately), since the value of this likelihood ratio is smaller (in absolute value) when a is larger, it's more difficult for an intermediary to detect shirking when the producer is required to work very hard. Similarly, it would be relatively easy to infer that the agent shirked when realized output y very different from expected output, since in this case the likelihood ratio is relatively large.

With this specification of production, it's simple to show that the equilibrium contract for commodity i takes the form of the broker contract, a point we make precise in the following proposition.

Proposition 1. *If output of commodity i is distributed according to (4), and aggregate income is constant, then in equilibrium the compensation awarded a producer of commodity i will almost certainly be a linear function of the revenue derived from the sale of his output.*

Proof. We begin with the following simple lemma. Given our maintained assumption of addilog preferences, this is an immediate corollary to Theorem 1 of Jewitt (1988).

Lemma 1. *Fix the aggregate supply shock ϵ_i. If, given this ϵ_i,*

- **C1:** $\int_0^y F^i(y|a,\epsilon_i)dy$ *is a non-increasing, convex function of a for all y;*
- **C2:** $\int y f^i(y|a,\epsilon_i)dy$ *is a non-decreasing, concave function of a; and*
- **C3:** $f_a^i(y|a,\epsilon_i)/f^i(y|a,\epsilon_i)$ *is a non-decreasing, concave function of y for all a,*

then with probability one the contract which solves the problem facing a firm of type i will feature a contract which satisfies

$$w = \lambda^i + \mu^i \frac{f_a^i(y|a,\epsilon_i)}{f^i(y|a,\epsilon_i)} \tag{5}$$

for some non-negative constants (λ^i, μ^i).

Note that the assumed distribution $f^i(y|a,\epsilon_i)$ satisfies the conditions of the lemma. Thus, for almost all agents the contract which solves the problem facing a firm of type i will feature a compensation schedule of the form

$$w = \lambda^i + \frac{\mu^i}{a}\big(\frac{y}{\mathrm{E}(y|a,\epsilon_i)} - 1\big).$$

Since agents are homogeneous, each agent will be advised to take the same action a. Thus, aggregate supply of the good i will be $\bar{y}_i(\epsilon_i) = \sigma_i \mathrm{E}(y|a,\epsilon_i)$, so that

$$w = \lambda^i + \frac{\mu^i}{a}\big(\sigma_i \frac{y}{\bar{y}_i(\epsilon_i)} - 1\big).$$

In this form, it's very clear that the contract features relative performance incentives, since the aggregate shock now enters the contract only via its influence on average production $\bar{y}_i(\epsilon_i)$.

To write compensation as a function of revenue, recall that addilog preferences imply that aggregate demand for commodity i is given by $\bar{c}_i = \alpha_i \bar{w}/p_i$; since in equilibrium supply equals demand, this implies that

$$w = \left(\lambda^i - \frac{\mu^i}{a}\right) + \frac{\mu^i \sigma_i}{a \alpha_i \bar{w}} p_i y,$$

which is a linear function of revenue, as required. □

Because firm profits are a strictly decreasing function of producers' *ex ante* expected utility u, it's clear that a u exists which generates zero expected firm profits. Thus, we have completely characterized the competitive equilibrium generated by the production technology in (4). This equilibrium is one in which firms use market price in their equilibrium compensation schedules. Moreover, price shows up in compensation via market revenue, which in turn enters compensation linearly.

The equilibrium contract of Proposition 1 looks very much like the "broker contract" discussed earlier, except that the term representing a "base payment" $\left(\lambda^i - \frac{\mu^i}{a}\right)$ will not generally be zero. In the stylized broker contract of Figure 1, there is no base payment. There are two possible explanations for this discrepancy: first, although the base payment term in an optimal relative performance scheme will not generally be zero, it nevertheless *can* be zero. Thus, our model is not inconsistent with the observed contract structure. Second, it's possible that our stylized broker contract is misspecified. That is, there are a wide variety of ways in which a producer might receive some type of base payment (e.g., signing bonus, disaster insurance, other sources of income), but where this payment is not an apparent feature of the contract between a producer and his intermediary.

Note also that relative performance incentives are a common feature of environments not normally associated with "contract" production. For example, consider a producer of some grain commodity who carries crop disaster insurance, and who sells his crop to a local elevator. Normally, the elevator will base her price on the nearby futures price (if one exists), on the cost of handling and storing the producer's grain, and on a desired profit margin. This setting is typical of many commodity grains (e.g., #2 yellow corn) in the United States, and is one where production incentives are quite closely aligned with those predicted by equation (4).

5. Conclusion

We've argued that when common shocks affect production for many farmers, then the information that there *was* a common shock ought to be used in some sort of relative performance evaluation scheme. However, although relative performance incentives are explicitly built into contracts for the provision of a variety of processing commodities (e.g., hogs, broilers, processing tomatoes), such explicit incentives are not an apparent feature of contracts governing production of many items for which a broadly reported "market price"—either for the commodity which is traded under contract, or for a closely related commodity—exists (e.g, fresh fruits and vegetables, specialty grains).

In these latter contracts, producers are typically exposed to considerable price risk associated with annual variation in market conditions. We've argued elsewhere (Hueth & Ligon 1999) that exposure to *idiosyncratic* price risk may provide valuable incentives to produce high quality output for fresh-market commodities (such a scheme is infeasible for a processing commodity); however, this argument cannot explain the exposure to *aggregate* price risk which is frequently observed for these commodities.

We resolve this pair of puzzles by observing that while exposure to aggregate price risk cannot help in the provision of incentives for quality, so long as the equilibrium expected price is a declining function of aggregate output (i.e., so long as demand curves are downward sloping, roughly), and so long as output increases in the face of favorable supply shocks, then exposing producers to variation in aggregate prices has the effect of implicitly providing the relative performance incentives predicted by theory.

Our model suggests that explicit relative performance schemes are of potentially greater value in environments where there is no well defined market price (that contains information about aggregate supply conditions). The observation that commodities for which spot markets have all but disappeared (e.g., broilers, many processing commodities) are the primary users of explicit relative performance schemes seems consistent with this prediction.

Our general conclusion that exposure to aggregate price risk provides a form of relative performance incentives seems quite robust. However, in general there's no reason to suppose that the relative performance incentives provided by price through its impact on revenue will be optimal. Nonetheless, when revenue is a sufficient statistic for a producer's action such a scheme will be optimal; we've provided a simple example of a distribution which has this property, and shown that for this distribution (and with log preferences) compensation should in fact be a *linear* function of revenue.

Whether or not it's strictly optimal, exposing producers of commodities such as fresh fruit or vegetables to aggregate price risk seems likely to provide useful relative performance incentives. Of course, there are other commodities for which this is apt to be a less successful strategy. First, consider extending the model we've presented above so that production shocks are fairly local in their effects, but the commodity is exchanged at a uniform price over large distances. For example, if an adverse shock affects production in California but not in Florida, then exposing producers to aggregate price risk would offer perverse incentives to producers in Florida, and attenuated incentives to California producers. Second, if producer compensation depends on aggregate prices, then it depends indirectly on both aggregate supply shocks *and* aggregate demand shocks; it would never be optimal to expose a risk-averse producer to demand shocks.

This variation in the suitability of using aggregate price risk according to commodity suggests a way in which one could go about testing our hypothesis that such risk sometimes provides useful relative performance incentives. In particular, non-perishable commodities produced in many different locations and traded over large distances may be poor candidates for the sort of scheme we have in mind, if supply shocks are local. Similarly, relative performance incentives for commodities with very uncertain demand should not be provided by exposing producers to price risk. On the hand, most agricultural commodities have a fairly stable demand, and the fresh fruits and vegetables procured under the broker contracts we study are both perishable, and mostly produced in a single state (California).

If exposure to aggregate price risk is important in the provision of incentives for fresh fruit and vegetable producers, then this line of research may have important implications for the design of government subsidized insurance schemes, such as the revenue-insurance scheme discussed by Hennessy, Babcock, & Hayes (1997). In particular, if existing contracts are optimal, then the introduction of such schemes would be expected to reduce the efficiency of producers. On the other hand, the optimal contracts discussed in Section 4 feature a base payment which doesn't depend on aggregate shocks, and it may be possible to interpret proposed revenue insurance schemes as providing something resembling such a non-contingent payment. Thus, combining a revenue insurance payment (which should be contingent only on aggregate shocks) with a flat-rate tax on revenues might be a way to provide the optimal incentives to producers, even in the complete absence of a contract with an intermediary.

References

Chambers, R. G. *Information, incentives and the design of agricultural policies.* In *Handbook of Agricultural Economics*, edited by B. Gardner & G. Rausser. North Holland, (1999) Forthcoming.

Dimitri, C. & E. Lichtenberg. *Quality verification in contracts for agricultural commodities*, (1999). Unpublished Manuscript.

Grossman, S. J. & O. D. Hart. *An analysis of the principal-agent problem. Econometrica* 51(1):(1983) 7–45.

Hamilton, N. D. *A Farmer's Legal Guide to Production Contracts.* Philadelphia, PA: Farm Journal, Inc., (1995).

Harris, M. & R. M. Townsend. *Resource allocation under asymmetric information. Econometrica* 49:(1981) 33–64.

Hennessy, D., B. Babcock, & D. Hayes. *Budgetary and producer welfare effects of revenue insurance. American Journal of Agricultural Economics* 79(3):(1997) 1024–34.

Holmström, B. *Moral hazard in teams. Bell Journal of Economics* 13:(1982) 324–340.

Hueth, B. & E. Ligon. *Producer price risk and quality measurement. American Journal of Agricultural Economics* 81(3):(1999) 512–524.

Hueth, B., E. Ligon, S. Wolf, & S. Wu. *Incentive instruments in fruit and vegetable contracts: Input control, monitoring, measuring, and price risk. Review of Agricultural Economics* In press.

Jewitt, I. *Justifying the first-order approach to principal-agent problems. Econometrica* 56(5):(1988) 1177–1190.

Judd, K. *The law of large numbers with a continuum of iid random variables. Journal of Economic Theory* 35:(1985) 19–25.

Kirman, A. *Measure theory with applications to economics.* In *Handbook of Mathematical Economics*, volume 1. Amsterdam: Elsevier, (1981) .

Knoeber, C. & W. Thurman. *A real game of chicken: Contracts, tournaments, and the production of broilers. Journal of Law, Economics and Organization* 5:(1989) 271–292.

Myerson, R. B. *Optimal coordination mechanisms in generalized principal-agent problems. Journal of Mathematical Economics* 10:(1982) 67–81.

Prescott, E. & R. Townsend. *Pareto optima and competitive equilibria with adverse selection and moral hazard. Econometrica* 52(1):(1984) 21–45.

Sheldon, I. M. *Contracting, imperfect information, and the food system. Review of Agricultural Economics* 18:(1996) 7–19.

Tsoulouhas, T. & T. Vukina. *Integrator contracts with many agents and bankruptcy. American Journal of Agricultural Economics* 81:(1999) 61–74.

Wolf, S., B. Hueth, & E. Ligon. *Policing mechanisms in agricultural contracts*, (1999). Unpublished Manuscript.

Iowa State University

University of California at Berkeley and Member, Giannini Foundation