PRODUCER PRICE RISK AND QUALITY MEASUREMENT

BRENT HUETH AND ETHAN LIGON

Abstract. The produce industry collectively solves an extremely complicated resource allocation problem in which risk-averse farmers grow a product whose market price is often quite unpredictable. Shippers or other intermediaries shield the farmer from much of this risk, permitting fairly efficient production. However, actual contracts between growers and shippers vary considerably across commodities in the residual price risk growers face. We hypothesize that imperfect quality measurement results in a moral hazard problem, and that idiosyncratic variation in the price of the produce reflects, in part, the consumer's assessment of quality. Thus, price provides additional information regarding quality, and as a consequence an efficient contract does not shield growers from all idiosyncratic price risk. We examine this hypothesis for the case of fresh-market tomatoes, and conclude that intermediaries' inability to perfectly observe quality is capable of explaining observed variation in the price risk tomato growers face.

Date: October 29, 1998.

1. Introduction

Producers of fruits and vegetables, especially for the fresh market, operate in an unusually risky economic environment. While these farmers face the same sorts of production risk common to much of agriculture, they also produce a perishable commodity whose price is subject to unusually large fluctuations. Some of this variation in prices is predictable (e.g. seasonal variation), though much of it is not, depending instead on unforeseeable shocks to both supply and demand, or on variation in the quality of produce.

Because producers often operate at a small, specialized scale, one would expect them to be particularly vulnerable to these sorts of risks. And indeed, there is indirect evidence that they are; various intermediaries (shippers, processors, brokers, cooperatives) write contracts with producers which often shield them from much production and price risk. However, such contracts ordinarily are written so that producers and intermediaries both face significant risk. This is somewhat surprising. In a world with complete markets, a risk neutral intermediary ought to insure risk-averse producers against all idiosyncratic risk, as it would be costless to do so (Wilson 1968).

A considerable literature in agricultural economics has examined the influence of risk on various aspects of farm level decisions [e.g., Just (1974), Just and Zilberman (1985)]. However, less attention has been focused on the sources or reasons for this risk. We can think of three reasons why producers might face risk. The first is simply that contracts may not be efficient—there may be unexploited gains to trade in insurance between growers and intermediaries. Yet the produce industry is astonishingly productive and efficient in other aspects, and so it seems doubtful that contracts are grossly inefficient at managing the risk faced by producers.

The second possibility is that either growers are risk neutral or that intermediaries are risk averse. There are reasons to doubt this—it seems unlikely that small, specialized family farms don't care about income risk, and if intermediaries aren't risk neutral, then there would seem to be rents one could earn in what appears to be a very competitive industry. Nonetheless, suppose that intermediaries are risk-averse, and seek to maximize the expected value of some concave function of profits. In Section 2.2 we demonstrate that an efficient contract between grower and intermediary would make the grower's compensation depend on profits realized by the intermediary, not on prices or production.

The third possibility is that there is a moral hazard problem in the fresh produce industry. This seems utterly plausible in the case of production risk—if an intermediary were to make payments to the grower which depended only on acreage planted and not on harvest, the grower would have a powerful incentive to underinvest in costly inputs and labor. However, the case seems much less clear when we consider price risk—why don't intermediaries make a payment which depends only on the quantity and quality of produce measured at the farmgate?

We hypothesize that the reason for price risk is that unobserved investments by the grower (e.g., labor effort, the costly application of fertilizers or pesticides) influence not only the quantity of the grower's output, but also its *quality*. By itself, this wouldn't necessarily expose the grower to price risk. The intermediary may be able to simply condition payments on the quality of the produce—if he can measure it. However, if this measurement is less than perfect, then the grower may well be exposed to price risk. To see why, consider a somewhat stylized description of fresh-market tomatoes (which we expand on in Section 3), drawn from discussions with shippers in the 'mature green' tomato industry in California.

At the beginning of the season, a shipper signs an enforceable contract with a grower. The contract specifies what actions are expected of each party at different points during the season, and specifies a compensation schedule for the grower. Typically, the shipper will supply some inputs and may do the actual harvesting. At harvest time, tomatoes are boxed, and identified by grower and the field they came from; usually this identification is preserved until the tomatoes are purchased by the consumer. The shipper then markets the tomatoes, possibly to the final consumer, but more usually to some other customer (e.g., supermarket, restaurant, broker) in the chain which stands between the grower and consumer. We refer to the price the shipper receives for the tomatoes as the 'downstream price,' or just as the 'price' where we don't fear ambiguity.

During the course of the growing season the grower takes some costly actions (irrigation, pesticide application, etc.) which affect the quality of the tomato a consumer will eventually see at the retail level (and which will consequently affect the price of the produce paid by intermediaries). The shipper doesn't observe all of these actions directly; however, after the harvest the shipper can attempt to measure the quality of the tomatoes (or perhaps a sample of these tomatoes) by examining them for color, rot, size, and so on.

The shipper may attempt to measure quality for two reasons. First, he may want to predict the price his customers will be willing to pay for a particular lot. This prediction may be more or less accurate, but since at the time of shipping the tomatoes are still green, it's unlikely that the shipper can perfectly predict the price his customers will be willing to pay, as there may be considerable uncertainty over how the tomatoes will look when they're ripened. Second, he may want to reward growers for their efforts to produce high-quality tomatoes. Because the shipper can't observe these efforts directly, the best he can do is to try to infer something about the grower's efforts, and conducting some sort of quality measurement may help with this problem of inference.

Of course, the shipper isn't the only one who may be interested in measuring the quality. The shippers' customers are also likely to measure quality (whether formally or informally) and once the tomatoes reach the retail level, consumers may attempt to measure quality by quite intensively examining each individual tomato. Tomatoes which don't measure up to the customer's standards may have to be discounted in order to sell, and so the price at which the tomato eventually sells reveals something about quality, and something about the efforts of the grower to produce a high quality tomato. Although the shipper is unlikely to sell directly to the final consumer, the fact that the retail price will eventually depend to some extent on consumers' assessment of quality means that the downstream price paid by the shipper's customers will depend on the customer's beliefs about the price consumers eventually *will* pay, and thus on the quality assessment of the customer.

The quality assessment of the shipper's customers will often reveal something about the effort of the original grower beyond the information revealed by any quality measurement undertaken by the shipper. At the height of the packing season, the shipper may pack hundreds of tons of tomatoes in a day, and inspection of individual tomatoes is apt to be minimal. However, any given customer will usually buy only a small fraction of the tomatoes sold by the shipper, and so may be able to manage a more intensive inspection. Also, the customer's quality measurement occurs at a later point in time than any measurement done by the shipper, often after the tomatoes have been ripened, and so may reveal aspects of the tomatoes' cultivation which weren't apparent at harvest. Because the shipper can't be sure of quality when he takes delivery, it may be sensible to ask the grower to share the loss if the price at which the produce sells turns out to be low; by exposing the grower to this sort of price risk, the shipper provides some additional incentive for the grower to produce high quality tomatoes.

If quality is not perfectly observable, then standard arguments from contract theory tell us, roughly, that variation in the compensation received by the grower should depend only on variables which contain information regarding quality (Holmström 1979). One set of such variables are physical measurements of various attributes of the produce (or of a sample of the produce); another is the downstream price.When the grower's compensation depends on this price, this is evidence that quality measurement undertaken by the intermediary is imperfect, and that the customer's quality measurement (reflected in prices) has some informational value that an efficient contract will exploit.

In what follows, we investigate the hypothesis that price risk is used to motivate grower attention to quality. We first develop a simple model of contracting in the produce industry, and then compare the theoretical contracts derived from our model with actual contracts from the 'mature green' tomato industry in California.

2. Model

We consider a model in which an intermediary contracts with a risk averse grower to produce a commodity which may vary in quality. We assume that the grower can take costly actions to control the quality produced. If quality is observable and verifiable, we would expect the grower's compensation to depend only on the *ex ante* terms of the contract, and on the *ex post* quality of produce (and perhaps on the *ex post* profits of the intermediary).

In order to abstract from yield risk,[1] we assume that a grower produces a single unit of some agricultural commodity. The grower controls the quality, q, of the commodity,but

[1]The arguments we'll advance hold for *any* fixed yield, so this restriction is made without loss of generality. Of course, if private actions influence yield, this raises some additional interesting contracting issues, but these have been fairly exhaustively covered in the existing literature on principal-agent models (Stiglitz 1974; Holmström 1979).

increased quality comes only at a cost $z(q)$ to the grower, where q is some real number. The cost function $z(q)$ is denominated in the grower's utils, and is increasing, convex, and continuously differentiable. Having produced a commodity of quality q, the grower could choose to sell the commodity; the price received for this commodity is taken to be some random variable $p \in P \subset \mathbb{R}$. The distribution of p is some $F(p|q)$. We assume that the density $f(p|q)$ exists, is strictly positive for all $p \in P$, and is a continuously differentiable function of q. Note that, for simplicity, the distribution of p doesn't depend on aggregate market conditions. The price p depends on the beliefs of the consumer regarding quality q. Thus, the distribution F summarizes the relationship between these beliefs and prices. The grower receives some compensation w, which he values according to some utility function $U(w)$, assumed to be strictly increasing, strictly concave, and continuously differentiable.

2.1. **No Contracting.** We consider three different possible ways to organize the marketing of the crop. In the first, the grower sells directly to the wholesale market. This is the standard marketing arrangement assumed in farm-household models (Singh, Squire, and Strauss 1986) and in much of the literature on agricultural risk. In this case, the problem facing the grower is to choose quality so as to maximize expected utility.[2] Because the grower does his own marketing, compensation is simply equal to the price he receives, and so he solves

$$\max_q \int U(p) f(p|q) dp - z(q) \tag{1}$$

The solution to this problem involves choosing q so as to satisfy the first order condition of this problem,

$$\int U(p) \frac{f_q(p|q)}{f(p|q)} f(p|q) dp = z'(q).$$

The marginal cost of taking action q appears on the right hand side of this expression, and the marginal benefit to the grower appears on the left. From this expression we see that the marginal benefit is just the grower's expected utility, where the distribution of p has been

[2]For simplicity, we assume that the farmer's preferences are separable between quality and consumption. Relaxing this assumption wouldn't affect the general tenor of much of the discussion which follows, though it certainly is important to many of the standard analytical results in the contracting literature. See Gjesdal (1982) for an insightful discussion and example of the consequences of relaxing separability.

adjusted by the likelihood ratio $f_q(p|q)/f(p|q)$. Because $U(\cdot)$ is concave, Jensen's inequality implies that the grower is worse off than he would be if he simply received the expected price with certainty, *ceteris paribus.* The disutility associated with risk faced by the grower will generally distort the grower's choice of q (relative to the q chosen by a profit maximizing grower), though the direction of the distortion depends on the curvature of the integrand $U(p)f_q(p|q)/f(p|q)$ relative to the curvature of the marginal product $pf_q(p|q)/f(p|q)$.

2.2. **Contracting with Full Information.** Because the grower faces all the price risk in the problem above, one might suppose there to be scope for some risk-sharing intermediary. This intermediary could assume a number of guises; it might be a firm, a growers' cooperative, or a futures market. For concreteness, we'll refer to the intermediary as a "shipper." Thus, we consider a second way of organizing marketing, in which a shipper writes some enforceable contract with the grower prior to planting. In the simplest version of the model, the grower and shipper agree on some level of quality, and on some form of payment for the grower.

We imagine that in the absence of a contract with the grower, the shipper earns profits (presumably from contracts with other growers) of $\pi \in \Pi$. These profits are taken to be a random variable, and may or may not be independent of p (which, recall, is the price received for the produce of a *particular* grower). We write the joint conditional density of p and π as $h(p, \pi|q)$; as with the conditional density $f(\cdot|q)$, h is asssumed to be a continuously differentiable function of q, with $h(p, \pi|q) > 0$ for all $(p, \pi) \in P \times \Pi$. The shipper's total profits are denoted by $\pi^* = \pi + p - w$, where w is the compensation given to the grower. To avoid imposing any artificial structure on this compensation, we imagine that the shipper is free to specify a different payment to the grower for every possible realization of p and π; we denote this contingent payment by $w(p, \pi)$.

The shipper values profits (including profits earned from dealings with the grower) according to some increasing, weakly concave function $V(\cdot)$. Because the shipper must induce the grower to actually sign the contract, the expected utility of the grower if he signs the contract must be greater than or equal to the grower's expected utility if he *doesn't* sign the contract. We suppose this reservation utility to be some number $\underline{U}$, and express this constraint by

$$\int U(w(p,\pi))h(p,\pi|q)dp\,d\pi - z(q) \geq \underline{U}. \tag{2}$$

Thus, in designing the contract, the shipper decides what level of quality she wants, and how best to compensate the grower by solving the problem

$$\max_{q,\{w(p,\pi)\}} \int [V\,(\pi + p - w(p,\pi))]h(p,\pi|q)dp\,d\pi \tag{3}$$

subject to the individual rationality constraint (2). Note that the grower's compensation can vary across every possible state (p,π); we're imposing absolutely no *ad hoc* structure on the optimal compensation schedule. The solution to this problem is Pareto efficient, and computing expected profits over a range of possible values for $\underline{U}$ traces out the Pareto frontier of efficient allocations for this full information environment.

Working with the first order conditions from this problem, we see that for all pairs (p,π),

$$\frac{V'(\pi^*)}{U'(w(p,\pi))} = \theta, \tag{4}$$

where θ is the Lagrange multiplier associated with the constraint (2), and is thus a function of the grower's reservation utility $\underline{U}$. Though we write the grower's compensation w in its general form as a function of (p,π), the striking thing about equation (4) is that the shipper chooses to specify a contract in which the grower's compensation *doesn't* actually depend directly on the price p. This follows because (given $\underline{U}$) θ is a constant; thus the grower's compensation depends only on $\underline{U}$ and the realized profits of the shipper, π^*. The grower is fully insured against price risk, except insofar as variation in prices affects the profits of the intermediary. Because risk is efficiently allocated, outcomes in this environment are Pareto optimal, and production is also fully efficient.

In thinking about the risk growers face, there are two cases to consider. In the first case, the shipper has full information on quality, but is risk averse. In this case, the grower will share the risk associated with variation in the shipper's profits, but will face no additional risk associated with variation in the price of the grower's own produce—in effect, the shipper maximizes her utility by insuring the grower against price risk. In the second case, the shipper continues to have full information, but is risk neutral (which in this setting means that she maximizes her expected profits). In this case $V'(\cdot)$ is a constant, and so from (4) the grower's compensation must also be a constant; the grower faces no risk of any sort.

Although this model is stylized and extremely simple, the observation that the intermediary will bear all price risk is remarkably robust. Though to our knowledge this precise observation hasn't been previously made, it's clear from earlier work on insurance that this feature survives the addition of uncertainty in production, whether over quality or quantity (Wilson 1968); survives more elaborate, non-separable utility functions (Scheinkman 1984); and is preserved in a multi-period version of this model (Townsend 1982).

2.3. **Contracting with Private Information.** Because we observe contracts between growers and shippers which expose risk-averse growers to price risk, our model must so far must be missing something important. We need to introduce some sort of friction to keep the shipper from insuring the grower against all idiosyncratic price risk. A promising sort of friction is private information regarding quality. Suppose the grower chooses q, but the quality of the produce can't be observed by the shipper. If the grower bears some of the price risk, this will help to provide incentives to choose high quality.

In this section, we capture the possibility of unobservable quality differences by having the shipper *recommend* some level of quality q, and require that this choice be *incentive compatible*. That is, it must be in the grower's best interests to actually produce the recommended quality, or

$$q \in \operatorname*{argmax}_{\hat{q}} \int U(w(p,\pi))h(p,\pi|\hat{q})dp\, d\pi - z(\hat{q}). \tag{5}$$

Accordingly, the shipper solves the contracting problem by maximizing equation (3) subject to (2) and (5). Now, if the shipper offers the grower a constant compensation w, the grower will respond by choosing the lowest possible quality of produce; this is clearly inefficient. However, if the shipper is only able to observe price, the efficient contract will typically expose the grower to a great deal of price risk. In particular, so long as the production problem is suitably concave and the grower is sufficiently risk-averse (Jewitt 1988), any interior solution to the contracting problem will satisfy

$$\frac{V'(\pi^*)}{U'(w(p,\pi))} = \theta + \mu \frac{h_q(p,\pi|q)}{h(p,\pi|q)}, \tag{6}$$

where θ is the Lagrange multiplier associated with the participation constraint (2), and μ is the multiplier associated with the incentive compatibility constraint (5). Note that when the incentive compatibility constraint isn't binding, grower's compensation is constant by (6). If equation (5) is binding, compensation depends on the market price via the likelihood ratio $h_q(p,\pi|q)/h(p,\pi|q)$.

The intermediary may be able to increase expected profits by engaging in some sort of costly quality measurement, and using the results of this measurement to modify the payment made to the grower. Call this quality measurement some random variable $r \in R$, and suppose that r is governed by the joint density $\psi(p,\pi,r|q)$. The first order conditions for this problem would simply replace the likelihood ratio in (6) with the new likelihood ratio $\psi_q(p,\pi,r|q)/\psi(p,\pi,r|q)$. In the special case of independence between r and (p,π), we can write the density of r as some $g(r|q)$, so that the conditional joint pdf of p, π and r can be written as $\psi(p,\pi,r|q) = h(p,\pi|q)g(r|q)$. In this case equation (6) is replaced by

$$\frac{V'(\pi^*)}{U'(w(p,\pi,r))} = \theta + \mu \left(\frac{h_q(p,\pi|q)}{h(p,\pi|q)} + \frac{g_q(r|q)}{g(r|q)} \right). \tag{7}$$

Quality measurement is valuable so long as there exists a compensation rule $w(p,\pi,r)$ which makes at least one party strictly better off than the compensation rule $w(p,\pi)$. From (6), such a rule will exist so long as the quality measure is informative; that is, so long as $g_q(r|q) \neq 0$ for some q and some r [Holmström (1979), Proposition 3]. Now compare equations (4) and (7); by estimating the two likelihood ratios $h_q(p,\pi|q)/h(p,\pi|q)$ and $g_q(r|q)/g(r|q)$ one could calculate not only the efficiency loss due to imperfect quality measurement, but also the value of any quality measurement used for a given commodity (see the example of Section 3.2). When the variance of $g_q(r|q)/g(r|q))$ is large, realizations of r are likely to convey a great deal of information about the grower's choice of q. Thus, for example, if r is nearly a deterministic (monotonic) function of q (i.e., if knowing r permits one to know q quite precisely), then the shipper can offer a "forcing" contract which, for the recommended quality, provides the grower with a payment sufficiently high to satisfy (2), while imposing a severe punishment for any other very different quality.

3. Tomato Contracts

In Section 3.1 we first give a characterization of an actual contract widely used in the mature green fresh market tomato industry of California. Features of this contract include a minimum level of compensation for growers, and a family of increasing (in downstream price) compensation schedules, which are indexed by some measure of quality undertaken by the shipper (in the form of a cull). These actual contracts provide considerable insurance to growers, but at the same time expose the growers to different levels of price risk, depending on the measured quality of the growers' tomatoes.

Next, in Section 3.2 we construct a simple example of the sort of contracting problem with private information described in Section 2.3, meant to approximate the contracting problem facing fresh tomato packer/shippers (henceforth shippers). Note that while in some sense we 'know' the structure of these contracts (because we observe the actual contracts), we don't impose *any* structure directly on the contracts which emerge from the theoretical model. Instead, we consider a simple model in which agents have very simple, restrictive preferences and operate a very simple, restrictive technology, but in which quality can be measured only imperfectly. We then ask whether it's possible that the optimal set of contractual arrangements (drawn from the set of *all* possible arrangements) in this simple, restrictive environment exhibits the stylized features that we observe in a real-world industry. The answer turns out to be an only slightly qualified "yes."

3.1. **A Real Fresh-Market Contract.** We begin by describing a particular contract commonly observed in the "mature green" tomato industry, known as Joint Venture Agreements (JVA). We obtained such a contract from a large California shipper. The contract specifies a base payment B which depends only on information known to both parties at the time the contract is signed. Let y denote the quantity of tomatoes produced by the grower. After receiving the grower's crop, the shipper discards any tomatoes which are below a minimum size, are off-color, or exhibit obvious defects. The discarded tomatoes are called "culls." We denote the number of culled tomatoes by some function $\ell(r, y)$. The shipper then markets the remaining tomatoes $(y - \ell(r,y))$, and receives some price p, so that gross revenues for the shipper are $p(y - \ell(r,y))$. We call the grower's compensation under the JVA some $\hat{w}$, which is given by the schedule

$$\hat{w} = B + \frac{1}{2} \max \left\{ \left[(p+b)(y - \ell(r,y)) - x(y) \right] (1-\tau), 0 \right\}. \tag{8}$$

The function $x(y)$ is meant to cover the shipper's costs of picking, packing, hauling, and marketing. Note that $x(y)$ is deterministic, and specified explicitly in the contract signed before planting. The number b is a "bonus," specified in advance and paid on every carton of tomatoes shipped. The number τ is a "sales commission" assessed on net sales. The shipper's ability to choose B, b, $\ell(r,y)$, $x(y)$, and τ gives her considerable flexibility in specifying the contract,[3] and makes a putatively 'simple' contract actually rather complex.

It is very likely that the JVA has implicit provisions, as well as the explicit ones in the written contract. One example of an implicit provision which we know is present in the JVA has to do with the function $\ell(r,y)$. This function is not specified in the written contract, and yet we know from discussions with the shipper that in fact considerable resources are expended on the culling process, and that growers are aware of this. Other examples of implicit provisions include the civil liability of each party in the event of a problem with the tomatoes downstream (having to do with, e.g., the health consequences of pesticide residue, or bacterial contamination).

For simplicity, we assume that the quality measurement variable r can take on values of only "high" or "low," and that the cull function takes the simple form

$$\ell(r,y) = \begin{cases} \beta y & \text{if } r \text{ is low,} \\ 0 & \text{if } r \text{ is high.} \end{cases}$$

The interpretation of the cull function is that if measured quality is high, then the cull is zero; while if measured quality is low, then some proportion β of the total yield is discarded.To focus on the role of price risk, we hold yields fixed, and choose the normalization $y = 1$.

[3]In early discussions, shippers to whom we've spoken have tended to maintain that the JVA simply shares risks and returns between grower and shipper on a "50–50" basis. Examination of outcomes predicted by equation (8) reveal that this isn't even remotely true. Only the second term in equation (8)—itself a rather complicated function—is shared in this manner.

Then assuming that the grower's net revenues are non-negative, the compensation schedule (8) can be rewritten as some

(9) $$\hat{w} = \alpha_0 + \alpha_2 p - \begin{cases} \alpha_1\beta + \alpha_3\beta p & \text{if } r \text{ is low} \\ 0 & \text{if } r \text{ is high,} \end{cases}$$

where the parameters α_i, $i = 0, 1, 2, 3$ are each nonnegative. This way of writing the compensation schedule makes it apparent that quality measurement is used to select one of a family of schedules with which to compensate the grower, and that given our simplifying assumptions regarding $\ell(r, y)$, each of these schedules is piecewise linear (and two in number).[4]

The grower's compensation under the provisions of the observed JVA is shown in Figure 1, holding yields fixed. In the figure, the dashed lines are the grower's compensation under the actual JVA, for each of two different levels of cull. The solid line is simply a 45 degree line; when the realized price and compensation lie below the 45 degree line, the shipper makes a profit, while when realized price and compensation lie above the 45 degree line, the shipper bears a loss. In a competitive shipping industry, shippers try to woo growers by offering them high expected utilities $\underline{U}$; thus, in equilibrium, $\underline{U}$ must be such that expected price must equal expected compensation, and so the expected values of both price and compensation must lie on the line segment AB in Figure 1. In Figure 1, we've chosen $\beta = 0.43$ (the manner in which this value was chosen is described below), so that when the cull is high, the grower faces only about 27 per cent of the price risk.

Although Figure 1 pins down some important aspects of grower's compensation, the variety of implicit provisions which may be present in the contract means that the formula determining grower's compensation may be considerably more complicated (and flexible)

[4]To pin down the parameters in this schedule, we take advantage of the relationship between parameters in (8) and (9). Under the JVA, the sales commission τ is six per cent, and so $\alpha_2 = (1 - \tau)/2 = 0.47$, indicating that the grower faces 47 per cent of the price risk when measured quality is high. The bonus b offered by the shipper in the actual contract is roughly equal to one seventh of the average price per carton of tomatoes in recent years, and so for illustrative purposes the bonus is set so that b is equal to one seventh of the expected price, so that $\alpha_1 = \alpha_2 \mathrm{E}p/7$ (note that because the distribution of prices depends on q, $\mathrm{E}p$ is endogenous, and so is determined only after we solve the theoretical model we present below). The parameter α_3 is equal to $-\alpha_2$. We can't assign a number in advance to the share of price risk faced by the grower when measured quality is low, because this share is is $\alpha_2(1 - \beta)$, and we don't have direct information on the value of β (we regard α_0 and β as free parameters in what follows, and effectively 'estimate' these values by matching our theoretical contract to the actual JVA contract).

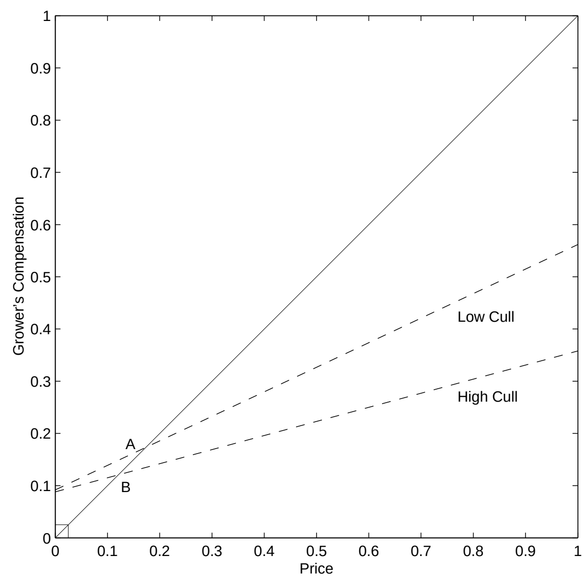

FIGURE 1. Grower's Actual Compensation Under the JVA

than the written contract suggests. However, with the exception of the cull, conversations with the shipper suggested that in most years there were no important implicit provisions which affected grower compensation. Nonetheless, it seems correct to regard the compensation schedule (8) and Figure 1 as no more than an approximation to the 'true' compensation schedule.

3.2. **A Theoretical Fresh-Market Contract.** Section 3.1 characterized an actual fresh-market tomato contract called a JVA. In this section, we specify precise preferences and distributions (up to a vector of parameters) so that we can compute the optimal contracts from the theoretical model of imperfect quality measurement which we presented in Section 2.3. We'd like for our model to be able to reproduce some of the stylized features of the JVA discussed above.

We assume that the tomato shipper is risk neutral, and that the utility of the grower is given by $U(w) - z(q) = \log(w) - \alpha q$. Note that here the distribution of π is immaterial to the form of the final contract because of the assumed risk-neutrality of the intermediary.[5] The support of the price distribution is $P = [0, \bar{p}]$; the support of the quality measure distribution

[5]The optimal contract is the solution to equations (2), (7), and (5). When the intermediary is risk neutral, $V'(\pi^*)$ is a constant, so that π appears nowhere in these equations.

is $R = [0, \bar{r}]$. The conditional distribution of price is permitted to be dependent on measured quality, with the joint cdf of price and quality measure given by

$$\Psi(p, r|q) = \left(\frac{\kappa_p/q}{\kappa_p/q - p + \bar{p}}\right)^{\gamma_p + \delta_p r} \left(\frac{\kappa_r/q}{\kappa_r/q - r + \bar{r}}\right)^{\gamma_r}. \tag{10}$$

With a change of variable, it's easy to see that this distribution function is a simple generalization of a multivariate logistic distribution.[6]

The parameter δ_p captures possible conditional dependence of p on r, which seems to be important. Empirically, this importance is reflected in the actual JVA contract by the interaction between p and r (reflected by the term $\alpha_3\beta p$ in (9)) which helps to determine growers' compensation. Theoretically, while the shipper's customer can engage in their own quality measurement in order to assess quality (and determine price), additional information regarding quality can be had from the shipper's measurement (perhaps reflected in grade), which suffices to induce dependence.

With the primitives of preferences and technology in hand, we again turn our attention to the problem defined by the binding individual rationality constraint (2), the first order conditions (7), and the first order conditions associated with the grower's problem (5). Because the first order approach is valid, any interior solution to these equations is unique (Rogerson 1985).[7] The solution to this problem depends on seven parameters of the theoretical contract: $(\alpha, \kappa_p, \kappa_r, \gamma_p, \gamma_r, \delta_p, \underline{U})$.[8] We assume that intermediation in the tomato industry is competitive,[9] which determines one of these parameters. In particular, the requirement that

[6]This distribution has a number of attractive properties for our purposes. It has compact support, which makes numerical work relatively simple. Because the likelihood ratio $\psi_q(p, r|q)/\psi(p, r|q)$ is a nondecreasing function of (p, r) for all $q > 0$, and because Ψ is convex for all q, the so-called "first order approach" is valid (Sinclair-Desgagné 1994); as a consequence, the shipper can never gain by asking the grower to randomize quality. Aside from guaranteeing that the first order approach is valid, these conditions have a nice interpretation in this production context—the monotone likelihood ratio implies that expected returns (prices) are an increasing function of quality, while the convexity of Ψ implies a sort of stochastic diminishing marginal product of q (more precisely, expectations of concave functions of p are themselves concave functions of q).

[7]Rather than solving this problem directly, we discretize the environment and translate the problem into a linear program (in probabilities), as suggested by Myerson (1982) or Prescott and Townsend (1984).

[8]We fix $\bar{p}$ and $\bar{r}$ at one. This is not completely innocuous (i.e., not simply a choice of *numeraire*), because the quantities $\bar{p}$ and $\bar{r}$ help to determine the pdf of prices and measured quality. The results we report seem to be fairly insensitive to this choice, however.

[9]This seems to be a reasonable assumption for fresh-market tomatoes; barriers to the entry of shippers are low, and there are many shippers in the state of California. Intermediation for other commodities may not

Theoretical Contract		
Parameter	Value	Interpretation
α	0.6588	Disutility of providing quality.
κ_p	0.2493	Governs marginal influence of quality on price.
κ_r	0.4231	Governs marginal informativeness of quality measure.
γ_p	0.1432	Governs rate at which marginal influence of quality diminishes.
γ_r	0.1281	Governs rate at which marginal informativeness of quality measure diminishes.
δ_p	1.1185	Governs interaction between measured quality and price.
Actual JVA Contract		
Parameter	Value	Interpretation
α_0	0.0867	Net base payment awarded growers.
β	0.4259	Proportion of tomatoes culled when measured quality is low.

TABLE 1. Estimated contractual parameters

the intermediary's expected profits are zero means that the grower receives all the expected surplus from the contract; if this weren't so, a shipper who received some positive expected surplus (profit) by offering a prospective grower $\underline{U}$ would have to contend with entrants willing to offer just slightly more.

After pinning down $\underline{U}$, we're left with six free parameters in the theoretical contract, and two in the JVA. For the JVA, these include the base payment net of the grower's share of input costs (α_0), and the cull parameter β.

The values of the parameters from the theoretical contract and the JVA parameters are chosen to minimize the distance between compensation under JVA and compensation under the theoretical contract. Our distance metric is given by the mean square error between the actual contract given a low cull and our theoretical contract given a low cull, plus the mean square error between the actual and theoretical contracts given a high cull, or

$$\int_0^{\bar{p}} (\hat{w}(p, \text{low}) - w(p, \text{low}))^2 dp + \int_0^{\bar{p}} (\hat{w}(p, \text{high}) - w(p, \text{high}))^2 dp.$$

The resulting parameter vector for the theoretical contract is given in Table 1.

Along with the actual JVA, Figure 2 displays the grower's compensation (solid lines) under the theoretical contract as a function of r and p, given the equilibrium choice of q. The grower's expected compensation given prices (but before the cull is determined by r) is

be competitive; for example, Durham and Sexton (1992) concludes that oligopolistic rents probably accrue to California's tomato processors.

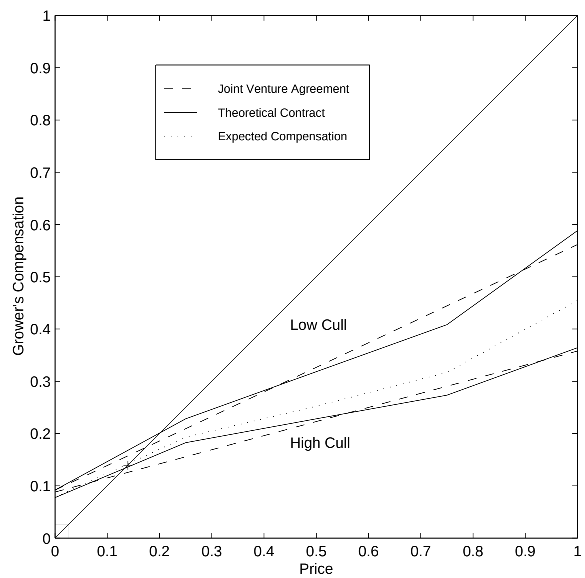

FIGURE 2. Grower's Actual & Theoretical Compensation

given by the dotted line. The point at which the 45 degree line intersects the dotted line is the expected price/compensation pair before the cull is known, $(0.1262, 0.1262)$. Note that expected price must be equal to expected compensation because expected profits are zero in competitive equilibrium.

Now compare the grower's compensation schedule under the JVA (the dashed lines in Figure 2) with that predicted by the theoretical model (the solid lines). We regard this comparison as something of a success for the model, since the model is plainly capable of reproducing the stylized features of actual contracts which we consider important—the grower receives a base payment, a considerable amount of insurance, yet receives a lower compensation when measured quality is low, and faces a considerable degree of price risk. Beyond matching these qualitative features, we'd be inclined to regard the theoretical schedule as a reasonably good 'match' to the actual schedule.[10] The fact that we're able to come close to

[10]On the other hand, it's somewhat hard to know what 'close' means in this setting—we don't have data on actual outcomes, and we don't have any way to define something like a sampling distribution for our measure of distance, the way we would if this were a statistical problem. Because our model is unable to *exactly* replicate the explicit JVA contract, a sterner interpretation of our results is that our model should be rejected—something in it, maybe preferences, maybe the distributions, maybe the basic structure—is fundamentally mis-specified. Our defense against this charge is basically that no one else has yet proposed a simple, well-specified behavioral model which can capture even the broad qualitative features of the observed contract, and so that our model performs very well when compared with any proposed alternative.

matching the actual contract is important, as it is by no means apparent that the solution to equations (2), (5), and (7) would look anything at all like the JVA contract in (8). While some authors (e.g., Hart and Holmstöm (1987)) have argued that the solution to agency problems like the one studied in this paper result in contract structures that are far more complicated than most real-world contracts, here we show that in fact the solution may be no more complicated than the real-world contract.

Returning to our interpretation of the contract, we can see that the shipper has two ways of mitigating the moral hazard problem associated with the grower's choice of q. The use of each of these tools is evident under both the JVA and the theoretical contract. First, the grower receives a higher compensation when the realized price is high; in this example, both the actual and the theoretical contracts expose the grower to about 47 per cent of the price risk when the cull is low. This level of price risk induces the grower to provide a higher level of quality by letting him share the benefits of this higher quality, in the form of higher expected prices. Second, the shipper can attempt to measure quality directly, and offers higher compensation when measured quality is high (cull is low). These incentives induce the grower to produce tomatoes of high quality, despite the fact that the shipper can't observe quality directly.

It's interesting to compare this contract with the set of arrangements that would prevail if quality were directly observable, or if there were no good measures of quality. Our our estimated theoretical model allows us to evaluate these counterfactual contracts. Figure 3 shows the full information contract (where there is no need for quality measurement), and the hidden-quality contracts with and without quality measurement. The underlying space is the quality q chosen by the grower, along with a certainty-equivalent compensation. Because the grower likes compensation and dislikes supplying quality, he prefers points in the northwest of this graph. The (slightly) convex, upward sloping lines in the figure are grower's indifference curves.The concave function in the figure is a production frontier, which gives the expected price for any given q.

The requirement of competitive intermediation implies that expected compensation for the growers should lie on the concave production frontier, whatever the contract. The certainty equivalent compensation corresponding to expected compensation will always lie somewhat

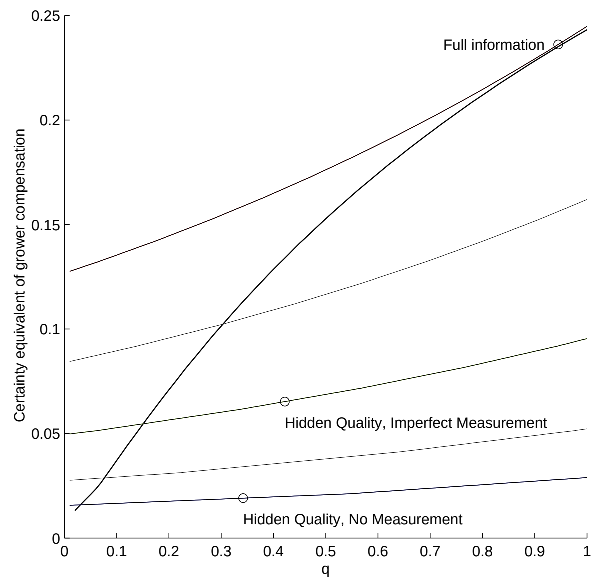

FIGURE 3. Tomato Contracts and Grower's Preferences

below this frontier for risk-averse growers. The vertical distance between the certainty-equivalent and the production frontier in can be interpreted as the largest insurance premium the grower would be willing to pay if insurance against price risk were available.

Figure 3 provides information about the value of quality measurement in the shipper's contract and about the efficiency loss from hidden quality. We assume the parameters reported in Table 1; that is, those parameters estimated under the null hypothesis of hidden quality with imperfect quality measurement. When there is full information, as in the model of Section 2.2, the grower receives a certain compensation of 0.23, and provides a level of quality of 0.95. When quality is hidden, but the shipper uses a cull or other measurement to imperfectly measure quality, then certainty-equivalent compensation falls to 0.065, while quality falls to 0.42. Finally, when there is no quality measurement at all, so that producers' compensation depends only on realized price, certainty-equivalent compensation falls to 0.019, and quality falls to 0.34.

The loss of producer surplus associated with these different information problems can be measured in terms of the size of the transfer which would be necessary to return growers to the indifference curve associated with the full information outcome. A grower in a world of hidden quality with imperfect quality measurement would require a transfer of 0.096, or 146 per cent of his certainty-equivalent income. A grower in a world of hidden quality with no measurement would require a transfer of 0.133, or hugely more than the certainty-equivalent compensation he receives.

One way to understand why imperfect quality measurement leads to an efficiency loss, relative to the full-information environment, is to refer back to equation (7). There, the information associated with quality measurement is determined by the distribution of the likelihood ratio $g_q(r|q)/g(r|q)$, and a mean-preserving spread in this distribution leads to more efficient outcomes (Kim 1995). In this example, an increase in the parameter γ_r leads to a mean preserving spread in the distribution of the likelihood ratio $\psi_q(p, \pi, r|q)/\psi(p, \pi, r|q)$, and this increases efficiency by making it easier for the shipper to infer quality through measurement.

There are several examples of the sorts of changes in quality measurement technology which we might interpret as increases in γ_r. For example, tomato shippers might adapt the very sophisticated quality measurement technologies used by many stone-fruit shippers, some of whom use computers and imaging technology to measure the color, size, and weight of every piece of fruit they ship. The corresponding improvement in the shipper's ability to sell differentiated lots of produce also gives much more detailed information about the quality of the produce from a given grower. This additional information can be used to improve the incentives for the grower to provide higher quality fruit, leading to an increase in the efficiency of the contract.

4. Conclusion

Many farmers who market their produce under contract with an intermediary face price risk in the sense that their compensation depends on the price paid in some downstream market.

In Section 2, we've developed three different models of contracting in the produce industry that help us understand the variation in price risk that different growers face. First, we consider a grower who does all of his own marketing, and thus faces 100 per cent of price risk. Second, we consider a grower who signs a contract with an intermediary who can observe the costly quality decisions made by the grower. In this full-information model, growers never face idiosyncratic price risk, but if intermediaries are risk averse, a grower's compensation may depend on the profits of the intermediary, and will not depend directly on the idiosyncratic price fetched by the grower's own produce. Third, we consider the contractual arrangements one might observe if quality isn't directly observable but influences price. For this model, we show that under quite general conditions growers' compensation will depend on price. Thus, each model produces a different prediction about the amount of price risk growers should face.

In Section 3, we calculate the non-yield related risks faced by fresh-market tomato growers by analyzing an actual contract called a "Joint Venture Agreement." Under this contract, the grower faces at most 47 per cent of the price risk, suggesting that the fresh-market tomato industry is best characterized using our model of unobservable quality. We construct a simple example using a particular family of conditional distributions related to the logistic distribution. The example features competitive risk-neutral intermediaries and growers with logarithmic preferences over consumption and linear preferences over quality. We then choose a preference parameter and several parameters of the price/quality measure distribution and specify an environment in which the optimal contract is (nearly) the same as the actual JVA contract.

Our analysis ignores three additional issues that require further study. First, in many real-world contracts between growers and shippers, market intermediaries often control some of their growers' inputs. Goodhue (1998) considers this issue in the case of broiler contracts. Second, though it's possible to interpret price variation in our model as stemming from both idiosyncratic and aggregate sources, in this paper we've chosen to focus on idiosyncratic price risk. A careful consideration of aggregate sources of risk requires thinking about numerous growers and intermediaries, and is a focus of our ongoing research. Third, fresh-market intermediaries may be able to influence the price a grower receives by mishandling his produce or by not working hard to find a high price. This fact, and that a typical shipper sells the produce of many growers, may help explain why relative performance evaluation (a common incentive instrument in livestock and processing contracts) is not used in the fresh produce industry. Finally, we have treated the quality-measurement technology and the intermediary's ability to identify a particular grower's produce as though these were exogenous. In practice, intermediaries may have some control over which of these instruments are used. For example, a processor with liability concerns might choose to arrange her production facility to allow tracking of processed product back to a particular grower. Once such a system is in place, some degree of price risk can be assigned to the grower, and the need for quality measurement may be reduced.

References

Durham, C. A. and R. J. Sexton (1992). Oligopsony potential in agriculture: Residual supply estimation in California's processing tomato market. *American Journal of Agricultural Economics 74*(4), 962–72.

Gjesdal, F. (1982). Information and incentives: The agency information problem. *Review of Economic Studies 49*, 373–390.

Goodhue, R. E. (1998). Broiler production contracts as a multi-agent problem: Common risk, incentives and heterogeneity. Manuscript.

Hart, O. and B. Holmstöm (1987). The theory of contracts. In T. F. Bewley (Ed.), *Advances in Economic Theory: Fifth World Congress*. Cambridge: Cambridge University Press.

Holmström, B. (1979). Moral hazard and observability. *Bell Journal of Economics 10*, 74–91.

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

Just, R. E. (1974). An investigation of the importance of risk in farmers' decisions. *American Journal of Agricultural Economics 56*, 14–25.

Just, R. E. and D. Zilberman (1985). Risk aversion, technology choice, and equity effects of agricultural policy. *American Journal of Agricultural Economics 67*(2), 435–440.

Kim, S. K. (1995). Efficiency of an information system in an agency model. *Econometrica 63*(1), 89–102.

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

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

Rogerson, W. P. (1985). Repeated moral hazard. *Econometrica 53*(1), 69–76.

Scheinkman, J. A. (1984). General equilibrium models of economic fluctuations: A survey of theory. Manuscript.

Sinclair-Desgagné, B. (1994, March). The first-order approach to multi-signal principal-agent problems. *Econometrica 62*(2), 459–466.

Singh, I., L. Squire, and J. Strauss (1986). *Agricultural Household Models: Extensions, Applications and Policy*. Baltimore: John Hopkins University Press.

Stiglitz, J. E. (1974). Incentives and risk sharing in sharecropping. *Review of Economic Studies 41*, 219–255.

Townsend, R. M. (1982). Optimal multiperiod contracts and the gain from enduring relationships under private information. *Journal of Political Economy 90*, 1166–1185.

Wilson, R. (1968). The theory of syndicates. *Econometrica 36*, 119–132.