AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

378.794
G43455
WP-840

Working Paper Series

WORKING PAPER NO. 840

PRODUCER PRICE RISK AND QUALITY MEASUREMENT

by

Brent Hueth and Ethan Ligon

Waite Library
Dept. of Applied Economics
University of Minnesota
1994 Buford Ave - 232 Clapp
St. Paul MN 55108-6040 USA

DEPARTMENT OF AGRICULTURAL AND
RESOURCE ECONOMICS

BERKELEY

California Agricultural Experiment Station

University of California

378.794
643455
WP-840

**DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS AND POLICY
DIVISION OF AGRICULTURE AND NATURAL RESOURCES
UNIVERSITY OF CALIFORNIA AT BERKELEY**

WORKING PAPER NO. 840

PRODUCER PRICE RISK AND QUALITY MEASUREMENT

by

Brent Hueth and Ethan Ligon

Assistant Post Graduate Researcher and Assistant Professor of Agricultural and Resource Economics and Policy, respectively.

**California Agricultural Experiment Station
Giannini Foundation of Agricultural Economics
May, 1998**

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 provides additional information regarding quality. As a consequence, an efficient contract does not shield growers from all price risk. We examine this hypothesis for the case of fresh-market and processing tomatoes, and conclude that unobserved quality is capable of explaining observed variation in the price risk tomato growers face.

Date: April 27, 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.

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 seldom shield the producer from all risk. This is somewhat surprising. In a competitive market for intermediation, a risk neutral intermediary ought to bear all risk, as it would be costless for it to do so.

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. Then 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 observed 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 observe it. If the intermediary is unable to observe quality directly, he may seek to infer it; objectively by measuring a variety of attributes of the produce, and perhaps also by more subjective means. However, if this inference is less than perfect, then the grower may well be exposed to price risk.

To see why poorly measured quality may expose the grower to price risk, consider the case of wine grapes. For premium wines, the vintner may care a great deal about the quality of the grapes he purchases, and the quality of the grapes is very likely to effect the price eventually paid by the consumer. The grower may well have a fair amount of control over some of these qualities—by modifying his cultivation practices, he may be able to impart various flavors which eventually show up in the wine. However, some of these quality attributes may be extremely difficult to measure at the farm gate, and may only become apparent after the wine is mature, years after the grape has been cultivated. Because the vintner may be uncertain as to whether or not the grapes he has purchased do in fact have some of the qualities advertised, it may be optimal to ask the grower to pay some of the costs of poor expression of these qualities by making his payment contingent on the final bottle price of the wine.

Is this exposure to price risk a bad thing? In the example we've just given, it appears to be a possibly optimal response to an information problem inherent in the uncertainties of vinification. Nonetheless, it's an inefficient outcome relative to what one would expect if all the quality aspects of a grape were observable at the farm gate—a risk-averse grower is always willing to take some cut in the price of his produce in exchange for a reduction in uncertainty.

Because we observe contracts which seem to expose the grower to idiosyncratic price risk, we hypothesize that quality is not perfectly observable. In this case, 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 price eventually paid by the consumer. When the grower's compensation depends on price, this is evidence that quality measurement undertaken by the intermediary is imperfect.

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 our model with actual contracts from the fresh and processing tomato industries 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, then 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.

In order to abstract from yield risk, we assume that a grower produces a single unit of some agricultural commodity. The grower controls the quality, q , of the commodity, but increased quality comes only at a cost $c(q)$ to the grower, where q is some real number. The cost function $c(q)$ 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 on the wholesale market; on this market the prevailing price of the commodity is 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 best way to think of p is as some hedonic price of the characteristics q , along with some error. 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.¹ Because the grower does his own marketing, compensation is simply equal to the price he receives, and so he solves

$$(1) \quad \max_q \int U(p) f(p|q) dp - c(q)$$

The solution to this problem involves choosing q so as to satisfy

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

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 $U(p)f_q(p|q)/f(p|q)$.

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 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 assumed to be a continuously differentiable function of q , with $h(p, \pi|q) > 0$ for all $(p, \pi) \in P \times \Pi$. 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

¹For 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.

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

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

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

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

subject to the individual rationality constraint (2).² The solution to this problem is first best, 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) ,

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

where θ is the Lagrange multiplier associated with the constraint (2).

The striking thing about equation (4) is that because θ is a constant, the shipper chooses to specify a contract in which the grower's compensation doesn't actually depend directly on the price p ; the grower's compensation depends only on $\underline{U}$ and the realized profits of the shipper.

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 may be risk averse. In this case, the grower will share the risk associated

²This leaves open the question of what it is that determines the grower's reservation utility $\underline{U}$. If the grower has the option of doing his own marketing, then certainly $\underline{U}$ must be at least as great as the expected utility the grower receives from solving the problem in Section 2.1. Another, more interesting way to determine $\underline{U}$ is to make some assumptions regarding the organization of the shipping industry. In particular, if shippers maximize expected profits in a competitive industry, then these expected profits must be zero. Because grower's preferences are monotone and concave, this is enough to determine a unique $\underline{U}$.

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. If we think about obvious directions in which one might wish to extend this model, we see that this feature survives the addition of uncertainty in production, whether over quality or quantity; survives more elaborate, non-separable utility functions; and is preserved in a multi-period version of this model.

2.3. Contracting with Private Information. Because we observe contracts between growers and shippers which expose risk-averse growers to price risk, our model 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 choice. 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 our third model, we capture the possibility of unobservable quality differences by having the shipper *recommend* some level of quality, 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

$$(5) \quad q \in \operatorname{argmax}_q \int U(w(p, \pi)) h(p, \pi | q) dp d\pi - c(q).$$

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 able to only 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

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

where λ 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, then we recover the constant compensation for growers from (4); when (5) is binding, then 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 pdf $\psi(p, \pi, r|q)$. 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 pdf 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

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

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]. 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 unobservable quality, but also the value of any quality measurement used for a given commodity.

3. AN EXAMPLE

In this section we construct a simple example of the sort of contracting problem with private information described above, meant to approximate the contracting problem facing fresh tomato packer/shippers (henceforth shippers). The optimal contract which solves our stylized problem turns out to resemble the sorts of contracts actually used by fresh tomato shippers. A small modification to the example approximates the problem facing a canner who contracts for processing tomatoes; this optimal contract also reproduces several interesting features of actual contracts.

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 an example of 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 $p(y - \ell(r, y))$. The grower's compensation is then given by

$$(8) \quad B + \frac{1}{2} \max \{ [(p + b)(y - \ell(r, y)) - x(y)] (1 - \tau), 0 \}.$$

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.

It is very likely that the JVA has implicit provisions, as well as 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).

To model the JVA, we assume that the quality measurement variable r can take on values of only zero or one, and that the cull function takes the simple form

$$\ell(r, y) = \begin{cases} \beta y & \text{if } r = 0, \\ 0 & \text{if } r = 1. \end{cases}$$

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. With these simplifications, and the assumption that the grower's net revenues are non-negative, the grower's compensation under the JVA can be rewritten as some

$$(9) \quad \alpha_0 + \alpha_2 p - \begin{cases} \alpha_1 \beta + \alpha_3 \beta p & \text{if } r = 0 \\ 0 & \text{if } r = 1, \end{cases}$$

where the parameters α_i , $i = 0, 1, 2, 3$ are each nonnegative. The bonus b offered by the shipper 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 (where the equilibrium expected price is determined from the model below). 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 competitive equilibrium, 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.

Under the JVA, the sales commission τ is six per cent, and so the slope of the upper line is $0.47 [(1 - \tau)/2]$, indicating that the grower faces 47 per cent of the price risk when measured quality is high. The slope of the lower line is equal to $0.47(1 - \beta)$. In Figure 1, $\beta = 0.44$ (the manner in which this value was chosen is described below), so that when the cull is high, the grower faces only about 26 per cent of the price risk.

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) 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.

3.2. A Theoretical Fresh-Market Contract. We'd like for our model to be able to reproduce some of the stylized features of the JVA discussed above. We construct an example assuming that the tomato shipper is risk neutral, and that the utility of the grower is given by

$$U(w) - c(q) = \log(w) - \alpha q.$$

FIGURE 1. Grower's Actual Compensation Under the JVA

Note that here the distribution of π is immaterial, because of the assumed risk-neutrality of the intermediary. The support of the price distribution is $P = [0, \bar{p}]$; the support of the quality measure distribution 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

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

With a change of variable, it's easy to see that this distribution function is a simple generalization of a multivariate logistic distribution. For our purposes, this distribution has a number of attractive properties. 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 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 returns to scale.

The conditional dependence of p on r is permitted for two reasons. First, it seems likely that consumers' perceptions of value may be shaped in part by the grade and marketing of produce. Second, given the logarithmic utility assumed for growers, some such dependence is necessary to capture the interaction between p and r (reflected by the term $\alpha_3\beta p$ in (9)) which helps to determine growers' compensation.

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.³ The solution to this problem depends on seven parameters of the theoretical contract: $(\alpha, \kappa_p, \kappa_r, \gamma_p, \gamma_r, \delta_p, \underline{U})$.⁴ We assume that intermediation in the tomato industry is competitive, which determines one of these parameters. In particular, the requirement that the intermediary's expected profits are zero means that the grower's reservation utility $\underline{U}$ is endogenous. This leaves 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 (the intercept in Figure 1), and the cull parameter β . We choose these eight parameters so as to make the two contracts approximate each other.⁵

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 (before the cull is determined by r) is given by the dotted line. The point at which the 45 degree line intersects the dotted

³Rather 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).

⁴We 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.

⁵More specifically, the values of these parameters and the JVA parameters are chosen to minimize the distance between compensation under JVA and compensation under the theoretical contract, where our distance metric is the sum of squared deviations between the two schedules at a fixed set of prices. The resulting parameter vector for the theoretical contract, ordered as above, is (0.6965, 0.2300, 0.3982, 0.1502, 0.1225, 1.1112).

FIGURE 2. Grower's Actual & Theoretical Compensation

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.

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 which would prevail if quality were directly observable. In this first-best world, the grower chooses a much larger q ; 0.84 rather than 0.42. The requirement of competitive equilibrium among intermediaries means that expected compensation must continue to be equal to expected price. When quality is observable the risk-averse grower faces no uncertainty. Accordingly, the grower receives a compensation of 0.21 with certainty. The welfare improvement in this case is noticeable; -2.15 utils instead of -3.09. The welfare loss associated with private information can also be measured in terms of certainty-equivalent grower incomes. A grower under the hidden quality contract would experience a welfare improvement equivalent to a 113 per cent increase in certain income if he could move to the competitive full information contract described above.

3.3. Real Processing Contracts. The fresh-market contracts described above rely on the intermediary's ability to observe the price at which a given grower's produce is sold. This is reasonable in the case of fresh market produce, where the grower can often be identified even by the eventual consumer,⁶ but is probably difficult for many processors. A manufacturer of tomato paste, for example, is apt to use the produce of many growers in a given lot of paste. Accordingly, the price at which the paste is sold reveals little about the quality chosen by the original grower.

How do actual processing contracts compare with the fresh market contract described above? The California Tomato Growers Association (CTGA) annually provides a summary of the contracts and terms offered by the principal processors in the state (California Tomato Growers Association 1996). In 1996, this summary included information on the contracts offered by 26 different processors. Nearly all of these contracts take a standard form, which can be written as

$$\max\{(a + m(r))(y - d(r)), 0\}.$$

The variable y is again the grower's gross tonnage, and a is a base price, negotiated in advance. The two functions $m(r)$ and $d(r)$ are "premium" functions and "deduction" functions, respectively. Premia are awarded for certain measured quality characteristics, such as a high percentage of soluble solids, and are measured in dollars per "net ton." Deductions also depend on measured quality characteristics. The most

⁶Cartons of tomatoes or other produce are usually marked by the shipper with a bar code which identifies the grower, and often the particular field or block from whence the produce originated.

common characteristics which result in deductions are "material other than tomatoes," poor color (high Agtron readings), mold, or worms. Tomatoes afflicted with any of these flaws are separated from 'good' tomatoes and weighed; some multiple of this weight is subtracted from y to arrive at the net tonnage. Precisely what multiple is used typically depends on the weight of each of the different flaws, so that $d(r)$ is very often a quite complicated, nonlinear function.

3.4. Theoretical Processing Contracts. To model the market for processing tomatoes, we add to our earlier model the requirement that that a grower's compensation be independent of realized price. The realized price still depends on the quality choices of many growers, but these growers are assumed to be 'small' enough that the price reveals nothing regarding an *individual* grower's quality choice. With this additional restriction, a grower's compensation does not, of course, depend on the realized price (though the processor's profits do). Nonetheless, the processor would like to induce the grower to choose high quality. Because the processor cannot observe quality, and because price reveals nothing about an individual grower's choice of quality, the processor must rely only on quality measurement or grading to provide incentives.

We want to compare optimal contracts in the fresh and processed tomato industry. To make comparisons simple, we compute the optimal processing contract assuming the same preferences and technologies estimated from the fresh market example above. As in our earlier example, we assume that processing is competitive, so that the processor's expected profits are zero (for some empirical evidence supporting this assumption, see Durham and Sexton (1992)). For the same parameters given above, the processor's reliance on quality measurement is quite extreme. The quality measure r can take on only one of two values, zero or one. The simulated processor contract, then, gives the grower compensation of 0.467 if $r = 1$, and 0.053 if $r = 0$.⁷

Figure 3 shows the processing contract, along with the full information and joint venture contracts. The underlying space is the quality q

⁷We do some injustice to the processing contract by supposing that it relies only on the same observable quality measures as do fresh market tomatoes. In reality, a processor typically undertakes to measure quality much more formally than do fresh-market producers. On the other hand, fresh-market shippers *could* choose to take the same measurements that processors do, but they do not; presumably the combination of a culling line and observations of price more than make up for less intensive formal quality measurements.

FIGURE 3. Tomato Contracts and Grower's Preferences

chosen by the grower, along with a certainty-equivalent compensation.⁸ Because the grower likes compensation and dislikes supplying quality, he prefers points to the northwest of this graph. The (slightly) convex, upward sloping lines in the figure are grower's indifference curves, each curve corresponding to the level of utility associated with one of the contracts discussed above. The concave function in the figure is a production frontier, which gives the expected price for any given q .

The requirements of competitive equilibrium imply 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 below this frontier for risk-averse growers. The distance between the certainty-equivalent and the production frontier in Figure 3 can be interpreted as the largest insurance premium the grower would be willing to pay if insurance against price risk were available.

⁸For this example, the grower's certainty equivalent is calculated as the number ξ such that $\ln(\xi) = \int \ln(w(p, r))\psi(p, r|q) = \zeta(q)$. Thus, the slopes of the indifference curves in Figure 3 are given by $\xi\zeta'(q)$.

The fact that the processor cannot condition compensation on realized price affects more than the compensation scheme itself. Faced with these weaker incentives, the grower choose a level of quality equal to only 32 per cent of quality at the full information optimum, and receives a certainty equivalent compensation of only 0.06.

The conventional wisdom is that produce for processing applications is often of lower quality than produce sold on the fresh market. There may be good demand-side explanations for this phenomenon; certainly there's no reason to require processing tomatoes to be visually appealing, and the size of tomato may matter less than it would for the fresh market. However, the analysis here offers an alternative, supply-side explanation for lower quality in processing—because growers' compensation can't depend on the downstream price of the tomatoes, growers' incentives to provide high quality produce are weaker than those for fresh-market growers. As a consequence, the average price for processing tomatoes is lower than the price for fresh-market tomatoes; 0.08 versus 0.13 in the simulation here.

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. Other farmers face no price risk, but their compensation depends on quality measurement.

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. Compensation 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 may be 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.

The processed tomato market differs from the fresh-market tomato market in two interesting ways. First, because processors typically blend tomatoes from many different growers to make their product, these intermediaries cannot use price to draw inferences regarding the quality chosen by growers. Second, grower compensation in the processed tomato market doesn't depend on price. We think that these two facts are related. We add to our computed example the feature that the intermediary learns nothing about quality by observing price. The optimal contract in this new environment specifies that grower compensation should depend only on measured quality, and not on the realization of prices, thus replicating this feature of processing tomato contracts.

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 (1997) considers this issue in the case of broiler contracts. Second, in real markets there are important sources of aggregate price risk which an extended model could consider. Third, 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.

Waite Library
Dept. of Applied Economics
University of Minnesota
1994 Buford Ave - 232 ClaOff
St. Paul MN 55108-6040 USA