Traceability, Moral Hazard, and Food Safety

Starbird, S. A.[1], Amanor-Boadu, V.[2] and Roberts, T.[3]

[1] Santa Clara University, Santa Clara, California, USA
[2] Kansas State University, Manhattan, Kansas, USA
[3]Economic Research Service, USDA, Washington, DC, USA

Abstract— **Errors in traceability can significantly impact the moral hazard associated with producing safe food. The effect of moral hazard depends on the proportion of unsafe food costs that can be allocated to the responsible producer, which depends on the efficiency of the traceability system. In this paper, we develop a model that identifies the minimum level of traceability needed to mitigate moral hazard and motivate suppliers to produce safe food. Regulators and consumer can use the results of this research to design regulations and contracts that mitigate moral hazard and motivate producers to deliver safe food.**

Keywords— **Food safety, traceability, moral hazard.**

I. INTRODUCTION

Moral hazard is the effect of insulation from risk on the behaviour of individuals and organizations. In food supply chains, producers are often insulated from the economic risks of producing unsafe food. These risks include the social and personal cost of food-borne illnesses and the loss of trust in the food system. When producers are insulated from the risk of unsafe food, they have less of an incentive to produce safe food. Exposing producers to the risk of unsafe food mitigates the moral hazard, motivates suppliers to work harder to produce safe food, and increases the safety of the food supply.

One of the essential prerequisites to mitigating moral hazard is an efficient traceability system. A traceability system collects and maintains information about the origin of a food product. An efficient traceability system is one in which the source of unsafe food can be quickly and accurately identified. Inefficient traceability systems exhibit significant error in the identification of the supplier responsible for unsafe food ("traceback error"). Traceback error is common in the food supply chain because food is often comingled during production and consumption, and because illnesses often occur several days after the victim consumes the food.

In this paper, we develop a model that incorporates sampling error, diagnostic error, and traceback error into a measure of the producer's utility from producing safe and unsafe food. We use the model to identify the maximum level of traceback error that will mitigate moral hazard and motivate producers to deliver safe food. We also identify the conditions under which the traceability system has no effect on the behaviour of suppliers. Our results suggest that efforts directed toward reducing errors in sampling and traceback and subsidizing safe production practices are likely to have the greatest impact on the safety of the food supply.

II. MORAL HAZARD & THE COST OF UNSAFE FOOD

A. Moral hazard

Insulation from risk, commonly called moral hazard, can have a significant effect on behaviour. The concept of moral hazard developed in the insurance industry where companies noticed that people who are insured behave differently than individuals who are not insured. For example, the existence of deposit insurance has been shown to affect the behaviour of bank managers [1] and the existence of reinsurance markets has been shown to influence the behaviour of insurance companies themselves [2].

In economics, principal-agent models are often used to represent economic relationships that exhibit moral hazard. In these models, the principal's payoff depends on the behaviour of the agent, but the principal cannot observe the agent's behaviour – only the outcome. The agent decides whether to shirk or to exert the effort that maximizes the principal's objective. The moral hazard arises because the agent prefers to exert less effort, but the principal wants the agent to work harder.

In the food supply chain, the economic relationship between suppliers and buyers exhibits moral hazard because there are may food attributes that are difficult for buyers to evaluate. Buyers cannot see whether or not a food product was grown under organic conditions, for example, or whether or not a product is contaminated. The buyer must rely on the supplier to exert the effort required to make sure food is uncontaminated. Suppliers who are insulated from the risk of unsafe food will be tempted to shirk their responsibility to provide safe food.

B. Mitigating Moral Hazard and Traceability

Multiple mechanisms have evolved to mitigate moral hazard in the food supply chain. For example, third-party certifications are often used to ensure that suppliers are providing hidden food attributes. Organic production and some kinds of ISO certifications fall into this category. Product guarantees and performance bonds are other mechanisms for mitigating moral hazard.

The best way to mitigate moral hazard, however, is to expose the supplier to the costs of unsafe food. Mechanisms that expose the producer to the cost of unsafe food are difficult to implement because they require that the producer of the unsafe food is identifiable and can be made to pay the cost of unsafe food. In order for a supplier to be identifiable, the source of the food must be traceable.

Traceability is the ability to trace a product and its attributes through the supply chain [3]. Traceability systems can be classified by breadth, depth, and the precision of the information that is collected and distributed. The amount of information that moves along the supply chain with a product is called the breadth. This information could include when the product was harvested, where it was grown, what inputs were used in the production process, etc. The depth of the traceability system refers to how far along the supply chain information is retained. The length of the supply chain can be measured by the number of processing stages and by the number of days or weeks since harvest. Finally, the quality of the information in the traceability system is the precision of the system. More precise information comes from smaller tracking units (lots), more accurate sampling procedures, and better diagnostic methods.

Traceability systems have several functions related to food safety [4]. First, traceability systems allow the impact of a food safety failure to be controlled. If the source of unsafe food can be identified, the food can be recalled and removed from the market before the costs escalate. Second, traceability systems allow for verification of quality attributes. As the product moves along the supply chain, information is collected and is available to potential buyers. Finally, traceability systems allow the cost of unsafe food to be allocated to the individual or firm responsible for the failure. This last function of traceability system has the greatest impact on mitigating moral hazard and providing an incentive to exert the effort needed to produce safe food.

C. Food Safety Events that Generate Cost

Unsafe food causes a number of different kinds of costly events. Some of these costs are paid by the supplier, some are paid by the consumer, and some are paid by all the suppliers in the industry. We identify three categories of unsafe food costs: test failure costs, direct failure costs, and indirect failure costs.

The cost of pulling a contaminated lot out of the supply chain before it moves to the next stage is called test failure cost. The magnitude of test failure cost depends on the disposition of the contaminated lot. The test failure cost is larger if a contaminated lot is scrapped and smaller if the contaminated lot is sold in a secondary market where safety is less of an issue (like pet food).

If a contaminated lot passes inspection, reaches consumers, and causes an illness, it generates a direct failure or indirect failure cost. When the contaminated lot can be traced to the responsible producer, the producer pays a direct failure cost. When the contaminated lot cannot be traced to the responsible producer, all producers in the industry pay an indirect failure cost

Direct failure costs can be significant. The producer may be held liable for the cost of the illnesses, be required to recall all potentially contaminated food, and be fined by government agencies if shown to be negligent. Indirect failure costs can also be significant, but they are often shared by many firms in the industry. Indirect costs include lost sales, equity erosion, and loss of goodwill. For example, the source of the recent outbreak of Salmonella in the United States could not be identified for many months. Before the contamination was attributed to fresh Jalapeno peppers from Mexico, the government incorrectly identified tomatoes as a possible culprit. U.S. tomato producers faced massive losses even though the source of the contamination was eventually identified as jalapeño peppers grown in Mexico.

The magnitude of unsafe food costs depends on whether the food causes an outbreak of illness and how many people become ill. In the US, Buzby *et al* [5] showed that even when the producer is identified, the actual direct cost to the producer is rarely significant. The producer's insulation from the external failure cost is an example of moral hazard.

Figure 1 depicts the events that lead to test failure costs, direct failure costs, and indirect failure costs. In the figure, nodes represent random events that influence the final cost to the unsafe food producer, which is represented by the final square node. The end node labelled "no cost" means that the producer does not pay any extra costs because of the unsafe food. The end node labelled "Test failure" means that the lot failed the safety test and did not move past that point in the supply chain. The end node labelled "Direct failure" means that the contaminated lot caused an illness and the origin of the lot can be traced to the producer. Finally, the end node labelled "Indirect failure" represents the not uncommon case in which a contaminated lot causes an illness but the origin of the lot cannot be identified.

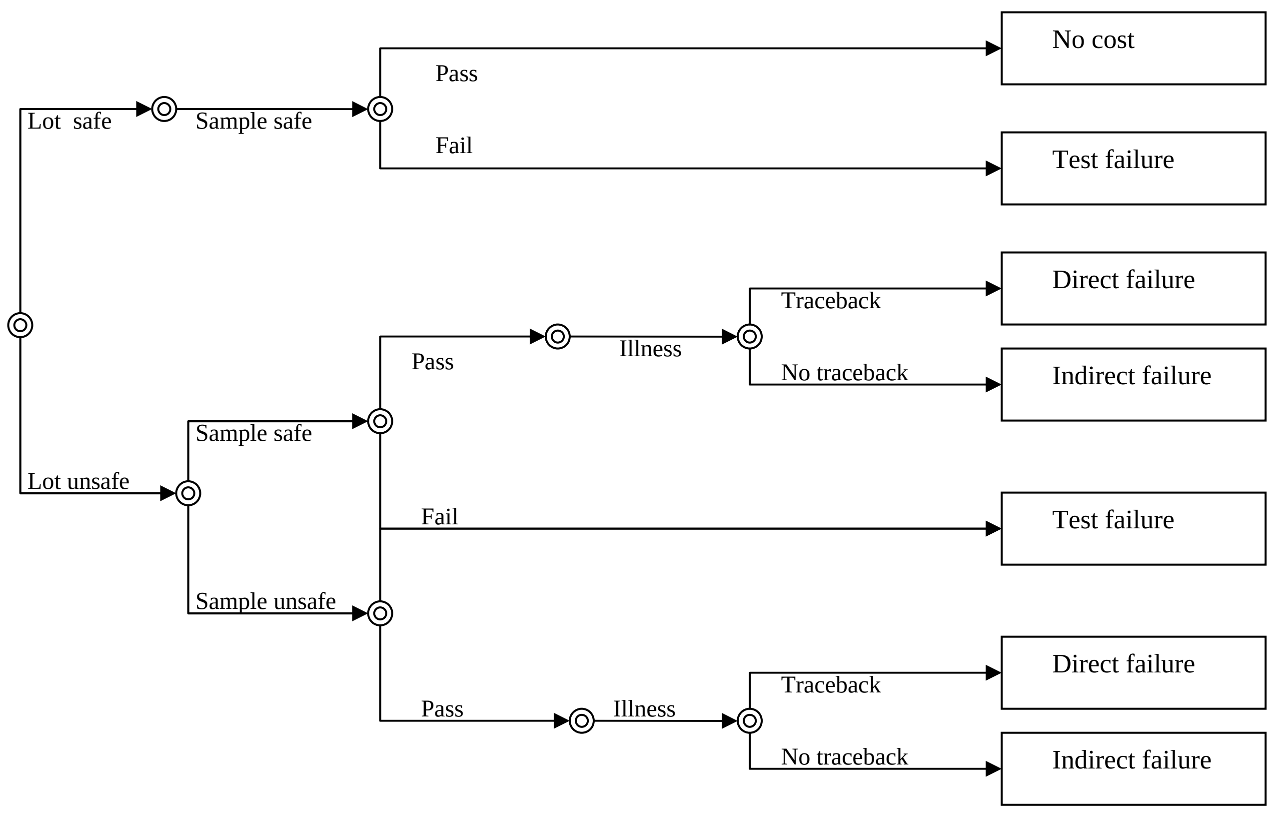

Figure 1. Events Leading to Unsafe Food Costs in the Supply Chain.

III. EXPECTED COSTS & ECONOMIC INCENTIVES

In the preceding section, we identified a number of different unsafe food events that generate costs. Of course, producers don't know which series of events will occur and so they make decisions based on expected cost. Expected cost, and the producer's expected utility, depends on the magnitude of the cost and the probability that the cost occurs.

A. The Probability of Unsafe Food Costs

The random events depicted in Figure 1 have quantifiable probabilities that depend on the errors in sampling, diagnosis, and traceback, and on the probability that unsafe food results in an illness. Of the three errors, diagnostic error is the best understood and most thoroughly studied. Sampling error and traceback error are poorly understood in the food supply chain.

Sampling error refers to the probability that a sample does not represent the characteristics of a lot. For example, if an uncontaminated sample is drawn from a contaminated lot, then a sampling error has occurred. (In this analysis we assume that it is impossible to draw a contaminated sample from an uncontaminated lot.) Sampling error is well-defined for many kinds of experiments, but it is poorly defined, and rarely studied, in the analysis of food safety. We define ε as the probability that a sample is uncontaminated given a lot is contaminated.

Diagnostic error refers to the probability that a test for contamination provides an incorrect result. There are two types of diagnostic error: false-positive and false-negative. A false-positive error shows contamination when a sample is uncontaminated and a false-negative error shows no contamination when a sample is contaminated. False-positive errors mean that safe food is unjustly condemned and false-negative errors mean that unsafe food moves along the supply chain. We define α as the probability that at test shows no contamination given a sample is contaminated and β as the probability that a test shows contamination given a sample is uncontaminated.

Traceback error refers to the probability that unsafe food that causes an illness cannot be traced to the responsible producer. Traceability systems in the US are poorly developed and the source of most food borne illnesses is never identified. When the source is successfully identified, it may take weeks or months for government agents to identify the producer responsible for the contaminated food. We define γ as the probability that the source of a contaminated lot that causes an illness cannot be traced to its source.

The relationship between these error rates and the costly events depicted in Figure 1 are shown in Table 1. Each of the error rates is assumed to be independent of the others. For simplicity, we assume that every contaminated lot that passes the safety test results in an illness (in reality, of course, many contaminated products never cause an illness).

Table 1. Error rates and the probability of costly food safety events.

Event	**Probability**
Contaminated lot	q
Uncontaminated lot	$1-q$
Error rates	
Sampling error	ε
False-negative error	α
False-positive error	β
Traceback error	γ
Food Safety Events	
No cost (nc)	$(1-\beta)(1-q)$
Test failure (tf)	$\beta(1-q)+\beta\varepsilon q+(1-\alpha)(1-\varepsilon)q$
Direct failure (df)	$\gamma[(1-\beta)\varepsilon q+\alpha(1-\varepsilon)q]$
Indirect failure (if)	$(1-\gamma)[(1-\beta)\varepsilon q+\alpha(1-\varepsilon)q]$

In Table 1, we can see that the probability that the producer incurs "No cost" is the product of the probability that a test shows no contamination when the lot is uncontaminated (the test specificity = 1-β) and the probability that a lot is uncontaminated (1-q). It is interesting to note that the probability of no cost is independent of the sampling and traceback errors.

B. The Risk of Illness

From a public safety perspective, one of the most important risks is the risk that an unsafe lot passes the safety test and causes an illness. If we assume that all contaminated lots that pass inspection cause an illness, then the risk of illness is the sum of the probability that a contaminated sample drawn from a contaminated lot is evaluated correctly, and the probability that an uncontaminated sample drawn from a contaminated lot is evaluated incorrectly. The probability that a contaminated lot causes an illness is, therefore:

$$\delta = \varepsilon(1-\beta) + (1-\varepsilon)\alpha \tag{1}$$

It is important to note that this probability depends on the sampling error, which depends on sampling frequency, sample size, and collection methods. In most food supply chains, the sampling error is poorly understood, and rarely studied.

C. Moral Hazard and the Incentive Restriction

The evaluation of moral hazard in a commercial transaction involves defining utility functions for a buyer (principal) and for a supplier (agent). The agent selects an effort level that maximizes her utility, the agent's effort results in an outcome that impacts the principal's utility, and the principal prefers outcomes from higher effort. The agent experiences a disutility from working harder and is, therefore, motivated to shirk. The goal of the solving this kind of problem is to identify a contract, defined as a transfer payment from the principal to agent, which will motivate the agent to work harder.

One of the first steps in evaluating the moral hazard in a commercial transaction is defining the participation constraint and the incentive compatibility constraint. The participation constraint (also called the individual rationality condition) defines the conditions under which an agent will be willing to supply the principal. In order to be willing, the agent must be able to earn more than some minimum level of utility. In most analyses, the minimum is set to zero for simplicity.

The second constraint, the incentive compatibility constraint (also called the incentive restriction), ensures that the agent maximizes utility. An agent won't select higher levels of effort unless it optimizes the agent's utility. The challenge is to link the agent's effort to the agent's payoff in a way that is consistent with the principal's preferences. Macho-Stadler and Perez-Castrillo [6] provide and excellent introduction to the moral hazard problem and the definition of these constraints. Their analysis confirms the importance of the individual rationality condition and the incentive restriction. In cases where information is symmetrical, the individual rationality condition defines the principal's optimal policy, but when information is asymmetrical, the incentive restriction defines the best course of action for the principal.

D. Incentive Compatibility and Traceback Error

We restrict our analysis to the effect of traceback error on the incentive compatibility constraint. We know that information is asymmetrical and that the principal will strive to find a contract that will induce the agent to exert high effort. Our questions are related to how traceability – as represented by traceback error – influences the feasibility of motivating suppliers to deliver safe lots instead of unsafe lots. In other words, how does traceability influence the principal's ability to mitigate the supplier's moral hazard.

We begin the analysis by defining the agent's expected utility function. We assume that the contamination rate is a function of effort and that the agent is risk neutral. Under these assumptions, the agent's decision variable becomes the contamination rate, q, and the agent's expected utility function becomes:

$$\begin{aligned} E[U(q)] = p_{nc}(q)w - p_{tf}(q)C_{tf} \\ - p_{df}(q)(C_{df} - w) \\ - p_{if}(q)(C_{if} - w) - C(q) \end{aligned} \tag{2}$$

In (2), p_{nc}, p_{tf}, p_{df}, and p_{if} represent the probabilities of the "No cost," "Test failure," "Direct failure," and "Inspection failure" events, respectively, described in Figure 1. These event probabilities are defined in Table 1 and are, of course, a function of the contamination rate q. The costs C_{tf}, C_{df}, and C_{if} represent the cost of the test failure, direct failure, and indirect failure events. The transfer payment from the principal, w, is exogenous for the purposes of this analysis. We assume, perhaps optimistically, that the supplier knows the cost of producing lots $C(q)$ with the selected contamination rate q.

In the case where the contamination rate and the effort are continuous, the incentive restriction is:

$$q \in Arg \max_{q_o} \left\{ E[U(q_o)] \right\} \tag{3}$$

In other words, for a given transfer from the principal (w) and given costs and probabilities for test failure, direct failure, and inspection failure, the supplier will select a contamination rate that maximizes the supplier's utility function. When effort and contamination rate are continuous, however, the evaluation of incentive restrictions becomes complex and, for most analyses, unmanageable [6].

In the interest of analytical tractability, we assume that the supplier is choosing between only two levels of effort, low and high. When a supplier exerts high effort, the contamination rate of the supplier's lots is relatively low and when the supplier exerts low effort, the contamination rate of the supplier's lots is relatively high. We define q^L as the low contamination rate generated by high effort and q^H as the high contamination rate generated by low effort. Producing low contamination rate lots costs $C(q^L)$ and producing high contamination rate lots costs $C(q^H)$. Of course, we assume that $C(q^L) > C(q^H)$.

Under these assumptions, the incentive restriction becomes:

$$E\left[U(q^L)\right] \geq E\left[U(q^H)\right] \tag{4}$$

The agent's expected utility at low contamination should be higher than the expected utility at high contamination.

The questions that naturally arise are when is condition (4) true? How does traceback error affect condition (4)? When is condition (4) always true and when is condition (4) never true? When condition (4) is always true, then high contamination will never be selected by the supplier and when condition (4) is never true, then low contamination will never be selected by the supplier.

IV. ANALYSIS OF THE INCENTIVE COMPATIBILITY CONSTRAINT

In the last section we developed a model that relates the contamination rate to the supplier's risk neutral utility. In this section, we use this model to analyze the connection between the effectiveness of the traceability system to the mitigation of the agent's moral hazard. The effectiveness of the traceability system is measured by the traceback error.

A. Traceback Error and Expected Utility

We can uncover the relationship between traceback error and expected utility by substituting (2) into condition (4). This yields:

$$\begin{aligned} &p_{nc}^L w - p_{tf}^L C_{tf} - p_{df}^L (C_{df} - w) \\ &\qquad - p_{if}^L (C_{if} - w) - C^L \\ &\qquad \geq \\ &p_{nc}^H w - p_{tf}^H C_{tf} - p_{df}^H (C_{df} - w) \\ &\qquad - p_{if}^H (C_{if} - w) - C^H \end{aligned} \tag{5}$$

where superscripts indicate the argument is either q^L or q^H. The no cost and test failure events are independent of the traceback error, but the direct failure and indirect failure events depend on traceback. If we isolate γ, the traceback error, and rearrange (5), we find:

$$\frac{(p_{nc}^{L}-p_{nc}^{H})w-(p_{tf}^{L}-p_{tf}^{H})C_{tf}-(C^{L}-C^{H})+\delta(q^{H}-q^{L})(C_{df}-w)}{\delta(q^{H}-q^{L})(C_{df}-C_{if})}\geq\gamma \quad (6)$$

Equation (6) identifies the maximum level of traceback error that is consistent with the incentive restriction. In other words, if (6) is true, then the supplier will elect to exert the effort to produce low contamination product and if (6) is not true, then the supplier will elect to deliver high contamination product.

Figure 2 illustrates the impact of traceback error on the incentive restriction. (The parameter values shown in Table 2 were used in the developing the Figure 2.) Figure 2 shows that the expected utility functions at low contamination and high contamination can intersect. The traceback error at which the two lines cross is defined by the LHS of (6). To the right of the intersection, the supplier will elect to deliver the higher level of contamination, because the traceback error insulates the supplier from the risk of illness caused by the contamination. To the left of the intersection, the supplier will elect to deliver the lower level of contamination, because the traceback error is too low to insulate the supplier from the cost of unsafe food.

Figure 2 shows the intersection of the two expected utility lines, however, it is possible that the two lines never intersect. Under certain conditions low contamination always generates more utility that high contamination and under other conditions high contamination always generates more utility than low

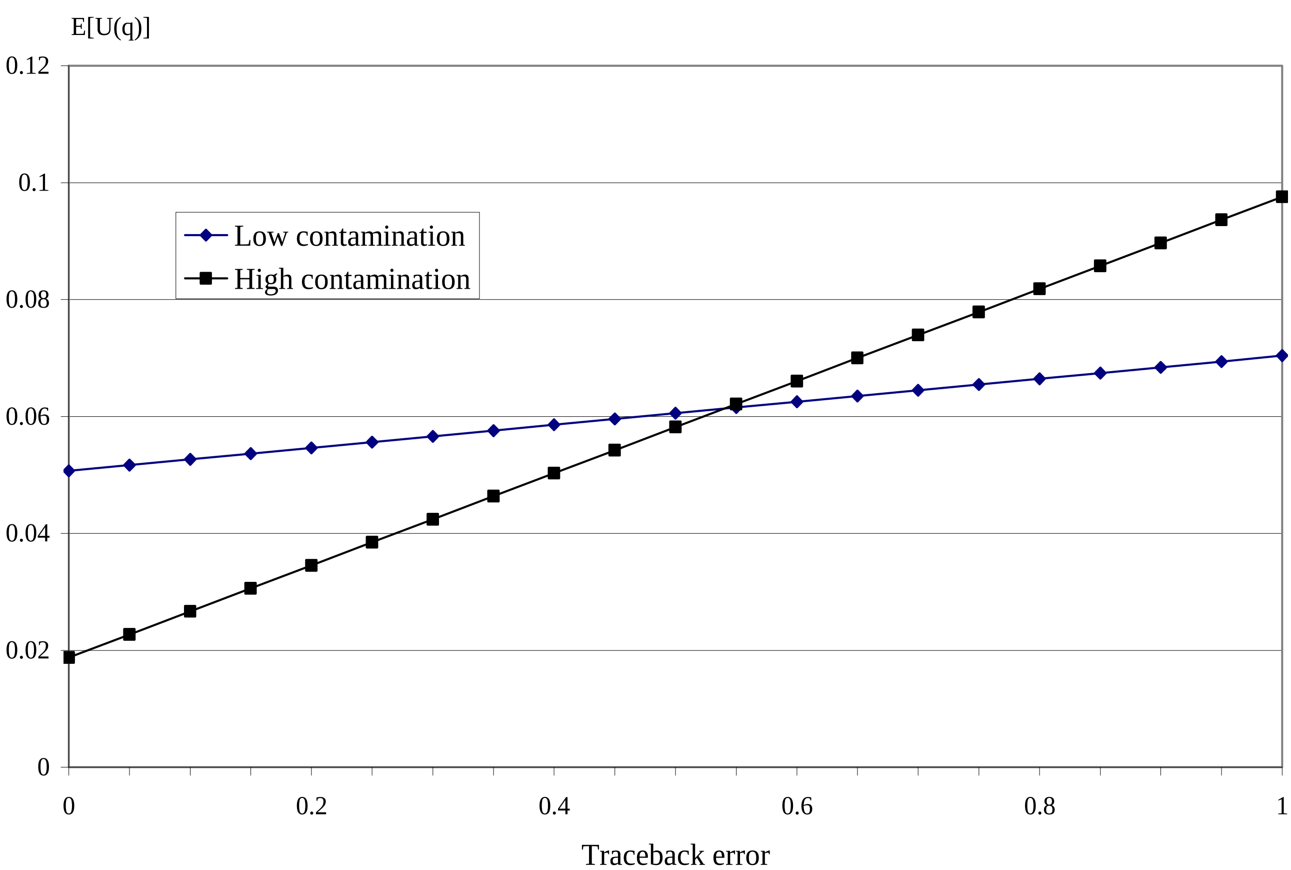

Figure 2. The Relationship between Traceback Error and Expected Utility

contamination. In the next section we identify these conditions.

B. Irrelevant Traceback

It is possible that the effectiveness of the traceability system has no effect on the expected utility of the agent. This occurs when the expected utility of low contamination is greater than the expected utility of high contamination, or vice versa, for all possible values of the traceback error, *γ*.

Regardless of the traceback error, low contamination will never generate more expected utility than high contamination when the LHS of (6) is less than or equal to 0. The LHS of (6) is less than or equal to 0 when $C_{df} \leq C_{if}$, i.e. direct failure is not as costly as indirect failure (an unlikely occurrence) or when:

$$(p_{nc}^L - p_{nc}^H)w - (p_{tf}^L - p_{tf}^H)C_{tf} + \delta(q^H - q^L)(C_{df} - w) \leq (C^L - C^H) \quad (7)$$

The LHS of (7) is the incremental expected utility of the transfer payments, test failure costs, and direct failure costs associated with reducing the contamination rate. The RHS of (7) is the incremental increase in the production cost. Condition (7) implies that the benefit of the reduction in the contamination rate is less than the production cost increase, and so the agent should not exert the effort to make the food safer.

We can perform a similar analysis to identify the conditions under which low contamination is always better for the agent, regardless of the traceability system. This is the case when the LHS of (6) is greater than or equal to 1. The LHS of (6) is greater than or equal to 1 when:

$$(p_{nc}^L - p_{nc}^H)w - (p_{tf}^L - p_{tf}^H)C_{tf} + \delta(q^H - q^L)(C_{if} - w) \geq (C^L - C^H) \quad (8)$$

The difference between conditions (7) and (8), besides the inequality, is the third term on the LHS. Condition (7) depends on the cost of direct failure and in condition (8) depends on the cost of indirect failure. Condition (7) implies that if the cost of direct failure is very low, then it doesn't matter what the traceability level is, because being identified as the source will not be very costly. Similarly, condition (8) implies that if the indirect failure cost is high, it does not really matter what the traceback error is, because the agent is paying a significant cost even though she is never identified as the source of the unsafe food.

The analysis in this section provides guidance to regulators and consumers (principals) who are trying to develop supply chains in which agents are motivated to produce low contamination product.

Table 2. Parameter values used in the development of Figure 2.

Contamination-dependent Parameters	**Symbol**	**Value**
Low contamination rate	q^L	0.5%
High contamination rate	q^H	2.0%
Low contamination production cost	C^L	0.90
High contamination production cost	C^H	0.85
Error rates		
Sampling error	ε	50%
False-negative error	α	0.5%
False-positive error	β	2.0%
Transfers & Costs		
Transfer from principal	w	1.00
Test failure	C_{tf}	0.10
Direct failure	C_{df}	10.0
Indirect failure	C_{if}	2.00

C. *Paying to Mitigate Moral Hazard*

The agent's expected utility is influenced by the parameters listed in Table 2. The principal has direct control over the transfer payment, *w*. How much does the principal have to pay to mitigate moral hazard and make low contamination the optimal policy for the agent? We can answer this question by rearranging (6) to isolate *w*. Doing so yields:

$$w \geq \frac{(p_{tf}^L - p_{tf}^H)C_{tf} + (C^L - C^H) - \delta(q^H - q^L)((1-\gamma)C_{df} + \gamma C_{if})}{(p_{nc}^L - p_{nc}^H) - \delta(q^H - q^L)} \quad (9)$$

In (9), the RHS represents the minimum transfer payment needed to mitigate moral hazard in the presence of traceback error when the incentive compatibility constraint is binding. We define w_{ic} as the RHS of (9)

When traceback error is 100%, w_{ic} is at its maximum and when traceback error is 0%, w_{ic} is at its minimum. For example, using the parameter values in Table 2, and assuming the traceback error is 100%, the principal will need to pay w_{ic} = 4.72 to mitigate moral hazard in the supply chain. As the traceback error declines, the minimum transfer also declines.

At some point the incentive compatibility constraint ceases to become active and the participation constraint becomes binding. When w_{ic} reaches 0, any transfer payment from the principal will satisfy the agent's incentive restriction. However, the participation constraint will become active before that happens. If we assume that the minimum expected utility for agent participation is 0, and rearrange (2) for q^L, we find:

$$w \geq \frac{p_{tf}^L C_{tf} + C^L + p_{df}^L C_{df} + p_{if}^L C_{if}}{p_{nc}^L + p_{df}^L + p_{if}^L} \quad (10)$$

So the transfer price that satisfies the incentive compatibility AND the participation constraint is the maximum of w_{ic} and the RHS of (10), which we denote w_{pc}. More formally:

$$w_{\min} = \max\{w_{ic}, w_{pc}\} \quad (11)$$

We can think of the amount by which w_{ic} exceeds w_{min} as the traceback inefficiency premium, τ. Or,

$$\tau = \max\{0, w_{ic} - w_{pc}\} \quad (12)$$

The traceback inefficiency premium represents the amount that the principal must pay the agent to mitigate moral hazard because of the inefficiency in the traceability system.

V. IMPLICATIONS AND CONCLUSIONS

In the preceding section, we identified the conditions under which low contamination satisfies the agent's incentive restriction. The principal can design a contract that mitigates the moral hazard inherent to the transaction if the principal can motivate the agent to produce safer food. We also identified the conditions under which the agent's incentive restriction will never be satisfied and the conditions under with the incentive restriction are always satisfied.

Regulators and buyers can use this information to develop terms of trade that mitigate moral hazard associated with unsafe food. The error rates, unsafe food costs, production costs and transfer payment all influence the existence of moral hazard. Some of these parameters are easier to manipulate than others in the effort to increase the utility of producing safer food. Below we discuss several of these parameters and the potential for creating a supply chain in which low contamination is optimal.

Error rates. Advances in testing technology have reduced diagnostic errors to almost zero. For example, false-positive and false-negative errors for the Dupont BAX system used to test for *E. coli* in ground beef were found to be negligible in independent tests [7]. Sampling error and, of course, traceback error could be quite a bit larger than diagnostic error. No studies designed to estimate these errors have been performed in the US. Since diagnostic error is already close to 0, there is significant potential to manipulate the terms of trade by reducing traceback and sampling error. Many private food service companies in the US have already begun to control traceback error by limiting the number of suppliers and requiring suppliers to be certified for safety and quality.

Unsafe food costs. Some of the unsafe food costs described above can be controlled by the buyer and some cannot. The cost of test failure is, for the most part, outside the control of the buyer. Scrap costs and salvage values are, for the most part, driven by market conditions and are not controllable. Direct failure costs have more potential for control. Direct failure costs include fines, recalls, and class action suits, all of which can be manipulated by regulators and buyers to mitigate moral hazard. Indirect failure costs can also be controlled, but by fining an industry, for example, many innocent producers will pay a portion of the costs generated by one unsafe producer.

Production costs. Production costs offer significant opportunity to mitigate moral hazard. By subsidizing the production of safe food the incremental cost declines and condition (8) is more likely to hold. Government support for safe food production technology could have a meaningful impact on the safety of the food supply.

In conclusion, our analysis reveals that inefficiencies in the traceability system significantly impact the existence of moral hazard in the food supply chain. The existence of moral hazard influences the incentive to produce safe food. High traceback error means there is less of incentive to produce safer food and low traceback error supports the production of safe food. Under certain circumstances, traceback efficiency is irrelevant with respect to the supplier's incentive to produce safe food. Regulators who wish to manipulate the terms of trade in order to mitigate moral hazard should focus their attention on the cost of producing safe food, the sampling and traceback error, and the cost of direct failures. These parameters have the greatest potential for improving the safety of the food supply.

REFERENCES

1. Grossman R (1992) Deposit insurance, regulation, and moral hazard in the thrift industry: evidence from the 1930s. American Economic Review 82: 800-821.
2. Doherty N and Smetters K (2005) Moral hazard in reinsurance markets. J of Risk and Insurance 72:375-391.
3. Golan E, Krissoff B, Kuchler F, Calvin L, Nelson K, and Price G (2004) Traceability in the US food supply: economic theory and industry studies. Agricultural Economic Report 830, Economic Research Service, USDA. Washington DC, USA.
4. Hobbs JE (2004) Information asymmetry and the role of traceability systems. Agribusiness 20:397-415.
5. Buzby JC, Frenzen PD, and Rasco B (2001) Product liability and microbial foodborne illness. Agricultural Economics Report 799, Economic Research Service, USDA. Washington DC, USA.
6. Macho-Stadler I, Perez- Castrillo JD (2001) An introduction to the economics of information: incentives and contracts, Second edition, Oxford University Press, New York.
7. Hochberg AM, Gerhardt PNM, Cao TK, Ocasio W, Barbour WM, and Mrozinski PM (2000) Sensitivity and specificity of the test kit BAX for screening *E. coli* O157:H7 in ground beef: independent laboratory study. Journal of the AOAC International 83:1349-56.