Ripoffs, Lemons, and Reputation Formation in Agency Relationships: A Laboratory Market Study
Author(s): Douglas V. Dejong, Robert Forsythe, Russell J. Lundholm
Source: *The Journal of Finance*, Vol. 40, No. 3, Papers and Proceedings of the Forty-Third Annual Meeting American Finance Association, Dallas, Texas, December 28-30, 1984 (Jul., 1985), pp. 809-820
Published by: Blackwell Publishing for the American Finance Association
Stable URL: http://www.jstor.org/stable/2327806
Accessed: 26/11/2009 08:43

THE JOURNAL OF FINANCE • VOL. XL, NO. 3 • JULY 1985

Ripoffs, Lemons, and Reputation Formation in Agency Relationships: A Laboratory Market Study

DOUGLAS V. DEJONG, ROBERT FORSYTHE, and RUSSELL J. LUNDHOLM*

ABSTRACT

This paper examines the effect of the moral hazard problem in an agency relationship where the principal cannot observe the level of service provided by the agent. Using data from laboratory markets, we demonstrate that the presence of moral hazard leads to shirking by agents. However, this "lemons" phenomenon occurs only about one-half of the time. While there is evidence of reputation effects in these markets, seemingly reputable agents are often able to use opportunities for false advertising to their advantage and "ripoff" principals.

THERE HAS BEEN CONSIDERABLE discussion about the price and quality of services provided within certain principal-agent relationships. For example, the Securities and Exchange Commission and Congress have focused a great deal of attention on the price and quality of outside audit services to investors. The Securities and Exchange Commission and the Department of Justice have considered the regulation of investment banking contracts because of their concern over the price and level of services provided. Finally, a large and growing literature has focused attention on the conflict between owners and managers and the problem of shirking by managers.

The above examples share a feature common to most agency relationships in that the services provided by the agent have a hidden characteristic. Specifically, the principal cannot costlessly observe the level of service provided by the agent and, thus, cannot determine whether an unfavorable outcome is due to the agent's shirking or due to an unfortunate occurrence of a state of nature. A large theoretical literature has developed (e.g., Shavell [9]), which proposes remedies for dealing with the moral hazard problem that arises in these situations. Yet, to our knowledge, there has been no direct empirical test demonstrating that the moral hazard problem does in fact lead to shirking by agents. Furthermore, the theoretical literature provides a number of competing models about the market's ability to provide Pareto optimal levels of service in the presence of moral hazard.

* Assistant Professor, Associate Professor, and Ph.D. student, respectively, at the University of Iowa. We would like to thank the Department of Accounting at the University of Iowa for providing the funds for the experiments, and W. Uecker and participants of the Accounting Workshops at the University of Iowa and University of Oklahoma for their comments and suggestions. Financial support for DeJong was provided, in part, by the Arthur Young Faculty Fellow Program at the University of Iowa.

In this paper, we examine the impact of the moral hazard problem in agency relationships using the data from laboratory markets. This data will then be used to test the predictions of three competing models. In the following section, we outline the experimental design and give a description of the laboratory markets. In Section II, we present the three models which give competing price and quality predictions for the laboratory markets we examine. The third section contains a discussion of our results and the final section is a summary.

I. The Laboratory Markets

In our laboratory environment, each principal faced a loss whose probability could be reduced by purchasing services from an agent. If a principal purchased such services, the higher the quality of service provided by the agent, the lower the probability of loss faced by the principal. However, a principal was never able to directly observe the quality of services provided by the agent.

The procedures employed in the laboratory markets were an extension of the procedures used in DeJong, Forsythe, and Uecker [4], hereafter DFU. Each market was conducted on the computer using the Prime system.[1] Each market consisted of three agents and four principals who were student volunteers at the University of Iowa. All agents were identical with regard to the parameters they faced and principals also faced identical parameters. Trades took place via a sealed offer auction over a sequence of market periods.

At the beginning of a market period, each agent could submit a sealed offer to each principal. The offer specified a price and type (level of service). After receiving all offers, each principal separately chose which agent, if any, he was going to purchase from. This information was then publicly disclosed but the price and type agreed upon were not revealed. After learning which of their offers had been accepted, agents were then asked to enter which type (level of service) they were actually going to deliver to each principal. The type delivered then determined the probability of loss faced by the principal.[2] Based on this probability, the principal and his agent were told whether or not a loss had occurred.

Since the principal always bore the loss, profits depended upon whether a loss occurred. The agent never bore the loss and, thus, profits were simply the difference between the fee received from the principal and the cost of delivering the level of service provided. After each principal and agent calculated their profits, the market period was concluded. Participants then proceeded to a new trading period which was a strict replication of the previous period(s).

[1] The instructions are available from the authors upon request.

[2] The empirical frequencies associated with the loss were preselected to conform to the theoretical frequencies. This procedure ensures that the sequence is "representative" and provides further control over the parameters of the laboratory markets and the probabilities used by participants. This also eliminates some of the potential difficulties reported in DFU where we presented some limited evidence that nonrepresentative samples affected market convergence. However, a potential difficulty with this technique arises if subjects are able to learn that the random sequence they are observing has been preselected. To minimize this possibility, we used inexperienced subjects and randomized the placement of losses within each sequence of ten deliveries for each type of service. There is no evidence that subjects acted on the basis of any probability distribution other than the one they had been told they would be observing.

Table I

Parameter Values for Experiments

Level of Service y	Probability of Loss $P(y)$	Expected Loss[a] $P(y)l$	Change in Prob. of Loss $\Delta P(y)$	Change in Expected Loss $-\Delta P(y)l$	Cost of Service $C(y)$	Change in Cost $\Delta C(y)$	Total Expected Cost[b]	Expected Surplus[c]
Set One								
No purchase	1.0	80¢			0			0
1	0.8	64¢	−0.2	16¢	8¢	8¢	72¢	8¢
3*	0.5	40¢	−0.3	24¢	20¢	12¢	60¢	20¢
5	0.1	8¢	−0.4	32¢	65¢	45¢	73¢	7¢
Set Two								
No purchase	1.0	80¢			0			0
1	0.6	48¢	−0.4	32¢	5¢	5¢	53¢	27¢
3*	0.2	16¢	−0.4	32¢	17¢	12¢	33¢	47¢
5	0.05	4¢	−0.15	12¢	41¢	24¢	45¢	35¢

[a] Loss (l) is 80¢ throughout.
[b] $P(y)l + C(y)$.
[c] 80¢ less total expected cost.
* Efficient level of quality.

Individual principals and agents received an initial endowment at the beginning of each new period of $1.30 and $0.50, respectively. At the conclusion of the market experiment, participants were paid an amount equal to the sum of their profits plus initial endowments for all the trading periods in the market. Payments ranged from $10.00 to $25.00.

Four laboratory markets were conducted using the two different parameter sets presented in Table I. Markets 1 and 2 used parameter set one, and Markets 3 and 4 used parameter set two.[3] The use of two distinct parameter sets provides a more robust test of the models presented in the next section.

II. Competing Models

In this section, we present three alternative models which address the moral hazard problem in a principal-agent framework. The first two models are static in nature. In both models, principals are assumed to use price as a signal of the quality (or level) of service an agent will provide. What distinguishes one model from the other is the assumption each makes about an agent's ability to mask deliveries of low quality services by pricing such deliveries as though they were

[3] In Market 3, which was the first of the four markets conducted, an agent indicated at the conclusion of the experiment that he thought principals knew his costs. This led us to design a set of practice calculations to verify that participants understood the trading procedures, the information available in the market, and their profit calculations. After reading the instructions, principals were asked to complete these practice calculations in the other three experiments. Market 3 is included here since there is no evidence that this agent's erroneous presumption had any effect on market performance. Further, Market 3 has been previously reported as an example of a moral hazard experiment in the conclusion of DFU.

higher quality services. The third model is a reputation model which is, by definition, a dynamic model. In the reputation model, principals attempt to infer the reliability of different agents, and agents are concerned with how the level of service delivered affects future sales.

As in DFU, we assume risk neutral behavior on the part of both principal and agent. Thus, the principal is an expected wealth maximizer who is exposed to a potential loss. The principal has an opportunity to hire an agent to perform a service, the purpose of which is to reduce the probability of the loss. The higher the level (or quality) of service provided by the agent the lower the probability of loss. The agent also maximizes his expected wealth which is increased by the fee he receives for his services but is reduced by service activity because it is costly. Consequently, the agent has an aversion to providing the service.[4]

In this situation, let

l = the amount of the loss

$y \in A$ = the level (or quality) of services provided by the agent where A is the set of all possible levels

$C(y)$ = a twice differentiable, increasing and strictly convex cost function where $C(y)$ represents the agent's cost of providing service level y

$P(y) \equiv P(y; l)$ = the probability of a loss given level of service y. This probability is assumed to be twice differentiable, decreasing and convex

$1 - P(y)$ = the probability of no loss

r = the contract fee between the principal and agent and

α = the probability that the agent will provide the lowest level of service, $\underline{y}$.

Under the assumption of risk neutrality, risk-bearing is not an issue, and the Pareto optimal, or first best, solution is characterized solely by production efficiency. Specifically, optimality requires that the quality of services provided by the agent is the level that minimizes expected total social costs, $C(y) + P(y)l$. That is, the optimal level of service provided, y^*, is such that

$$-P'(y^*)l = C'(y^*) \tag{1}$$

where $P'(\cdot)$ and $C'(\cdot)$ represent the first derivatives of $P(\cdot)$ and $C(\cdot)$. Finally, for convenience (but with no loss of generality), we limit the possible actions that an agent may take to three, where $A = \{\underline{y}, y^*, \bar{y}\}$ and $\underline{y} < y^* < \bar{y}$. The levels $\underline{y}$, y^*, *and* $\bar{y}$ correspond to type 1, type 3, and type 5, respectively, in Table I. For both parameter sets, y^* (type 3) minimizes expected total cost.

[4] The situation described here resembles the auditor-investor problem where investors in a firm face a loss if an auditor is not retained, and the level of audit services provided by the auditor affects the probability of that loss. The investment banking problem is symmetric to the auditor-investor problem in that a firm perceives an expected gain if an investment banker is engaged and the size of the expected gain depends on the level of service the investment banker provides.

A. The Lemons Model[5]

In this model, principals are aware that they are at the mercy of the agents but are unable to do anything about it. Once hired, a profit-maximizing agent will always deliver the lowest quality level of service, $\underline{y}$, (a "lemon") no matter what level of service the agent has contracted to provide. This occurs because $\underline{y}$ is the level of service with the lowest cost and thus maximizes the agent's profits on a contract. Further, in a competitive market for agents, this level of service will be priced so that the expected profit of each agent is zero, i.e., $r = C(\underline{y})$.

Using the parameters in Table I, the numerical predictions for parameter set one are $r = \$0.08$, $\alpha = 1$ and for set two the predictions are $r = \$0.05$, $\alpha = 1$.

B. The Ripoffs Model[6]

In this model, principals attempt to use the price of an offer to infer the level of service that the agent intends to provide. Realizing this, agents price all their deliveries as though they were high quality services. An agent does not price his offer for less than what it would cost him to provide the quality level offered regardless of the quality he plans to deliver. Otherwise, principals will infer that he is planning to deliver a lower quality and will not make a purchase. A "ripoff" is said to occur when an agent delivers a low quality level of service after offering a higher quality level for the same price that he (and other agents) charge for delivering the higher quality level.

We assume that agents will always offer to provide the socially optimal level, y^*, regardless of the quality the agent plans to deliver,[7] and, in equilibrium, these units will sell for price $r = C(y^*)$. Each agent then must choose the frequency with which to deliver units of optimal quality level y^* and "ripoffs" where $\underline{y}$ is delivered. Notice that $\bar{y}$ is not delivered, because an agent would earn negative profits on each sale of $\bar{y}$. As defined previously, α denotes the probability that the agent will provide $\underline{y}$ and therefore $(1 - \alpha)$ is the probability that he will provide y^*.

While agents are assumed to choose α so as to maximize their expected profits, the frequency with which agents deliver units of quality level $\underline{y}$ at price $r = C(y^*)$ is limited by the principals' decision about remaining in the market. Specifically, principals are assumed to know each agent's choice of α, and principals will remain in the market as long as the expected cost of purchasing a unit from an agent does not exceed the expected cost incurred if he does not purchase. Thus,

[5] This model is similar in spirit to the Arrow [1] model.

[6] This model is similar in spirit to the Cooper and Ross [2] model.

[7] Cooper and Ross [2] assume that some positive fraction of the principals have full information regarding the type of unit an agent will deliver. These informed principals will then purchase only the socially optimal level of service, y^*, which, in equilibrium, will sell for $r = C(y^*)$. So that uninformed principals cannot infer quality from price, all units, including ripoffs, sell for this price. To operationalize this model to the environment we consider, we must assume that agents will offer the socially optimal level since we do not provide full information to any of the principals.

an agent chooses α to

$$\begin{aligned} &\underset{\alpha}{\text{maximize}} && r - (1-\alpha)C(y^*) - \alpha C(\underline{y}) \\ &\text{subject to} && (1-\alpha)(r + P(y^*)l) + \alpha(r + P(\underline{y})l) \le P(0)l \\ & && \alpha \in [0, 1]. \end{aligned}$$

The constraint in the problem ensures that α will not be chosen so as to drive the principal from the market. The solution to this problem is given by

$$\alpha = \min\left[\frac{P(0)l - P(y^*)l - r}{P(\underline{y})l - P(y^*)l}, 1\right]. \tag{2}$$

That is, since each delivery of a $\underline{y}$ level of service increases the agent's profits by $C(y^*) - C(\underline{y})$, he will deliver units of service level $\underline{y}$ until principals are just indifferent between staying in or dropping out of the market or until $\alpha = 1$.

In this model, agents earn positive profits in equilibrium. For this to be a sustainable equilibrium, the constraint in the maximization problem is critical since it requires that agents understand that increasing the frequency of supplying units of level $\underline{y}$ beyond α will drive principals from the market.

The numerical predictions for parameter set one are $r = \$0.20$ and $\alpha = \$0.83$ [from Equation (2)]. For set two, the predictions are $r = \$0.17$ and $\alpha = 1$.

C. A Reputation Model

The lemons and ripoffs models are static in nature. That is, principals and agents are assumed to maximize wealth within a single period framework. In the markets we examine, agents had the opportunity to make repeat sales over a finite time horizon. Klein and Leffler [5] present a model in which reputable agents receive a quality-assuring price which is in excess of their costs of providing high quality services. Furthermore, this quality-assuring price remains above their costs as other reputable agents enter the market because these entrants must incur reputation-building costs to earn this quality-assuring price. They incur these costs by offering and delivering high quality services at low quality prices until they have established their reputation. In equilibrium, the present value of this reputation-building strategy is zero.

The difficulty of applying this reputation model to the markets studied here is that the model relies upon an infinite time horizon. Specifically, with a finite horizon, reputations will never form in a Nash equilibrium. The logic behind this statement is a backward induction argument (e.g., Selten [9]) which begins by observing that in the final period, a static model must apply and agents will behave according to the assumptions of the static model. Principals are assumed to be aware of this and thus do not use each agent's performance in the next to last period as an indicator of what agents will deliver in the last period. Agents are aware that their deliveries in the next to last period have no effect on their reputation and thus the static model again applies in that period. This argument is then used to demonstrate that principal and agent behavior is the same two periods from the end as it was one period from the end, etc., until it is argued that a repeated, finite horizon model will yield the same equilibrium as a static model. Furthermore, using this logic, either of the two static models, lemons or ripoffs, apply.

Alternatively, Kreps et al. [6] model this situation as a game with incomplete information where the principal is unsure of the "type" of agent he has contracted with. In particular, if a principal initially believes that there is a small but positive probability that an agent will honor his commitment and the agent believes there is a small but positive probability that a repeat sale will occur if the commitment is honored, reputation could develop even though there is a finite horizon. Then, an agent, by behaving reputably may receive a quality-assuring price and, by behaving disreputably, believes that he will lose his repeat sales. Thus, the agent may find it in his interest to develop his reputation if his profits from cooperating in this fashion exceed those from behaving as in one of the static models. This cooperation will continue until the principal discerns that the agent is no longer reputable.

Computation of equilibrium strategies in this case is difficult, and it is not possible to provide numerical predictions since they depend upon a principal's subjective probability of contracting with a reputable agent and these probabilities cannot be measured. However, it is possible to provide qualitative evidence in support of a reputation model. Specifically, support for the following two possibilities would constitute positive evidence in support of a reputation model:

1. The frequency of type 1 deliveries is significantly less than the frequencies predicted by either static model.
2. Agents receive a premium for selling and delivering quality level y^* (type 3). This is consistent with the notion of a quality-assuring price.

III. Results

The data from Markets 1 to 4 are presented in Figures 1 to 4. Each figure is organized into three panels with each panel displaying the time series of deliveries of a particular quality type, i.e., type 1, type 3, or type 5. The symbols (1, 3, and 5) on each panel correspond to the types that were offered.

We will present our results as a series of propositions. Each proposition will take the form of a statement as to which model is consistent with the behavior observed in the markets studied. Following the statement of each proposition, its validity will be demonstrated in the form of an argument which presents the supporting evidence.

PROPOSITION 1. *The ripoffs model outperforms the lemons model.*

Argument: The trend of prices in Market 1 (Figure 1) is closer to the \$0.20 prediction of the ripoffs model than the \$0.08 prediction of the lemons model. In Market 2 (Figure 2), however, the prices observed are closer to the lemons model. In Markets 3 and 4 (Figures 3 and 4), the ripoffs model is also supported since the prices seem to have converged to the cost of the level of service offered and not to the cost of the level of service delivered as predicted by the lemons model.

As an estimate of α (the probability that an agent delivers the lowest level of service), the frequency of type 1 deliveries is used, regardless of the types that

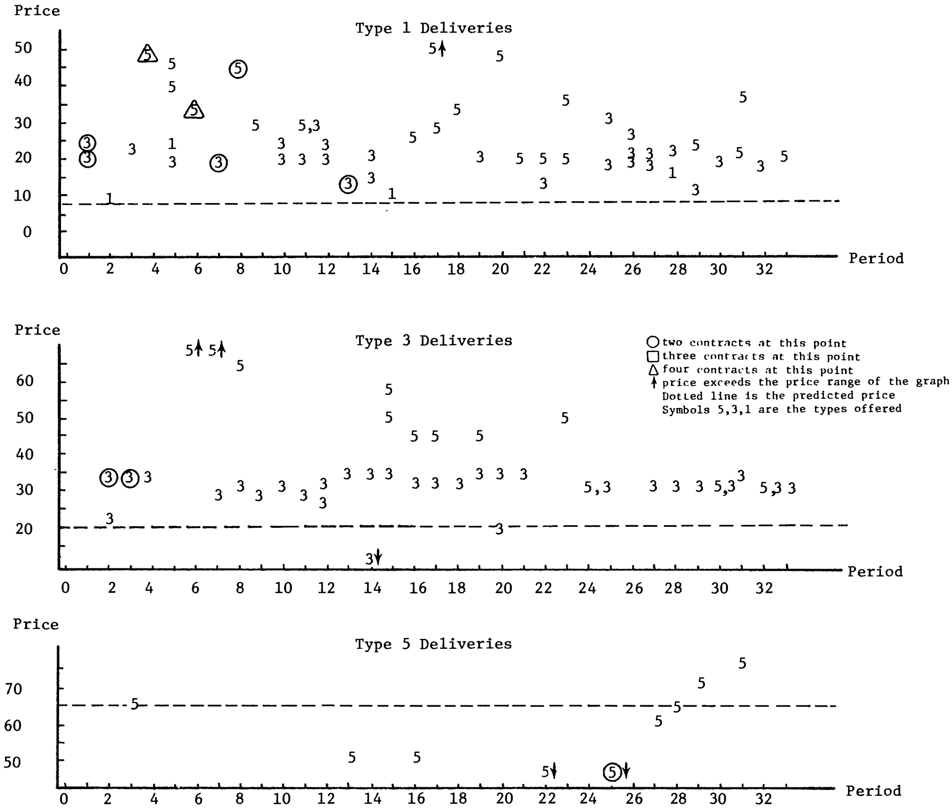

Figure 1. Market 1: Parameter Set One

were offered. In each of the four markets, the frequency of type 1 deliveries was 52.6, 68.2, 51.1, and 30.1%, respectively. The average frequency of type 1's is 51.1%. For parameter set one (Markets 1 and 2), the ripoffs model prediction of 0.83 is closer to the observed frequency than the lemons model prediction of 100%. For parameter set two (Markets 3 and 4), neither model performs well since both predict a frequency of 100%.

Although the ripoffs model outperforms the lemons model, its performance still leaves much to be desired. While it predicts the equilibrium prices in at least 3 of the 4 markets reasonably well, its performance in predicting α is poor. Specifically, averaging the absolute difference between the observed and predicted frequency of α across the four markets yields an average error of 41% for the ripoffs model. Further, the ripoffs model predicts that the frequency of lemons delivered will be higher for the markets with parameter set two than for the markets with parameter set one. To the contrary, the observed frequency of lemons delivered is lower in the markets using parameter set two (51.5 and 30.1%) than in markets using parameter set one (52.5 and 68.2%). Consequently, we turn our attention to a reputations model.

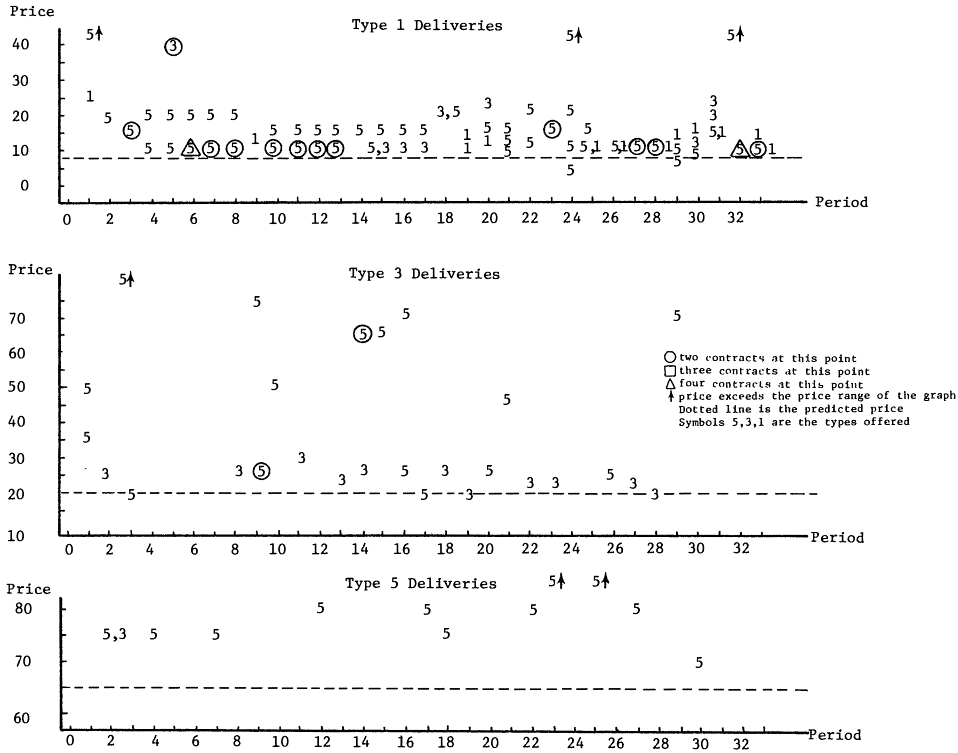

Figure 2. Market 2: Parameter Set One

PROPOSITION 2. *In all markets, there is support for a reputation model. In particular, the frequency of type 1 deliveries is significantly less than the frequency predicted by either static model, and there is evidence of a quality-assuring price.*

Argument: Far less than 100% type 1 deliveries were observed in all four markets. In fact, the frequencies observed are lower than the frequencies predicted by either static model for either parameter set at a 0.01 level of significance. This supports the reputation model and is consistent with the assumption that agents maintain some concern about establishing a reputation for quality.

To test for evidence in support of a quality-assuring price, we used a multinomial logit model to estimate the probability that an agent would deliver a level of service lower than offered when this probability is a function of the level of service offered and the contract price.[8] In all four markets, the coefficient on price was negative and significant at the 0.02 level. This implies that as price increases the probability that the agent delivers a lower quality than offered diminishes, which is consistent with the notion of a quality-assuring price.

Anecdotal evidence in support of a reputation model is also found in Market 1. One principal-agent pair repeatedly contracted with each other from periods 7

[8] For a full description of the multinomial logit model see DeJong et al. [3].

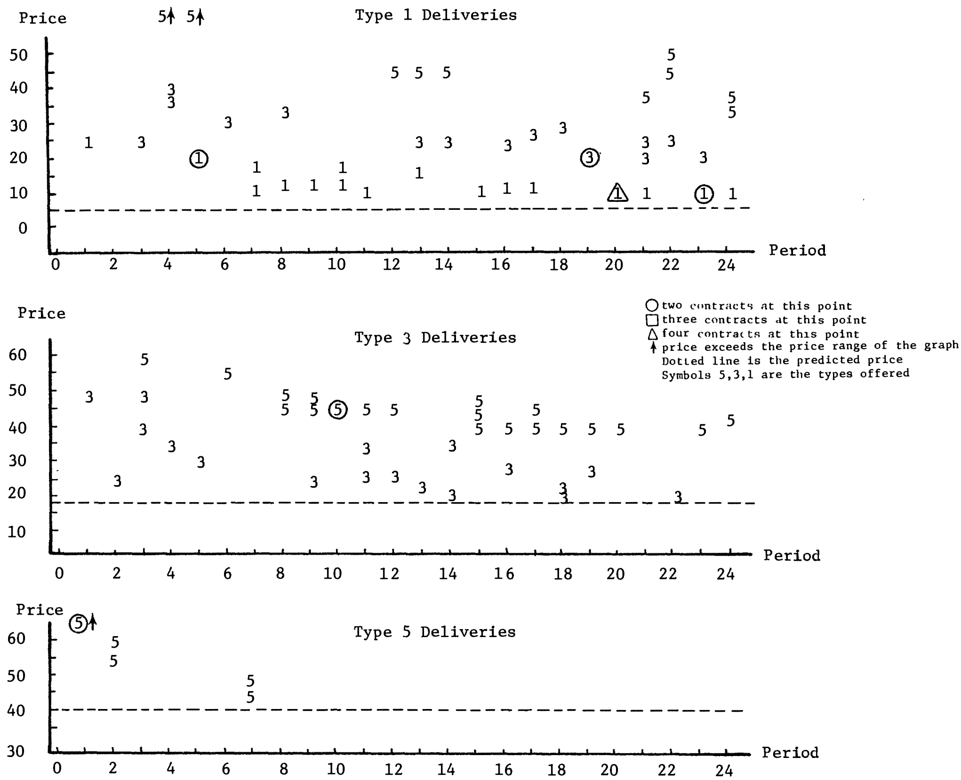

Figure 3. Market 3: Parameter Set Two

through 21, 24, and 27 through 32. In each of these instances, the agent offered and delivered a type 3 at a price of 8 to 15 cents above the competitive price prediction. The principal continued to pay this premium in spite of type 3 offers from other agents at lower prices. A second piece of evidence can be seen in the "Type 5 Deliveries" panel of Figure 1. All sales and deliveries of this quality level from period 22 onward are made by the same agent. This agent seems to be pursuing a reputation building strategy since the initial sales are well below the agent's cost of providing that level of service. Eventually, the agent was able to make sales at a premium over cost in an attempt to cover his earlier losses.

While the argument presented for Proposition 2 shows that higher prices *tended* to bring forth a higher quality of service, it is clear from the figures that agents failed to build and maintain reputations necessary to eliminate type 1 deliveries. We next pursue a potential implication of the failure to establish reputations necessary to control the adverse effects of moral hazard by examining the data to see if price depends only on the level of service delivered or, alternatively, on the level of service offered. If the types of services offered by agents have no impact on price, then any misrepresentations made by agents are harmless. If, however, the level of service offered does effect the price received by agents, then these false advertisements redistribute income from principals to agents and, as such, constitute "ripoffs."

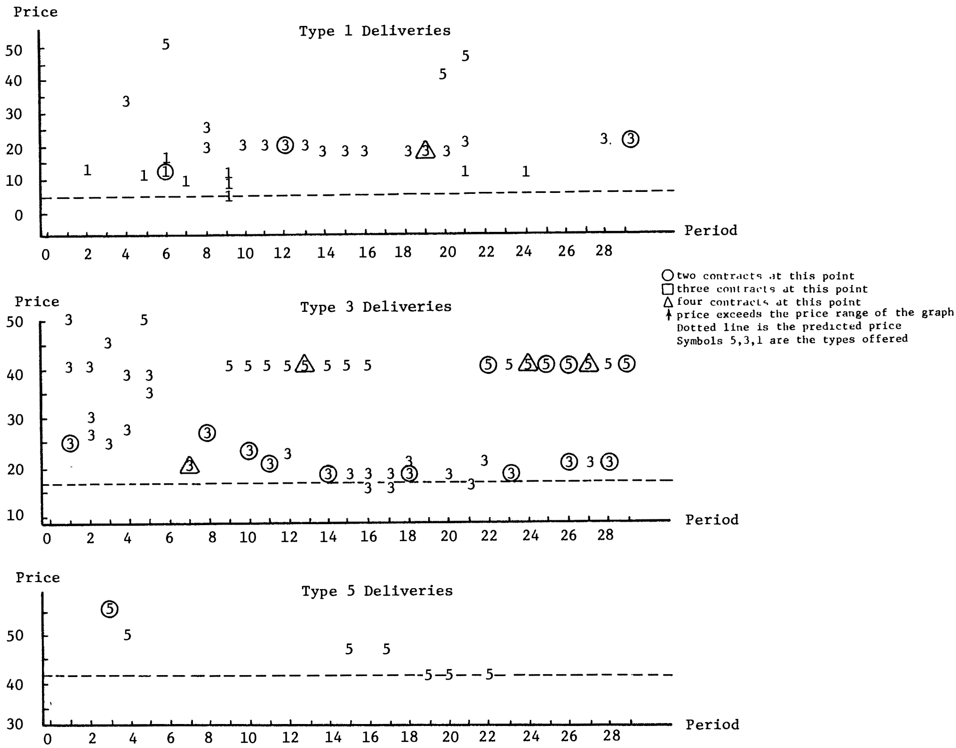

Figure 4. Market 4: Parameter Set Two

PROPOSITION 3. *Both the level of service offered and level of service delivered have a significant effect on the prices received by agents.*

Argument: To test this proposition we performed an ANOVA, with price as the dependent variable and type offered and type delivered as the main effects. In all but Market 2, the type offered explains a significant portion of the variance in price at the 0.01 level. The type-delivered variable also explains a significant portion of the variance in all four markets at the 0.01 level. Thus, it appears that while ripoffs do indeed occur, principals are not completely fooled by agents.

IV. Summary

We have demonstrated that the presence of moral hazard does indeed lead to the provision of nonoptimal levels of services in agency relationships. Except for a related laboratory market study by Lynch et al. [7], who examine markets for goods whose characteristics were known to buyers immediately after purchase, we know of no other direct empirical evidence dealing with the effects of moral hazard. In our study, the "lemons" phenomenon only occurs about 50% of the time, and there is evidence which suggests that this is due to agents' concern for the reputations. However, agents do not establish the reputations required to eliminate the adverse effect of moral hazard. In fact, in three of the four markets studied, agents were able to use opportunities for false advertising to their advantage and "ripoff" principals.

Given the evidence provided in this study, we may now begin to examine some of the remedies proposed for dealing with the effects of moral hazard. In DeJong et al. [3], we have examined the impact of a negligence liability rule on these environments where agents could, for a cost, investigate and perfectly determine an agent's actions if a loss occurred. If the agent delivered a "lemon," the agent was held liable for the loss. We found that this liability rule almost completely eliminated the adverse effects of moral hazard as long as the outcomes from all investigations were publicly announced. In a future study, we plan to construct a set of laboratory markets to examine the effect of risk-sharing incentive contracts on the moral hazard problem.

REFERENCES

1. K. J. Arrow. *Essays in the Theory of Risk-Bearing.* Markham, Chicago, 1971.
2. R. Cooper and T. W. Ross. "Prices, Product Qualities and Asymmetric Information: The Competitive Case." *Review of Economic Studies* 51 (April 1984), 197–207.
3. D. V. DeJong, R. Forsythe, R. J. Lundholm, and W. C. Uecker. "A Laboratory Investigation of the Moral Hazard Problem in Agency Relationships." Institute of Accounting Research Working Paper 8407, University of Iowa, December 1984.
4. D. V. DeJong, R. Forsythe, and W. C. Uecker. "The Methodology of Laboratory Markets and Its Implications for Agency Research in Accounting and Auditing." *Journal of Accounting Research* (Autumn 1985), forthcoming.
5. B. Klein and K. B. Leffler. "The Role of Market Forces in Assuring Contractual Performance." *Journal of Political Economy* 89 (August 1981), 615–41.
6. D. M. Kreps, P. Milgrom, J. Roberts, and R. Wilson. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory* 17 (August 1982), 245–52.
7. M. Lynch, R. M. Miller, C. R. Plott, and R. Porter. "Product Quality, Information Efficiency and Regulations in Experimental Markets." Unpublished manuscript, 1984.
8. R. Selten. "The Chain-Store Paradox." *Theory and Decision* 9 (April 1978), 127–59.
9. S. Shavell. "Risk Sharing and Incentives in the Principal and Agent Relationship." *Bell Journal of Economics* 10 (Spring 1979), 55–73.

DISCUSSION

HAIM MENDELSON*: This paper studies a laboratory market where an agent is paid by a principal to provide a service of unknown quality. It is useful to place this model in the context of studies on markets with quality variation: the quality variation may be either exogenous or *endogenous*; the setting may be single-period, *finite-multiperiod*, or infinite-period; quality may be set once and for all (in the first period) or *adjusted dynamically*; and the market structure may range between monopolistic and *competitive* [the *italics* identify the specific choices made by DeJong, Forsythe, and Lundholm (DFL)]. Clearly, the above suggests a number of alternative settings for future research and for interesting comparisons; for example, if quality is set once-and-for-all, we obtain a clear rationale for reputation formation—how will this effect be reflected in laboratory results?

* Graduate School of Management, University of Rochester.