

Lecture Note: Monitoring, Measurement and Risk

David H. Autor
MIT 14.661, Fall 2003

November 13, 2003

1 INTRODUCTION

So far, we have toyed with issues of contracting in our discussions of training (both general and specific), holdup and shirking. The goal of this lecture is to make this set of issues more precise. We'll take three cuts at the contracting problem

1. The canonical Principal-Agent problem, which formalizes the trade-off between risk and incentives (review). This model is about uncertainty of outcomes with no uncertainty about appropriate actions.
2. The Principal-Agent problem with uncertainty about appropriate actions and asymmetric information (Baker, 1992).
3. The Principal-Agent problem with uncertainty about appropriate actions and asymmetric information: the choice of input- versus output-based contracts (Prendergast, 2002)

2 CANONICAL PRINCIPAL-AGENT PROBLEM

- An optimal compensation scheme must accomplish two things:
 1. Induce a given worker to put forth the appropriate level of effort (incentive compatible)
 2. Induce the 'right workers' to accept the contract (selection, individual rationality)\
- Under risk neutrality, the optimal contract is linear.

$$w = \alpha + \beta q,$$

where w is the wage, q is output, and α, β are chosen contracting parameters.

- Output depends on effort and luck. Normalize effort e so that one unit of effort produces one unit of output in expectation,

$$q = e + \nu,$$

where ν is luck or measurement error.

- Workers cost of effort is $c(e)$, with $c'(\cdot), c''(\cdot) > 0$.
- Worker's labor supply function is

$$\max_e E[\alpha + \beta(e + \nu) - c(e)],$$

with FOC

$$c'(e) = \beta. \tag{1}$$

- Firm's problem is therefore

$$\max_{\alpha, \beta} E(q) - (\alpha + \beta e) = \max_{\alpha, \beta} e - (\alpha + \beta e)$$

subject to Individual Rationality (IR) constraint

$$\alpha + \beta e > c(e).$$

- Substituting IR into firm's maximization problem (assuming holds with equality) gives

$$\max_{\alpha, \beta} e - c(e),$$

with FOC

$$\frac{\partial}{\partial \beta} = [1 - c'(e)] \frac{\partial e}{\partial \beta} = 0 \tag{2}$$

(and $\partial e / \partial \alpha = 0$, so the 2nd condition redundant).

- Combining (1) with (2) implies that $\beta^* = 1$. The firm pays the worker the entire residual profit. Given this, the optimal α – which is generally negative – is chosen to satisfy the IR constraint.
- Hence, firm rents the worker the job at price α , and pays piece rate 1.
- However, if the worker is risk averse, then this result must be modified: There is now a trade-off between incentives and risk:
 - A flat rate contract where $w = E[\alpha + \beta(e + \nu) - c(e)]$ provides the best insurance and the worst incentives

- A pure piece rate contract with $w = \alpha + q$ provides the best incentives and worst insurance
- For any concave utility function, we'll therefore have $0 \leq \beta_{RA} < \beta_{RN} = 1$, where RA and RN stand for Risk Averse and Risk Neutral respectively, where β_{RA} is declining in the agent's risk aversion (and α will therefore be increasing).
- The only “agency problem” in this setup is that a firm cannot simultaneously provide efficient incentives and efficient insurance to a risk averse agent.
- The important comparative static in this model is that, holding agent risk aversion constant, the ‘riskier’ is the environment (greater σ_ν^2), the lower is β_{RA} . Although risk has no effect on the optimal β for a risk neutral party, higher risk makes the contract less attractive to a risk averse agent, and the firm may have to compensate the agent heavily (through α) to satisfy the IR.
- Notice that you can think of this trade-off in terms of input versus output monitoring.
 - The extreme case of input monitoring is when you pay the worker if he sits at his desk, regardless of what he produces.
 - The extreme case of output monitoring is that you pay him q , regardless of what he does.
- The PA model implies that output monitoring is the ideal case for a risk neutral agent.

3 ADDING IMPERFECT OUTPUT MONITORING TO THIS MODEL

- It in its canonical form, the PA model takes the problem of incentives as purely a problem of mitigating the impacts of random noise. This suggests that the problem might be relatively solvable. For example, if the principal knew the agent's cost function $c(e)$ and could monitor effort perfectly, she could write an optimal ‘input contract’ based on the agent supplying e^* that solves $c(e^*) = 1$. So, in this simple form, perhaps it's not such a deep problem.

- In fact, the only distortion in the equilibrium of the PA model is that incentives might be insufficiently ‘strong’ for first best output given that agents are risk averse.
- Could a contract be ‘too strong,’ in that it generates too much output? In the PA model, the answer is no so long as $\beta \leq 1$, which will always be true (and this is a strict inequality of the agent is risk averse).
- But intuition – and many colorful examples – suggest that the answer must be yes. See for example the figure from Oyer, 1999, QJE on quarterly sales at competing Tandem and Stratus computer companies. Even though their fiscal years end in different quarters, each has large sales jumps at end of fiscal year. Clearly, this is not business cyclicity. Most likely explanation: gaming by managers.
- This idea was first formally articulated in a management journal article by Steven Kerr, 1975, “On the Folly of Rewarding A, While Hoping for B.” At the time of its publication, economists did not have a language for Kerr’s idea. Baker’s 1992 *JPE* formalizes Kerr’s insight.

3.1 SKETCH:

- There are numerous tasks that the agent could engage in at any time.
- These tasks, in aggregate affect the thing that the principle wishes to maximize $V(e, \epsilon)$.
- Which activities will have the highest marginal product is unknown ex ante (a random variable, function of ϵ).
- But these values are observed privately by the agent at the time of production.
- This would not create an incentive problem *except* that a contract cannot be written for $V(e, \epsilon)$.
- Instead, a contract must be written for $P(e, \epsilon)$, which is observable/verifiable.

- [Why? We can only reward what we can measure. We cannot measure ‘maximizing shareholder value.’ We can measure ‘burgers flipper per hour.’ So, that’s what you will paid for.]
- This disjuncture between $V(e, \epsilon)$, $P(e, \epsilon)$ gives rise to a set of contracts that can be either too weak or too strong, even with risk neutral agents.
- [Note that if we assumed there was no private information, this would be equivalent to assuming that the principal could contract directly over $V(e, \epsilon)$, since would be able to specify optimal efforts given the opportunity set.]

3.2 SETUP IN FULL

- P’s objective f’n $V(e, \epsilon)$ is not contractible.
- It is a f’n of e , agent’s actions, and ϵ , a set of random variables that completely characterize the state of the world (*‘SOW’*).
- Assume an arbitrary contractible performance measure $P(e, \epsilon)$, which is also a f’n of e and *SOW*.
- Assume risk neutrality – complicates without added insight. Key point is that we may get $\beta < 1$ despite risk neutrality.
- Consider only linear incentive contracts:

$$w = \alpha + \beta P(e, \epsilon) \tag{3}$$

- Neither P nor A knows ϵ before signing contract.
- Information: Agent is asymmetrically well informed about ϵ when actions are to be performed, and this affects her optimal action choice. (Seems sensible: Agent directly observes opportunity set at time actions must be taken; Principal is on a lounge chair gazing out at the Caribbean sea.)
- This setup implies that effort is a random variable, a function of ϵ .

- Non-triviality: At least some elements of ϵ affect marginal product of effort on both $V(\cdot), P(\cdot)$.
- Contracts are binding – cannot be renege (otherwise IR constraint also random)
- The standard deviation of V_e (the marginal product of effort) with respect to ϵ , defined as σ_{ve} is a measure of the amount of valuable information that the agent possesses. When σ_{ve} low, variations in opportunity set observed by agent are not particularly important for the principal’s objectives. And v.v. when σ_{ve} high. Be clear: marginal product of effort could be high or low in either cases, but if σ_{ve} is low, there is little reason to vary effort.
- Similarly σ_{pe} is a measure of valuable info A has about performance measure $P(\cdot)$.
- Consider e as a unidimensional variable, although could be written as vector.
- Time line:
 1. Agent signs binding contact specifying $P(\cdot), \alpha, \beta$.
 2. ϵ is unknown, but its distribution is known
 3. After signing, before choice of e , all elements of ϵ revealed to agent.

3.3 MAXIMIZATION

- Normalize $P(\cdot)$ so that the average dollar value of an incremental unit of performance measure is 1. Hence, optimal piece rate is 1 under a first-best contract:

$$E [P_e(e, \epsilon)] = E [V_e(e, \epsilon)]. \tag{4}$$

The expected marginal product of effort on performance measure equal to expected marginal product of effort on value.

- Standard assumptions on effort cost: $c' > 0, c'' > 0$.

$$w = \alpha + \beta P(e, \epsilon) - c(e). \tag{5}$$

- IR constraint with outside opportunity set H :

$$E[\alpha + \beta P(e, \epsilon) - c(e)] \geq H, \quad (6)$$

- Effort choice

$$\beta P_e(e^*, \epsilon) = c'(e^*). \quad (7)$$

- Principal's maximization

$$\max_{\alpha, \beta} E[V(e^*, \epsilon) - \alpha - \beta P(e, \epsilon)], \quad (8)$$

subject to (6) and (7).

- Substituting the IR constraint into the maximization gives

$$\max_{\beta} E[V(e^*, \epsilon) - H + \beta P(e, \epsilon) - c(e) - \beta P(e, \epsilon)],$$

and differentiating WRT β yields

$$E[V_e e_{\beta}^*] - E[c'(e) e_{\beta}^*] = 0,$$

where $e_{\beta}^* = \partial e^* / \partial \beta$.

- Substituting in equation (7) for $c'(e^*)$ and rearranging yields β^* ,

$$\beta^* = \frac{E[V_e e_{\beta}^*]}{E[P_e e_{\beta}^*]}. \quad (9)$$

Notice that if there were no random variables involved – that is optimal effort is not a function of ϵ – then $\beta^* = 1$, and there is no contracting problem.

- To get an expression for e_{β}^* , differentiate (7):

$$\begin{aligned} \frac{\partial}{\partial \beta} \beta P_e(e^*, \epsilon) - c'(e^*) &= P_e + \beta P_{ee} e_{\beta}^* - c'' e_{\beta}^* \\ e_{\beta}^* &= \frac{P_e}{c'' - \beta P_{ee}}. \end{aligned} \quad (10)$$

- Now, substitute back into (9) to get

$$\beta^* = \frac{E \left[V_e \left(\frac{P_e}{c'' - \beta P_{ee}} \right) \right]}{E \left[P_e \left(\frac{P_e}{c'' - \beta P_{ee}} \right) \right]}.$$

Assuming $(c'' - \beta P_{ee})$ is constant ('2nd order Taylor approximation') reduces this to

$$\beta^* = \frac{E [V_e P_e]}{E [P_e^2]}. \quad (11)$$

- Assume WLOG that $E [P_e] = E [V_e] = 1$ at e^* and recalling that $E [A \cdot B] = E [A] \cdot E [B] + \text{cov}(A, B)$, we can write

$$\beta^* = \frac{1 + \text{cov}(V_e, P_e)}{1 + \text{var}(P_e)}. \quad (12)$$

The correlation coefficient between two variables a, b is defined as

$$\rho_{ab} = \frac{\sigma_{ab}}{\sigma_a \sigma_b} \quad (13)$$

Using this, we can rewrite as (12)

$$\beta^* = \left(\frac{\sigma_{v_e p_e}}{\sigma_{v_e} \sigma_{p_e}} \right) \frac{\sigma_{v_e} \sigma_{p_e} + 1}{1 + \sigma_{p_e}^2} = \frac{\rho \sigma_{v_e} \sigma_{p_e} + 1}{1 + \sigma_{p_e}^2}. \quad (14)$$

- This the coefficient from a bivariate regression without an intercept...

- Key observations

1. If $\sigma_{v_e} = \sigma_{p_e}$ and $\rho = 1$, then $\beta^* = 1$. That is, if V_e, P_e are perfectly correlated with identical variances, then we are back in the first-best world. (Measures don't need to be identical for this to be true, just proportional.)
2. Correlation of V_e, P_e is a key determinant of optimal piece rate. If the marginal product of effort on a performance measure is strongly correlated with the marginal product of effort on value, then the agent, who chooses effort level based on the value of P_e , will choose high levels of effort when V_e is high and low levels when V_e low.
3. The piece rate β serves two functions:
 - One is to get the agent to exert effort on average

- Two is to get the agent to use her superior information in choosing her effort level.
4. This means that the piece rate can still be positive even when V_e, P_e are negatively correlated (so long as the marginal product of effort is positive on average).
 Example, from Baker: 2 states of world, equally likely. In one state, $P_e = 5, V_e = 10$ and in the other $P_e = 10, V_e = 5$. So, the measures are perfectly negative correlated. Would you want $\beta < 0$ in this case? No, because you still want positive effort on average. Marginal products are 7.5 on average, and optimal piece rate is 0.8.
- Comparison with standard PA model:
 - If A risk neutral, there is no conflict of interest between P and A; first best achievable.
 - Here, a conflict of interest arises – even with risk neutrality – because the piece rate does not perfectly align A’s incentives with P’s objectives (unless the correlation is one).
 - So, even risk neutral agent facing a β of 1 could be induced to work *too hard* if the marginal value of effort for P is lower than the marginal value of effort for V. This would be inefficiently costly for the agent and principal. (Note that if $c'' = 0$, this would not be true. Convexity of effort function means that marginal cost of effort eventually rises faster than marginal product.)
 - Trade-off: the Principal will reduce effort variation by reducing piece rate – reduces incentives for effort but this is efficient.

3.4 OBSERVABLE EFFORT CASE

- When effort is observable, Principal could condition the agent’s payoff on her effort choice.
- In the standard PA model, this would be first best efficient. If the Principal exerted first-best effort and this could be verified, pay could be independent of actual output – could offer full insurance and full incentives. So, an input contract with monitoring solves incentive and participation constraint perfectly.

- In the Baker model, that result will not hold. Because of the information asymmetry, the Principal still will not know the optimal level of effort because does not observe realization of ϵ . Can still reward e on average if observable, however.
- Consider this contract

$$w = \alpha + \beta_1 E[P|e] + \beta_2 (P - E[P|e]). \quad (15)$$

This contract pays two separate piece rates: one for expected output conditional on effort, the other for the realization of output. Each piece rate serves a purpose.

1. β_1 causes the principal to exert the right level of effort on average
2. β_2 provides incentives to adjust effort level according to superior information about marginal productivity (due to observation of ϵ).

- Optimal solutions for these parameters is

$$\beta_1^* = \frac{E[V_e]}{E[P_e]} = 1,$$

$$\beta_2^* = \rho \frac{\sigma_{v_e}}{\sigma_{p_e}} = \frac{\sigma_{v_e p_e}}{\sigma_{v_e} \sigma_{p_e}} \cdot \frac{\sigma_{v_e}}{\sigma_{p_e}} = \frac{\sigma_{v_e p_e}}{\sigma_{v_e}^2},$$

which are the intercept and slope coefficients from a regression of V on P and a constant.

- So, observing effort allows the principal to improve the contract. (This contract would produce shirking if effort was not observable.)
- How do we know that this contract is more efficient than (14)?
 - In the prior example, we could never have $\beta^* < 0$, if the marginal product of effort was positive on average. That's even true if V_e and P_e are perfectly negatively correlated.
 - In this example, β_1^* provides incentives for average effort and β_2^* provides incentives for marginal effort. So, we could easily have a case where $\beta_1^* > 0$ and $\beta_2^* < 0$. And if we did, this would dominate the single instrument case where one β must trade off between solving these two problems simultaneously.
- This is an insightful and important paper (and elegant).

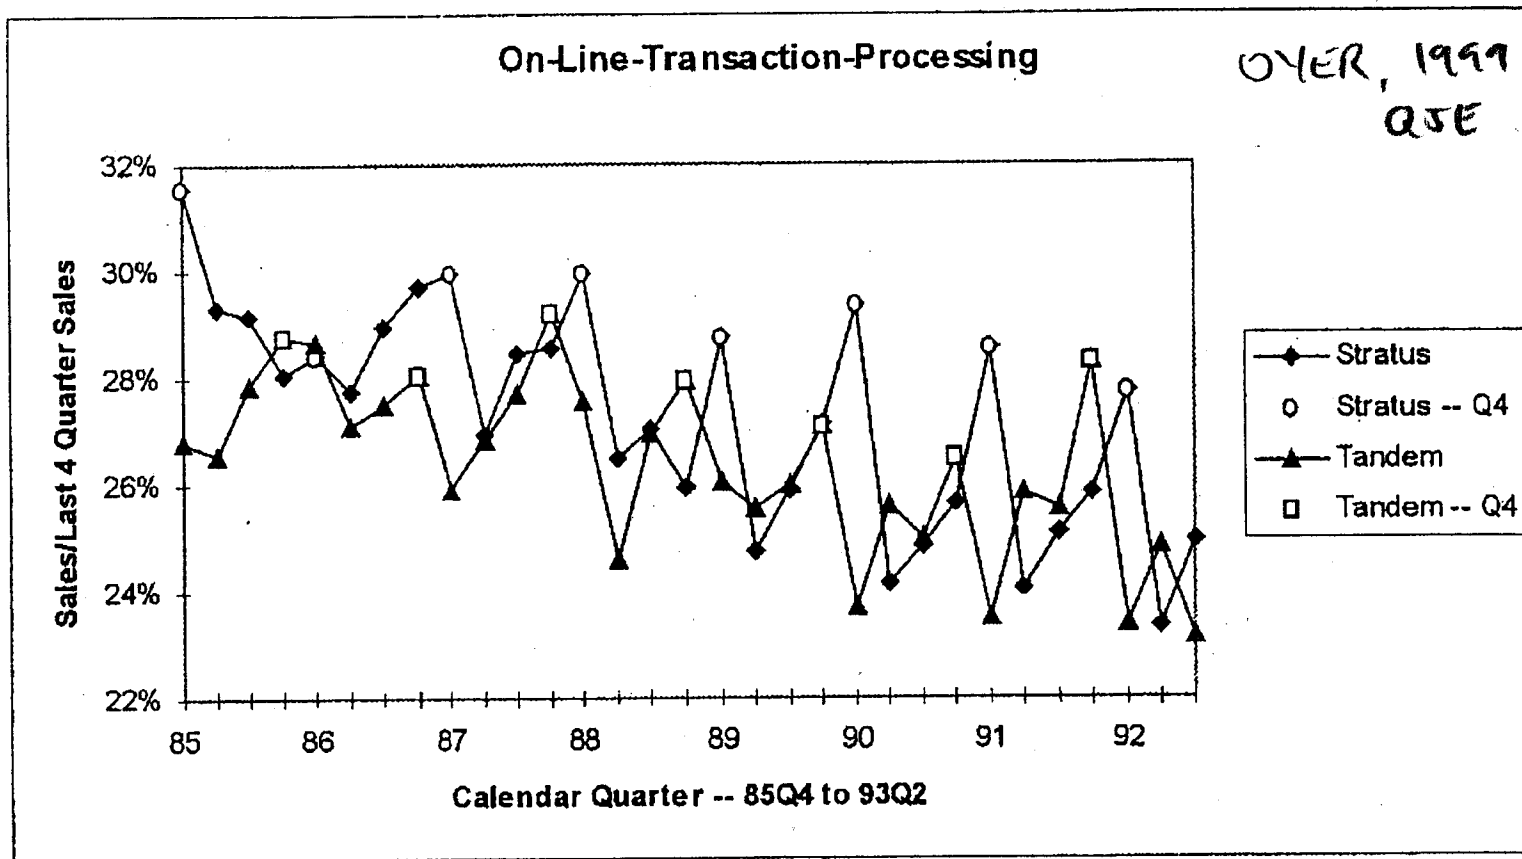
4 PRENDERGAST 2002 JPE

- I will not develop the details of this paper. But it provides another insight into why the PA model is too parsimonious (and also an empirical failure).
- The contrast here is again between input and output monitoring. Recall from the PA model that input monitoring is preferred where there is high risk to the risk-averse agent and output monitoring when there is low risk.
- In the Prendergast model, there is again a large dimensional action space and again private information observed by the Agent about marginal products. Another key idea: input monitoring is cheaper than output monitoring. Why would this be true? Seems that the answer will vary with context.
- Under these assumptions, input monitoring will be preferred by Principal when she has a clear idea of the right actions to take. This reduces monitoring costs at not a great loss in efficiency.
- Problem: precisely in the ‘riskiest’ settings where Principal does not have good information about what actions to take (e.g., sending a sales representative to a foreign country to develop a new line of business versus siting a McDonalds on the interstate).
- In these uncertain cases, the principal may choose output monitoring so as not to constrain agents’ superior information about actions. But this will shift risk to agents when risk is high – opposite of PA model.
- The details of this model are not especially elegant but the point seems good. If you need references on “order statistics,” please ask me.

5 CONCLUSION

As this lecture may underscore, the contracting issues we visited in the lectures on training investment and efficiency wage incentives are fairly rudimentary relative to the depth of the subject. As a labor economist, you should understand the fundamentals of incentive and

mechanism design, at least at the introductory level here. If you want further depth, I highly recommend the Organizational Economics and Contract Theory classes taught by Bob Gibbons.



Fiscal Year Ends: Tandem - September; Stratus - December