Dynamic Principal-Agent Theory

Baomin Dong[*]

Abstract. Studies on dynamic optimal incentive contracting are reviewed. Attention is directed into a few topics including: renegotiation proofness, convergence to the first-best, linearity, career concerns, team issues, ratchet effect and comparable performance information.

---

[*] This is a survey about recent literature on dynamic principal-agent theory.

Dynamic principal-agent theory, in most cases, is also viewed as an intertemporal moral hazard problem. There are two important distinct features compared with static moral hazard problems.

First, the incentive wage that the agent receives at period t, like in the static case, depends on his effort and on a shock that he does not control; it is therefore a stochastic income stream for him. As any ordinary consumer who has a concave utility function and receives a random income stream, he will want to smooth his consumption, if he can, by saving and borrowing . The study of intertemporal moral hazard therefore should not abstract from the conditions under which the agent can access credit markets.

Secondly, the repetition of a moral hazard problem can create endogenously generated information, mostly for the agent. This will be the case if the technology or the preferences of the agent in any given period depend on his actions in earlier periods. The dynamic moral hazard problem then is complicated by an intertemporal adverse selection problem too.

A question of both theoretical interest and practical importance is under what circumstances the principal and agent can then do better by committing themselves in advance to a long-term contract rather than by negotiating short-term contracts at the start of each period.

1 A basic static model

Several of the main issues can be illustrated quite simply in the context of sharecropping. Three standard sharecropping contracts are: wage labor, which imposes no risk on the agent; crop sharing, which shares risk between the principal and the agent; and fixed-payment land rental, which leaves the agent with all the crop risk. The classic agency model, which emphasizes the tradeoff between incentives and insurance, implies that where there is greater crop risk there should also be more risk sharing - more used of fixed wages and crop sharing rather than land rental. Even if the principal is risk averse and the agent is risk neutral, at optimal, the principal should still undertake some risk in order to provide more incentives to the agent. Evidence consistent with this prediction can be easily found in many developing countries. This intuition about tradeoff between incentives and insurance effect is demonstrated in a simple model below.

The classic model in agency theory involves an agent who takes an action a to produce output of value x. The principal owns the output but contracts to share it with the agent by paying a wage contingent on output, $w(x)$. The production function is linear, $x = a + \varepsilon$, where ε is a normally distributed noise term with zero mean and variance σ^2. The incentive contract is linear, $w(x) = s + bx$, where the intercept s is the fixed payment and the slope b is the piece rate. The agent's utility function is $U(x) = -e^{-rx}$, where $r > 0$ is the agent's coefficient of absolute risk aversion and $w_n = w - c(a)$ is the agent's net payoff: the realized wage minus the convex disutility of action $c(a)$. The principal is risk neutral and so he maximizes the expected value of profit, $x - w$.

Given a contract $w(x) = s + bx$, the agent's problem is to choose an action to maximize the expected utility

$$\int_{\varepsilon} -e^{-r[s+b(a+\varepsilon)-c(a)]}\phi(\varepsilon)d\varepsilon = -e^{-r[s+ba-c(a)]}\int_{\varepsilon} -e^{-rb\varepsilon}\phi(\varepsilon)d\varepsilon;$$

where $\phi(\varepsilon)$ is the normal density function. The agent's optimal action, denoted $a^*(b)$, solves $c'(a) = b$. The agent's maximized expected utility is therefore

$$-e^{-r[s+ba^*(b)-c(a^*(b))]}\int_{\varepsilon} -e^{-rb\varepsilon}\phi(\varepsilon)d\varepsilon = -e^{-r\{s+ba^*(b)-c(a^*(b))-(\frac{1}{2}rb^2\sigma^2)\}};$$

the agent's certainty equivalent is

$$CE(s;b) = s + ba^*(b) - c(a^*(b)) - (\frac{1}{2}rb^2\sigma^2);$$

That is, the agent's certainty equivalent from the contract $w(y) = s + bx$ is the expected wage minus the cost of effort minus the cost of bearing risk.

The principal's expected profit is

$$E\pi(s;b) = (1-b)a^*(b) - s;$$

so the total social surplus depends on b but not on s:

Then the efficient contract piece rate, denoted b^*, is the slope that maximizes the total surplus $TS(b)$. If the parties agreed to a contract with some other slope then both parties could be made better off by switching to a contract with slope b^* and choosing an appropriate value of s to distribute the increased total surplus. The first-order condition for b^* is $a^{*\prime} - c'a^{*\prime} - rb\sigma^2 = 0$. Because $c'[a^*(b)] = b$, we have $a^{*\prime} = 1/c''$ and hence

$$b^* = \frac{1}{1 + r\sigma^2 c''};$$

b^* is smaller if the agent is more risk averse (r is higher) or there is more uncertainty in production (a' is higher) or marginal disutility increases more quickly (c'' is higher).

From this static model, we can see that an incentive contract between the principal and the agent is a tradeoff between providing incentive to the agent and insuring him. In other words, if the agent bears too much risk relative to that of the optimal contract, he would work less hard to counter balance the disutility from extra risk-taking; however if the agent is insured too much, he will lose the incentive to work at optimal also because the extra income from a good outcome is not enough to cover his effort exerted to obtain a possible good outcome.

2 Renegotiation after Effort

Fudenberg and Tirole (1990) show that even one-period moral hazard problems can have dynamic aspects. They start from a simple agency problem and consider the renegotiation proof optimal contract. Assume the technology is again $x = a + \varepsilon$; where ε is an observational noise with mean zero. At the optimum the Principal announces a wage schedule $w^*(x)$, the agent makes an effort a^* and expects to get a random wage given by $w^*(a^* + \varepsilon)$, while the principal gets a surplus $a^* + \varepsilon - w^*(a^* + \varepsilon)$.

The function w^* results from a trade-off between incentives and risk-sharing. Now consider a point in time when the agent has made effort a^* but the outcome x has not been observed yet. The function w^* then has played its part in providing incentives to the agent, and only risk-sharing matters now. But if we assume as usual that the principal is risk-neutral and the agent is risk-averse, the risk-sharing properties of the function w^* cannot be optimal; the optimum would be a fixed wage so that the principal insures the agent perfectly against the risk represented by the shock. This argument shows that once the agent has made effort a^*, the parties would gain by renegotiating toward a perfect insurance contract that gives all risk to the principal. The optimal contract therefore is not robust to renegotiation. Obviously, if the agent anticipates that his wage schedule will be renegotiated to a constant wage after he chooses his effort, he will choose the least costly action. The contract w^* does not play any incentive role any more.

So when renegotiation is permitted, the principal then will have to consider how to prevent renegotiation or how to remove the cost caused by renegotiation or potential renegotiation. Consider if the time interval between the choice of action and observation of the outcome is long, this problem would be serious, the principal and the agent are trapped into a bilateral monopoly situation.

If this type of renegotiation is possible, then it must be taken into account when the principal is designing the optimal contract. Consider first the simplest case where the action can only take two values, $a = 0$ and $a = 1$. At the renegotiation date the principal faces two possible types of agent: one who chose $a = 1$, and one who chose $a = 0$; he must therefore solve an adverse selection problem. Like in the insurance model, he will offer two different wage schedules. It can be shown, that the wage schedule designed for the low effort agent insures him perfectly.

3 Infinitely Repeated Moral Hazard

When the game is repeated infinitely, the classical problem of statistical inference comes to mind: The principal tries to infer the action a from observing the outcome x. It therefore seems likely that a law of large numbers should apply: If the interaction between the principal and the agent is repeated infinitely, the principal will observe a large number of outcomes; he will be able to infer the action with great precision and to punish the agent very strongly if the latter does not choose the optimal action. In the limit the principal should be able to implement the first-best optimal action. Rubinstein and Yaari (1983) show that this intuition is correct when neither the principal nor the agent has a preference for the present. Assume the technology is again linear

$$x_t = a + \varepsilon_t$$

within each period t, where the ε_t are i.i.d. noises with mean zero and a finite variance σ^2. Let a^* be the first-best optimal action. If the agent chooses a^* in each period, then by the law of large numbers, the average

$$\frac{1}{t}\sum_{\tau=1}^{t}(x_\tau - a^*)$$

will go to zero almost surely as t goes to infinity. To induce the agent to choose action a^* in each period, the principal can then punish the agent if the absolute value of this average is greater than some positive threshold, indicating that the agent has deviated relatively often. The difficult point is how to choose this threshold: it should go to zero as t goes to infinity in order to take advantage of the law of large numbers, but it should not vanish too fast; otherwise the agent would be punished too often, which would not be good for risk-sharing.

The appropriate tool for this problem is the law of the iterated logarithm, which bounds the large deviations from the law of large numbers. Let λ be any real number greater than 1, and let

$$\delta_t = \frac{\left(\sum_{\tau=1}^{t} \varepsilon_\tau\right)/t}{\sqrt{2\lambda\sigma^2 \ln\ln t/t}};$$

Then the law of the iterated logarithm states that

$$\Pr(\lim_{t\to\infty} \sup \delta_t < 1) = 1$$

The policy consisting in choosing $\lambda > 1$ and punishing the agent at date t if

$$\left|\frac{1}{t}\sum_{\tau=1}^{t}(x_\tau - a^*)\right| > \sqrt{\frac{2\lambda\sigma^2 \ln\ln t}{t}};$$

This implements the first-best action if the punishment is harsh enough and the interaction is repeated infinitely. Note that if the agent does choose a^* in each period, then he will be punished with vanishing probability. The trouble with this result is that it rests on two crucial assumptions: the interaction is infinitely repeated and both agents are extremely patient.

If the interaction between the principal and the agent is repeated over a finite horizon. Then the argument developed above fails completely and the optimum is clearly second-best. We show it next.

4 Finitely Repeated Moral Hazard

Assume that the interaction between the principal and the agent lasts for T periods. The principal's utility function is

$$\sum_{t=1}^{T} \delta^{t-1}(x_t - w_t)$$

while that of the agent is

$$\sum_{t=1}^{T} \delta^{t-1}(u(c_t) - a_t)$$

where c_t is the consumption of the agent at t and δ is the common discount parameter. Note wage and consumption are not the same if the agent has access to a credit market.

Assume that the outcome in period t only depends on the action chosen in the same period. If, for instance, the outcome in period t also depends on the action at a_{t-1}, which is only observed by the agent, then the latter would have an informational advantage on the Principal at the beginning of period t, since he would have a finer knowledge of the period-t technology.

Chiappori-Macho-Rey-Salanié (1994) show that two conditions are necessary for the optimal long-term contract to be implementable via spot contract: first, the long-term optimum must be renegotiation-proof; second, spot contracts should provide efficient consumption smoothing. Availability of credit to the agents affects these two prerequisites.

4.1 No Access to Credit

If the agent has no access to the credit market, would the principal artificially smooth up the agent's income in order to improve the tradeoff between incentives and insurance? Malcomson and Spinnewyn(1988) show that if efficient contracting does not require long-term commitment from the principal, or the agent, then making contracts contingent on previous periods does not improve the tradeoff between incentives and risk sharing.

For reasons discussed by Sappington(1983), however, courts may not enforce very high penalties on an agent. In the employment context, for example, the penalties an employer can impose on an employee are typically limited whatever the circumstances.

To overcome this, the principal may use a long-term contract to penalize the agent for a bad outcome in one period using low payoffs over a number of subsequent periods. So may commitment by the principal that allows withholding of payoffs for good outcomes as a hostage to ensure good performance in the future, that is, bonding. But limits on enforceability may also permit the agent to break a long-term contract on payment of some penalty. Therefore, if that penalty implies a payoff to the agent no lower than the lowest enforceable payoff in a one-period contract, then limitations on the payoffs do not change the above results.

Even without moral hazard, of course, optimal consumption paths under uncertainty exhibit memory in the sense that an economic agent with a strictly concave, time-separable utility function will smooth consumption over time. This raises the question of whether the memory result of Rogerson(1985) arises only because of the consumption smoothing that would occur even without moral hazard or because it also improves the efficiency of incentives in the presence of moral hazard.

By not requiring the agent to consume all compensation right away, Malcomson and Spinnewyn(1988) show that, whenever short-term contracts are as efficient as long-term contracts, this memory result does indeed arise simply from the efficiency of smoothing consumption and not because it permits an improvement in the tradeoff between risk-sharing and incentives. An optimal short-term contract requires no more information about the outcomes of previous periods than is provided by the agent's wealth. Moreover, it is independent of whether the agent acquired that wealth from a previous contract or from some other source, for example, inheritance. Thus the memory result arises from a pure wealth effect.

Rey and Salanie (1987) have also showed the equivalence of short-term and long-term contracts. However in their simple two period model, the agent still want to save if he could, given the second period wage depends on first period's outcome. The key is the agent's risk aversion and they show by applying Jensen's inequality. The model also appears in a discussion paper of Chiappori-Macho-Rey-Salanié(1994). In their model, it is assumed that the agent cannot save or borrow, so his consumption equals his wage within any period. An immediate application of the dynamic programming principle shows that full commitment coincides with long-term commitment in this model: The agent's characteristics are fully known to the principal when the contract is signed, and the principal therefore can choose the optimal sequence of wage schedules without ever feeling the need to adapt it to the arrival of new information. Since the agent has a concave utility function, he will want to have a smooth consumption over time. If the outcome x_t is particularly high because a favorable shock took place in period t, the principal therefore will have to spread this positive shock over several periods so as to smooth the agent's consumption stream. He will do this by increasing the wage he gives to the agent in all future periods. Thus the wage given in any period t will have to depend not only on the current outcome x_t but also on the sequence of past outcomes. This property, which Rogerson (1985) called the memory effect[1], therefore is a simple consequence of the need for the principal to smooth the agent's consumption at the full commitment optimum. In the absence of commitment the principal cannot spread the effect of a shock on x_t over several periods. Therefore the period-t wage can only depend on the current outcome x_t, and the optimal sequence of spot contracts is memoryless. This clearly involves an important efficiency loss.

Assume that $T = 2$, and denote w_i the first-period wage associated with first-period outcome x_i, and second-period wage w_j. A contract is called memoryless if the second period wage is independent of the first-period outcome:

$$\forall (i, i'; j) \in I, \quad w_{ij} = w_{i'j}.$$

The agent chooses $(1 + n)$ contingent actions; let a_0 be the action chosen in the first period and a_i be the action chosen in the second period when the first period outcome is x_i. Then the contract is characterized by the following program:

$$\max_{(w_j),(w_{ij}),a_0,(a_i)} \sum_{i=1} p_i(a_0)[h(x_i - w_i) + \sum_{j=1} p_j(a_i) h(x_j - w_{ij})],$$

which is equivalent to minimizing the wage bill for all i; i.e.,

$$\min_{(w_j),(w_{ij}),a_0,(a_i)} \sum_{i=1} p_i(a_0)[w_i + \sum_{j=1} p_j(a_i) w_{ij}],$$

subject to

$$(a_0, (a_i)) \in \arg\max_{\tilde{a}_0,(\tilde{a}_i)} \sum_{i=1} p_i(\tilde{a}_0)[U(w_i, \tilde{a}_0) + \sum_{j=1} p_j(\tilde{a}_i) U(w_{ij}, \tilde{a}_i)],$$

$$\sum_{i=1} p_i(a_0)[U(w_i, a_0) + \sum_{j=1} p_j(a_i) U(w_{ij}, a_i)] \geq 2\underline{U},$$

---

[1] Lambert (1983) and Rogerson (1985) have considered the properties of Pareto efficient consumption paths in finitely repeated principal-agent models by analysing the compensation to the agent when all compensation is used for immediate consumption. Their central result is that "memory plays a role in every Pareto-optimal contract" (Rogerson) in the sense that "the agent's compensation in one period depends on his performance in that period and his performance in prior periods" (Lambert).

where p is the probability of outcomes as a function of action taken by the agent, h is the principal's utility, U is the agent's utility and $\underline{U}$ is the agent's reservation utility level. The first-order condition[2] for this problem is

$$\frac{1}{U'(w_i; a_0)} = \sum_{i=1}^{n} \frac{p_j(a_i)}{U'(w_{ij}; a_i)},$$

which reveals the memory effect. By Jensen's inequality[3]

$$\sum_{i=1} \frac{p_j(a_i)}{U'(w_{ij}; a_i)} \geq \frac{1}{\sum_{i=1} p_j(a_i) U'(w_{ij}; a_0)},$$

and hence

$$U'(w_i; a_i) \leq \sum_{i=1} p_j(a_i) U'(w_{ij}; a_0);$$

---

[2] It can be established that the solution is the same under the assumption that the agent has monitorable access to the credit market. The agent can save and borrow as he wishes. Assume that the principal can observe the agent's savings. The problem turns to be

$$\max_{(w_j);(w_{ij});a_0;(a_i)} \sum_{i=1} p_i(a_0)[h(x_i - w_i - t_i) + \sum_{j=1} p_j(a_i) h(x_j - w_{ij} + t_i)];$$

such that

$$(a_0; (a_i)) \in \arg\max_{\tilde{a}_0;(\tilde{a}_i)} \sum_{i=1} p_i(\tilde{a}_0)[U(w_i - s_i; \tilde{a}_0) + \sum_{j=1} p_j(\tilde{a}_i) U(w_{ij} + s_i; \tilde{a}_i)];$$

$$\sum_{i=1} p_i(a_0)[U(w_i - s_i; a_0) + \sum_{j=1} p_j(a_i) U(w_{ij} + s_i; a_i)] \geq 2\underline{U};$$

This situation where the agent's savings can be monitored is equivalent to a 'no access' situation where the principal has access to credit. Spot contracts can achieve long-run efficiency when the agent has access to a perfect credit market and the principal can monitor this access. Controlling the agent's borrowing and saving allows the principal to introduce memory in spot contracts, since the reservation utility u at the beginnning of second period depends on the agent's first-period savings.

The principal will implement the long-term optimum by offering the wage schedule $\tilde{w}_i = c_i^* + \tilde{s}_i$ and specifying savings $\tilde{s}_i$ in the first period and then offering the wage schedule $\tilde{w}_{ij} = c_{ij}^* - \tilde{s}_i$ in the second period.

[3] $E[f(x)] \geq f(E[x])$ if X is a random variable and $f(\cdot)$ is a convex function.

which implies that the agent would like to save if he can or save more than the contract allows him to[4].

Therefore the argument casts doubts on "constrained access" case: it would be hard to imagine how the principal could constrain the agent's savings. It might be the case with sharecropping in developing countries where the sharecropper produces a perishable good only, or in the case of labor contract where the reward is non monetary. However in general, free access to credit market in more interesting.

4.2 Free Access to Credit

Assume neither the agent's consumption nor his savings are observed by the principal at the end of the first period[5].

Let s_{T-1} denote the savings of the agent in period $(T-1)$ (which depend of course on the whole past history) and $r = 1/\delta - 1$ the market interest rate. Then the utility function of the agent in period T, expressed as a function of the wage he gets from the principal, is

$$U(w_T, s_{T-1}) = u(w_T + (1+r)s_{T-1}),$$

which depends on his past savings s_{T-1}.

Since the adverse selection problem emerges on top of the moral hazard problem[6], the full commitment optimum is not renegotiation proof. Chiappori-Macho-Rey-Salanié (1994) show the following striking result: if the long-term optimum only involves pure strategies, then it can only implement the cost-minimizing action from the second period onward. To see this, assume that $T = 2$ and use the same notation as above, with in addition s_j as the savings when the first-period outcome is j and a_0 the optimal action in the first period. Now assume that the optimal contract implements a_i in the second period after the first-period outcome was i. If a_i is not the cost-minimizing action[7], then at least one second-period incentive constraint

[4] If the agent saves s, then the agent's utility is $U(w_i - s) + \sum_j p_j(a_i)U(w_{ij} + s)$.

[5] In the second period, the agent's wealth is his private information. And indeed the non-observability of savings arises adverse selection.

[6] The exception is CARA utility which excludes the wealth effect and hence eliminates adverse selection. Fudenberg, Holmstrom and Milgrom(1990) discussed this case.

[7] It is proved by contradiction.

must be binding: There exists an a' such that

$$\sum_j p_j(a_i)u(w_{ij} + \frac{s_i}{\delta}) - a_i = \sum_j p_j(a')u(w_{ij} + \frac{s_i}{\delta}) - a';$$

Let s' be the optimal savings when the agent chooses a', i.e.,

$$s' = \arg\max_s u(w_i - s) + \delta(\sum_j p_j(a')u(w_{ij} + \frac{s_i}{\delta}) - a');$$

Now assume that instead of responding to the optimal contract $(w_i; w_{ij})$ with $(a_0; s_i; a_i)$, the agent responds with $(a_0'; s_i'; a_i')$ where $a_0' \neq a_0$ unless $s_i = s_i'$ and $a_i = a_i'$. This improves the agent's expected utility.

$$\begin{aligned} & \sum_j p_j(a_0)(u(w_j - s_j) - a_0 + \delta(\sum_k p_k(a_j)u(w_{jk} + \frac{s_j}{\delta}) - a_j)) \\ = & \sum_j p_j(a_0')(u(w_j - s_j) - a_0' + \delta(\sum_k p_k(a_j')u(w_{jk} + \frac{s_j}{\delta}) - a_j')) \\ < & \sum_j p_j(a_0')(u(w_j - s_j') - a_0' + \delta(\sum_k p_k(a_j')u(w_{jk} + \frac{s_j'}{\delta}) - a_j')); \end{aligned}$$

The inequality(generically by definition that s_j' is a better choice than s_j given the action taken was a_j') implies that the action the agent will take tomorrow does influence his saving today. Since this inequality violates the first-period incentive constraint, the premise that a_i was not the cost-minimizing action must be wrong and this completes the proof. As we can see from the proof, adverse selection does arise and is easier to manipulate when the agent has access to the credit market. Then two points emerged: first, ex post efficiency must require that the incentive compatibility be binding at the beginnning of the second period; second, the level of savings must vary according to the level of effort which the agent plans to choose. The intuition is that when compensation is also based on previous outcomes, the agent would save more today to enjoy more leisure and/or against possible negative shocks tomorrow.

5 Linearity in the dynamic optimal incentive contracts

In the basic model in section 1, Mirrlees (1974) argued that the best linear contract, $w = s + b^*x$, is inferior to various non-linear contracts. In particular, a step-function contract [8] can perform very well, approaching the twin goals of full incentives and full insurance in the limit. The optimal contract then is linear only under very special assumptions about the utility function and the conditional distribution of output.

Holmstrom and Milgrom (1987) reinterpreted the classic agency model so as to rescue linear contracts. In their model, rather than a single action that influences a single outcome, Holmstrom and Milgrom envision a sequence of actions influencing a corresponding sequence of outcomes. There are no connections across days[9] and all past outcomes are observed before the next day's action is chosen. The output x from the classic model is interpreted as the aggregate output for the year in the sequential-action model, $x = \sum_i x_i$.

Suppose that each day's outcome takes one of two values - say L or H. Then according to Mirrlees (1974), a one-day optimal incentive contract is simply a pair of wages, w_H is paid if the outcome is H; w_L,. if L. Suppose that the agent labors under the same one-day contract for all the days of the year. If there are T days in the year and the agent produces H on N of these days then the aggregate output for the year is $x = TL + N(H - L)$ and the aggregate wage for the year is $w = Tw_L + N(w_H - w_L)$. Thus, $N = (y - TL)/(H - L)$ and

$$w = \frac{T(Hw_L - Lw_H)}{H - L} + \frac{w_H - w_L}{H - L}x = s + bx,$$

Then it is established that under the same spot step-contract, over time, the aggregate wage is a linear function of the aggregate output which breaks the step-contract itself. Given several other assumptions, Holmstrorn and Milgrom(1987) show not only that it is optimal for the agent to work under a constant spot step-contract but also that the optimal piece rate in the aggregate representation of this contract[10] is b^*, just as in the static agency model.

Holmtrom and Milgrom warn us that the linear contract is quite robust in a dynamic context. Intuitively, a step-function contract of the kind studied by Mirrlees induces no effort once the agent's aggregate output to date passes the hurdle y_0 in a dynamic model. Generally, if the incentive contract for the year is a non-linear function of year-end aggregate output then the worker's

---

[8] the agent earns w_H, if $y > y_0$, but $w_L < w_H$, if $y < y_0$.

[9] So the outcomes are series uncorrelated, or in other words, independent of actions taken in other periods.

[10] $w = s + bx$.

incentives change from day to day, depending on the aggregate output to date which is hard to meet principal's purpose.

Holmstrom and Milgrom(1987) first consider a multi-period model in which the incentive contract repeats T times. The agent's minimum certainty equivalent is normalized to zero and assured by the principal. The problem is then

$$\max_{\{p^t\};s} E\{v(\sum_{t=1}^{T} \pi^t - s(X^T))\}; \quad \text{subject to}$$

$$\{p^t\} \in \arg\max E\{u(s(X^T) - \sum_{t=1}^{T} c(p^t(X^{T-1}); \theta^t))\}.$$

$$E\{u(s(X^T) - \sum_{t=1}^{T} c(p^t(X^{T-1}); \theta^t))\} \geq u(0)$$

where π is the total profit, x is the outcome, v is the principal's utility and p^t is the efforts privately chosen by the agent and affect the probablity distribution over states of nature. $\{p^t(X^{T-1})\}$ is a stochastic process called strategy. Utility functions are exponential.

To solve it, using dynamic programming, fix a compensation rule $s(X^T)$ first and define a value function

$$V_t = E\{u[s(X^T) - \sum_{t=1}^{T} c(p^t(X^{T-1}); \theta^t)] \mid X^t\}.$$

They proved the following theorem on the linearity of the incentive contract.

Theorem 1 An optimal compensation rule is

$$s(X^T) = \sum_{t=1}^{T} s(X^t; p^*) = s(p^*) \cdot A^T$$

where p^* is any single-period optimum, and A^T represents the vector $(A_1^t, \ldots, A_M^t)$ and A_i^t is the number of times in the first t periods the ith outcome occurs.

The proof is by induction. For exponential utility, the theorem just restates the definitions of p^* and s^*.

The sharing rule is a linear function of the aggregate output. The theorem also says the optimal sharing rule in a multi-period environment is based on the aggregate information A^T; which is not a sufficient statistic for agent's full strategy, however it is a sufficient statistic for deviations in that class, and so an optimal compensation scheme can be based on it.

However they addressed that theorem 1 does not say that the optimal scheme is necessarily linear on profits. Nevertheless, there is a special case, that is, two outcomes yielding two different profit levels.

Before going to continuous-time model, several remarks are made. First, some sufficient conditions were made in obtaining the results: exponential utility and a history- and time-independent technology. These assure that the T single period problems faced by the principal are identical and so have identical solutions; second, availibility of the credit and discounting are not considered here; finally, the timing of the information is a crucial aspect of the formulation.

Holmstrom and Milgrom further analyzed the linearity in a continuous time model. Consider the Brownian motion,

$$dZ = \mu dt + dB;$$

where the agent controls the drift rate μ and B is a driftless N-dimentional vector Brownian motion with covariance matrix Σ. Then it can be shown that,

Theorem 2 The stochastic process $\{\mu(t); 0 \le t \le 1\}$ is implemented with certainty equivalent w by sharing rule only if

$$\begin{aligned} s(Z^1) &= w + \int_0^1 c(\mu(t))dt + \{\int_0^1 c'(\mu(t))^T dZ - \int_0^1 c'(\mu(t))^T \mu(t)dt\} \\ &\quad + \frac{r}{2}\int_0^1 c'(\mu(t))^T \Sigma c'(\mu(t))dt. \end{aligned}$$

Now in order to find the optimal sharing rule, define the principal's Brownian problem accordingly, which is, the principal chooses instruction $\{\mu(t); 0 \le t \le 1\}$ and a sharing rule s which implements it and maximizes the principal's aggregate utility subject to the constraint that agent's certainty equivalent is no less than zero. Then,

Theorem 3 suppose that a $\mu^* \in M$ that maximizes principal's certainty equivalent, then an optimal solution to the principal's Brownian problem is to instruct the agent to set effort at μ^*, and set

$$s(Z^1) = c(\mu^*) + c'(\mu^*)^T(Z(1) - \mu^*) + \frac{r}{2}c'(\mu^*)^T \Sigma c'(\mu^*).$$

The proof follows from previous theorems and Ito's lemma for the exponential function. Therefore the agent's compensation scheme is a linear in profits. Once again, the timing is crucial for the optimality of linearity. A slight change in the timing of the agent's information reverses the conclusion. If we assume the agent observes next period's diffusion before choosing effort, the unique sharing rule that induces any particular sequence of pure actions does not approach the first-best outcome.

The other main conclusion beside the linearity result is that one need not always use all of the information available for an optimal incentive contract. It is linked with the aggregation property, as if the optimal incentive scheme works with aggregate accounting information, then it would be naturally useless to spend resources in obtaining detailed information on the agent's efforts.

Holmstrom and Milgrom(1987) claims that computational ease gives linearity substantial methodological value. They gave some examples. An intuitive example that can be treated in team production is relative evaluation. Assume that in addition to z, the principal observes another signal y, which can be a market index, or the output of some other agents, and could be included in the contract, then, straightforward computation illustrates that observing this additional signal y is equivalent to a reduciton in the variance of z if they are correlated. Indeed, if they are perfectly correlated, the first-best outcome can be attained and the the optimal contract filters out all uncertainties.

The second example is allocation of effort. Consider the agent can allocate his effort between two activities, both Brownian motions. If they are independent and the principal can only observe the total of the two, then the agent will allocate his effort symmetrically. However if the outcomes of the two activities can be observed separately, the optimal scheme is still linear, $s(z_1, z_2) = \alpha_1 z_1 + \alpha_2 z_2 + \beta$, but α_1 need not equal α_2, indeed, the only case for which the two are set equal is if z_1 and z_2 have the same variance. As an example of the case, if the agent can allocate time between reducing costs or increasing revenues and is equally effective at both, but if the revenue increasing is subject to more exogenous variance than the cost reducing, then the incentive scheme should not be based on profit alone; it should reward cost reductions more highly than revenue increases.

The above applications are not ad hoc dynamic models but the underlying optimal incentive contract, the linear contract, is derived from a dynamic formulation in the sense that the agent chooses strategy sequentially.

6 Career concerns

Career concerns are the concerns about the effect of current performance on future compensation. The optimal compensation scheme should maximize total incentives which is the combination of the implicit incentives from career concerns and the explicit incentives from the wage contract. Intuitively, optimal compensation should be stronger for agents who are close to retirement as career concerns are weaker for those. Fama (1980) first addressed this effect and argues that incentive contracts are not necessary because managers are disciplined by the managerial labor market. Holmstrom (1982, 1999) showed that with the absence of contracts, agents work too hard in the beginning of their career and not hard enough in later years. Gibbons and Murphy (1992) derived a model on this and found empirical support for this prediction between chief excutive compensation and stock market performance. In a formal model that incorprates career concerns, they gave an optimal contract which optimizes total incentives.

Consider an agent works for T periods. Output is

$$y_t = \eta + a_t + \epsilon_t.$$

where $\eta \sim N(m_0, \sigma_0^2)$ and $\epsilon_t \sim N(0, \sigma_\epsilon^2)$. The agent maximizes T-period aggregate exponential utility, with discounting. Assuming a linear contract $w_t(y_t) = c_t + b_t y_t$ where c_t is a fixed component of the wage; in a two-period model, by applying some well-known Bayesian results, the conditional distribution of η given first period output y_1 is then normal with mean

$$m_1 = \frac{\sigma_\epsilon^2 m_0 + \sigma_0^2 (y_1 - \hat{a}_1)}{\sigma_\epsilon^2 + \sigma_0^2},$$

and variance

$$\sigma_1^2 = \frac{\sigma_\epsilon^2 \sigma_0^2}{\sigma_\epsilon^2 + \sigma_0^2},$$

Given exponential utility, using a second-order Taylor approximation to optimize agent's utility, one obtains

$$b_2 = \frac{1}{1 + r(\sigma_\varepsilon^2 + \sigma_0^2)g''[a_2^*(b_2)]};$$

where $g(a)$ is agent's disutility. Given the assumption that $g''' > 0$; b_2^* decreases with both risk aversion and uncertainty.

For a T-period model, there are three equations that characterize the labor market.

$$c_t = (1 - b_t)E\{y_t \mid y_1, \ldots, y_{t-1}\};$$

$$E\{y_t \mid y_1, \ldots, y_{t-1}\} = m_{t-1} + \hat{a}_t;$$

$$m_{t-1}(y_1, \ldots, y_{t-1}; \hat{a}_1, \ldots, \hat{a}_{t-1}) = \frac{\sigma_\varepsilon^2 m_0 + \sigma_0^2 \sum_{\tau=1}^{t-1} (y_\tau - \hat{a}_\tau)}{\sigma_\varepsilon^2 + (t-1)\sigma_0^2}.$$

The first says that because of the competitiveness of the employers' market, agents receive all the expected surplus. First order conditions yield

$$\begin{aligned} g'(a_t^*) &= b_t + \sum_{\tau=t+1}^{T} \delta^{\tau-t} \frac{\partial c_\tau}{\partial a_t} \\ &= b_t + \sum_{\tau=t+1}^{T} \delta^{\tau-t}(1 - b_\tau)\frac{\sigma_0^2}{\sigma_\varepsilon^2 + (\tau-1)\sigma_0^2}. \end{aligned}$$

solving a first order condition similar to a two-period problem yields the optimal first-period contract slope $b_1^*(b_2, \ldots, b_T)$;

$$\begin{aligned} b_1 &= \frac{1}{1 + r(\sigma_\varepsilon^2 + \sigma_0^2)g''[a_1^*]} - \sum_{t=2}^{T} \delta^{\tau-t}(1 - b_t)\frac{\sigma_0^2}{\sigma_\varepsilon^2 + (\tau-1)\sigma_0^2} \\ &\quad - \frac{r\sigma_0^2 g''(a_1^*) \sum_{t=2} \delta^{t-1} B_t}{1 + r(\sigma_\varepsilon^2 + \sigma_0^2)g''[a_1^*]}. \end{aligned}$$

Applying it recursively yields the optimal piece rates $(b_1^*, \ldots, b_T^*)$; beginning with the last period T. And it is easy to prove that $b_t^* < b_{t+1}^* \ \forall\, t < T$.

Therefore, in an optimal incentive contract, the contractual incentives increase monotonically with t and so are the strongest to those about to retire.

Therefore in their model, explicit incentives play a role to balance undesired career concerns and optimize total incentives from the contract and from career concerns.

Gibbons and Murphy find further evidence of career concerns for CEOs from the data on CEOs' tenure. The data they took was by following all CEOs listed in the Executive Compensation Surveys published in Forbes from 1971 to 1989 which includes 2972 executives serving in 1493 of the naiton's largest corporations during the fiscal years 1970-99. In the completed-spells subsample, the average CEO's salary and bonus has grown by 6.6 percent per year over the sample period in 1988-constant dollars. Finally, the idea of optimal incentive contracts which neutralize career concern incentives and maxmize total incentives can be applied to promotions.

7 Dynamics in teams

7.1 Reputational concerns in teams

Jeon(1996) studies the reputation control on moral hazard in teams. Two organizational issues addressed here are grouping agents in teams and sharing team output among agents. The main conclusions are grouping junior team members with experienced team members are beneficial or, efficient to the team and equal sharing in teams which explains teams normally ignore obvious differences in members' attributes and share equally.

In a simple dynamic multiple agency model with two periods and two agents, two features are obtained: first, the incentive provided by reputational concerns in teams is weaker because of free-riding problem;l and the second, an agent's reputaional concerns are stronger if his own ability is less known or if his partner's ability is better known. These results were obtained from directly applying some results in Bayesian statistical inference. Therefore the first organizational suggestion is that an inter-generational grouping with both younger and older generations in teams, is more efficient than an intra-generational grouping with the same generation in teams. This statement reads like a common wisdom however Jeon emphasizes the old's value as partners for the young not as transferring knowledge and experience, but to stimulate the young to work harder to help updating the principal's belief on his ability.

Another issue Jeon addressed is output sharing arrangement. He suggests equal sharing in a team disregarding the difference in memebers' attributes. The virtue of equal sharing, is to gain from aligning incentives, even though equal sharing cannot change the sum of two agents' reputational coefficients, it results in reducing the gap and hence increasing the sum of net gains from two agents' efforts.

7.2 Teams with both adverse selection and moral hazard

McAfee and McMillan (1991) consider a tean subject to both adverse selection and moral hazard, where adverse selection here refers to that each memeber's ability is known only to himself and moral hazard refers to that effort cannot be observed. They suggest a payment scheme that is linear on total output.

In the case the moral hazard problem can be completely solved, the principal would offer to pay each of the n agents 100 percent of any marginal increase in team output. It gives each agent appropriate incentive to exert effort but it results in the principals total variable payment being n times the value of output. To balance it, there must be a negative fixed payment part which is set equal to the expected value of output minus the agent's production cost. The agents make zero rents on average. Now when adverse selection is added to the model, the agents will earn informaitonal rents. The principal extracts some of these informational rents by reducing the marginal payments below 100 percent. As the variable payments are reduced, the fixed payments ecome less negative. If the uncertainty about the agents' abilities were sufficiently dispersed, the marginal payment would become so small that they would sum to less than one, and the fixed payments would become positive and this looks like a more conventional payment scheme of salary plus commissions.

The reason why the principal pays less than full marginal product is that he would not be able to distinguish between high-ability and low-ability agents if he does not do so. McAfee and McMillan illustrate that in order to induce the agents' correct incentive to work in a team with both adverse selection and moral hazard, two conditions must hold, that is, agent i's effort increases with other agents' self-reported ability and the agent's share of marginal team output (if the payment scheme is based on total output) increases with his self-reported ability. Then with these two conditions, they prove that the optimal contract can be based on team output only. The rationale is the tradeoff between overstate and understate one's ability, for example, a low-ability agent is not tempted to overstate his ability because he is then faced with a payment function that is relatively sensitive to his effort.

Surprisingly, the principal is indifferent when he monitors individual contributions and when he obseves only total output, that is, no free-riding problem. This is clearly seen in the case of no adverse selection because no-monitoring contract pays each individudal at the margin 100 percent of marginal team output. Adding adverse selection does not break this equivalence because the rents that an agent receives as a result of private informaiton are completely determined by the contribution function abd this function can be induced by either monitored or nonmonitored contracts.

The authours agree that inefficiencies are undoubtedly inherent in team production as opposed to this model, however their model suggests that the source of team problem is not the unobservability of team members' efforts or abilities per se. The source must be sought elsewhere like team members' risk aversion or collusion among team members where these features are assumed away in their model.

8 First-order approach in a continuous-time principal-agent problem

Continuous time models differ from their discrete-time counterparts in that the agent is allowed to control the outcome process continuously during the contract time period. How would an optimal incentive contract formulation be in this framework? Sufficient conditions are derived in Schattler and Sung(1993) within the first-order approach framework.

With continuous time, the production function is modeled by a stochastic differential equation. The problem is then related to a stochastic optimal control problem by relaxing the agent's incentive compatibility constraint to the first-order conditions. The conditions, similar to the results obtained by Holmstrom and Milgrom(1987), which are only necessary conditions, take the form of a "semi-martingale representation" for the wage in terms of the agent's optimal control. This representation describes the pay as the sum of four components which can be interpreted as (i) the agent's opportunity cost, (ii) compensation for the agent's actual monetary cost of controlling the drift, (iii) compensation error due to the fact that incentive pay is based on the principal's observation of the realized outcome rather than on the agent's control effort, and (iv), a risk premium for the compensation error.

Consider a stochastic functional differential equation of the form

$$dX_t = f(t, X, u)dt + \sigma(t, X)dB_t.$$

where the agent controls the drift rate but has no control over the diffusion rate and B is a standard Brownian motion. f and σ are unknown functionals, and depend on the history not on the future. The agent has exponential utility

$$E[-\exp\{-r(S(X)-\int_0^1 c(t, X, u)dt)\}].$$

over all admissible effort control u, subject to the stochastic process described above. Following Holmstrom and Milgrom(define certainty equivalent), and apply Ito's lemma, the first-order necessary conditions for optimality in the agent's problem allow to derive the following representation for the agent's incentive pay:

$$\begin{aligned} s(X) &= w + \int_0^1 c(t, X, u)dt + \int_0^1 \nabla_u c(t, X, u)\cdot D_u^{-1} f(t, X, u)\sigma(t, X)dB_t \\ &\quad + \frac{r}{2}\int_0^1 \| \nabla_u c(t, X, u)\cdot D_u^{-1} f(t, X, u)\sigma(t, X)\|^2\, dt. \end{aligned}$$

Using the martingale approach, they outline the model for a general form of maximization problem which applied to bothe the agent's and the principal's relaxed problems.

It is shown that every control is implemetable[11] with a unique sharing rule if a certain function, called Hamiltonian, is formed solely from the data of the problem and is strictly convex in the control. In this case the first-order approach is always valid, and the class of compensation function which implements a given control will be much smaller than that in a static model.

---

[11] If the representation corresponding to a specific control is assigned to the agent, and if the agent actually uses the same control to maximize his utility, then such controls are called implementable.

Sung (1995) presents a continuous-time principal-agent model where the agent can control the mean and variance of outcome. It is shown that the optimal compensation contract is linear in final outcome alone even when the agent is allowed to control variance. Conditions under which no conflicts between the principal and the agent over the choice of variance due to unobservability of variance are characterized. When the principal is risk neutral, a sufficient condition for his irrelevance of the observability of variance is that the agent's cost function is convex and additively sparable in mean and variance. With an additively separable cost function, the agent can control the variance only to reduce the noise of the compensation without affecting the cost of controlling the mean. But reducing the noise simply leads to an improvement of the agent's incentives to work, which is what the risk-neutral principal wants. Thus the principal's and the agent's interests over the variance are alighned.

Managerial contract serves two goals: to provide incentives for managerial effort, and to induce the manager to choose particular projects. When the project choice is unobservable, the manager may choose either a riskier high net present value project or a safer low NPV project than the investors would desired. Given unobservability, the manager has an incentive to shoose a safer project at the expense of NPV, because the manager takes all the risk but is compensated only by a fraction of the realized profit. Thus the investors would have to reduce the sensitivity of the contract to induce the manager to take risky projects with higher NPV. Based on these observations, Sung (1995) tries to explain why optimal contracts have very low sensitivities documented in Jensen and Murphy (1990) by their empirical findings.

The derivation is quite technical however it merits the fact that the distribution of outcomes of the agent's actions cannot be ordered in terms of first-order stochastic dominance.

In Schattler and Sung (1997), they relate the existing first-order approaches for discrete- and continuous-time problems by considering continuous-time principal-agent problems as limiting cases of multi-period discrete-time formulations. An important difference between the discrete-time and continuous-time formulations is that the principal in the continuous time can design a contract based on intermediate outcomes as well as the final outcome. This feature in effect enables the principal to behave as if the principal could stop the outcome process anytime before the contact expires to revise the salary schedule. Because of this feature, salary schedule can be approximated by a linear function, the principal will not gain much by considering other complicated nonlinear salary schedules.

The first-order approach to the principal-agent problem with the class of linear contracts is much easier than that with all possible class of functions. Therefore, the first-order approach withe continuous-time formulation is easier to be justified. For the same reason, the principal-agent problem in a continuous-time framework can yield a solution even when the discrete-time counterpart fails.

9 The Ratchet effect and performance comparisons

Earlier research[12] shows that comparative performance information (CPI) can improve incentives and efficiency in the principal-agent relationship. The intuition is that since the comparison increases the precision of the estimated effort and therefore improves the terms of the tradeoff between insurance and efficiency in designing incentives for a risk-averse agent. In a dynamic setting, implicit incentives, namely, the ratchet effect, can be important even when explicit incentives are provided. To see it, first start from a two period model.

Consider one risk-neutral agent who works two periods and has the utility function

$$U = w_1 - C(e_1) + w_2 - C(e_2).$$

The ouput in period t is

$$x_t = e_t + a + u_t.$$

where a is a time invariant characterisitc of the agent, say, ability, and u_t is a transient shock. e_t is the effort level that privately is chosen by the agent.

Solving the model by backward induction, we know $e_2 = 0$. Therefore, $w_2 = E[x_2 \mid x_1] = E[a \mid x_1]$. A standard result shows that

$$w_2 = \tau(x_1 - \hat{e}_1)$$

where $\hat{e}_1$ is the labor market's conjecture about period 1 effort and $\tau = \frac{Var(a)}{Var(a)+Var(u)}$. Therefore first-order conditions show that the equilibrium effort level is lower than the first best.

---

[12] Holmstrom(1982), Mookherjee (1984).

Now introduce CPI, it remains true that $e_{2i} = 0$ and w_{2i} is still equal to the conditional expectation on a. Given the assumptions, the variables a, x_{1i} and x_{1j} have a multivariate normal distribution with covariance matrix proportional to

$$\begin{pmatrix} \tau & \tau & \eta\tau \\ \tau & 1 & \kappa \\ \eta\tau & \kappa & 1 \end{pmatrix};$$

where $\eta = corr(a_i; a_j)$ and $\rho = corr(u_{ti}; u_{tj})$ and $\kappa = (1 - \tau)\rho + \eta\tau$: Bayesian updating yields

$$w_{2i} = E(a_i \mid x_{1i}; x_{1j}) = \frac{\tau}{1-\kappa^2}[(1-\eta\kappa)(x_{1i} - \hat{e}_{1i}) + (\eta - \kappa)(x_{1j} - \hat{e}_{1j})]:$$

The first-order condition yields

$$C'(e_{1i}) = \tau\left(\frac{1-\eta\kappa}{1-\kappa^2}\right) \equiv a \leq 1:$$

Proposition 4 In managerial career concern model, effort incentives and efficiency are greater with performance comparisons than without iff $\kappa(\eta - \rho) < 0$:

Proof. By comparing with the no CPI case.

The ratchet effect is illustrated in a model of implicit and explicit incentives below. Assume that there are one principal P and two agents A_k as before: Output of A_k; $k = i; j$; in period $t = 1; 2$; at any firm, is

$$x_{tk} = e_{tk} + a_k + u_{tk}:$$

where

$$\begin{aligned} a_k &\sim N(0; \tau\sigma^2); \; u_{tk} \sim N(0; (1-\tau)\sigma^2); \; k = i; j: \\ corr(a_i; a_j) &= \eta; \quad corr(u_{ti}; u_{tj}) = \rho \end{aligned}$$

therefore

$$\begin{aligned} var(x_{tk}) &= \sigma^2; \; k = i; j: \\ var(x_{ti} &\mid x_{tj}) = (1-\kappa^2)\sigma^2 \\ var(x_{2i} &\mid x_{1j}) = (1-\tau^2)\sigma^2 \\ var(x_{2i} &\mid x_{1i}; x_{1j}; x_{2j}) = (1-\tau)(1+\gamma)(1-\rho^2)\sigma^2 \end{aligned}$$

where

$$\gamma = \tau[1 + 2(1 - \tau)\frac{N}{D}] \in [0,1]$$

and

$$\begin{aligned} N &= (\rho - \eta)(\rho + \eta\tau); \\ D &= (1 - \rho^2)(1 + \tau) + 2\tau N. \end{aligned}$$

Before deriving optimal incentive contracts, two assumptions are made: 1. only one-period contracts are enforceable; 2. one-period contracts take the linear form[13]

$$w_{ti} = \alpha_t + \beta_t x_{ti} + \epsilon_t x_{tj}.$$

First consider the effect of CPI in a static world, the principal minimizes the welfare loss relative to the first-best, which under exponential utility, is

$$\begin{aligned} l &= 1/2[(1 - \beta)^2 + r\mathrm{var}(\beta x_i + \epsilon x_j)] \\ &= 1/2[(1 - \beta)^2 + r\sigma^2(\beta^2 + \epsilon^2 + 2\beta\epsilon\kappa)] \end{aligned}$$

the optimal β is then given by

$$\beta = \frac{1}{1 + r\sigma^2(1 - \kappa^2)}$$

and for any given β, the optimal ϵ minimizes the variance of w_i;

$$\epsilon = -\beta[\frac{\mathrm{cov}(x_i, x_j)}{\mathrm{var}(x_j)}] = -\beta\kappa;$$

then it can be easily shown that the welfare loss under the wage schedule based on both agents' outputs is reduced. This is the improvement in the tradeoff between insurance effect and incentive effect: with CPI in the contract, the principal can filter out some of the uncertainty, either about ability of the agents or the transient shocks, and hence lower the risk cost of providing incentives to the agents.

Now, consider the two-period model. The conditional expectation of x_{2i} can be written in the form

$$E(x_{2i} \mid x_{1i}, x_{1j}, x_{2j}) = \hat{\mu}_{2i} + \gamma(x_{1i} - \hat{\mu}_{1i}) + \delta_1(x_{1j} - \hat{\mu}_{1i}) + \delta_2(x_{2j} - \hat{\mu}_{2j});$$

---

[13] see section on linearity of the dynamic incentive contract in this survey.

Principal chooses ϵ_2 to minimize $\mathrm{var}(w_{2i} \mid x_{1i}, x_{1j})$; therefore

$$\epsilon_2 = -\beta_2\left[\frac{\mathrm{cov}(x_{2i}, x_{2j} \mid x_{1j}, x_{1i})}{\mathrm{var}(x_{2i} \mid x_{1i}, x_{1j})}\right] = -\beta_2\delta_2,$$

and it yields

$$\begin{aligned}\mathrm{var}(w_{2i} \mid x_{1i}, x_{1j}) &= (\beta_2)^2\mathrm{var}(x_{2i} \mid x_{1i}, x_{1j}, x_{2j}) \\ &= (\beta_2)^2(1-\tau)(1+\gamma)(1-\rho^2)\sigma^2,\end{aligned}$$

Then by minimizing expected welfare loss by the principal, the optimal effort incentives in period 2 are given by

$$\beta_2 = \frac{1}{1 + r\sigma^2(1-\tau)(1+\gamma)(1-\rho^2)},$$

Constructing the agents' certainty equivalent, and deriving the variance of $w_{1i} + w_{2i}$, one obtains principal's loss function as

$$\begin{aligned} l &= 1/2\{(1-\hat{\beta}_1)^2 + (1-\beta_2)^2 + r\sigma^2[(\hat{\beta}_1 + \beta_2\gamma)^2(1-\eta^2) \\ &\quad +(\beta_2)^2(1-\tau)(1+\gamma)(1-\rho^2)]\}.\end{aligned}$$

Meyer and Vickers show that if one introduces a bargainning power parameter $b \in [0,1]$ to the agents which generates the reputation effect[14], although a change in b alters the strength of reputation effect, the welfare impact of this change can be costlessly offset by approperiate changes in β_1 and ϵ_1 which leave total incentives unchanged. However unlike the reputation effect, changes in the ratchet effect $\beta_2\gamma$ do affect welfare. Because when $\beta_2\gamma$ changes, the adjustments to β_1 would be costly though feasible.

[14] By approperiately constructing the agent's certainty equivalent, the second period wage can be written

$$w_{2i} = \mathrm{constant} + bE(a_i \mid x_{1i}, x_{1j}) + \beta_2[x_{2i} - E(x_{2i} \mid x_{1i}, x_{1j}, x_{2j})];$$

hence the overall wage $w_{1i} + w_{2i}$ is given by

$$\begin{aligned}\tilde{\beta}_1 = \beta_1 + b\frac{\partial}{\partial x_{1i}}E(a_i \mid x_{1i}, x_{1j}) - \beta_2\frac{\partial}{\partial x_{1i}}E(x_{2i} \mid x_{1i}, x_{1j}, x_{2j}) \\ = \beta_1 + b\alpha - \beta_2\gamma.\end{aligned}$$

where the term $b\alpha$ is interpreted as bargainning power/reputation effect which is the first type of implicit incentives, and the term $\beta_2\gamma$ is understood as ratchet effect which is another well-known implicit incentive documented in the literature.

So by introducing sequentially optimal explicit contracts and risk averse agents, the reputation effect no longer has significance: only the ratchet effect will have influences on agents' effort levels and welfare[15].

Finally welfare loss minimization yields a result which says the larger the ratchet effect, the mosre costly it is to provide first-period incentives and the larger the first period's share to the overall welfare loss.

To summarize, in the simple career concern model, no explicit incentives are designed and CPI influences incentives and henceforth efficiency only through implicit incentives, specifically, only the reputation effect. Depending on the correlation between agents' intrinsic characterisitics relative to that between the transitory shocks affecting their performances, CPI can raise or lower the efficiency.

In a two-period model in which explicit incentives are choice variables and contracts are sequentially optimal, given contracts take the linear form and risk averse agents, CPI can significantly influence the ratchet effect, which in this case[16], is the only determinant of overall incentives and so of welfare. If the correlation between the transitory shocks to agents' performances is not too great relative to that between their intrinsic characterisitcs, then CPI reduces the ratchet effect. However, the cost of reducing the ratchet effect might outweigh the the benefit from the improvement of insurance and incentives by introducing CPI in an explicit contract. Then, if the principal cannot commit ex ante on not readjust the contract terms in the future, it would be essential to evaluate the impact of CPI on ratchet effect before declaring an explicit contract with a CPI term.

10 A dynamic job matching model

There are three distinct features in labor market in a dynamic context: (1). wages rise with tenure in general, (2). quits are negatively correlated with tenure, in other words, people quit with a higher probablity in their earlier careers, (3) the subsequent quit is negatively correlated with the current wage. Jovanoiv (1979) tries to explain these features within a dynamic

---

[15] Where is in the pure implicit CPI incentive model discussed in the beginning, reputation effects, or career concerns, can increase incentives and welfare under certain conditions.

[16] When explicit contracts can be designed.

programming model.

Consider infinite horizon, let Q denote the expected present value of wages who was unemployed last period and who behaves optimally. The output in each period is

$$x = \theta + \varepsilon;$$

where ε is a random shock and θ is the agent's ability. Labor market is perfectly competitive on the employers' side, so the agent receives wage

$$E[\theta \mid (\theta + \varepsilon)];$$

θ and ε are i.i.d. random variables, both are normally distributed as

$$\theta \sim N(\mu; \sigma_0^2); \qquad \varepsilon \sim N(0; 1):$$

After the first period outcome, both principal and agent are about to make inferences on θ, using Bayes' law: Then the posterior distribution of θ is normal, with mean and variance

$$\begin{aligned} m_1 &= \mu + \frac{\sigma^2}{\sigma^2 + 1}(x - \mu); \\ \sigma_1^2 &= \frac{\sigma_0^2}{\sigma_0^2 + 1}: \end{aligned}$$

It follows

$$m_1 \sim N(\mu; \frac{\sigma_0^4}{\sigma_0^2 + 1}):$$

and since

$$\frac{\sigma_0^4}{\sigma_0^2 + 1} < \sigma_0^2;$$

therefore both agent and principal make more precise inference on agent's ability over time.

The agent maximizes the expected present value of wages. The Bellman's functional equation is

$$V(\theta) = \max\{\theta + \beta V(\theta)\};$$

The solution is

$$V(\theta) = \begin{cases} \theta + \beta V(\theta) & \text{for } \theta > \bar{\theta} \\ \beta Q & \text{for } \theta \leq \bar{\theta} \end{cases}.$$

The optimal policy is a reservation wage policy: accept offers $\theta > \bar{\theta}$; and reject offers $\theta \leq \bar{\theta}$: where $\bar{\theta}$ solves

$$\frac{\bar{\theta}}{1-\beta} = \beta Q.$$

Since rejections exist, the second period Bellman's equation becomes

$$V(m_1) = \max\{(m_1 + \beta \int V(\theta)dF(\theta \mid m_1; \sigma_1^2); \beta Q\}.$$

where m_1 is the current wage offer. From the solution to the program one can show

$$\bar{m}_1 + \beta\frac{\bar{\theta}}{1-\beta}\int_0^{\bar{\theta}} dF(\theta \mid \bar{m}_1; \sigma_1^2) + \frac{\beta}{1-\beta}\int_{\bar{\theta}}^{\infty} \theta dF(\theta \mid \bar{m}_1; \sigma_1^2)$$
$$= \frac{\bar{\theta}}{1-\beta} = \frac{\bar{\theta}}{1-\beta}\int_0^{\bar{\theta}} dF(\theta \mid \bar{m}_1; \sigma_1^2)$$
$$+\frac{\bar{\theta}}{1-\beta}\int_{\bar{\theta}}^{\infty} dF(\theta \mid \bar{m}_1; \sigma_1^2)$$

and it implies

$$\bar{\theta} - \bar{m}_1 = \frac{\beta}{1-\beta}\int_{\bar{\theta}}^{\infty} (\theta - \bar{\theta})dF(\theta \mid \bar{m}_1; \sigma_1^2)$$

therefore

$$\bar{\theta} > \bar{m}_1.$$

So the agent turns to be selective over time.

Other results are: the probablity that a previously unemployed agent accepts an offer is

$$\Pr\{m_1 > \bar{m}_1\} = \int_{\bar{m}_1}^{\infty} dG[m_1 \mid \mu; \sigma_0^4/(1+\sigma_0^2)].$$

where $G[m_1 \mid \mu; \sigma_0^4/(1+\sigma_0^2))$ is normal distribution of m_1.

The mean wage of those who were hired in the first period is

$$\bar{w}_1 = \frac{\int_{\bar{m}_1}^{\infty} m_1 dG[m_1 \mid \mu; \sigma_0^4/(1+\sigma_0^2)]}{\int_{\bar{m}_1}^{\infty} dG[m_1 \mid \mu; \sigma_0^4/(1+\sigma_0^2)]}.$$

whereas the mean wage of those who are in the second period of tenure is

$$\bar{w}_2 = \frac{\int_{\bar{m}_1}^{\infty}\int_{\bar{\theta}}^{\infty} \theta_1 dF(\theta \mid m_1; \sigma_1^2) dG[m_1 \mid \mu; \sigma_0^4/(1+\sigma_0^2)]}{\int_{\bar{m}_1}^{\infty}\int_{\bar{\theta}}^{\infty} dF(\theta \mid m_1; \sigma_1^2) dG[m_1 \mid \mu; \sigma_0^4/(1+\sigma_0^2)]}.$$

and it can be proved that

$$\bar{w}_1 < \bar{w}_2.$$

so that tenure salary increases over time.

11 Conclusion

We covered a wide range of issues associated with compensation and incentives in a broad context of dynamic setting in this survey. The past two decades have witnessed the progress towards understadning two limited aspects of employment relationships: incentive pay in a long run relationship and careers in organizations. Several other concrete questions are also analysed, including payment scheme, sharing and grouping in a team, robustness of linearity, incentive pay when project selection is involved. Theoretical models aim to explain or comply empirical findings such as low sensitivity of the incentive contract, increasing incentive pay (sometimes promotion) over time, equal sharing in a team despite the fact that team members's ability widely spread, and alike. Most models do gain insights in those issues.

Much of the theory starts from a conventional story and subjects to usual assumptions. However some fundamental questions need work: what are the exact jobs an agent entitled with? Do the agent's abilities increase substantially over time? How does the job market really affect the agent's incentive(rather than focusing only on what goes on inside the firms)? If these areas can be carefully explored, then we would have a better picture of optimal labor contracts.

References

Aghion, P., and P. Bolton. 1987. Contracts as a barrier to entry, American Economic Review 77:388-401.

Caillaud, B., B. Jullien, and P. Picard. 1995. Competing vertical structures: Precommitment and renegotiation. Econometrica 63:621 6.

Chiappori, P.-A., 1. Macho, P. Rey, and B. Salanid. 1994. Repeated moral hazard: The role of memory, commitment, and the access to credit markets. European Economic Review, 38:1527-53.

Dewatripont, M. 1988. Commitment through renegotiation-proof contracts with third parties. Review of Economic Studies 55:377-90.

Dewatripont, M. 1989. Renegotiation and information revelation over time: The case of optimal labor contracts. Quarterly Journal of Economics 104:589-619.

Dewatripont, M., I. Jewitt and J. Tirole. 1999. The economics of career concerns, Part I: comparing information structure Review of Economic Studies 66:183-198.

Durrett, R. 1996. Stochastic Calculus: a Practial Introduction. CRC Press.

Fershtman, C., and K. judd. 1987. Equilibrium incentives in oligopoly. American Economic Review 77:927-40.

Freixas, X., R. Guesnerie, and J. Tirole. 1985. Planning under incomplete information and the ratchet e®ect. Review of Economic Studies 52:173-92.

Fudenberg, D., B. Holmstrom, and P. Milgrom. 1990. Short-term contracts and long-term agency relationships. Journal of Econmmic Theory 51:1-31.

Fudenberg, D., and J. Tirole. 1990. Moral hazard and renegotiation in agency contracts. Econometrica 58:1279-1320.

Hart, 0., and J. Tirole. 1988. Contract renegotiation and coasian dynamics. Review of Economic Studies 55:509-40.

Holmtrom, B. 1999. Managerial incentive problems: a dynamic perspective. Review of Economic Studies. 66, 169-182.

Holmtrom, B., and P. Milgrom. 1987. Aggregation and linearity in the provision of intertemporal incentives. Econometrica 55:303-28.

Gibbons, R., and K. Murphy. 1992. Optimal incentive contracts in the presence of career concerns: theory and evidence, Journal of Political Economy, 468-505.

Jeon., S., 1996. Moral hazard and reputational concerns in teams: implications for organizational choice. International Journal of Industrial Organization, 14:297-315.

Katz, M. 1991, Game-playing agents: Unobservable contracts as precommitment. Rand Journal of Economics 22:307-28.

Laffont, J.-J., and J. Tirole. 1987. Comparative statics of the optimal dynamic incentives contract. European Economics Review 31:901-26.

Laffont, J.-J., and J. Tirole. 1990. Adverse selection and renegotiation in procurement. Review of Economic Studies 75:597-626.

Laffont, J.-J., and J. Tirole. 1993. A Theory of Incentives in Procurement and Regulation. Cambridge: MIT Press.

Jovanovic, B., 1979. Job matching and the theory of turnover. Journal of Political Economy 87(5):972-990.

Malcolmson, j., and F. Spintiewyn. 1988. The multiperiod Principal-Agent problem. Review of Economic Studies 55:391 08.

McAfee, R., and J.McMillan. 1991. Optimal contract for teams. International Economic Review Vol 33, No. 3 561-577.

Meyer, M., 1994. The dynamics of learning with team production: implications for task assignment, Quarterly Journal of Economics, 1157-1184.

Meyer, M., and J. Vickers., 1997, performance comparisons and dynamic incentives. Journal of Political Economy. 105: 547-581.

Rnovaes, W., and L. Zingales. 1998. Bureaucracy as a mechanism to generate information. memio.

Rey, P., and B. Salanie. 1990. Long-term, short-term and renegotiation: On the value of commitment in contracting. Econometrica 58:597-619.

Rey, P., and B. Salanie.1996. On the value of commitment in contracting with asymmetric information. Econometrica 64:1395-1414.

Rogerson, W. 1985. Repeated moral hazard. Econometrica 53:69-76.

Rubinstein, A., and M. Yaari. 1983. Insurance and moral hazard. Journal of Economic Theory 14:441-52.

Schattler, H., and J. Sung. 1993. The first-order approach to the continuous-time principal-agent problem with exponential utility. Journal of Economic Theory 61:331-371.

Schattler, H., and J. Sung. 1997. On optimal sharing rules in discrete- and continuous-time principal-agent problems with exponential utility. Journal of Economic Dynamics and Control 21:551-574.

Sung, J. 1995. Linearity with project selection and controllable diffusion rate in continuous-time principal-agent problems. RAND Journal of Economics 26, No. 4:720-743.